

ECONOMETRICS

BRUCE E. HANSEN
©2000, 2015¹

University of Wisconsin
Department of Economics

This Revision: January 16, 2015
Comments Welcome

¹This manuscript may be printed and reproduced for individual or instructional use, but may not be printed for commercial purposes.

Contents

Preface	viii
1 Introduction	1
1.1 What is Econometrics?	1
1.2 The Probability Approach to Econometrics	1
1.3 Econometric Terms and Notation	2
1.4 Observational Data	3
1.5 Standard Data Structures	4
1.6 Sources for Economic Data	5
1.7 Econometric Software	7
1.8 Reading the Manuscript	7
1.9 Common Symbols	8
2 Conditional Expectation and Projection	9
2.1 Introduction	9
2.2 The Distribution of Wages	9
2.3 Conditional Expectation	11
2.4 Log Differences*	13
2.5 Conditional Expectation Function	14
2.6 Continuous Variables	15
2.7 Law of Iterated Expectations	16
2.8 CEF Error	18
2.9 Intercept-Only Model	19
2.10 Regression Variance	19
2.11 Best Predictor	20
2.12 Conditional Variance	21
2.13 Homoskedasticity and Heteroskedasticity	22
2.14 Regression Derivative	23
2.15 Linear CEF	24
2.16 Linear CEF with Nonlinear Effects	25
2.17 Linear CEF with Dummy Variables	26
2.18 Best Linear Predictor	28
2.19 Linear Predictor Error Variance	34
2.20 Regression Coefficients	35
2.21 Regression Sub-Vectors	35
2.22 Coefficient Decomposition	36
2.23 Omitted Variable Bias	37
2.24 Best Linear Approximation	38
2.25 Normal Regression	38
2.26 Regression to the Mean	39
2.27 Reverse Regression	40
2.28 Limitations of the Best Linear Predictor	41

2.29	Random Coefficient Model	41
2.30	Causal Effects	43
2.31	Expectation: Mathematical Details*	47
2.32	Existence and Uniqueness of the Conditional Expectation*	49
2.33	Identification*	50
2.34	Technical Proofs*	51
	Exercises	55
3	The Algebra of Least Squares	57
3.1	Introduction	57
3.2	Random Samples	57
3.3	Sample Means	58
3.4	Least Squares Estimator	58
3.5	Solving for Least Squares with One Regressor	59
3.6	Solving for Least Squares with Multiple Regressors	60
3.7	Illustration	62
3.8	Least Squares Residuals	62
3.9	Model in Matrix Notation	63
3.10	Projection Matrix	65
3.11	Orthogonal Projection	66
3.12	Estimation of Error Variance	67
3.13	Analysis of Variance	68
3.14	Regression Components	68
3.15	Residual Regression	70
3.16	Prediction Errors	71
3.17	Influential Observations	72
3.18	Normal Regression Model	74
3.19	CPS Data Set	76
3.20	Programming	78
3.21	Technical Proofs*	82
	Exercises	83
4	Least Squares Regression	86
4.1	Introduction	86
4.2	Sample Mean	86
4.3	Linear Regression Model	87
4.4	Mean of Least-Squares Estimator	88
4.5	Variance of Least Squares Estimator	89
4.6	Gauss-Markov Theorem	91
4.7	Residuals	92
4.8	Estimation of Error Variance	93
4.9	Mean-Square Forecast Error	95
4.10	Covariance Matrix Estimation Under Homoskedasticity	96
4.11	Covariance Matrix Estimation Under Heteroskedasticity	97
4.12	Standard Errors	100
4.13	Computation	101
4.14	Measures of Fit	102
4.15	Empirical Example	103
4.16	Multicollinearity	105
4.17	Normal Regression Model	108
	Exercises	110

5	An Introduction to Large Sample Asymptotics	112
5.1	Introduction	112
5.2	Asymptotic Limits	112
5.3	Convergence in Probability	114
5.4	Weak Law of Large Numbers	115
5.5	Almost Sure Convergence and the Strong Law*	116
5.6	Vector-Valued Moments	117
5.7	Convergence in Distribution	118
5.8	Higher Moments	120
5.9	Functions of Moments	121
5.10	Delta Method	123
5.11	Stochastic Order Symbols	124
5.12	Uniform Stochastic Bounds*	126
5.13	Semiparametric Efficiency	127
5.14	Technical Proofs*	130
	Exercises	134
6	Asymptotic Theory for Least Squares	135
6.1	Introduction	135
6.2	Consistency of Least-Squares Estimator	136
6.3	Asymptotic Normality	137
6.4	Joint Distribution	141
6.5	Consistency of Error Variance Estimators	144
6.6	Homoskedastic Covariance Matrix Estimation	144
6.7	Heteroskedastic Covariance Matrix Estimation	145
6.8	Summary of Covariance Matrix Notation	147
6.9	Alternative Covariance Matrix Estimators*	147
6.10	Functions of Parameters	148
6.11	Asymptotic Standard Errors	151
6.12	t statistic	153
6.13	Confidence Intervals	154
6.14	Regression Intervals	155
6.15	Forecast Intervals	157
6.16	Wald Statistic	158
6.17	Homoskedastic Wald Statistic	158
6.18	Confidence Regions	159
6.19	Semiparametric Efficiency in the Projection Model	160
6.20	Semiparametric Efficiency in the Homoskedastic Regression Model*	162
6.21	Uniformly Consistent Residuals*	163
6.22	Asymptotic Leverage*	164
	Exercises	166
7	Restricted Estimation	169
7.1	Introduction	169
7.2	Constrained Least Squares	170
7.3	Exclusion Restriction	171
7.4	Minimum Distance	172
7.5	Asymptotic Distribution	173
7.6	Efficient Minimum Distance Estimator	174
7.7	Exclusion Restriction Revisited	175
7.8	Variance and Standard Error Estimation	177
7.9	Misspecification	177

7.10	Nonlinear Constraints	179
7.11	Inequality Restrictions	180
7.12	Constrained MLE	181
7.13	Technical Proofs*	181
	Exercises	183
8	Hypothesis Testing	185
8.1	Hypotheses	185
8.2	Acceptance and Rejection	186
8.3	Type I Error	187
8.4	t tests	187
8.5	Type II Error and Power	188
8.6	Statistical Significance	189
8.7	P-Values	190
8.8	t-ratios and the Abuse of Testing	192
8.9	Wald Tests	192
8.10	Homoskedastic Wald Tests	194
8.11	Criterion-Based Tests	194
8.12	Minimum Distance Tests	195
8.13	Minimum Distance Tests Under Homoskedasticity	196
8.14	F Tests	197
8.15	Likelihood Ratio Test	198
8.16	Problems with Tests of NonLinear Hypotheses	199
8.17	Monte Carlo Simulation	202
8.18	Confidence Intervals by Test Inversion	204
8.19	Power and Test Consistency	205
8.20	Asymptotic Local Power	207
8.21	Asymptotic Local Power, Vector Case	210
8.22	Technical Proofs*	211
	Exercises	213
9	Regression Extensions	215
9.1	NonLinear Least Squares	215
9.2	Generalized Least Squares	218
9.3	Testing for Heteroskedasticity	221
9.4	Testing for Omitted NonLinearity	221
9.5	Least Absolute Deviations	222
9.6	Quantile Regression	224
	Exercises	227
10	The Bootstrap	229
10.1	Definition of the Bootstrap	229
10.2	The Empirical Distribution Function	229
10.3	Nonparametric Bootstrap	231
10.4	Bootstrap Estimation of Bias and Variance	231
10.5	Percentile Intervals	232
10.6	Percentile-t Equal-Tailed Interval	234
10.7	Symmetric Percentile-t Intervals	234
10.8	Asymptotic Expansions	235
10.9	One-Sided Tests	237
10.10	Symmetric Two-Sided Tests	238
10.11	Percentile Confidence Intervals	239

10.12 Bootstrap Methods for Regression Models	240
Exercises	242
11 NonParametric Regression	243
11.1 Introduction	243
11.2 Binned Estimator	243
11.3 Kernel Regression	245
11.4 Local Linear Estimator	246
11.5 Nonparametric Residuals and Regression Fit	247
11.6 Cross-Validation Bandwidth Selection	249
11.7 Asymptotic Distribution	252
11.8 Conditional Variance Estimation	255
11.9 Standard Errors	255
11.10 Multiple Regressors	256
12 Series Estimation	259
12.1 Approximation by Series	259
12.2 Splines	259
12.3 Partially Linear Model	261
12.4 Additively Separable Models	261
12.5 Uniform Approximations	261
12.6 Runge's Phenomenon	263
12.7 Approximating Regression	263
12.8 Residuals and Regression Fit	266
12.9 Cross-Validation Model Selection	266
12.10 Convergence in Mean-Square	267
12.11 Uniform Convergence	268
12.12 Asymptotic Normality	269
12.13 Asymptotic Normality with Undersmoothing	270
12.14 Regression Estimation	271
12.15 Kernel Versus Series Regression	272
12.16 Technical Proofs	272
13 Generalized Method of Moments	278
13.1 Overidentified Linear Model	278
13.2 GMM Estimator	279
13.3 Distribution of GMM Estimator	280
13.4 Estimation of the Efficient Weight Matrix	281
13.5 GMM: The General Case	282
13.6 Over-Identification Test	282
13.7 Hypothesis Testing: The Distance Statistic	283
13.8 Conditional Moment Restrictions	284
13.9 Bootstrap GMM Inference	285
Exercises	287
14 Empirical Likelihood	289
14.1 Non-Parametric Likelihood	289
14.2 Asymptotic Distribution of EL Estimator	291
14.3 Overidentifying Restrictions	292
14.4 Testing	293
14.5 Numerical Computation	294

15	Endogeneity	296
15.1	Instrumental Variables	297
15.2	Reduced Form	298
15.3	Identification	299
15.4	Estimation	299
15.5	Special Cases: IV and 2SLS	299
15.6	Bekker Asymptotics	301
15.7	Identification Failure	302
	Exercises	304
16	Univariate Time Series	306
16.1	Stationarity and Ergodicity	306
16.2	Autoregressions	308
16.3	Stationarity of AR(1) Process	309
16.4	Lag Operator	309
16.5	Stationarity of AR(k)	310
16.6	Estimation	310
16.7	Asymptotic Distribution	311
16.8	Bootstrap for Autoregressions	312
16.9	Trend Stationarity	312
16.10	Testing for Omitted Serial Correlation	313
16.11	Model Selection	314
16.12	Autoregressive Unit Roots	314
17	Multivariate Time Series	316
17.1	Vector Autoregressions (VARs)	316
17.2	Estimation	317
17.3	Restricted VARs	317
17.4	Single Equation from a VAR	317
17.5	Testing for Omitted Serial Correlation	318
17.6	Selection of Lag Length in an VAR	318
17.7	Granger Causality	319
17.8	Cointegration	319
17.9	Cointegrated VARs	320
18	Limited Dependent Variables	322
18.1	Binary Choice	322
18.2	Count Data	323
18.3	Censored Data	324
18.4	Sample Selection	325
19	Panel Data	327
19.1	Individual-Effects Model	327
19.2	Fixed Effects	327
19.3	Dynamic Panel Regression	329
20	Nonparametric Density Estimation	330
20.1	Kernel Density Estimation	330
20.2	Asymptotic MSE for Kernel Estimates	332

A Matrix Algebra	335
A.1 Notation	335
A.2 Matrix Addition	336
A.3 Matrix Multiplication	336
A.4 Trace	337
A.5 Rank and Inverse	338
A.6 Determinant	339
A.7 Eigenvalues	340
A.8 Positive Definiteness	341
A.9 Matrix Calculus	342
A.10 Kronecker Products and the Vec Operator	342
A.11 Vector and Matrix Norms	343
A.12 Matrix Inequalities	343
B Probability	348
B.1 Foundations	348
B.2 Random Variables	350
B.3 Expectation	350
B.4 Gamma Function	351
B.5 Common Distributions	352
B.6 Multivariate Random Variables	354
B.7 Conditional Distributions and Expectation	356
B.8 Transformations	358
B.9 Normal and Related Distributions	359
B.10 Inequalities	361
B.11 Maximum Likelihood	364
C Numerical Optimization	369
C.1 Grid Search	369
C.2 Gradient Methods	369
C.3 Derivative-Free Methods	371

Preface

This book is intended to serve as the textbook for a first-year graduate course in econometrics. It can be used as a stand-alone text, or be used as a supplement to another text.

Students are assumed to have an understanding of multivariate calculus, probability theory, linear algebra, and mathematical statistics. A prior course in undergraduate econometrics would be helpful, but not required. Two excellent undergraduate textbooks are Wooldridge (2009) and Stock and Watson (2010).

For reference, some of the basic tools of matrix algebra, probability, and statistics are reviewed in the Appendix.

For students wishing to deepen their knowledge of matrix algebra in relation to their study of econometrics, I recommend *Matrix Algebra* by Abadir and Magnus (2005).

An excellent introduction to probability and statistics is *Statistical Inference* by Casella and Berger (2002). For those wanting a deeper foundation in probability, I recommend Ash (1972) or Billingsley (1995). For more advanced statistical theory, I recommend Lehmann and Casella (1998), van der Vaart (1998), Shao (2003), and Lehmann and Romano (2005).

For further study in econometrics beyond this text, I recommend Davidson (1994) for asymptotic theory, Hamilton (1994) for time-series methods, Wooldridge (2002) for panel data and discrete response models, and Li and Racine (2007) for nonparametrics and semiparametric econometrics. Beyond these texts, the *Handbook of Econometrics* series provides advanced summaries of contemporary econometric methods and theory.

The end-of-chapter exercises are important parts of the text and are meant to help teach students of econometrics. Answers are not provided, and this is intentional.

I would like to thank Ying-Ying Lee for providing research assistance in preparing some of the empirical examples presented in the text.

As this is a manuscript in progress, some parts are quite incomplete, and there are many topics which I plan to add. In general, the earlier chapters are the most complete while the later chapters need significant work and revision.

Chapter 1

Introduction

1.1 What is Econometrics?

The term “econometrics” is believed to have been crafted by Ragnar Frisch (1895-1973) of Norway, one of the three principal founders of the Econometric Society, first editor of the journal *Econometrica*, and co-winner of the first Nobel Memorial Prize in Economic Sciences in 1969. It is therefore fitting that we turn to Frisch’s own words in the introduction to the first issue of *Econometrica* to describe the discipline.

A word of explanation regarding the term econometrics may be in order. Its definition is implied in the statement of the scope of the [Econometric] Society, in Section I of the Constitution, which reads: “The Econometric Society is an international society for the advancement of economic theory in its relation to statistics and mathematics.... Its main object shall be to promote studies that aim at a unification of the theoretical-quantitative and the empirical-quantitative approach to economic problems....”

But there are several aspects of the quantitative approach to economics, and no single one of these aspects, taken by itself, should be confounded with econometrics. Thus, econometrics is by no means the same as economic statistics. Nor is it identical with what we call general economic theory, although a considerable portion of this theory has a definitely quantitative character. Nor should econometrics be taken as synonymous with the application of mathematics to economics. Experience has shown that each of these three view-points, that of statistics, economic theory, and mathematics, is a necessary, but not by itself a sufficient, condition for a real understanding of the quantitative relations in modern economic life. It is the *unification* of all three that is powerful. And it is this unification that constitutes econometrics.

Ragnar Frisch, *Econometrica*, (1933), 1, pp. 1-2.

This definition remains valid today, although some terms have evolved somewhat in their usage. Today, we would say that econometrics is the unified study of economic models, mathematical statistics, and economic data.

Within the field of econometrics there are sub-divisions and specializations. **Econometric theory** concerns the development of tools and methods, and the study of the properties of econometric methods. **Applied econometrics** is a term describing the development of quantitative economic models and the application of econometric methods to these models using economic data.

1.2 The Probability Approach to Econometrics

The unifying methodology of modern econometrics was articulated by Trygve Haavelmo (1911-1999) of Norway, winner of the 1989 Nobel Memorial Prize in Economic Sciences, in his seminal

paper “The probability approach in econometrics”, *Econometrica* (1944). Haavelmo argued that quantitative economic models must necessarily be *probability models* (by which today we would mean *stochastic*). Deterministic models are blatantly inconsistent with observed economic quantities, and it is incoherent to apply deterministic models to non-deterministic data. Economic models should be explicitly designed to incorporate randomness; stochastic errors should not be simply added to deterministic models to make them random. Once we acknowledge that an economic model is a probability model, it follows naturally that an appropriate tool way to quantify, estimate, and conduct inferences about the economy is through the powerful theory of mathematical statistics. The appropriate method for a quantitative economic analysis follows from the probabilistic construction of the economic model.

Haavelmo’s probability approach was quickly embraced by the economics profession. Today no quantitative work in economics shuns its fundamental vision.

While all economists embrace the probability approach, there has been some evolution in its implementation.

The **structural approach** is the closest to Haavelmo’s original idea. A probabilistic economic model is specified, and the quantitative analysis performed under the assumption that the economic model is correctly specified. Researchers often describe this as “taking their model seriously.” The structural approach typically leads to likelihood-based analysis, including maximum likelihood and Bayesian estimation.

A criticism of the structural approach is that it is misleading to treat an economic model as correctly specified. Rather, it is more accurate to view a model as a useful abstraction or approximation. In this case, how should we interpret structural econometric analysis? The **quasi-structural approach** to inference views a structural economic model as an approximation rather than the truth. This theory has led to the concepts of the pseudo-true value (the parameter value defined by the estimation problem), the quasi-likelihood function, quasi-MLE, and quasi-likelihood inference.

Closely related is the **semiparametric approach**. A probabilistic economic model is partially specified but some features are left unspecified. This approach typically leads to estimation methods such as least-squares and the Generalized Method of Moments. The semiparametric approach dominates contemporary econometrics, and is the main focus of this textbook.

Another branch of quantitative structural economics is the **calibration approach**. Similar to the quasi-structural approach, the calibration approach interprets structural models as approximations and hence inherently false. The difference is that the calibrationist literature rejects mathematical statistics (deeming classical theory as inappropriate for approximate models) and instead selects parameters by matching model and data moments using non-statistical *ad hoc*¹ methods.

1.3 Econometric Terms and Notation

In a typical application, an econometrician has a set of repeated measurements on a set of variables. For example, in a labor application the variables could include weekly earnings, educational attainment, age, and other descriptive characteristics. We call this information the **data**, **dataset**, or **sample**.

We use the term **observations** to refer to the distinct repeated measurements on the variables. An individual observation often corresponds to a specific economic unit, such as a person, household, corporation, firm, organization, country, state, city or other geographical region. An individual observation could also be a measurement at a point in time, such as quarterly GDP or a daily interest rate.

¹*Ad hoc* means “for this purpose” – a method designed for a specific problem – and not based on a generalizable principle.

Economists typically denote variables by the italicized roman characters y , x , and/or z . The convention in econometrics is to use the character y to denote the variable to be explained, while the characters x and z are used to denote the conditioning (explaining) variables.

Following mathematical convention, real numbers (elements of the real line \mathbb{R} , also called **scalars**) are written using lower case italics such as y , and vectors (elements of \mathbb{R}^k) by lower case bold italics such as \mathbf{x} , e.g.

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_k \end{pmatrix}.$$

Upper case bold italics such as \mathbf{X} are used for matrices.

We denote the number of observations by the natural number n , and subscript the variables by the index i to denote the individual observation, e.g. y_i , \mathbf{x}_i and \mathbf{z}_i . In some contexts we use indices other than i , such as in time-series applications where the index t is common and T is used to denote the number of observations. In panel studies we typically use the double index it to refer to individual i at a time period t .

The i^{th} observation is the set $(y_i, \mathbf{x}_i, \mathbf{z}_i)$. The sample is the set $\{(y_i, \mathbf{x}_i, \mathbf{z}_i) : i = 1, \dots, n\}$.

It is proper mathematical practice to use upper case X for random variables and lower case x for realizations or specific values. Since we use upper case to denote matrices, the distinction between random variables and their realizations is not rigorously followed in econometric notation. Thus the notation y_i will in some places refer to a random variable, and in other places a specific realization. This is an undesirable but there is little to be done about it without terrifically complicating the notation. Hopefully there will be no confusion as the use should be evident from the context.

We typically use Greek letters such as β , θ and σ^2 to denote unknown parameters of an econometric model, and will use boldface, e.g. $\boldsymbol{\beta}$ or $\boldsymbol{\theta}$, when these are vector-valued. Estimates are typically denoted by putting a hat “ $\hat{\cdot}$ ”, tilde “ $\tilde{\cdot}$ ” or bar “ $\bar{\cdot}$ ” over the corresponding letter, e.g. $\hat{\beta}$ and $\tilde{\beta}$ are estimates of β .

The covariance matrix of an econometric estimator will typically be written using the capital boldface \mathbf{V} , often with a subscript to denote the estimator, e.g. $\mathbf{V}_{\hat{\boldsymbol{\beta}}} = \text{var}(\hat{\boldsymbol{\beta}})$ as the covariance matrix for $\hat{\boldsymbol{\beta}}$. Hopefully without causing confusion, we will use the notation $\mathbf{V}_{\boldsymbol{\beta}} = \text{avar}(\hat{\boldsymbol{\beta}})$ to denote the asymptotic covariance matrix of $\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})$ (the variance of the asymptotic distribution). Estimates will be denoted by appending hats or tildes, e.g. $\hat{\mathbf{V}}_{\boldsymbol{\beta}}$ is an estimate of $\mathbf{V}_{\boldsymbol{\beta}}$.

1.4 Observational Data

A common econometric question is to quantify the impact of one set of variables on another variable. For example, a concern in labor economics is the returns to schooling – the change in earnings induced by increasing a worker’s education, holding other variables constant. Another issue of interest is the earnings gap between men and women.

Ideally, we would use **experimental** data to answer these questions. To measure the returns to schooling, an experiment might randomly divide children into groups, mandate different levels of education to the different groups, and then follow the children’s wage path after they mature and enter the labor force. The differences between the groups would be direct measurements of the effects of different levels of education. However, experiments such as this would be widely

condemned as immoral! Consequently, in economics non-laboratory experimental data sets are typically narrow in scope.

Instead, most economic data is **observational**. To continue the above example, through data collection we can record the level of a person's education and their wage. With such data we can measure the joint distribution of these variables, and assess the joint dependence. But from observational data it is difficult to infer **causality**, as we are not able to manipulate one variable to see the direct effect on the other. For example, a person's level of education is (at least partially) determined by that person's choices. These factors are likely to be affected by their personal abilities and attitudes towards work. The fact that a person is highly educated suggests a high level of ability, which suggests a high relative wage. This is an alternative explanation for an observed positive correlation between educational levels and wages. High ability individuals do better in school, and therefore choose to attain higher levels of education, and their high ability is the fundamental reason for their high wages. The point is that multiple explanations are consistent with a positive correlation between schooling levels and education. Knowledge of the joint distribution alone may not be able to distinguish between these explanations.

Most economic data sets are observational, not experimental. This means that all variables must be treated as random and possibly jointly determined.

This discussion means that it is difficult to infer causality from observational data alone. Causal inference requires identification, and this is based on strong assumptions. We will discuss these issues on occasion throughout the text.

1.5 Standard Data Structures

There are three major types of economic data sets: cross-sectional, time-series, and panel. They are distinguished by the dependence structure across observations.

Cross-sectional data sets have one observation per individual. Surveys are a typical source for cross-sectional data. In typical applications, the individuals surveyed are persons, households, firms or other economic agents. In many contemporary econometric cross-section studies the sample size n is quite large. It is conventional to assume that cross-sectional observations are mutually independent. Most of this text is devoted to the study of cross-section data.

Time-series data are indexed by time. Typical examples include macroeconomic aggregates, prices and interest rates. This type of data is characterized by serial dependence so the random sampling assumption is inappropriate. Most aggregate economic data is only available at a low frequency (annual, quarterly or perhaps monthly) so the sample size is typically much smaller than in cross-section studies. The exception is financial data where data are available at a high frequency (weekly, daily, hourly, or by transaction) so sample sizes can be quite large.

Panel data combines elements of cross-section and time-series. These data sets consist of a set of individuals (typically persons, households, or corporations) surveyed repeatedly over time. The common modeling assumption is that the individuals are mutually independent of one another, but a given individual's observations are mutually dependent. This is a modified random sampling environment.

Data Structures
<ul style="list-style-type: none"> • Cross-section • Time-series • Panel

Many contemporary econometric applications combine elements of cross-section, time-series, and panel data modeling. These include models of spatial correlation and clustering.

As we mentioned above, most of this text will be devoted to cross-sectional data under the assumption of mutually independent observations. By mutual independence we mean that the i^{th} observation $(y_i, \mathbf{x}_i, \mathbf{z}_i)$ is independent of the j^{th} observation $(y_j, \mathbf{x}_j, \mathbf{z}_j)$ for $i \neq j$. (Sometimes the label “independent” is misconstrued. It is a statement about the relationship between observations i and j , not a statement about the relationship between y_i and \mathbf{x}_i and/or \mathbf{z}_i .)

Furthermore, if the data is randomly gathered, it is reasonable to model each observation as a random draw from the same probability distribution. In this case we say that the data are **independent and identically distributed** or **iid**. We call this a **random sample**. For most of this text we will assume that our observations come from a random sample.

Definition 1.5.1 *The observations $(y_i, \mathbf{x}_i, \mathbf{z}_i)$ are a **random sample** if they are mutually independent and identically distributed (**iid**) across $i = 1, \dots, n$.*

In the random sampling framework, we think of an individual observation $(y_i, \mathbf{x}_i, \mathbf{z}_i)$ as a realization from a joint probability distribution $F(y, \mathbf{x}, \mathbf{z})$ which we can call the **population**. This “population” is infinitely large. This abstraction can be a source of confusion as it does not correspond to a physical population in the real world. It is an abstraction since the distribution F is unknown, and the goal of statistical inference is to learn about features of F from the sample. The *assumption* of random sampling provides the mathematical foundation for treating economic statistics with the tools of mathematical statistics.

The random sampling framework was a major intellectual breakthrough of the late 19th century, allowing the application of mathematical statistics to the social sciences. Before this conceptual development, methods from mathematical statistics had not been applied to economic data as the latter was viewed as non-random. The random sampling framework enabled economic samples to be treated as random, a necessary precondition for the application of statistical methods.

1.6 Sources for Economic Data

Fortunately for economists, the internet provides a convenient forum for dissemination of economic data. Many large-scale economic datasets are available without charge from governmental agencies. An excellent starting point is the Resources for Economists Data Links, available at rfe.org. From this site you can find almost every publically available economic data set. Some specific data sources of interest include

- Bureau of Labor Statistics
- US Census

- Current Population Survey
- Survey of Income and Program Participation
- Panel Study of Income Dynamics
- Federal Reserve System (Board of Governors and regional banks)
- National Bureau of Economic Research
- U.S. Bureau of Economic Analysis
- CompuStat
- International Financial Statistics

Another good source of data is from authors of published empirical studies. Most journals in economics require authors of published papers to make their datasets generally available. For example, in its instructions for submission, *Econometrica* states:

Econometrica has the policy that all empirical, experimental and simulation results must be replicable. Therefore, authors of accepted papers must submit data sets, programs, and information on empirical analysis, experiments and simulations that are needed for replication and some limited sensitivity analysis.

The *American Economic Review* states:

All data used in analysis must be made available to any researcher for purposes of replication.

The *Journal of Political Economy* states:

It is the policy of the *Journal of Political Economy* to publish papers only if the data used in the analysis are clearly and precisely documented and are readily available to any researcher for purposes of replication.

If you are interested in using the data from a published paper, first check the journal's website, as many journals archive data and replication programs online. Second, check the website(s) of the paper's author(s). Most academic economists maintain webpages, and some make available replication files complete with data and programs. If these investigations fail, email the author(s), politely requesting the data. You may need to be persistent.

As a matter of professional etiquette, all authors absolutely have the obligation to make their data and programs available. Unfortunately, many fail to do so, and typically for poor reasons. The irony of the situation is that it is typically in the best interests of a scholar to make as much of their work (including all data and programs) freely available, as this only increases the likelihood of their work being cited and having an impact.

Keep this in mind as you start your own empirical project. Remember that as part of your end product, you will need (and want) to provide all data and programs to the community of scholars. The greatest form of flattery is to learn that another scholar has read your paper, wants to extend your work, or wants to use your empirical methods. In addition, public openness provides a healthy incentive for transparency and integrity in empirical analysis.

1.7 Econometric Software

Economists use a variety of econometric, statistical, and programming software.

STATA (www.stata.com) is a powerful statistical program with a broad set of pre-programmed econometric and statistical tools. It is quite popular among economists, and is continuously being updated with new methods. It is an excellent package for most econometric analysis, but is limited when you want to use new or less-common econometric methods which have not yet been programmed.

R (www.r-project.org), GAUSS (www.aptech.com), MATLAB (www.mathworks.com), and Ox (www.oxmetrics.net) are high-level matrix programming languages with a wide variety of built-in statistical functions. Many econometric methods have been programmed in these languages and are available on the web. The advantage of these packages is that you are in complete control of your analysis, and it is easier to program new methods than in STATA. Some disadvantages are that you have to do much of the programming yourself, programming complicated procedures takes significant time, and programming errors are hard to prevent and difficult to detect and eliminate. Of these languages, Gauss used to be quite popular among econometricians, but currently Matlab is more popular. A smaller but growing group of econometricians are enthusiastic fans of R, which of these languages is uniquely open-source, user-contributed, and best of all, completely free!

For highly-intensive computational tasks, some economists write their programs in a standard programming language such as Fortran or C. This can lead to major gains in computational speed, at the cost of increased time in programming and debugging.

As these different packages have distinct advantages, many empirical economists end up using more than one package. As a student of econometrics, you will learn at least one of these packages, and probably more than one.

1.8 Reading the Manuscript

I have endeavored to use a unified notation and nomenclature. The development of the material is cumulative, with later chapters building on the earlier ones. Never-the-less, every attempt has been made to make each chapter self-contained, so readers can pick and choose topics according to their interests.

To fully understand econometric methods, it is necessary to have a mathematical understanding of its mechanics, and this includes the mathematical proofs of the main results. Consequently, this text is self-contained, with nearly all results proved with full mathematical rigor. The mathematical development and proofs aim at brevity and conciseness (sometimes described as mathematical elegance), but also at pedagogy. To understand a mathematical proof, it is not sufficient to simply *read* the proof, you need to follow it, and re-create it for yourself.

Never-the-less, many readers will not be interested in each mathematical detail, explanation, or proof. This is okay. To use a method it may not be necessary to understand the mathematical details. Accordingly I have placed the more technical mathematical proofs and details in chapter appendices. These appendices and other technical sections are marked with an asterisk (*). These sections can be skipped without any loss in exposition.

1.9 Common Symbols

y	scalar
\mathbf{x}	vector
\mathbf{X}	matrix
\mathbb{R}	real line
\mathbb{R}^k	Euclidean k space
$\mathbb{E}(y)$	mathematical expectation
$\text{var}(y)$	variance
$\text{cov}(x, y)$	covariance
$\text{var}(\mathbf{x})$	covariance matrix
$\text{corr}(x, y)$	correlation
Pr	probability
\longrightarrow	limit
\xrightarrow{p}	convergence in probability
\xrightarrow{d}	convergence in distribution
$\text{plim}_{n \rightarrow \infty}$	probability limit
$N(\mu, \sigma^2)$	normal distribution
$N(0, 1)$	standard normal distribution
χ_k^2	chi-square distribution with k degrees of freedom
\mathbf{I}_n	identity matrix
$\text{tr } \mathbf{A}$	trace
\mathbf{A}'	matrix transpose
\mathbf{A}^{-1}	matrix inverse
$\mathbf{A} > 0, \mathbf{A} \geq 0$	positive definite, positive semi-definite
$\ \mathbf{a}\ $	Euclidean norm
$\ \mathbf{A}\ $	matrix (Frobenius) norm
\approx	approximate equality
$\stackrel{def}{=}$	definitional equality
\sim	is distributed as
\log	natural logarithm

Chapter 2

Conditional Expectation and Projection

2.1 Introduction

The most commonly applied econometric tool is least-squares estimation, also known as **regression**. As we will see, least-squares is a tool to estimate an approximate conditional mean of one variable (the **dependent variable**) given another set of variables (the **regressors, conditioning variables, or covariates**).

In this chapter we abstract from estimation, and focus on the probabilistic foundation of the conditional expectation model and its projection approximation.

2.2 The Distribution of Wages

Suppose that we are interested in wage rates in the United States. Since wage rates vary across workers, we cannot describe wage rates by a single number. Instead, we can describe wages using a probability distribution. Formally, we view the wage of an individual worker as a random variable *wage* with the **probability distribution**

$$F(u) = \Pr(\text{wage} \leq u).$$

When we say that a person's wage is random we mean that we do not know their wage before it is measured, and we treat observed wage rates as realizations from the distribution F . Treating unobserved wages as random variables and observed wages as realizations is a powerful mathematical abstraction which allows us to use the tools of mathematical probability.

A useful thought experiment is to imagine dialing a telephone number selected at random, and then asking the person who responds to tell us their wage rate. (Assume for simplicity that all workers have equal access to telephones, and that the person who answers your call will respond honestly.) In this thought experiment, the wage of the person you have called is a single draw from the distribution F of wages in the population. By making many such phone calls we can learn the distribution F of the entire population.

When a distribution function F is differentiable we define the probability density function

$$f(u) = \frac{d}{du}F(u).$$

The density contains the same information as the distribution function, but the density is typically easier to visually interpret.

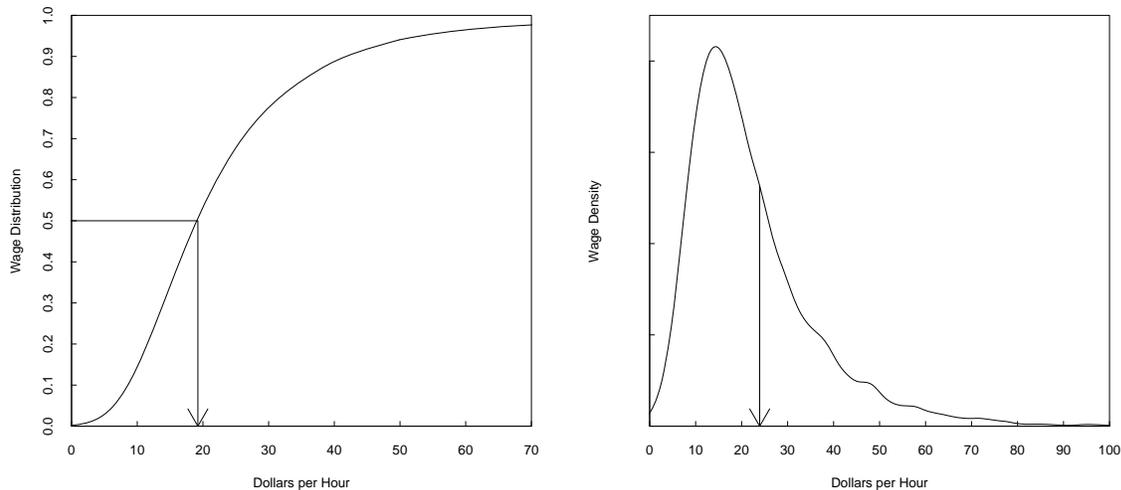


Figure 2.1: Wage Distribution and Density. All full-time U.S. workers

In Figure 2.1 we display estimates¹ of the probability distribution function (on the left) and density function (on the right) of U.S. wage rates in 2009. We see that the density is peaked around \$15, and most of the probability mass appears to lie between \$10 and \$40. These are ranges for typical wage rates in the U.S. population.

Important measures of central tendency are the median and the mean. The **median** m of a continuous² distribution F is the unique solution to

$$F(m) = \frac{1}{2}.$$

The median U.S. wage (\$19.23) is indicated in the left panel of Figure 2.1 by the arrow. The median is a robust³ measure of central tendency, but it is tricky to use for many calculations as it is not a linear operator.

The **expectation** or **mean** of a random variable y with density f is

$$\mu = \mathbb{E}(y) = \int_{-\infty}^{\infty} uf(u)du.$$

Here we have used the common and convenient convention of using the single character y to denote a random variable, rather than the more cumbersome label *wage*. A general definition of the mean is presented in Section 2.31. The mean U.S. wage (\$23.90) is indicated in the right panel of Figure 2.1 by the arrow.

We sometimes use the notation $\mathbb{E}y$ instead of $\mathbb{E}(y)$ when the variable whose expectation is being taken is clear from the context. There is no distinction in meaning.

The mean is a convenient measure of central tendency because it is a linear operator and arises naturally in many economic models. A disadvantage of the mean is that it is not robust⁴ especially in the presence of substantial skewness or thick tails, which are both features of the wage

¹The distribution and density are estimated nonparametrically from the sample of 50,742 full-time non-military wage-earners reported in the March 2009 Current Population Survey. The wage rate is constructed as annual individual wage and salary earnings divided by hours worked.

²If F is not continuous the definition is $m = \inf\{u : F(u) \geq \frac{1}{2}\}$

³The median is not sensitive to perturbations in the tails of the distribution.

⁴The mean is sensitive to perturbations in the tails of the distribution.

distribution as can be seen easily in the right panel of Figure 2.1. Another way of viewing this is that 64% of workers earn less than the mean wage of \$23.90, suggesting that it is incorrect to describe the mean as a “typical” wage rate.

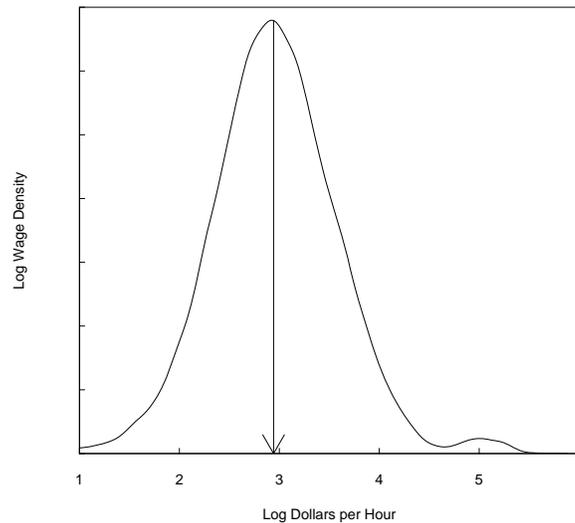


Figure 2.2: Log Wage Density

In this context it is useful to transform the data by taking the natural logarithm⁵. Figure 2.2 shows the density of log hourly wages $\log(\text{wage})$ for the same population, with its mean 2.95 drawn in with the arrow. The density of log wages is much less skewed and fat-tailed than the density of the level of wages, so its mean

$$\mathbb{E}(\log(\text{wage})) = 2.95$$

is a much better (more robust) measure⁶ of central tendency of the distribution. For this reason, wage regressions typically use log wages as a dependent variable rather than the level of wages.

Another useful way to summarize the probability distribution $F(u)$ is in terms of its quantiles. For any $\alpha \in (0, 1)$, the α^{th} quantile of the continuous⁷ distribution F is the real number q_α which satisfies

$$F(q_\alpha) = \alpha.$$

The quantile function q_α , viewed as a function of α , is the inverse of the distribution function F . The most commonly used quantile is the median, that is, $q_{0.5} = m$. We sometimes refer to quantiles by the percentile representation of α , and in this case they are often called percentiles, e.g. the median is the 50th percentile.

2.3 Conditional Expectation

We saw in Figure 2.2 the density of log wages. Is this distribution the same for all workers, or does the wage distribution vary across subpopulations? To answer this question, we can compare wage distributions for different groups – for example, men and women. The plot on the left in Figure 2.3 displays the densities of log wages for U.S. men and women with their means (3.05 and 2.81) indicated by the arrows. We can see that the two wage densities take similar shapes but the density for men is somewhat shifted to the right with a higher mean.

⁵Throughout the text, we will use $\log(y)$ or $\log y$ to denote the natural logarithm of y .

⁶More precisely, the geometric mean $\exp(\mathbb{E}(\log w)) = \19.11 is a robust measure of central tendency.

⁷If F is not continuous the definition is $q_\alpha = \inf\{u : F(u) \geq \alpha\}$

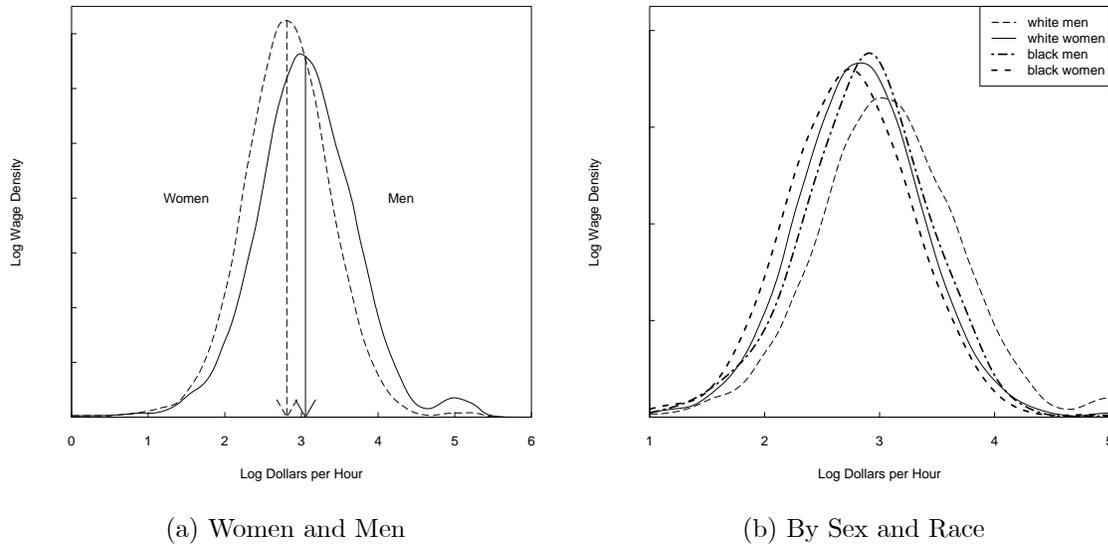


Figure 2.3: Log Wage Density by Sex and Race

The values 3.05 and 2.81 are the mean log wages in the subpopulations of men and women workers. They are called the **conditional means** (or **conditional expectations**) of log wages given sex. We can write their specific values as

$$\mathbb{E}(\log(\text{wage}) \mid \text{sex} = \text{man}) = 3.05 \quad (2.1)$$

$$\mathbb{E}(\log(\text{wage}) \mid \text{sex} = \text{woman}) = 2.81. \quad (2.2)$$

We call these means *conditional* as they are conditioning on a fixed value of the variable *sex*. While you might not think of a person's sex as a random variable, it is random from the viewpoint of econometric analysis. If you randomly select an individual, the sex of the individual is unknown and thus random. (In the population of U.S. workers, the probability that a worker is a woman happens to be 43%.) In observational data, it is most appropriate to view all measurements as random variables, and the means of subpopulations are then conditional means.

As the two densities in Figure 2.3 appear similar, a hasty inference might be that there is not a meaningful difference between the wage distributions of men and women. Before jumping to this conclusion let us examine the differences in the distributions of Figure 2.3 more carefully. As we mentioned above, the primary difference between the two densities appears to be their means. This difference equals

$$\begin{aligned} \mathbb{E}(\log(\text{wage}) \mid \text{sex} = \text{man}) - \mathbb{E}(\log(\text{wage}) \mid \text{sex} = \text{woman}) &= 3.05 - 2.81 \\ &= 0.24 \end{aligned} \quad (2.3)$$

A difference in expected log wages of 0.24 implies an average 24% difference between the wages of men and women, which is quite substantial. (For an explanation of logarithmic and percentage differences see Section 2.4.)

Consider further splitting the men and women subpopulations by race, dividing the population into whites, blacks, and other races. We display the log wage density functions of four of these groups on the right in Figure 2.3. Again we see that the primary difference between the four density functions is their central tendency.

	men	women
white	3.07	2.82
black	2.86	2.73
other	3.03	2.86

Table 2.1: Mean Log Wages by Sex and Race

Focusing on the means of these distributions, Table 2.1 reports the mean log wage for each of the six sub-populations.

The entries in Table 2.1 are the conditional means of $\log(\text{wage})$ given *sex* and *race*. For example

$$\mathbb{E}(\log(\text{wage}) \mid \text{sex} = \text{man}, \text{race} = \text{white}) = 3.07$$

and

$$\mathbb{E}(\log(\text{wage}) \mid \text{sex} = \text{woman}, \text{race} = \text{black}) = 2.73$$

One benefit of focusing on conditional means is that they reduce complicated distributions to a single summary measure, and thereby facilitate comparisons across groups. Because of this simplifying property, conditional means are the primary interest of regression analysis and are a major focus in econometrics.

Table 2.1 allows us to easily calculate average wage differences between groups. For example, we can see that the wage gap between men and women continues after disaggregation by race, as the average gap between white men and white women is 25%, and that between black men and black women is 13%. We also can see that there is a race gap, as the average wages of blacks are substantially less than the other race categories. In particular, the average wage gap between white men and black men is 21%, and that between white women and black women is 9%.

2.4 Log Differences*

A useful approximation for the natural logarithm for small x is

$$\log(1+x) \approx x. \tag{2.4}$$

This can be derived from the infinite series expansion of $\log(1+x)$:

$$\begin{aligned} \log(1+x) &= x - \frac{x^2}{2} + \frac{x^3}{3} - \frac{x^4}{4} + \cdots \\ &= x + O(x^2). \end{aligned}$$

The symbol $O(x^2)$ means that the remainder is bounded by Ax^2 as $x \rightarrow 0$ for some $A < \infty$. A plot of $\log(1+x)$ and the linear approximation x is shown in Figure 2.4. We can see that $\log(1+x)$ and the linear approximation x are very close for $|x| \leq 0.1$, and reasonably close for $|x| \leq 0.2$, but the difference increases with $|x|$.

Now, if y^* is $c\%$ greater than y , then

$$y^* = (1 + c/100)y.$$

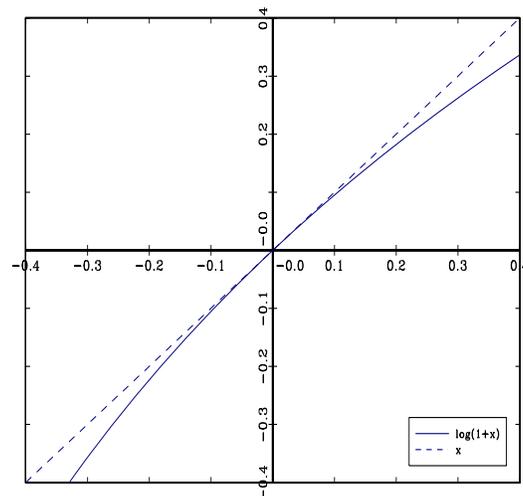
Taking natural logarithms,

$$\log y^* = \log y + \log(1 + c/100)$$

or

$$\log y^* - \log y = \log(1 + c/100) \approx \frac{c}{100}$$

where the approximation is (2.4). This shows that 100 multiplied by the difference in logarithms is approximately the percentage difference between y and y^* , and this approximation is quite good for $|c| \leq 10$.

Figure 2.4: $\log(1+x)$

2.5 Conditional Expectation Function

An important determinant of wage levels is education. In many empirical studies economists measure educational attainment by the number of years of schooling, and we will write this variable as *education*⁸.

The conditional mean of log wages given *sex*, *race*, and *education* is a single number for each category. For example

$$\mathbb{E}(\log(\text{wage}) \mid \text{sex} = \text{man}, \text{race} = \text{white}, \text{education} = 12) = 2.84$$

We display in Figure 2.5 the conditional means of $\log(\text{wage})$ for white men and white women as a function of *education*. The plot is quite revealing. We see that the conditional mean is increasing in years of education, but at a different rate for schooling levels above and below nine years. Another striking feature of Figure 2.5 is that the gap between men and women is roughly constant for all education levels. As the variables are measured in logs this implies a constant average percentage gap between men and women regardless of educational attainment.

In many cases it is convenient to simplify the notation by writing variables using single characters, typically y , x and/or z . It is conventional in econometrics to denote the dependent variable (e.g. $\log(\text{wage})$) by the letter y , a conditioning variable (such as *sex*) by the letter x , and multiple conditioning variables (such as *race*, *education* and *sex*) by the subscripted letters x_1, x_2, \dots, x_k .

Conditional expectations can be written with the generic notation

$$\mathbb{E}(y \mid x_1, x_2, \dots, x_k) = m(x_1, x_2, \dots, x_k).$$

We call this the **conditional expectation function** (CEF). The CEF is a function of (x_1, x_2, \dots, x_k) as it varies with the variables. For example, the conditional expectation of $y = \log(\text{wage})$ given $(x_1, x_2) = (\text{sex}, \text{race})$ is given by the six entries of Table 2.1. The CEF is a function of $(\text{gender}, \text{race})$ as it varies across the entries.

For greater compactness, we will typically write the conditioning variables as a vector in \mathbb{R}^k :

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_k \end{pmatrix}. \quad (2.5)$$

⁸Here, *education* is defined as years of schooling beyond kindergarten. A high school graduate has *education*=12, a college graduate has *education*=16, a Master's degree has *education*=18, and a professional degree (medical, law or PhD) has *education*=20.

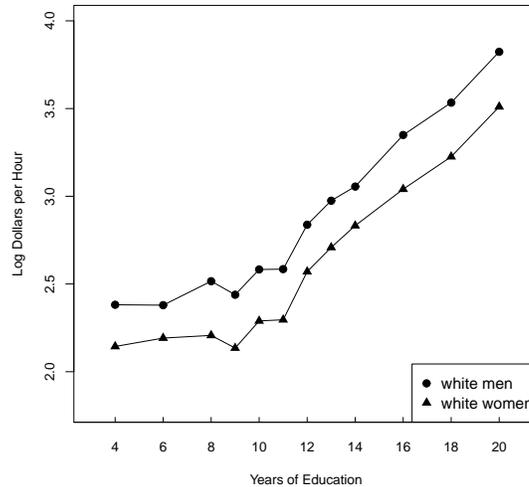


Figure 2.5: Mean Log Wage as a Function of Years of Education

Here we follow the convention of using lower case bold italics \mathbf{x} to denote a vector. Given this notation, the CEF can be compactly written as

$$\mathbb{E}(y \mid \mathbf{x}) = m(\mathbf{x}).$$

The CEF $\mathbb{E}(y \mid \mathbf{x})$ is a random variable as it is a function of the random variable \mathbf{x} . It is also sometimes useful to view the CEF as a function of \mathbf{x} . In this case we can write $m(\mathbf{u}) = \mathbb{E}(y \mid \mathbf{x} = \mathbf{u})$, which is a function of the argument \mathbf{u} . The expression $\mathbb{E}(y \mid \mathbf{x} = \mathbf{u})$ is the conditional expectation of y , given that we know that the random variable \mathbf{x} equals the specific value \mathbf{u} . However, sometimes in econometrics we take a notational shortcut and use $\mathbb{E}(y \mid \mathbf{x})$ to refer to this function. Hopefully, the use of $\mathbb{E}(y \mid \mathbf{x})$ should be apparent from the context.

2.6 Continuous Variables

In the previous sections, we implicitly assumed that the conditioning variables are discrete. However, many conditioning variables are continuous. In this section, we take up this case and assume that the variables (y, \mathbf{x}) are continuously distributed with a joint density function $f(y, \mathbf{x})$.

As an example, take $y = \log(\text{wage})$ and $x = \text{experience}$, the number of years of potential labor market experience⁹. The contours of their joint density are plotted on the left side of Figure 2.6 for the population of white men with 12 years of education.

Given the joint density $f(y, \mathbf{x})$ the variable \mathbf{x} has the marginal density

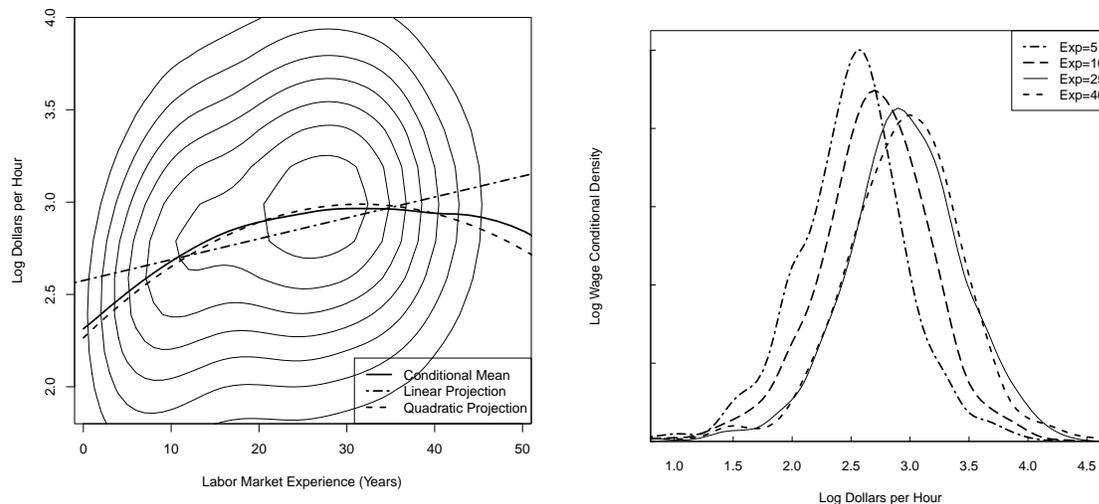
$$f_{\mathbf{x}}(\mathbf{x}) = \int_{\mathbb{R}} f(y, \mathbf{x}) dy.$$

For any \mathbf{x} such that $f_{\mathbf{x}}(\mathbf{x}) > 0$ the conditional density of y given \mathbf{x} is defined as

$$f_{y|\mathbf{x}}(y \mid \mathbf{x}) = \frac{f(y, \mathbf{x})}{f_{\mathbf{x}}(\mathbf{x})}. \quad (2.6)$$

The conditional density is a slice of the joint density $f(y, \mathbf{x})$ holding \mathbf{x} fixed. We can visualize this by slicing the joint density function at a specific value of \mathbf{x} parallel with the y -axis. For example,

⁹Here, *experience* is defined as potential labor market experience, equal to $\text{age} - \text{education} - 6$

(a) Joint density of $\log(\text{wage})$ and experience and conditional mean

(b) Conditional density

Figure 2.6: White men with $\text{education}=12$

take the density contours on the left side of Figure 2.6 and slice through the contour plot at a specific value of experience . This gives us the conditional density of $\log(\text{wage})$ for white men with 12 years of education and this level of experience. We do this for four levels of experience (5, 10, 25, and 40 years), and plot these densities on the right side of Figure 2.6. We can see that the distribution of wages shifts to the right and becomes more diffuse as experience increases from 5 to 10 years, and from 10 to 25 years experience, but there is little change from 25 to 40 years experience.

The CEF of y given \mathbf{x} is the mean of the conditional density (2.6)

$$m(\mathbf{x}) = \mathbb{E}(y | \mathbf{x}) = \int_{\mathbb{R}} y f_{y|\mathbf{x}}(y | \mathbf{x}) dy. \quad (2.7)$$

Intuitively, $m(\mathbf{x})$ is the mean of y for the idealized subpopulation where the conditioning variables are fixed at \mathbf{x} . This is idealized since \mathbf{x} is continuously distributed so this subpopulation is infinitely small.

In Figure 2.6 the CEF of $\log(\text{wage})$ given experience is plotted as the solid line. We can see that the CEF is a smooth but nonlinear function. The CEF is initially increasing in experience , flattens out around $\text{experience} = 30$, and then decreases for high levels of experience.

2.7 Law of Iterated Expectations

An extremely useful tool from probability theory is the **law of iterated expectations**. An important special case is the known as the Simple Law.

Theorem 2.7.1 Simple Law of Iterated Expectations

If $\mathbb{E}|y| < \infty$ then for any random vector \mathbf{x} ,

$$\mathbb{E}(\mathbb{E}(y | \mathbf{x})) = \mathbb{E}(y)$$

The simple law states that the expectation of the conditional expectation is the unconditional expectation. In other words, the average of the conditional averages is the unconditional average. When \mathbf{x} is discrete

$$\mathbb{E}(\mathbb{E}(y | \mathbf{x})) = \sum_{j=1}^{\infty} \mathbb{E}(y | \mathbf{x}_j) \Pr(\mathbf{x} = \mathbf{x}_j)$$

and when \mathbf{x} is continuous

$$\mathbb{E}(\mathbb{E}(y | \mathbf{x})) = \int_{\mathbb{R}^k} \mathbb{E}(y | \mathbf{x}) f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}.$$

Going back to our investigation of average log wages for men and women, the simple law states that

$$\begin{aligned} & \mathbb{E}(\log(\text{wage}) | \text{gender} = \text{man}) \Pr(\text{gender} = \text{man}) \\ & + \mathbb{E}(\log(\text{wage}) | \text{gender} = \text{woman}) \Pr(\text{gender} = \text{woman}) \\ & = \mathbb{E}(\log(\text{wage})). \end{aligned}$$

Or numerically,

$$3.05 \times 0.57 + 2.79 \times 0.43 = 2.92.$$

The general law of iterated expectations allows two sets of conditioning variables.

Theorem 2.7.2 Law of Iterated Expectations

If $\mathbb{E}|y| < \infty$ then for any random vectors \mathbf{x}_1 and \mathbf{x}_2 ,

$$\mathbb{E}(\mathbb{E}(y | \mathbf{x}_1, \mathbf{x}_2) | \mathbf{x}_1) = \mathbb{E}(y | \mathbf{x}_1)$$

Notice the way the law is applied. The inner expectation conditions on \mathbf{x}_1 and \mathbf{x}_2 , while the outer expectation conditions only on \mathbf{x}_1 . The iterated expectation yields the simple answer $\mathbb{E}(y | \mathbf{x}_1)$, the expectation conditional on \mathbf{x}_1 alone. Sometimes we phrase this as: “The smaller information set wins.”

As an example

$$\begin{aligned} & \mathbb{E}(\log(\text{wage}) | \text{gender} = \text{man}, \text{race} = \text{white}) \Pr(\text{race} = \text{white} | \text{gender} = \text{man}) \\ & + \mathbb{E}(\log(\text{wage}) | \text{gender} = \text{man}, \text{race} = \text{black}) \Pr(\text{race} = \text{black} | \text{gender} = \text{man}) \\ & + \mathbb{E}(\log(\text{wage}) | \text{gender} = \text{man}, \text{race} = \text{other}) \Pr(\text{race} = \text{other} | \text{gender} = \text{man}) \\ & = \mathbb{E}(\log(\text{wage}) | \text{gender} = \text{man}) \end{aligned}$$

or numerically

$$3.07 \times 0.84 + 2.86 \times 0.08 + 3.03 \times 0.08 = 3.05.$$

A property of conditional expectations is that when you condition on a random vector \mathbf{x} you can effectively treat it as if it is constant. For example, $\mathbb{E}(\mathbf{x} | \mathbf{x}) = \mathbf{x}$ and $\mathbb{E}(g(\mathbf{x}) | \mathbf{x}) = g(\mathbf{x})$ for any function $g(\cdot)$. The general property is known as the Conditioning Theorem.

Theorem 2.7.3 Conditioning Theorem

If

$$\mathbb{E}|g(\mathbf{x})y| < \infty \tag{2.8}$$

then

$$\mathbb{E}(g(\mathbf{x})y | \mathbf{x}) = g(\mathbf{x}) \mathbb{E}(y | \mathbf{x}) \tag{2.9}$$

and

$$\mathbb{E}(g(\mathbf{x})y) = \mathbb{E}(g(\mathbf{x}) \mathbb{E}(y | \mathbf{x})). \tag{2.10}$$

The proofs of Theorems 2.7.1, 2.7.2 and 2.7.3 are given in Section 2.34.

2.8 CEF Error

The CEF error e is defined as the difference between y and the CEF evaluated at the random vector \mathbf{x} :

$$e = y - m(\mathbf{x}).$$

By construction, this yields the formula

$$y = m(\mathbf{x}) + e. \quad (2.11)$$

In (2.11) it is useful to understand that the error e is derived from the joint distribution of (y, \mathbf{x}) , and so its properties are derived from this construction.

A key property of the CEF error is that it has a conditional mean of zero. To see this, by the linearity of expectations, the definition $m(\mathbf{x}) = \mathbb{E}(y | \mathbf{x})$ and the Conditioning Theorem

$$\begin{aligned} \mathbb{E}(e | \mathbf{x}) &= \mathbb{E}((y - m(\mathbf{x})) | \mathbf{x}) \\ &= \mathbb{E}(y | \mathbf{x}) - \mathbb{E}(m(\mathbf{x}) | \mathbf{x}) \\ &= m(\mathbf{x}) - m(\mathbf{x}) \\ &= 0. \end{aligned}$$

This fact can be combined with the law of iterated expectations to show that the unconditional mean is also zero.

$$\mathbb{E}(e) = \mathbb{E}(\mathbb{E}(e | \mathbf{x})) = \mathbb{E}(0) = 0.$$

We state this and some other results formally.

Theorem 2.8.1 *Properties of the CEF error*

If $\mathbb{E}|y| < \infty$ then

1. $\mathbb{E}(e | \mathbf{x}) = 0$.
2. $\mathbb{E}(e) = 0$.
3. If $\mathbb{E}|y|^r < \infty$ for $r \geq 1$ then $\mathbb{E}|e|^r < \infty$.
4. For any function $h(\mathbf{x})$ such that $\mathbb{E}|h(\mathbf{x})e| < \infty$ then $\mathbb{E}(h(\mathbf{x})e) = 0$.

The proof of the third result is deferred to Section 2.34.

The fourth result, whose proof is left to Exercise 2.3, implies that e is uncorrelated with any function of the regressors.

The equations

$$\begin{aligned} y &= m(\mathbf{x}) + e \\ \mathbb{E}(e | \mathbf{x}) &= 0 \end{aligned}$$

together imply that $m(\mathbf{x})$ is the CEF of y given \mathbf{x} . It is important to understand that this is not a restriction. These equations hold true by definition.

The condition $\mathbb{E}(e | \mathbf{x}) = 0$ is implied by the definition of e as the difference between y and the CEF $m(\mathbf{x})$. The equation $\mathbb{E}(e | \mathbf{x}) = 0$ is sometimes called a conditional mean restriction, since the conditional mean of the error e is restricted to equal zero. The property is also sometimes called **mean independence**, for the conditional mean of e is 0 and thus independent of \mathbf{x} . However, it does not imply that the distribution of e is independent of \mathbf{x} . Sometimes the assumption “ e is

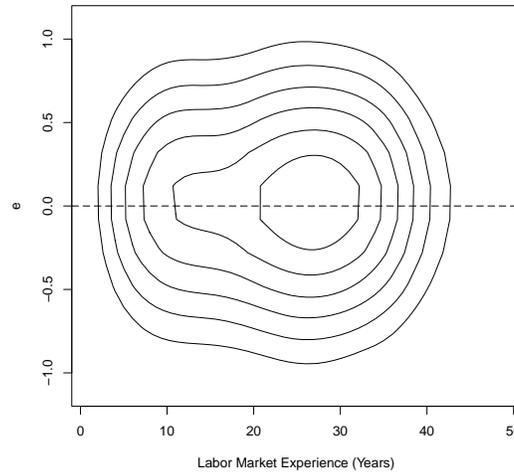


Figure 2.7: Joint density of CEF error e and $experience$ for white men with $education=12$.

independent of \mathbf{x} ” is added as a convenient simplification, but it is not generic feature of the conditional mean. Typically and generally, e and \mathbf{x} are jointly dependent, even though the conditional mean of e is zero.

As an example, the contours of the joint density of e and $experience$ are plotted in Figure 2.7 for the same population as Figure 2.6. The error e has a conditional mean of zero for all values of $experience$, but the shape of the conditional distribution varies with the level of $experience$.

As a simple example of a case where x and e are mean independent yet dependent, let $e = x\varepsilon$ where x and ε are independent $N(0, 1)$. Then conditional on x , the error e has the distribution $N(0, x^2)$. Thus $\mathbb{E}(e | x) = 0$ and e is mean independent of x , yet e is not fully independent of x . Mean independence does not imply full independence.

2.9 Intercept-Only Model

A special case of the regression model is when there are no regressors \mathbf{x} . In this case $m(\mathbf{x}) = \mathbb{E}(y) = \mu$, the unconditional mean of y . We can still write an equation for y in the regression format:

$$y = \mu + e$$

$$\mathbb{E}(e) = 0$$

This is useful for it unifies the notation.

2.10 Regression Variance

An important measure of the dispersion about the CEF function is the unconditional variance of the CEF error e . We write this as

$$\sigma^2 = \text{var}(e) = \mathbb{E}\left((e - \mathbb{E}e)^2\right) = \mathbb{E}(e^2).$$

Theorem 2.8.1.3 implies the following simple but useful result.

Theorem 2.10.1 *If $\mathbb{E}y^2 < \infty$ then $\sigma^2 < \infty$.*

We can call σ^2 the regression variance or the variance of the regression error. The magnitude of σ^2 measures the amount of variation in y which is not “explained” or accounted for in the conditional mean $\mathbb{E}(y | \mathbf{x})$.

The regression variance depends on the regressors \mathbf{x} . Consider two regressions

$$\begin{aligned} y &= \mathbb{E}(y | \mathbf{x}_1) + e_1 \\ y &= \mathbb{E}(y | \mathbf{x}_1, \mathbf{x}_2) + e_2. \end{aligned}$$

We write the two errors distinctly as e_1 and e_2 as they are different – changing the conditioning information changes the conditional mean and therefore the regression error as well.

In our discussion of iterated expectations, we have seen that by increasing the conditioning set, the conditional expectation reveals greater detail about the distribution of y . What is the implication for the regression error?

It turns out that there is a simple relationship. We can think of the conditional mean $\mathbb{E}(y | \mathbf{x})$ as the “explained portion” of y . The remainder $e = y - \mathbb{E}(y | \mathbf{x})$ is the “unexplained portion”. The simple relationship we now derive shows that the variance of this unexplained portion decreases when we condition on more variables. This relationship is monotonic in the sense that increasing the amount of information always decreases the variance of the unexplained portion.

Theorem 2.10.2 *If $\mathbb{E}y^2 < \infty$ then*

$$\text{var}(y) \geq \text{var}(y - \mathbb{E}(y | \mathbf{x}_1)) \geq \text{var}(y - \mathbb{E}(y | \mathbf{x}_1, \mathbf{x}_2)).$$

Theorem 2.10.2 says that the variance of the difference between y and its conditional mean (weakly) decreases whenever an additional variable is added to the conditioning information.

The proof of Theorem 2.10.2 is given in Section 2.34.

2.11 Best Predictor

Suppose that given a realized value of \mathbf{x} , we want to create a prediction or forecast of y . We can write any predictor as a function $g(\mathbf{x})$ of \mathbf{x} . The prediction error is the realized difference $y - g(\mathbf{x})$. A non-stochastic measure of the magnitude of the prediction error is the expectation of its square

$$\mathbb{E}(y - g(\mathbf{x}))^2. \tag{2.12}$$

We can define the best predictor as the function $g(\mathbf{x})$ which minimizes (2.12). What function is the best predictor? It turns out that the answer is the CEF $m(\mathbf{x})$. This holds regardless of the joint distribution of (y, \mathbf{x}) .

To see this, note that the mean squared error of a predictor $g(\mathbf{x})$ is

$$\begin{aligned} \mathbb{E}(y - g(\mathbf{x}))^2 &= \mathbb{E}(e + m(\mathbf{x}) - g(\mathbf{x}))^2 \\ &= \mathbb{E}e^2 + 2\mathbb{E}(e(m(\mathbf{x}) - g(\mathbf{x}))) + \mathbb{E}(m(\mathbf{x}) - g(\mathbf{x}))^2 \\ &= \mathbb{E}e^2 + \mathbb{E}(m(\mathbf{x}) - g(\mathbf{x}))^2 \\ &\geq \mathbb{E}e^2 \\ &= \mathbb{E}(y - m(\mathbf{x}))^2 \end{aligned}$$

where the first equality makes the substitution $y = m(\mathbf{x}) + e$ and the third equality uses Theorem 2.8.1.4. The right-hand-side after the third equality is minimized by setting $g(\mathbf{x}) = m(\mathbf{x})$, yielding

the inequality in the fourth line. The minimum is finite under the assumption $\mathbb{E}y^2 < \infty$ as shown by Theorem 2.10.1.

We state this formally in the following result.

Theorem 2.11.1 *Conditional Mean as Best Predictor*

If $\mathbb{E}y^2 < \infty$, then for any predictor $g(\mathbf{x})$,

$$\mathbb{E}(y - g(\mathbf{x}))^2 \geq \mathbb{E}(y - m(\mathbf{x}))^2$$

where $m(\mathbf{x}) = \mathbb{E}(y | \mathbf{x})$.

It may be helpful to consider this result in the context of the intercept-only model

$$\begin{aligned} y &= \mu + e \\ \mathbb{E}(e) &= 0. \end{aligned}$$

Theorem 2.11.1 shows that the best predictor for y (in the class of constants) is the unconditional mean $\mu = \mathbb{E}(y)$, in the sense that the mean minimizes the mean squared prediction error.

2.12 Conditional Variance

While the conditional mean is a good measure of the location of a conditional distribution, it does not provide information about the spread of the distribution. A common measure of the dispersion is the **conditional variance**.

Definition 2.12.1 If $\mathbb{E}y^2 < \infty$, the **conditional variance** of y given \mathbf{x} is

$$\begin{aligned} \sigma^2(\mathbf{x}) &= \text{var}(y | \mathbf{x}) \\ &= \mathbb{E}\left((y - \mathbb{E}(y | \mathbf{x}))^2 | \mathbf{x}\right) \\ &= \mathbb{E}(e^2 | \mathbf{x}). \end{aligned}$$

Generally, $\sigma^2(\mathbf{x})$ is a non-trivial function of \mathbf{x} and can take any form subject to the restriction that it is non-negative. One way to think about $\sigma^2(\mathbf{x})$ is that it is the conditional mean of e^2 given \mathbf{x} .

The variance is in a different unit of measurement than the original variable. To convert the variance back to the same unit of measure we define the **conditional standard deviation** as its square root $\sigma(\mathbf{x}) = \sqrt{\sigma^2(\mathbf{x})}$.

As an example of how the conditional variance depends on observables, compare the conditional log wage densities for men and women displayed in Figure 2.3. The difference between the densities is not purely a location shift, but is also a difference in spread. Specifically, we can see that the density for men's log wages is somewhat more spread out than that for women, while the density for women's wages is somewhat more peaked. Indeed, the conditional standard deviation for men's wages is 3.05 and that for women is 2.81. So while men have higher average wages, they are also somewhat more dispersed.

The unconditional error variance and the conditional variance are related by the law of iterated expectations

$$\sigma^2 = \mathbb{E}(e^2) = \mathbb{E}(\mathbb{E}(e^2 | \mathbf{x})) = \mathbb{E}(\sigma^2(\mathbf{x})).$$

That is, the unconditional error variance is the average conditional variance.

Given the conditional variance, we can define a rescaled error

$$\varepsilon = \frac{e}{\sigma(\mathbf{x})}. \quad (2.13)$$

We can calculate that since $\sigma(\mathbf{x})$ is a function of \mathbf{x}

$$\mathbb{E}(\varepsilon | \mathbf{x}) = \mathbb{E}\left(\frac{e}{\sigma(\mathbf{x})} | \mathbf{x}\right) = \frac{1}{\sigma(\mathbf{x})}\mathbb{E}(e | \mathbf{x}) = 0$$

and

$$\text{var}(\varepsilon | \mathbf{x}) = \mathbb{E}(\varepsilon^2 | \mathbf{x}) = \mathbb{E}\left(\frac{e^2}{\sigma^2(\mathbf{x})} | \mathbf{x}\right) = \frac{1}{\sigma^2(\mathbf{x})}\mathbb{E}(e^2 | \mathbf{x}) = \frac{\sigma^2(\mathbf{x})}{\sigma^2(\mathbf{x})} = 1.$$

Thus ε has a conditional mean of zero, and a conditional variance of 1.

Notice that (2.13) can be rewritten as

$$e = \sigma(\mathbf{x})\varepsilon.$$

and substituting this for e in the CEF equation (2.11), we find that

$$y = m(\mathbf{x}) + \sigma(\mathbf{x})\varepsilon. \quad (2.14)$$

This is an alternative (mean-variance) representation of the CEF equation.

Many econometric studies focus on the conditional mean $m(\mathbf{x})$ and either ignore the conditional variance $\sigma^2(\mathbf{x})$, treat it as a constant $\sigma^2(\mathbf{x}) = \sigma^2$, or treat it as a nuisance parameter (a parameter not of primary interest). This is appropriate when the primary variation in the conditional distribution is in the mean, but can be short-sighted in other cases. Dispersion is relevant to many economic topics, including income and wealth distribution, economic inequality, and price dispersion. Conditional dispersion (variance) can be a fruitful subject for investigation.

The perverse consequences of a narrow-minded focus on the mean has been parodied in a classic joke:

An economist was standing with one foot in a bucket of boiling water and the other foot in a bucket of ice. When asked how he felt, he replied, “On average I feel just fine.”

Clearly, the economist in question ignored variance!

2.13 Homoskedasticity and Heteroskedasticity

An important special case obtains when the conditional variance $\sigma^2(\mathbf{x})$ is a constant and independent of \mathbf{x} . This is called **homoskedasticity**.

Definition 2.13.1 *The error is **homoskedastic** if $\mathbb{E}(e^2 | \mathbf{x}) = \sigma^2$ does not depend on \mathbf{x} .*

In the general case where $\sigma^2(\mathbf{x})$ depends on \mathbf{x} we say that the error e is **heteroskedastic**.

Definition 2.13.2 *The error is **heteroskedastic** if $\mathbb{E}(e^2 | \mathbf{x}) = \sigma^2(\mathbf{x})$ depends on \mathbf{x} .*

It is helpful to understand that the concepts homoskedasticity and heteroskedasticity concern the conditional variance, not the unconditional variance. By definition, the unconditional variance σ^2 is a constant and independent of the regressors \mathbf{x} . So when we talk about the variance as a function of the regressors, we are talking about the conditional variance $\sigma^2(\mathbf{x})$.

Some older or introductory textbooks describe heteroskedasticity as the case where “the variance of e varies across observations”. This is a poor and confusing definition. It is more constructive to understand that heteroskedasticity means that the conditional variance $\sigma^2(\mathbf{x})$ depends on observables.

Older textbooks also tend to describe homoskedasticity as a component of a correct regression specification, and describe heteroskedasticity as an exception or deviance. This description has influenced many generations of economists, but it is unfortunately backwards. The correct view is that heteroskedasticity is generic and “standard”, while homoskedasticity is unusual and exceptional. The default in empirical work should be to assume that the errors are heteroskedastic, not the converse.

In apparent contradiction to the above statement, we will still frequently impose the homoskedasticity assumption when making theoretical investigations into the properties of estimation and inference methods. The reason is that in many cases homoskedasticity greatly simplifies the theoretical calculations, and it is therefore quite advantageous for teaching and learning. It should always be remembered, however, that homoskedasticity is never imposed because it is believed to be a correct feature of an empirical model, but rather because of its simplicity.

2.14 Regression Derivative

One way to interpret the CEF $m(\mathbf{x}) = \mathbb{E}(y | \mathbf{x})$ is in terms of how marginal changes in the regressors \mathbf{x} imply changes in the conditional mean of the response variable y . It is typical to consider marginal changes in a single regressor, say x_1 , holding the remainder fixed. When a regressor x_1 is continuously distributed, we define the marginal effect of a change in x_1 , holding the variables x_2, \dots, x_k fixed, as the partial derivative of the CEF

$$\frac{\partial}{\partial x_1} m(x_1, \dots, x_k).$$

When x_1 is discrete we define the marginal effect as a discrete difference. For example, if x_1 is binary, then the marginal effect of x_1 on the CEF is

$$m(1, x_2, \dots, x_k) - m(0, x_2, \dots, x_k).$$

We can unify the continuous and discrete cases with the notation

$$\nabla_1 m(\mathbf{x}) = \begin{cases} \frac{\partial}{\partial x_1} m(x_1, \dots, x_k), & \text{if } x_1 \text{ is continuous} \\ m(1, x_2, \dots, x_k) - m(0, x_2, \dots, x_k), & \text{if } x_1 \text{ is binary.} \end{cases}$$

Collecting the k effects into one $k \times 1$ vector, we define the **regression derivative** with respect to \mathbf{x} :

$$\nabla m(\mathbf{x}) = \begin{bmatrix} \nabla_1 m(\mathbf{x}) \\ \nabla_2 m(\mathbf{x}) \\ \vdots \\ \nabla_k m(\mathbf{x}) \end{bmatrix}$$

When all elements of \mathbf{x} are continuous, then we have the simplification $\nabla m(\mathbf{x}) = \frac{\partial}{\partial \mathbf{x}} m(\mathbf{x})$, the vector of partial derivatives.

There are two important points to remember concerning our definition of the regression derivative.

First, the effect of each variable is calculated holding the other variables constant. This is the **ceteris paribus** concept commonly used in economics. But in the case of a regression derivative, the conditional mean does not literally hold *all else* constant. It only holds constant the variables included in the conditional mean. This means that the regression derivative depends on which regressors are included. For example, in a regression of wages on education, experience, race and sex, the regression derivative with respect to education shows the marginal effect of education on mean wages, holding constant experience, race and sex. But it does not hold constant an individual's unobservable characteristics (such as ability), nor variables not included in the regression (such as the quality of education).

Second, the regression derivative is the change in the conditional expectation of y , not the change in the actual value of y for an individual. It is tempting to think of the regression derivative as the change in the actual value of y , but this is not a correct interpretation. The regression derivative $\nabla m(\mathbf{x})$ is the change in the actual value of y only if the error e is unaffected by the change in the regressor \mathbf{x} . We return to a discussion of causal effects in Section 2.30.

2.15 Linear CEF

An important special case is when the CEF $m(\mathbf{x}) = \mathbb{E}(y | \mathbf{x})$ is linear in \mathbf{x} . In this case we can write the mean equation as

$$m(\mathbf{x}) = x_1\beta_1 + x_2\beta_2 + \cdots + x_k\beta_k + \beta_{k+1}.$$

Notationally it is convenient to write this as a simple function of the vector \mathbf{x} . An easy way to do so is to augment the regressor vector \mathbf{x} by listing the number “1” as an element. We call this the “constant” and the corresponding coefficient is called the “intercept”. Equivalently, specify that the final element¹⁰ of the vector \mathbf{x} is $x_k = 1$. Thus (2.5) has been redefined as the $k \times 1$ vector

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_{k-1} \\ 1 \end{pmatrix}. \quad (2.15)$$

With this redefinition, the CEF is

$$\begin{aligned} m(\mathbf{x}) &= x_1\beta_1 + x_2\beta_2 + \cdots + \beta_k \\ &= \mathbf{x}'\boldsymbol{\beta} \end{aligned} \quad (2.16)$$

¹⁰The order doesn't matter. It could be any element.

where

$$\boldsymbol{\beta} = \begin{pmatrix} \beta_1 \\ \vdots \\ \beta_k \end{pmatrix} \quad (2.17)$$

is a $k \times 1$ coefficient vector. This is the **linear CEF model**. It is also often called the **linear regression model**, or the regression of y on \mathbf{x} .

In the linear CEF model, the regression derivative is simply the coefficient vector. That is

$$\nabla m(\mathbf{x}) = \boldsymbol{\beta}.$$

This is one of the appealing features of the linear CEF model. The coefficients have simple and natural interpretations as the marginal effects of changing one variable, holding the others constant.

<p>Linear CEF Model</p> $y = \mathbf{x}'\boldsymbol{\beta} + e$ $\mathbb{E}(e \mid \mathbf{x}) = 0$
--

If in addition the error is homoskedastic, we call this the homoskedastic linear CEF model.

<p>Homoskedastic Linear CEF Model</p> $y = \mathbf{x}'\boldsymbol{\beta} + e$ $\mathbb{E}(e \mid \mathbf{x}) = 0$ $\mathbb{E}(e^2 \mid \mathbf{x}) = \sigma^2$

2.16 Linear CEF with Nonlinear Effects

The linear CEF model of the previous section is less restrictive than it might appear, as we can include as regressors nonlinear transformations of the original variables. In this sense, the linear CEF framework is flexible and can capture many nonlinear effects.

For example, suppose we have two scalar variables x_1 and x_2 . The CEF could take the quadratic form

$$m(x_1, x_2) = x_1\beta_1 + x_2\beta_2 + x_1^2\beta_3 + x_2^2\beta_4 + x_1x_2\beta_5 + \beta_6. \quad (2.18)$$

This equation is quadratic in the regressors (x_1, x_2) yet linear in the coefficients $\boldsymbol{\beta} = (\beta_1, \dots, \beta_6)'$. We will descriptively call (2.18) a **quadratic CEF**, and yet (2.18) is also a **linear CEF** in the sense of being linear in the coefficients. The key is to understand that (2.18) is quadratic in the variables (x_1, x_2) yet linear in the coefficients $\boldsymbol{\beta}$.

To simplify the expression, we define the transformations $x_3 = x_1^2$, $x_4 = x_2^2$, $x_5 = x_1x_2$, and $x_6 = 1$, and redefine the regressor vector as $\mathbf{x} = (x_1, \dots, x_6)'$. With this redefinition,

$$m(x_1, x_2) = \mathbf{x}'\boldsymbol{\beta}$$

which is linear in $\boldsymbol{\beta}$. For most econometric purposes (estimation and inference on $\boldsymbol{\beta}$) the linearity in $\boldsymbol{\beta}$ is all that is important.

An exception is in the analysis of regression derivatives. In nonlinear equations such as (2.18), the regression derivative should be defined with respect to the original variables, not with respect to the transformed variables. Thus

$$\begin{aligned}\frac{\partial}{\partial x_1}m(x_1, x_2) &= \beta_1 + 2x_1\beta_3 + x_2\beta_5 \\ \frac{\partial}{\partial x_2}m(x_1, x_2) &= \beta_2 + 2x_2\beta_4 + x_1\beta_5\end{aligned}$$

We see that in the model (2.18), the regression derivatives are not a simple coefficient, but are functions of several coefficients plus the levels of (x_1, x_2) . Consequently it is difficult to interpret the coefficients individually. It is more useful to interpret them as a group.

We typically call β_5 the **interaction effect**. Notice that it appears in both regression derivative equations, and has a symmetric interpretation in each. If $\beta_5 > 0$ then the regression derivative with respect to x_1 is increasing in the level of x_2 (and the regression derivative with respect to x_2 is increasing in the level of x_1), while if $\beta_5 < 0$ the reverse is true. It is worth noting that this symmetry is an artificial implication of the quadratic equation (2.18), and is not a general feature of nonlinear conditional means $m(x_1, x_2)$.

2.17 Linear CEF with Dummy Variables

When all regressors take a finite set of values, it turns out the CEF can be written as a linear function of regressors.

This simplest example is a **binary** variable, which takes only two distinct values. For example, the variable *sex* takes only the values *man* and *woman*. Binary variables are extremely common in econometric applications, and are alternatively called **dummy variables** or **indicator variables**.

Consider the simple case of a single binary regressor. In this case, the conditional mean can only take two distinct values. For example,

$$\mathbb{E}(y \mid sex) = \begin{cases} \mu_0 & \text{if } sex=man \\ \mu_1 & \text{if } sex=woman \end{cases}$$

To facilitate a mathematical treatment, we typically record dummy variables with the values $\{0, 1\}$. For example

$$x_1 = \begin{cases} 0 & \text{if } sex=man \\ 1 & \text{if } sex=woman \end{cases} \quad (2.19)$$

Given this notation we can write the conditional mean as a linear function of the dummy variable x_1 , that is

$$\mathbb{E}(y \mid x_1) = \beta_1 x_1 + \beta_2$$

where $\beta_1 = \mu_1 - \mu_0$ and $\beta_2 = \mu_0$. In this simple regression equation the intercept β_2 is equal to the conditional mean of y for the $x_1 = 0$ subpopulation (men) and the slope β_1 is equal to the difference in the conditional means between the two subpopulations.

Equivalently, we could have defined x_1 as

$$x_1 = \begin{cases} 1 & \text{if } sex=man \\ 0 & \text{if } sex=woman \end{cases} \quad (2.20)$$

In this case, the regression intercept is the mean for women (rather than for men) and the regression slope has switched signs. The two regressions are equivalent but the interpretation of the coefficients has changed. Therefore it is always important to understand the precise definitions of the variables, and illuminating labels are helpful. For example, labelling x_1 as “sex” does not help distinguish

between definitions (2.19) and (2.20). Instead, it is better to label x_1 as “women” or “female” if definition (2.19) is used, or as “men” or “male” if (2.20) is used.

Now suppose we have two dummy variables x_1 and x_2 . For example, $x_2 = 1$ if the person is married, else $x_2 = 0$. The conditional mean given x_1 and x_2 takes at most four possible values:

$$\mathbb{E}(y \mid x_1, x_2) = \begin{cases} \mu_{00} & \text{if } x_1 = 0 \text{ and } x_2 = 0 & (\text{unmarried men}) \\ \mu_{01} & \text{if } x_1 = 0 \text{ and } x_2 = 1 & (\text{married men}) \\ \mu_{10} & \text{if } x_1 = 1 \text{ and } x_2 = 0 & (\text{unmarried women}) \\ \mu_{11} & \text{if } x_1 = 1 \text{ and } x_2 = 1 & (\text{married women}) \end{cases}$$

In this case we can write the conditional mean as a linear function of x_1 , x_2 and their product x_1x_2 :

$$\mathbb{E}(y \mid x_1, x_2) = \beta_1x_1 + \beta_2x_2 + \beta_3x_1x_2 + \beta_4$$

where $\beta_1 = \mu_{10} - \mu_{00}$, $\beta_2 = \mu_{01} - \mu_{00}$, $\beta_3 = \mu_{11} - \mu_{10} - \mu_{01} + \mu_{00}$, and $\beta_4 = \mu_{00}$.

We can view the coefficient β_1 as the effect of sex on expected log wages for unmarried wage earners, the coefficient β_2 as the effect of marriage on expected log wages for men wage earners, and the coefficient β_3 as the difference between the effects of marriage on expected log wages among women and among men. Alternatively, it can also be interpreted as the difference between the effects of sex on expected log wages among married and non-married wage earners. Both interpretations are equally valid. We often describe β_3 as measuring the **interaction** between the two dummy variables, or the **interaction effect**, and describe $\beta_3 = 0$ as the case when the interaction effect is zero.

In this setting we can see that the CEF is linear in the three variables (x_1, x_2, x_1x_2) . Thus to put the model in the framework of Section 2.15, we would define the regressor $x_3 = x_1x_2$ and the regressor vector as

$$\mathbf{x} = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ 1 \end{pmatrix}.$$

So even though we started with only 2 dummy variables, the number of regressors (including the intercept) is 4.

If there are 3 dummy variables x_1, x_2, x_3 , then $\mathbb{E}(y \mid x_1, x_2, x_3)$ takes at most $2^3 = 8$ distinct values and can be written as the linear function

$$\mathbb{E}(y \mid x_1, x_2, x_3) = \beta_1x_1 + \beta_2x_2 + \beta_3x_3 + \beta_4x_1x_2 + \beta_5x_1x_3 + \beta_6x_2x_3 + \beta_7x_1x_2x_3 + \beta_8$$

which has eight regressors including the intercept.

In general, if there are p dummy variables x_1, \dots, x_p then the CEF $\mathbb{E}(y \mid x_1, x_2, \dots, x_p)$ takes at most 2^p distinct values, and can be written as a linear function of the 2^p regressors including x_1, x_2, \dots, x_p and all cross-products. This might be excessive in practice if p is modestly large. In the next section we will discuss projection approximations which yield more parsimonious parameterizations.

We started this section by saying that the conditional mean is linear whenever all regressors take only a finite number of possible values. How can we see this? Take a **categorical** variable, such as *race*. For example, we earlier divided race into three categories. We can record categorical variables using numbers to indicate each category, for example

$$x_3 = \begin{cases} 1 & \text{if } \textit{white} \\ 2 & \text{if } \textit{black} \\ 3 & \text{if } \textit{other} \end{cases}$$

When doing so, the values of x_3 have no meaning in terms of magnitude, they simply indicate the relevant category.

When the regressor is categorical the conditional mean of y given x_3 takes a distinct value for each possibility:

$$\mathbb{E}(y \mid x_3) = \begin{cases} \mu_1 & \text{if } x_3 = 1 \\ \mu_2 & \text{if } x_3 = 2 \\ \mu_3 & \text{if } x_3 = 3 \end{cases}$$

This is not a linear function of x_3 itself, but it can be made a linear function by constructing dummy variables for two of the three categories. For example

$$x_4 = \begin{cases} 1 & \text{if } \textit{black} \\ 0 & \text{if } \textit{not black} \end{cases}$$

$$x_5 = \begin{cases} 1 & \text{if } \textit{other} \\ 0 & \text{if } \textit{not other} \end{cases}$$

In this case, the categorical variable x_3 is equivalent to the pair of dummy variables (x_4, x_5) . The explicit relationship is

$$x_3 = \begin{cases} 1 & \text{if } x_4 = 0 \text{ and } x_5 = 0 \\ 2 & \text{if } x_4 = 1 \text{ and } x_5 = 0 \\ 3 & \text{if } x_4 = 0 \text{ and } x_5 = 1 \end{cases}$$

Given these transformations, we can write the conditional mean of y as a linear function of x_4 and x_5

$$\mathbb{E}(y \mid x_3) = \mathbb{E}(y \mid x_4, x_5) = \beta_1 x_4 + \beta_2 x_5 + \beta_3$$

We can write the CEF as either $\mathbb{E}(y \mid x_3)$ or $\mathbb{E}(y \mid x_4, x_5)$ (they are equivalent), but it is only linear as a function of x_4 and x_5 .

This setting is similar to the case of two dummy variables, with the difference that we have not included the interaction term $x_4 x_5$. This is because the event $\{x_4 = 1 \text{ and } x_5 = 1\}$ is empty by construction, so $x_4 x_5 = 0$ by definition.

2.18 Best Linear Predictor

While the conditional mean $m(\mathbf{x}) = \mathbb{E}(y \mid \mathbf{x})$ is the best predictor of y among all functions of \mathbf{x} , its functional form is typically unknown. In particular, the linear CEF model is empirically unlikely to be accurate unless \mathbf{x} is discrete and low-dimensional so all interactions are included. Consequently in most cases it is more realistic to view the linear specification (2.16) as an approximation. In this section we derive a specific approximation with a simple interpretation.

Theorem 2.11.1 showed that the conditional mean $m(\mathbf{x})$ is the best predictor in the sense that it has the lowest mean squared error among all predictors. By extension, we can define an approximation to the CEF by the linear function with the lowest mean squared error among all linear predictors.

For this derivation we require the following regularity condition.

Assumption 2.18.1

1. $\mathbb{E}y^2 < \infty$.
2. $\mathbb{E}\|\mathbf{x}\|^2 < \infty$.
3. $\mathbf{Q}_{\mathbf{x}\mathbf{x}} = \mathbb{E}(\mathbf{x}\mathbf{x}')$ is positive definite.

In Assumption 2.18.1.2 we use the notation $\|\mathbf{x}\| = (\mathbf{x}'\mathbf{x})^{1/2}$ to denote the Euclidean length of the vector \mathbf{x} .

The first two parts of Assumption 2.18.1 imply that the variables y and \mathbf{x} have finite means, variances, and covariances. The third part of the assumption is more technical, and its role will become apparent shortly. It is equivalent to imposing that the columns of $\mathbf{Q}_{\mathbf{x}\mathbf{x}} = \mathbb{E}(\mathbf{x}\mathbf{x}')$ are linearly independent, or equivalently that the matrix $\mathbf{Q}_{\mathbf{x}\mathbf{x}}$ is invertible.

A linear predictor for y is a function of the form $\mathbf{x}'\boldsymbol{\beta}$ for some $\boldsymbol{\beta} \in \mathbb{R}^k$. The mean squared prediction error is

$$S(\boldsymbol{\beta}) = \mathbb{E}(y - \mathbf{x}'\boldsymbol{\beta})^2.$$

The **best linear predictor** of y given \mathbf{x} , written $\mathcal{P}(y | \mathbf{x})$, is found by selecting the vector $\boldsymbol{\beta}$ to minimize $S(\boldsymbol{\beta})$.

Definition 2.18.1 *The Best Linear Predictor of y given \mathbf{x} is*

$$\mathcal{P}(y | \mathbf{x}) = \mathbf{x}'\boldsymbol{\beta}$$

where $\boldsymbol{\beta}$ minimizes the mean squared prediction error

$$S(\boldsymbol{\beta}) = \mathbb{E}(y - \mathbf{x}'\boldsymbol{\beta})^2.$$

The minimizer

$$\boldsymbol{\beta} = \underset{\boldsymbol{\beta} \in \mathbb{R}^k}{\operatorname{argmin}} S(\boldsymbol{\beta}) \tag{2.21}$$

is called the **Linear Projection Coefficient**.

We now calculate an explicit expression for its value.

The mean squared prediction error can be written out as a quadratic function of $\boldsymbol{\beta}$:

$$S(\boldsymbol{\beta}) = \mathbb{E}y^2 - 2\boldsymbol{\beta}'\mathbb{E}(\mathbf{x}y) + \boldsymbol{\beta}'\mathbb{E}(\mathbf{x}\mathbf{x}')\boldsymbol{\beta}.$$

The quadratic structure of $S(\boldsymbol{\beta})$ means that we can solve explicitly for the minimizer. The first-order condition for minimization (from Appendix A.9) is

$$\mathbf{0} = \frac{\partial}{\partial \boldsymbol{\beta}} S(\boldsymbol{\beta}) = -2\mathbb{E}(\mathbf{x}y) + 2\mathbb{E}(\mathbf{x}\mathbf{x}')\boldsymbol{\beta}. \tag{2.22}$$

Rewriting (2.22) as

$$2\mathbb{E}(\mathbf{x}y) = 2\mathbb{E}(\mathbf{x}\mathbf{x}')\boldsymbol{\beta}$$

and dividing by 2, this equation takes the form

$$\mathbf{Q}_{\mathbf{x}y} = \mathbf{Q}_{\mathbf{x}\mathbf{x}}\boldsymbol{\beta} \tag{2.23}$$

where $\mathbf{Q}_{\mathbf{x}y} = \mathbb{E}(\mathbf{x}y)$ is $k \times 1$ and $\mathbf{Q}_{\mathbf{x}\mathbf{x}} = \mathbb{E}(\mathbf{x}\mathbf{x}')$ is $k \times k$. The solution is found by inverting the matrix $\mathbf{Q}_{\mathbf{x}\mathbf{x}}$, and is written

$$\boldsymbol{\beta} = \mathbf{Q}_{\mathbf{x}\mathbf{x}}^{-1} \mathbf{Q}_{\mathbf{x}y}$$

or

$$\boldsymbol{\beta} = (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1} \mathbb{E}(\mathbf{x}y). \tag{2.24}$$

It is worth taking the time to understand the notation involved in the expression (2.24). $\mathbf{Q}_{\mathbf{x}\mathbf{x}}$ is a $k \times k$ matrix and $\mathbf{Q}_{\mathbf{x}y}$ is a $k \times 1$ column vector. Therefore, alternative expressions such as $\frac{\mathbb{E}(\mathbf{x}y)}{\mathbb{E}(\mathbf{x}\mathbf{x}')}$

or $\mathbb{E}(\mathbf{x}y)(\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1}$ are incoherent and incorrect. We also can now see the role of Assumption 2.18.1.3. It is necessary in order for the solution (2.24) to exist. Otherwise, there would be multiple solutions to the equation (2.23).

We now have an explicit expression for the best linear predictor:

$$\mathcal{P}(y | \mathbf{x}) = \mathbf{x}' (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1} \mathbb{E}(\mathbf{x}y).$$

This expression is also referred to as the **linear projection** of y on \mathbf{x} .

The **projection error** is

$$e = y - \mathbf{x}'\boldsymbol{\beta}. \tag{2.25}$$

This equals the error from the regression equation when (and only when) the conditional mean is linear in \mathbf{x} , otherwise they are distinct.

Rewriting, we obtain a decomposition of y into linear predictor and error

$$y = \mathbf{x}'\boldsymbol{\beta} + e. \tag{2.26}$$

In general we call equation (2.26) or $\mathbf{x}'\boldsymbol{\beta}$ the best linear predictor of y given \mathbf{x} , or the linear projection of y on \mathbf{x} . Equation (2.26) is also often called the **regression** of y on \mathbf{x} but this can sometimes be confusing as economists use the term *regression* in many contexts. (Recall that we said in Section 2.15 that the linear CEF model is also called the linear regression model.)

An important property of the projection error e is

$$\mathbb{E}(\mathbf{x}e) = \mathbf{0}. \tag{2.27}$$

To see this, using the definitions (2.25) and (2.24) and the matrix properties $\mathbf{A}\mathbf{A}^{-1} = \mathbf{I}$ and $\mathbf{I}\mathbf{a} = \mathbf{a}$,

$$\begin{aligned} \mathbb{E}(\mathbf{x}e) &= \mathbb{E}(\mathbf{x}(y - \mathbf{x}'\boldsymbol{\beta})) \\ &= \mathbb{E}(\mathbf{x}y) - \mathbb{E}(\mathbf{x}\mathbf{x}') (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1} \mathbb{E}(\mathbf{x}y) \\ &= \mathbf{0} \end{aligned} \tag{2.28}$$

as claimed.

Equation (2.27) is a set of k equations, one for each regressor. In other words, (2.27) is equivalent to

$$\mathbb{E}(x_j e) = 0 \tag{2.29}$$

for $j = 1, \dots, k$. As in (2.15), the regressor vector \mathbf{x} typically contains a constant, e.g. $x_k = 1$. In this case (2.29) for $j = k$ is the same as

$$\mathbb{E}(e) = 0. \tag{2.30}$$

Thus the projection error has a mean of zero when the regressor vector contains a constant. (When \mathbf{x} does not have a constant, (2.30) is not guaranteed. As it is desirable for e to have a zero mean, this is a good reason to always include a constant in any regression model.)

It is also useful to observe that since $\text{cov}(x_j, e) = \mathbb{E}(x_j e) - \mathbb{E}(x_j)\mathbb{E}(e)$, then (2.29)-(2.30) together imply that the variables x_j and e are uncorrelated.

This completes the derivation of the model. We summarize some of the most important properties.

Theorem 2.18.1 Properties of Linear Projection Model

Under Assumption 2.18.1,

1. The moments $\mathbb{E}(\mathbf{x}\mathbf{x}')$ and $\mathbb{E}(\mathbf{x}y)$ exist with finite elements.

2. The Linear Projection Coefficient (2.21) exists, is unique, and equals

$$\boldsymbol{\beta} = (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1} \mathbb{E}(\mathbf{x}y).$$

3. The best linear predictor of y given \mathbf{x} is

$$\mathcal{P}(y | \mathbf{x}) = \mathbf{x}' (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1} \mathbb{E}(\mathbf{x}y).$$

4. The projection error $e = y - \mathbf{x}'\boldsymbol{\beta}$ exists and satisfies

$$\mathbb{E}(e^2) < \infty$$

and

$$\mathbb{E}(\mathbf{x}e) = \mathbf{0}.$$

5. If \mathbf{x} contains an constant, then

$$\mathbb{E}(e) = 0.$$

6. If $\mathbb{E}|y|^r < \infty$ and $\mathbb{E}\|\mathbf{x}\|^r < \infty$ for $r \geq 2$ then $\mathbb{E}|e|^r < \infty$.

A complete proof of Theorem 2.18.1 is given in Section 2.34.

It is useful to reflect on the generality of Theorem 2.18.1. The only restriction is Assumption 2.18.1. Thus for any random variables (y, \mathbf{x}) with finite variances we can define a linear equation (2.26) with the properties listed in Theorem 2.18.1. Stronger assumptions (such as the linear CEF model) are not necessary. In this sense the linear model (2.26) exists quite generally. However, it is important not to misinterpret the generality of this statement. The linear equation (2.26) is defined as the best linear predictor. It is not necessarily a conditional mean, nor a parameter of a structural or causal economic model.

Linear Projection Model

$$y = \mathbf{x}'\boldsymbol{\beta} + e.$$

$$\mathbb{E}(\mathbf{x}e) = \mathbf{0}$$

$$\boldsymbol{\beta} = (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1} \mathbb{E}(\mathbf{x}y)$$

We illustrate projection using three log wage equations introduced in earlier sections.

For our first example, we consider a model with the two dummy variables for sex and race similar to Table 2.1. As we learned in Section 2.17, the entries in this table can be equivalently expressed by a linear CEF. For simplicity, let's consider the CEF of $\log(\text{wage})$ as a function of *Black* and *Female*.

$$\mathbb{E}(\log(\text{wage}) | \text{Black}, \text{Female}) = -0.20\text{Black} - 0.24\text{Female} + 0.10\text{Black} \times \text{Female} + 3.06. \quad (2.31)$$

This is a CEF as the variables are binary and all interactions are included.

Now consider a simpler model omitting the interaction effect. This is the linear projection on the variables *Black* and *Female*

$$\mathcal{P}(\log(wage) \mid Black, Female) = -0.15Black - 0.23Female + 3.06. \quad (2.32)$$

What is the difference? The full CEF (2.31) shows that the race gap is differentiated by sex: it is 20% for black men (relative to non-black men) and 10% for black women (relative to non-black women). The projection model (2.32) simplifies this analysis, calculating an average 15% wage gap for blacks, ignoring the role of sex. Notice that this is despite the fact that the sex variable is included in (2.32).

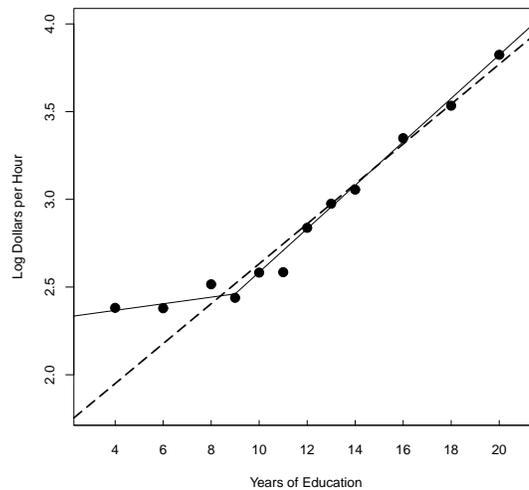


Figure 2.8: Projections of $\log(wage)$ onto Education

For our second example we consider the CEF of log wages as a function of years of education for white men which was illustrated in Figure 2.5 and is repeated in Figure 2.8. Superimposed on the figure are two projections. The first (given by the dashed line) is the linear projection of log wages on years of education

$$\mathcal{P}(\log(wage) \mid Education) = 0.11Education + 1.5$$

This simple equation indicates an average 11% increase in wages for every year of education. An inspection of the Figure shows that this approximation works well for $education \geq 9$, but under-predicts for individuals with lower levels of education. To correct this imbalance we use a linear spline equation which allows different rates of return above and below 9 years of education:

$$\begin{aligned} \mathcal{P}(\log(wage) \mid Education, (Education - 9) \times 1 (Education > 9)) \\ = 0.02Education + 0.10 \times (Education - 9) \times 1 (Education > 9) + 2.3 \end{aligned}$$

This equation is displayed in Figure 2.8 using the solid line, and appears to fit much better. It indicates a 2% increase in mean wages for every year of education below 9, and a 12% increase in mean wages for every year of education above 9. It is still an approximation to the conditional mean but it appears to be fairly reasonable.

For our third example we take the CEF of log wages as a function of years of experience for white men with 12 years of education, which was illustrated in Figure 2.6 and is repeated as the solid line in Figure 2.9. Superimposed on the figure are two projections. The first (given by the dot-dashed line) is the linear projection on experience

$$\mathcal{P}(\log(wage) \mid Experience) = 0.011Experience + 2.5$$

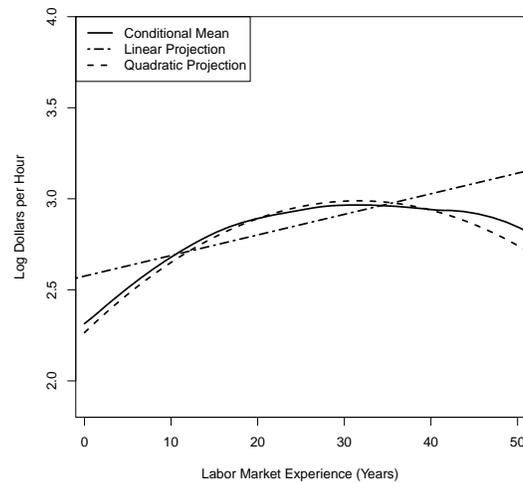


Figure 2.9: Linear and Quadratic Projections of $\log(wage)$ onto Experience

and the second (given by the dashed line) is the linear projection on experience and its square

$$\mathcal{P}(\log(wage) \mid Experience) = 0.046Experience - 0.0007Experience^2 + 2.3.$$

It is fairly clear from an examination of Figure 2.9 that the first linear projection is a poor approximation. It over-predicts wages for young and old workers, and under-predicts for the rest. Most importantly, it misses the strong downturn in expected wages for older wage-earners. The second projection fits much better. We can call this equation a **quadratic projection** since the function is quadratic in *experience*.

Invertibility and Identification

The linear projection coefficient $\beta = (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1}\mathbb{E}(\mathbf{x}y)$ exists and is unique as long as the $k \times k$ matrix $\mathbf{Q}_{\mathbf{x}\mathbf{x}} = \mathbb{E}(\mathbf{x}\mathbf{x}')$ is invertible. The matrix $\mathbf{Q}_{\mathbf{x}\mathbf{x}}$ is sometimes called the **design matrix**, as in experimental settings the researcher is able to control $\mathbf{Q}_{\mathbf{x}\mathbf{x}}$ by manipulating the distribution of the regressors \mathbf{x} .

Observe that for any non-zero $\alpha \in \mathbb{R}^k$,

$$\alpha' \mathbf{Q}_{\mathbf{x}\mathbf{x}} \alpha = \mathbb{E}(\alpha' \mathbf{x} \mathbf{x}' \alpha) = \mathbb{E}(\alpha' \mathbf{x})^2 \geq 0$$

so $\mathbf{Q}_{\mathbf{x}\mathbf{x}}$ by construction is positive semi-definite. The assumption that it is positive definite means that this is a strict inequality, $\mathbb{E}(\alpha' \mathbf{x})^2 > 0$. Equivalently, there cannot exist a non-zero vector α such that $\alpha' \mathbf{x} = 0$ identically. This occurs when redundant variables are included in \mathbf{x} . Positive semi-definite matrices are invertible if and only if they are positive definite. When $\mathbf{Q}_{\mathbf{x}\mathbf{x}}$ is invertible then $\beta = (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1}\mathbb{E}(\mathbf{x}y)$ exists and is uniquely defined. In other words, in order for β to be uniquely defined, we must exclude the degenerate situation of redundant variables.

Theorem 2.18.1 shows that the linear projection coefficient β is **identified** (uniquely determined) under Assumption 2.18.1. The key is invertibility of $\mathbf{Q}_{\mathbf{x}\mathbf{x}}$. Otherwise, there is no unique solution to the equation

$$\mathbf{Q}_{\mathbf{x}\mathbf{x}} \beta = \mathbf{Q}_{\mathbf{x}y}. \quad (2.33)$$

When $\mathbf{Q}_{\mathbf{x}\mathbf{x}}$ is not invertible there are multiple solutions to (2.33), all of which yield an equivalent best linear predictor $\mathbf{x}'\beta$. In this case the coefficient β is **not identified** as it does not have a unique value. Even so, the best linear predictor $\mathbf{x}'\beta$ still identified. One solution is to set

$$\beta = (\mathbb{E}(\mathbf{x}\mathbf{x}'))^- \mathbb{E}(\mathbf{x}y)$$

where \mathbf{A}^- denotes the generalized inverse of \mathbf{A} (see Appendix A.5).

2.19 Linear Predictor Error Variance

As in the CEF model, we define the error variance as

$$\sigma^2 = \mathbb{E}(e^2).$$

Setting $Q_{yy} = \mathbb{E}(y^2)$ and $\mathbf{Q}_{y\mathbf{x}} = \mathbb{E}(y\mathbf{x}')$ we can write σ^2 as

$$\begin{aligned} \sigma^2 &= \mathbb{E}(y - \mathbf{x}'\beta)^2 \\ &= \mathbb{E}y^2 - 2\mathbb{E}(y\mathbf{x}')\beta + \beta'\mathbb{E}(\mathbf{x}\mathbf{x}')\beta \\ &= Q_{yy} - 2\mathbf{Q}_{y\mathbf{x}}\mathbf{Q}_{\mathbf{x}\mathbf{x}}^{-1}\mathbf{Q}_{\mathbf{x}y} + \mathbf{Q}_{y\mathbf{x}}\mathbf{Q}_{\mathbf{x}\mathbf{x}}^{-1}\mathbf{Q}_{\mathbf{x}\mathbf{x}}\mathbf{Q}_{\mathbf{x}\mathbf{x}}^{-1}\mathbf{Q}_{\mathbf{x}y} \\ &= Q_{yy} - \mathbf{Q}_{y\mathbf{x}}\mathbf{Q}_{\mathbf{x}\mathbf{x}}^{-1}\mathbf{Q}_{\mathbf{x}y} \\ &\stackrel{\text{def}}{=} Q_{yy \cdot \mathbf{x}}. \end{aligned} \quad (2.34)$$

One useful feature of this formula is that it shows that $Q_{yy \cdot \mathbf{x}} = Q_{yy} - \mathbf{Q}_{y\mathbf{x}}\mathbf{Q}_{\mathbf{x}\mathbf{x}}^{-1}\mathbf{Q}_{\mathbf{x}y}$ equals the variance of the error from the linear projection of y on \mathbf{x} .

2.20 Regression Coefficients

Sometimes it is useful to separate the constant from the other regressors, and write the linear projection equation in the format

$$y = \mathbf{x}'\boldsymbol{\beta} + \alpha + e \quad (2.35)$$

where α is the intercept and \mathbf{x} does not contain a constant.

Taking expectations of this equation, we find

$$\mathbb{E}y = \mathbb{E}\mathbf{x}'\boldsymbol{\beta} + \mathbb{E}\alpha + \mathbb{E}e$$

or

$$\mu_y = \mu_x'\boldsymbol{\beta} + \alpha$$

where $\mu_y = \mathbb{E}y$ and $\mu_x = \mathbb{E}\mathbf{x}$, since $\mathbb{E}(e) = 0$ from (2.30). (While \mathbf{x} does not contain a constant, the equation does so (2.30) still applies.) Rearranging, we find

$$\alpha = \mu_y - \mu_x'\boldsymbol{\beta}.$$

Subtracting this equation from (2.35) we find

$$y - \mu_y = (\mathbf{x} - \mu_x)'\boldsymbol{\beta} + e, \quad (2.36)$$

a linear equation between the centered variables $y - \mu_y$ and $\mathbf{x} - \mu_x$. (They are centered at their means, so are mean-zero random variables.) Because $\mathbf{x} - \mu_x$ is uncorrelated with e , (2.36) is also a linear projection, thus by the formula for the linear projection model,

$$\begin{aligned} \boldsymbol{\beta} &= (\mathbb{E}((\mathbf{x} - \mu_x)(\mathbf{x} - \mu_x)'))^{-1} \mathbb{E}((\mathbf{x} - \mu_x)(y - \mu_y)) \\ &= \text{var}(\mathbf{x})^{-1} \text{cov}(\mathbf{x}, y) \end{aligned}$$

a function only of the covariances¹¹ of \mathbf{x} and y .

Theorem 2.20.1 *In the linear projection model*

$$y = \mathbf{x}'\boldsymbol{\beta} + \alpha + e,$$

then

$$\alpha = \mu_y - \mu_x'\boldsymbol{\beta} \quad (2.37)$$

and

$$\boldsymbol{\beta} = \text{var}(\mathbf{x})^{-1} \text{cov}(\mathbf{x}, y). \quad (2.38)$$

2.21 Regression Sub-Vectors

Let the regressors be partitioned as

$$\mathbf{x} = \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \end{pmatrix}. \quad (2.39)$$

¹¹The **covariance matrix** between vectors \mathbf{x} and \mathbf{z} is $\text{cov}(\mathbf{x}, \mathbf{z}) = \mathbb{E}((\mathbf{x} - \mathbb{E}\mathbf{x})(\mathbf{z} - \mathbb{E}\mathbf{z})')$. The (co)variance matrix of the vector \mathbf{x} is $\text{var}(\mathbf{x}) = \text{cov}(\mathbf{x}, \mathbf{x}) = \mathbb{E}((\mathbf{x} - \mathbb{E}\mathbf{x})(\mathbf{x} - \mathbb{E}\mathbf{x})')$.

We can write the projection of y on \mathbf{x} as

$$\begin{aligned} y &= \mathbf{x}'\boldsymbol{\beta} + e \\ &= \mathbf{x}'_1\boldsymbol{\beta}_1 + \mathbf{x}'_2\boldsymbol{\beta}_2 + e \\ \mathbb{E}(\mathbf{x}e) &= \mathbf{0}. \end{aligned} \tag{2.40}$$

In this section we derive formula for the sub-vectors $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$.

Partition \mathbf{Q}_{xx} conformably with \mathbf{x}

$$\mathbf{Q}_{xx} = \begin{bmatrix} \mathbf{Q}_{11} & \mathbf{Q}_{12} \\ \mathbf{Q}_{21} & \mathbf{Q}_{22} \end{bmatrix} = \begin{bmatrix} \mathbb{E}(\mathbf{x}_1\mathbf{x}'_1) & \mathbb{E}(\mathbf{x}_1\mathbf{x}'_2) \\ \mathbb{E}(\mathbf{x}_2\mathbf{x}'_1) & \mathbb{E}(\mathbf{x}_2\mathbf{x}'_2) \end{bmatrix}$$

and similarly \mathbf{Q}_{xy}

$$\mathbf{Q}_{xy} = \begin{bmatrix} \mathbf{Q}_{1y} \\ \mathbf{Q}_{2y} \end{bmatrix} = \begin{bmatrix} \mathbb{E}(\mathbf{x}_1y) \\ \mathbb{E}(\mathbf{x}_2y) \end{bmatrix}.$$

By the partitioned matrix inversion formula (A.4)

$$\mathbf{Q}_{xx}^{-1} = \begin{bmatrix} \mathbf{Q}_{11} & \mathbf{Q}_{12} \\ \mathbf{Q}_{21} & \mathbf{Q}_{22} \end{bmatrix}^{-1} \stackrel{def}{=} \begin{bmatrix} \mathbf{Q}^{11} & \mathbf{Q}^{12} \\ \mathbf{Q}^{21} & \mathbf{Q}^{22} \end{bmatrix} = \begin{bmatrix} \mathbf{Q}_{11.2}^{-1} & -\mathbf{Q}_{11.2}^{-1}\mathbf{Q}_{12}\mathbf{Q}_{22}^{-1} \\ -\mathbf{Q}_{22.1}^{-1}\mathbf{Q}_{21}\mathbf{Q}_{11}^{-1} & \mathbf{Q}_{22.1}^{-1} \end{bmatrix}. \tag{2.41}$$

where $\mathbf{Q}_{11.2} \stackrel{def}{=} \mathbf{Q}_{11} - \mathbf{Q}_{12}\mathbf{Q}_{22}^{-1}\mathbf{Q}_{21}$ and $\mathbf{Q}_{22.1} \stackrel{def}{=} \mathbf{Q}_{22} - \mathbf{Q}_{21}\mathbf{Q}_{11}^{-1}\mathbf{Q}_{12}$. Thus

$$\begin{aligned} \boldsymbol{\beta} &= \begin{pmatrix} \boldsymbol{\beta}_1 \\ \boldsymbol{\beta}_2 \end{pmatrix} \\ &= \begin{bmatrix} \mathbf{Q}_{11.2}^{-1} & -\mathbf{Q}_{11.2}^{-1}\mathbf{Q}_{12}\mathbf{Q}_{22}^{-1} \\ -\mathbf{Q}_{22.1}^{-1}\mathbf{Q}_{21}\mathbf{Q}_{11}^{-1} & \mathbf{Q}_{22.1}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{Q}_{1y} \\ \mathbf{Q}_{2y} \end{bmatrix} \\ &= \begin{pmatrix} \mathbf{Q}_{11.2}^{-1}(\mathbf{Q}_{1y} - \mathbf{Q}_{12}\mathbf{Q}_{22}^{-1}\mathbf{Q}_{2y}) \\ \mathbf{Q}_{22.1}^{-1}(\mathbf{Q}_{2y} - \mathbf{Q}_{21}\mathbf{Q}_{11}^{-1}\mathbf{Q}_{1y}) \end{pmatrix} \\ &= \begin{pmatrix} \mathbf{Q}_{11.2}^{-1}\mathbf{Q}_{1y.2} \\ \mathbf{Q}_{22.1}^{-1}\mathbf{Q}_{2y.1} \end{pmatrix} \end{aligned}$$

We have shown that

$$\begin{aligned} \boldsymbol{\beta}_1 &= \mathbf{Q}_{11.2}^{-1}\mathbf{Q}_{1y.2} \\ \boldsymbol{\beta}_2 &= \mathbf{Q}_{22.1}^{-1}\mathbf{Q}_{2y.1} \end{aligned}$$

2.22 Coefficient Decomposition

In the previous section we derived formulae for the coefficient sub-vectors $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$. We now use these formulae to give a useful interpretation of the coefficients in terms of an iterated projection.

Take equation (2.40) for the case $\dim(x_1) = 1$ so that $\boldsymbol{\beta}_1 \in \mathbb{R}$.

$$y = x_1\boldsymbol{\beta}_1 + \mathbf{x}'_2\boldsymbol{\beta}_2 + e. \tag{2.42}$$

Now consider the projection of x_1 on \mathbf{x}_2 :

$$\begin{aligned} x_1 &= \mathbf{x}'_2\boldsymbol{\gamma}_2 + u_1 \\ \mathbb{E}(\mathbf{x}_2u_1) &= \mathbf{0}. \end{aligned}$$

From (2.24) and (2.34), $\boldsymbol{\gamma}_2 = \mathbf{Q}_{22}^{-1}\mathbf{Q}_{21}$ and $\mathbb{E}u_1^2 = \mathbf{Q}_{11.2} = \mathbf{Q}_{11} - \mathbf{Q}_{12}\mathbf{Q}_{22}^{-1}\mathbf{Q}_{21}$. We can also calculate that

$$\mathbb{E}(u_1y) = \mathbb{E}((x_1 - \boldsymbol{\gamma}'_2\mathbf{x}_2)y) = \mathbb{E}(x_1y) - \boldsymbol{\gamma}'_2\mathbb{E}(\mathbf{x}_2y) = \mathbf{Q}_{1y} - \mathbf{Q}_{12}\mathbf{Q}_{22}^{-1}\mathbf{Q}_{2y} = \mathbf{Q}_{1y.2}.$$

We have found that

$$\beta_1 = \mathbf{Q}_{11 \cdot 2}^{-1} \mathbf{Q}_{1y \cdot 2} = \frac{\mathbb{E}(u_1 y)}{\mathbb{E}u_1^2}$$

the coefficient from the simple regression of y on u_1 .

What this means is that in the multivariate projection equation (2.42), the coefficient β_1 equals the projection coefficient from a regression of y on u_1 , the error from a projection of x_1 on the other regressors \mathbf{x}_2 . The error u_1 can be thought of as the component of x_1 which is not linearly explained by the other regressors. Thus the coefficient β_1 equals the linear effect of x_1 on y , after stripping out the effects of the other variables.

There was nothing special in the choice of the variable x_1 . This derivation applies symmetrically to all coefficients in a linear projection. Each coefficient equals the simple regression of y on the error from a projection of that regressor on all the other regressors. Each coefficient equals the linear effect of that variable on y , after linearly controlling for all the other regressors.

2.23 Omitted Variable Bias

Again, let the regressors be partitioned as in (2.39). Consider the projection of y on \mathbf{x}_1 only. Perhaps this is done because the variables \mathbf{x}_2 are not observed. This is the equation

$$\begin{aligned} y &= \mathbf{x}'_1 \gamma_1 + u \\ \mathbb{E}(\mathbf{x}_1 u) &= \mathbf{0}. \end{aligned} \tag{2.43}$$

Notice that we have written the coefficient on \mathbf{x}_1 as γ_1 rather than β_1 and the error as u rather than e . This is because (2.43) is different than (2.40). Goldberger (1991) introduced the catchy labels **long regression** for (2.40) and **short regression** for (2.43) to emphasize the distinction.

Typically, $\beta_1 \neq \gamma_1$, except in special cases. To see this, we calculate

$$\begin{aligned} \gamma_1 &= (\mathbb{E}(\mathbf{x}_1 \mathbf{x}'_1))^{-1} \mathbb{E}(\mathbf{x}_1 y) \\ &= (\mathbb{E}(\mathbf{x}_1 \mathbf{x}'_1))^{-1} \mathbb{E}(\mathbf{x}_1 (\mathbf{x}'_1 \beta_1 + \mathbf{x}'_2 \beta_2 + e)) \\ &= \beta_1 + (\mathbb{E}(\mathbf{x}_1 \mathbf{x}'_1))^{-1} \mathbb{E}(\mathbf{x}_1 \mathbf{x}'_2) \beta_2 \\ &= \beta_1 + \mathbf{\Gamma} \beta_2 \end{aligned}$$

where

$$\mathbf{\Gamma} = (\mathbb{E}(\mathbf{x}_1 \mathbf{x}'_1))^{-1} \mathbb{E}(\mathbf{x}_1 \mathbf{x}'_2)$$

is the coefficient matrix from a projection of \mathbf{x}_2 on \mathbf{x}_1 .

Observe that $\gamma_1 = \beta_1 + \mathbf{\Gamma} \beta_2 \neq \beta_1$ unless $\mathbf{\Gamma} = \mathbf{0}$ or $\beta_2 = \mathbf{0}$. Thus the short and long regressions have different coefficients on \mathbf{x}_1 . They are the same only under one of two conditions. First, if the projection of \mathbf{x}_2 on \mathbf{x}_1 yields a set of zero coefficients (they are uncorrelated), or second, if the coefficient on \mathbf{x}_2 in (2.40) is zero. In general, the coefficient in (2.43) is γ_1 rather than β_1 . The difference $\mathbf{\Gamma} \beta_2$ between γ_1 and β_1 is known as **omitted variable bias**. It is the consequence of omission of a relevant correlated variable.

To avoid omitted variables bias the standard advice is to include all potentially relevant variables in estimated models. By construction, the general model will be free of such bias. Unfortunately in many cases it is not feasible to completely follow this advice as many desired variables are not observed. In this case, the possibility of omitted variables bias should be acknowledged and discussed in the course of an empirical investigation.

For example, suppose y is log wages, x_1 is education, and x_2 is intellectual ability. It seems reasonable to suppose that education and intellectual ability are positively correlated (highly able individuals attain higher levels of education) which means $\mathbf{\Gamma} > 0$. It also seems reasonable to suppose that conditional on education, individuals with higher intelligence will earn higher wages

on average, so that $\beta_2 > 0$. This implies that $\Gamma\beta_2 > 0$ and $\gamma_1 = \beta_1 + \Gamma\beta_2 > \beta_1$. Therefore, it seems reasonable to expect that in a regression of wages on education with ability omitted, the coefficient on education is higher than in a regression where ability is included. In other words, in this context the omitted variable biases the regression coefficient upwards.

2.24 Best Linear Approximation

There are alternative ways we could construct a linear approximation $\mathbf{x}'\boldsymbol{\beta}$ to the conditional mean $m(\mathbf{x})$. In this section we show that one alternative approach turns out to yield the same answer as the best linear predictor.

We start by defining the mean-square approximation error of $\mathbf{x}'\boldsymbol{\beta}$ to $m(\mathbf{x})$ as the expected squared difference between $\mathbf{x}'\boldsymbol{\beta}$ and the conditional mean $m(\mathbf{x})$

$$d(\boldsymbol{\beta}) = \mathbb{E} (m(\mathbf{x}) - \mathbf{x}'\boldsymbol{\beta})^2. \quad (2.44)$$

The function $d(\boldsymbol{\beta})$ is a measure of the deviation of $\mathbf{x}'\boldsymbol{\beta}$ from $m(\mathbf{x})$. If the two functions are identical then $d(\boldsymbol{\beta}) = 0$, otherwise $d(\boldsymbol{\beta}) > 0$. We can also view the mean-square difference $d(\boldsymbol{\beta})$ as a density-weighted average of the function $(m(\mathbf{x}) - \mathbf{x}'\boldsymbol{\beta})^2$, since

$$d(\boldsymbol{\beta}) = \int_{\mathbb{R}^k} (m(\mathbf{x}) - \mathbf{x}'\boldsymbol{\beta})^2 f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}$$

where $f_{\mathbf{x}}(\mathbf{x})$ is the marginal density of \mathbf{x} .

We can then define the best linear approximation to the conditional $m(\mathbf{x})$ as the function $\mathbf{x}'\boldsymbol{\beta}$ obtained by selecting $\boldsymbol{\beta}$ to minimize $d(\boldsymbol{\beta})$:

$$\boldsymbol{\beta} = \underset{\boldsymbol{\beta} \in \mathbb{R}^k}{\operatorname{argmin}} d(\boldsymbol{\beta}). \quad (2.45)$$

Similar to the best linear predictor we are measuring accuracy by expected squared error. The difference is that the best linear predictor (2.21) selects $\boldsymbol{\beta}$ to minimize the expected squared prediction error, while the best linear approximation (2.45) selects $\boldsymbol{\beta}$ to minimize the expected squared approximation error.

Despite the different definitions, it turns out that the best linear predictor and the best linear approximation are identical. By the same steps as in (2.18) plus an application of conditional expectations we can find that

$$\boldsymbol{\beta} = (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1} \mathbb{E}(\mathbf{x}m(\mathbf{x})) \quad (2.46)$$

$$= (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1} \mathbb{E}(\mathbf{x}y) \quad (2.47)$$

(see Exercise 2.19). Thus (2.45) equals (2.21). We conclude that the definition (2.45) can be viewed as an alternative motivation for the linear projection coefficient.

2.25 Normal Regression

Suppose the variables (y, \mathbf{x}) are jointly normally distributed. Consider the best linear predictor of y given \mathbf{x}

$$y = \mathbf{x}'\boldsymbol{\beta} + e$$

$$\boldsymbol{\beta} = (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1} \mathbb{E}(\mathbf{x}y).$$

Since the error e is a linear transformation of the normal vector (y, \mathbf{x}) , it follows that (e, \mathbf{x}) is jointly normal, and since they are jointly normal and uncorrelated (since $\mathbb{E}(\mathbf{x}e) = 0$) they are also independent (see Appendix B.9). Independence implies that

$$\mathbb{E}(e | \mathbf{x}) = \mathbb{E}(e) = 0$$

and

$$\mathbb{E}(e^2 | \mathbf{x}) = \mathbb{E}(e^2) = \sigma^2$$

which are properties of a homoskedastic linear CEF.

We have shown that when (y, \mathbf{x}) are jointly normally distributed, they satisfy a normal linear CEF

$$y = \mathbf{x}'\boldsymbol{\beta} + e$$

where

$$e \sim N(0, \sigma^2)$$

is independent of \mathbf{x} .

This is an alternative (and traditional) motivation for the linear CEF model. This motivation has limited merit in econometric applications since economic data is typically non-normal.

2.26 Regression to the Mean

The term **regression** originated in an influential paper by Francis Galton published in 1886, where he examined the joint distribution of the stature (height) of parents and children. Effectively, he was estimating the conditional mean of children's height given their parent's height. Galton discovered that this conditional mean was approximately linear with a slope of $2/3$. This implies that *on average* a child's height is more mediocre (average) than his or her parent's height. Galton called this phenomenon **regression to the mean**, and the label **regression** has stuck to this day to describe most conditional relationships.

One of Galton's fundamental insights was to recognize that if the marginal distributions of y and x are the same (e.g. the heights of children and parents in a stable environment) then the regression slope in a linear projection is always less than one.

To be more precise, take the simple linear projection

$$y = x\beta + \alpha + e \tag{2.48}$$

where y equals the height of the child and x equals the height of the parent. Assume that y and x have the same mean, so that $\mu_y = \mu_x = \mu$. Then from (2.37)

$$\alpha = (1 - \beta)\mu$$

so we can write the linear projection (2.48) as

$$\mathcal{P}(y | x) = (1 - \beta)\mu + x\beta.$$

This shows that the projected height of the child is a weighted average of the population average height μ and the parent's height x , with the weight equal to the regression slope β . When the height distribution is stable across generations, so that $\text{var}(y) = \text{var}(x)$, then this slope is the simple correlation of y and x . Using (2.38)

$$\beta = \frac{\text{cov}(x, y)}{\text{var}(x)} = \text{corr}(x, y).$$

By the properties of correlation (e.g. equation (B.7) in the Appendix), $-1 \leq \text{corr}(x, y) \leq 1$, with $\text{corr}(x, y) = 1$ only in the degenerate case $y = x$. Thus if we exclude degeneracy, β is strictly less than 1.

This means that on average a child's height is more mediocre (closer to the population average) than the parent's.

Sir Francis Galton

Sir Francis Galton (1822-1911) of England was one of the leading figures in late 19th century statistics. In addition to inventing the concept of regression, he is credited with introducing the concepts of correlation, the standard deviation, and the bivariate normal distribution. His work on heredity made a significant intellectual advance by examining the joint distributions of observables, allowing the application of the tools of mathematical statistics to the social sciences.

A common error – known as the **regression fallacy** – is to infer from $\beta < 1$ that the population is **converging**, meaning that its variance is declining towards zero. This is a fallacy because we derived the implication $\beta < 1$ under the assumption of constant means and variances. So certainly $\beta < 1$ does not imply that the variance y is less than than the variance of x .

Another way of seeing this is to examine the conditions for convergence in the context of equation (2.48). Since x and e are uncorrelated, it follows that

$$\text{var}(y) = \beta^2 \text{var}(x) + \text{var}(e).$$

Then $\text{var}(y) < \text{var}(x)$ if and only if

$$\beta^2 < 1 - \frac{\text{var}(e)}{\text{var}(x)}$$

which is not implied by the simple condition $|\beta| < 1$.

The regression fallacy arises in related empirical situations. Suppose you sort families into groups by the heights of the parents, and then plot the average heights of each subsequent generation over time. If the population is stable, the regression property implies that the plots lines will converge – children’s height will be more average than their parents. The regression fallacy is to incorrectly conclude that the population is converging. A message to be learned from this example is that such plots are misleading for inferences about convergence.

The regression fallacy is subtle. It is easy for intelligent economists to succumb to its temptation. A famous example is *The Triumph of Mediocrity in Business* by Horace Secrist, published in 1933. In this book, Secrist carefully and with great detail documented that in a sample of department stores over 1920-1930, when he divided the stores into groups based on 1920-1921 profits, and plotted the average profits of these groups for the subsequent 10 years, he found clear and persuasive evidence for convergence “toward mediocrity”. Of course, there was no discovery – regression to the mean is a necessary feature of stable distributions.

2.27 Reverse Regression

Galton noticed another interesting feature of the bivariate distribution. There is nothing special about a regression of y on x . We can also regress x on y . (In his heredity example this is the best linear predictor of the height of parents given the height of their children.) This regression takes the form

$$x = y\beta^* + \alpha^* + e^*. \tag{2.49}$$

This is sometimes called the **reverse regression**. In this equation, the coefficients α^* , β^* and error e^* are defined by linear projection. In a stable population we find that

$$\beta^* = \text{corr}(x, y) = \beta$$

$$\alpha^* = (1 - \beta) \mu = \alpha$$

which are exactly the same as in the projection of y on x ! The intercept and slope have exactly the same values in the forward and reverse projections!

While this algebraic discovery is quite simple, it is counter-intuitive. Instead, a common yet mistaken guess for the form of the reverse regression is to take the equation (2.48), divide through by β and rewrite to find the equation

$$x = y \frac{1}{\beta} - \frac{\alpha}{\beta} - \frac{1}{\beta} e \quad (2.50)$$

suggesting that the projection of x on y should have a slope coefficient of $1/\beta$ instead of β , and intercept of $-\alpha/\beta$ rather than α . What went wrong? Equation (2.50) is perfectly valid, because it is a simple manipulation of the valid equation (2.48). The trouble is that (2.50) is neither a CEF nor a linear projection. Inverting a projection (or CEF) does not yield a projection (or CEF). Instead, (2.49) is a valid projection, not (2.50).

In any event, Galton's finding was that when the variables are standardized, the slope in both projections (y on x , and x and y) equals the correlation, and both equations exhibit regression to the mean. It is not a causal relation, but a natural feature of all joint distributions.

2.28 Limitations of the Best Linear Predictor

Let's compare the linear projection and linear CEF models.

From Theorem 2.8.1.4 we know that the CEF error has the property $\mathbb{E}(\mathbf{x}e) = \mathbf{0}$. Thus a linear CEF is a linear projection. However, the converse is not true as the projection error does not necessarily satisfy $\mathbb{E}(e | \mathbf{x}) = 0$. Furthermore, the linear projection may be a poor approximation to the CEF.

To see these points in a simple example, suppose that the true process is $y = x + x^2$ with $x \sim N(0, 1)$. In this case the true CEF is $m(x) = x + x^2$ and there is no error. Now consider the linear projection of y on x and a constant, namely the model $y = \beta x + \alpha + u$. Since $x \sim N(0, 1)$ then x and x^2 are uncorrelated the linear projection takes the form $\mathcal{P}(y | x) = x + 1$. This is quite different from the true CEF $m(x) = x + x^2$. The projection error equals $e = x^2 - 1$, which is a deterministic function of x , yet is uncorrelated with x . We see in this example that a projection error need not be a CEF error, and a linear projection can be a poor approximation to the CEF.

Another defect of linear projection is that it is sensitive to the marginal distribution of the regressors when the conditional mean is non-linear. We illustrate the issue in Figure 2.10 for a constructed¹² joint distribution of y and x . The solid line is the non-linear CEF of y given x . The data are divided in two – Group 1 and Group 2 – which have different marginal distributions for the regressor x , and Group 1 has a lower mean value of x than Group 2. The separate linear projections of y on x for these two groups are displayed in the Figure by the dashed lines. These two projections are distinct approximations to the CEF. A defect with linear projection is that it leads to the incorrect conclusion that the effect of x on y is different for individuals in the two groups. This conclusion is incorrect because in fact there is no difference in the conditional mean function. The apparent difference is a by-product of a linear approximation to a non-linear mean, combined with different marginal distributions for the conditioning variables.

2.29 Random Coefficient Model

A model which is notationally similar to but conceptually distinct from the linear CEF model is the linear random coefficient model. It takes the form

$$y = \mathbf{x}'\boldsymbol{\eta}$$

¹²The x in Group 1 are $N(2, 1)$ and those in Group 2 are $N(4, 1)$, and the conditional distribution of y given x is $N(m(x), 1)$ where $m(x) = 2x - x^2/6$.

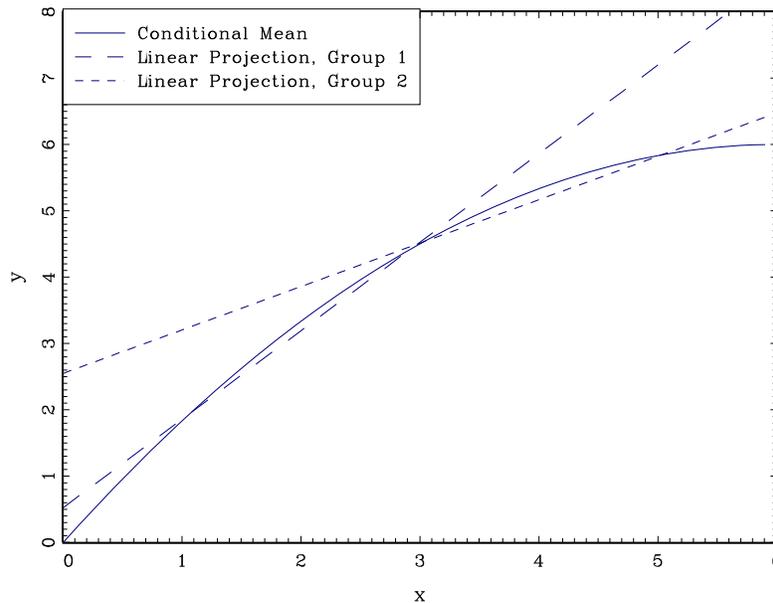


Figure 2.10: Conditional Mean and Two Linear Projections

where the individual-specific coefficient $\boldsymbol{\eta}$ is random and independent of \boldsymbol{x} . For example, if \boldsymbol{x} is years of schooling and y is log wages, then $\boldsymbol{\eta}$ is the individual-specific returns to schooling. If a person obtains an extra year of schooling, $\boldsymbol{\eta}$ is the actual change in their wage. The random coefficient model allows the returns to schooling to vary in the population. Some individuals might have a high return to education (a high $\boldsymbol{\eta}$) and others a low return, possibly 0, or even negative.

In the linear CEF model the regressor coefficient equals the regression derivative – the change in the conditional mean due to a change in the regressors, $\boldsymbol{\beta} = \nabla m(\boldsymbol{x})$. This is not the effect on a given individual, it is the effect on the population average. In contrast, in the random coefficient model, the random vector $\boldsymbol{\eta} = \nabla(\boldsymbol{x}'\boldsymbol{\eta})$ is the true causal effect – the change in the response variable y itself due to a change in the regressors.

It is interesting, however, to discover that the linear random coefficient model implies a linear CEF. To see this, let $\boldsymbol{\beta}$ and $\boldsymbol{\Sigma}$ denote the mean and covariance matrix of $\boldsymbol{\eta}$:

$$\begin{aligned}\boldsymbol{\beta} &= \mathbb{E}(\boldsymbol{\eta}) \\ \boldsymbol{\Sigma} &= \text{var}(\boldsymbol{\eta})\end{aligned}$$

and then decompose the random coefficient as

$$\boldsymbol{\eta} = \boldsymbol{\beta} + \boldsymbol{u}$$

where \boldsymbol{u} is distributed independently of \boldsymbol{x} with mean zero and covariance matrix $\boldsymbol{\Sigma}$. Then we can write

$$\mathbb{E}(y \mid \boldsymbol{x}) = \boldsymbol{x}'\mathbb{E}(\boldsymbol{\eta} \mid \boldsymbol{x}) = \boldsymbol{x}'\mathbb{E}(\boldsymbol{\eta}) = \boldsymbol{x}'\boldsymbol{\beta}$$

so the CEF is linear in \boldsymbol{x} , and the coefficients $\boldsymbol{\beta}$ equal the mean of the random coefficient $\boldsymbol{\eta}$.

We can thus write the equation as a linear CEF

$$y = \boldsymbol{x}'\boldsymbol{\beta} + e \tag{2.51}$$

where $e = \boldsymbol{x}'\boldsymbol{u}$ and $\boldsymbol{u} = \boldsymbol{\eta} - \boldsymbol{\beta}$. The error is conditionally mean zero:

$$\mathbb{E}(e \mid \boldsymbol{x}) = 0.$$

Furthermore

$$\begin{aligned}\text{var}(e \mid \mathbf{x}) &= \mathbf{x}' \text{var}(\boldsymbol{\eta}) \mathbf{x} \\ &= \mathbf{x}' \boldsymbol{\Sigma} \mathbf{x}\end{aligned}$$

so the error is conditionally heteroskedastic with its variance a quadratic function of \mathbf{x} .

Theorem 2.29.1 *In the linear random coefficient model $y = \mathbf{x}'\boldsymbol{\eta}$ with $\boldsymbol{\eta}$ independent of \mathbf{x} , $\mathbb{E}\|\mathbf{x}\|^2 < \infty$, and $\mathbb{E}\|\boldsymbol{\eta}\|^2 < \infty$, then*

$$\begin{aligned}\mathbb{E}(y \mid \mathbf{x}) &= \mathbf{x}'\boldsymbol{\beta} \\ \text{var}(y \mid \mathbf{x}) &= \mathbf{x}'\boldsymbol{\Sigma}\mathbf{x}\end{aligned}$$

where $\boldsymbol{\beta} = \mathbb{E}(\boldsymbol{\eta})$ and $\boldsymbol{\Sigma} = \text{var}(\boldsymbol{\eta})$.

2.30 Causal Effects

So far we have avoided the concept of causality, yet often the underlying goal of an econometric analysis is to uncover a causal relationship between variables. It is often of great interest to understand the causes and effects of decisions, actions, and policies. For example, we may be interested in the effect of class sizes on test scores, police expenditures on crime rates, climate change on economic activity, years of schooling on wages, institutional structure on growth, the effectiveness of rewards on behavior, the consequences of medical procedures for health outcomes, or any variety of possible causal relationships. In each case, the goal is to understand what is the actual effect on the outcome y due to a change in the input x . We are not just interested in the conditional mean or linear projection, we would like to know the actual change.

Two inherent barriers are that the causal effect is typically specific to an individual and that it is unobserved.

Consider the effect of schooling on wages. The causal effect is the actual difference a person would receive in wages if we could change their level of education *holding all else constant*. This is specific to each individual as their employment outcomes in these two distinct situations is individual. The causal effect is unobserved because the most we can observe is their actual level of education and their actual wage, but not the counterfactual wage if their education had been different.

To be even more specific, suppose that there are two individuals, Jennifer and George, and both have the possibility of being high-school graduates or college graduates, but both would have received different wages given their choices. For example, suppose that Jennifer would have earned \$10 an hour as a high-school graduate and \$20 an hour as a college graduate while George would have earned \$8 as a high-school graduate and \$12 as a college graduate. In this example the causal effect of schooling is \$10 a hour for Jennifer and \$4 an hour for George. The causal effects are specific to the individual and neither causal effect is observed.

A variable x_1 can be said to have a causal effect on the response variable y if the latter changes when all other inputs are held constant. To make this precise we need a mathematical formulation. We can write a full model for the response variable y as

$$y = h(x_1, \mathbf{x}_2, \mathbf{u}) \tag{2.52}$$

where x_1 and \mathbf{x}_2 are the observed variables, \mathbf{u} is an $\ell \times 1$ unobserved random factor, and h is a functional relationship. This framework includes as a special case the random coefficient model

(2.29) studied earlier. We define the causal effect of x_1 within this model as the change in y due to a change in x_1 holding the other variables \mathbf{x}_2 and \mathbf{u} constant.

Definition 2.30.1 *In the model (2.52) the **causal effect** of x_1 on y is*

$$C(x_1, \mathbf{x}_2, \mathbf{u}) = \nabla_1 h(x_1, \mathbf{x}_2, \mathbf{u}), \quad (2.53)$$

the change in y due to a change in x_1 , holding \mathbf{x}_2 and \mathbf{u} constant.

To understand this concept, imagine taking a single individual. As far as our structural model is concerned, this person is described by their observables x_1 and \mathbf{x}_2 and their unobservables \mathbf{u} . In a wage regression the unobservables would include characteristics such as the person's abilities, skills, work ethic, interpersonal connections, and preferences. The causal effect of x_1 (say, education) is the change in the wage as x_1 changes, holding constant all other observables **and** unobservables.

It may be helpful to understand that (2.53) is a definition, and does not necessarily describe causality in a fundamental or experimental sense. Perhaps it would be more appropriate to label (2.53) as a **structural effect** (the effect within the structural model).

Sometimes it is useful to write this relationship as a potential outcome function

$$y(x_1) = h(x_1, \mathbf{x}_2, \mathbf{u})$$

where the notation implies that $y(x_1)$ is holding \mathbf{x}_2 and \mathbf{u} constant.

A popular example arises in the analysis of treatment effects with a binary regressor x_1 . Let $x_1 = 1$ indicate treatment (e.g. a medical procedure) and $x_1 = 0$ indicate non-treatment. In this case $y(x_1)$ can be written

$$\begin{aligned} y(0) &= h(0, \mathbf{x}_2, \mathbf{u}) \\ y(1) &= h(1, \mathbf{x}_2, \mathbf{u}). \end{aligned}$$

In the literature on treatment effects, it is common to refer to $y(0)$ and $y(1)$ as the latent outcomes associated with non-treatment and treatment, respectively. That is, for a given individual, $y(0)$ is the health outcome if there is no treatment, and $y(1)$ is the health outcome if there is treatment. The causal effect of treatment for the individual is the change in their health outcome due to treatment – the change in y as we hold both \mathbf{x}_2 and \mathbf{u} constant:

$$C(\mathbf{x}_2, \mathbf{u}) = y(1) - y(0).$$

This is random (a function of \mathbf{x}_2 and \mathbf{u}) as both potential outcomes $y(0)$ and $y(1)$ are different across individuals.

In a sample, we cannot observe both outcomes from the same individual, we only observe the realized value

$$y = \begin{cases} y(0) & \text{if } x_1 = 0 \\ y(1) & \text{if } x_1 = 1. \end{cases}$$

As the causal effect varies across individuals and is not observable, it cannot be measured on the individual level. We therefore focus on aggregate causal effects, in particular what is known as the average causal effect.

Definition 2.30.2 In the model (2.52) the *average causal effect* of x_1 on y conditional on \mathbf{x}_2 is

$$\begin{aligned} ACE(x_1, \mathbf{x}_2) &= \mathbb{E}(C(x_1, \mathbf{x}_2, \mathbf{u}) \mid x_1, \mathbf{x}_2) \\ &= \int_{\mathbb{R}^\ell} \nabla_1 h(x_1, \mathbf{x}_2, \mathbf{u}) f(\mathbf{u} \mid x_1, \mathbf{x}_2) d\mathbf{u} \end{aligned} \quad (2.54)$$

where $f(\mathbf{u} \mid x_1, \mathbf{x}_2)$ is the conditional density of \mathbf{u} given x_1, \mathbf{x}_2 .

We can think of the average causal effect $ACE(x_1, \mathbf{x}_2)$ as the average effect in the general population. In our Jennifer & George schooling example given earlier, supposing that half of the population are Jennifer's and the other half George's, then the average causal effect of college is $(10+4)/2 = \$7$ an hour. This is not the individual causal effect, it is the average of the causal effect across all individuals in the population. Given data on only educational attainment and wages, the ACE of \$7 is the best we can hope to learn.

When we conduct a regression analysis (that is, consider the regression of observed wages on educational attainment) we might hope that the regression reveals the average causal effect. Technically, that the regression derivative (the coefficient on education) equals the ACE. Is this the case? In other words, what is the relationship between the average causal effect $ACE(x_1, \mathbf{x}_2)$ and the regression derivative $\nabla_1 m(x_1, \mathbf{x}_2)$? Equation (2.52) implies that the CEF is

$$\begin{aligned} m(x_1, \mathbf{x}_2) &= \mathbb{E}(h(x_1, \mathbf{x}_2, \mathbf{u}) \mid x_1, \mathbf{x}_2) \\ &= \int_{\mathbb{R}^\ell} h(x_1, \mathbf{x}_2, \mathbf{u}) f(\mathbf{u} \mid x_1, \mathbf{x}_2) d\mathbf{u}, \end{aligned}$$

the average causal equation, averaged over the conditional distribution of the unobserved component \mathbf{u} .

Applying the marginal effect operator, the regression derivative is

$$\begin{aligned} \nabla_1 m(x_1, \mathbf{x}_2) &= \int_{\mathbb{R}^\ell} \nabla_1 h(x_1, \mathbf{x}_2, \mathbf{u}) f(\mathbf{u} \mid x_1, \mathbf{x}_2) d\mathbf{u} \\ &\quad + \int_{\mathbb{R}^\ell} h(x_1, \mathbf{x}_2, \mathbf{u}) \nabla_1 f(\mathbf{u} \mid x_1, \mathbf{x}_2) d\mathbf{u} \\ &= ACE(x_1, \mathbf{x}_2) + \int_{\mathbb{R}^\ell} h(x_1, \mathbf{x}_2, \mathbf{u}) \nabla_1 f(\mathbf{u} \mid x_1, \mathbf{x}_2) d\mathbf{u}. \end{aligned} \quad (2.55)$$

Equation (2.55) shows that in general, the regression derivative does not equal the average causal effect. The difference is the second term on the right-hand-side of (2.55). The regression derivative and ACE equal in the special case when this term equals zero, which occurs when $\nabla_1 f(\mathbf{u} \mid x_1, \mathbf{x}_2) = 0$, that is, when the conditional density of \mathbf{u} given (x_1, \mathbf{x}_2) does not depend on x_1 . When this condition holds then the regression derivative equals the ACE, which means that regression analysis can be interpreted causally, in the sense that it uncovers average causal effects.

The condition is sufficiently important that it has a special name in the treatment effects literature.

Definition 2.30.3 *Conditional Independence Assumption (CIA).* Conditional on \mathbf{x}_2 , the random variables x_1 and \mathbf{u} are statistically independent.

The CIA implies $f(\mathbf{u} \mid x_1, \mathbf{x}_2) = f(\mathbf{u} \mid \mathbf{x}_2)$ does not depend on x_1 , and thus $\nabla_1 f(\mathbf{u} \mid x_1, \mathbf{x}_2) = 0$. Thus the CIA implies that $\nabla_1 m(x_1, \mathbf{x}_2) = ACE(x_1, \mathbf{x}_2)$, the regression derivative equals the average causal effect.

Theorem 2.30.1 *In the structural model (2.52), the Conditional Independence Assumption implies*

$$\nabla_1 m(x_1, \mathbf{x}_2) = ACE(x_1, \mathbf{x}_2)$$

the regression derivative equals the average causal effect for x_1 on y conditional on \mathbf{x}_2 .

This is a fascinating result. It shows that whenever the unobservable is independent of the treatment variable (after conditioning on appropriate regressors) the regression derivative equals the average causal effect. In this case, the CEF has causal economic meaning, giving strong justification to estimation of the CEF. Our derivation also shows the critical role of the CIA. If CIA fails, then the equality of the regression derivative and ACE fails.

This theorem is quite general. It applies equally to the treatment-effects model where x_1 is binary or to more general settings where x_1 is continuous.

It is also helpful to understand that the CIA is weaker than full independence of \mathbf{u} from the regressors (x_1, \mathbf{x}_2) . The CIA was introduced precisely as a minimal sufficient condition to obtain the desired result. Full independence implies the CIA and implies that each regression derivative equals that variable's average causal effect, but full independence is not necessary in order to causally interpret a subset of the regressors.

To illustrate, let's return to our education example involving a population with equal numbers of Jennifer's and George's. Recall that Jennifer earns \$10 as a high-school graduate and \$20 as a college graduate (and so has a causal effect of \$10) while George earns \$8 as a high-school graduate and \$12 as a college graduate (so has a causal effect of \$4). Given this information, the average causal effect of college is \$7, which is what we hope to learn from a regression analysis.

Now suppose that while in high school all students take an aptitude test, and if a student gets a high (H) score he or she goes to college with probability 3/4, and if a student gets a low (L) score he or she goes to college with probability 1/4. Suppose further that Jennifer's get an aptitude score of H with probability 3/4, while George's get a score of H with probability 1/4. Given this situation, 62.5% of Jennifer's will go to college¹³, while 37.5% of George's will go to college¹⁴.

An econometrician who randomly samples 32 individuals and collects data on educational attainment and wages will find the following wage distribution:

	\$8	\$10	\$12	\$20	Mean
High-School Graduate	10	6	0	0	\$8.75
College Graduate	0	0	6	10	\$17.00

Let *college* denote a dummy variable taking the value of 1 for a college graduate, otherwise 0. Thus the regression of wages on college attendance takes the form

$$\mathbb{E}(\text{wage} \mid \text{college}) = 8.25\text{college} + 8.75.$$

The coefficient on the college dummy, \$8.25, is the regression derivative, and the implied wage effect of college attendance. But \$8.25 overstates the average causal effect of \$7. The reason is because

¹³ $\Pr(\text{College} \mid \text{Jennifer}) = \Pr(\text{College} \mid H) \Pr(H \mid \text{Jennifer}) + \Pr(\text{College} \mid L) \Pr(L \mid \text{Jennifer}) = (3/4)^2 + (1/4)^2$

¹⁴ $\Pr(\text{College} \mid \text{George}) = \Pr(\text{College} \mid H) \Pr(H \mid \text{George}) + \Pr(\text{College} \mid L) \Pr(L \mid \text{George}) = (3/4)(1/4) + (1/4)(3/4)$

the CIA fails. In this model the unobservable \mathbf{u} is the individual's type (Jennifer or George) which is not independent of the regressor x_1 (education), since Jennifer is more likely to go to college than George. Since Jennifer's causal effect is higher than George's, the regression derivative overstates the ACE. The coefficient \$8.25 is not the average benefit of college attendance, rather it is the observed difference in realized wages in a population whose decision to attend college is correlated with their individual causal effect. At the risk of repeating myself, in this example, \$8.25 is the true regression derivative, it is the difference in average wages between those with a college education and those without. It is not, however, the average causal effect of college education in the population.

This does not mean that it is impossible to estimate the ACE. The key is conditioning on the appropriate variables. The CIA says that we need to find a variable x_2 such that conditional on x_2 , \mathbf{u} and x_1 (type and education) are independent. In this example a variable which will achieve this is the aptitude test score. The decision to attend college was based on the test score, not on an individual's type. Thus educational attainment and type are independent once we condition on the test score.

This also alters the ACE. Notice that Definition 2.30.2 is a function of x_2 (the test score). Among the students who receive a high test score, 3/4 are Jennifer's and 1/4 are George's. Thus the ACE for students with a score of H is $(3/4) \times 10 + (1/4) \times 4 = \8.50 . Among the students who receive a low test score, 1/4 are Jennifer's and 3/4 are George's. Thus the ACE for students with a score of L is $(1/4) \times 10 + (3/4) \times 4 = \5.50 . The ACE varies between these two observable groups (those with high test scores and those with low test scores). Again, we would hope to be able to learn the ACE from a regression analysis, this time from a regression of wages on education and test scores.

To see this in the wage distribution, suppose that the econometrician collects data on the aptitude test score as well as education and wages. Given a random sample of 32 individuals we would expect to find the following wage distribution:

	\$8	\$10	\$12	\$20	Mean
High-School Graduate + High Test Score	1	3	0	0	\$9.50
College Graduate + High Test Score	0	0	3	9	\$18.00
High-School Graduate + Low Test Score	9	3	0	0	\$8.50
College Graduate + Low Test Score	0	0	3	1	\$14.00

Define the dummy variable *highscore* which takes the value 1 for students who received a high test score, else zero. The regression of wages on college attendance and test scores (with interactions) takes the form

$$\mathbb{E}(\text{wage} \mid \text{college}, \text{highscore}) = 1.00\text{highscore} + 5.50\text{college} + 3.00\text{highscore} \times \text{college} + 8.50.$$

The coefficient on *college*, \$5.50, is the regression derivative of college attendance for those with low test scores, and the sum of this coefficient with the interaction coefficient, \$8.50, is the regression derivative for college attendance for those with high test scores. These equal the average causal effect.

In this example, by conditioning on the aptitude test score, the average causal effect of education on wages can be learned from a regression analysis. What this shows is that by conditioning on the proper variables, it may be possible to achieve the CIA, in which case regression analysis measures average causal effects.

2.31 Expectation: Mathematical Details*

We define the **mean** or **expectation** $\mathbb{E}y$ of a random variable y as follows. If y is discrete on the set $\{\tau_1, \tau_2, \dots\}$ then

$$\mathbb{E}y = \sum_{j=1}^{\infty} \tau_j \Pr(y = \tau_j),$$

and if y is continuous with density f then

$$\mathbb{E}y = \int_{-\infty}^{\infty} yf(y)dy.$$

We can unify these definitions by writing the expectation as the Lebesgue integral with respect to the distribution function F

$$\mathbb{E}y = \int_{-\infty}^{\infty} ydF(y). \quad (2.56)$$

In the event that the integral (2.56) is not finite, separately evaluate the two integrals

$$I_1 = \int_0^{\infty} ydF(y) \quad (2.57)$$

$$I_2 = - \int_{-\infty}^0 ydF(y). \quad (2.58)$$

If $I_1 = \infty$ and $I_2 < \infty$ then it is typical to define $\mathbb{E}y = \infty$. If $I_1 < \infty$ and $I_2 = \infty$ then we define $\mathbb{E}y = -\infty$. However, if both $I_1 = \infty$ and $I_2 = \infty$ then $\mathbb{E}y$ is undefined. If

$$\mathbb{E}|y| = \int_{-\infty}^{\infty} |y| dF(y) = I_1 + I_2 < \infty$$

then $\mathbb{E}y$ exists and is finite. In this case it is common to say that the mean $\mathbb{E}y$ is “well-defined”.

More generally, y has a finite r 'th moment if

$$\mathbb{E}|y|^r < \infty. \quad (2.59)$$

By Liapunov's Inequality (B.20), (2.59) implies $\mathbb{E}|y|^s < \infty$ for all $1 \leq s \leq r$. Thus, for example, if the fourth moment is finite then the first, second and third moments are also finite.

It is common in econometric theory to assume that the variables, or certain transformations of the variables, have finite moments of a certain order. How should we interpret this assumption? How restrictive is it?

One way to visualize the importance is to consider the class of Pareto densities given by

$$f(y) = ay^{-a-1}, \quad y > 1.$$

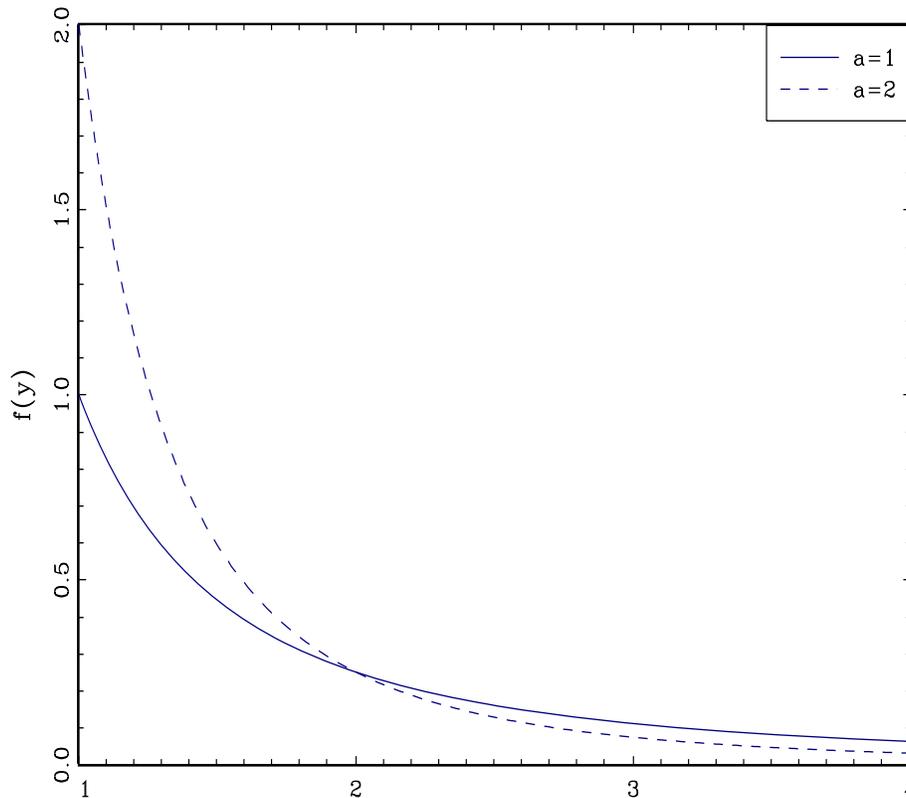
The parameter a of the Pareto distribution indexes the rate of decay of the tail of the density. Larger a means that the tail declines to zero more quickly. See Figure 2.11 below where we show the Pareto density for $a = 1$ and $a = 2$. The parameter a also determines which moments are finite. We can calculate that

$$\mathbb{E}|y|^r = \begin{cases} a \int_1^{\infty} y^{r-a-1} dy = \frac{a}{a-r} & \text{if } r < a \\ \infty & \text{if } r \geq a. \end{cases}$$

This shows that if y is Pareto distributed with parameter a , then the r 'th moment of y is finite if and only if $r < a$. Higher a means higher finite moments. Equivalently, the faster the tail of the density declines to zero, the more moments are finite.

This connection between tail decay and finite moments is not limited to the Pareto distribution. We can make a similar analysis using a tail bound. Suppose that y has density $f(y)$ which satisfies the bound $f(y) \leq A|y|^{-a-1}$ for some $A < \infty$ and $a > 0$. Since $f(y)$ is bounded below a scale of a Pareto density, its tail behavior is similarly bounded. This means that for $r < a$

$$\mathbb{E}|y|^r = \int_{-\infty}^{\infty} |y|^r f(y) dy \leq \int_{-1}^1 f(y) dy + 2A \int_1^{\infty} y^{r-a-1} dy \leq 1 + \frac{2A}{a-r} < \infty.$$

Figure 2.11: Pareto Densities, $a = 1$ and $a = 2$

Thus if the tail of the density declines at the rate $|y|^{-a-1}$ or faster, then y has finite moments up to (but not including) a . Broadly speaking, the restriction that y has a finite r^{th} moment means that the tail of y 's density declines to zero faster than y^{-r-1} . The faster decline of the tail means that the probability of observing an extreme value of y is a more rare event.

We complete this section by adding an alternative representation of expectation in terms of the distribution function.

Theorem 2.31.1 For any non-negative random variable y

$$\mathbb{E}y = \int_0^\infty \Pr(y > u) du$$

Proof of Theorem 2.31.1: Let $F^*(x) = \Pr(y > x) = 1 - F(x)$, where $F(x)$ is the distribution function. By integration by parts

$$\mathbb{E}y = \int_0^\infty y dF(y) = - \int_0^\infty y dF^*(y) = - [yF^*(y)]_0^\infty + \int_0^\infty F^*(y) dy = \int_0^\infty \Pr(y > u) du$$

as stated. ■

2.32 Existence and Uniqueness of the Conditional Expectation*

In Sections 2.3 and 2.6 we defined the conditional mean when the conditioning variables \mathbf{x} are discrete and when the variables (y, \mathbf{x}) have a joint density. We have explored these cases because

these are the situations where the conditional mean is easiest to describe and understand. However, the conditional mean exists quite generally without appealing to the properties of either discrete or continuous random variables.

To justify this claim we now present a deep result from probability theory. What it says is that the conditional mean exists for all joint distributions (y, \mathbf{x}) for which y has a finite mean.

Theorem 2.32.1 Existence of the Conditional Mean

If $\mathbb{E}|y| < \infty$ then there exists a function $m(\mathbf{x})$ such that for all measurable sets \mathcal{X}

$$\mathbb{E}(1(\mathbf{x} \in \mathcal{X})y) = \mathbb{E}(1(\mathbf{x} \in \mathcal{X})m(\mathbf{x})). \quad (2.60)$$

The function $m(\mathbf{x})$ is almost everywhere unique, in the sense that if $h(\mathbf{x})$ satisfies (2.60), then there is a set S such that $\Pr(S) = 1$ and $m(\mathbf{x}) = h(\mathbf{x})$ for $\mathbf{x} \in S$. The function $m(\mathbf{x})$ is called the **conditional mean** and is written $m(\mathbf{x}) = \mathbb{E}(y | \mathbf{x})$.

See, for example, Ash (1972), Theorem 6.3.3.

The conditional mean $m(\mathbf{x})$ defined by (2.60) specializes to (2.7) when (y, \mathbf{x}) have a joint density. The usefulness of definition (2.60) is that Theorem 2.32.1 shows that the conditional mean $m(\mathbf{x})$ exists for all finite-mean distributions. This definition allows y to be discrete or continuous, for \mathbf{x} to be scalar or vector-valued, and for the components of \mathbf{x} to be discrete or continuously distributed.

2.33 Identification*

A critical and important issue in structural econometric modeling is identification, meaning that a parameter is uniquely determined by the distribution of the observed variables. It is relatively straightforward in the context of the unconditional and conditional mean, but it is worthwhile to introduce and explore the concept at this point for clarity.

Let F denote the distribution of the observed data, for example the distribution of the pair (y, x) . Let \mathcal{F} be a collection of distributions F . Let θ be a parameter of interest (for example, the mean $\mathbb{E}y$).

Definition 2.33.1 A parameter $\theta \in \mathbb{R}$ is identified on \mathcal{F} if for all $F \in \mathcal{F}$, there is a uniquely determined value of θ .

Equivalently, θ is identified if we can write it as a mapping $\theta = g(F)$ on the set \mathcal{F} . The restriction to the set \mathcal{F} is important. Most parameters are identified only on a strict subset of the space of all distributions.

Take, for example, the mean $\mu = \mathbb{E}y$. It is uniquely determined if $\mathbb{E}|y| < \infty$, so it is clear that μ is identified for the set $\mathcal{F} = \left\{ F : \int_{-\infty}^{\infty} |y| dF(y) < \infty \right\}$. However, μ is also well defined when it is either positive or negative infinity. Hence, defining I_1 and I_2 as in (2.57) and (2.58), we can deduce that μ is identified on the set $\mathcal{F} = \{F : \{I_1 < \infty\} \cup \{I_2 < \infty\}\}$.

Next, consider the conditional mean. Theorem 2.32.1 demonstrates that $\mathbb{E}|y| < \infty$ is a sufficient condition for identification.

Theorem 2.33.1 Identification of the Conditional Mean

If $\mathbb{E}|y| < \infty$, the conditional mean $m(\mathbf{x}) = \mathbb{E}(y | \mathbf{x})$ is identified almost everywhere.

It might seem as if identification is a general property for parameters, so long as we exclude degenerate cases. This is true for moments of observed data, but not necessarily for more complicated models. As a case in point, consider the context of censoring. Let y be a random variable with distribution F . Instead of observing y , we observe y^* defined by the censoring rule

$$y^* = \begin{cases} y & \text{if } y \leq \tau \\ \tau & \text{if } y > \tau \end{cases}.$$

That is, y^* is capped at the value τ . A common example is income surveys, where income responses are “top-coded”, meaning that incomes above the top code τ are recorded as equalling the top code. The observed variable y^* has distribution

$$F^*(u) = \begin{cases} F(u) & \text{for } u \leq \tau \\ 1 & \text{for } u \geq \tau. \end{cases}$$

We are interested in features of the distribution F not the censored distribution F^* . For example, we are interested in the mean wage $\mu = \mathbb{E}(y)$. The difficulty is that we cannot calculate μ from F^* except in the trivial case where there is no censoring $\Pr(y \geq \tau) = 0$. Thus the mean μ is not generically identified from the censored distribution.

A typical solution to the identification problem is to assume a parametric distribution. For example, let \mathcal{F} be the set of normal distributions $y \sim N(\mu, \sigma^2)$. It is possible to show that the parameters (μ, σ^2) are identified for all $F \in \mathcal{F}$. That is, if we know that the uncensored distribution is normal, we can uniquely determine the parameters from the censored distribution. This is often called **parametric identification** as identification is restricted to a parametric class of distributions. In modern econometrics this is generally viewed as a second-best solution, as identification has been achieved only through the use of an arbitrary and unverifiable parametric assumption.

A pessimistic conclusion might be that it is impossible to identify parameters of interest from censored data without parametric assumptions. Interestingly, this pessimism is unwarranted. It turns out that we can identify the quantiles q_α of F for $\alpha \leq \Pr(y \leq \tau)$. For example, if 20% of the distribution is censored, we can identify all quantiles for $\alpha \in (0, 0.8)$. This is often called **nonparametric identification** as the parameters are identified without restriction to a parametric class.

What we have learned from this little exercise is that in the context of censored data, moments can only be parametrically identified, while (non-censored) quantiles are nonparametrically identified. Part of the message is that a study of identification can help focus attention on what can be learned from the data distributions available.

2.34 Technical Proofs*

Proof of Theorem 2.7.1: For convenience, assume that the variables have a joint density $f(y, \mathbf{x})$. Since $\mathbb{E}(y | \mathbf{x})$ is a function of the random vector \mathbf{x} only, to calculate its expectation we integrate with respect to the density $f_{\mathbf{x}}(\mathbf{x})$ of \mathbf{x} , that is

$$\mathbb{E}(\mathbb{E}(y | \mathbf{x})) = \int_{\mathbb{R}^k} \mathbb{E}(y | \mathbf{x}) f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x}.$$

Substituting in (2.7) and noting that $f_{y|\mathbf{x}}(y|\mathbf{x})f_{\mathbf{x}}(\mathbf{x}) = f(y, \mathbf{x})$, we find that the above expression equals

$$\int_{\mathbb{R}^k} \left(\int_{\mathbb{R}} y f_{y|\mathbf{x}}(y|\mathbf{x}) dy \right) f_{\mathbf{x}}(\mathbf{x}) d\mathbf{x} = \int_{\mathbb{R}^k} \int_{\mathbb{R}} y f(y, \mathbf{x}) dy d\mathbf{x} = \mathbb{E}(y)$$

the unconditional mean of y . ■

Proof of Theorem 2.7.2: Again assume that the variables have a joint density. It is useful to observe that

$$f(y|\mathbf{x}_1, \mathbf{x}_2) f(\mathbf{x}_2|\mathbf{x}_1) = \frac{f(y, \mathbf{x}_1, \mathbf{x}_2) f(\mathbf{x}_1, \mathbf{x}_2)}{f(\mathbf{x}_1, \mathbf{x}_2)} = f(y, \mathbf{x}_2|\mathbf{x}_1), \quad (2.61)$$

the density of (y, \mathbf{x}_2) given \mathbf{x}_1 . Here, we have abused notation and used a single symbol f to denote the various unconditional and conditional densities to reduce notational clutter.

Note that

$$\mathbb{E}(y | \mathbf{x}_1, \mathbf{x}_2) = \int_{\mathbb{R}} y f(y|\mathbf{x}_1, \mathbf{x}_2) dy. \quad (2.62)$$

Integrating (2.62) with respect to the conditional density of \mathbf{x}_2 given \mathbf{x}_1 , and applying (2.61) we find that

$$\begin{aligned} \mathbb{E}(\mathbb{E}(y | \mathbf{x}_1, \mathbf{x}_2) | \mathbf{x}_1) &= \int_{\mathbb{R}^{k_2}} \mathbb{E}(y | \mathbf{x}_1, \mathbf{x}_2) f(\mathbf{x}_2|\mathbf{x}_1) d\mathbf{x}_2 \\ &= \int_{\mathbb{R}^{k_2}} \left(\int_{\mathbb{R}} y f(y|\mathbf{x}_1, \mathbf{x}_2) dy \right) f(\mathbf{x}_2|\mathbf{x}_1) d\mathbf{x}_2 \\ &= \int_{\mathbb{R}^{k_2}} \int_{\mathbb{R}} y f(y|\mathbf{x}_1, \mathbf{x}_2) f(\mathbf{x}_2|\mathbf{x}_1) dy d\mathbf{x}_2 \\ &= \int_{\mathbb{R}^{k_2}} \int_{\mathbb{R}} y f(y, \mathbf{x}_2|\mathbf{x}_1) dy d\mathbf{x}_2 \\ &= \mathbb{E}(y | \mathbf{x}_1) \end{aligned}$$

as stated. ■

Proof of Theorem 2.7.3:

$$\mathbb{E}(g(\mathbf{x}) y | \mathbf{x}) = \int_{\mathbb{R}} g(\mathbf{x}) y f_{y|\mathbf{x}}(y|\mathbf{x}) dy = g(\mathbf{x}) \int_{\mathbb{R}} y f_{y|\mathbf{x}}(y|\mathbf{x}) dy = g(\mathbf{x}) \mathbb{E}(y | \mathbf{x})$$

This is (2.9). The assumption that $\mathbb{E}|g(\mathbf{x}) y| < \infty$ is required for the first equality to be well-defined. Equation (2.10) follows by applying the Simple Law of Iterated Expectations to (2.9). ■

Proof of Theorem 2.10.2: The assumption that $\mathbb{E}y^2 < \infty$ implies that all the conditional expectations below exist.

Set $z = \mathbb{E}(y | \mathbf{x}_1, \mathbf{x}_2)$. By the conditional Jensen's inequality (B.13),

$$(\mathbb{E}(z | \mathbf{x}_1))^2 \leq \mathbb{E}(z^2 | \mathbf{x}_1).$$

Taking unconditional expectations, this implies

$$\mathbb{E}(\mathbb{E}(y | \mathbf{x}_1))^2 \leq \mathbb{E}\left(\mathbb{E}(y | \mathbf{x}_1, \mathbf{x}_2)^2\right).$$

Similarly,

$$(\mathbb{E}y)^2 \leq \mathbb{E}\left(\mathbb{E}(y | \mathbf{x}_1)^2\right) \leq \mathbb{E}\left(\mathbb{E}(y | \mathbf{x}_1, \mathbf{x}_2)^2\right). \quad (2.63)$$

The variables y , $\mathbb{E}(y \mid \mathbf{x}_1)$ and $\mathbb{E}(y \mid \mathbf{x}_1, \mathbf{x}_2)$ all have the same mean $\mathbb{E}y$, so the inequality (2.63) implies that the variances are ranked monotonically:

$$0 \leq \text{var}(\mathbb{E}(y \mid \mathbf{x}_1)) \leq \text{var}(\mathbb{E}(y \mid \mathbf{x}_1, \mathbf{x}_2)). \quad (2.64)$$

Next, for $\mu = \mathbb{E}y$ observe that

$$\mathbb{E}(y - \mathbb{E}(y \mid \mathbf{x}))(\mathbb{E}(y \mid \mathbf{x}) - \mu) = \mathbb{E}(y - \mathbb{E}(y \mid \mathbf{x}))(\mathbb{E}(y \mid \mathbf{x}) - \mu) = 0$$

so the decomposition

$$y - \mu = y - \mathbb{E}(y \mid \mathbf{x}) + \mathbb{E}(y \mid \mathbf{x}) - \mu$$

satisfies

$$\text{var}(y) = \text{var}(y - \mathbb{E}(y \mid \mathbf{x})) + \text{var}(\mathbb{E}(y \mid \mathbf{x})). \quad (2.65)$$

The monotonicity of the variances of the conditional mean (2.64) applied to the variance decomposition (2.65) implies the reverse monotonicity of the variances of the differences, completing the proof. ■

Proof of Theorem 2.8.1. Applying Minkowski's Inequality (B.19) to $e = y - m(\mathbf{x})$,

$$(\mathbb{E}|e|^r)^{1/r} = (\mathbb{E}|y - m(\mathbf{x})|^r)^{1/r} \leq (\mathbb{E}|y|^r)^{1/r} + (\mathbb{E}|m(\mathbf{x})|^r)^{1/r} < \infty,$$

where the two parts on the right-hand are finite since $\mathbb{E}|y|^r < \infty$ by assumption and $\mathbb{E}|m(\mathbf{x})|^r < \infty$ by the Conditional Expectation Inequality (B.14). The fact that $(\mathbb{E}|e|^r)^{1/r} < \infty$ implies $\mathbb{E}|e|^r < \infty$. ■

Proof of Theorem 2.18.1. For part 1, by the Expectation Inequality (B.15), (A.9) and Assumption 2.18.1,

$$\|\mathbb{E}(\mathbf{x}\mathbf{x}')\| \leq \mathbb{E}\|\mathbf{x}\mathbf{x}'\| = \mathbb{E}\|\mathbf{x}\|^2 < \infty.$$

Similarly, using the Expectation Inequality (B.15), the Cauchy-Schwarz Inequality (B.17) and Assumption 2.18.1,

$$\|\mathbb{E}(\mathbf{x}y)\| \leq \mathbb{E}\|\mathbf{x}y\| \leq \left(\mathbb{E}\|\mathbf{x}\|^2\right)^{1/2} (\mathbb{E}y^2)^{1/2} < \infty.$$

Thus the moments $\mathbb{E}(\mathbf{x}y)$ and $\mathbb{E}(\mathbf{x}\mathbf{x}')$ are finite and well defined.

For part 2, the coefficient $\boldsymbol{\beta} = (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1}\mathbb{E}(\mathbf{x}y)$ is well defined since $(\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1}$ exists under Assumption 2.18.1.

Part 3 follows from Definition 2.18.1 and part 2.

For part 4, first note that

$$\begin{aligned} \mathbb{E}e^2 &= \mathbb{E}(y - \mathbf{x}'\boldsymbol{\beta})^2 \\ &= \mathbb{E}y^2 - 2\mathbb{E}(y\mathbf{x}')\boldsymbol{\beta} + \boldsymbol{\beta}'\mathbb{E}(\mathbf{x}\mathbf{x}')\boldsymbol{\beta} \\ &= \mathbb{E}y^2 - 2\mathbb{E}(y\mathbf{x}')(\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1}\mathbb{E}(\mathbf{x}y) \\ &\leq \mathbb{E}y^2 \\ &< \infty. \end{aligned}$$

The first inequality holds because $\mathbb{E}(y\mathbf{x}')(\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1}\mathbb{E}(\mathbf{x}y)$ is a quadratic form and therefore necessarily non-negative. Second, by the Expectation Inequality (B.15), the Cauchy-Schwarz Inequality (B.17) and Assumption 2.18.1,

$$\|\mathbb{E}(\mathbf{x}e)\| \leq \mathbb{E}\|\mathbf{x}e\| = \left(\mathbb{E}\|\mathbf{x}\|^2\right)^{1/2} (\mathbb{E}e^2)^{1/2} < \infty.$$

It follows that the expectation $\mathbb{E}(\mathbf{x}e)$ is finite, and is zero by the calculation (2.28).

For part 6, Applying Minkowski's Inequality (B.19) to $e = y - \mathbf{x}'\boldsymbol{\beta}$,

$$\begin{aligned}(\mathbb{E}|e|^r)^{1/r} &= (\mathbb{E}|y - \mathbf{x}'\boldsymbol{\beta}|^r)^{1/r} \\ &\leq (\mathbb{E}|y|^r)^{1/r} + (\mathbb{E}|\mathbf{x}'\boldsymbol{\beta}|^r)^{1/r} \\ &\leq (\mathbb{E}|y|^r)^{1/r} + (\mathbb{E}\|\mathbf{x}\|^r)^{1/r} \|\boldsymbol{\beta}\| \\ &< \infty,\end{aligned}$$

the final inequality by assumption. ■

Exercises

Exercise 2.1 Find $\mathbb{E}(\mathbb{E}(\mathbb{E}(y \mid \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) \mid \mathbf{x}_1, \mathbf{x}_2) \mid \mathbf{x}_1)$.

Exercise 2.2 If $\mathbb{E}(y \mid x) = a + bx$, find $\mathbb{E}(yx)$ as a function of moments of x .

Exercise 2.3 Prove Theorem 2.8.1.4 using the law of iterated expectations.

Exercise 2.4 Suppose that the random variables y and x only take the values 0 and 1, and have the following joint probability distribution

	$x = 0$	$x = 1$
$y = 0$.1	.2
$y = 1$.4	.3

Find $\mathbb{E}(y \mid x)$, $\mathbb{E}(y^2 \mid x)$ and $\text{var}(y \mid x)$ for $x = 0$ and $x = 1$.

Exercise 2.5 Show that $\sigma^2(\mathbf{x})$ is the best predictor of e^2 given \mathbf{x} :

- Write down the mean-squared error of a predictor $h(\mathbf{x})$ for e^2 .
- What does it mean to be predicting e^2 ?
- Show that $\sigma^2(\mathbf{x})$ minimizes the mean-squared error and is thus the best predictor.

Exercise 2.6 Use $y = m(\mathbf{x}) + e$ to show that

$$\text{var}(y) = \text{var}(m(\mathbf{x})) + \sigma^2$$

Exercise 2.7 Show that the conditional variance can be written as

$$\sigma^2(\mathbf{x}) = \mathbb{E}(y^2 \mid \mathbf{x}) - (\mathbb{E}(y \mid \mathbf{x}))^2.$$

Exercise 2.8 Suppose that y is discrete-valued, taking values only on the non-negative integers, and the conditional distribution of y given \mathbf{x} is Poisson:

$$\Pr(y = j \mid \mathbf{x}) = \frac{\exp(-\mathbf{x}'\boldsymbol{\beta}) (\mathbf{x}'\boldsymbol{\beta})^j}{j!}, \quad j = 0, 1, 2, \dots$$

Compute $\mathbb{E}(y \mid \mathbf{x})$ and $\text{var}(y \mid \mathbf{x})$. Does this justify a linear regression model of the form $y = \mathbf{x}'\boldsymbol{\beta} + e$?

Hint: If $\Pr(y = j) = \frac{\exp(-\lambda)\lambda^j}{j!}$, then $\mathbb{E}y = \lambda$ and $\text{var}(y) = \lambda$.

Exercise 2.9 Suppose you have two regressors: x_1 is binary (takes values 0 and 1) and x_2 is categorical with 3 categories (A, B, C). Write $\mathbb{E}(y \mid x_1, x_2)$ as a linear regression.

Exercise 2.10 True or False. If $y = x\beta + e$, $x \in \mathbb{R}$, and $\mathbb{E}(e \mid x) = 0$, then $\mathbb{E}(x^2e) = 0$.

Exercise 2.11 True or False. If $y = x\beta + e$, $x \in \mathbb{R}$, and $\mathbb{E}(xe) = 0$, then $\mathbb{E}(x^2e) = 0$.

Exercise 2.12 True or False. If $y = \mathbf{x}'\boldsymbol{\beta} + e$ and $\mathbb{E}(e \mid \mathbf{x}) = 0$, then e is independent of \mathbf{x} .

Exercise 2.13 True or False. If $y = \mathbf{x}'\boldsymbol{\beta} + e$ and $\mathbb{E}(xe) = \mathbf{0}$, then $\mathbb{E}(e \mid \mathbf{x}) = 0$.

Exercise 2.14 True or False. If $y = \mathbf{x}'\boldsymbol{\beta} + e$, $\mathbb{E}(e | \mathbf{x}) = 0$, and $\mathbb{E}(e^2 | \mathbf{x}) = \sigma^2$, a constant, then e is independent of \mathbf{x} .

Exercise 2.15 Consider the intercept-only model $y = \alpha + e$ defined as the best linear predictor. Show that $\alpha = \mathbb{E}(y)$.

Exercise 2.16 Let x and y have the joint density $f(x, y) = \frac{3}{2}(x^2 + y^2)$ on $0 \leq x \leq 1$, $0 \leq y \leq 1$. Compute the coefficients of the best linear predictor $y = \alpha + \beta x + e$. Compute the conditional mean $m(x) = \mathbb{E}(y | x)$. Are the best linear predictor and conditional mean different?

Exercise 2.17 Let x be a random variable with $\mu = \mathbb{E}x$ and $\sigma^2 = \text{var}(x)$. Define

$$g(x | \mu, \sigma^2) = \left(\begin{array}{c} x - \mu \\ (x - \mu)^2 - \sigma^2 \end{array} \right).$$

Show that $\mathbb{E}g(x | m, s) = 0$ if and only if $m = \mu$ and $s = \sigma^2$.

Exercise 2.18 Suppose that

$$\mathbf{x} = \begin{pmatrix} 1 \\ x_2 \\ x_3 \end{pmatrix}$$

and $x_3 = \alpha_1 + \alpha_2 x_2$ is a linear function of x_2 .

- Show that $\mathbf{Q}_{\mathbf{x}\mathbf{x}} = \mathbb{E}(\mathbf{x}\mathbf{x}')$ is not invertible.
- Use a linear transformation of \mathbf{x} to find an expression for the best linear predictor of y given \mathbf{x} . (Be explicit, do not just use the generalized inverse formula.)

Exercise 2.19 Show (2.46)-(2.47), namely that for

$$d(\boldsymbol{\beta}) = \mathbb{E}(m(\mathbf{x}) - \mathbf{x}'\boldsymbol{\beta})^2$$

then

$$\begin{aligned} \boldsymbol{\beta} &= \underset{\boldsymbol{\beta} \in \mathbb{R}^k}{\text{argmin}} d(\boldsymbol{\beta}) \\ &= (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1} \mathbb{E}(\mathbf{x}m(\mathbf{x})) \\ &= (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1} \mathbb{E}(\mathbf{x}y). \end{aligned}$$

Hint: To show $\mathbb{E}(\mathbf{x}m(\mathbf{x})) = \mathbb{E}(\mathbf{x}y)$ use the law of iterated expectations.

Exercise 2.20 Verify that (2.60) holds with $m(\mathbf{x})$ defined in (2.7) when (y, \mathbf{x}) have a joint density $f(y, \mathbf{x})$.

Chapter 3

The Algebra of Least Squares

3.1 Introduction

In this chapter we introduce the popular least-squares estimator. Most of the discussion will be algebraic, with questions of distribution and inference deferred to later chapters.

3.2 Random Samples

In Section 2.18 we derived and discussed the best linear predictor of y given \mathbf{x} for a pair of random variables $(y, \mathbf{x}) \in \mathbb{R} \times \mathbb{R}^k$, and called this the linear projection model. We are now interested in **estimating** the parameters of this model, in particular the projection coefficient

$$\boldsymbol{\beta} = (\mathbb{E}(\mathbf{x}\mathbf{x}'))^{-1} \mathbb{E}(\mathbf{x}y).$$

We can estimate $\boldsymbol{\beta}$ from observational data which includes joint measurements on the variables (y, \mathbf{x}) . For example, supposing we are interested in estimating a wage equation, we would use a dataset with observations on wages (or weekly earnings), education, experience (or age), and demographic characteristics (gender, race, location). One possible dataset is the Current Population Survey (CPS), a survey of U.S. households which includes questions on employment, income, education, and demographic characteristics.

Notationally we wish to emphasize when we are discussing observations. Typically in econometrics we denote observations by appending a subscript i which runs from 1 to n , thus the i^{th} observation is (y_i, \mathbf{x}_i) , and n denotes the sample size. The dataset is then $\{(y_i, \mathbf{x}_i); i = 1, \dots, n\}$.

From the viewpoint of empirical analysis, a dataset is a array of numbers often organized as a table, where the columns of the table correspond to distinct variables and the rows correspond to distinct observations. For empirical analysis, the dataset and observations are fixed in the sense that they are numbers presented to the researcher. For statistical analysis we need to view the dataset as random, or more precisely as a realization of a random process. For cross-sectional studies, the most common approach is to treat the individual observations as independent draws from an underlying population F . When the observations are realizations of independent and identically distributed random variables, we say that the data is a random sample.

Assumption 3.2.1 *The observations $\{(y_1, \mathbf{x}_1), \dots, (y_i, \mathbf{x}_i), \dots, (y_n, \mathbf{x}_n)\}$ are a random sample.*

With a random sample, the ordering of the data is irrelevant. There is nothing special about any specific observation or ordering. You can permute the order of the observations and no information is gained or lost.

As most economic data sets are not literally the result of a random experiment, the random sampling framework is best viewed as an approximation rather than being literally true.

The linear projection model applies to the random observations (y_i, \mathbf{x}_i) . This means that the probability model for the observations is the same as that described in Section 2.18. We can write the model as

$$y_i = \mathbf{x}'_i \boldsymbol{\beta} + e_i \quad (3.1)$$

where the linear projection coefficient $\boldsymbol{\beta}$ is defined as

$$\boldsymbol{\beta} = \underset{\boldsymbol{\beta} \in \mathbb{R}^k}{\operatorname{argmin}} S(\boldsymbol{\beta}), \quad (3.2)$$

the minimizer of the expected squared error

$$S(\boldsymbol{\beta}) = \mathbb{E} (y_i - \mathbf{x}'_i \boldsymbol{\beta})^2, \quad (3.3)$$

and has the explicit solution

$$\boldsymbol{\beta} = (\mathbb{E} (\mathbf{x}_i \mathbf{x}'_i))^{-1} \mathbb{E} (\mathbf{x}_i y_i). \quad (3.4)$$

3.3 Sample Means

Consider the intercept-only model

$$\begin{aligned} y_i &= \mu + e_i \\ \mathbb{E}(e_i) &= 0. \end{aligned}$$

In this case the regression parameter is the unconditional mean $\mu = \mathbb{E}(y_i)$.

The standard estimator of a population mean is the sample mean, namely

$$\hat{\mu} = \frac{1}{n} \sum_{i=1}^n y_i.$$

The sample mean is the empirical analog of the population mean, and is the conventional estimator in the lack of other information about μ or the distribution of y . We call $\hat{\mu}$ the **moment estimator** for μ .

Indeed, whenever we have a parameter which can be written as the expectation of a function of random variables, a natural estimator of the parameter is the moment estimator, which is the sample mean of the corresponding function of the observations. For example, for $\mu_2 = \mathbb{E}(y_i^2)$ the moment estimator is $\hat{\mu}_2 = \frac{1}{n} \sum_{i=1}^n y_i^2$, and for $\theta = \mathbb{E}(y_{1i} y_{2i})$ the moment estimator is $\hat{\theta} = \frac{1}{n} \sum_{i=1}^n y_{1i} y_{2i}$.

3.4 Least Squares Estimator

The linear projection coefficient $\boldsymbol{\beta}$ is defined in (3.2) as the minimizer of the expected squared error $S(\boldsymbol{\beta})$ defined in (3.3). For given $\boldsymbol{\beta}$, the expected squared error is the expectation of the squared error $(y_i - \mathbf{x}'_i \boldsymbol{\beta})^2$. The moment estimator of $S(\boldsymbol{\beta})$ is the sample average:

$$\begin{aligned} S_n(\boldsymbol{\beta}) &= \frac{1}{n} \sum_{i=1}^n (y_i - \mathbf{x}'_i \boldsymbol{\beta})^2 \\ &= \frac{1}{n} SSE_n(\boldsymbol{\beta}) \end{aligned} \quad (3.5)$$

where

$$SSE_n(\boldsymbol{\beta}) = \sum_{i=1}^n (y_i - \mathbf{x}'_i \boldsymbol{\beta})^2$$

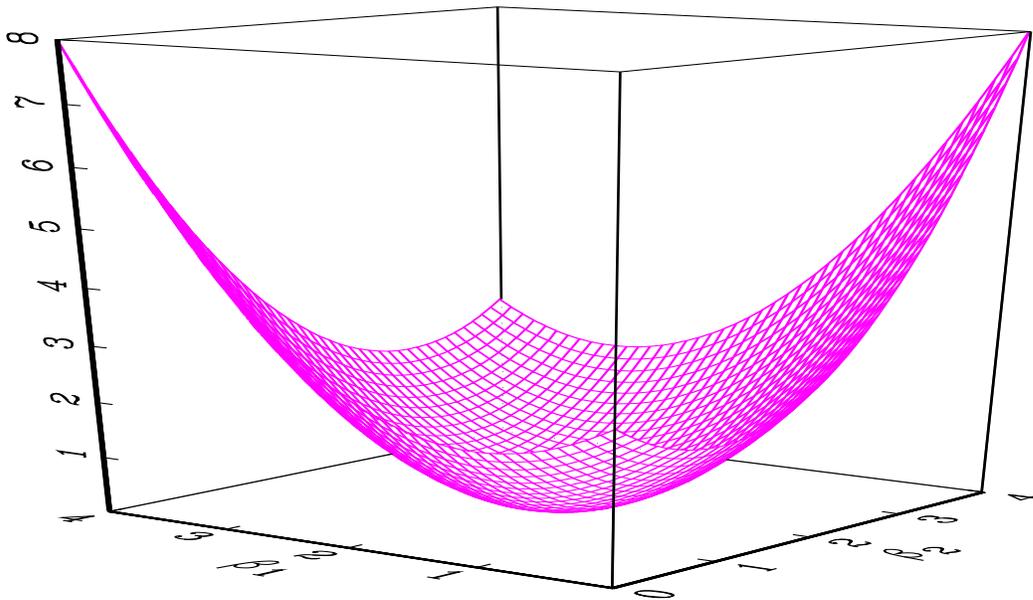


Figure 3.1: Sum-of-Squared Errors Function

is called the **sum-of-squared-errors** function.

Since $S_n(\boldsymbol{\beta})$ is a sample average, we can interpret it as an estimator of the expected squared error $S(\boldsymbol{\beta})$. Examining $S_n(\boldsymbol{\beta})$ as a function of $\boldsymbol{\beta}$ therefore is informative about how $S(\boldsymbol{\beta})$ varies with $\boldsymbol{\beta}$. The projection coefficient coefficient minimizes $S(\boldsymbol{\beta})$, an analog estimator minimizes (3.5):

$$\hat{\boldsymbol{\beta}} = \underset{\boldsymbol{\beta} \in \mathbb{R}^k}{\operatorname{argmin}} S_n(\boldsymbol{\beta}).$$

Alternatively, as $S_n(\boldsymbol{\beta})$ is a scale multiple of $SSE_n(\boldsymbol{\beta})$, we may equivalently define $\hat{\boldsymbol{\beta}}$ as the minimizer of $SSE_n(\boldsymbol{\beta})$. Hence $\hat{\boldsymbol{\beta}}$ is commonly called the **least-squares (LS) (or ordinary least squares (OLS)) estimator** of $\boldsymbol{\beta}$. Here, as is common in econometrics, we put a hat “ $\hat{}$ ” over the parameter $\boldsymbol{\beta}$ to indicate that $\hat{\boldsymbol{\beta}}$ is a sample estimate of $\boldsymbol{\beta}$. This is a helpful convention, as just by seeing the symbol $\hat{\boldsymbol{\beta}}$ we can immediately interpret it as an estimator (because of the hat), and as an estimator of a parameter labelled $\boldsymbol{\beta}$. Sometimes when we want to be explicit about the estimation method, we will write $\hat{\boldsymbol{\beta}}_{\text{ols}}$ to signify that it is the OLS estimator. It is also common to see the notation $\hat{\boldsymbol{\beta}}_n$, where the subscript “ n ” indicates that the estimator depends on the sample size n .

It is important to understand the distinction between population parameters such as $\boldsymbol{\beta}$ and sample estimates such as $\hat{\boldsymbol{\beta}}$. The population parameter $\boldsymbol{\beta}$ is a non-random feature of the population while the sample estimate $\hat{\boldsymbol{\beta}}$ is a random feature of a random sample. $\boldsymbol{\beta}$ is fixed, while $\hat{\boldsymbol{\beta}}$ varies across samples.

To visualize the quadratic function $S_n(\boldsymbol{\beta})$, Figure 3.1 displays an example sum-of-squared errors function $SSE_n(\boldsymbol{\beta})$ for the case $k = 2$. The least-squares estimator $\hat{\boldsymbol{\beta}}$ is the the pair $(\hat{\beta}_1, \hat{\beta}_2)$ minimizing this function.

3.5 Solving for Least Squares with One Regressor

For simplicity, we start by considering the case $k = 1$ so that the coefficient β is a scalar. Then the sum of squared errors is a simple quadratic

$$\begin{aligned} SSE_n(\beta) &= \sum_{i=1}^n (y_i - x_i\beta)^2 \\ &= \left(\sum_{i=1}^n y_i^2 \right) - 2\beta \left(\sum_{i=1}^n x_i y_i \right) + \beta^2 \left(\sum_{i=1}^n x_i^2 \right). \end{aligned}$$

The OLS estimator $\hat{\beta}$ minimizes this function. From elementary algebra we know that the minimizer of the quadratic function $a - 2bx + cx^2$ is $x = b/c$. Thus the minimizer of $SSE_n(\beta)$ is

$$\hat{\beta} = \frac{\sum_{i=1}^n x_i y_i}{\sum_{i=1}^n x_i^2}. \quad (3.6)$$

The intercept-only model is the special case $x_i = 1$. In this case we find

$$\hat{\beta} = \frac{\sum_{i=1}^n 1 y_i}{\sum_{i=1}^n 1^2} = \frac{1}{n} \sum_{i=1}^n y_i = \bar{y}, \quad (3.7)$$

the sample mean of y_i . Here, as is common, we put a bar “ $\bar{}$ ” over y to indicate that the quantity is a sample mean. This calculation shows that the OLS estimator in the intercept-only model is the sample mean.

3.6 Solving for Least Squares with Multiple Regressors

We now consider the case with $k \geq 1$ so that the coefficient β is a vector.

To solve for $\hat{\beta}$, expand the SSE function to find

$$SSE_n(\beta) = \sum_{i=1}^n y_i^2 - 2\beta' \sum_{i=1}^n \mathbf{x}_i y_i + \beta' \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \beta.$$

This is a quadratic expression in the vector argument β . The first-order-condition for minimization of $SSE_n(\beta)$ is

$$0 = \frac{\partial}{\partial \beta} SSE_n(\hat{\beta}) = -2 \sum_{i=1}^n \mathbf{x}_i y_i + 2 \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \hat{\beta}. \quad (3.8)$$

We have written this using a single expression, but it is actually a system of k equations with k unknowns (the elements of $\hat{\beta}$).

The solution for $\hat{\beta}$ may be found by solving the system of k equations in (3.8). We can write this solution compactly using matrix algebra. Inverting the $k \times k$ matrix $\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i'$ we find an explicit formula for the least-squares estimator

$$\hat{\beta} = \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \right)^{-1} \left(\sum_{i=1}^n \mathbf{x}_i y_i \right). \quad (3.9)$$

This is the natural estimator of the best linear projection coefficient β defined in (3.2), and can also be called the linear projection estimator.

We see that (3.9) simplifies to the expression (3.6) when $k = 1$. The expression (3.9) is a notationally simple generalization but requires a careful attention to vector and matrix manipulations.

Alternatively, equation (3.4) writes the projection coefficient β as an explicit function of the population moments \mathbf{Q}_{xy} and \mathbf{Q}_{xx} . Their moment estimators are the sample moments

$$\begin{aligned} \hat{\mathbf{Q}}_{xy} &= \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i y_i \\ \hat{\mathbf{Q}}_{xx} &= \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i'. \end{aligned}$$

The moment estimator of β replaces the population moments in (3.4) with the sample moments:

$$\begin{aligned}\widehat{\beta} &= \widehat{\mathbf{Q}}_{xx}^{-1} \widehat{\mathbf{Q}}_{xy} \\ &= \left(\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}'_i \right)^{-1} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i y_i \right) \\ &= \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}'_i \right)^{-1} \left(\sum_{i=1}^n \mathbf{x}_i y_i \right)\end{aligned}$$

which is identical with (3.9).

Least Squares Estimation

Definition 3.6.1 *The least-squares estimator $\widehat{\beta}$ is*

$$\widehat{\beta} = \underset{\beta \in \mathbb{R}^k}{\operatorname{argmin}} S_n(\beta)$$

where

$$S_n(\beta) = \frac{1}{n} \sum_{i=1}^n (y_i - \mathbf{x}'_i \beta)^2$$

and has the solution

$$\widehat{\beta} = \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}'_i \right)^{-1} \left(\sum_{i=1}^n \mathbf{x}_i y_i \right).$$

Adrien-Marie Legendre

The method of least-squares was first published in 1805 by the French mathematician Adrien-Marie Legendre (1752-1833). Legendre proposed least-squares as a solution to the algebraic problem of solving a system of equations when the number of equations exceeded the number of unknowns. This was a vexing and common problem in astronomical measurement. As viewed by Legendre, (3.1) is a set of n equations with k unknowns. As the equations cannot be solved exactly, Legendre's goal was to select β to make the set of errors as small as possible. He proposed the sum of squared error criterion, and derived the algebraic solution presented above. As he noted, the first-order conditions (3.8) is a system of k equations with k unknowns, which can be solved by "ordinary" methods. Hence the method became known as **Ordinary Least Squares** and to this day we still use the abbreviation OLS to refer to Legendre's estimation method.

3.7 Illustration

We illustrate the least-squares estimator in practice with the data set used to generate the estimates from Chapter 2. This is the March 2009 Current Population Survey, which has extensive information on the U.S. population. This data set is described in more detail in Section 3.19. For this illustration, we use the sub-sample of married (spouse present) black female wages earners with 12 years potential work experience. This sub-sample has 20 observations. Let y_i be log wages and \mathbf{x}_i be years of education and an intercept. Then

$$\sum_{i=1}^n \mathbf{x}_i y_i = \begin{pmatrix} 995.86 \\ 62.64 \end{pmatrix},$$

$$\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' = \begin{pmatrix} 5010 & 314 \\ 314 & 20 \end{pmatrix},$$

and

$$\left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \right)^{-1} = \begin{pmatrix} 0.0125 & -0.196 \\ -0.196 & 3.124 \end{pmatrix}$$

Thus

$$\begin{aligned} \hat{\boldsymbol{\beta}} &= \begin{pmatrix} 0.0125 & -0.196 \\ -0.196 & 3.124 \end{pmatrix} \begin{pmatrix} 995.86 \\ 62.64 \end{pmatrix} \\ &= \begin{pmatrix} 0.155 \\ 0.698 \end{pmatrix}. \end{aligned} \tag{3.10}$$

We often write the estimated equation using the format

$$\widehat{\log(Wage)} = 0.155 \textit{ education} + 0.698. \tag{3.11}$$

An interpretation of the estimated equation is that each year of education is associated with an 16% increase in mean wages.

Equation (3.11) is called a **bivariate regression** as there are only two variables. A **multivariate regression** has two or more regressors, and allows a more detailed investigation. Let's take an example similar to (3.11) but include all levels of experience. This time, we use the sub-sample of single (never married) asian men, which has 268 observations. Including as regressors years of potential work experience (*experience*) and its square ($experience^2/100$) (we divide by 100 to simplify reporting), we obtain the estimates

$$\widehat{\log(Wage)} = 0.143 \textit{ education} + 0.036 \textit{ experience} - 0.071 \textit{ experience}^2/100 + 0.575. \tag{3.12}$$

These estimates suggest a 14% increase in mean wages per year of education, holding experience constant.

3.8 Least Squares Residuals

As a by-product of estimation, we define the **fitted value**

$$\hat{y}_i = \mathbf{x}_i' \hat{\boldsymbol{\beta}}$$

and the **residual**

$$\hat{e}_i = y_i - \hat{y}_i = y_i - \mathbf{x}_i' \hat{\boldsymbol{\beta}}. \tag{3.13}$$

Sometimes \hat{y}_i is called the predicted value, but this is a misleading label. The fitted value \hat{y}_i is a function of the entire sample, including y_i , and thus cannot be interpreted as a valid prediction of y_i . It is thus more accurate to describe \hat{y}_i as a *fitted* rather than a *predicted* value.

Note that $y_i = \hat{y}_i + \hat{e}_i$ and

$$y_i = \mathbf{x}'_i \hat{\boldsymbol{\beta}} + \hat{e}_i. \quad (3.14)$$

We make a distinction between the **error** e_i and the **residual** \hat{e}_i . The error e_i is unobservable while the residual \hat{e}_i is a by-product of estimation. These two variables are frequently mislabeled, which can cause confusion.

Equation (3.8) implies that

$$\sum_{i=1}^n \mathbf{x}_i \hat{e}_i = \mathbf{0}. \quad (3.15)$$

To see this by a direct calculation, using (3.13) and (3.9),

$$\begin{aligned} \sum_{i=1}^n \mathbf{x}_i \hat{e}_i &= \sum_{i=1}^n \mathbf{x}_i (y_i - \mathbf{x}'_i \hat{\boldsymbol{\beta}}) \\ &= \sum_{i=1}^n \mathbf{x}_i y_i - \sum_{i=1}^n \mathbf{x}_i \mathbf{x}'_i \hat{\boldsymbol{\beta}} \\ &= \sum_{i=1}^n \mathbf{x}_i y_i - \sum_{i=1}^n \mathbf{x}_i \mathbf{x}'_i \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}'_i \right)^{-1} \left(\sum_{i=1}^n \mathbf{x}_i y_i \right) \\ &= \sum_{i=1}^n \mathbf{x}_i y_i - \sum_{i=1}^n \mathbf{x}_i y_i \\ &= \mathbf{0}. \end{aligned}$$

When \mathbf{x}_i contains a constant, an implication of (3.15) is

$$\frac{1}{n} \sum_{i=1}^n \hat{e}_i = 0. \quad (3.16)$$

Thus the residuals have a sample mean of zero and the sample correlation between the regressors and the residual is zero. These are algebraic results, and hold true for all linear regression estimates.

3.9 Model in Matrix Notation

For many purposes, including computation, it is convenient to write the model and statistics in matrix notation. The linear equation (2.26) is a system of n equations, one for each observation. We can stack these n equations together as

$$\begin{aligned} y_1 &= \mathbf{x}'_1 \boldsymbol{\beta} + e_1 \\ y_2 &= \mathbf{x}'_2 \boldsymbol{\beta} + e_2 \\ &\vdots \\ y_n &= \mathbf{x}'_n \boldsymbol{\beta} + e_n. \end{aligned}$$

Now define

$$\mathbf{y} = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{pmatrix}, \quad \mathbf{X} = \begin{pmatrix} \mathbf{x}'_1 \\ \mathbf{x}'_2 \\ \vdots \\ \mathbf{x}'_n \end{pmatrix}, \quad \mathbf{e} = \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_n \end{pmatrix}.$$

Observe that \mathbf{y} and \mathbf{e} are $n \times 1$ vectors, and \mathbf{X} is an $n \times k$ matrix. Then the system of n equations can be compactly written in the single equation

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}. \quad (3.17)$$

Sample sums can be written in matrix notation. For example

$$\begin{aligned} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' &= \mathbf{X}'\mathbf{X} \\ \sum_{i=1}^n \mathbf{x}_i y_i &= \mathbf{X}'\mathbf{y}. \end{aligned}$$

Therefore the least-squares estimator can be written as

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1} (\mathbf{X}'\mathbf{y}). \quad (3.18)$$

The matrix version of (3.14) and estimated version of (3.17) is

$$\mathbf{y} = \mathbf{X}\hat{\boldsymbol{\beta}} + \hat{\mathbf{e}},$$

or equivalently the residual vector is

$$\hat{\mathbf{e}} = \mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}.$$

Using the residual vector, we can write (3.15) as

$$\mathbf{X}'\hat{\mathbf{e}} = \mathbf{0}. \quad (3.20)$$

Using matrix notation we have simple expressions for most estimators. This is particularly convenient for computer programming, as most languages allow matrix notation and manipulation.

Important Matrix Expressions

$$\begin{aligned} \mathbf{y} &= \mathbf{X}\boldsymbol{\beta} + \mathbf{e} \\ \hat{\boldsymbol{\beta}} &= (\mathbf{X}'\mathbf{X})^{-1} (\mathbf{X}'\mathbf{y}) \\ \hat{\mathbf{e}} &= \mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}} \\ \mathbf{X}'\hat{\mathbf{e}} &= \mathbf{0}. \end{aligned}$$

Early Use of Matrices

The earliest known treatment of the use of matrix methods to solve simultaneous systems is found in Chapter 8 of the Chinese text *The Nine Chapters on the Mathematical Art*, written by several generations of scholars from the 10th to 2nd century BCE.

3.10 Projection Matrix

Define the matrix

$$\mathbf{P} = \mathbf{X} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'.$$

Observe that

$$\mathbf{P}\mathbf{X} = \mathbf{X} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{X} = \mathbf{X}.$$

This is a property of a **projection matrix**. More generally, for any matrix \mathbf{Z} which can be written as $\mathbf{Z} = \mathbf{X}\mathbf{\Gamma}$ for some matrix $\mathbf{\Gamma}$ (we say that \mathbf{Z} lies in the **range space** of \mathbf{X}), then

$$\mathbf{P}\mathbf{Z} = \mathbf{P}\mathbf{X}\mathbf{\Gamma} = \mathbf{X} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{X}\mathbf{\Gamma} = \mathbf{X}\mathbf{\Gamma} = \mathbf{Z}.$$

As an important example, if we partition the matrix \mathbf{X} into two matrices \mathbf{X}_1 and \mathbf{X}_2 so that

$$\mathbf{X} = [\mathbf{X}_1 \quad \mathbf{X}_2],$$

then $\mathbf{P}\mathbf{X}_1 = \mathbf{X}_1$. (See Exercise 3.7.)

The matrix \mathbf{P} is **symmetric** and **idempotent**¹. To see that it is symmetric,

$$\begin{aligned} \mathbf{P}' &= (\mathbf{X} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}')' \\ &= (\mathbf{X}')' \left((\mathbf{X}'\mathbf{X})^{-1} \right)' (\mathbf{X})' \\ &= \mathbf{X} \left((\mathbf{X}'\mathbf{X})' \right)^{-1} \mathbf{X}' \\ &= \mathbf{X} \left((\mathbf{X})' (\mathbf{X}')' \right)^{-1} \mathbf{X}' \\ &= \mathbf{P}. \end{aligned}$$

To establish that it is idempotent, the fact that $\mathbf{P}\mathbf{X} = \mathbf{X}$ implies that

$$\begin{aligned} \mathbf{P}\mathbf{P} &= \mathbf{P}\mathbf{X} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}' \\ &= \mathbf{X} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}' \\ &= \mathbf{P}. \end{aligned}$$

The matrix \mathbf{P} has the property that it creates the fitted values in a least-squares regression:

$$\mathbf{P}\mathbf{y} = \mathbf{X} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{y} = \mathbf{X}\hat{\boldsymbol{\beta}} = \hat{\mathbf{y}}.$$

Because of this property, \mathbf{P} is also known as the “hat matrix”.

A special example of a projection matrix occurs when $\mathbf{X} = \mathbf{1}$ is an n -vector of ones. Then

$$\begin{aligned} \mathbf{P}_1 &= \mathbf{1} (\mathbf{1}'\mathbf{1})^{-1} \mathbf{1}' \\ &= \frac{1}{n} \mathbf{1}\mathbf{1}'. \end{aligned}$$

Note that

$$\begin{aligned} \mathbf{P}_1\mathbf{y} &= \mathbf{1} (\mathbf{1}'\mathbf{1})^{-1} \mathbf{1}'\mathbf{y} \\ &= \mathbf{1}\bar{y} \end{aligned}$$

creates an n -vector whose elements are the sample mean \bar{y} of y_i .

¹A matrix \mathbf{P} is symmetric if $\mathbf{P}' = \mathbf{P}$. A matrix \mathbf{P} is idempotent if $\mathbf{P}\mathbf{P} = \mathbf{P}$. See Appendix A.8.

The i 'th diagonal element of $\mathbf{P} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$ is

$$h_{ii} = \mathbf{x}'_i (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i \quad (3.21)$$

which is called the **leverage** of the i 'th observation.

Some useful properties of the the matrix \mathbf{P} and the leverage values h_{ii} are now summarized.

Theorem 3.10.1

$$\sum_{i=1}^n h_{ii} = \text{tr } \mathbf{P} = k \quad (3.22)$$

and

$$0 \leq h_{ii} \leq 1. \quad (3.23)$$

To show (3.22),

$$\begin{aligned} \text{tr } \mathbf{P} &= \text{tr} \left(\mathbf{X} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}' \right) \\ &= \text{tr} \left((\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{X} \right) \\ &= \text{tr} (\mathbf{I}_k) \\ &= k. \end{aligned}$$

See Appendix A.4 for definition and properties of the trace operator. The proof of (3.23) is deferred to Section 3.21.

3.11 Orthogonal Projection

Define

$$\begin{aligned} \mathbf{M} &= \mathbf{I}_n - \mathbf{P} \\ &= \mathbf{I}_n - \mathbf{X} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}' \end{aligned}$$

where \mathbf{I}_n is the $n \times n$ identity matrix. Note that

$$\mathbf{M}\mathbf{X} = (\mathbf{I}_n - \mathbf{P})\mathbf{X} = \mathbf{X} - \mathbf{P}\mathbf{X} = \mathbf{X} - \mathbf{X} = \mathbf{0}.$$

Thus \mathbf{M} and \mathbf{X} are orthogonal. We call \mathbf{M} an **orthogonal projection matrix** or an **annihilator matrix** due to the property that for any matrix \mathbf{Z} in the range space of \mathbf{X} then

$$\mathbf{M}\mathbf{Z} = \mathbf{Z} - \mathbf{P}\mathbf{Z} = \mathbf{0}.$$

For example, $\mathbf{M}\mathbf{X}_1 = \mathbf{0}$ for any subcomponent \mathbf{X}_1 of \mathbf{X} , and $\mathbf{M}\mathbf{P} = \mathbf{0}$ (see Exercise 3.7).

The orthogonal projection matrix \mathbf{M} has many similar properties with \mathbf{P} , including that \mathbf{M} is symmetric ($\mathbf{M}' = \mathbf{M}$) and idempotent ($\mathbf{M}\mathbf{M} = \mathbf{M}$). Similarly to (3.22) we can calculate

$$\text{tr } \mathbf{M} = n - k. \quad (3.24)$$

(See Exercise 3.9.) While \mathbf{P} creates fitted values, \mathbf{M} creates least-squares residuals:

$$\mathbf{M}\mathbf{y} = \mathbf{y} - \mathbf{P}\mathbf{y} = \mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}} = \hat{\mathbf{e}}. \quad (3.25)$$

As discussed in the previous section, a special example of a projection matrix occurs when $\mathbf{X} = \mathbf{1}$ is an n -vector of ones, so that $\mathbf{P}_1 = \mathbf{1}(\mathbf{1}'\mathbf{1})^{-1}\mathbf{1}'$. Similarly, set

$$\begin{aligned}\mathbf{M}_1 &= \mathbf{I}_n - \mathbf{P}_1 \\ &= \mathbf{I}_n - \mathbf{1}(\mathbf{1}'\mathbf{1})^{-1}\mathbf{1}'.\end{aligned}$$

While \mathbf{P}_1 creates a vector of sample means, \mathbf{M}_1 creates demeaned values:

$$\mathbf{M}_1\mathbf{y} = \mathbf{y} - \mathbf{1}\bar{y}.$$

For simplicity we will often write the right-hand-side as $\mathbf{y} - \bar{y}$. The i 'th element is $y_i - \bar{y}$, the **demeaned** value of y_i .

We can also use (3.25) to write an alternative expression for the residual vector. Substituting $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}$ into $\hat{\mathbf{e}} = \mathbf{M}\mathbf{y}$ and using $\mathbf{M}\mathbf{X} = \mathbf{0}$ we find

$$\hat{\mathbf{e}} = \mathbf{M}\mathbf{y} = \mathbf{M}(\mathbf{X}\boldsymbol{\beta} + \mathbf{e}) = \mathbf{M}\mathbf{e} \quad (3.26)$$

which is free of dependence on the regression coefficient $\boldsymbol{\beta}$.

3.12 Estimation of Error Variance

The error variance $\sigma^2 = \mathbb{E}e_i^2$ is a moment, so a natural estimator is a moment estimator. If e_i were observed we would estimate σ^2 by

$$\tilde{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n e_i^2. \quad (3.27)$$

However, this is infeasible as e_i is not observed. In this case it is common to take a two-step approach to estimation. The residuals \hat{e}_i are calculated in the first step, and then we substitute \hat{e}_i for e_i in expression (3.27) to obtain the feasible estimator

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n \hat{e}_i^2. \quad (3.28)$$

In matrix notation, we can write (3.27) and (3.28) as

$$\tilde{\sigma}^2 = n^{-1}\mathbf{e}'\mathbf{e}$$

and

$$\hat{\sigma}^2 = n^{-1}\hat{\mathbf{e}}'\hat{\mathbf{e}}. \quad (3.29)$$

Recall the expressions $\hat{\mathbf{e}} = \mathbf{M}\mathbf{y} = \mathbf{M}\mathbf{e}$ from (3.25) and (3.26). Applied to (3.29) we find

$$\begin{aligned}\hat{\sigma}^2 &= n^{-1}\hat{\mathbf{e}}'\hat{\mathbf{e}} \\ &= n^{-1}\mathbf{y}'\mathbf{M}\mathbf{M}\mathbf{y} \\ &= n^{-1}\mathbf{y}'\mathbf{M}\mathbf{y} \\ &= n^{-1}\mathbf{e}'\mathbf{M}\mathbf{e}\end{aligned}$$

the third equality since $\mathbf{M}\mathbf{M} = \mathbf{M}$.

An interesting implication is that

$$\begin{aligned}\tilde{\sigma}^2 - \hat{\sigma}^2 &= n^{-1}\mathbf{e}'\mathbf{e} - n^{-1}\mathbf{e}'\mathbf{M}\mathbf{e} \\ &= n^{-1}\mathbf{e}'\mathbf{P}\mathbf{e} \\ &\geq 0.\end{aligned}$$

The final inequality holds because \mathbf{P} is positive semi-definite and $\mathbf{e}'\mathbf{P}\mathbf{e}$ is a quadratic form. This shows that the feasible estimator $\hat{\sigma}^2$ is numerically smaller than the idealized estimator (3.27).

3.13 Analysis of Variance

Another way of writing (3.25) is

$$\mathbf{y} = \mathbf{P}\mathbf{y} + \mathbf{M}\mathbf{y} = \hat{\mathbf{y}} + \hat{\mathbf{e}}. \quad (3.30)$$

This decomposition is **orthogonal**, that is

$$\hat{\mathbf{y}}' \hat{\mathbf{e}} = (\mathbf{P}\mathbf{y})' (\mathbf{M}\mathbf{y}) = \mathbf{y}' \mathbf{P} \mathbf{M} \mathbf{y} = 0.$$

It follows that

$$\mathbf{y}' \mathbf{y} = \hat{\mathbf{y}}' \hat{\mathbf{y}} + 2\hat{\mathbf{y}}' \hat{\mathbf{e}} + \hat{\mathbf{e}}' \hat{\mathbf{e}} = \hat{\mathbf{y}}' \hat{\mathbf{y}} + \hat{\mathbf{e}}' \hat{\mathbf{e}}$$

or

$$\sum_{i=1}^n y_i^2 = \sum_{i=1}^n \hat{y}_i^2 + \sum_{i=1}^n \hat{e}_i^2.$$

Subtracting \bar{y} from both sides of (3.30) we obtain

$$\mathbf{y} - \mathbf{1}\bar{y} = \hat{\mathbf{y}} - \mathbf{1}\bar{y} + \hat{\mathbf{e}}$$

This decomposition is also orthogonal when \mathbf{X} contains a constant, as

$$(\hat{\mathbf{y}} - \mathbf{1}\bar{y})' \hat{\mathbf{e}} = \hat{\mathbf{y}}' \hat{\mathbf{e}} - \bar{y} \mathbf{1}' \hat{\mathbf{e}} = 0$$

under (3.16). It follows that

$$(\mathbf{y} - \mathbf{1}\bar{y})' (\mathbf{y} - \mathbf{1}\bar{y}) = (\hat{\mathbf{y}} - \mathbf{1}\bar{y})' (\hat{\mathbf{y}} - \mathbf{1}\bar{y}) + \hat{\mathbf{e}}' \hat{\mathbf{e}}$$

or

$$\sum_{i=1}^n (y_i - \bar{y})^2 = \sum_{i=1}^n (\hat{y}_i - \bar{y})^2 + \sum_{i=1}^n \hat{e}_i^2.$$

This is commonly called the **analysis-of-variance** formula for least squares regression.

A commonly reported statistic is the **coefficient of determination** or **R-squared**:

$$R^2 = \frac{\sum_{i=1}^n (\hat{y}_i - \bar{y})^2}{\sum_{i=1}^n (y_i - \bar{y})^2} = 1 - \frac{\sum_{i=1}^n \hat{e}_i^2}{\sum_{i=1}^n (y_i - \bar{y})^2}.$$

It is often described as the fraction of the sample variance of y_i which is explained by the least-squares fit. R^2 is a crude measure of regression fit. We have better measures of fit, but these require a statistical (not just algebraic) analysis and we will return to these issues later. One deficiency with R^2 is that it increases when regressors are added to a regression (see Exercise 3.16) so the “fit” can be always increased by increasing the number of regressors.

3.14 Regression Components

Partition

$$\mathbf{X} = [\mathbf{X}_1 \quad \mathbf{X}_2]$$

and

$$\boldsymbol{\beta} = \begin{pmatrix} \boldsymbol{\beta}_1 \\ \boldsymbol{\beta}_2 \end{pmatrix}.$$

Then the regression model can be rewritten as

$$\mathbf{y} = \mathbf{X}_1 \boldsymbol{\beta}_1 + \mathbf{X}_2 \boldsymbol{\beta}_2 + \mathbf{e}. \quad (3.31)$$

The OLS estimator of $\boldsymbol{\beta} = (\boldsymbol{\beta}'_1, \boldsymbol{\beta}'_2)'$ is obtained by regression of \mathbf{y} on $\mathbf{X} = [\mathbf{X}_1 \ \mathbf{X}_2]$ and can be written as

$$\mathbf{y} = \mathbf{X}\hat{\boldsymbol{\beta}} + \hat{\mathbf{e}} = \mathbf{X}_1\hat{\boldsymbol{\beta}}_1 + \mathbf{X}_2\hat{\boldsymbol{\beta}}_2 + \hat{\mathbf{e}}. \quad (3.32)$$

We are interested in algebraic expressions for $\hat{\boldsymbol{\beta}}_1$ and $\hat{\boldsymbol{\beta}}_2$.

The algebra for the estimator is identical as that for the population coefficients as presented in Section 2.21.

Partition $\hat{\mathbf{Q}}_{xx}$ and $\hat{\mathbf{Q}}_{xy}$ as

$$\hat{\mathbf{Q}}_{xx} = \begin{bmatrix} \hat{\mathbf{Q}}_{11} & \hat{\mathbf{Q}}_{12} \\ \hat{\mathbf{Q}}_{21} & \hat{\mathbf{Q}}_{22} \end{bmatrix} = \begin{bmatrix} \frac{1}{n}\mathbf{X}'_1\mathbf{X}_1 & \frac{1}{n}\mathbf{X}'_1\mathbf{X}_2 \\ \frac{1}{n}\mathbf{X}'_2\mathbf{X}_1 & \frac{1}{n}\mathbf{X}'_2\mathbf{X}_2 \end{bmatrix}$$

and similarly \mathbf{Q}_{xy}

$$\hat{\mathbf{Q}}_{xy} = \begin{bmatrix} \hat{\mathbf{Q}}_{1y} \\ \hat{\mathbf{Q}}_{2y} \end{bmatrix} = \begin{bmatrix} \frac{1}{n}\mathbf{X}'_1\mathbf{y} \\ \frac{1}{n}\mathbf{X}'_2\mathbf{y} \end{bmatrix}.$$

By the partitioned matrix inversion formula (A.4)

$$\hat{\mathbf{Q}}_{xx}^{-1} = \begin{bmatrix} \hat{\mathbf{Q}}_{11} & \hat{\mathbf{Q}}_{12} \\ \hat{\mathbf{Q}}_{21} & \hat{\mathbf{Q}}_{22} \end{bmatrix}^{-1} \stackrel{def}{=} \begin{bmatrix} \hat{\mathbf{Q}}^{11} & \hat{\mathbf{Q}}^{12} \\ \hat{\mathbf{Q}}^{21} & \hat{\mathbf{Q}}^{22} \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{Q}}_{11.2}^{-1} & -\hat{\mathbf{Q}}_{11.2}^{-1}\hat{\mathbf{Q}}_{12}\hat{\mathbf{Q}}_{22}^{-1} \\ -\hat{\mathbf{Q}}_{22.1}^{-1}\hat{\mathbf{Q}}_{21}\hat{\mathbf{Q}}_{11}^{-1} & \hat{\mathbf{Q}}_{22.1}^{-1} \end{bmatrix} \quad (3.33)$$

where $\hat{\mathbf{Q}}_{11.2} = \hat{\mathbf{Q}}_{11} - \hat{\mathbf{Q}}_{12}\hat{\mathbf{Q}}_{22}^{-1}\hat{\mathbf{Q}}_{21}$ and $\hat{\mathbf{Q}}_{22.1} = \hat{\mathbf{Q}}_{22} - \hat{\mathbf{Q}}_{21}\hat{\mathbf{Q}}_{11}^{-1}\hat{\mathbf{Q}}_{12}$.

Thus

$$\begin{aligned} \hat{\boldsymbol{\beta}} &= \begin{pmatrix} \hat{\boldsymbol{\beta}}_1 \\ \hat{\boldsymbol{\beta}}_2 \end{pmatrix} \\ &= \begin{bmatrix} \hat{\mathbf{Q}}_{11.2}^{-1} & -\hat{\mathbf{Q}}_{11.2}^{-1}\hat{\mathbf{Q}}_{12}\hat{\mathbf{Q}}_{22}^{-1} \\ -\hat{\mathbf{Q}}_{22.1}^{-1}\hat{\mathbf{Q}}_{21}\hat{\mathbf{Q}}_{11}^{-1} & \hat{\mathbf{Q}}_{22.1}^{-1} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{Q}}_{1y} \\ \hat{\mathbf{Q}}_{2y} \end{bmatrix} \\ &= \begin{pmatrix} \hat{\mathbf{Q}}_{11.2}^{-1}\hat{\mathbf{Q}}_{1y.2} \\ \hat{\mathbf{Q}}_{22.1}^{-1}\hat{\mathbf{Q}}_{2y.1} \end{pmatrix} \end{aligned}$$

Now

$$\begin{aligned} \hat{\mathbf{Q}}_{11.2} &= \hat{\mathbf{Q}}_{11} - \hat{\mathbf{Q}}_{12}\hat{\mathbf{Q}}_{22}^{-1}\hat{\mathbf{Q}}_{21} \\ &= \frac{1}{n}\mathbf{X}'_1\mathbf{X}_1 - \frac{1}{n}\mathbf{X}'_1\mathbf{X}_2 \left(\frac{1}{n}\mathbf{X}'_2\mathbf{X}_2 \right)^{-1} \frac{1}{n}\mathbf{X}'_2\mathbf{X}_1 \\ &= \frac{1}{n}\mathbf{X}'_1\mathbf{M}_2\mathbf{X}_1 \end{aligned}$$

where

$$\mathbf{M}_2 = \mathbf{I}_n - \mathbf{X}_2(\mathbf{X}'_2\mathbf{X}_2)^{-1}\mathbf{X}'_2$$

is the orthogonal projection matrix for \mathbf{X}_2 . Similarly $\hat{\mathbf{Q}}_{22.1} = \frac{1}{n}\mathbf{X}'_2\mathbf{M}_1\mathbf{X}_2$ where

$$\mathbf{M}_1 = \mathbf{I}_n - \mathbf{X}_1(\mathbf{X}'_1\mathbf{X}_1)^{-1}\mathbf{X}'_1$$

is the orthogonal projection matrix for \mathbf{X}_1 . Also

$$\begin{aligned}\widehat{\mathbf{Q}}_{1y \cdot 2} &= \widehat{\mathbf{Q}}_{1y} - \widehat{\mathbf{Q}}_{12} \widehat{\mathbf{Q}}_{22}^{-1} \widehat{\mathbf{Q}}_{2y} \\ &= \frac{1}{n} \mathbf{X}'_1 \mathbf{y} - \frac{1}{n} \mathbf{X}'_1 \mathbf{X}_2 \left(\frac{1}{n} \mathbf{X}'_2 \mathbf{X}_2 \right)^{-1} \frac{1}{n} \mathbf{X}'_2 \mathbf{y} \\ &= \frac{1}{n} \mathbf{X}'_1 \mathbf{M}_2 \mathbf{y}\end{aligned}$$

and $\widehat{\mathbf{Q}}_{2y \cdot 1} = \frac{1}{n} \mathbf{X}'_2 \mathbf{M}_1 \mathbf{y}$.

Therefore

$$\widehat{\boldsymbol{\beta}}_1 = (\mathbf{X}'_1 \mathbf{M}_2 \mathbf{X}_1)^{-1} (\mathbf{X}'_1 \mathbf{M}_2 \mathbf{y}) \quad (3.34)$$

and

$$\widehat{\boldsymbol{\beta}}_2 = (\mathbf{X}'_2 \mathbf{M}_1 \mathbf{X}_2)^{-1} (\mathbf{X}'_2 \mathbf{M}_1 \mathbf{y}). \quad (3.35)$$

These are algebraic expressions for the sub-coefficient estimates from (3.32).

3.15 Residual Regression

As first recognized by Frisch and Waugh (1933), expressions (3.34) and (3.35) can be used to show that the least-squares estimators $\widehat{\boldsymbol{\beta}}_1$ and $\widehat{\boldsymbol{\beta}}_2$ can be found by a two-step regression procedure.

Take (3.35). Since \mathbf{M}_1 is idempotent, $\mathbf{M}_1 = \mathbf{M}_1 \mathbf{M}_1$ and thus

$$\begin{aligned}\widehat{\boldsymbol{\beta}}_2 &= (\mathbf{X}'_2 \mathbf{M}_1 \mathbf{X}_2)^{-1} (\mathbf{X}'_2 \mathbf{M}_1 \mathbf{y}) \\ &= (\mathbf{X}'_2 \mathbf{M}_1 \mathbf{M}_1 \mathbf{X}_2)^{-1} (\mathbf{X}'_2 \mathbf{M}_1 \mathbf{M}_1 \mathbf{y}) \\ &= (\widetilde{\mathbf{X}}'_2 \widetilde{\mathbf{X}}_2)^{-1} (\widetilde{\mathbf{X}}'_2 \widetilde{\mathbf{e}}_1)\end{aligned}$$

where

$$\widetilde{\mathbf{X}}_2 = \mathbf{M}_1 \mathbf{X}_2$$

and

$$\widetilde{\mathbf{e}}_1 = \mathbf{M}_1 \mathbf{y}.$$

Thus the coefficient estimate $\widehat{\boldsymbol{\beta}}_2$ is algebraically equal to the least-squares regression of $\widetilde{\mathbf{e}}_1$ on $\widetilde{\mathbf{X}}_2$. Notice that these two are \mathbf{y} and \mathbf{X}_2 , respectively, premultiplied by \mathbf{M}_1 . But we know that multiplication by \mathbf{M}_1 is equivalent to creating least-squares residuals. Therefore $\widetilde{\mathbf{e}}_1$ is simply the least-squares residual from a regression of \mathbf{y} on \mathbf{X}_1 , and the columns of $\widetilde{\mathbf{X}}_2$ are the least-squares residuals from the regressions of the columns of \mathbf{X}_2 on \mathbf{X}_1 .

We have proven the following theorem.

Theorem 3.15.1 Frisch-Waugh-Lovell

In the model (3.31), the OLS estimator of $\boldsymbol{\beta}_2$ and the OLS residuals $\widehat{\mathbf{e}}$ may be equivalently computed by either the OLS regression (3.32) or via the following algorithm:

1. Regress \mathbf{y} on \mathbf{X}_1 , obtain residuals $\widetilde{\mathbf{e}}_1$;
2. Regress \mathbf{X}_2 on \mathbf{X}_1 , obtain residuals $\widetilde{\mathbf{X}}_2$;
3. Regress $\widetilde{\mathbf{e}}_1$ on $\widetilde{\mathbf{X}}_2$, obtain OLS estimates $\widehat{\boldsymbol{\beta}}_2$ and residuals $\widehat{\mathbf{e}}$.

In some contexts, the FWL theorem can be used to speed computation, but in most cases there is little computational advantage to using the two-step algorithm.

This result is a direct analogy of the coefficient representation obtained in Section 2.22. The result obtained in that section concerned the population projection coefficients, the result obtained here concern the least-squares estimates. The key message is the same. In the least-squares regression (3.32), the estimated coefficient $\hat{\beta}_2$ numerically equals the regression of \mathbf{y} on the regressors \mathbf{X}_2 , only after the regressors \mathbf{X}_1 have been linearly projected out. Similarly, the coefficient estimate $\hat{\beta}_1$ numerically equals the regression of \mathbf{y} on the regressors \mathbf{X}_1 , after the regressors \mathbf{X}_2 have been linearly projected out. This result can be very insightful when interpreting regression coefficients.

A common application of the FWL theorem, which you may have seen in an introductory econometrics course, is the demeaning formula for regression. Partition $\mathbf{X} = [\mathbf{X}_1 \ \mathbf{X}_2]$ where $\mathbf{X}_1 = \mathbf{1}$ is a vector of ones and \mathbf{X}_2 is a matrix of observed regressors. In this case,

$$\mathbf{M}_1 = \mathbf{I}_n - \mathbf{1}(\mathbf{1}'\mathbf{1})^{-1}\mathbf{1}'.$$

Observe that

$$\widetilde{\mathbf{X}}_2 = \mathbf{M}_1\mathbf{X}_2 = \mathbf{X}_2 - \overline{\mathbf{X}}_2$$

and

$$\tilde{\mathbf{y}} = \mathbf{M}_1\mathbf{y} = \mathbf{y} - \bar{y}$$

are the “demeaned” variables. The FWL theorem says that $\hat{\beta}_2$ is the OLS estimate from a regression of $y_i - \bar{y}$ on $x_{2i} - \bar{x}_2$:

$$\hat{\beta}_2 = \left(\sum_{i=1}^n (x_{2i} - \bar{x}_2)(x_{2i} - \bar{x}_2)' \right)^{-1} \left(\sum_{i=1}^n (x_{2i} - \bar{x}_2)(y_i - \bar{y}) \right).$$

Thus the OLS estimator for the slope coefficients is a regression with demeaned data.

Ragnar Frisch

Ragnar Frisch (1895-1973) was co-winner with Jan Tinbergen of the first Nobel Memorial Prize in Economic Sciences in 1969 for their work in developing and applying dynamic models for the analysis of economic problems. Frisch made a number of foundational contributions to modern economics beyond the Frisch-Waugh-Lovell Theorem, including formalizing consumer theory, production theory, and business cycle theory.

3.16 Prediction Errors

The least-squares residual \hat{e}_i are not true prediction errors, as they are constructed based on the full sample including y_i . A proper prediction for y_i should be based on estimates constructed using only the other observations. We can do this by defining the **leave-one-out** OLS estimator of β as that obtained from the sample of $n - 1$ observations *excluding* the i 'th observation:

$$\begin{aligned} \hat{\beta}_{(-i)} &= \left(\frac{1}{n-1} \sum_{j \neq i} \mathbf{x}_j \mathbf{x}_j' \right)^{-1} \left(\frac{1}{n-1} \sum_{j \neq i} \mathbf{x}_j y_j \right) \\ &= \left(\mathbf{X}'_{(-i)} \mathbf{X}_{(-i)} \right)^{-1} \mathbf{X}_{(-i)} \mathbf{y}_{(-i)}. \end{aligned} \tag{3.36}$$

Here, $\mathbf{X}_{(-i)}$ and $\mathbf{y}_{(-i)}$ are the data matrices omitting the i 'th row. The leave-one-out predicted value for y_i is

$$\tilde{y}_i = \mathbf{x}'_i \widehat{\boldsymbol{\beta}}_{(-i)},$$

and the **leave-one-out residual** or **prediction error** or **prediction residual** is

$$\tilde{e}_i = y_i - \tilde{y}_i.$$

A convenient alternative expression for $\widehat{\boldsymbol{\beta}}_{(-i)}$ (derived in Section 3.21) is

$$\widehat{\boldsymbol{\beta}}_{(-i)} = \widehat{\boldsymbol{\beta}} - (1 - h_{ii})^{-1} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i \hat{e}_i \quad (3.37)$$

where h_{ii} are the leverage values as defined in (3.21).

Using (3.37) we can simplify the expression for the prediction error:

$$\begin{aligned} \tilde{e}_i &= y_i - \mathbf{x}'_i \widehat{\boldsymbol{\beta}}_{(-i)} \\ &= y_i - \mathbf{x}'_i \widehat{\boldsymbol{\beta}} + (1 - h_{ii})^{-1} \mathbf{x}'_i (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i \hat{e}_i \\ &= \hat{e}_i + (1 - h_{ii})^{-1} h_{ii} \hat{e}_i \\ &= (1 - h_{ii})^{-1} \hat{e}_i. \end{aligned} \quad (3.38)$$

To write this in vector notation, define

$$\begin{aligned} \mathbf{M}^* &= (\mathbf{I}_n - \text{diag}\{h_{11}, \dots, h_{nn}\})^{-1} \\ &= \text{diag}\{(1 - h_{11})^{-1}, \dots, (1 - h_{nn})^{-1}\}. \end{aligned} \quad (3.39)$$

Then (3.38) is equivalent to

$$\tilde{\mathbf{e}} = \mathbf{M}^* \hat{\mathbf{e}}. \quad (3.40)$$

A convenient feature of this expression is that it shows that computation of the full vector of prediction errors $\tilde{\mathbf{e}}$ is based on a simple linear operation, and does not really require n separate estimations.

One use of the prediction errors is to estimate the out-of-sample mean squared error

$$\begin{aligned} \tilde{\sigma}^2 &= \frac{1}{n} \sum_{i=1}^n \tilde{e}_i^2 \\ &= \frac{1}{n} \sum_{i=1}^n (1 - h_{ii})^{-2} \hat{e}_i^2. \end{aligned} \quad (3.41)$$

This is also known as the **sample mean squared prediction error**. Its square root $\tilde{\sigma} = \sqrt{\tilde{\sigma}^2}$ is the **prediction standard error**.

3.17 Influential Observations

Another use of the leave-one-out estimator is to investigate the impact of **influential observations**, sometimes called **outliers**. We say that observation i is influential if its omission from the sample induces a substantial change in a parameter estimate of interest.

For illustration, consider Figure 3.2 which shows a scatter plot of random variables (y_i, x_i) . The 25 observations shown with the open circles are generated by $x_i \sim U[1, 10]$ and $y_i \sim N(x_i, 4)$. The 26th observation shown with the filled circle is $x_{26} = 9$, $y_{26} = 0$. (Imagine that $y_{26} = 0$ was incorrectly recorded due to a mistaken key entry.) The Figure shows both the least-squares fitted line from the full sample and that obtained after deletion of the 26th observation from the sample. In this example we can see how the 26th observation (the “outlier”) greatly tilts the least-squares

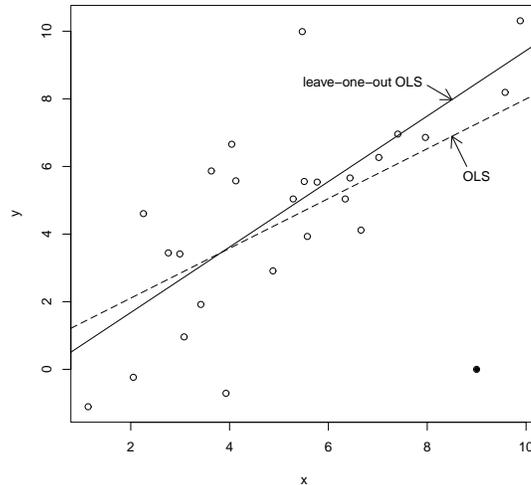


Figure 3.2: Impact of an influential observation on the least-squares estimator

fitted line towards the 26th observation. In fact, the slope coefficient decreases from 0.97 (which is close to the true value of 1.00) to 0.56, which is substantially reduced. Neither y_{26} nor x_{26} are unusual values relative to their marginal distributions, so this outlier would not have been detected from examination of the marginal distributions of the data. The change in the slope coefficient of -0.41 is meaningful and should raise concern to an applied economist.

From (3.37)-(3.38) we know that

$$\begin{aligned}\widehat{\boldsymbol{\beta}} - \widehat{\boldsymbol{\beta}}_{(-i)} &= (1 - h_{ii})^{-1} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i \hat{e}_i \\ &= (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i \tilde{e}_i.\end{aligned}\tag{3.42}$$

By direct calculation of this quantity for each observation i , we can directly discover if a specific observation i is influential for a coefficient estimate of interest.

For a general assessment, we can focus on the predicted values. The difference between the full-sample and leave-one-out predicted values is

$$\begin{aligned}\hat{y}_i - \tilde{y}_i &= \mathbf{x}_i' \widehat{\boldsymbol{\beta}} - \mathbf{x}_i' \widehat{\boldsymbol{\beta}}_{(-i)} \\ &= \mathbf{x}_i' (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i \tilde{e}_i \\ &= h_{ii} \tilde{e}_i\end{aligned}$$

which is a simple function of the leverage values h_{ii} and prediction errors \tilde{e}_i . Observation i is influential for the predicted value if $|h_{ii}\tilde{e}_i|$ is large, which requires that both h_{ii} and $|\tilde{e}_i|$ are large.

One way to think about this is that a large leverage value h_{ii} gives the potential for observation i to be influential. A large h_{ii} means that observation i is unusual in the sense that the regressor \mathbf{x}_i is far from its sample mean. We call an observation with large h_{ii} a **leverage point**. A leverage point is not necessarily influential as the latter also requires that the prediction error \tilde{e}_i is large.

To determine if any individual observations are influential in this sense, several diagnostics have been proposed (some names include DFITS, Cook's Distance, and Welsch Distance). Unfortunately, from a statistical perspective it is difficult to recommend these diagnostics for applications as they are not based on statistical theory. Probably the most relevant measure is the change in the coefficient estimates given in (3.42). The ratio of these changes to the coefficient's standard error is called its DFBETA, and is a postestimation diagnostic available in STATA. While there is no magic threshold, the concern is whether or not an individual observation meaningfully changes an

estimated coefficient of interest. A simple diagnostic for influential observations is to calculate

$$Influence = \max_{1 \leq i \leq n} |\hat{y}_i - \tilde{y}_i| = \max_{1 \leq i \leq n} |h_{ii}\tilde{e}_i|.$$

This is the largest (absolute) change in the predicted value due to a single observation. If this diagnostic is large relative to the distribution of y_i , it may indicate that that observation is influential.

If an observation is determined to be influential, what should be done? As a common cause of influential observations is data entry error, the influential observations should be examined for evidence that the observation was mis-recorded. Perhaps the observation falls outside of permitted ranges, or some observables are inconsistent (for example, a person is listed as having a job but receives earnings of \$0). If it is determined that an observation is incorrectly recorded, then the observation is typically deleted from the sample. This process is often called “cleaning the data”. The decisions made in this process involve an fair amount of individual judgement. When this is done it is proper empirical practice to document such choices. (It is useful to keep the source data in its original form, a revised data file after cleaning, and a record describing the revision process. This is especially useful when revising empirical work at a later date.)

It is also possible that an observation is correctly measured, but unusual and influential. In this case it is unclear how to proceed. Some researchers will try to alter the specification to properly model the influential observation. Other researchers will delete the observation from the sample. The motivation for this choice is to prevent the results from being skewed or determined by individual observations, but this practice is viewed skeptically by many researchers who believe it reduces the integrity of reported empirical results.

For an empirical illustration, consider the log wage regression (3.12) for single asian males. This regression, which has 268 observations, has $Influence = 0.29$. This means that the most influential observation, when deleted, changes the predicted (fitted) value of the dependent variable $\log(Wage)$ by 0.29, or equivalently the wage by 29%. This is a meaningful change and suggests further investigation. We examine the influential observation, and find that its leverage h_{ii} is 0.33, which is disturbingly large. (Recall that the leverage values are all positive and sum to one. One third of the leverage in this sample of 268 observations is contained in just this single observation!) Examining further, we find that this individual is 65 years old with 8 years education, so that his potential experience is 51 years. This is the highest experience in the subsample – the next highest is 41 years. The large leverage is due to his unusual characteristics (very low education and very high experience) within this sample. Essentially, regression (3.12) is attempting to estimate the conditional mean at $experience = 51$ with only one observation, so it is not surprising that this observation determines the fit and is thus influential. A reasonable conclusion is the regression function can only be estimated over a smaller range of $experience$. We restrict the sample to individuals with less than 45 years experience, re-estimate, and obtain the following estimates.

$$\widehat{\log(Wage)} = 0.144 \textit{education} + 0.043 \textit{experience} - 0.095 \textit{experience}^2/100 + 0.531. \quad (3.43)$$

For this regression, we calculate that $Influence = 0.11$, which is greatly reduced relative to the regression (3.12). Comparing (3.43) with (3.12), the slope coefficient for education is essentially unchanged, but the coefficients in experience and its square have slightly increased.

By eliminating the influential observation, equation (3.43) can be viewed as a more robust estimate of the conditional mean for most levels of $experience$. Whether to report (3.12) or (3.43) in an application is largely a matter of judgment.

3.18 Normal Regression Model

The normal regression model is the linear regression model under the restriction that the error e_i is independent of \mathbf{x}_i and has the distribution $N(0, \sigma^2)$. We can write this as

$$e_i \mid \mathbf{x}_i \sim N(0, \sigma^2).$$

This assumption implies

$$y_i \mid \mathbf{x}_i \sim N(\mathbf{x}_i' \boldsymbol{\beta}, \sigma^2).$$

Normal regression is a parametric model, where likelihood methods can be used for estimation, testing, and distribution theory.

The log-likelihood function for the normal regression model is

$$\begin{aligned} \log L(\boldsymbol{\beta}, \sigma^2) &= \sum_{i=1}^n \log \left(\frac{1}{(2\pi\sigma^2)^{1/2}} \exp \left(-\frac{1}{2\sigma^2} (y_i - \mathbf{x}_i' \boldsymbol{\beta})^2 \right) \right) \\ &= -\frac{n}{2} \log(2\pi) - \frac{n}{2} \log(\sigma^2) - \frac{1}{2\sigma^2} SSE_n(\boldsymbol{\beta}). \end{aligned} \quad (3.44)$$

The maximum likelihood estimator (MLE) $(\hat{\boldsymbol{\beta}}_{\text{mle}}, \hat{\sigma}_{\text{mle}}^2)$ maximizes $\log L(\boldsymbol{\beta}, \sigma^2)$. Since the latter is a function of $\boldsymbol{\beta}$ only through the sum of squared errors $SSE_n(\boldsymbol{\beta})$, maximizing the likelihood is identical to minimizing $SSE_n(\boldsymbol{\beta})$. Hence

$$\hat{\boldsymbol{\beta}}_{\text{mle}} = \hat{\boldsymbol{\beta}}_{\text{ols}},$$

the MLE for $\boldsymbol{\beta}$ equals the OLS estimator. Due to this equivalence, the least squares estimator $\hat{\boldsymbol{\beta}}_{\text{ols}}$ is often called the MLE.

We can also find the MLE for σ^2 . Plugging $\hat{\boldsymbol{\beta}}$ into the log-likelihood we obtain

$$\log L(\hat{\boldsymbol{\beta}}_{\text{mle}}, \sigma^2) = -\frac{n}{2} \log(2\pi) - \frac{n}{2} \log(\sigma^2) - \frac{SSE_n(\hat{\boldsymbol{\beta}}_{\text{mle}})}{2\sigma^2}.$$

Maximization with respect to σ^2 yields the first-order condition

$$\frac{\partial}{\partial \sigma^2} \log L(\hat{\boldsymbol{\beta}}_{\text{mle}}, \hat{\sigma}^2) = -\frac{n}{2\hat{\sigma}^2} + \frac{1}{2(\hat{\sigma}^2)^2} SSE_n(\hat{\boldsymbol{\beta}}_{\text{mle}}) = 0.$$

Solving for $\hat{\sigma}^2$ yields the MLE for σ^2

$$\hat{\sigma}_{\text{mle}}^2 = \frac{SSE_n(\hat{\boldsymbol{\beta}}_{\text{mle}})}{n} = \frac{1}{n} \sum_{i=1}^n \hat{e}_i^2$$

which is the same as the moment estimator (3.28).

Plugging the estimates into (3.44) we obtain the maximized log-likelihood

$$\log L(\hat{\boldsymbol{\beta}}_{\text{mle}}, \hat{\sigma}_{\text{mle}}^2) = -\frac{n}{2} (\log(2\pi) + 1) - \frac{n}{2} \log(\hat{\sigma}_{\text{mle}}^2). \quad (3.45)$$

The log-likelihood (or the negative log-likelihood) is typically reported as a measure of fit.

It may seem surprising that the MLE $\hat{\boldsymbol{\beta}}_{\text{mle}}$ is numerically equal to the OLS estimator, despite emerging from quite different motivations. It is not completely accidental. The least-squares estimator minimizes a particular sample loss function – the sum of squared error criterion – and most loss functions are equivalent to the likelihood of a specific parametric distribution, in this case the normal regression model. In this sense it is not surprising that the least-squares estimator can be motivated as either the minimizer of a sample loss function or as the maximizer of a likelihood function.

Carl Friedrich Gauss

The mathematician Carl Friedrich Gauss (1777-1855) proposed the normal regression model, and derived the least squares estimator as the maximum likelihood estimator for this model. He claimed to have discovered the method in 1795 at the age of eighteen, but did not publish the result until 1809. Interest in Gauss's approach was reinforced by Laplace's simultaneous discovery of the central limit theorem, which provided a justification for viewing random disturbances as approximately normal.

3.19 CPS Data Set

In this section we describe the data set used in the empirical illustrations.

The Current Population Survey (CPS) is a monthly survey of about 57,000 U.S. households conducted by the Bureau of the Census of the Bureau of Labor Statistics. The CPS is the primary source of information on the labor force characteristics of the U.S. population. The survey covers employment, earnings, educational attainment, income, poverty, health insurance coverage, job experience, voting and registration, computer usage, veteran status, and other variables. Details can be found at www.census.gov/cps and dataferrett.census.gov.

From the March 2009 survey we extracted the individuals with non-allocated variables who were full-time employed (defined as those who had worked at least 36 hours per week for at least 48 weeks the past year), and excluded those in the military. This sample has 50,742 individuals. We extracted 14 variables from the CPS on these individuals and created the data files `cps09mar.dta` (Stata format) and `cps09mar.txt` (text format). The variables are:

1. age: years, capped at 85
2. female: 1 if female, 0 otherwise
3. hisp: 1 if Spanish, Hispanic, or Latino, 0 otherwise
4. education

0	Less than 1st grade
4	1st, 2nd, 3rd, or 4th grade
6	5th or 6th grade
8	7th or 8th grade
9	9th grade
10	10th grade
11	11th grade or 12th grade with no high school diploma
12	High school graduate, high school diploma or equivalent
13	Some college but no degree
14	Associate degree in college, including occupation/vocation programs
16	Bachelor's degree or equivalent (BA, AB, BS)
18	Master's degree (MA, MS MENG, MED, MSW, MBA)
20	Professional degree or Doctorate degree (MD, DDS, DVM, LLB, JD, PHD, EDD)
5. earnings: total annual wage and salary earnings

6. hours: number of hours worked per week
7. week: number of weeks worked per year
8. union: 1 for member of a labor union, 0 otherwise
9. uncov: 1 if covered by a union or employee association contract, 0 otherwise
10. region
 - 1 Northeast
 - 2 Midwest
 - 3 South
 - 4 West
11. Race
 - 1 White only
 - 2 Black only
 - 3 American Indian, Alaskan Native (AI) only
 - 4 Asian only
 - 5 Hawaiian/Pacific Islander (HP) only
 - 6 White-Black
 - 7 White-AI
 - 8 White-Asian
 - 9 White-HP
 - 10 Black-AI
 - 11 Black-Asian
 - 12 Black-HP
 - 13 AI-Asian
 - 14 Asian-HP
 - 15 White-Black-AI
 - 16 White-Black-Asian
 - 17 White-AI-Asian
 - 18 White-Asian-HP
 - 19 White-Black-AI-Asian
 - 20 2 or 3 races
 - 21 4 or 5 races
12. marital
 - 1 Married - civilian spouse present
 - 2 Married - Armed Forces spouse present
 - 3 Married - spouse absent (except separated)
 - 4 Widowed
 - 5 Divorced
 - 6 Separated
 - 7 Never married

3.20 Programming

Most packages allow both interactive programming (where you enter commands one-by-one) and batch programming (where you run a pre-written sequence of commands from a file). Interactive programming can be useful for exploratory analysis, but eventually all work should be executed in batch mode. This is the best way to control and document your work.

Batch programs are text files where each line executes a single command. For Stata, this file needs to have the filename extension “.do”, and for Matlab “.m”, while for Gauss and R there are no specific naming requirements.

To execute a program file, you type a command within the program.

Stata: `do chapter3` executes the file *chapter3.do*

Gauss: `run chapter3.prg` executes the file *chapter3.prg*

Matlab: `run chapter3` executes the file *chapter3.m*

R: `source("chapter3.r")` executes the file *chapter3.r*

When writing batch files, it is useful to include comments for documentation and readability.

We illustrate programming files for Stata, Gauss, R, and Matlab, which execute a portion of the empirical illustrations from Sections 3.7 and 3.17.

Stata do File

```
*      Clear memory and load the data
clear
use cps09mar.dta
*      Generate transformations
gen wage=ln(earnings/(hours*week))
gen experience = age - education - 6
gen exp2 = (experience^2)/100
*      Create indicator for subsamples
gen mbf = (race == 2) & (marital <= 2) & (female == 1)
gen sam = (race == 4) & (marital == 7) & (female == 0)
*      Regressions
reg wage education if (mbf == 1) & (experience == 12)
reg wage education experience exp2 if sam == 1
*      Leverage and influence
predict leverage,hat
predict e,residual
gen d=e*leverage/(1-leverage)
summarize d if sam ==1
```

Gauss Program File

```

/* Load the data and create subsamples */
load dat[50742,12]=cps09mar.txt;
experience=dat[:,1]-dat[:,4]-6;
mbf=(dat[:,11].==2).*(dat[:,12].<=2).*(dat[:,2].==1).*(experience.==12);
sam=(dat[:,11].==4).*(dat[:,12].==7).*(dat[:,2].==0);
dat1=selif(dat,mbf);
dat2=selif(dat,sam);
/* First regression */
y=ln(dat1[:,5]./(dat1[:,6].*dat1[:,7]));
x=dat1[:,4]~ones(rows(dat1),1);
beta=invpd(x'x)*(x'y);
beta;
/* Second regression */
y=ln(dat2[:,5]./(dat2[:,6].*dat2[:,7]));
experience=dat2[:,1]-dat2[:,4]-6;
exp2 = (experience.^2)/100;
x=dat2[:,4]~experience~exp2~ones(rows(dat2),1);
beta=invpd(x'x)*(x'y);
beta;
/* Create leverage and influence */
e=y-x*beta;
leverage=sumc((x.*(x*invpd(x'x))))');
d=leverage.*e./(1-leverage);
"Influence " maxc(abs(d));

```

R Program File

```

# Load the data and create subsamples
dat <- read.table("cps09mar.txt")
experience <- dat[,1]-dat[,4]-6
mbf <- (dat[,11]==2)&(dat[,12]<=2)&(dat[,2]==1)&(experience==12)
sam <- (dat[,11]==4)&(dat[,12]==7)&(dat[,2]==0)
dat1 <- dat[mbf,]
dat2 <- dat[sam,]
# First regression
y <- as.matrix(log(dat1[,5]/(dat1[,6]*dat1[,7])))
x <- cbind(dat1[,4],matrix(1,nrow(dat1),1))
beta <- solve(t(x)%*%x,t(x)%*%y)
print(beta)
# Second regression
y <- as.matrix(log(dat2[,5]/(dat2[,6]*dat2[,7])))
experience <- dat2[,1]-dat2[,4]-6
exp2 <- (experience^2)/100
x <- cbind(dat2[,4],experience,exp2,matrix(1,nrow(dat2),1))
beta <- solve(t(x)%*%x,t(x)%*%y)
print(beta)
# Create leverage and influence
e <- y-x*beta
leverage <- rowSums(x*(x*solve(t(x)%*%x)))
r <- e/(1-leverage)
d <- leverage*e/(1-leverage)
print(max(abs(d)))

```

Matlab Program File

```

% Load the data and create subsamples
load cps09mar.txt;
dat=cps09mar;
experience=dat(:,1)-dat(:,4)-6;
mbf = (dat(:,11)==2)&(dat(:,12)<=2)&(dat(:,2)==1)&(experience==12);
sam = (dat(:,11)==4)&(dat(:,12)==7)&(dat(:,2)==0);
dat1=dat(mbf,:);
dat2=dat(sam,:);
% First regression
y=log(dat1(:,5)./(dat1(:,6).*dat1(:,7)));
x=[dat1(:,4),ones(length(dat1),1)];
beta=inv(x'*x)*(x'*y);
display(beta);
% Second regression
y=log(dat2(:,5)./(dat2(:,6).*dat2(:,7)));
experience=dat2(:,1)-dat2(:,4)-6;
exp2 = (experience.^2)/100;
x=[dat2(:,4),experience,exp2,ones(length(dat2),1)];
beta=inv(x'*x)*(x'*y);
display(beta);
% Create leverage and influence
e=y-x*beta;
leverage=sum((x.*(x*inv(x'*x))))';
d=leverage.*e./(1-leverage);
influence=max(abs(d));
display(influence);

```

3.21 Technical Proofs*

Proof of Theorem 3.10.1, equation (3.23): First, $h_{ii} = \mathbf{x}'_i (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i \geq 0$ since it is a quadratic form and $\mathbf{X}'\mathbf{X} > 0$. Next, since h_{ii} is the i 'th diagonal element of the projection matrix $\mathbf{P} = \mathbf{X} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}$, then

$$h_{ii} = \mathbf{s}'\mathbf{P}\mathbf{s}$$

where

$$\mathbf{s} = \begin{pmatrix} 0 \\ \vdots \\ 1 \\ \vdots \\ 0 \end{pmatrix}$$

is a unit vector with a 1 in the i 'th place (and zeros elsewhere).

By the spectral decomposition of the idempotent matrix \mathbf{P} (see equation (A.5))

$$\mathbf{P} = \mathbf{B}' \begin{bmatrix} \mathbf{I}_k & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{B}$$

where $\mathbf{B}'\mathbf{B} = \mathbf{I}_n$. Thus letting $\mathbf{b} = \mathbf{B}\mathbf{s}$ denote the i 'th column of \mathbf{B} , and partitioning $\mathbf{b}' = (\mathbf{b}'_1 \quad \mathbf{b}'_2)$ then

$$\begin{aligned} h_{ii} &= \mathbf{s}'\mathbf{B}' \begin{bmatrix} \mathbf{I}_k & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{B}\mathbf{s} \\ &= \mathbf{b}'_1 \begin{bmatrix} \mathbf{I}_k & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{b}_1 \\ &= \mathbf{b}'_1 \mathbf{b}_1 \\ &\leq \mathbf{b}'\mathbf{b} \\ &= 1 \end{aligned}$$

the final equality since \mathbf{b} is the i 'th column of \mathbf{B} and $\mathbf{B}'\mathbf{B} = \mathbf{I}_n$. We have shown that $h_{ii} \leq 1$, establishing (3.23). ■

Proof of Equation (3.37). The Sherman–Morrison formula (A.3) from Appendix A.5 states that for nonsingular \mathbf{A} and vector \mathbf{b}

$$(\mathbf{A} - \mathbf{b}\mathbf{b}')^{-1} = \mathbf{A}^{-1} + (1 - \mathbf{b}'\mathbf{A}^{-1}\mathbf{b})^{-1} \mathbf{A}^{-1}\mathbf{b}\mathbf{b}'\mathbf{A}^{-1}.$$

This implies

$$(\mathbf{X}'\mathbf{X} - \mathbf{x}_i\mathbf{x}'_i)^{-1} = (\mathbf{X}'\mathbf{X})^{-1} + (1 - h_{ii})^{-1} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i\mathbf{x}'_i (\mathbf{X}'\mathbf{X})^{-1}$$

and thus

$$\begin{aligned} \widehat{\boldsymbol{\beta}}_{(-i)} &= (\mathbf{X}'\mathbf{X} - \mathbf{x}_i\mathbf{x}'_i)^{-1} (\mathbf{X}'\mathbf{y} - \mathbf{x}_iy_i) \\ &= (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{y} - (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_iy_i \\ &\quad + (1 - h_{ii})^{-1} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i\mathbf{x}'_i (\mathbf{X}'\mathbf{X})^{-1} (\mathbf{X}'\mathbf{y} - \mathbf{x}_iy_i) \\ &= \widehat{\boldsymbol{\beta}} - (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_iy_i + (1 - h_{ii})^{-1} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i (\mathbf{x}'_i\widehat{\boldsymbol{\beta}} - h_{ii}y_i) \\ &= \widehat{\boldsymbol{\beta}} - (1 - h_{ii})^{-1} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i \left((1 - h_{ii})y_i - \mathbf{x}'_i\widehat{\boldsymbol{\beta}} + h_{ii}y_i \right) \\ &= \widehat{\boldsymbol{\beta}} - (1 - h_{ii})^{-1} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i\hat{\epsilon}_i \end{aligned}$$

the third equality making the substitutions $\widehat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{y}$ and $h_{ii} = \mathbf{x}'_i (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i$, and the remainder collecting terms. ■

Exercises

Exercise 3.1 Let y be a random variable with $\mu = \mathbb{E}y$ and $\sigma^2 = \text{var}(y)$. Define

$$g(y, \mu, \sigma^2) = \begin{pmatrix} y - \mu \\ (y - \mu)^2 - \sigma^2 \end{pmatrix}.$$

Let $(\hat{\mu}, \hat{\sigma}^2)$ be the values such that $\bar{g}_n(\hat{\mu}, \hat{\sigma}^2) = \mathbf{0}$ where $\bar{g}_n(m, s) = n^{-1} \sum_{i=1}^n g(y_i, m, s)$. Show that $\hat{\mu}$ and $\hat{\sigma}^2$ are the sample mean and variance.

Exercise 3.2 Consider the OLS regression of the $n \times 1$ vector \mathbf{y} on the $n \times k$ matrix \mathbf{X} . Consider an alternative set of regressors $\mathbf{Z} = \mathbf{X}\mathbf{C}$, where \mathbf{C} is a $k \times k$ non-singular matrix. Thus, each column of \mathbf{Z} is a mixture of some of the columns of \mathbf{X} . Compare the OLS estimates and residuals from the regression of \mathbf{y} on \mathbf{X} to the OLS estimates from the regression of \mathbf{y} on \mathbf{Z} .

Exercise 3.3 Using matrix algebra, show $\mathbf{X}'\hat{\mathbf{e}} = \mathbf{0}$.

Exercise 3.4 Let $\hat{\mathbf{e}}$ be the OLS residual from a regression of \mathbf{y} on $\mathbf{X} = [\mathbf{X}_1 \ \mathbf{X}_2]$. Find $\mathbf{X}_2'\hat{\mathbf{e}}$.

Exercise 3.5 Let $\hat{\mathbf{e}}$ be the OLS residual from a regression of \mathbf{y} on \mathbf{X} . Find the OLS coefficient from a regression of $\hat{\mathbf{e}}$ on \mathbf{X} .

Exercise 3.6 Let $\hat{\mathbf{y}} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}$. Find the OLS coefficient from a regression of $\hat{\mathbf{y}}$ on \mathbf{X} .

Exercise 3.7 Show that if $\mathbf{X} = [\mathbf{X}_1 \ \mathbf{X}_2]$ then $\mathbf{P}\mathbf{X}_1 = \mathbf{X}_1$ and $\mathbf{M}\mathbf{X}_1 = \mathbf{0}$.

Exercise 3.8 Show that \mathbf{M} is idempotent: $\mathbf{M}\mathbf{M} = \mathbf{M}$.

Exercise 3.9 Show that $\text{tr } \mathbf{M} = n - k$.

Exercise 3.10 Show that if $\mathbf{X} = [\mathbf{X}_1 \ \mathbf{X}_2]$ and $\mathbf{X}_1'\mathbf{X}_2 = 0$ then $\mathbf{P} = \mathbf{P}_1 + \mathbf{P}_2$.

Exercise 3.11 Show that when \mathbf{X} contains a constant, $\frac{1}{n} \sum_{i=1}^n \hat{y}_i = \bar{y}$.

Exercise 3.12 A dummy variable takes on only the values 0 and 1. It is used for categorical data, such as an individual's gender. Let \mathbf{d}_1 and \mathbf{d}_2 be vectors of 1's and 0's, with the i 'th element of \mathbf{d}_1 equaling 1 and that of \mathbf{d}_2 equaling 0 if the person is a man, and the reverse if the person is a woman. Suppose that there are n_1 men and n_2 women in the sample. Consider fitting the following three equations by OLS

$$\mathbf{y} = \mu + \mathbf{d}_1\alpha_1 + \mathbf{d}_2\alpha_2 + \mathbf{e} \tag{3.46}$$

$$\mathbf{y} = \mathbf{d}_1\alpha_1 + \mathbf{d}_2\alpha_2 + \mathbf{e} \tag{3.47}$$

$$\mathbf{y} = \mu + \mathbf{d}_1\phi + \mathbf{e} \tag{3.48}$$

Can all three equations (3.46), (3.47), and (3.48) be estimated by OLS? Explain if not.

- Compare regressions (3.47) and (3.48). Is one more general than the other? Explain the relationship between the parameters in (3.47) and (3.48).
- Compute $\boldsymbol{\iota}'\mathbf{d}_1$ and $\boldsymbol{\iota}'\mathbf{d}_2$, where $\boldsymbol{\iota}$ is an $n \times 1$ vector of ones.
- Letting $\boldsymbol{\alpha} = (\alpha_1 \ \alpha_2)'$, write equation (3.47) as $\mathbf{y} = \mathbf{X}\boldsymbol{\alpha} + \mathbf{e}$. Consider the assumption $\mathbb{E}(\mathbf{x}_i e_i) = 0$. Is there any content to this assumption in this setting?

Exercise 3.13 Let \mathbf{d}_1 and \mathbf{d}_2 be defined as in the previous exercise.

(a) In the OLS regression

$$\mathbf{y} = \mathbf{d}_1\hat{\gamma}_1 + \mathbf{d}_2\hat{\gamma}_2 + \hat{\mathbf{u}},$$

show that $\hat{\gamma}_1$ is the sample mean of the dependent variable among the men of the sample (\bar{y}_1), and that $\hat{\gamma}_2$ is the sample mean among the women (\bar{y}_2).

(b) Let \mathbf{X} ($n \times k$) be an additional matrix of regressors. Describe in words the transformations

$$\begin{aligned}\mathbf{y}^* &= \mathbf{y} - \mathbf{d}_1\bar{y}_1 - \mathbf{d}_2\bar{y}_2 \\ \mathbf{X}^* &= \mathbf{X} - \mathbf{d}_1\bar{\mathbf{x}}_1' - \mathbf{d}_2\bar{\mathbf{x}}_2'\end{aligned}$$

where $\bar{\mathbf{x}}_1$ and $\bar{\mathbf{x}}_2$ are the $k \times 1$ means of the regressors for men and women, respectively.

(c) Compare $\tilde{\boldsymbol{\beta}}$ from the OLS regression

$$\mathbf{y}^* = \mathbf{X}^*\tilde{\boldsymbol{\beta}} + \tilde{\mathbf{e}}$$

with $\hat{\boldsymbol{\beta}}$ from the OLS regression

$$\mathbf{y} = \mathbf{d}_1\hat{\alpha}_1 + \mathbf{d}_2\hat{\alpha}_2 + \mathbf{X}\hat{\boldsymbol{\beta}} + \hat{\mathbf{e}}.$$

Exercise 3.14 Let $\hat{\boldsymbol{\beta}}_n = (\mathbf{X}'_n\mathbf{X}_n)^{-1}\mathbf{X}'_n\mathbf{y}_n$ denote the OLS estimate when \mathbf{y}_n is $n \times 1$ and \mathbf{X}_n is $n \times k$. A new observation ($y_{n+1}, \mathbf{x}_{n+1}$) becomes available. Prove that the OLS estimate computed using this additional observation is

$$\hat{\boldsymbol{\beta}}_{n+1} = \hat{\boldsymbol{\beta}}_n + \frac{1}{1 + \mathbf{x}'_{n+1}(\mathbf{X}'_n\mathbf{X}_n)^{-1}\mathbf{x}_{n+1}} (\mathbf{X}'_n\mathbf{X}_n)^{-1}\mathbf{x}_{n+1} (y_{n+1} - \mathbf{x}'_{n+1}\hat{\boldsymbol{\beta}}_n).$$

Exercise 3.15 Prove that R^2 is the square of the sample correlation between \mathbf{y} and $\hat{\mathbf{y}}$.

Exercise 3.16 Consider two least-squares regressions

$$\mathbf{y} = \mathbf{X}_1\tilde{\boldsymbol{\beta}}_1 + \tilde{\mathbf{e}}$$

and

$$\mathbf{y} = \mathbf{X}_1\hat{\boldsymbol{\beta}}_1 + \mathbf{X}_2\hat{\boldsymbol{\beta}}_2 + \hat{\mathbf{e}}.$$

Let R_1^2 and R_2^2 be the R -squared from the two regressions. Show that $R_2^2 \geq R_1^2$. Is there a case (explain) when there is equality $R_2^2 = R_1^2$?

Exercise 3.17 Show that $\tilde{\sigma}^2 \geq \hat{\sigma}^2$. Is equality possible?

Exercise 3.18 For which observations will $\hat{\boldsymbol{\beta}}_{(-i)} = \hat{\boldsymbol{\beta}}$?

Exercise 3.19 Use the data set from Section 3.19 and the sub-sample used for equation (3.43) (see Section 3.20) for data construction)

1. Estimate equation (3.43) and compute the equation R^2 and sum of squared errors.
2. Re-estimate the slope on education using the residual regression approach. Regress $\log(\text{Wage})$ on experience and its square, regress education on experience and its square, and the residuals on the residuals. Report the estimates from this final regression, along with the equation R^2 and sum of squared errors. Does the slope coefficient equal the value in (3.43)? Explain.
3. Do the R^2 and sum-of-squared errors from parts 1 and 2 equal? Explain.

Exercise 3.20 Estimate equation (3.43) as in part 1 of the previous question. Let \hat{e}_i be the OLS residual, \hat{y}_i the predicted value from the regression, x_{1i} be education and x_{2i} be experience. Numerically calculate the following:

- (a) $\sum_{i=1}^n \hat{e}_i$
- (b) $\sum_{i=1}^n x_{1i} \hat{e}_i$
- (c) $\sum_{i=1}^n x_{2i} \hat{e}_i$
- (d) $\sum_{i=1}^n x_{1i}^2 \hat{e}_i$
- (e) $\sum_{i=1}^n x_{2i}^2 \hat{e}_i$
- (f) $\sum_{i=1}^n \hat{y}_i \hat{e}_i$
- (g) $\sum_{i=1}^n \hat{e}_i^2$

Are these calculations consistent with the theoretical properties of OLS? Explain.

Exercise 3.21 Use the data set from Section 3.19.

1. Estimate a log wage regression for the subsample of white male Hispanics. In addition to education, experience, and its square, include a set of binary variables for regions and marital status. For regions, you create dummy variables for Northeast, South and West so that Midwest is the excluded group. For marital status, create variables for married, widowed or divorced, and separated, so that single (never married) is the excluded group.
2. Repeat this estimation using a different econometric package. Compare your results. Do they agree?

Chapter 4

Least Squares Regression

4.1 Introduction

In this chapter we investigate some finite-sample properties of least-squares applied to a random sample in the the linear regression model. In particular, we calculate the finite-sample mean and covariance matrix and propose standard errors for the coefficient estimates.

4.2 Sample Mean

To start with the simplest setting, we first consider the intercept-only model

$$y_i = \mu + e_i$$
$$\mathbb{E}(e_i) = 0.$$

which is equivalent to the regression model with $k = 1$ and $x_i = 1$. In the intercept model, $\mu = \mathbb{E}(y_i)$ is the mean of y_i . (See Exercise 2.15.) The least-squares estimator $\hat{\mu} = \bar{y}$ equals the sample mean as shown in equation (3.7).

We now calculate the mean and variance of the estimator \bar{y} . Since the sample mean is a linear function of the observations, its expectation is simple to calculate

$$\mathbb{E}(\bar{y}) = \mathbb{E}\left(\frac{1}{n} \sum_{i=1}^n y_i\right) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}y_i = \mu.$$

This shows that the expected value of least-squares estimator (the sample mean) equals the projection coefficient (the population mean). An estimator with the property that its expectation equals the parameter it is estimating is called **unbiased**.

Definition 4.2.1 An estimator $\hat{\theta}$ for θ is **unbiased** if $\mathbb{E}\hat{\theta} = \theta$.

We next calculate the variance of the estimator \bar{y} . Making the substitution $y_i = \mu + e_i$ we find

$$\bar{y} - \mu = \frac{1}{n} \sum_{i=1}^n e_i.$$

Then

$$\begin{aligned}
 \text{var}(\bar{y}) &= \mathbb{E}(\bar{y} - \mu)^2 \\
 &= \mathbb{E}\left(\frac{1}{n} \sum_{i=1}^n e_i\right) \left(\frac{1}{n} \sum_{j=1}^n e_j\right) \\
 &= \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \mathbb{E}(e_i e_j) \\
 &= \frac{1}{n^2} \sum_{i=1}^n \sigma^2 \\
 &= \frac{1}{n} \sigma^2.
 \end{aligned}$$

The second-to-last inequality is because $\mathbb{E}(e_i e_j) = \sigma^2$ for $i = j$ yet $\mathbb{E}(e_i e_j) = 0$ for $i \neq j$ due to independence.

We have shown that $\text{var}(\bar{y}) = \frac{1}{n} \sigma^2$. This is the familiar formula for the variance of the sample mean.

4.3 Linear Regression Model

We now consider the linear regression model. Throughout the remainder of this chapter we maintain the following.

Assumption 4.3.1 *Linear Regression Model*

The observations (y_i, \mathbf{x}_i) come from a random sample and satisfy the linear regression equation

$$y_i = \mathbf{x}_i' \boldsymbol{\beta} + e_i \quad (4.1)$$

$$\mathbb{E}(e_i \mid \mathbf{x}_i) = 0. \quad (4.2)$$

The variables have finite second moments

$$\mathbb{E} y_i^2 < \infty,$$

$$\mathbb{E} \|\mathbf{x}_i\|^2 < \infty,$$

and an invertible design matrix

$$\mathbf{Q}_{\mathbf{x}\mathbf{x}} = \mathbb{E}(\mathbf{x}_i \mathbf{x}_i') > 0.$$

We will consider both the general case of heteroskedastic regression, where the conditional variance

$$\mathbb{E}(e_i^2 \mid \mathbf{x}_i) = \sigma^2(\mathbf{x}_i) = \sigma_i^2$$

is unrestricted, and the specialized case of homoskedastic regression, where the conditional variance is constant. In the latter case we add the following assumption.

Assumption 4.3.2 Homoskedastic Linear Regression Model

In addition to Assumption 4.3.1,

$$\mathbb{E}(e_i^2 | \mathbf{x}_i) = \sigma^2(\mathbf{x}_i) = \sigma^2 \quad (4.3)$$

is independent of \mathbf{x}_i .

4.4 Mean of Least-Squares Estimator

In this section we show that the OLS estimator is unbiased in the linear regression model. This calculation can be done using either summation notation or matrix notation. We will use both.

First take summation notation. Observe that under (4.1)-(4.2)

$$\mathbb{E}(y_i | \mathbf{X}) = \mathbb{E}(y_i | \mathbf{x}_i) = \mathbf{x}_i' \boldsymbol{\beta}. \quad (4.4)$$

The first equality states that the conditional expectation of y_i given $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ only depends on \mathbf{x}_i , since the observations are independent across i . The second equality is the assumption of a linear conditional mean.

Using definition (3.9), the conditioning theorem, the linearity of expectations, (4.4), and properties of the matrix inverse,

$$\begin{aligned} \mathbb{E}(\hat{\boldsymbol{\beta}} | \mathbf{X}) &= \mathbb{E}\left(\left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i'\right)^{-1} \left(\sum_{i=1}^n \mathbf{x}_i y_i\right) | \mathbf{X}\right) \\ &= \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i'\right)^{-1} \mathbb{E}\left(\left(\sum_{i=1}^n \mathbf{x}_i y_i\right) | \mathbf{X}\right) \\ &= \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i'\right)^{-1} \sum_{i=1}^n \mathbb{E}(\mathbf{x}_i y_i | \mathbf{X}) \\ &= \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i'\right)^{-1} \sum_{i=1}^n \mathbf{x}_i \mathbb{E}(y_i | \mathbf{X}) \\ &= \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i'\right)^{-1} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \boldsymbol{\beta} \\ &= \boldsymbol{\beta}. \end{aligned}$$

Now let's show the same result using matrix notation. (4.4) implies

$$\mathbb{E}(\mathbf{y} | \mathbf{X}) = \begin{pmatrix} \vdots \\ \mathbb{E}(y_i | \mathbf{X}) \\ \vdots \end{pmatrix} = \begin{pmatrix} \vdots \\ \mathbf{x}_i' \boldsymbol{\beta} \\ \vdots \end{pmatrix} = \mathbf{X} \boldsymbol{\beta}. \quad (4.5)$$

Similarly

$$\mathbb{E}(\mathbf{e} | \mathbf{X}) = \begin{pmatrix} \vdots \\ \mathbb{E}(e_i | \mathbf{X}) \\ \vdots \end{pmatrix} = \begin{pmatrix} \vdots \\ \mathbb{E}(e_i | \mathbf{x}_i) \\ \vdots \end{pmatrix} = \mathbf{0}. \quad (4.6)$$

Using definition (3.18), the conditioning theorem, the linearity of expectations, (4.5), and the properties of the matrix inverse,

$$\begin{aligned}\mathbb{E}(\widehat{\boldsymbol{\beta}} | \mathbf{X}) &= \mathbb{E}\left((\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{y} | \mathbf{X}\right) \\ &= (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbb{E}(\mathbf{y} | \mathbf{X}) \\ &= (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{X}\boldsymbol{\beta} \\ &= \boldsymbol{\beta}.\end{aligned}$$

At the risk of belaboring the derivation, another way to calculate the same result is as follows. Insert $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}$ into the formula (3.18) for $\widehat{\boldsymbol{\beta}}$ to obtain

$$\begin{aligned}\widehat{\boldsymbol{\beta}} &= (\mathbf{X}'\mathbf{X})^{-1} (\mathbf{X}'(\mathbf{X}\boldsymbol{\beta} + \mathbf{e})) \\ &= (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{X}\boldsymbol{\beta} + (\mathbf{X}'\mathbf{X})^{-1} (\mathbf{X}'\mathbf{e}) \\ &= \boldsymbol{\beta} + (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{e}.\end{aligned}\tag{4.7}$$

This is a useful linear decomposition of the estimator $\widehat{\boldsymbol{\beta}}$ into the true parameter $\boldsymbol{\beta}$ and the stochastic component $(\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{e}$. Once again, we can calculate that

$$\begin{aligned}\mathbb{E}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta} | \mathbf{X}) &= \mathbb{E}\left((\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{e} | \mathbf{X}\right) \\ &= (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbb{E}(\mathbf{e} | \mathbf{X}) \\ &= \mathbf{0}.\end{aligned}$$

Regardless of the method, we have shown that $\mathbb{E}(\widehat{\boldsymbol{\beta}} | \mathbf{X}) = \boldsymbol{\beta}$. Applying the law of iterated expectations, we find that

$$\mathbb{E}(\widehat{\boldsymbol{\beta}}) = \mathbb{E}\left(\mathbb{E}(\widehat{\boldsymbol{\beta}} | \mathbf{X})\right) = \boldsymbol{\beta}.$$

We have shown the following theorem.

Theorem 4.4.1 Mean of Least-Squares Estimator

In the linear regression model (Assumption 4.3.1)

$$\mathbb{E}(\widehat{\boldsymbol{\beta}} | \mathbf{X}) = \boldsymbol{\beta}\tag{4.8}$$

and

$$\mathbb{E}(\widehat{\boldsymbol{\beta}}) = \boldsymbol{\beta}.\tag{4.9}$$

Equation (4.9) says that the estimator $\widehat{\boldsymbol{\beta}}$ is unbiased for $\boldsymbol{\beta}$, meaning that the distribution of $\widehat{\boldsymbol{\beta}}$ is centered at $\boldsymbol{\beta}$. Equation (4.8) says that the estimator is conditionally unbiased, which is a stronger result. It says that $\widehat{\boldsymbol{\beta}}$ is unbiased for any realization of the regressor matrix \mathbf{X} .

4.5 Variance of Least Squares Estimator

In this section we calculate the conditional variance of the OLS estimator.

For any $r \times 1$ random vector \mathbf{Z} define the $r \times r$ covariance matrix

$$\begin{aligned}\text{var}(\mathbf{Z}) &= \mathbb{E}(\mathbf{Z} - \mathbb{E}\mathbf{Z})(\mathbf{Z} - \mathbb{E}\mathbf{Z})' \\ &= \mathbb{E}\mathbf{Z}\mathbf{Z}' - (\mathbb{E}\mathbf{Z})(\mathbb{E}\mathbf{Z})'\end{aligned}$$

and for any pair (\mathbf{Z}, \mathbf{X}) define the conditional covariance matrix

$$\text{var}(\mathbf{Z} | \mathbf{X}) = \mathbb{E}((\mathbf{Z} - \mathbb{E}(\mathbf{Z} | \mathbf{X}))(\mathbf{Z} - \mathbb{E}(\mathbf{Z} | \mathbf{X}))' | \mathbf{X}).$$

We define

$$\mathbf{V}_{\hat{\boldsymbol{\beta}}} \stackrel{\text{def}}{=} \text{var}(\hat{\boldsymbol{\beta}} | \mathbf{X})$$

the conditional covariance matrix of the regression coefficient estimates. We now derive its form.

The conditional covariance matrix of the $n \times 1$ regression error \mathbf{e} is the $n \times n$ matrix

$$\mathbf{D} = \mathbb{E}(\mathbf{e}\mathbf{e}' | \mathbf{X}).$$

The i 'th diagonal element of \mathbf{D} is

$$\mathbb{E}(e_i^2 | \mathbf{X}) = \mathbb{E}(e_i^2 | \mathbf{x}_i) = \sigma_i^2$$

while the ij 'th off-diagonal element of \mathbf{D} is

$$\mathbb{E}(e_i e_j | \mathbf{X}) = \mathbb{E}(e_i | \mathbf{x}_i) \mathbb{E}(e_j | \mathbf{x}_j) = 0.$$

where the first equality uses independence of the observations (Assumption 1.5.1) and the second is (4.2). Thus \mathbf{D} is a diagonal matrix with i 'th diagonal element σ_i^2 :

$$\mathbf{D} = \text{diag}(\sigma_1^2, \dots, \sigma_n^2) = \begin{pmatrix} \sigma_1^2 & 0 & \cdots & 0 \\ 0 & \sigma_2^2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \sigma_n^2 \end{pmatrix}. \quad (4.10)$$

In the special case of the linear homoskedastic regression model (4.3), then

$$\mathbb{E}(e_i^2 | \mathbf{x}_i) = \sigma_i^2 = \sigma^2$$

and we have the simplification

$$\mathbf{D} = \mathbf{I}_n \sigma^2.$$

In general, however, \mathbf{D} need not necessarily take this simplified form.

For any $n \times r$ matrix $\mathbf{A} = \mathbf{A}(\mathbf{X})$,

$$\text{var}(\mathbf{A}'\mathbf{y} | \mathbf{X}) = \text{var}(\mathbf{A}'\mathbf{e} | \mathbf{X}) = \mathbf{A}'\mathbf{D}\mathbf{A}. \quad (4.11)$$

In particular, we can write $\hat{\boldsymbol{\beta}} = \mathbf{A}'\mathbf{y}$ where $\mathbf{A} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}$ and thus

$$\mathbf{V}_{\hat{\boldsymbol{\beta}}} = \text{var}(\hat{\boldsymbol{\beta}} | \mathbf{X}) = \mathbf{A}'\mathbf{D}\mathbf{A} = (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{D}\mathbf{X} (\mathbf{X}'\mathbf{X})^{-1}.$$

It is useful to note that

$$\mathbf{X}'\mathbf{D}\mathbf{X} = \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \sigma_i^2,$$

a weighted version of $\mathbf{X}'\mathbf{X}$.

In the special case of the linear homoskedastic regression model, $\mathbf{D} = \mathbf{I}_n \sigma^2$, so $\mathbf{X}'\mathbf{D}\mathbf{X} = \mathbf{X}'\mathbf{X}\sigma^2$, and the variance matrix simplifies to

$$\mathbf{V}_{\hat{\boldsymbol{\beta}}} = (\mathbf{X}'\mathbf{X})^{-1} \sigma^2.$$

Theorem 4.5.1 Variance of Least-Squares Estimator

In the linear regression model (Assumption 4.3.1)

$$\begin{aligned} \mathbf{V}_{\hat{\beta}} &= \text{var}(\hat{\beta} | \mathbf{X}) \\ &= (\mathbf{X}'\mathbf{X})^{-1} (\mathbf{X}'\mathbf{D}\mathbf{X}) (\mathbf{X}'\mathbf{X})^{-1} \end{aligned} \quad (4.12)$$

where \mathbf{D} is defined in (4.10).

In the homoskedastic linear regression model (Assumption 4.3.2)

$$\mathbf{V}_{\hat{\beta}} = (\mathbf{X}'\mathbf{X})^{-1} \sigma^2.$$

4.6 Gauss-Markov Theorem

Now consider the class of estimators of β which are linear functions of the vector \mathbf{y} , and thus can be written as

$$\tilde{\beta} = \mathbf{A}'\mathbf{y}$$

where \mathbf{A} is an $n \times k$ function of \mathbf{X} . As noted before, the least-squares estimator is the special case obtained by setting $\mathbf{A} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}$. What is the best choice of \mathbf{A} ? The Gauss-Markov theorem, which we now present, says that the least-squares estimator is the best choice among linear unbiased estimators when the errors are homoskedastic, in the sense that the least-squares estimator has the smallest variance among all unbiased linear estimators.

To see this, since $\mathbb{E}(\mathbf{y} | \mathbf{X}) = \mathbf{X}\beta$, then for any linear estimator $\tilde{\beta} = \mathbf{A}'\mathbf{y}$ we have

$$\mathbb{E}(\tilde{\beta} | \mathbf{X}) = \mathbf{A}'\mathbb{E}(\mathbf{y} | \mathbf{X}) = \mathbf{A}'\mathbf{X}\beta,$$

so $\tilde{\beta}$ is unbiased if (and only if) $\mathbf{A}'\mathbf{X} = \mathbf{I}_k$. Furthermore, we saw in (4.11) that

$$\text{var}(\tilde{\beta} | \mathbf{X}) = \text{var}(\mathbf{A}'\mathbf{y} | \mathbf{X}) = \mathbf{A}'\mathbf{D}\mathbf{A} = \mathbf{A}'\mathbf{A}\sigma^2$$

the last equality using the homoskedasticity assumption $\mathbf{D} = \mathbf{I}_n\sigma^2$. The “best” unbiased linear estimator is obtained by finding the matrix \mathbf{A}_0 satisfying $\mathbf{A}_0'\mathbf{X} = \mathbf{I}_k$ such that $\mathbf{A}_0'\mathbf{A}_0$ is minimized in the positive definite sense, in that for any other matrix \mathbf{A} satisfying $\mathbf{A}'\mathbf{X} = \mathbf{I}_k$, then $\mathbf{A}'\mathbf{A} - \mathbf{A}_0'\mathbf{A}_0$ is positive semi-definite.

Theorem 4.6.1 Gauss-Markov

1. In the homoskedastic linear regression model (Assumption 4.3.2), the best (minimum-variance) unbiased linear estimator is the least-squares estimator

$$\hat{\beta} = (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{y}$$

2. In the linear regression model (Assumption 4.3.1), the best unbiased linear estimator is

$$\tilde{\beta} = (\mathbf{X}'\mathbf{D}^{-1}\mathbf{X})^{-1} \mathbf{X}'\mathbf{D}^{-1}\mathbf{y} \quad (4.13)$$

The first part of the Gauss-Markov theorem is a limited efficiency justification for the least-squares estimator. The justification is limited because the class of models is restricted to homoskedastic linear regression and the class of potential estimators is restricted to linear unbiased estimators. This latter restriction is particularly unsatisfactory as the theorem leaves open the possibility that a non-linear or biased estimator could have lower mean squared error than the least-squares estimator.

The second part of the theorem shows that in the (heteroskedastic) linear regression model, within the class of linear unbiased estimators the best estimator is not least-squares but is (4.13). This is called the **Generalized Least Squares** (GLS) estimator. The GLS estimator is infeasible as the matrix \mathbf{D} is unknown. This result does not suggest a practical alternative to least-squares. We return to the issue of feasible implementation of GLS in Section 9.2.

We give a proof of the first part of the theorem below, and leave the proof of the second part for Exercise 4.3.

Proof of Theorem 4.6.1.1. Let \mathbf{A} be any $n \times k$ function of \mathbf{X} such that $\mathbf{A}'\mathbf{X} = \mathbf{I}_k$. The variance of the least-squares estimator is $(\mathbf{X}'\mathbf{X})^{-1}\sigma^2$ and that of $\mathbf{A}'\mathbf{y}$ is $\mathbf{A}'\mathbf{A}\sigma^2$. It is sufficient to show that the difference $\mathbf{A}'\mathbf{A} - (\mathbf{X}'\mathbf{X})^{-1}$ is positive semi-definite. Set $\mathbf{C} = \mathbf{A} - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}$. Note that $\mathbf{X}'\mathbf{C} = \mathbf{0}$. Then we calculate that

$$\begin{aligned} \mathbf{A}'\mathbf{A} - (\mathbf{X}'\mathbf{X})^{-1} &= \left(\mathbf{C} + \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\right)' \left(\mathbf{C} + \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\right) - (\mathbf{X}'\mathbf{X})^{-1} \\ &= \mathbf{C}'\mathbf{C} + \mathbf{C}'\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} + (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{C} \\ &\quad + (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1} - (\mathbf{X}'\mathbf{X})^{-1} \\ &= \mathbf{C}'\mathbf{C}. \end{aligned}$$

The matrix $\mathbf{C}'\mathbf{C}$ is positive semi-definite (see Appendix A.8) as required.

4.7 Residuals

What are some properties of the residuals $\hat{e}_i = y_i - \mathbf{x}'_i\hat{\boldsymbol{\beta}}$ and prediction errors $\tilde{e}_i = y_i - \mathbf{x}'_i\hat{\boldsymbol{\beta}}_{(-i)}$, at least in the context of the linear regression model?

Recall from (3.26) that we can write the residuals in vector notation as

$$\hat{\mathbf{e}} = \mathbf{M}\mathbf{e}$$

where $\mathbf{M} = \mathbf{I}_n - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$ is the orthogonal projection matrix. Using the properties of conditional expectation

$$\mathbb{E}(\hat{\mathbf{e}} | \mathbf{X}) = \mathbb{E}(\mathbf{M}\mathbf{e} | \mathbf{X}) = \mathbf{M}\mathbb{E}(\mathbf{e} | \mathbf{X}) = \mathbf{0}$$

and

$$\text{var}(\hat{\mathbf{e}} | \mathbf{X}) = \text{var}(\mathbf{M}\mathbf{e} | \mathbf{X}) = \mathbf{M} \text{var}(\mathbf{e} | \mathbf{X}) \mathbf{M} = \mathbf{M}\mathbf{D}\mathbf{M} \quad (4.14)$$

where \mathbf{D} is defined in (4.10).

We can simplify this expression under the assumption of conditional homoskedasticity

$$\mathbb{E}(e_i^2 | \mathbf{x}_i) = \sigma^2.$$

In this case (4.14) simplifies to

$$\text{var}(\hat{\mathbf{e}} | \mathbf{X}) = \mathbf{M}\sigma^2. \quad (4.15)$$

In particular, for a single observation i , we can find the (conditional) variance of \hat{e}_i by taking the i^{th} diagonal element of (4.16). Since the i^{th} diagonal element of \mathbf{M} is $1 - h_{ii}$ as defined in (3.21) we obtain

$$\text{var}(\hat{e}_i | \mathbf{X}) = \mathbb{E}(\hat{e}_i^2 | \mathbf{X}) = (1 - h_{ii})\sigma^2. \quad (4.16)$$

As this variance is a function of h_{ii} and hence \mathbf{x}_i , the residuals \hat{e}_i are heteroskedastic even if the errors e_i are homoskedastic.

Similarly, recall from (3.40) that the prediction errors $\tilde{e}_i = (1 - h_{ii})^{-1}\hat{e}_i$ can be written in vector notation as $\tilde{\mathbf{e}} = \mathbf{M}^*\hat{\mathbf{e}}$ where \mathbf{M}^* is a diagonal matrix with i^{th} diagonal element $(1 - h_{ii})^{-1}$. Thus $\tilde{\mathbf{e}} = \mathbf{M}^*\mathbf{M}\mathbf{e}$. We can calculate that

$$\mathbb{E}(\tilde{\mathbf{e}} | \mathbf{X}) = \mathbf{M}^*\mathbf{M}\mathbb{E}(\mathbf{e} | \mathbf{X}) = \mathbf{0}$$

and

$$\text{var}(\tilde{\mathbf{e}} | \mathbf{X}) = \mathbf{M}^*\mathbf{M}\text{var}(\mathbf{e} | \mathbf{X})\mathbf{M}\mathbf{M}^* = \mathbf{M}^*\mathbf{M}\mathbf{D}\mathbf{M}\mathbf{M}^*$$

which simplifies under homoskedasticity to

$$\begin{aligned} \text{var}(\tilde{\mathbf{e}} | \mathbf{X}) &= \mathbf{M}^*\mathbf{M}\mathbf{M}\mathbf{M}^*\sigma^2 \\ &= \mathbf{M}^*\mathbf{M}\mathbf{M}^*\sigma^2. \end{aligned}$$

The variance of the i^{th} prediction error is then

$$\begin{aligned} \text{var}(\tilde{e}_i | \mathbf{X}) &= \mathbb{E}(\tilde{e}_i^2 | \mathbf{X}) \\ &= (1 - h_{ii})^{-1}(1 - h_{ii})(1 - h_{ii})^{-1}\sigma^2 \\ &= (1 - h_{ii})^{-1}\sigma^2. \end{aligned}$$

A residual with constant conditional variance can be obtained by rescaling. The **standardized residuals** are

$$\bar{e}_i = (1 - h_{ii})^{-1/2}\hat{e}_i, \quad (4.17)$$

and in vector notation

$$\bar{\mathbf{e}} = (\bar{e}_1, \dots, \bar{e}_n)' = \mathbf{M}^{*1/2}\mathbf{M}\mathbf{e}.$$

From our above calculations, under homoskedasticity,

$$\text{var}(\bar{\mathbf{e}} | \mathbf{X}) = \mathbf{M}^{*1/2}\mathbf{M}\mathbf{M}^{*1/2}\sigma^2$$

and

$$\text{var}(\bar{e}_i | \mathbf{X}) = \mathbb{E}(\bar{e}_i^2 | \mathbf{X}) = \sigma^2 \quad (4.18)$$

and thus these standardized residuals have the same bias and variance as the original errors when the latter are homoskedastic.

4.8 Estimation of Error Variance

The error variance $\sigma^2 = \mathbb{E}e_i^2$ can be a parameter of interest, even in a heteroskedastic regression or a projection model. σ^2 measures the variation in the “unexplained” part of the regression. Its method of moments estimator (MME) is the sample average of the squared residuals:

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n \hat{e}_i^2$$

and equals the MLE in the normal regression model (3.28).

In the linear regression model we can calculate the mean of $\hat{\sigma}^2$. From (3.26), the properties of projection matrices and the trace operator, observe that

$$\hat{\sigma}^2 = \frac{1}{n} \hat{\mathbf{e}}' \hat{\mathbf{e}} = \frac{1}{n} \mathbf{e}' \mathbf{M} \mathbf{M} \mathbf{e} = \frac{1}{n} \mathbf{e}' \mathbf{M} \mathbf{e} = \frac{1}{n} \text{tr}(\mathbf{e}' \mathbf{M} \mathbf{e}) = \frac{1}{n} \text{tr}(\mathbf{M} \mathbf{e} \mathbf{e}').$$

Then

$$\begin{aligned} \mathbb{E}(\hat{\sigma}^2 | \mathbf{X}) &= \frac{1}{n} \text{tr}(\mathbb{E}(\mathbf{M} \mathbf{e} \mathbf{e}' | \mathbf{X})) \\ &= \frac{1}{n} \text{tr}(\mathbf{M} \mathbb{E}(\mathbf{e} \mathbf{e}' | \mathbf{X})) \\ &= \frac{1}{n} \text{tr}(\mathbf{M} \mathbf{D}). \end{aligned} \tag{4.19}$$

Adding the assumption of conditional homoskedasticity $\mathbb{E}(e_i^2 | \mathbf{x}_i) = \sigma^2$, so that $\mathbf{D} = \mathbf{I}_n \sigma^2$, then (4.19) simplifies to

$$\begin{aligned} \mathbb{E}(\hat{\sigma}^2 | \mathbf{X}) &= \frac{1}{n} \text{tr}(\mathbf{M} \sigma^2) \\ &= \sigma^2 \left(\frac{n-k}{n} \right), \end{aligned}$$

the final equality by (3.24). This calculation shows that $\hat{\sigma}^2$ is biased towards zero. The order of the bias depends on k/n , the ratio of the number of estimated coefficients to the sample size.

Another way to see this is to use (4.16). Note that

$$\begin{aligned} \mathbb{E}(\hat{\sigma}^2 | \mathbf{X}) &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}(\hat{e}_i^2 | \mathbf{X}) \\ &= \frac{1}{n} \sum_{i=1}^n (1 - h_{ii}) \sigma^2 \\ &= \left(\frac{n-k}{n} \right) \sigma^2 \end{aligned} \tag{4.20}$$

the last equality using Theorem 3.10.1.

Since the bias takes a scale form, a classic method to obtain an unbiased estimator is by rescaling the estimator. Define

$$s^2 = \frac{1}{n-k} \sum_{i=1}^n \hat{e}_i^2. \tag{4.21}$$

By the above calculation,

$$\mathbb{E}(s^2 | \mathbf{X}) = \sigma^2 \tag{4.22}$$

so

$$\mathbb{E}(s^2) = \sigma^2$$

and the estimator s^2 is unbiased for σ^2 . Consequently, s^2 is known as the “bias-corrected estimator” for σ^2 and in empirical practice s^2 is the most widely used estimator for σ^2 .

Interestingly, this is not the only method to construct an unbiased estimator for σ^2 . An estimator constructed with the standardized residuals \bar{e}_i from (4.17) is

$$\bar{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n \bar{e}_i^2 = \frac{1}{n} \sum_{i=1}^n (1 - h_{ii})^{-1} \hat{e}_i^2. \tag{4.23}$$

You can show (see Exercise 4.6) that

$$\mathbb{E}(\bar{\sigma}^2 | \mathbf{X}) = \sigma^2 \quad (4.24)$$

and thus $\bar{\sigma}^2$ is unbiased for σ^2 (in the homoskedastic linear regression model).

When k/n is small (typically, this occurs when n is large), the estimators $\hat{\sigma}^2$, s^2 and $\bar{\sigma}^2$ are likely to be close. However, if not then s^2 and $\bar{\sigma}^2$ are generally preferred to $\hat{\sigma}^2$. Consequently it is best to use one of the bias-corrected variance estimators in applications.

4.9 Mean-Square Forecast Error

A major purpose of estimated regressions is to predict out-of-sample values. Consider an out-of-sample observation $(y_{n+1}, \mathbf{x}_{n+1})$ where \mathbf{x}_{n+1} will be observed but not y_{n+1} . Given the coefficient estimate $\hat{\boldsymbol{\beta}}$ the standard point estimate of $\mathbb{E}(y_{n+1} | \mathbf{x}_{n+1}) = \mathbf{x}'_{n+1}\boldsymbol{\beta}$ is $\tilde{y}_{n+1} = \mathbf{x}'_{n+1}\hat{\boldsymbol{\beta}}$. The forecast error is the difference between the actual value y_{n+1} and the point forecast, $\tilde{e}_{n+1} = y_{n+1} - \tilde{y}_{n+1}$. The mean-squared forecast error (MSFE) is

$$MSFE_n = \mathbb{E}\tilde{e}_{n+1}^2.$$

In the linear regression model, $\tilde{e}_{n+1} = e_{n+1} - \mathbf{x}'_{n+1}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})$, so

$$\begin{aligned} MSFE_n &= \mathbb{E}e_{n+1}^2 - 2\mathbb{E}\left(e_{n+1}\mathbf{x}'_{n+1}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})\right) \\ &\quad + \mathbb{E}\left(\mathbf{x}'_{n+1}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})'\mathbf{x}_{n+1}\right). \end{aligned} \quad (4.25)$$

The first term in (4.25) is σ^2 . The second term in (4.25) is zero since $e_{n+1}\mathbf{x}'_{n+1}$ is independent of $\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}$ and both are mean zero. The third term in (4.25) is

$$\begin{aligned} &\text{tr}\left(\mathbb{E}(\mathbf{x}_{n+1}\mathbf{x}'_{n+1})\mathbb{E}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})'\right) \\ &= \text{tr}\left(\mathbb{E}(\mathbf{x}_{n+1}\mathbf{x}'_{n+1})\mathbb{E}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})'\right) \\ &= \text{tr}\left(\mathbb{E}(\mathbf{x}_{n+1}\mathbf{x}'_{n+1})\mathbb{E}\mathbf{V}_{\hat{\boldsymbol{\beta}}}\right) \\ &= \mathbb{E}\text{tr}\left((\mathbf{x}_{n+1}\mathbf{x}'_{n+1})\mathbf{V}_{\hat{\boldsymbol{\beta}}}\right) \\ &= \mathbb{E}\left(\mathbf{x}'_{n+1}\mathbf{V}_{\hat{\boldsymbol{\beta}}}\mathbf{x}_{n+1}\right) \end{aligned} \quad (4.26)$$

where we use the fact that \mathbf{x}_{n+1} is independent of $\hat{\boldsymbol{\beta}}$ and use the fact $\mathbf{V}_{\hat{\boldsymbol{\beta}}} = \mathbb{E}\left((\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})' | \mathbf{X}\right)$.

Thus

$$MSFE_n = \sigma^2 + \mathbb{E}\left(\mathbf{x}'_{n+1}\mathbf{V}_{\hat{\boldsymbol{\beta}}}\mathbf{x}_{n+1}\right).$$

Under conditional homoskedasticity, this simplifies to

$$MSFE_n = \sigma^2 \left(1 + \mathbb{E}\left(\mathbf{x}'_{n+1}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{x}_{n+1}\right)\right).$$

A simple estimator for the MSFE is obtained by averaging the squared prediction errors (3.41)

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n \hat{e}_i^2$$

where $\tilde{e}_i = y_i - \mathbf{x}'_i \widehat{\boldsymbol{\beta}}_{(-i)} = \hat{e}_i(1 - h_{ii})^{-1}$. Indeed, we can calculate that

$$\begin{aligned} \mathbb{E}\tilde{\sigma}^2 &= \mathbb{E}\tilde{e}_i^2 \\ &= \mathbb{E}\left(e_i - \mathbf{x}'_i \left(\widehat{\boldsymbol{\beta}}_{(-i)} - \boldsymbol{\beta}\right)\right)^2 \\ &= \sigma^2 + \mathbb{E}\left(\mathbf{x}'_i \left(\widehat{\boldsymbol{\beta}}_{(-i)} - \boldsymbol{\beta}\right) \left(\widehat{\boldsymbol{\beta}}_{(-i)} - \boldsymbol{\beta}\right)' \mathbf{x}_i\right). \end{aligned}$$

By the same calculations as in (4.26) we find

$$\mathbb{E}\tilde{\sigma}^2 = \sigma^2 + \mathbb{E}\left(\mathbf{x}'_i \mathbf{V}_{\widehat{\boldsymbol{\beta}}_{(-i)}} \mathbf{x}_i\right) = MSFE_{n-1}.$$

This is the MSFE based on a sample of size $n - 1$, rather than size n . The difference arises because the in-sample prediction errors \tilde{e}_i for $i \leq n$ are calculated using an effective sample size of $n - 1$, while the out-of sample prediction error \tilde{e}_{n+1} is calculated from a sample with the full n observations. Unless n is very small we should expect $MSFE_{n-1}$ (the MSFE based on $n - 1$ observations) to be close to $MSFE_n$ (the MSFE based on n observations). Thus $\tilde{\sigma}^2$ is a reasonable estimator for $MSFE_n$.

Theorem 4.9.1 MSFE

In the linear regression model (Assumption 4.3.1)

$$MSFE_n = \mathbb{E}\tilde{e}_{n+1}^2 = \sigma^2 + \mathbb{E}\left(\mathbf{x}'_{n+1} \mathbf{V}_{\widehat{\boldsymbol{\beta}}} \mathbf{x}_{n+1}\right)$$

where $\mathbf{V}_{\widehat{\boldsymbol{\beta}}} = \text{var}\left(\widehat{\boldsymbol{\beta}} \mid \mathbf{X}\right)$. Furthermore, $\tilde{\sigma}^2$ defined in (3.41) is an unbiased estimator of $MSFE_{n-1}$:

$$\mathbb{E}\tilde{\sigma}^2 = MSFE_{n-1}$$

4.10 Covariance Matrix Estimation Under Homoskedasticity

For inference, we need an estimate of the covariance matrix $\mathbf{V}_{\widehat{\boldsymbol{\beta}}}$ of the least-squares estimator. In this section we consider the homoskedastic regression model (Assumption 4.3.2).

Under homoskedasticity, the covariance matrix takes the relatively simple form

$$\mathbf{V}_{\widehat{\boldsymbol{\beta}}} = (\mathbf{X}'\mathbf{X})^{-1} \sigma^2$$

which is known up to the unknown scale σ^2 . In Section 4.8 we discussed three estimators of σ^2 . The most commonly used choice is s^2 , leading to the classic covariance matrix estimator

$$\widehat{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}}^0 = (\mathbf{X}'\mathbf{X})^{-1} s^2. \quad (4.27)$$

Since s^2 is conditionally unbiased for σ^2 , it is simple to calculate that $\widehat{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}}^0$ is conditionally unbiased for $\mathbf{V}_{\widehat{\boldsymbol{\beta}}}$ under the assumption of homoskedasticity:

$$\begin{aligned} \mathbb{E}\left(\widehat{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}}^0 \mid \mathbf{X}\right) &= (\mathbf{X}'\mathbf{X})^{-1} \mathbb{E}\left(s^2 \mid \mathbf{X}\right) \\ &= (\mathbf{X}'\mathbf{X})^{-1} \sigma^2 \\ &= \mathbf{V}_{\widehat{\boldsymbol{\beta}}}. \end{aligned}$$

This estimator was the dominant covariance matrix estimator in applied econometrics for many years, and is still the default method in most regression packages.

If the estimator (4.27) is used, but the regression error is heteroskedastic, it is possible for $\widehat{\mathbf{V}}_{\hat{\beta}}^0$ to be quite biased for the correct covariance matrix $\mathbf{V}_{\hat{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}(\mathbf{X}'\mathbf{D}\mathbf{X})(\mathbf{X}'\mathbf{X})^{-1}$. For example, suppose $k = 1$ and $\sigma_i^2 = x_i^2$ with $\mathbb{E}x_i = 0$. The ratio of the true variance of the least-squares estimator to the expectation of the variance estimator is

$$\frac{\mathbf{V}_{\hat{\beta}}}{\mathbb{E}\left(\widehat{\mathbf{V}}_{\hat{\beta}}^0 \mid \mathbf{X}\right)} = \frac{\sum_{i=1}^n x_i^4}{\sigma^2 \sum_{i=1}^n x_i^2} \simeq \frac{\mathbb{E}x_i^4}{(\mathbb{E}x_i^2)^2} = \kappa.$$

(Notice that we use the fact that $\sigma_i^2 = x_i^2$ implies $\sigma^2 = \mathbb{E}\sigma_i^2 = \mathbb{E}x_i^2$.) The constant κ is the standardized fourth moment (or kurtosis) of the regressor x_i , and can be any number greater than one. For example, if $x_i \sim N(0, \sigma^2)$ then $\kappa = 3$, so the true variance $\mathbf{V}_{\hat{\beta}}$ is three times larger than the expected homoskedastic estimator $\widehat{\mathbf{V}}_{\hat{\beta}}^0$. But κ can be much larger. Suppose, for example, that $x_i \sim \chi_1^2 - 1$. In this case $\kappa = 15$, so that the true variance $\mathbf{V}_{\hat{\beta}}$ is fifteen times larger than the expected homoskedastic estimator $\widehat{\mathbf{V}}_{\hat{\beta}}^0$. While this is an extreme and constructed example, the point is that the classic covariance matrix estimator (4.27) may be quite biased when the homoskedasticity assumption fails.

4.11 Covariance Matrix Estimation Under Heteroskedasticity

In the previous section we showed that the classic covariance matrix estimator can be highly biased if homoskedasticity fails. In this section we show how to construct covariance matrix estimators which do not require homoskedasticity.

Recall that the general form for the covariance matrix is

$$\mathbf{V}_{\hat{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}(\mathbf{X}'\mathbf{D}\mathbf{X})(\mathbf{X}'\mathbf{X})^{-1}.$$

This depends on the unknown matrix \mathbf{D} which we can write as

$$\begin{aligned} \mathbf{D} &= \text{diag}(\sigma_1^2, \dots, \sigma_n^2) \\ &= \mathbb{E}(\mathbf{e}\mathbf{e}' \mid \mathbf{X}) \\ &= \mathbb{E}(\mathbf{D}_0 \mid \mathbf{X}) \end{aligned}$$

where $\mathbf{D}_0 = \text{diag}(e_1^2, \dots, e_n^2)$. Thus \mathbf{D}_0 is a conditionally unbiased estimator for \mathbf{D} . If the squared errors e_i^2 were observable, we could construct the unbiased estimator

$$\begin{aligned} \widehat{\mathbf{V}}_{\hat{\beta}}^{ideal} &= (\mathbf{X}'\mathbf{X})^{-1}(\mathbf{X}'\mathbf{D}_0\mathbf{X})(\mathbf{X}'\mathbf{X})^{-1} \\ &= (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' e_i^2 \right) (\mathbf{X}'\mathbf{X})^{-1}. \end{aligned}$$

Indeed,

$$\begin{aligned} \mathbb{E}\left(\widehat{\mathbf{V}}_{\hat{\beta}}^{ideal} \mid \mathbf{X}\right) &= (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \mathbb{E}(e_i^2 \mid \mathbf{X}) \right) (\mathbf{X}'\mathbf{X})^{-1} \\ &= (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \sigma_i^2 \right) (\mathbf{X}'\mathbf{X})^{-1} \\ &= (\mathbf{X}'\mathbf{X})^{-1}(\mathbf{X}'\mathbf{D}\mathbf{X})(\mathbf{X}'\mathbf{X})^{-1} \\ &= \mathbf{V}_{\hat{\beta}} \end{aligned}$$

verifying that $\widehat{\mathbf{V}}_{\widehat{\beta}}^{ideal}$ is unbiased for $\mathbf{V}_{\widehat{\beta}}$

Since the errors e_i^2 are unobserved, $\widehat{\mathbf{V}}_{\widehat{\beta}}^{ideal}$ is not a feasible estimator. However, we can replace the errors e_i with the least-squares residuals \widehat{e}_i . Making this substitution we obtain the estimator

$$\widehat{\mathbf{V}}_{\widehat{\beta}}^W = (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \widehat{e}_i^2 \right) (\mathbf{X}'\mathbf{X})^{-1}. \quad (4.28)$$

We know, however, that \widehat{e}_i^2 is biased towards zero. To estimate the variance σ^2 the unbiased estimator s^2 scales the moment estimator $\widehat{\sigma}^2$ by $n/(n-k)$. Making the same adjustment we obtain the estimator

$$\widehat{\mathbf{V}}_{\widehat{\beta}} = \left(\frac{n}{n-k} \right) (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \widehat{e}_i^2 \right) (\mathbf{X}'\mathbf{X})^{-1}. \quad (4.29)$$

While the scaling by $n/(n-k)$ is ad hoc, it is recommended over the unscaled estimator (4.28).

Alternatively, we could use the prediction errors \widetilde{e}_i or the standardized residuals \bar{e}_i , yielding the estimators

$$\begin{aligned} \widetilde{\mathbf{V}}_{\widehat{\beta}} &= (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \widetilde{e}_i^2 \right) (\mathbf{X}'\mathbf{X})^{-1} \\ &= (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{i=1}^n (1 - h_{ii})^{-2} \mathbf{x}_i \mathbf{x}_i' \widehat{e}_i^2 \right) (\mathbf{X}'\mathbf{X})^{-1} \end{aligned} \quad (4.30)$$

and

$$\begin{aligned} \bar{\mathbf{V}}_{\widehat{\beta}} &= (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \bar{e}_i^2 \right) (\mathbf{X}'\mathbf{X})^{-1} \\ &= (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{i=1}^n (1 - h_{ii})^{-1} \mathbf{x}_i \mathbf{x}_i' \widehat{e}_i^2 \right) (\mathbf{X}'\mathbf{X})^{-1}. \end{aligned} \quad (4.31)$$

The four estimators $\widehat{\mathbf{V}}_{\widehat{\beta}}^W$, $\widehat{\mathbf{V}}_{\widehat{\beta}}$, $\widetilde{\mathbf{V}}_{\widehat{\beta}}$, and $\bar{\mathbf{V}}_{\widehat{\beta}}$ are collectively called **robust, heteroskedasticity-consistent**, or **heteroskedasticity-robust** covariance matrix estimators. The estimator $\widehat{\mathbf{V}}_{\widehat{\beta}}$ was first developed by Eicker (1963) and introduced to econometrics by White (1980), and is sometimes called the **Eicker-White** or **White** covariance matrix estimator. The scaled estimator $\widehat{\mathbf{V}}_{\widehat{\beta}}$ is the default robust covariance matrix estimator implemented in Stata. The estimator $\widetilde{\mathbf{V}}_{\widehat{\beta}}$ was introduced by Andrews (1991) based on the principle of leave-one-out cross-validation (and is implemented using the `vce(hc3)` option in Stata). The estimator $\bar{\mathbf{V}}_{\widehat{\beta}}$ was introduced by Horn, Horn and Duncan (1975) (and is implemented using the `vce(hc2)` option in Stata).

Since $(1 - h_{ii})^{-2} > (1 - h_{ii})^{-1} > 1$ it is straightforward to show that

$$\widehat{\mathbf{V}}_{\widehat{\beta}}^W < \bar{\mathbf{V}}_{\widehat{\beta}} < \widetilde{\mathbf{V}}_{\widehat{\beta}} \quad (4.32)$$

(See Exercise 4.7). The inequality $\mathbf{A} < \mathbf{B}$ when applied to matrices means that the matrix $\mathbf{B} - \mathbf{A}$ is positive definite.

In general, the bias of the covariance matrix estimators is quite complicated, but they greatly

simplify under the assumption of homoskedasticity (4.3). For example, using (4.16),

$$\begin{aligned}
 \mathbb{E}\left(\widehat{\mathbf{V}}_{\beta}^W \mid \mathbf{X}\right) &= (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \mathbb{E}(\hat{\epsilon}_i^2 \mid \mathbf{X}) \right) (\mathbf{X}'\mathbf{X})^{-1} \\
 &= (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' (1 - h_{ii}) \sigma^2 \right) (\mathbf{X}'\mathbf{X})^{-1} \\
 &= (\mathbf{X}'\mathbf{X})^{-1} \sigma^2 - (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' h_{ii} \right) (\mathbf{X}'\mathbf{X})^{-1} \sigma^2 \\
 &\leq (\mathbf{X}'\mathbf{X})^{-1} \sigma^2 \\
 &= \mathbf{V}_{\beta}.
 \end{aligned}$$

This calculation shows that $\widehat{\mathbf{V}}_{\beta}^W$ is biased towards zero.

Similarly, (again under homoskedasticity) we can calculate that $\widetilde{\mathbf{V}}_{\beta}$ is biased away from zero, specifically

$$\mathbb{E}\left(\widetilde{\mathbf{V}}_{\beta} \mid \mathbf{X}\right) \geq (\mathbf{X}'\mathbf{X})^{-1} \sigma^2 \quad (4.33)$$

while the estimator $\overline{\mathbf{V}}_{\beta}$ is unbiased

$$\mathbb{E}\left(\overline{\mathbf{V}}_{\beta} \mid \mathbf{X}\right) = (\mathbf{X}'\mathbf{X})^{-1} \sigma^2. \quad (4.34)$$

(See Exercise 4.8.)

It might seem rather odd to compare the bias of heteroskedasticity-robust estimators under the assumption of homoskedasticity, but it does give us a baseline for comparison.

We have introduced five covariance matrix estimators, $\widehat{\mathbf{V}}_{\beta}^0$, $\widehat{\mathbf{V}}_{\beta}^W$, $\widehat{\mathbf{V}}_{\beta}$, $\widetilde{\mathbf{V}}_{\beta}$, and $\overline{\mathbf{V}}_{\beta}$. Which should you use? The classic estimator $\widehat{\mathbf{V}}_{\beta}^0$ is typically a poor choice, as it is only valid under the unlikely homoskedasticity restriction. For this reason it is not typically used in contemporary econometric research. Unfortunately, standard regression packages set their default choice as $\widehat{\mathbf{V}}_{\beta}^0$, so users must intentionally select a robust covariance matrix estimator.

Of the four robust estimators, $\widehat{\mathbf{V}}_{\beta}^W$ and $\widehat{\mathbf{V}}_{\beta}$ are the most commonly used, and in particular $\widehat{\mathbf{V}}_{\beta}$ is the default robust covariance matrix option in Stata. However, $\widetilde{\mathbf{V}}_{\beta}$ and $\overline{\mathbf{V}}_{\beta}$ are preferred based on their improved bias. As $\widetilde{\mathbf{V}}_{\beta}$ and $\overline{\mathbf{V}}_{\beta}$ are simple to implement, this should not be a barrier.

Halbert L. White

Hal White (1950-2012) of the United States was an influential econometrician of recent years. His 1980 paper on heteroskedasticity-consistent covariance matrix estimation for many years has been the most cited paper in economics. His research was central to the movement to view econometric models as approximations, and to the drive for increased mathematical rigor in the discipline. In addition to being a highly prolific and influential scholar, he also co-founded the economic consulting firm Bates White.

4.12 Standard Errors

A variance estimator such as $\widehat{\mathbf{V}}_{\widehat{\beta}}$ is an estimate of the variance of the distribution of $\widehat{\beta}$. A more easily interpretable measure of spread is its square root – the standard deviation. This is so important when discussing the distribution of parameter estimates, we have a special name for estimates of their standard deviation.

Definition 4.12.1 A *standard error* $s(\widehat{\beta})$ for a real-valued estimator $\widehat{\beta}$ is an estimate of the standard deviation of the distribution of $\widehat{\beta}$.

When β is a vector with estimate $\widehat{\beta}$ and covariance matrix estimate $\widehat{\mathbf{V}}_{\widehat{\beta}}$, standard errors for individual elements are the square roots of the diagonal elements of $\widehat{\mathbf{V}}_{\widehat{\beta}}$. That is,

$$s(\widehat{\beta}_j) = \sqrt{\widehat{\mathbf{V}}_{\widehat{\beta}_j}} = \sqrt{[\widehat{\mathbf{V}}_{\widehat{\beta}}]_{jj}}.$$

As we discussed in the previous section, there are multiple possible covariance matrix estimators, so standard errors are not unique. It is therefore important to understand what formula and method is used by an author when studying their work. It is also important to understand that a particular standard error may be relevant under one set of model assumptions, but not under another set of assumptions.

To illustrate, we return to the log wage regression (3.11) of Section 3.7. We calculate that $s^2 = 0.160$. Therefore the homoskedastic covariance matrix estimate is

$$\widehat{\mathbf{V}}_{\widehat{\beta}}^0 = \begin{pmatrix} 5010 & 314 \\ 314 & 20 \end{pmatrix}^{-1} 0.215 = \begin{pmatrix} 0.002 & -0.031 \\ -0.031 & 0.499 \end{pmatrix}.$$

We also calculate that

$$\sum_{i=1}^n (1 - h_{ii})^{-1} \mathbf{x}_i \mathbf{x}_i' \widehat{e}_i^2 = \begin{pmatrix} 763.26 & 48.51 \\ 48.51 & 3.11 \end{pmatrix}.$$

Therefore the Horn-Horn-Duncan covariance matrix estimate is

$$\begin{aligned} \overline{\mathbf{V}}_{\widehat{\beta}} &= \begin{pmatrix} 5010 & 314 \\ 314 & 20 \end{pmatrix}^{-1} \begin{pmatrix} 763.26 & 48.51 \\ 48.51 & 3.11 \end{pmatrix} \begin{pmatrix} 5010 & 314 \\ 314 & 20 \end{pmatrix}^{-1} \\ &= \begin{pmatrix} 0.001 & -0.015 \\ -0.015 & 0.243 \end{pmatrix}. \end{aligned} \tag{4.35}$$

The standard errors are the square roots of the diagonal elements of these matrices. A conventional format to write the estimated equation with standard errors is

$$\widehat{\log(\text{Wage})} = \begin{matrix} 0.155 \\ (0.031) \end{matrix} \text{Education} + \begin{matrix} 0.698 \\ (0.493) \end{matrix}.$$

Alternatively, standard errors could be calculated using the other formulae. We report the different standard errors in the following table.

	Education	Intercept
Homoskedastic (4.27)	0.045	0.707
White (4.28)	0.029	0.461
Scaled White (4.29)	0.030	0.486
Andrews (4.30)	0.033	0.527
Horn-Horn-Duncan (4.31)	0.031	0.493

The homoskedastic standard errors are noticeably different (larger, in this case) than the others, but the four robust standard errors are quite close to one another.

4.13 Computation

We illustrate methods to compute standard errors for equation (3.12) extending the code of Section 3.20.

Stata do File (continued)

```
* Homoskedastic formula (4.27):
reg wage education experience exp2 if (mnwf == 1)
* Scaled White formula (4.29):
reg wage education experience exp2 if (mnwf == 1), r
* Andrews formula (4.30):
reg wage education experience exp2 if (mnwf == 1), vce(hc3)
* Horn-Horn-Duncan formula (4.31):
reg wage education experience exp2 if (mnwf == 1), vce(hc2)
```

Gauss Program File (continued)

```
n=rows(y);
k=cols(x);
a=n/(n-k);
sig2=(e'e)/(n-k);
u1=x.*e;
u2=x.*(e./(1-leverage));
u3=x.*(e./sqrt(1-leverage));
xx=inv(x'x);
v0=xx*sig2;
v1=xx*(u1'u1)*xx;
v1a=a*xx*(u1'u1)*xx;
v2=xx*(u2'u2)*xx;
v3=xx*(u3'u3)*xx
s0=sqrt(diag(v0));           @ Homoskedastic formula @
s1=sqrt(diag(v1));           @ White formula @
s1a=sqrt(diag(v1a));         @ Scaled White formula @
s2=sqrt(diag(v2));           @ Andrews formula @
s3=sqrt(diag(v3));           @ Horn-Horn-Duncan formula @
```

R Program File (continued)

```

n <- nrow(y)
k <- ncol(x)
a <- n/(n-k)
sig2 <- (t(e) %*% e)/(n-k)
u1 <- x*(e%*%matrix(1,1,k))
u2 <- x*((e/(1-leverage))%*%matrix(1,1,k))
u3 <- x*((e/sqrt(1-leverage))%*%matrix(1,1,k))
v0 <- xx*sig2
xx <- solve(t(x)%*%x)
v1 <- xx %*% (t(u1)%*%u1) %*% xx
v1a <- a * xx %*% (t(u1)%*%u1) %*% xx
v2 <- xx %*% (t(u2)%*%u2) %*% xx
v3 <- xx %*% (t(u3)%*%u3) %*% xx
s0 <- sqrt(diag(v0))      # Homoskedastic formula
s1 <- sqrt(diag(v1))      # White formula
s1a <- sqrt(diag(v1a))    # Scaled White formula
s2 <- sqrt(diag(v2))      # Andrews formula
s3 <- sqrt(diag(v3))      # Horn-Horn-Duncan formula

```

Matlab Program File (continued)

```

[n,k]=size(x);
a=n/(n-k);
sig2=(e'*e)/(n-k);
u1=x.*(e*ones(1,k));
u2=x.*((e./(1-leverage))*ones(1,k));
u3=x.*((e./sqrt(1-leverage))*ones(1,k));
xx=inv(x'*x);
v0=xx*sig2;
v1=xx*(u1'*u1)*xx;
v1a=a*xx*(u1'*u1)*xx;
v2=xx*(u2'*u2)*xx;
v3=xx*(u3'*u3)*xx;
s0=sqrt(diag(v0));      # Homoskedastic formula
s1=sqrt(diag(v1));      # White formula
s1a=sqrt(diag(v1a));    # Scaled White formula
s2=sqrt(diag(v2));      # Andrews formula
s3=sqrt(diag(v3));      # Horn-Horn-Duncan formula

```

4.14 Measures of Fit

As we described in the previous chapter, a commonly reported measure of regression fit is the regression R^2 defined as

$$R^2 = 1 - \frac{\sum_{i=1}^n \hat{e}_i^2}{\sum_{i=1}^n (y_i - \bar{y})^2} = 1 - \frac{\hat{\sigma}_y^2}{\hat{\sigma}_y^2}$$

where $\hat{\sigma}_y^2 = n^{-1} \sum_{i=1}^n (y_i - \bar{y})^2$. R^2 can be viewed as an estimator of the population parameter

$$\rho^2 = \frac{\text{var}(\mathbf{x}'_i \boldsymbol{\beta})}{\text{var}(y_i)} = 1 - \frac{\sigma^2}{\sigma_y^2}.$$

However, $\hat{\sigma}^2$ and $\hat{\sigma}_y^2$ are biased estimators. Theil (1961) proposed replacing these by the unbiased versions s^2 and $\tilde{\sigma}_y^2 = (n-1)^{-1} \sum_{i=1}^n (y_i - \bar{y})^2$ yielding what is known as **R-bar-squared** or **adjusted R-squared**:

$$\bar{R}^2 = 1 - \frac{s^2}{\tilde{\sigma}_y^2} = 1 - \frac{(n-1) \sum_{i=1}^n \hat{e}_i^2}{(n-k) \sum_{i=1}^n (y_i - \bar{y})^2}.$$

While \bar{R}^2 is an improvement on R^2 , a much better improvement is

$$\tilde{R}^2 = 1 - \frac{\sum_{i=1}^n \tilde{e}_i^2}{\sum_{i=1}^n (y_i - \bar{y})^2} = 1 - \frac{\tilde{\sigma}^2}{\tilde{\sigma}_y^2}$$

where \tilde{e}_i are the prediction errors (3.38) and $\tilde{\sigma}^2$ is the MSPE from (3.41). As described in Section (4.9), $\tilde{\sigma}^2$ is a good estimator of the out-of-sample mean-squared forecast error, so \tilde{R}^2 is a good estimator of the percentage of the forecast variance which is explained by the regression forecast. In this sense, \tilde{R}^2 is a good measure of fit.

One problem with R^2 , which is partially corrected by \bar{R}^2 and fully corrected by \tilde{R}^2 , is that R^2 necessarily increases when regressors are added to a regression model. This occurs because R^2 is a negative function of the sum of squared residuals which cannot increase when a regressor is added. In contrast, \bar{R}^2 and \tilde{R}^2 are non-monotonic in the number of regressors. \tilde{R}^2 can even be negative, which occurs when an estimated model predicts worse than a constant-only model.

In the statistical literature the MSPE $\tilde{\sigma}^2$ is known as the **leave-one-out cross validation** criterion, and is popular for model comparison and selection, especially in high-dimensional (non-parametric) contexts. It is equivalent to use \tilde{R}^2 or $\tilde{\sigma}^2$ to compare and select models. Models with high \tilde{R}^2 (or low $\tilde{\sigma}^2$) are better models in terms of expected out of sample squared error. In contrast, R^2 cannot be used for model selection, as it necessarily increases when regressors are added to a regression model. \bar{R}^2 is also an inappropriate choice for model selection (it tends to select models with too many parameters), though a justification of this assertion requires a study of the theory of model selection. Unfortunately, \bar{R}^2 is routinely used by some economists, possibly as a hold-over from previous generations.

In summary, it is recommended to calculate and report \tilde{R}^2 and/or $\tilde{\sigma}^2$ in regression analysis, and omit R^2 and \bar{R}^2 .

Henri Theil

Henri Theil (1924-2000) of Holland invented \bar{R}^2 and two-stage least squares, both of which are routinely seen in applied econometrics. He also wrote an early influential advanced textbook on econometrics (Theil, 1971).

4.15 Empirical Example

We again return to our wage equation, but use a much larger sample of all individuals with at least 12 years of education. For regressors we include years of education, potential work experience, experience squared, and dummy variable indicators for the following: female, female union member,

male union member, married female¹, married male, formerly married female², formerly married male, hispanic, black, American Indian, Asian, and mixed race³. The available sample is 46,943 so the parameter estimates are quite precise and reported in Table 4.1. For standard errors we use the unbiased Horn-Horn-Duncan formula.

Table 4.1 displays the parameter estimates in a standard tabular format. The table clearly states the estimation method (OLS), the dependent variable ($\log(\text{Wage})$), and the regressors are clearly labeled. Both parameter estimates and standard errors are reported for all coefficients. In addition to the coefficient estimates, the table also reports the estimated error standard deviation and the sample size. These are useful summary measures of fit which aid readers.

Table 4.1
OLS Estimates of Linear Equation for $\log(\text{Wage})$

	$\hat{\beta}$	$s(\hat{\beta})$
Education	0.117	0.001
Experience	0.033	0.001
Experience ² /100	-0.056	0.002
Female	-0.098	0.011
Female Union Member	0.023	0.020
Male Union Member	0.095	0.020
Married Female	0.016	0.010
Married Male	0.211	0.010
Formerly Married Female	-0.006	0.012
Formerly Married Male	0.083	0.015
Hispanic	-0.108	0.008
Black	-0.096	0.008
American Indian	-0.137	0.027
Asian	-0.038	0.013
Mixed Race	-0.041	0.021
Intercept	0.909	0.021
$\hat{\sigma}$	0.565	
Sample Size	46,943	

Note: Standard errors are heteroskedasticity-consistent (Horn-Horn-Duncan formula)

As a general rule, it is advisable to always report standard errors along with parameter estimates. This allows readers to assess the precision of the parameter estimates, and as we will discuss in later chapters, form confidence intervals and t-tests for individual coefficients if desired.

The results in Table 4.1 confirm our earlier findings that the return to a year of education is approximately 12%, the return to experience is concave, that single women earn approximately 10% less than single men, and blacks earn about 10% less than whites. In addition, we see that Hispanics earn about 11% less than whites, American Indians 14% less, and Asians and Mixed races about 4% less. We also see there are wage premiums for men who are members of a labor union (about 10%), married (about 21%) or formerly married (about 8%), but no similar premiums are apparent for women.

¹Defining “married” as marital code 1, 2, or 3.

²Defining “formerly married” as marital code 4, 5, or 6.

³Race code 6 or higher.

4.16 Multicollinearity

If $\mathbf{X}'\mathbf{X}$ is singular, then $(\mathbf{X}'\mathbf{X})^{-1}$ and $\widehat{\boldsymbol{\beta}}$ are not defined. This situation is called **strict multicollinearity**, as the columns of \mathbf{X} are linearly dependent, i.e., there is some $\boldsymbol{\alpha} \neq \mathbf{0}$ such that $\mathbf{X}\boldsymbol{\alpha} = \mathbf{0}$. Most commonly, this arises when sets of regressors are included which are identically related. For example, if \mathbf{X} includes both the logs of two prices and the log of the relative prices, $\log(p_1)$, $\log(p_2)$ and $\log(p_1/p_2)$, for then $\mathbf{X}'\mathbf{X}$ will necessarily be singular. When this happens, the applied researcher quickly discovers the error as the statistical software will be unable to construct $(\mathbf{X}'\mathbf{X})^{-1}$. Since the error is discovered quickly, this is rarely a *problem* for applied econometric practice.

The more relevant situation is **near multicollinearity**, which is often called “multicollinearity” for brevity. This is the situation when the $\mathbf{X}'\mathbf{X}$ matrix is *near* singular, when the columns of \mathbf{X} are *close* to linearly dependent. This definition is not precise, because we have not said what it means for a matrix to be “near singular”. This is one difficulty with the definition and interpretation of multicollinearity.

One potential complication of near singularity of matrices is that the numerical reliability of the calculations may be reduced. In practice this is rarely an important concern, except when the number of regressors is very large.

A more relevant implication of near multicollinearity is that individual coefficient estimates will be imprecise. We can see this most simply in a homoskedastic linear regression model with two regressors

$$y_i = x_{1i}\beta_1 + x_{2i}\beta_2 + e_i,$$

and

$$\frac{1}{n}\mathbf{X}'\mathbf{X} = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}.$$

In this case

$$\text{var}(\widehat{\boldsymbol{\beta}} | \mathbf{X}) = \frac{\sigma^2}{n} \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}^{-1} = \frac{\sigma^2}{n(1-\rho^2)} \begin{pmatrix} 1 & -\rho \\ -\rho & 1 \end{pmatrix}.$$

The correlation ρ indexes collinearity, since as ρ approaches 1 the matrix becomes singular. We can see the effect of collinearity on precision by observing that the variance of a coefficient estimate $\sigma^2 [n(1-\rho^2)]^{-1}$ approaches infinity as ρ approaches 1. Thus the more “collinear” are the regressors, the worse the precision of the individual coefficient estimates.

What is happening is that when the regressors are highly dependent, it is statistically difficult to disentangle the impact of β_1 from that of β_2 . As a consequence, the precision of individual estimates are reduced. The imprecision, however, will be reflected by large standard errors, so there is no distortion in inference.

Some earlier textbooks overemphasized a concern about multicollinearity. A very amusing parody of these texts appeared in Chapter 23.3 of Goldberger’s *A Course in Econometrics* (1991), which is reprinted below. To understand his basic point, you should notice how the estimation variance $\sigma^2 [n(1-\rho^2)]^{-1}$ depends equally and symmetrically on the correlation ρ and the sample size n .

Arthur S. Goldberger

Art Goldberger (1930-2009) was one of the most distinguished members of the Department of Economics at the University of Wisconsin. His PhD thesis developed an early macroeconometric forecasting model (known as the Klein-Goldberger model) but most of his career focused on microeconomic issues. He was the leading pioneer of what has been called the Wisconsin Tradition of empirical work – a combination of formal econometric theory with a careful critical analysis of empirical work. Goldberger wrote a series of highly regarded and influential graduate econometric textbooks, including including *Econometric Theory* (1964), *Topics in Regression Analysis* (1968), and *A Course in Econometrics* (1991).

Micronumerosity

Arthur S. Goldberger

A Course in Econometrics (1991), Chapter 23.3

Econometrics texts devote many pages to the problem of multicollinearity in multiple regression, but they say little about the closely analogous problem of small sample size in estimating a univariate mean. Perhaps that imbalance is attributable to the lack of an exotic polysyllabic name for “small sample size.” If so, we can remove that impediment by introducing the term *micronumerosity*.

Suppose an econometrician set out to write a chapter about small sample size in sampling from a univariate population. Judging from what is now written about multicollinearity, the chapter might look like this:

1. *Micronumerosity*

The extreme case, “exact micronumerosity,” arises when $n = 0$, in which case the sample estimate of μ is not unique. (Technically, there is a violation of the rank condition $n > 0$: the matrix 0 is singular.) The extreme case is easy enough to recognize. “Near micronumerosity” is more subtle, and yet very serious. It arises when the rank condition $n > 0$ is barely satisfied. Near micronumerosity is very prevalent in empirical economics.

2. *Consequences of micronumerosity*

The consequences of micronumerosity are serious. Precision of estimation is reduced. There are two aspects of this reduction: estimates of μ may have large errors, and not only that, but $V_{\bar{y}}$ will be large.

Investigators will sometimes be led to accept the hypothesis $\mu = 0$ because $\bar{y}/\hat{\sigma}_{\bar{y}}$ is small, even though the true situation may be not that $\mu = 0$ but simply that the sample data have not enabled us to pick μ up.

The estimate of μ will be very sensitive to sample data, and the addition of a few more observations can sometimes produce drastic shifts in the sample mean.

The true μ may be sufficiently large for the null hypothesis $\mu = 0$ to be rejected, even though $V_{\bar{y}} = \sigma^2/n$ is large because of micronumerosity. But if the true μ is small (although nonzero) the hypothesis $\mu = 0$ may mistakenly be accepted.

3. *Testing for micronumerosity*

Tests for the presence of micronumerosity require the judicious use of various fingers. Some researchers prefer a single finger, others use their toes, still others let their thumbs rule.

A generally reliable guide may be obtained by counting the number of observations. Most of the time in econometric analysis, when n is close to zero, it is also far from infinity.

Several test procedures develop critical values n^* , such that micronumerosity is a problem only if n is smaller than n^* . But those procedures are questionable.

4. *Remedies for micronumerosity*

If micronumerosity proves serious in the sense that the estimate of μ has an unsatisfactorily low degree of precision, we are in the statistical position of not being able to make bricks without straw. The remedy lies essentially in the acquisition, if possible, of larger samples from the same population.

But more data are no remedy for micronumerosity if the additional data are simply “more of the same.” So obtaining lots of small samples from the same population will not help.

4.17 Normal Regression Model

In the special case of the normal linear regression model introduced in Section 3.18, we can derive exact sampling distributions for the least-squares estimator, residuals, and variance estimator.

In particular, under the normality assumption $e_i \mid \mathbf{x}_i \sim N(0, \sigma^2)$ then we have the multivariate implication

$$\mathbf{e} \mid \mathbf{X} \sim N(\mathbf{0}, \mathbf{I}_n \sigma^2).$$

That is, the error vector \mathbf{e} is independent of \mathbf{X} and is normally distributed. Since linear functions of normals are also normal, this implies that conditional on \mathbf{X}

$$\begin{pmatrix} \hat{\boldsymbol{\beta}} - \boldsymbol{\beta} \\ \hat{\boldsymbol{\varepsilon}} \end{pmatrix} = \begin{pmatrix} (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}' \\ \mathbf{M} \end{pmatrix} \mathbf{e} \sim N\left(0, \begin{pmatrix} \sigma^2 (\mathbf{X}'\mathbf{X})^{-1} & 0 \\ 0 & \sigma^2 \mathbf{M} \end{pmatrix}\right)$$

where $\mathbf{M} = \mathbf{I}_n - \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$. Since uncorrelated normal variables are independent, it follows that $\hat{\boldsymbol{\beta}}$ is independent of any function of the OLS residuals including the estimated error variance s^2 or $\hat{\sigma}^2$ or prediction errors $\tilde{\boldsymbol{\varepsilon}}$.

The spectral decomposition (see equation (A.5)) of \mathbf{M} yields

$$\mathbf{M} = \mathbf{H} \begin{bmatrix} \mathbf{I}_{n-k} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{H}'$$

where $\mathbf{H}'\mathbf{H} = \mathbf{I}_n$. Let $\mathbf{u} = \sigma^{-1}\mathbf{H}'\mathbf{e} \sim \mathbf{N}(\mathbf{0}, \mathbf{H}'\mathbf{H}) \sim \mathbf{N}(\mathbf{0}, \mathbf{I}_n)$. Then

$$\begin{aligned} \frac{n\hat{\sigma}^2}{\sigma^2} &= \frac{(n-k)s^2}{\sigma^2} \\ &= \frac{1}{\sigma^2} \hat{\mathbf{e}}' \hat{\mathbf{e}} \\ &= \frac{1}{\sigma^2} \mathbf{e}' \mathbf{M} \mathbf{e} \\ &= \frac{1}{\sigma^2} \mathbf{e}' \mathbf{H} \begin{bmatrix} \mathbf{I}_{n-k} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{H}' \mathbf{e} \\ &= \mathbf{u}' \begin{bmatrix} \mathbf{I}_{n-k} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{u} \\ &\sim \chi_{n-k}^2, \end{aligned}$$

a chi-square distribution with $n - k$ degrees of freedom.

Furthermore, if standard errors are calculated using the homoskedastic formula (4.27)

$$\frac{\hat{\beta}_j - \beta_j}{s(\hat{\beta}_j)} = \frac{\hat{\beta}_j - \beta_j}{s \sqrt{[(\mathbf{X}'\mathbf{X})^{-1}]_{jj}}} \sim \frac{\mathbf{N}\left(0, \sigma^2 [(\mathbf{X}'\mathbf{X})^{-1}]_{jj}\right)}{\sqrt{\frac{\sigma^2}{n-k} \chi_{n-k}^2} \sqrt{[(\mathbf{X}'\mathbf{X})^{-1}]_{jj}}} = \frac{\mathbf{N}(0, 1)}{\sqrt{\frac{\chi_{n-k}^2}{n-k}}} \sim t_{n-k}$$

a t distribution with $n - k$ degrees of freedom.

Theorem 4.17.1 Normal Regression

In the linear regression model (Assumption 4.3.1) if e_i is independent of \mathbf{x}_i and distributed $\mathbf{N}(0, \sigma^2)$ then

- $\hat{\boldsymbol{\beta}} - \boldsymbol{\beta} \sim \mathbf{N}\left(\mathbf{0}, \sigma^2 (\mathbf{X}'\mathbf{X})^{-1}\right)$
- $\frac{n\hat{\sigma}^2}{\sigma^2} = \frac{(n-k)s^2}{\sigma^2} \sim \chi_{n-k}^2$
- $\frac{\hat{\beta}_j - \beta_j}{s(\hat{\beta}_j)} \sim t_{n-k}$

These are the exact finite-sample distributions of the least-squares estimator and variance estimators, and are the basis for traditional inference in linear regression.

While elegant, the difficulty in applying Theorem 4.17.1 is that the normality assumption is too restrictive to be empirically plausible, and therefore inference based on Theorem 4.17.1 has no guarantee of accuracy. We develop an alternative inference theory based on large sample (asymptotic) approximations in the following chapter.

William Gosset

William S. Gosset (1876-1937) of England is most famous for his derivation of the student's t distribution, published in the paper "The probable error of a mean" in 1908. At the time, Gosset worked at Guinness Brewery, which prohibited its employees from publishing in order to prevent the possible loss of trade secrets. To circumvent this barrier, Gosset published under the pseudonym "Student". Consequently, this famous distribution is known as the student's t rather than Gosset's t!

Exercises

Exercise 4.1 Explain the difference between $\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i'$ and $\mathbb{E}(\mathbf{x}_i \mathbf{x}_i')$.

Exercise 4.2 True or False. If $y_i = x_i \beta + e_i$, $x_i \in \mathbb{R}$, $\mathbb{E}(e_i | x_i) = 0$, and \hat{e}_i is the OLS residual from the regression of y_i on x_i , then $\sum_{i=1}^n x_i^2 \hat{e}_i = 0$.

Exercise 4.3 Prove Theorem 4.6.1.2.

Exercise 4.4 In a linear model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}, \quad \mathbb{E}(\mathbf{e} | \mathbf{X}) = 0, \quad \text{var}(\mathbf{e} | \mathbf{X}) = \sigma^2 \boldsymbol{\Omega}$$

with $\boldsymbol{\Omega}$ a known function of \mathbf{X} , the GLS estimator is

$$\tilde{\boldsymbol{\beta}} = (\mathbf{X}'\boldsymbol{\Omega}^{-1}\mathbf{X})^{-1} (\mathbf{X}'\boldsymbol{\Omega}^{-1}\mathbf{y}),$$

the residual vector is $\hat{\mathbf{e}} = \mathbf{y} - \mathbf{X}\tilde{\boldsymbol{\beta}}$, and an estimate of σ^2 is

$$s^2 = \frac{1}{n-k} \hat{\mathbf{e}}'\boldsymbol{\Omega}^{-1}\hat{\mathbf{e}}.$$

- Find $\mathbb{E}(\tilde{\boldsymbol{\beta}} | \mathbf{X})$.
- Find $\text{var}(\tilde{\boldsymbol{\beta}} | \mathbf{X})$.
- Prove that $\hat{\mathbf{e}} = \mathbf{M}_1 \mathbf{e}$, where $\mathbf{M}_1 = \mathbf{I} - \mathbf{X}(\mathbf{X}'\boldsymbol{\Omega}^{-1}\mathbf{X})^{-1}\mathbf{X}'\boldsymbol{\Omega}^{-1}$.
- Prove that $\mathbf{M}_1'\boldsymbol{\Omega}^{-1}\mathbf{M}_1 = \boldsymbol{\Omega}^{-1} - \boldsymbol{\Omega}^{-1}\mathbf{X}(\mathbf{X}'\boldsymbol{\Omega}^{-1}\mathbf{X})^{-1}\mathbf{X}'\boldsymbol{\Omega}^{-1}$.
- Find $\mathbb{E}(s^2 | \mathbf{X})$.
- Is s^2 a reasonable estimator for σ^2 ?

Exercise 4.5 Let (y_i, \mathbf{x}_i) be a random sample with $\mathbb{E}(\mathbf{y} | \mathbf{X}) = \mathbf{X}\boldsymbol{\beta}$. Consider the **Weighted Least Squares** (WLS) estimator of $\boldsymbol{\beta}$

$$\tilde{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{W}\mathbf{X})^{-1} (\mathbf{X}'\mathbf{W}\mathbf{y})$$

where $\mathbf{W} = \text{diag}(w_1, \dots, w_n)$ and $w_i = x_{ji}^{-2}$, where x_{ji} is one of the \mathbf{x}_i .

- In which contexts would $\tilde{\boldsymbol{\beta}}$ be a good estimator?
- Using your intuition, in which situations would you expect that $\tilde{\boldsymbol{\beta}}$ would perform better than OLS?

Exercise 4.6 Show (4.24) in the homoskedastic regression model.

Exercise 4.7 Prove (4.32).

Exercise 4.8 Show (4.33) and (4.34) in the homoskedastic regression model.

Exercise 4.9 Let $\mu = \mathbb{E}(y_i)$, $\sigma^2 = \mathbb{E}(y_i - \mu)^2$ and $\mu_3 = \mathbb{E}(y_i - \mu)^3$ and consider the sample mean $\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$. Find $\mathbb{E}(\bar{y} - \mu)^3$ as a function of μ , σ^2 , μ_3 and n .

Exercise 4.10 Take the simple regression model $y_i = x_i\beta + e_i$, $x_i \in \mathbb{R}$, $\mathbb{E}(e_i | x_i) = 0$. Define $\sigma_i^2 = \mathbb{E}(e_i^2 | x_i)$ and $\mu_{3i} = \mathbb{E}(e_i^3 | x_i)$ and consider the OLS coefficient $\hat{\beta}$. Find $\mathbb{E}\left(\left(\hat{\beta} - \beta\right)^3 | \mathbf{X}\right)$.

Exercise 4.11 Continue the empirical analysis in Exercise 3.19.

1. Calculate standard errors using the homoskedasticity formula and using the four covariance matrices from Section 4.11.
2. Repeat in your second programming language. Are they identical?

Exercise 4.12 Continue the empirical analysis in Exercise 3.21. Calculate standard errors using the Horn-Horn-Duncan method. Repeat in your second programming language. Are they identical?

Chapter 5

An Introduction to Large Sample Asymptotics

5.1 Introduction

In Chapter 4 we derived the mean and variance of the least-squares estimator in the context of the linear regression model, but this is not a complete description of the sampling distribution, nor sufficient for inference (confidence intervals and hypothesis testing) on the unknown parameters. Furthermore, the theory does not apply in the context of the linear projection model, which is more relevant for empirical applications.

To illustrate the situation with an example, let y_i and x_i be drawn from the joint density

$$f(x, y) = \frac{1}{2\pi xy} \exp\left(-\frac{1}{2}(\log y - \log x)^2\right) \exp\left(-\frac{1}{2}(\log x)^2\right)$$

and let $\hat{\beta}$ be the slope coefficient estimate from a least-squares regression of y_i on x_i and a constant. Using simulation methods, the density function of $\hat{\beta}$ was computed and plotted in Figure 5.1 for sample sizes of $n = 25$, $n = 100$ and $n = 800$. The vertical line marks the true projection coefficient.

From the figure we can see that the density functions are dispersed and highly non-normal. As the sample size increases the density becomes more concentrated about the population coefficient. Is there a simple way to characterize the sampling distribution of $\hat{\beta}$?

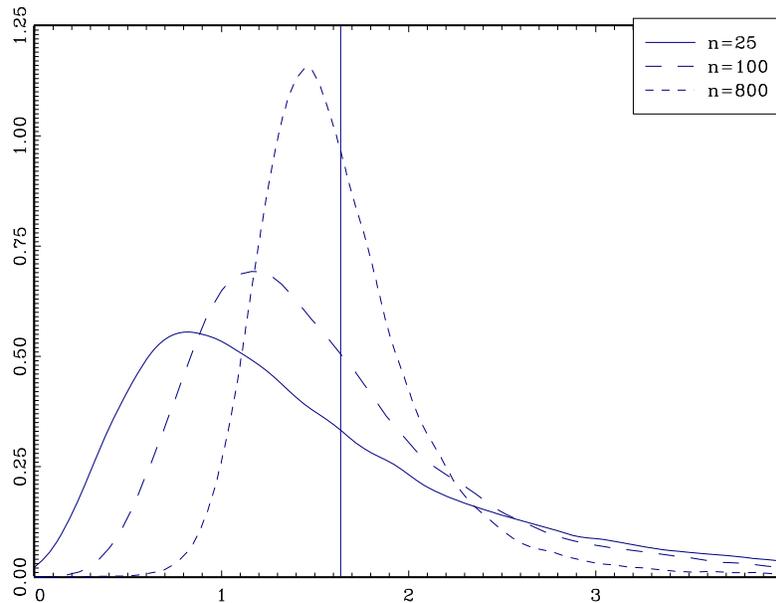
In principle the sampling distribution of $\hat{\beta}$ is a function of the joint distribution of (y_i, x_i) and the sample size n , but in practice this function is extremely complicated so it is not feasible to analytically calculate the exact distribution of $\hat{\beta}$ except in very special cases. Therefore we typically rely on approximation methods.

The most widely used and versatile method is asymptotic theory, which approximates sampling distributions by taking the limit of the finite sample distribution as the sample size n tends to infinity. It is important to understand that this is an approximation technique, as the asymptotic distributions are used to assess the finite sample distributions of our estimators in actual practical samples. The primary tools of asymptotic theory are the weak law of large numbers (WLLN), central limit theorem (CLT), and continuous mapping theorem (CMT). With these tools we can approximate the sampling distributions of most econometric estimators.

In this chapter we provide a concise summary. It will be useful for most students to review this material, even if most is familiar.

5.2 Asymptotic Limits

“Asymptotic analysis” is a method of approximation obtained by taking a suitable limit. There is more than one method to take limits, but the most common is to take the limit of the sequence

Figure 5.1: Sampling Density of $\hat{\beta}$

of sampling distributions as the sample size tends to positive infinity, written “as $n \rightarrow \infty$.” It is not meant to be interpreted literally, but rather as an approximating device.

The first building block for asymptotic analysis is the concept of a limit of a sequence.

Definition 5.2.1 A sequence a_n has the **limit** a , written $a_n \rightarrow a$ as $n \rightarrow \infty$, or alternatively as $\lim_{n \rightarrow \infty} a_n = a$, if for all $\delta > 0$ there is some $n_\delta < \infty$ such that for all $n \geq n_\delta$, $|a_n - a| \leq \delta$.

In words, a_n has the limit a if the sequence gets closer and closer to a as n gets larger. If a sequence has a limit, that limit is unique (a sequence cannot have two distinct limits). If a_n has the limit a , we also say that a_n **converges** to a as $n \rightarrow \infty$.

Not all sequences have limits. For example, the sequence $\{1, 2, 1, 2, 1, 2, \dots\}$ does not have a limit. It is therefore sometimes useful to have a more general definition of limits which always exist, and these are the limit superior and limit inferior of sequence

Definition 5.2.2 $\liminf_{n \rightarrow \infty} a_n \stackrel{def}{=} \lim_{n \rightarrow \infty} \inf_{m \geq n} a_m$

Definition 5.2.3 $\limsup_{n \rightarrow \infty} a_n \stackrel{def}{=} \lim_{n \rightarrow \infty} \sup_{m \geq n} a_m$

The limit inferior and limit superior always exist, and equal when the limit exists. In the example given earlier, the limit inferior of $\{1, 2, 1, 2, 1, 2, \dots\}$ is 1, and the limit superior is 2.

5.3 Convergence in Probability

A sequence of numbers may converge to a limit, but what about a sequence of random variables? For example, consider a sample mean $\bar{y} = n^{-1} \sum_{i=1}^n y_i$ based on an random sample of n observations. As n increases, the distribution of \bar{y} changes. In what sense can we describe the “limit” of \bar{y} ? In what sense does it converge?

Since \bar{y} is a random variable, we cannot directly apply the deterministic concept of a sequence of numbers. Instead, we require a definition of convergence which is appropriate for random variables. There are more than one such definition, but the most commonly used is called convergence in probability.

Definition 5.3.1 A random variable $z_n \in \mathbb{R}$ **converges in probability** to z as $n \rightarrow \infty$, denoted $z_n \xrightarrow{p} z$, or alternatively $\text{plim}_{n \rightarrow \infty} z_n = z$, if for all $\delta > 0$,

$$\lim_{n \rightarrow \infty} \Pr(|z_n - z| \leq \delta) = 1. \quad (5.1)$$

We call z the **probability limit** (or **plim**) of z_n .

The definition looks quite abstract, but it formalizes the concept of a sequence of random variables concentrating about a point. The event $\{|z_n - z| \leq \delta\}$ occurs when z_n is within δ of the point z . $\Pr(|z_n - z| \leq \delta)$ is the probability of this event – that z_n is within δ of the point z . Equation (5.1) states that this probability approaches 1 as the sample size n increases. The definition of convergence in probability requires that this holds for any δ . So for any small interval about z the distribution of z_n concentrates within this interval for large n .

You may notice that the definition concerns the *distribution* of the random variables z_n , not their realizations. Furthermore, notice that the definition uses the concept of a conventional (deterministic) limit, but the latter is applied to a sequence of probabilities, not directly to the random variables z_n or their realizations.

Two comments about the notation are worth mentioning. First, it is conventional to write the convergence symbol as \xrightarrow{p} where the “ p ” above the arrow indicates that the convergence is “in probability”. You should try and adhere to this notation, and not simply write $z_n \rightarrow z$. Second, it is important to include the phrase “as $n \rightarrow \infty$ ” to be specific about how the limit is obtained.

A common mistake to confuse convergence in probability with convergence in expectation:

$$\mathbb{E}z_n \longrightarrow \mathbb{E}z. \quad (5.2)$$

They are related but distinct concepts. Neither (5.1) nor (5.2) implies the other.

To see the distinction it might be helpful to think through a stylized example. Consider a discrete random variable z_n which takes the value 0 with probability $1 - n^{-1}$ and the value $a_n \neq 0$ with probability n^{-1} , or

$$\begin{aligned} \Pr(z_n = 0) &= 1 - \frac{1}{n} \\ \Pr(z_n = a_n) &= \frac{1}{n}. \end{aligned} \quad (5.3)$$

In this example the probability distribution of z_n concentrates at zero as n increases, regardless of the sequence a_n . You can check that $z_n \xrightarrow{p} 0$ as $n \rightarrow \infty$.

In this example we can also calculate that the expectation of z_n is

$$\mathbb{E}z_n = \frac{a_n}{n}.$$

Despite the fact that z_n converges in probability to zero, its expectation will not decrease to zero unless $a_n/n \rightarrow 0$. If a_n diverges to infinity at a rate equal to n (or faster) then $\mathbb{E}z_n$ will not converge to zero. For example, if $a_n = n$, then $\mathbb{E}z_n = 1$ for all n , even though $z_n \xrightarrow{p} 0$. This example might seem a bit artificial, but the point is that the concepts of convergence in probability and convergence in expectation are distinct, so it is important not to confuse one with the other.

Another common source of confusion with the notation surrounding probability limits is that the expression to the right of the arrow “ \xrightarrow{p} ” must be free of dependence on the sample size n . Thus expressions of the form “ $z_n \xrightarrow{p} c_n$ ” are notationally meaningless and should not be used.

5.4 Weak Law of Large Numbers

In large samples we expect parameter estimates to be close to the population values. For example, in Section 4.2 we saw that the sample mean \bar{y} is unbiased for $\mu = \mathbb{E}y$ and has variance σ^2/n . As n gets large its variance decreases and thus the distribution of \bar{y} concentrates about the population mean μ . It turns out that this implies that the sample mean converges in probability to the population mean.

When y has a finite variance there is a fairly straightforward proof by applying Chebyshev’s inequality.

Theorem 5.4.1 Chebyshev’s Inequality. For any random variable z_n and constant $\delta > 0$

$$\Pr(|z_n - \mathbb{E}z_n| > \delta) \leq \frac{\text{var}(z_n)}{\delta^2}.$$

Chebyshev’s inequality is terrifically important in asymptotic theory. While its proof is a technical exercise in probability theory, it is quite simple so we discuss it forthwith. Let $F_n(u)$ denote the distribution of $z_n - \mathbb{E}z_n$. Then

$$\Pr(|z_n - \mathbb{E}z_n| > \delta) = \Pr\left((z_n - \mathbb{E}z_n)^2 > \delta^2\right) = \int_{\{u^2 > \delta^2\}} dF_n(u).$$

The integral is over the event $\{u^2 > \delta^2\}$, so that the inequality $1 \leq \frac{u^2}{\delta^2}$ holds throughout. Thus

$$\int_{\{u^2 > \delta^2\}} dF_n(u) \leq \int_{\{u^2 > \delta^2\}} \frac{u^2}{\delta^2} dF_n(u) \leq \int \frac{u^2}{\delta^2} dF_n(u) = \frac{\mathbb{E}(z_n - \mathbb{E}z_n)^2}{\delta^2} = \frac{\text{var}(z_n)}{\delta^2},$$

which establishes the desired inequality.

Applied to the sample mean \bar{y} , Chebyshev’s inequality shows that for any $\delta > 0$

$$\Pr(|\bar{y} - \mathbb{E}\bar{y}| > \delta) \leq \frac{\sigma^2}{n\delta^2}.$$

For fixed σ^2 and δ , the bound on the right-hand-side shrinks to zero as $n \rightarrow \infty$. Thus the probability that \bar{y} is within δ of $\mathbb{E}\bar{y} = \mu$ approaches 1 as n gets large, or

$$\lim_{n \rightarrow \infty} \Pr(|\bar{y} - \mu| \leq \delta) = 1.$$

This means that \bar{y} converges in probability to μ as $n \rightarrow \infty$.

This result is called the **weak law of large numbers**. Our derivation assumed that y has a finite variance, but with a more careful proof all that is necessary is a finite mean.

Theorem 5.4.2 Weak Law of Large Numbers (WLLN)

If y_i are independent and identically distributed and $\mathbb{E}|y| < \infty$, then as $n \rightarrow \infty$,

$$\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i \xrightarrow{p} \mathbb{E}(y).$$

The proof of Theorem 5.4.2 is presented in Section 5.14.

The WLLN shows that the estimator \bar{y} converges in probability to the true population mean μ . In general, an estimator which converges in probability to the population value is called **consistent**.

Definition 5.4.1 An estimator $\hat{\theta}$ of a parameter θ is **consistent** if $\hat{\theta} \xrightarrow{p} \theta$ as $n \rightarrow \infty$.

Theorem 5.4.3 If y_i are independent and identically distributed and $\mathbb{E}|y| < \infty$, then $\hat{\mu} = \bar{y}$ is consistent for the population mean μ .

Consistency is a good property for an estimator to possess. It means that for any given data distribution, there is a sample size n sufficiently large such that the estimator $\hat{\theta}$ will be arbitrarily close to the true value θ with high probability. The theorem does not tell us, however, how large this n has to be. Thus the theorem does not give practical guidance for empirical practice. Still, it is a minimal property for an estimator to be considered a “good” estimator, and provides a foundation for more useful approximations.

5.5 Almost Sure Convergence and the Strong Law*

Convergence in probability is sometimes called **weak convergence**. A related concept is **almost sure convergence**, also known as **strong convergence**. (In probability theory the term “almost sure” means “with probability equal to one”. An event which is random but occurs with probability equal to one is said to be **almost sure**.)

Definition 5.5.1 A random variable $z_n \in \mathbb{R}$ **converges almost surely** to z as $n \rightarrow \infty$, denoted $z_n \xrightarrow{a.s.} z$, if for every $\delta > 0$

$$\Pr \left(\lim_{n \rightarrow \infty} |z_n - z| \leq \delta \right) = 1. \quad (5.4)$$

The convergence (5.4) is stronger than (5.1) because it computes the probability of a limit rather than the limit of a probability. Almost sure convergence is stronger than convergence in probability in the sense that $z_n \xrightarrow{a.s.} z$ implies $z_n \xrightarrow{p} z$.

In the example (5.3) of Section 5.3, the sequence z_n converges in probability to zero for any sequence a_n , but this is not sufficient for z_n to converge almost surely. In order for z_n to converge to zero almost surely, it is necessary that $a_n \rightarrow 0$.

In the random sampling context the sample mean can be shown to converge almost surely to the population mean. This is called the **strong law of large numbers**.

Theorem 5.5.1 Strong Law of Large Numbers (SLLN)

If y_i are independent and identically distributed and $\mathbb{E}|y| < \infty$, then as $n \rightarrow \infty$,

$$\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i \xrightarrow{\text{a.s.}} \mathbb{E}(y).$$

The proof of the SLLN is technically quite advanced so is not presented here. For a proof see Billingsley (1995, Section 22) or Ash (1972, Theorem 7.2.5).

The WLLN is sufficient for most purposes in econometrics, so we will not use the SLLN in this text.

5.6 Vector-Valued Moments

Our preceding discussion focused on the case where y is real-valued (a scalar), but nothing important changes if we generalize to the case where $\mathbf{y} \in \mathbb{R}^m$ is a vector. To fix notation, the elements of \mathbf{y} are

$$\mathbf{y} = \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_m \end{pmatrix}.$$

The population mean of \mathbf{y} is just the vector of marginal means

$$\boldsymbol{\mu} = \mathbb{E}(\mathbf{y}) = \begin{pmatrix} \mathbb{E}(y_1) \\ \mathbb{E}(y_2) \\ \vdots \\ \mathbb{E}(y_m) \end{pmatrix}.$$

When working with random vectors \mathbf{y} it is convenient to measure their magnitude by their Euclidean length or Euclidean norm

$$\|\mathbf{y}\| = (y_1^2 + \cdots + y_m^2)^{1/2}.$$

In vector notation we have

$$\|\mathbf{y}\|^2 = \mathbf{y}'\mathbf{y}.$$

It turns out that it is equivalent to describe finiteness of moments in terms of the Euclidean norm of a vector or all individual components.

Theorem 5.6.1 For $\mathbf{y} \in \mathbb{R}^m$, $\mathbb{E}\|\mathbf{y}\| < \infty$ if and only if $\mathbb{E}|y_j| < \infty$ for $j = 1, \dots, m$.

The $m \times m$ variance matrix of \mathbf{y} is

$$\mathbf{V} = \text{var}(\mathbf{y}) = \mathbb{E}((\mathbf{y} - \boldsymbol{\mu})(\mathbf{y} - \boldsymbol{\mu})').$$

\mathbf{V} is often called a variance-covariance matrix. You can show that the elements of \mathbf{V} are finite if $\mathbb{E}\|\mathbf{y}\|^2 < \infty$.

A random sample $\{\mathbf{y}_1, \dots, \mathbf{y}_n\}$ consists of n observations of independent and identically distributed draws from the distribution of \mathbf{y} . (Each draw is an m -vector.) The vector sample mean

$$\bar{\mathbf{y}} = \frac{1}{n} \sum_{i=1}^n \mathbf{y}_i = \begin{pmatrix} \bar{y}_1 \\ \bar{y}_2 \\ \vdots \\ \bar{y}_m \end{pmatrix}$$

is the vector of sample means of the individual variables.

Convergence in probability of a vector can be defined as convergence in probability of all elements in the vector. Thus $\bar{\mathbf{y}} \xrightarrow{p} \boldsymbol{\mu}$ if and only if $\bar{y}_j \xrightarrow{p} \mu_j$ for $j = 1, \dots, m$. Since the latter holds if $\mathbb{E}|y_j| < \infty$ for $j = 1, \dots, m$, or equivalently $\mathbb{E}\|\mathbf{y}\| < \infty$, we can state this formally as follows.

Theorem 5.6.2 Weak Law of Large Numbers (WLLN) for random vectors

If \mathbf{y}_i are independent and identically distributed and $\mathbb{E}\|\mathbf{y}\| < \infty$, then as $n \rightarrow \infty$,

$$\bar{\mathbf{y}} = \frac{1}{n} \sum_{i=1}^n \mathbf{y}_i \xrightarrow{p} \mathbb{E}(\mathbf{y}).$$

5.7 Convergence in Distribution

The WLLN is a useful first step, but does not give an approximation to the distribution of an estimator. A **large-sample** or **asymptotic** approximation can be obtained using the concept of convergence in distribution.

Definition 5.7.1 Let \mathbf{z}_n be a random vector with distribution $F_n(\mathbf{u}) = \Pr(\mathbf{z}_n \leq \mathbf{u})$. We say that \mathbf{z}_n **converges in distribution** to \mathbf{z} as $n \rightarrow \infty$, denoted $\mathbf{z}_n \xrightarrow{d} \mathbf{z}$, if for all \mathbf{u} at which $F(\mathbf{u}) = \Pr(\mathbf{z} \leq \mathbf{u})$ is continuous, $F_n(\mathbf{u}) \rightarrow F(\mathbf{u})$ as $n \rightarrow \infty$.

When $\mathbf{z}_n \xrightarrow{d} \mathbf{z}$, it is common to refer to \mathbf{z} as the **asymptotic distribution** or **limit distribution** of \mathbf{z}_n .

When the limit distribution \mathbf{z} is degenerate (that is, $\Pr(\mathbf{z} = \mathbf{c}) = 1$ for some \mathbf{c}) we can write the convergence as $\mathbf{z}_n \xrightarrow{d} \mathbf{c}$, which is equivalent to convergence in probability, $\mathbf{z}_n \xrightarrow{p} \mathbf{c}$.

The typical path to establishing convergence in distribution is through the central limit theorem (CLT), which states that a standardized sample average converges in distribution to a normal random vector.

Theorem 5.7.1 Lindeberg–Lévy Central Limit Theorem (CLT). If \mathbf{y}_i are independent and identically distributed and $\mathbb{E}\|\mathbf{y}\|^2 < \infty$, then as $n \rightarrow \infty$

$$\sqrt{n}(\bar{\mathbf{y}} - \boldsymbol{\mu}) = \frac{1}{\sqrt{n}} \sum_{i=1}^n (\mathbf{y}_i - \boldsymbol{\mu}) \xrightarrow{d} \mathbf{N}(\mathbf{0}, \mathbf{V})$$

where $\boldsymbol{\mu} = \mathbb{E}\mathbf{y}$ and $\mathbf{V} = \mathbb{E}((\mathbf{y} - \boldsymbol{\mu})(\mathbf{y} - \boldsymbol{\mu})')$.

The standardized sum $\mathbf{z}_n = \sqrt{n}(\bar{\mathbf{y}}_n - \boldsymbol{\mu})$ has mean zero and variance \mathbf{V} . What the CLT adds is that the variable \mathbf{z}_n is also approximately normally distributed, and that the normal approximation improves as n increases.

The CLT is one of the most powerful and mysterious results in statistical theory. It shows that the simple process of averaging induces normality. The first version of the CLT (for the number of heads resulting from many tosses of a fair coin) was established by the French mathematician Abraham de Moivre in an article published in 1733. This was extended to cover an approximation to the binomial distribution in 1812 by Pierre-Simon Laplace in his book *Théorie Analytique des Probabilités*, and the most general statements are credited to articles by the Russian mathematician Aleksandr Lyapunov (1901) and the Finnish mathematician Jarl Waldemar Lindeberg (1920, 1922). The above statement is known as the classic (or Lindeberg–Lévy) CLT due to contributions by Lindeberg (1920) and the French mathematician Paul Pierre Lévy.

A more general version which does not require the restriction to identical distributions was provided by Lindeberg (1922).

Theorem 5.7.2 Lindeberg Central Limit Theorem (CLT). Suppose that y_i are independent but not necessarily identically distributed with finite means $\mu_i = \mathbb{E}y_i$ and variances $\sigma_i^2 = \mathbb{E}(y_i - \mu_i)^2$. Set $\nu_n^2 = \sum_{i=1}^n \sigma_i^2$. If for all $\varepsilon > 0$

$$\lim_{n \rightarrow \infty} \frac{1}{\nu_n^2} \sum_{i=1}^n \mathbb{E}(y_i - \mu_i)^2 \mathbf{1}(|y_i - \mu_i| \geq \varepsilon \nu_n) = 0 \quad (5.5)$$

then

$$\frac{1}{\nu_n} \sum_{i=1}^n (y_i - \mu_i) \xrightarrow{d} \mathbf{N}(0, 1).$$

Equation (5.5) is known as Lindeberg’s condition. A standard method to verify (5.5) is via Lyapunov’s condition: For some $\delta > 0$

$$\lim_{n \rightarrow \infty} \frac{1}{\nu_n^{2+\delta}} \sum_{i=1}^n \mathbb{E}(y_i - \mu_i)^{2+\delta} = 0. \quad (5.6)$$

It is easy to verify that (5.6) implies (5.5), and (5.6) is often easy to verify. For example, if $\sup_i \mathbb{E}(y_i - \mu_i)^3 \leq \kappa < \infty$ and $\inf_i \sigma_i^2 \geq c > 0$ then

$$\frac{1}{\nu_n^3} \sum_{i=1}^n \mathbb{E}(y_i - \mu_i)^3 \leq \frac{n\kappa}{(nc)^{3/2}} \rightarrow 0$$

so (5.6) is satisfied.

5.8 Higher Moments

Often we want to estimate a parameter $\boldsymbol{\mu}$ which is the expected value of a transformation of a random vector \mathbf{y} . That is, $\boldsymbol{\mu}$ can be written as

$$\boldsymbol{\mu} = \mathbb{E}\mathbf{h}(\mathbf{y})$$

for some function $\mathbf{h} : \mathbb{R}^m \rightarrow \mathbb{R}^k$. For example, the second moment of y is $\mathbb{E}y^2$, the k 'th is $\mathbb{E}y^k$, the moment generating function is $\mathbb{E}\exp(ty)$, and the distribution function is $\mathbb{E}1\{y \leq x\}$.

Estimating parameters of this form fits into our previous analysis by defining the random variable $\mathbf{z} = \mathbf{h}(\mathbf{y})$ for then $\boldsymbol{\mu} = \mathbb{E}\mathbf{z}$ is just a simple moment of \mathbf{z} . This suggests the moment estimator

$$\hat{\boldsymbol{\mu}} = \frac{1}{n} \sum_{i=1}^n \mathbf{z}_i = \frac{1}{n} \sum_{i=1}^n \mathbf{h}(\mathbf{y}_i).$$

For example, the moment estimator of $\mathbb{E}y^k$ is $n^{-1} \sum_{i=1}^n y_i^k$, that of the moment generating function is $n^{-1} \sum_{i=1}^n \exp(ty_i)$, and for the distribution function the estimator is $n^{-1} \sum_{i=1}^n 1\{y_i \leq x\}$.

Since $\hat{\boldsymbol{\mu}}$ is a sample average, and transformations of iid variables are also iid, the asymptotic results of the previous sections immediately apply.

Theorem 5.8.1 *If \mathbf{y}_i are independent and identically distributed, $\boldsymbol{\mu} = \mathbb{E}\mathbf{h}(\mathbf{y})$, and $\mathbb{E}\|\mathbf{h}(\mathbf{y})\| < \infty$, then for $\hat{\boldsymbol{\mu}} = \frac{1}{n} \sum_{i=1}^n \mathbf{h}(\mathbf{y}_i)$, as $n \rightarrow \infty$, $\hat{\boldsymbol{\mu}} \xrightarrow{p} \boldsymbol{\mu}$.*

Theorem 5.8.2 *If \mathbf{y}_i are independent and identically distributed, $\boldsymbol{\mu} = \mathbb{E}\mathbf{h}(\mathbf{y})$, and $\mathbb{E}\|\mathbf{h}(\mathbf{y})\|^2 < \infty$, then for $\hat{\boldsymbol{\mu}} = \frac{1}{n} \sum_{i=1}^n \mathbf{h}(\mathbf{y}_i)$, as $n \rightarrow \infty$,*

$$\sqrt{n}(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}) \xrightarrow{d} \mathbf{N}(\mathbf{0}, \mathbf{V})$$

where $\mathbf{V} = \mathbb{E}((\mathbf{h}(\mathbf{y}) - \boldsymbol{\mu})(\mathbf{h}(\mathbf{y}) - \boldsymbol{\mu})')$.

Theorems 5.8.1 and 5.8.2 show that the estimate $\hat{\boldsymbol{\mu}}$ is consistent for $\boldsymbol{\mu}$ and asymptotically normally distributed, so long as the stated moment conditions hold.

A word of caution. Theorems 5.8.1 and 5.8.2 give the impression that it is possible to estimate any moment of y . Technically this is the case so long as that moment is finite. What is hidden by the notation, however, is that estimates of high order moments can be quite imprecise. For example, consider the sample 8th moment $\hat{\mu}_8 = \frac{1}{n} \sum_{i=1}^n y_i^8$, and suppose for simplicity that y is $\mathbf{N}(0, 1)$. Then we can calculate¹ that $\text{var}(\hat{\mu}_8) = n^{-1}2,016,000$, which is immense, even for large n ! In general, higher-order moments are challenging to estimate because their variance depends upon even higher moments which can be quite large in some cases.

¹By the formula for the variance of a mean $\text{var}(\hat{\mu}_8) = n^{-1}(\mathbb{E}y^{16} - (\mathbb{E}y^8)^2)$. Since y is $\mathbf{N}(0, 1)$, $\mathbb{E}y^{16} = 15!! = 2,027,025$ and $\mathbb{E}y^8 = 7!! = 105$ where $k!! = k(k-2)\cdots$ is the double factorial for odd k .

5.9 Functions of Moments

We now expand our investigation and consider estimation of parameters which can be written as a continuous function of $\boldsymbol{\mu} = \mathbb{E}\mathbf{h}(\mathbf{y})$. That is, the parameter of interest can be written as

$$\boldsymbol{\beta} = \mathbf{g}(\boldsymbol{\mu}) = \mathbf{g}(\mathbb{E}\mathbf{h}(\mathbf{y})) \quad (5.7)$$

for some functions $\mathbf{g} : \mathbb{R}^k \rightarrow \mathbb{R}^\ell$ and $\mathbf{h} : \mathbb{R}^m \rightarrow \mathbb{R}^k$.

As one example, the geometric mean of wages w is

$$\gamma = \exp(\mathbb{E}(\log(w))). \quad (5.8)$$

This is (5.7) with $g(u) = \exp(u)$ and $h(w) = \log(w)$.

A simple yet common example is the variance

$$\begin{aligned} \sigma^2 &= \mathbb{E}(w - \mathbb{E}w)^2 \\ &= \mathbb{E}w^2 - (\mathbb{E}w)^2. \end{aligned}$$

This is (5.7) with

$$\mathbf{h}(w) = \begin{pmatrix} w \\ w^2 \end{pmatrix}$$

and

$$g(\mu_1, \mu_2) = \mu_2 - \mu_1^2.$$

Similarly, the skewness of the wage distribution is

$$sk = \frac{\mathbb{E}(w - \mathbb{E}w)^3}{(\mathbb{E}(w - \mathbb{E}w)^2)^{3/2}}.$$

This is (5.7) with

$$\mathbf{h}(w) = \begin{pmatrix} w \\ w^2 \\ w^3 \end{pmatrix}$$

and

$$g(\mu_1, \mu_2, \mu_3) = \frac{\mu_3 - 3\mu_2\mu_1 + 2\mu_1^3}{(\mu_2 - \mu_1^2)^{3/2}}. \quad (5.9)$$

The parameter $\boldsymbol{\beta} = \mathbf{g}(\boldsymbol{\mu})$ is not a population moment, so it does not have a direct moment estimator. Instead, it is common to use a **plug-in estimate** formed by replacing the unknown $\boldsymbol{\mu}$ with its point estimate $\hat{\boldsymbol{\mu}}$ and then “plugging” this into the expression for $\boldsymbol{\beta}$. The first step is

$$\hat{\boldsymbol{\mu}} = \frac{1}{n} \sum_{i=1}^n \mathbf{h}(\mathbf{y}_i)$$

and the second step is

$$\hat{\boldsymbol{\beta}} = \mathbf{g}(\hat{\boldsymbol{\mu}}).$$

Again, the hat “ $\hat{}$ ” indicates that $\hat{\boldsymbol{\beta}}$ is a sample estimate of $\boldsymbol{\beta}$.

For example, the plug-in estimate of the geometric mean γ of the wage distribution from (5.8) is

$$\hat{\gamma} = \exp(\hat{\boldsymbol{\mu}})$$

with

$$\hat{\boldsymbol{\mu}} = \frac{1}{n} \sum_{i=1}^n \log(\text{wage}_i).$$

The plug-in estimate of the variance is

$$\begin{aligned}\hat{\sigma}^2 &= \frac{1}{n} \sum_{i=1}^n w_i^2 - \left(\frac{1}{n} \sum_{i=1}^n w_i \right)^2 \\ &= \frac{1}{n} \sum_{i=1}^n (w_i - \bar{w})^2.\end{aligned}$$

and that for the skewness is

$$\begin{aligned}\widehat{sk} &= \frac{\widehat{\mu}_3 - 3\widehat{\mu}_2\widehat{\mu}_1 + 2\widehat{\mu}_1^3}{(\widehat{\mu}_2 - \widehat{\mu}_1^2)^{3/2}} \\ &= \frac{\frac{1}{n} \sum_{i=1}^n (w_i - \bar{w})^3}{\left(\frac{1}{n} \sum_{i=1}^n (w_i - \bar{w})^2 \right)^{3/2}}\end{aligned}$$

where

$$\widehat{\mu}_j = \frac{1}{n} \sum_{i=1}^n w_i^j.$$

A useful property is that continuous functions are limit-preserving.

Theorem 5.9.1 *Continuous Mapping Theorem (CMT).* If $z_n \xrightarrow{p} c$ as $n \rightarrow \infty$ and $g(\cdot)$ is continuous at c , then $g(z_n) \xrightarrow{p} g(c)$ as $n \rightarrow \infty$.

The proof of Theorem 5.9.1 is given in Section 5.14.

For example, if $z_n \xrightarrow{p} c$ as $n \rightarrow \infty$ then

$$\begin{aligned}z_n + a &\xrightarrow{p} c + a \\ az_n &\xrightarrow{p} ac \\ z_n^2 &\xrightarrow{p} c^2\end{aligned}$$

as the functions $g(u) = u + a$, $g(u) = au$, and $g(u) = u^2$ are continuous. Also

$$\frac{a}{z_n} \xrightarrow{p} \frac{a}{c}$$

if $c \neq 0$. The condition $c \neq 0$ is important as the function $g(u) = a/u$ is not continuous at $u = 0$.

If \mathbf{y}_i are independent and identically distributed, $\boldsymbol{\mu} = \mathbb{E}\mathbf{h}(\mathbf{y})$, and $\mathbb{E}\|\mathbf{h}(\mathbf{y})\| < \infty$, then for $\widehat{\boldsymbol{\mu}} = \frac{1}{n} \sum_{i=1}^n \mathbf{h}(\mathbf{y}_i)$, as $n \rightarrow \infty$, $\widehat{\boldsymbol{\mu}} \xrightarrow{p} \boldsymbol{\mu}$. Applying the CMT, $\widehat{\boldsymbol{\beta}} = \mathbf{g}(\widehat{\boldsymbol{\mu}}) \xrightarrow{p} \mathbf{g}(\boldsymbol{\mu}) = \boldsymbol{\beta}$.

Theorem 5.9.2 If \mathbf{y}_i are independent and identically distributed, $\boldsymbol{\beta} = \mathbf{g}(\mathbb{E}\mathbf{h}(\mathbf{y}))$, $\mathbb{E}\|\mathbf{h}(\mathbf{y})\| < \infty$, and $\mathbf{g}(\mathbf{u})$ is continuous at $\mathbf{u} = \boldsymbol{\mu}$, then for $\widehat{\boldsymbol{\beta}} = \mathbf{g}(\frac{1}{n} \sum_{i=1}^n \mathbf{h}(\mathbf{y}_i))$, as $n \rightarrow \infty$, $\widehat{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\beta}$.

To apply Theorem 5.9.2 it is necessary to check if the function \mathbf{g} is continuous at $\boldsymbol{\mu}$. In our first example $g(u) = \exp(u)$ is continuous everywhere. It therefore follows from Theorem 5.6.2 and Theorem 5.9.2 that if $\mathbb{E}|\log(\text{wage})| < \infty$ then as $n \rightarrow \infty$, $\widehat{\gamma} \xrightarrow{p} \gamma$.

In the example of the variance, g is continuous for all $\boldsymbol{\mu}$. Thus if $\mathbb{E}w^2 < \infty$ then as $n \rightarrow \infty$, $\widehat{\sigma}^2 \xrightarrow{p} \sigma^2$.

In our third example g defined in (5.9) is continuous for all $\boldsymbol{\mu}$ such that $\text{var}(w) = \mu_2 - \mu_1^2 > 0$, which holds unless w has a degenerate distribution. Thus if $\mathbb{E}|w|^3 < \infty$ and $\text{var}(w) > 0$ then as $n \rightarrow \infty$, $\widehat{sk} \xrightarrow{p} sk$.

5.10 Delta Method

In this section we introduce two tools – an extended version of the CMT and the Delta Method – which allow us to calculate the asymptotic distribution of the parameter estimate $\widehat{\beta}$.

We first present an extended version of the continuous mapping theorem which allows convergence in distribution.

Theorem 5.10.1 *Continuous Mapping Theorem*

If $z_n \xrightarrow{d} z$ as $n \rightarrow \infty$ and $g : \mathbb{R}^m \rightarrow \mathbb{R}^k$ has the set of discontinuity points D_g such that $\Pr(z \in D_g) = 0$, then $g(z_n) \xrightarrow{d} g(z)$ as $n \rightarrow \infty$.

For a proof of Theorem 5.10.1 see Theorem 2.3 of van der Vaart (1998). It was first proved by Mann and Wald (1943) and is therefore sometimes referred to as the Mann-Wald Theorem.

Theorem 5.10.1 allows the function g to be discontinuous only if the probability at being at a discontinuity point is zero. For example, the function $g(u) = u^{-1}$ is discontinuous at $u = 0$, but if $z_n \xrightarrow{d} z \sim N(0, 1)$ then $\Pr(z = 0) = 0$ so $z_n^{-1} \xrightarrow{d} z^{-1}$.

A special case of the Continuous Mapping Theorem is known as Slutsky's Theorem.

Theorem 5.10.2 *Slutsky's Theorem*

If $z_n \xrightarrow{d} z$ and $c_n \xrightarrow{p} c$ as $n \rightarrow \infty$, then

1. $z_n + c_n \xrightarrow{d} z + c$
2. $z_n c_n \xrightarrow{d} z c$
3. $\frac{z_n}{c_n} \xrightarrow{d} \frac{z}{c}$ if $c \neq 0$

Even though Slutsky's Theorem is a special case of the CMT, it is a useful statement as it focuses on the most common applications – addition, multiplication, and division.

Despite the fact that the plug-in estimator $\widehat{\beta}$ is a function of $\widehat{\mu}$ for which we have an asymptotic distribution, Theorem 5.10.1 does not directly give us an asymptotic distribution for $\widehat{\beta}$. This is because $\widehat{\beta} = g(\widehat{\mu})$ is written as a function of $\widehat{\mu}$, not of the standardized sequence $\sqrt{n}(\widehat{\mu} - \mu)$. We need an intermediate step – a first order Taylor series expansion. This step is so critical to statistical theory that it has its own name – **The Delta Method**.

Theorem 5.10.3 *Delta Method:*

If $\sqrt{n}(\widehat{\mu} - \mu) \xrightarrow{d} \xi$, where $g(\mathbf{u})$ is continuously differentiable in a neighborhood of μ then as $n \rightarrow \infty$

$$\sqrt{n}(g(\widehat{\mu}) - g(\mu)) \xrightarrow{d} \mathbf{G}'\xi \quad (5.10)$$

where $\mathbf{G}(\mathbf{u}) = \frac{\partial}{\partial \mathbf{u}} g(\mathbf{u})'$ and $\mathbf{G} = \mathbf{G}(\mu)$. In particular, if $\xi \sim N(\mathbf{0}, \mathbf{V})$ then as $n \rightarrow \infty$

$$\sqrt{n}(g(\widehat{\mu}) - g(\mu)) \xrightarrow{d} N(0, \mathbf{G}'\mathbf{V}\mathbf{G}). \quad (5.11)$$

The Delta Method allows us to complete our derivation of the asymptotic distribution of the estimator $\widehat{\beta}$ of β .

By combining Theorems 5.8.2 and 5.10.3 we can find the asymptotic distribution of the plug-in estimator $\widehat{\beta}$.

Theorem 5.10.4 *If \mathbf{y}_i are independent and identically distributed, $\boldsymbol{\mu} = \mathbb{E}\mathbf{h}(\mathbf{y})$, $\beta = \mathbf{g}(\boldsymbol{\mu})$, $\mathbb{E}\|\mathbf{h}(\mathbf{y})\|^2 < \infty$, and $\mathbf{G}(\mathbf{u}) = \frac{\partial}{\partial \mathbf{u}}\mathbf{g}(\mathbf{u})'$ is continuous in a neighborhood of $\boldsymbol{\mu}$, then for $\widehat{\beta} = \mathbf{g}\left(\frac{1}{n}\sum_{i=1}^n \mathbf{h}(\mathbf{y}_i)\right)$, as $n \rightarrow \infty$*

$$\sqrt{n}(\widehat{\beta} - \beta) \xrightarrow{d} \mathbf{N}(\mathbf{0}, \mathbf{G}'\mathbf{V}\mathbf{G})$$

where $\mathbf{V} = \mathbb{E}((\mathbf{h}(\mathbf{y}) - \boldsymbol{\mu})(\mathbf{h}(\mathbf{y}) - \boldsymbol{\mu})')$ and $\mathbf{G} = \mathbf{G}(\boldsymbol{\mu})$.

Theorem 5.9.2 established the consistency of $\widehat{\beta}$ for β , and Theorem 5.10.4 established its asymptotic normality. It is instructive to compare the conditions required for these results. Consistency required that $\mathbf{h}(\mathbf{y})$ have a finite mean, while asymptotic normality requires that this variable have a finite variance. Consistency required that $\mathbf{g}(\mathbf{u})$ be continuous, while asymptotic normality required that $\mathbf{g}(\mathbf{u})$ be continuously differentiable.

5.11 Stochastic Order Symbols

It is convenient to have simple symbols for random variables and vectors which converge in probability to zero or are stochastically bounded. In this section we introduce some of the most commonly found notation.

It might be useful to review the common notation for non-random convergence and boundedness. Let x_n and a_n , $n = 1, 2, \dots$, be non-random sequences. The notation

$$x_n = o(1)$$

(pronounced “small oh-one”) is equivalent to $x_n \rightarrow 0$ as $n \rightarrow \infty$. The notation

$$x_n = o(a_n)$$

is equivalent to $a_n^{-1}x_n \rightarrow 0$ as $n \rightarrow \infty$. The notation

$$x_n = O(1)$$

(pronounced “big oh-one”) means that x_n is bounded uniformly in n : there exists an $M < \infty$ such that $|x_n| \leq M$ for all n . The notation

$$x_n = O(a_n)$$

is equivalent to $a_n^{-1}x_n = O(1)$.

We now introduce similar concepts for sequences of random variables. Let z_n and a_n , $n = 1, 2, \dots$ be sequences of random variables. (In most applications, a_n is non-random.) The notation

$$z_n = o_p(1)$$

(“small oh-P-one”) means that $z_n \xrightarrow{p} 0$ as $n \rightarrow \infty$. We also write

$$z_n = o_p(a_n)$$

if $a_n^{-1}z_n = o_p(1)$. For example, for any consistent estimator $\widehat{\beta}$ for β we can write

$$\widehat{\beta} = \beta + o_p(1).$$

Similarly, the notation $z_n = O_p(1)$ (“big oh-P-one”) means that z_n is bounded in probability. Precisely, for any $\varepsilon > 0$ there is a constant $M_\varepsilon < \infty$ such that

$$\limsup_{n \rightarrow \infty} \Pr(|z_n| > M_\varepsilon) \leq \varepsilon.$$

Furthermore, we write

$$z_n = O_p(a_n)$$

if $a_n^{-1}z_n = O_p(1)$.

$O_p(1)$ is weaker than $o_p(1)$ in the sense that $z_n = o_p(1)$ implies $z_n = O_p(1)$ but not the reverse. However, if $z_n = O_p(a_n)$ then $z_n = o_p(b_n)$ for any b_n such that $a_n/b_n \rightarrow 0$.

If a random vector converges in distribution $z_n \xrightarrow{d} z$ (for example, if $z \sim N(\mathbf{0}, \mathbf{V})$) then $z_n = O_p(1)$. It follows that for estimators $\widehat{\beta}$ which satisfy the convergence of Theorem 5.10.4 then we can write

$$\widehat{\beta} = \beta + O_p(n^{-1/2}).$$

In words, this statement says that the estimator $\widehat{\beta}$ equals the true coefficient β plus a random component which is shrinking to zero at the rate $n^{-1/2}$.

Another useful observation is that a random sequence with a bounded moment is stochastically bounded.

Theorem 5.11.1 *If z_n is a random vector which satisfies*

$$\mathbb{E} \|z_n\|^\delta = O(a_n)$$

for some sequence a_n and $\delta > 0$, then

$$z_n = O_p(a_n^{1/\delta}).$$

Similarly, $\mathbb{E} \|z_n\|^\delta = o(a_n)$ implies $z_n = o_p(a_n^{1/\delta})$.

This can be shown using Markov’s inequality (B.21). The assumptions imply that there is some $M < \infty$ such that $\mathbb{E} \|z_n\|^\delta \leq Ma_n$ for all n . For any ε set $B = \left(\frac{M}{\varepsilon}\right)^{1/\delta}$. Then

$$\Pr\left(a_n^{-1/\delta} \|z_n\| > B\right) = \Pr\left(\|z_n\|^\delta > \frac{Ma_n}{\varepsilon}\right) \leq \frac{\varepsilon}{Ma_n} \mathbb{E} \|z_n\|^\delta \leq \varepsilon$$

as required.

There are many simple rules for manipulating $o_p(1)$ and $O_p(1)$ sequences which can be deduced from the continuous mapping theorem or Slutsky’s Theorem. For example,

$$\begin{aligned} o_p(1) + o_p(1) &= o_p(1) \\ o_p(1) + O_p(1) &= O_p(1) \\ O_p(1) + O_p(1) &= O_p(1) \\ o_p(1)o_p(1) &= o_p(1) \\ o_p(1)O_p(1) &= o_p(1) \\ O_p(1)O_p(1) &= O_p(1) \end{aligned}$$

5.12 Uniform Stochastic Bounds*

For some applications it can be useful to obtain the stochastic order of the random variable

$$\max_{1 \leq i \leq n} |y_i|.$$

This is the magnitude of the largest observation in the sample $\{y_1, \dots, y_n\}$. If the support of the distribution of y_i is unbounded, then as the sample size n increases, the largest observation will also tend to increase. It turns out that there is a simple characterization.

Theorem 5.12.1 *Assume y_i are independent and identically distributed. If $\mathbb{E} |y|^r < \infty$, then as $n \rightarrow \infty$*

$$n^{-1/r} \max_{1 \leq i \leq n} |y_i| \xrightarrow{p} 0. \quad (5.12)$$

If $\mathbb{E} \exp(ty) < \infty$ for all $t < \infty$, then

$$(\log n)^{-1} \max_{1 \leq i \leq n} |y_i| \xrightarrow{p} 0. \quad (5.13)$$

The proof of Theorem 5.12.1 is presented in Section 5.14.

Equivalently, (5.12) can be written as

$$\max_{1 \leq i \leq n} |y_i| = o_p(n^{1/r}) \quad (5.14)$$

and (5.13) as

$$\max_{1 \leq i \leq n} |y_i| = o_p(\log n). \quad (5.15)$$

Equation (5.12) says that if y has r finite moments, then the largest observation will diverge at a rate slower than $n^{1/r}$. As r increases this rate decreases. Equation (5.13) shows that if we strengthen this to y having all finite moments and a finite moment generating function (for example, if y is normally distributed) then the largest observation will diverge slower than $\log n$. Thus the higher the moments, the slower the rate of divergence.

To simplify the notation, we write (5.14) as $y_i = o_p(n^{1/r})$ uniformly in $1 \leq i \leq n$, and similarly (5.15) as $y_i = o_p(\log n)$, uniformly in $1 \leq i \leq n$. It is important to understand when the O_p or o_p symbols are applied to subscript i random variables whether the convergence is pointwise in i , or is uniform in i in the sense of (5.14)-(5.15).

Theorem 5.12.1 applies to random vectors. If $\mathbb{E} \|\mathbf{y}\|^r < \infty$ then

$$\max_{1 \leq i \leq n} \|\mathbf{y}_i\| = o_p(n^{1/r}),$$

and if $\mathbb{E} \exp(\mathbf{t}'\mathbf{y}) < \infty$ for all $\|\mathbf{t}\| < \infty$ then

$$(\log n)^{-1} \max_{1 \leq i \leq n} \|\mathbf{y}_i\| \xrightarrow{p} 0. \quad (5.16)$$

5.13 Semiparametric Efficiency

In this section we argue that the sample mean $\hat{\boldsymbol{\mu}}$ and plug-in estimator $\hat{\boldsymbol{\beta}} = \mathbf{g}(\hat{\boldsymbol{\mu}})$ are efficient estimators of the parameters $\boldsymbol{\mu}$ and $\boldsymbol{\beta}$. Our demonstration is based on the rich but technically challenging theory of semiparametric efficiency bounds. An excellent accessible review has been provided by Newey (1990). We will also appeal to the asymptotic theory of maximum likelihood estimation (see Section B.11).

We start by examining the sample mean $\hat{\boldsymbol{\mu}}$, for the asymptotic efficiency of $\hat{\boldsymbol{\beta}}$ will follow from that of $\hat{\boldsymbol{\mu}}$.

Recall, we know that if $\mathbb{E}\|\mathbf{y}\|^2 < \infty$ then the sample mean has the asymptotic distribution $\sqrt{n}(\hat{\boldsymbol{\mu}} - \boldsymbol{\mu}) \xrightarrow{d} N(0, \mathbf{V})$. We want to know if $\hat{\boldsymbol{\mu}}$ is the best feasible estimator, or if there is another estimator with a smaller asymptotic variance. While it seems intuitively unlikely that another estimator could have a smaller asymptotic variance, how do we know that this is not the case?

When we ask if $\hat{\boldsymbol{\mu}}$ is the best estimator, we need to be clear about the class of models – the class of permissible distributions. For estimation of the mean $\boldsymbol{\mu}$ of the distribution of \mathbf{y} the broadest conceivable class is $\mathcal{L}_1 = \{F : \mathbb{E}\|\mathbf{y}\| < \infty\}$. This class is too broad for our current purposes, as $\hat{\boldsymbol{\mu}}$ is not asymptotically $N(0, \mathbf{V})$ for all $F \in \mathcal{L}_1$. A more realistic choice is $\mathcal{L}_2 = \{F : \mathbb{E}\|\mathbf{y}\|^2 < \infty\}$ – the class of finite-variance distributions. When we seek an efficient estimator of the mean $\boldsymbol{\mu}$ in the class of models \mathcal{L}_2 what we are seeking is the best estimator, given that all we know is that $F \in \mathcal{L}_2$.

To show that the answer is not immediately obvious, it might be helpful to review a setting where the sample mean is inefficient. Suppose that $y \in \mathbb{R}$ has the double exponential density $f(y | \mu) = 2^{-1/2} \exp(-|y - \mu| \sqrt{2})$. Since $\text{var}(y) = 1$ we see that the sample mean satisfies $\sqrt{n}(\hat{\mu} - \mu) \xrightarrow{d} N(0, 1)$. In this model the maximum likelihood estimator (MLE) $\tilde{\mu}$ for μ is the sample median. Recall from the theory of maximum likelihood that the MLE satisfies $\sqrt{n}(\tilde{\mu} - \mu) \xrightarrow{d} N(0, (\mathbb{E}S^2)^{-1})$ where $S = \frac{\partial}{\partial \mu} \log f(y | \mu) = -\sqrt{2} \text{sgn}(y - \mu)$ is the score. We can calculate that $\mathbb{E}S^2 = 2$ and thus conclude that $\sqrt{n}(\tilde{\mu} - \mu) \xrightarrow{d} N(0, 1/2)$. The asymptotic variance of the MLE is one-half that of the sample mean. Thus when the true density is known to be double exponential the sample mean is inefficient.

But the estimator which achieves this improved efficiency – the sample median – is not generically consistent for the population mean. It is inconsistent if the density is asymmetric or skewed. So the improvement comes at a great cost. Another way of looking at this is that the sample median is efficient in the class of densities $\{f(y | \mu) = 2^{-1/2} \exp(-|y - \mu| \sqrt{2})\}$ but unless it is known that this is the correct distribution class this knowledge is not very useful.

The relevant question is whether or not the sample mean is efficient when the form of the distribution is unknown. We call this setting **semiparametric** as the parameter of interest (the mean) is finite dimensional while the remaining features of the distribution are unspecified. In the semiparametric context an estimator is called **semiparametrically efficient** if it has the smallest asymptotic variance among all semiparametric estimators.

The mathematical trick is to reduce the semiparametric model to a set of parametric “submodels”. The Cramer-Rao variance bound can be found for each parametric submodel. The variance bound for the semiparametric model (the union of the submodels) is then defined as the supremum of the individual variance bounds.

Formally, suppose that the true density of \mathbf{y} is the unknown function $f(\mathbf{y})$ with mean $\boldsymbol{\mu} = \mathbb{E}\mathbf{y} = \int \mathbf{y}f(\mathbf{y})d\mathbf{y}$. A parametric submodel η for $f(\mathbf{y})$ is a density $f_\eta(\mathbf{y} | \boldsymbol{\theta})$ which is a smooth function of a parameter $\boldsymbol{\theta}$, and there is a true value $\boldsymbol{\theta}_0$ such that $f_\eta(\mathbf{y} | \boldsymbol{\theta}_0) = f(\mathbf{y})$. The index η indicates the submodels. The equality $f_\eta(\mathbf{y} | \boldsymbol{\theta}_0) = f(\mathbf{y})$ means that the submodel class passes through the true density, so the submodel is a true model. The class of submodels η and parameter $\boldsymbol{\theta}_0$ depend on the true density f . In the submodel $f_\eta(\mathbf{y} | \boldsymbol{\theta})$, the mean is $\boldsymbol{\mu}_\eta(\boldsymbol{\theta}) = \int \mathbf{y}f_\eta(\mathbf{y} | \boldsymbol{\theta}) d\mathbf{y}$ which varies with the parameter $\boldsymbol{\theta}$. Let $\eta \in \aleph$ be the class of all submodels for f .

Since each submodel η is parametric we can calculate the efficiency bound for estimation of $\boldsymbol{\mu}$ within this submodel. Specifically, given the density $f_\eta(\mathbf{y} \mid \boldsymbol{\theta})$ its likelihood score is

$$\mathbf{S}_\eta = \frac{\partial}{\partial \boldsymbol{\theta}} \log f_\eta(\mathbf{y} \mid \boldsymbol{\theta}),$$

so the Cramer-Rao lower bound for estimation of $\boldsymbol{\theta}$ is $(\mathbb{E} \mathbf{S}_\eta \mathbf{S}'_\eta)^{-1}$. Defining $\mathbf{M}_\eta = \frac{\partial}{\partial \boldsymbol{\theta}} \boldsymbol{\mu}_\eta(\boldsymbol{\theta}_0)'$, by Theorem B.11.5 the Cramer-Rao lower bound for estimation of $\boldsymbol{\mu}$ within the submodel η is $\mathbf{V}_\eta = \mathbf{M}'_\eta (\mathbb{E} \mathbf{S}_\eta \mathbf{S}'_\eta)^{-1} \mathbf{M}_\eta$.

As \mathbf{V}_η is the efficiency bound for the submodel class $f_\eta(\mathbf{y} \mid \boldsymbol{\theta})$, no estimator can have an asymptotic variance smaller than \mathbf{V}_η for any density $f_\eta(\mathbf{y} \mid \boldsymbol{\theta})$ in the submodel class, including the true density f . This is true for all submodels η . Thus the asymptotic variance of any semiparametric estimator cannot be smaller than \mathbf{V}_η for any conceivable submodel. Taking the supremum of the Cramer-Rao bounds lower from all conceivable submodels we define²

$$\overline{\mathbf{V}} = \sup_{\eta \in \mathfrak{N}} \mathbf{V}_\eta.$$

The asymptotic variance of any semiparametric estimator cannot be smaller than $\overline{\mathbf{V}}$, since it cannot be smaller than any individual \mathbf{V}_η . We call $\overline{\mathbf{V}}$ the **semiparametric asymptotic variance bound** or **semiparametric efficiency bound** for estimation of $\boldsymbol{\mu}$, as it is a lower bound on the asymptotic variance for any semiparametric estimator. If the asymptotic variance of a specific semiparametric estimator equals the bound $\overline{\mathbf{V}}$ we say that the estimator is **semiparametrically efficient**.

For many statistical problems it is quite challenging to calculate the semiparametric variance bound. However, in some cases there is a simple method to find the solution. Suppose that we can find a submodel η_0 whose Cramer-Rao lower bound satisfies $\mathbf{V}_{\eta_0} = \mathbf{V}_\mu$ where \mathbf{V}_μ is the asymptotic variance of a known semiparametric estimator. In this case, we can deduce that $\overline{\mathbf{V}} = \mathbf{V}_{\eta_0} = \mathbf{V}_\mu$. Otherwise there would exist another submodel η_1 whose Cramer-Rao lower bound satisfies $\mathbf{V}_{\eta_0} < \mathbf{V}_{\eta_1}$ but this would imply $\mathbf{V}_\mu < \mathbf{V}_{\eta_1}$ which contradicts the Cramer-Rao Theorem.

We now find this submodel for the sample mean $\hat{\boldsymbol{\mu}}$. Our goal is to find a parametric submodel whose Cramer-Rao bound for $\boldsymbol{\mu}$ is \mathbf{V} . This can be done by creating a tilted version of the true density. Consider the parametric submodel

$$f_\eta(\mathbf{y} \mid \boldsymbol{\theta}) = f(\mathbf{y}) (1 + \boldsymbol{\theta}' \mathbf{V}^{-1} (\mathbf{y} - \boldsymbol{\mu})) \quad (5.17)$$

where $f(\mathbf{y})$ is the true density and $\boldsymbol{\mu} = \mathbb{E} \mathbf{y}$. Note that

$$\int f_\eta(\mathbf{y} \mid \boldsymbol{\theta}) d\mathbf{y} = \int f(\mathbf{y}) d\mathbf{y} + \boldsymbol{\theta}' \mathbf{V}^{-1} \int f(\mathbf{y}) (\mathbf{y} - \boldsymbol{\mu}) d\mathbf{y} = 1$$

and for all $\boldsymbol{\theta}$ close to zero $f_\eta(\mathbf{y} \mid \boldsymbol{\theta}) \geq 0$. Thus $f_\eta(\mathbf{y} \mid \boldsymbol{\theta})$ is a valid density function. It is a parametric submodel since $f_\eta(\mathbf{y} \mid \boldsymbol{\theta}_0) = f(\mathbf{y})$ when $\boldsymbol{\theta}_0 = 0$. This parametric submodel has the mean

$$\begin{aligned} \boldsymbol{\mu}(\boldsymbol{\theta}) &= \int \mathbf{y} f_\eta(\mathbf{y} \mid \boldsymbol{\theta}) d\mathbf{y} \\ &= \int \mathbf{y} f(\mathbf{y}) d\mathbf{y} + \int f(\mathbf{y}) \mathbf{y} (\mathbf{y} - \boldsymbol{\mu})' \mathbf{V}^{-1} \boldsymbol{\theta} d\mathbf{y} \\ &= \boldsymbol{\mu} + \boldsymbol{\theta} \end{aligned}$$

which is a smooth function of $\boldsymbol{\theta}$.

²It is not obvious that this supremum exists, as \mathbf{V}_η is a matrix so there is not a unique ordering of matrices. However, in many cases (including the ones we study) the supremum exists and is unique.

Since

$$\frac{\partial}{\partial \boldsymbol{\theta}} \log f_{\eta}(\mathbf{y} \mid \boldsymbol{\theta}) = \frac{\partial}{\partial \boldsymbol{\theta}} \log (1 + \boldsymbol{\theta}' \mathbf{V}^{-1} (\mathbf{y} - \boldsymbol{\mu})) = \frac{\mathbf{V}^{-1} (\mathbf{y} - \boldsymbol{\mu})}{1 + \boldsymbol{\theta}' \mathbf{V}^{-1} (\mathbf{y} - \boldsymbol{\mu})}$$

it follows that the score function for $\boldsymbol{\theta}$ is

$$\mathbf{S}_{\eta} = \frac{\partial}{\partial \boldsymbol{\theta}} \log f_{\eta}(\mathbf{y} \mid \boldsymbol{\theta}_0) = \mathbf{V}^{-1} (\mathbf{y} - \boldsymbol{\mu}). \quad (5.18)$$

By Theorem B.11.3 the Cramer-Rao lower bound for $\boldsymbol{\theta}$ is

$$(\mathbb{E} (\mathbf{S}_{\eta} \mathbf{S}'_{\eta}))^{-1} = (\mathbf{V}^{-1} \mathbb{E} ((\mathbf{y} - \boldsymbol{\mu})(\mathbf{y} - \boldsymbol{\mu})') \mathbf{V}^{-1})^{-1} = \mathbf{V}. \quad (5.19)$$

The Cramer-Rao lower bound for $\boldsymbol{\mu}(\boldsymbol{\theta}) = \boldsymbol{\mu} + \boldsymbol{\theta}$ is also \mathbf{V} , and this equals the asymptotic variance of the moment estimator $\hat{\boldsymbol{\mu}}$. This was what we set out to show.

In summary, we have shown that in the submodel (5.17) the Cramer-Rao lower bound for estimation of $\boldsymbol{\mu}$ is \mathbf{V} which equals the asymptotic variance of the sample mean. This establishes the following result.

Proposition 5.13.1 *In the class of distributions $F \in \mathcal{L}_2$, the semiparametric variance bound for estimation of $\boldsymbol{\mu}$ is $\mathbf{V} = \text{var}(y_i)$, and the sample mean $\hat{\boldsymbol{\mu}}$ is a semiparametrically efficient estimator of the population mean $\boldsymbol{\mu}$.*

We call this result a proposition rather than a theorem as we have not attended to the regularity conditions.

It is a simple matter to extend this result to the plug-in estimator $\hat{\boldsymbol{\beta}} = \mathbf{g}(\hat{\boldsymbol{\mu}})$. We know from Theorem 5.10.4 that if $\mathbb{E} \|\mathbf{y}\|^2 < \infty$ and $\mathbf{g}(\mathbf{u})$ is continuously differentiable at $\mathbf{u} = \boldsymbol{\mu}$ then the plug-in estimator has the asymptotic distribution $\sqrt{n} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \xrightarrow{d} \text{N}(0, \mathbf{G}' \mathbf{V} \mathbf{G})$. We therefore consider the class of distributions

$$\mathcal{L}_2(\mathbf{g}) = \left\{ F : \mathbb{E} \|\mathbf{y}\|^2 < \infty, \mathbf{g}(\mathbf{u}) \text{ is continuously differentiable at } \mathbf{u} = \mathbb{E} \mathbf{y} \right\}.$$

For example, if $\boldsymbol{\beta} = \boldsymbol{\mu}_1 / \boldsymbol{\mu}_2$ where $\boldsymbol{\mu}_1 = \mathbb{E} y_1$ and $\boldsymbol{\mu}_2 = \mathbb{E} y_2$ then $\mathcal{L}_2(\mathbf{g}) = \{F : \mathbb{E} y_1^2 < \infty, \mathbb{E} y_2^2 < \infty, \text{ and } \mathbb{E} y_2 \neq 0\}$.

For any submodel η the Cramer-Rao lower bound for estimation of $\boldsymbol{\beta} = \mathbf{g}(\boldsymbol{\mu})$ is $\mathbf{G}' \mathbf{V}_{\eta} \mathbf{G}$ by Theorem B.11.5. For the submodel (5.17) this bound is $\mathbf{G}' \mathbf{V} \mathbf{G}$ which equals the asymptotic variance of $\hat{\boldsymbol{\beta}}$ from Theorem 5.10.4. Thus $\hat{\boldsymbol{\beta}}$ is semiparametrically efficient.

Proposition 5.13.2 *In the class of distributions $F \in \mathcal{L}_2(\mathbf{g})$ the semiparametric variance bound for estimation of $\boldsymbol{\beta} = \mathbf{g}(\boldsymbol{\mu})$ is $\mathbf{G}' \mathbf{V} \mathbf{G}$, and the plug-in estimator $\hat{\boldsymbol{\beta}} = \mathbf{g}(\hat{\boldsymbol{\mu}})$ is a semiparametrically efficient estimator of $\boldsymbol{\beta}$.*

The result in Proposition 5.13.2 is quite general. Smooth functions of sample moments are efficient estimators for their population counterparts. This is a very powerful result, as most econometric estimators can be written (or approximated) as smooth functions of sample means.

5.14 Technical Proofs*

In this section we provide proofs of some of the more technical points in the chapter. These proofs may only be of interest to more mathematically inclined.

Proof of Theorem 5.4.2: Without loss of generality, we can assume $\mathbb{E}(y_i) = 0$ by recentering y_i on its expectation.

We need to show that for all $\delta > 0$ and $\eta > 0$ there is some $N < \infty$ so that for all $n \geq N$, $\Pr(|\bar{y}| > \delta) \leq \eta$. Fix δ and η . Set $\varepsilon = \delta\eta/3$. Pick $C < \infty$ large enough so that

$$\mathbb{E}(|y_i| 1(|y_i| > C)) \leq \varepsilon \quad (5.20)$$

(where $1(\cdot)$ is the indicator function) which is possible since $\mathbb{E}|y_i| < \infty$. Define the random variables

$$\begin{aligned} w_i &= y_i 1(|y_i| \leq C) - \mathbb{E}(y_i 1(|y_i| \leq C)) \\ z_i &= y_i 1(|y_i| > C) - \mathbb{E}(y_i 1(|y_i| > C)) \end{aligned}$$

so that

$$\bar{y} = \bar{w} + \bar{z}$$

and

$$\mathbb{E}|\bar{y}| \leq \mathbb{E}|\bar{w}| + \mathbb{E}|\bar{z}|. \quad (5.21)$$

We now show that sum of the expectations on the right-hand-side can be bounded below 3ε .

First, by the Triangle Inequality (A.12) and the Expectation Inequality (B.15),

$$\begin{aligned} \mathbb{E}|z_i| &= \mathbb{E}|y_i 1(|y_i| > C) - \mathbb{E}(y_i 1(|y_i| > C))| \\ &\leq \mathbb{E}|y_i 1(|y_i| > C)| + |\mathbb{E}(y_i 1(|y_i| > C))| \\ &\leq 2\mathbb{E}|y_i 1(|y_i| > C)| \\ &\leq 2\varepsilon, \end{aligned} \quad (5.22)$$

and thus by the Triangle Inequality (A.12) and (5.22)

$$\mathbb{E}|\bar{z}| = \mathbb{E}\left|\frac{1}{n} \sum_{i=1}^n z_i\right| \leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}|z_i| \leq 2\varepsilon. \quad (5.23)$$

Second, by a similar argument

$$\begin{aligned} |w_i| &= |y_i 1(|y_i| \leq C) - \mathbb{E}(y_i 1(|y_i| \leq C))| \\ &\leq |y_i 1(|y_i| \leq C)| + |\mathbb{E}(y_i 1(|y_i| \leq C))| \\ &\leq 2|y_i 1(|y_i| \leq C)| \\ &\leq 2C \end{aligned} \quad (5.24)$$

where the final inequality is (5.20). Then by Jensen's Inequality (B.12), the fact that the w_i are iid and mean zero, and (5.24),

$$(\mathbb{E}|\bar{w}|)^2 \leq \mathbb{E}|\bar{w}|^2 = \frac{\mathbb{E}w_i^2}{n} = \frac{4C^2}{n} \leq \varepsilon^2 \quad (5.25)$$

the final inequality holding for $n \geq 4C^2/\varepsilon^2 = 36C^2/\delta^2\eta^2$. Equations (5.21), (5.23) and (5.25) together show that

$$\mathbb{E}|\bar{y}| \leq 3\varepsilon^2 \quad (5.26)$$

as desired.

Finally, by Markov's Inequality (B.21) and (5.26),

$$\Pr(|\bar{y}| > \delta) \leq \frac{\mathbb{E}|\bar{y}|}{\delta} \leq \frac{3\varepsilon}{\delta} = \eta,$$

the final equality by the definition of ε . We have shown that for any $\delta > 0$ and $\eta > 0$ then for all $n \geq 36C^2/\delta^2\eta^2$, $\Pr(|\bar{y}| > \delta) \leq \eta$, as needed. ■

Proof of Theorem 5.6.1: By Loève's c_r Inequality (A.19)

$$\|\mathbf{y}\| = \left(\sum_{j=1}^m y_j^2 \right)^{1/2} \leq \sum_{j=1}^m |y_j|.$$

Thus if $\mathbb{E}|y_j| < \infty$ for $j = 1, \dots, m$, then

$$\mathbb{E}\|\mathbf{y}\| \leq \sum_{j=1}^m \mathbb{E}|y_j| < \infty.$$

For the reverse inequality, the Euclidean norm of a vector is larger than the length of any individual component, so for any j , $|y_j| \leq \|\mathbf{y}\|$. Thus, if $\mathbb{E}\|\mathbf{y}\| < \infty$, then $\mathbb{E}|y_j| < \infty$ for $j = 1, \dots, m$. ■

Proof of Theorem 5.7.1: The moment bound $\mathbb{E}\mathbf{y}'_i \mathbf{y}_i < \infty$ is sufficient to guarantee that the elements of $\boldsymbol{\mu}$ and \mathbf{V} are well defined and finite. Without loss of generality, it is sufficient to consider the case $\boldsymbol{\mu} = \mathbf{0}$.

Our proof method is to calculate the characteristic function of $\sqrt{n}\bar{\mathbf{y}}_n$ and show that it converges pointwise to the characteristic function of $N(\mathbf{0}, \mathbf{V})$. By Lévy's Continuity Theorem (see Van der Vaart (2008) Theorem 2.13) this is sufficient to establish that $\sqrt{n}\bar{\mathbf{y}}_n$ converges in distribution to $N(\mathbf{0}, \mathbf{V})$.

For $\boldsymbol{\lambda} \in \mathbb{R}^m$, let $C(\boldsymbol{\lambda}) = \mathbb{E} \exp(i\boldsymbol{\lambda}' \mathbf{y}_i)$ denote the characteristic function of \mathbf{y}_i and set $c(\boldsymbol{\lambda}) = \log C(\boldsymbol{\lambda})$. Since \mathbf{y}_i has two finite moments the first and second derivatives of $C(\boldsymbol{\lambda})$ are continuous in $\boldsymbol{\lambda}$. They are

$$\begin{aligned} \frac{\partial}{\partial \boldsymbol{\lambda}} C(\boldsymbol{\lambda}) &= i\mathbb{E}(\mathbf{y}_i \exp(i\boldsymbol{\lambda}' \mathbf{y}_i)) \\ \frac{\partial^2}{\partial \boldsymbol{\lambda} \partial \boldsymbol{\lambda}'} C(\boldsymbol{\lambda}) &= i^2 \mathbb{E}(\mathbf{y}_i \mathbf{y}_i' \exp(i\boldsymbol{\lambda}' \mathbf{y}_i)). \end{aligned}$$

When evaluated at $\boldsymbol{\lambda} = \mathbf{0}$

$$\begin{aligned} C(\mathbf{0}) &= 1 \\ \frac{\partial}{\partial \boldsymbol{\lambda}} C(\mathbf{0}) &= i\mathbb{E}(\mathbf{y}_i) = \mathbf{0} \\ \frac{\partial^2}{\partial \boldsymbol{\lambda} \partial \boldsymbol{\lambda}'} C(\mathbf{0}) &= -\mathbb{E}(\mathbf{y}_i \mathbf{y}_i') = -\mathbf{V}. \end{aligned}$$

Furthermore,

$$\begin{aligned} c_{\boldsymbol{\lambda}}(\boldsymbol{\lambda}) &= \frac{\partial}{\partial \boldsymbol{\lambda}} c(\boldsymbol{\lambda}) = C(\boldsymbol{\lambda})^{-1} \frac{\partial}{\partial \boldsymbol{\lambda}} C(\boldsymbol{\lambda}) \\ c_{\boldsymbol{\lambda}\boldsymbol{\lambda}}(\boldsymbol{\lambda}) &= \frac{\partial^2}{\partial \boldsymbol{\lambda} \partial \boldsymbol{\lambda}'} c(\boldsymbol{\lambda}) = C(\boldsymbol{\lambda})^{-1} \frac{\partial^2}{\partial \boldsymbol{\lambda} \partial \boldsymbol{\lambda}'} C(\boldsymbol{\lambda}) - C(\boldsymbol{\lambda})^{-2} \frac{\partial}{\partial \boldsymbol{\lambda}} C(\boldsymbol{\lambda}) \frac{\partial}{\partial \boldsymbol{\lambda}'} C(\boldsymbol{\lambda}) \end{aligned}$$

so when evaluated at $\boldsymbol{\lambda} = \mathbf{0}$

$$\begin{aligned} c(\mathbf{0}) &= 0 \\ c_{\boldsymbol{\lambda}}(\mathbf{0}) &= \mathbf{0} \\ c_{\boldsymbol{\lambda}\boldsymbol{\lambda}}(\mathbf{0}) &= -\mathbf{V}. \end{aligned}$$

By a second-order Taylor series expansion of $c(\boldsymbol{\lambda})$ about $\boldsymbol{\lambda} = \mathbf{0}$,

$$c(\boldsymbol{\lambda}) = c(\mathbf{0}) + c_{\lambda}(\mathbf{0})' \boldsymbol{\lambda} + \frac{1}{2} \boldsymbol{\lambda}' c_{\lambda\lambda}(\boldsymbol{\lambda}^*) \boldsymbol{\lambda} = \frac{1}{2} \boldsymbol{\lambda}' c_{\lambda\lambda}(\boldsymbol{\lambda}^*) \boldsymbol{\lambda} \quad (5.27)$$

where $\boldsymbol{\lambda}^*$ lies on the line segment joining $\mathbf{0}$ and $\boldsymbol{\lambda}$.

We now compute $C_n(\boldsymbol{\lambda}) = E \exp(i \boldsymbol{\lambda}' \sqrt{n} \bar{\mathbf{y}}_n)$, the characteristic function of $\sqrt{n} \bar{\mathbf{y}}_n$. By the properties of the exponential function, the independence of the \mathbf{y}_i , the definition of $c(\boldsymbol{\lambda})$ and (5.27)

$$\begin{aligned} \log C_n(\boldsymbol{\lambda}) &= \log \mathbb{E} \exp \left(i \frac{1}{\sqrt{n}} \sum_{i=1}^n \boldsymbol{\lambda}' \mathbf{y}_i \right) \\ &= \log \mathbb{E} \prod_{i=1}^n \exp \left(i \frac{1}{\sqrt{n}} \boldsymbol{\lambda}' \mathbf{y}_i \right) \\ &= \log \prod_{i=1}^n \mathbb{E} \exp \left(i \frac{1}{\sqrt{n}} \boldsymbol{\lambda}' \mathbf{y}_i \right) \\ &= \sum_{i=1}^n \log \mathbb{E} \exp \left(i \frac{1}{\sqrt{n}} \boldsymbol{\lambda}' \mathbf{y}_i \right) \\ &= nc \left(\frac{\boldsymbol{\lambda}}{\sqrt{n}} \right) \\ &= \frac{1}{2} \boldsymbol{\lambda}' c_{\lambda\lambda}(\boldsymbol{\lambda}_n) \boldsymbol{\lambda} \end{aligned}$$

where $\boldsymbol{\lambda}_n$ lies on the line segment joining $\mathbf{0}$ and $\boldsymbol{\lambda}/\sqrt{n}$. Since $\boldsymbol{\lambda}_n \rightarrow \mathbf{0}$ and $c_{\lambda\lambda}(\boldsymbol{\lambda})$ is continuous, $c_{\lambda\lambda}(\boldsymbol{\lambda}_n) \rightarrow c_{\lambda\lambda}(\mathbf{0}) = -\mathbf{V}$. We thus find that as $n \rightarrow \infty$,

$$\log C_n(\boldsymbol{\lambda}) \rightarrow -\frac{1}{2} \boldsymbol{\lambda}' \mathbf{V} \boldsymbol{\lambda}$$

and

$$C_n(\boldsymbol{\lambda}) \rightarrow \exp \left(-\frac{1}{2} \boldsymbol{\lambda}' \mathbf{V} \boldsymbol{\lambda} \right)$$

which is the characteristic function of the $N(\mathbf{0}, \mathbf{V})$ distribution. This completes the proof. \blacksquare

Proof of Theorem 5.9.1: Since \mathbf{g} is continuous at \mathbf{c} , for all $\varepsilon > 0$ we can find a $\delta > 0$ such that if $\|\mathbf{z}_n - \mathbf{c}\| < \delta$ then $\|\mathbf{g}(\mathbf{z}_n) - \mathbf{g}(\mathbf{c})\| \leq \varepsilon$. Recall that $A \subset B$ implies $\Pr(A) \leq \Pr(B)$. Thus $\Pr(\|\mathbf{g}(\mathbf{z}_n) - \mathbf{g}(\mathbf{c})\| \leq \varepsilon) \geq \Pr(\|\mathbf{z}_n - \mathbf{c}\| < \delta) \rightarrow 1$ as $n \rightarrow \infty$ by the assumption that $\mathbf{z}_n \xrightarrow{p} \mathbf{c}$. Hence $\mathbf{g}(\mathbf{z}_n) \xrightarrow{p} \mathbf{g}(\mathbf{c})$ as $n \rightarrow \infty$. \blacksquare

Proof of Theorem 5.10.3: By a vector Taylor series expansion, for each element of \mathbf{g} ,

$$g_j(\boldsymbol{\theta}_n) = g_j(\boldsymbol{\theta}) + g_{j\boldsymbol{\theta}}(\boldsymbol{\theta}_{j_n}^*) (\boldsymbol{\theta}_n - \boldsymbol{\theta})$$

where $\boldsymbol{\theta}_{j_n}^*$ lies on the line segment between $\boldsymbol{\theta}_n$ and $\boldsymbol{\theta}$ and therefore converges in probability to $\boldsymbol{\theta}$. It follows that $a_{jn} = g_{j\boldsymbol{\theta}}(\boldsymbol{\theta}_{j_n}^*) - g_{j\boldsymbol{\theta}} \xrightarrow{p} 0$. Stacking across elements of \mathbf{g} , we find

$$\sqrt{n} (\mathbf{g}(\boldsymbol{\theta}_n) - \mathbf{g}(\boldsymbol{\theta})) = (\mathbf{G} + \mathbf{a}_n)' \sqrt{n} (\boldsymbol{\theta}_n - \boldsymbol{\theta}) \xrightarrow{d} \mathbf{G}' \boldsymbol{\xi}. \quad (5.28)$$

The convergence is by Theorem 5.10.1, as $\mathbf{G} + \mathbf{a}_n \xrightarrow{d} \mathbf{G}$, $\sqrt{n} (\boldsymbol{\theta}_n - \boldsymbol{\theta}) \xrightarrow{d} \boldsymbol{\xi}$, and their product is continuous. This establishes (5.10)

When $\boldsymbol{\xi} \sim N(\mathbf{0}, \mathbf{V})$, the right-hand-side of (5.28) equals

$$\mathbf{G}' \boldsymbol{\xi} = \mathbf{G}' N(\mathbf{0}, \mathbf{V}) = N(\mathbf{0}, \mathbf{G}' \mathbf{V} \mathbf{G})$$

establishing (5.11). ■

Proof of Theorem 5.12.1: First consider (5.12). Take any $\delta > 0$. The event $\{\max_{1 \leq i \leq n} |y_i| > \delta n^{1/r}\}$ means that at least one of the $|y_i|$ exceeds $\delta n^{1/r}$, which is the same as the event $\bigcup_{i=1}^n \{|y_i| > \delta n^{1/r}\}$ or equivalently $\bigcup_{i=1}^n \{|y_i|^r > \delta^r n\}$. Since the probability of the union of events is smaller than the sum of the probabilities,

$$\begin{aligned} \Pr\left(n^{-1/r} \max_{1 \leq i \leq n} |y_i| > \delta\right) &= \Pr\left(\bigcup_{i=1}^n \{|y_i|^r > \delta^r n\}\right) \\ &\leq \sum_{i=1}^n \Pr(|y_i|^r > n\delta^r) \\ &\leq \frac{1}{n\delta^r} \sum_{i=1}^n \mathbb{E}(|y_i|^r \mathbf{1}(|y_i|^r > n\delta^r)) \\ &= \frac{1}{\delta^r} \mathbb{E}(|y_1|^r \mathbf{1}(|y_1|^r > n\delta^r)) \end{aligned}$$

where the second inequality is the strong form of Markov's inequality (Theorem B.22) and the final equality is since the y_i are iid. Since $\mathbb{E}|y|^r < \infty$ this final expectation converges to zero as $n \rightarrow \infty$. This is because

$$\mathbb{E}|y|^r = \int |y|^r dF(y) < \infty$$

implies

$$\mathbb{E}(|y|^r \mathbf{1}(|y|^r > c)) = \int_{|y|^r > c} |y|^r dF(y) \rightarrow 0 \quad (5.29)$$

as $c \rightarrow \infty$. This establishes (5.12).

Now consider (5.13). Take any $\delta > 0$ and set $t = 1/\delta$. By a similar calculation

$$\begin{aligned} \Pr\left((\log n)^{-1} \max_{1 \leq i \leq n} |y_i| > \delta\right) &= \Pr\left(\bigcup_{i=1}^n \{\exp |ty_i| > \exp(t\delta \log n)\}\right) \\ &\leq \sum_{i=1}^n \Pr(\exp |ty_i| > n) \\ &\leq \mathbb{E}(\exp |ty| \mathbf{1}(\exp |ty| > n)) \end{aligned}$$

where the second line uses $\exp(t\delta \log n) = \exp(\log n) = n$. The assumption $\mathbb{E} \exp(ty) < \infty$ means $\mathbb{E}(\exp |ty| \mathbf{1}(\exp |ty| > n)) \rightarrow 0$ as $n \rightarrow \infty$ by the same argument as in (5.29). This establishes (5.13). ■

Exercises

Exercise 5.1 For the following sequences, find the liminf, limsup and limit (if it exists) as $n \rightarrow \infty$

1. $a_n = 1/n$
2. $a_n = \sin\left(\frac{\pi}{2}n\right)$
3. $a_n = \frac{1}{n} \sin\left(\frac{\pi}{2}n\right)$

Exercise 5.2 A weighted sample mean takes the form $\bar{y}^* = \frac{1}{n} \sum_{i=1}^n w_i y_i$ for some non-negative constants w_i satisfying $\frac{1}{n} \sum_{i=1}^n w_i = 1$. Assume y_i is iid.

1. Show that \bar{y}^* is unbiased for $\mu = \mathbb{E}y_i$.
2. Calculate $\text{var}(\bar{y}^*)$.
3. Show that a sufficient condition for $\bar{y}^* \xrightarrow{p} \mu$ is that $\frac{1}{n^2} \sum_{i=1}^n w_i^2 \rightarrow 0$.
4. Show that a sufficient condition for the condition in part 3 is $\max_{i \leq n} w_i = o(n)$.

Exercise 5.3 Take a random variable Z such that $\mathbb{E}Z = 0$ and $\text{var}(Z) = 1$. Use Chebyshev's inequality to find a δ such that $\Pr(|Z| > \delta) \leq 0.05$. Contrast this with the exact δ which solves $\Pr(|Z| > \delta) = 0.05$ when $Z \sim N(0, 1)$. Comment on the difference.

Exercise 5.4 Find the moment estimator $\hat{\mu}_3$ of $\mu_3 = \mathbb{E}y_i^3$ and show that $\sqrt{n}(\hat{\mu}_3 - \mu_3) \xrightarrow{d} N(0, v^2)$ for some v^2 . Write v^2 as a function of the moments of y_i .

Exercise 5.5 Suppose $z_n \xrightarrow{p} c$ as $n \rightarrow \infty$. Show that $z_n^2 \xrightarrow{p} c^2$ as $n \rightarrow \infty$ without using the definition of convergence in probability, without using the CMT.

Exercise 5.6 Suppose $\sqrt{n}(\hat{\mu} - \mu) \xrightarrow{d} N(0, v^2)$ and set $\beta = \mu^2$ and $\hat{\beta} = \hat{\mu}^2$.

1. Use the Delta Method to obtain an asymptotic distribution for $\sqrt{n}(\hat{\beta} - \beta)$.
2. Now suppose $\mu = 0$. Describe what happens to the asymptotic distribution from the previous part.
3. Improve on the previous answer. Under the assumption $\mu = 0$, find the asymptotic distribution for $n\hat{\beta} = n\hat{\mu}^2$.
4. Comment on the differences between the answers in parts 1 and 3.

Chapter 6

Asymptotic Theory for Least Squares

6.1 Introduction

It turns out that the asymptotic theory of least-squares estimation applies equally to the projection model and the linear CEF model, and therefore the results in this chapter will be stated for the broader projection model described in Section 2.18. Recall that the model is

$$y_i = \mathbf{x}'_i \boldsymbol{\beta} + e_i$$

for $i = 1, \dots, n$, where the linear projection $\boldsymbol{\beta}$ is

$$\boldsymbol{\beta} = (\mathbb{E}(\mathbf{x}_i \mathbf{x}'_i))^{-1} \mathbb{E}(\mathbf{x}_i y_i).$$

Some of the results of this section hold under random sampling (Assumption 1.5.1) and finite second moments (Assumption 2.18.1). We restate this condition here for clarity.

Assumption 6.1.1

1. *The observations (y_i, \mathbf{x}_i) , $i = 1, \dots, n$, are independent and identically distributed.*
2. $\mathbb{E}y^2 < \infty$.
3. $\mathbb{E} \|\mathbf{x}\|^2 < \infty$.
4. $\mathbf{Q}_{\mathbf{x}\mathbf{x}} = \mathbb{E}(\mathbf{x}\mathbf{x}')$ is positive definite.

Some of the results will require a strengthening to finite fourth moments.

Assumption 6.1.2 *In addition to Assumption 6.1.1, $\mathbb{E}y_i^4 < \infty$ and $\mathbb{E} \|\mathbf{x}_i\|^4 < \infty$.*

6.2 Consistency of Least-Squares Estimator

In this section we use the weak law of large numbers (WLLN, Theorem 5.4.2 and Theorem 5.6.2) and continuous mapping theorem (CMT, Theorem 5.9.1) to show that the least-squares estimator $\widehat{\boldsymbol{\beta}}$ is consistent for the projection coefficient $\boldsymbol{\beta}$.

This derivation is based on three key components. First, the OLS estimator can be written as a continuous function of a set of sample moments. Second, the WLLN shows that sample moments converge in probability to population moments. And third, the CMT states that continuous functions preserve convergence in probability. We now explain each step in brief and then in greater detail.

First, observe that the OLS estimator

$$\widehat{\boldsymbol{\beta}} = \left(\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \right)^{-1} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i y_i \right) = \widehat{\mathbf{Q}}_{xx}^{-1} \widehat{\mathbf{Q}}_{xy}$$

is a function of the sample moments $\widehat{\mathbf{Q}}_{xx} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i'$ and $\widehat{\mathbf{Q}}_{xy} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i y_i$.

Second, by an application of the WLLN these sample moments converge in probability to the population moments. Specifically, the fact that (y_i, \mathbf{x}_i) are mutually independent and identically distributed implies that any function of (y_i, \mathbf{x}_i) is iid, including $\mathbf{x}_i \mathbf{x}_i'$ and $\mathbf{x}_i y_i$. These variables also have finite expectations by Theorem 2.18.1.1. Under these conditions, the WLLN (Theorem 5.6.2) implies that as $n \rightarrow \infty$,

$$\widehat{\mathbf{Q}}_{xx} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \xrightarrow{p} \mathbb{E}(\mathbf{x}_i \mathbf{x}_i') = \mathbf{Q}_{xx} \quad (6.1)$$

and

$$\widehat{\mathbf{Q}}_{xy} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i y_i \xrightarrow{p} \mathbb{E}(\mathbf{x}_i y_i) = \mathbf{Q}_{xy}. \quad (6.2)$$

Third, the CMT (Theorem 5.9.1) allows us to combine these equations to show that $\widehat{\boldsymbol{\beta}}$ converges in probability to $\boldsymbol{\beta}$. Specifically, as $n \rightarrow \infty$,

$$\begin{aligned} \widehat{\boldsymbol{\beta}} &= \widehat{\mathbf{Q}}_{xx}^{-1} \widehat{\mathbf{Q}}_{xy} \\ &\xrightarrow{p} \mathbf{Q}_{xx}^{-1} \mathbf{Q}_{xy} \\ &= \boldsymbol{\beta}. \end{aligned} \quad (6.3)$$

We have shown that $\widehat{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\beta}$, as $n \rightarrow \infty$. In words, the OLS estimator converges in probability to the projection coefficient vector $\boldsymbol{\beta}$ as the sample size n gets large.

To fully understand the application of the CMT we walk through it in detail. We can write

$$\widehat{\boldsymbol{\beta}} = \mathbf{g}(\widehat{\mathbf{Q}}_{xx}, \widehat{\mathbf{Q}}_{xy})$$

where $\mathbf{g}(\mathbf{A}, \mathbf{b}) = \mathbf{A}^{-1} \mathbf{b}$ is a function of \mathbf{A} and \mathbf{b} . The function $\mathbf{g}(\mathbf{A}, \mathbf{b})$ is a continuous function of \mathbf{A} and \mathbf{b} at all values of the arguments such that \mathbf{A}^{-1} exists. Assumption 2.18.1 implies that \mathbf{Q}_{xx}^{-1} exists and thus $\mathbf{g}(\mathbf{A}, \mathbf{b})$ is continuous at $\mathbf{A} = \mathbf{Q}_{xx}$. This justifies the application of the CMT in (6.3).

For a slightly different demonstration of (6.3), recall that (4.7) implies that

$$\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta} = \widehat{\mathbf{Q}}_{xx}^{-1} \widehat{\mathbf{Q}}_{xe} \quad (6.4)$$

where

$$\widehat{\mathbf{Q}}_{xe} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i e_i.$$

The WLLN and (2.27) imply

$$\widehat{\mathbf{Q}}_{xe} \xrightarrow{p} \mathbb{E}(\mathbf{x}_i e_i) = \mathbf{0}. \quad (6.5)$$

Therefore

$$\begin{aligned} \widehat{\boldsymbol{\beta}} - \boldsymbol{\beta} &= \widehat{\mathbf{Q}}_{xx}^{-1} \widehat{\mathbf{Q}}_{xe} \\ &\xrightarrow{p} \mathbf{Q}_{xx}^{-1} \mathbf{0} \\ &= \mathbf{0} \end{aligned}$$

which is the same as $\widehat{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\beta}$.

Theorem 6.2.1 Consistency of Least-Squares

Under Assumption 6.1.1, $\widehat{\mathbf{Q}}_{xx} \xrightarrow{p} \mathbf{Q}_{xx}$, $\widehat{\mathbf{Q}}_{xy} \xrightarrow{p} \mathbf{Q}_{xy}$, $\widehat{\mathbf{Q}}_{xx}^{-1} \xrightarrow{p} \mathbf{Q}_{xx}^{-1}$, $\widehat{\mathbf{Q}}_{xe} \xrightarrow{p} \mathbf{0}$, and $\widehat{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\beta}$ as $n \rightarrow \infty$.

Theorem 6.2.1 states that the OLS estimator $\widehat{\boldsymbol{\beta}}$ converges in probability to $\boldsymbol{\beta}$ as n increases, and thus $\widehat{\boldsymbol{\beta}}$ is consistent for $\boldsymbol{\beta}$. In the stochastic order notation, Theorem 6.2.1 can be equivalently written as

$$\widehat{\boldsymbol{\beta}} = \boldsymbol{\beta} + o_p(1). \quad (6.6)$$

To illustrate the effect of sample size on the least-squares estimator consider the least-squares regression

$$\ln(\text{Wage}_i) = \beta_1 \text{Education}_i + \beta_2 \text{Experience}_i + \beta_3 \text{Experience}_i^2 + \beta_4 + e_i.$$

We use the sample of 24,344 white men from the March 2009 CPS. Randomly sorting the observations, and sequentially estimating the model by least-squares, starting with the first 5 observations, and continuing until the full sample is used, the sequence of estimates are displayed in Figure 6.1. You can see how the least-squares estimate changes with the sample size, but as the number of observations increases it settles down to the full-sample estimate $\widehat{\beta}_1 = 0.114$.

6.3 Asymptotic Normality

We started this chapter discussing the need for an approximation to the distribution of the OLS estimator $\widehat{\boldsymbol{\beta}}$. In Section 6.2 we showed that $\widehat{\boldsymbol{\beta}}$ converges in probability to $\boldsymbol{\beta}$. Consistency is a good first step, but in itself does not describe the distribution of the estimator. In this section we derive an approximation typically called the **asymptotic distribution**.

The derivation starts by writing the estimator as a function of sample moments. One of the moments must be written as a sum of zero-mean random vectors and normalized so that the central limit theorem can be applied. The steps are as follows.

Take equation (6.4) and multiply it by \sqrt{n} . This yields the expression

$$\sqrt{n}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}) = \left(\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \right)^{-1} \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n \mathbf{x}_i e_i \right). \quad (6.7)$$

This shows that the normalized and centered estimator $\sqrt{n}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})$ is a function of the sample average $\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i'$ and the normalized sample average $\frac{1}{\sqrt{n}} \sum_{i=1}^n \mathbf{x}_i e_i$. Furthermore, the latter has mean zero so the central limit theorem (CLT, Theorem 5.7.1) applies.

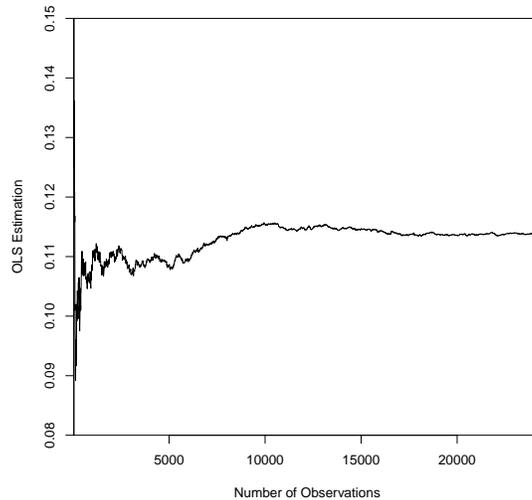


Figure 6.1: The least-squares estimator $\hat{\beta}_1$ as a function of sample size n

The product $\mathbf{x}_i e_i$ is iid (since the observations are iid) and mean zero (since $\mathbb{E}(\mathbf{x}_i e_i) = \mathbf{0}$). Define the $k \times k$ covariance matrix

$$\mathbf{\Omega} = \mathbb{E}(\mathbf{x}_i \mathbf{x}_i' e_i^2). \quad (6.8)$$

We require the elements of $\mathbf{\Omega}$ to be finite, written $\mathbf{\Omega} < \infty$. By the Expectation Inequality (B.15),

$$\|\mathbf{\Omega}\| \leq \mathbb{E} \|\mathbf{x}_i \mathbf{x}_i' e_i^2\| = \mathbb{E} \|\mathbf{x}_i e_i\|^2 = \mathbb{E} \|\mathbf{x}_i\|^2 e_i^2$$

or equivalently that $\mathbb{E} \|\mathbf{x}_i e_i\|^2 < \infty$. Using $\|\mathbf{x}_i e_i\|^2 = \|\mathbf{x}_i\|^2 e_i^2$ and the Cauchy-Schwarz Inequality (B.17),

$$\|\mathbf{\Omega}\| \leq \mathbb{E} \|\mathbf{x}_i \mathbf{x}_i' e_i^2\| = \mathbb{E} \|\mathbf{x}_i e_i\|^2 = \mathbb{E} \left(\|\mathbf{x}_i\|^2 e_i^2 \right) \leq \left(\mathbb{E} \|\mathbf{x}_i\|^4 \right)^{1/2} \left(\mathbb{E} e_i^4 \right)^{1/2} \quad (6.9)$$

which is finite if \mathbf{x}_i and e_i have finite fourth moments. As e_i is a linear combination of y_i and \mathbf{x}_i , it is sufficient that the observables have finite fourth moments (Theorem 2.18.1.6). We can then apply the CLT (Theorem 5.7.1).

Theorem 6.3.1 Under Assumption 6.1.2,

$$\|\mathbf{\Omega}\| \leq \mathbb{E} \|\mathbf{x}_i e_i\|^2 < \infty \quad (6.10)$$

and

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n \mathbf{x}_i e_i \xrightarrow{d} \mathbf{N}(\mathbf{0}, \mathbf{\Omega}) \quad (6.11)$$

as $n \rightarrow \infty$.

Putting together (6.1), (6.7), and (6.11),

$$\begin{aligned} \sqrt{n} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) &\xrightarrow{d} \mathbf{Q}_{xx}^{-1} \mathbf{N}(\mathbf{0}, \mathbf{\Omega}) \\ &= \mathbf{N}(\mathbf{0}, \mathbf{Q}_{xx}^{-1} \mathbf{\Omega} \mathbf{Q}_{xx}^{-1}) \end{aligned}$$

as $n \rightarrow \infty$, where the final equality follows from the property that linear combinations of normal vectors are also normal (Theorem B.9.1).

We have derived the asymptotic normal approximation to the distribution of the least-squares estimator.

Theorem 6.3.2 Asymptotic Normality of Least-Squares Estimator

Under Assumption 6.1.2, as $n \rightarrow \infty$

$$\sqrt{n}(\hat{\beta} - \beta) \xrightarrow{d} N(\mathbf{0}, \mathbf{V}_\beta)$$

where

$$\mathbf{V}_\beta = \mathbf{Q}_{xx}^{-1} \mathbf{\Omega} \mathbf{Q}_{xx}^{-1}, \quad (6.12)$$

$$\mathbf{Q}_{xx} = \mathbb{E}(\mathbf{x}_i \mathbf{x}_i'), \text{ and } \mathbf{\Omega} = \mathbb{E}(\mathbf{x}_i \mathbf{x}_i' e_i^2).$$

In the stochastic order notation, Theorem 6.3.2 implies that

$$\hat{\beta} = \beta + O_p(n^{-1/2}) \quad (6.13)$$

which is stronger than (6.6).

The matrix $\mathbf{V}_\beta = \mathbf{Q}_{xx}^{-1} \mathbf{\Omega} \mathbf{Q}_{xx}^{-1}$ is the variance of the asymptotic distribution of $\sqrt{n}(\hat{\beta} - \beta)$. Consequently, \mathbf{V}_β is often referred to as the **asymptotic covariance matrix** of $\hat{\beta}$. The expression $\mathbf{V}_\beta = \mathbf{Q}_{xx}^{-1} \mathbf{\Omega} \mathbf{Q}_{xx}^{-1}$ is called a **sandwich** form, as the matrix $\mathbf{\Omega}$ is sandwiched between two copies of \mathbf{Q}_{xx}^{-1} .

It is useful to compare the variance of the asymptotic distribution given in (6.12) and the finite-sample conditional variance in the CEF model as given in (4.12):

$$\mathbf{V}_{\hat{\beta}} = \text{var}(\hat{\beta} | \mathbf{X}) = (\mathbf{X}'\mathbf{X})^{-1} (\mathbf{X}'\mathbf{D}\mathbf{X}) (\mathbf{X}'\mathbf{X})^{-1}. \quad (6.14)$$

Notice that $\mathbf{V}_{\hat{\beta}}$ is the exact conditional variance of $\hat{\beta}$ and \mathbf{V}_β is the asymptotic variance of $\sqrt{n}(\hat{\beta} - \beta)$. Thus \mathbf{V}_β should be (roughly) n times as large as $\mathbf{V}_{\hat{\beta}}$, or $\mathbf{V}_\beta \approx n\mathbf{V}_{\hat{\beta}}$. Indeed, multiplying (6.14) by n and distributing, we find

$$n\mathbf{V}_{\hat{\beta}} = \left(\frac{1}{n}\mathbf{X}'\mathbf{X}\right)^{-1} \left(\frac{1}{n}\mathbf{X}'\mathbf{D}\mathbf{X}\right) \left(\frac{1}{n}\mathbf{X}'\mathbf{X}\right)^{-1}$$

which looks like an estimator of \mathbf{V}_β . Indeed, as $n \rightarrow \infty$

$$n\mathbf{V}_{\hat{\beta}} \xrightarrow{p} \mathbf{V}_\beta.$$

The expression $\mathbf{V}_{\hat{\beta}}$ is useful for practical inference (such as computation of standard errors and tests) since it is the variance of the estimator $\hat{\beta}$, while \mathbf{V}_β is useful for asymptotic theory as it is well defined in the limit as n goes to infinity. We will make use of both symbols and it will be advisable to adhere to this convention.

There is a special case where $\mathbf{\Omega}$ and \mathbf{V}_β simplify. We say that e_i is a **Homoskedastic Projection Error** when

$$\text{cov}(\mathbf{x}_i \mathbf{x}_i', e_i^2) = \mathbf{0}. \quad (6.15)$$

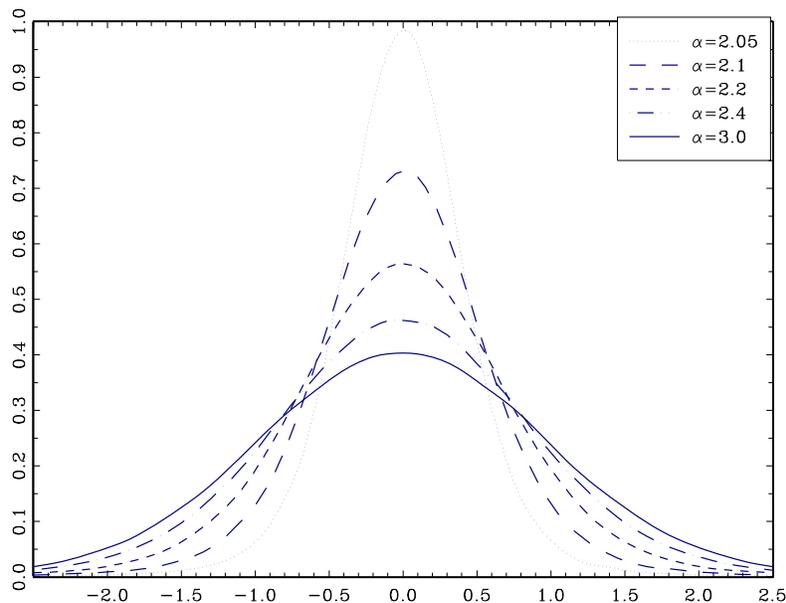


Figure 6.2: Density of Normalized OLS estimator with Double Pareto Error

Condition (6.15) holds in the homoskedastic linear regression model, but is somewhat broader. Under (6.15) the asymptotic variance formulae simplify as

$$\mathbf{\Omega} = \mathbb{E}(\mathbf{x}_i \mathbf{x}_i') \mathbb{E}(e_i^2) = \mathbf{Q}_{xx} \sigma^2 \quad (6.16)$$

$$\mathbf{V}_\beta = \mathbf{Q}_{xx}^{-1} \mathbf{\Omega} \mathbf{Q}_{xx}^{-1} = \mathbf{Q}_{xx}^{-1} \sigma^2 \equiv \mathbf{V}_\beta^0 \quad (6.17)$$

In (6.17) we define $\mathbf{V}_\beta^0 = \mathbf{Q}_{xx}^{-1} \sigma^2$ whether (6.15) is true or false. When (6.15) is true then $\mathbf{V}_\beta = \mathbf{V}_\beta^0$, otherwise $\mathbf{V}_\beta \neq \mathbf{V}_\beta^0$. We call \mathbf{V}_β^0 the **homoskedastic asymptotic covariance matrix**.

Theorem 6.3.2 states that the sampling distribution of the least-squares estimator, after rescaling, is approximately normal when the sample size n is sufficiently large. This holds true for all joint distributions of (y_i, \mathbf{x}_i) which satisfy the conditions of Assumption 6.1.2, and is therefore broadly applicable. Consequently, asymptotic normality is routinely used to approximate the finite sample distribution of $\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})$.

A difficulty is that for any fixed n the sampling distribution of $\hat{\boldsymbol{\beta}}$ can be arbitrarily far from the normal distribution. In Figure 5.1 we have already seen a simple example where the least-squares estimate is quite asymmetric and non-normal even for reasonably large sample sizes. The normal approximation improves as n increases, but how large should n be in order for the approximation to be useful? Unfortunately, there is no simple answer to this reasonable question. The trouble is that no matter how large is the sample size, the normal approximation is arbitrarily poor for some data distribution satisfying the assumptions. We illustrate this problem using a simulation. Let $y_i = \beta_1 x_i + \beta_2 + e_i$ where x_i is $N(0, 1)$, and e_i is independent of x_i with the Double Pareto density $f(e) = \frac{\alpha}{2} |e|^{-\alpha-1}$, $|e| \geq 1$. If $\alpha > 2$ the error e_i has zero mean and variance $\alpha/(\alpha - 2)$. As α approaches 2, however, its variance diverges to infinity. In this context the normalized least-squares slope estimator $\sqrt{n \frac{\alpha-2}{\alpha}} (\hat{\beta}_1 - \beta_1)$ has the $N(0, 1)$ asymptotic distribution for any $\alpha > 2$.

In Figure 6.2 we display the finite sample densities of the normalized estimator $\sqrt{n \frac{\alpha-2}{\alpha}} (\hat{\beta}_1 - \beta_1)$, setting $n = 100$ and varying the parameter α . For $\alpha = 3.0$ the density is very close to the $N(0, 1)$ density. As α diminishes the density changes significantly, concentrating most of the probability mass around zero.

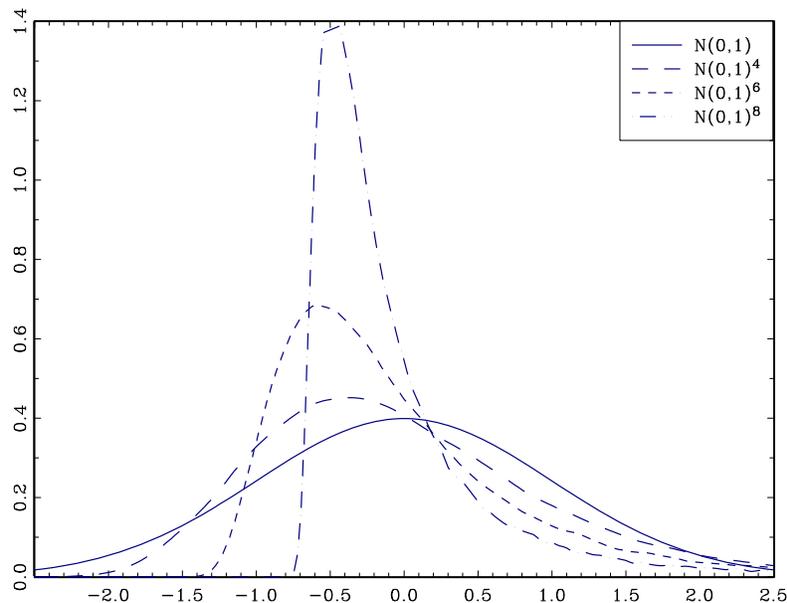


Figure 6.3: Density of Normalized OLS estimator with error process (6.18)

Another example is shown in Figure 6.3. Here the model is $y_i = \beta + e_i$ where

$$e_i = \frac{u_i^k - \mathbb{E}(u_i^k)}{\left(\mathbb{E}(u_i^{2k}) - (\mathbb{E}(u_i^k))^2\right)^{1/2}} \quad (6.18)$$

and $u_i \sim N(0, 1)$. We show the sampling distribution of $\sqrt{n}(\hat{\beta} - \beta)$ setting $n = 100$, for $k = 1, 4, 6$ and 8 . As k increases, the sampling distribution becomes highly skewed and non-normal. The lesson from Figures 6.2 and 6.3 is that the $N(0, 1)$ asymptotic approximation is never guaranteed to be accurate.

6.4 Joint Distribution

Theorem 6.3.2 gives the joint asymptotic distribution of the coefficient estimates. We can use the result to study the covariance between the coefficient estimates. For example, suppose $k = 2$ and write the estimates as $(\hat{\beta}_1, \hat{\beta}_2)$. For simplicity suppose that the regressors are mean zero. Then we can write

$$\mathbf{Q}_{xx} = \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix}$$

where σ_1^2 and σ_2^2 are the variances of x_{1i} and x_{2i} , and ρ is their correlation. If the error is homoskedastic, then the asymptotic variance matrix for $(\hat{\beta}_1, \hat{\beta}_2)$ is $\mathbf{V}_{\beta}^0 = \mathbf{Q}_{xx}^{-1}\sigma^2$. By the formula for inversion of a 2×2 matrix,

$$\mathbf{Q}_{xx}^{-1} = \frac{1}{\sigma_1^2\sigma_2^2(1-\rho^2)} \begin{bmatrix} \sigma_2^2 & -\rho\sigma_1\sigma_2 \\ -\rho\sigma_1\sigma_2 & \sigma_1^2 \end{bmatrix}.$$

Thus if x_{1i} and x_{2i} are positively correlated ($\rho > 0$) then $\hat{\beta}_1$ and $\hat{\beta}_2$ are negatively correlated (and vice-versa).

For illustration, Figure 6.4 displays the probability contours of the joint asymptotic distribution of $\hat{\beta}_1 - \beta_1$ and $\hat{\beta}_2 - \beta_2$ when $\beta_1 = \beta_2 = 0$, $\sigma_1^2 = \sigma_2^2 = \sigma^2 = 1$, and $\rho = 0.5$. The coefficient estimates are negatively correlated since the regressors are positively correlated. This means that if $\hat{\beta}_1$ is

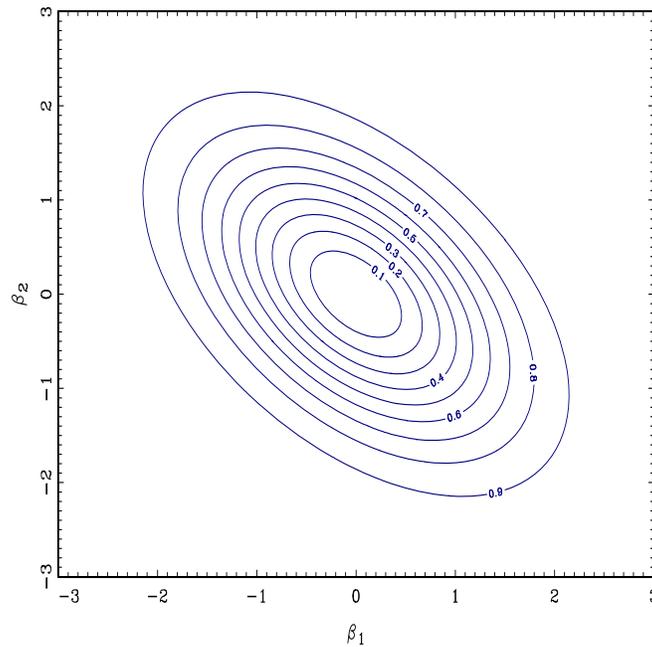


Figure 6.4: Contours of Joint Distribution of $(\hat{\beta}_1, \hat{\beta}_2)$, homoskedastic case

unusually negative, it is likely that $\hat{\beta}_2$ is unusually positive, or conversely. It is also unlikely that we will observe both $\hat{\beta}_1$ and $\hat{\beta}_2$ unusually large and of the same sign.

This finding that the correlation of the regressors is of opposite sign of the correlation of the coefficient estimates is sensitive to the assumption of homoskedasticity. If the errors are heteroskedastic then this relationship is not guaranteed.

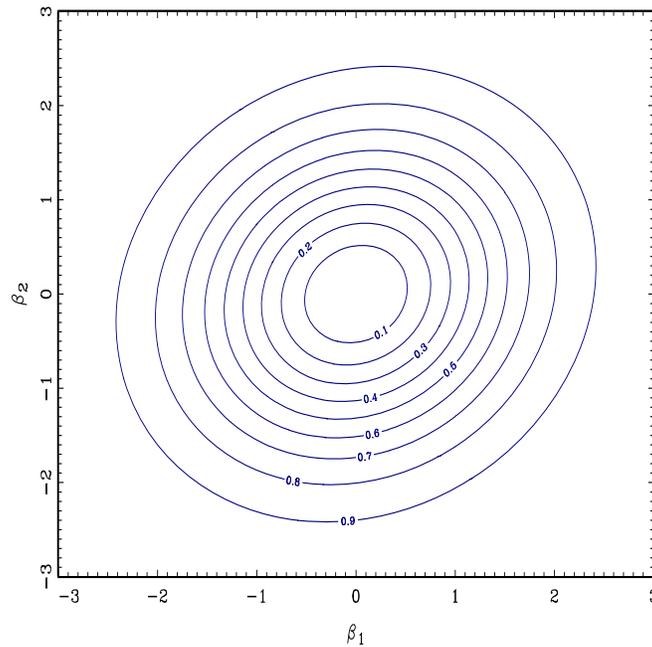
This can be seen through a simple constructed example. Suppose that x_{1i} and x_{2i} only take the values $\{-1, +1\}$, symmetrically, with $\Pr(x_{1i} = x_{2i} = 1) = \Pr(x_{1i} = x_{2i} = -1) = 3/8$, and $\Pr(x_{1i} = 1, x_{2i} = -1) = \Pr(x_{1i} = -1, x_{2i} = 1) = 1/8$. You can check that the regressors are mean zero, unit variance and correlation 0.5, which is identical with the setting displayed in Figure 6.4.

Now suppose that the error is heteroskedastic. Specifically, suppose that $\mathbb{E}(e_i^2 | x_{1i} = x_{2i}) = \frac{5}{4}$ and $\mathbb{E}(e_i^2 | x_{1i} \neq x_{2i}) = \frac{1}{4}$. You can check that $\mathbb{E}(e_i^2) = 1$, $\mathbb{E}(x_{1i}^2 e_i^2) = \mathbb{E}(x_{2i}^2 e_i^2) = 1$ and $\mathbb{E}(x_{1i} x_{2i} e_i^2) = \frac{7}{8}$. Therefore

$$\begin{aligned} \mathbf{V}_\beta &= \mathbf{Q}_{xx}^{-1} \mathbf{\Omega} \mathbf{Q}_{xx}^{-1} \\ &= \frac{9}{16} \begin{bmatrix} 1 & -\frac{1}{2} \\ -\frac{1}{2} & 1 \end{bmatrix} \begin{bmatrix} 1 & \frac{7}{8} \\ \frac{7}{8} & 1 \end{bmatrix} \begin{bmatrix} 1 & -\frac{1}{2} \\ -\frac{1}{2} & 1 \end{bmatrix} \\ &= \frac{4}{3} \begin{bmatrix} 1 & \frac{1}{4} \\ \frac{1}{4} & 1 \end{bmatrix}. \end{aligned}$$

Thus the coefficient estimates $\hat{\beta}_1$ and $\hat{\beta}_2$ are positively correlated (their correlation is 1/4.) The joint probability contours of their asymptotic distribution is displayed in Figure 6.5. We can see how the two estimates are positively associated.

What we found through this example is that in the presence of heteroskedasticity there is no simple relationship between the correlation of the regressors and the correlation of the parameter

Figure 6.5: Contours of Joint Distribution of $\hat{\beta}_1$ and $\hat{\beta}_2$, heteroskedastic case

estimates.

We can extend the above analysis to study the covariance between coefficient sub-vectors. For example, partitioning $\mathbf{x}'_i = (\mathbf{x}'_{1i}, \mathbf{x}'_{2i})$ and $\boldsymbol{\beta}' = (\boldsymbol{\beta}'_1, \boldsymbol{\beta}'_2)$, we can write the general model as

$$y_i = \mathbf{x}'_{1i}\boldsymbol{\beta}_1 + \mathbf{x}'_{2i}\boldsymbol{\beta}_2 + e_i$$

and the coefficient estimates as $\hat{\boldsymbol{\beta}}' = (\hat{\boldsymbol{\beta}}'_1, \hat{\boldsymbol{\beta}}'_2)$. Make the partitions

$$\mathbf{Q}_{xx} = \begin{bmatrix} \mathbf{Q}_{11} & \mathbf{Q}_{12} \\ \mathbf{Q}_{21} & \mathbf{Q}_{22} \end{bmatrix}, \quad \boldsymbol{\Omega} = \begin{bmatrix} \boldsymbol{\Omega}_{11} & \boldsymbol{\Omega}_{12} \\ \boldsymbol{\Omega}_{21} & \boldsymbol{\Omega}_{22} \end{bmatrix}. \quad (6.19)$$

From (2.41)

$$\mathbf{Q}_{xx}^{-1} = \begin{bmatrix} \mathbf{Q}_{11.2}^{-1} & -\mathbf{Q}_{11.2}^{-1}\mathbf{Q}_{12}\mathbf{Q}_{22}^{-1} \\ -\mathbf{Q}_{22.1}^{-1}\mathbf{Q}_{21}\mathbf{Q}_{11}^{-1} & \mathbf{Q}_{22.1}^{-1} \end{bmatrix}$$

where $\mathbf{Q}_{11.2} = \mathbf{Q}_{11} - \mathbf{Q}_{12}\mathbf{Q}_{22}^{-1}\mathbf{Q}_{21}$ and $\mathbf{Q}_{22.1} = \mathbf{Q}_{22} - \mathbf{Q}_{21}\mathbf{Q}_{11}^{-1}\mathbf{Q}_{12}$. Thus when the error is homoskedastic,

$$\text{cov}(\hat{\boldsymbol{\beta}}_1, \hat{\boldsymbol{\beta}}_2) = -\sigma^2 \mathbf{Q}_{11.2}^{-1} \mathbf{Q}_{12} \mathbf{Q}_{22}^{-1}$$

which is a matrix generalization of the two-regressor case.

In the general case, you can show that (Exercise 6.5)

$$\mathbf{V}_{\boldsymbol{\beta}} = \begin{bmatrix} \mathbf{V}_{11} & \mathbf{V}_{12} \\ \mathbf{V}_{21} & \mathbf{V}_{22} \end{bmatrix} \quad (6.20)$$

where

$$\mathbf{V}_{11} = \mathbf{Q}_{11.2}^{-1} (\boldsymbol{\Omega}_{11} - \mathbf{Q}_{12}\mathbf{Q}_{22}^{-1}\boldsymbol{\Omega}_{21} - \boldsymbol{\Omega}_{12}\mathbf{Q}_{22}^{-1}\mathbf{Q}_{21} + \mathbf{Q}_{12}\mathbf{Q}_{22}^{-1}\boldsymbol{\Omega}_{22}\mathbf{Q}_{22}^{-1}\mathbf{Q}_{21}) \mathbf{Q}_{11.2}^{-1} \quad (6.21)$$

$$\mathbf{V}_{21} = \mathbf{Q}_{22.1}^{-1} (\boldsymbol{\Omega}_{21} - \mathbf{Q}_{21}\mathbf{Q}_{11}^{-1}\boldsymbol{\Omega}_{11} - \boldsymbol{\Omega}_{22}\mathbf{Q}_{22}^{-1}\mathbf{Q}_{21} + \mathbf{Q}_{21}\mathbf{Q}_{11}^{-1}\boldsymbol{\Omega}_{12}\mathbf{Q}_{22}^{-1}\mathbf{Q}_{21}) \mathbf{Q}_{11.2}^{-1} \quad (6.22)$$

$$\mathbf{V}_{22} = \mathbf{Q}_{22.1}^{-1} (\boldsymbol{\Omega}_{22} - \mathbf{Q}_{21}\mathbf{Q}_{11}^{-1}\boldsymbol{\Omega}_{12} - \boldsymbol{\Omega}_{21}\mathbf{Q}_{11}^{-1}\mathbf{Q}_{12} + \mathbf{Q}_{21}\mathbf{Q}_{11}^{-1}\boldsymbol{\Omega}_{11}\mathbf{Q}_{11}^{-1}\mathbf{Q}_{12}) \mathbf{Q}_{22.1}^{-1} \quad (6.23)$$

Unfortunately, these expressions are not easily interpretable.

6.5 Consistency of Error Variance Estimators

Using the methods of Section 6.2 we can show that the estimators $\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n \hat{e}_i^2$ and $s^2 = \frac{1}{n-k} \sum_{i=1}^n \hat{e}_i^2$ are consistent for σ^2 .

The trick is to write the residual \hat{e}_i as equal to the error e_i plus a deviation term

$$\begin{aligned}\hat{e}_i &= y_i - \mathbf{x}_i' \hat{\boldsymbol{\beta}} \\ &= e_i + \mathbf{x}_i' \boldsymbol{\beta} - \mathbf{x}_i' \hat{\boldsymbol{\beta}} \\ &= e_i - \mathbf{x}_i' (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}).\end{aligned}$$

Thus the squared residual equals the squared error plus a deviation

$$\hat{e}_i^2 = e_i^2 - 2e_i \mathbf{x}_i' (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) + (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})' \mathbf{x}_i \mathbf{x}_i' (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}). \quad (6.24)$$

So when we take the average of the squared residuals we obtain the average of the squared errors, plus two terms which are (hopefully) asymptotically negligible.

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n e_i^2 - 2 \left(\frac{1}{n} \sum_{i=1}^n e_i \mathbf{x}_i' \right) (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) + (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})' \left(\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \right) (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}). \quad (6.25)$$

Indeed, the WLLN shows that

$$\begin{aligned}\frac{1}{n} \sum_{i=1}^n e_i^2 &\xrightarrow{p} \sigma^2 \\ \frac{1}{n} \sum_{i=1}^n e_i \mathbf{x}_i' &\xrightarrow{p} \mathbb{E}(e_i \mathbf{x}_i') = 0 \\ \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' &\xrightarrow{p} \mathbb{E}(\mathbf{x}_i \mathbf{x}_i') = \mathbf{Q}_{xx}\end{aligned}$$

and Theorem 6.2.1 shows that $\hat{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\beta}$. Hence (6.25) converges in probability to σ^2 , as desired.

Finally, since $n/(n-k) \rightarrow 1$ as $n \rightarrow \infty$, it follows that

$$s^2 = \left(\frac{n}{n-k} \right) \hat{\sigma}^2 \xrightarrow{p} \sigma^2.$$

Thus both estimators are consistent.

Theorem 6.5.1 Under Assumption 6.1.1, $\hat{\sigma}^2 \xrightarrow{p} \sigma^2$ and $s^2 \xrightarrow{p} \sigma^2$ as $n \rightarrow \infty$.

6.6 Homoskedastic Covariance Matrix Estimation

Theorem 6.3.2 shows that $\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})$ is asymptotically normal with asymptotic covariance matrix $\mathbf{V}_{\boldsymbol{\beta}}$. For asymptotic inference (confidence intervals and tests) we need a consistent estimate of $\mathbf{V}_{\boldsymbol{\beta}}$. Under homoskedasticity, $\mathbf{V}_{\boldsymbol{\beta}}$ simplifies to $\mathbf{V}_{\boldsymbol{\beta}}^0 = \mathbf{Q}_{xx}^{-1} \sigma^2$, and in this section we consider the simplified problem of estimating $\mathbf{V}_{\boldsymbol{\beta}}^0$.

The standard moment estimator of \mathbf{Q}_{xx} is $\widehat{\mathbf{Q}}_{xx}$ defined in (6.1), and thus an estimator for \mathbf{Q}_{xx}^{-1} is $\widehat{\mathbf{Q}}_{xx}^{-1}$. Also, the standard estimator of σ^2 is the unbiased estimator s^2 defined in (4.21). Thus a natural plug-in estimator for $\mathbf{V}_\beta^0 = \mathbf{Q}_{xx}^{-1}\sigma^2$ is $\widehat{\mathbf{V}}_\beta^0 = \widehat{\mathbf{Q}}_{xx}^{-1}s^2$.

Consistency of $\widehat{\mathbf{V}}_\beta^0$ for \mathbf{V}_β^0 follows from consistency of the moment estimates $\widehat{\mathbf{Q}}_{xx}$ and s^2 , and an application of the continuous mapping theorem. Specifically, Theorem 6.2.1 established that $\widehat{\mathbf{Q}}_{xx} \xrightarrow{p} \mathbf{Q}_{xx}$, and Theorem 6.5.1 established $s^2 \xrightarrow{p} \sigma^2$. The function $\mathbf{V}_\beta^0 = \mathbf{Q}_{xx}^{-1}\sigma^2$ is a continuous function of \mathbf{Q}_{xx} and σ^2 so long as $\mathbf{Q}_{xx} > 0$, which holds true under Assumption 6.1.1.4. It follows by the CMT that

$$\widehat{\mathbf{V}}_\beta^0 = \widehat{\mathbf{Q}}_{xx}^{-1}s^2 \xrightarrow{p} \mathbf{Q}_{xx}^{-1}\sigma^2 = \mathbf{V}_\beta^0$$

so that $\widehat{\mathbf{V}}_\beta^0$ is consistent for \mathbf{V}_β^0 , as desired.

Theorem 6.6.1 Under Assumption 6.1.1, $\widehat{\mathbf{V}}_\beta^0 \xrightarrow{p} \mathbf{V}_\beta^0$ as $n \rightarrow \infty$.

It is instructive to notice that Theorem 6.6.1 does not require the assumption of homoskedasticity. That is, $\widehat{\mathbf{V}}_\beta^0$ is consistent for \mathbf{V}_β^0 regardless if the regression is homoskedastic or heteroskedastic. However, $\mathbf{V}_\beta^0 = \mathbf{V}_\beta = \text{avar}(\widehat{\beta})$ only under homoskedasticity. Thus in the general case, $\widehat{\mathbf{V}}_\beta^0$ is consistent for a well-defined but non-useful object.

6.7 Heteroskedastic Covariance Matrix Estimation

Theorems 6.3.2 established that the asymptotic covariance matrix of $\sqrt{n}(\widehat{\beta} - \beta)$ is $\mathbf{V}_\beta = \mathbf{Q}_{xx}^{-1}\boldsymbol{\Omega}\mathbf{Q}_{xx}^{-1}$. We now consider estimation of this covariance matrix without imposing homoskedasticity. The standard approach is to use a plug-in estimator which replaces the unknowns with sample moments.

As described in the previous section, a natural estimator for \mathbf{Q}_{xx}^{-1} is $\widehat{\mathbf{Q}}_{xx}^{-1}$, where $\widehat{\mathbf{Q}}_{xx}$ defined in (6.1).

The moment estimator for $\boldsymbol{\Omega}$ is

$$\widehat{\boldsymbol{\Omega}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \widehat{e}_i^2, \quad (6.26)$$

leading to the plug-in covariance matrix estimator

$$\widehat{\mathbf{V}}_\beta^W = \widehat{\mathbf{Q}}_{xx}^{-1} \widehat{\boldsymbol{\Omega}} \widehat{\mathbf{Q}}_{xx}^{-1}. \quad (6.27)$$

You can check that $\widehat{\mathbf{V}}_\beta^W = n \widehat{\mathbf{V}}_{\widehat{\beta}}^W$ where $\widehat{\mathbf{V}}_{\widehat{\beta}}^W$ is the White covariance matrix estimator introduced in (4.28).

As shown in Theorem 6.2.1, $\widehat{\mathbf{Q}}_{xx}^{-1} \xrightarrow{p} \mathbf{Q}_{xx}^{-1}$, so we just need to verify the consistency of $\widehat{\boldsymbol{\Omega}}$. The key is to replace the squared residual \widehat{e}_i^2 with the squared error e_i^2 , and then show that the difference is asymptotically negligible.

Specifically, observe that

$$\begin{aligned} \widehat{\boldsymbol{\Omega}} &= \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \widehat{e}_i^2 \\ &= \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' e_i^2 + \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' (\widehat{e}_i^2 - e_i^2). \end{aligned} \quad (6.28)$$

The first term is an average of the iid random variables $\mathbf{x}_i \mathbf{x}_i' e_i^2$, and therefore by the WLLN converges in probability to its expectation, namely,

$$\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' e_i^2 \xrightarrow{p} \mathbb{E}(\mathbf{x}_i \mathbf{x}_i' e_i^2) = \mathbf{\Omega}.$$

Technically, this requires that $\mathbf{\Omega}$ has finite elements, which was shown in (6.10).

So to establish that $\widehat{\mathbf{\Omega}}$ is consistent for $\mathbf{\Omega}$ it remains to show that

$$\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' (\hat{e}_i^2 - e_i^2) \xrightarrow{p} 0. \quad (6.29)$$

There are multiple ways to do this. A reasonable straightforward yet slightly tedious derivation is to start by applying the Triangle Inequality (A.12)

$$\begin{aligned} \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' (\hat{e}_i^2 - e_i^2) \right\| &\leq \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i \mathbf{x}_i' (\hat{e}_i^2 - e_i^2)\| \\ &= \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i\|^2 |\hat{e}_i^2 - e_i^2|. \end{aligned} \quad (6.30)$$

Then recalling the expression for the squared residual (6.24), apply the Triangle Inequality and then the Schwarz Inequality (A.10) twice

$$\begin{aligned} |\hat{e}_i^2 - e_i^2| &\leq 2 |e_i \mathbf{x}_i' (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})| + (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})' \mathbf{x}_i \mathbf{x}_i' (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \\ &= 2 |e_i| |\mathbf{x}_i' (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})| + |(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})' \mathbf{x}_i|^2 \\ &\leq 2 |e_i| \|\mathbf{x}_i\| \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\| + \|\mathbf{x}_i\|^2 \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|^2. \end{aligned} \quad (6.31)$$

Combining (6.30) and (6.31), we find

$$\begin{aligned} \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' (\hat{e}_i^2 - e_i^2) \right\| &\leq 2 \left(\frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i\|^3 |e_i| \right) \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\| + \left(\frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i\|^4 \right) \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\|^2 \\ &= o_p(1). \end{aligned} \quad (6.32)$$

The expression is $o_p(1)$ because $\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}\| \xrightarrow{p} 0$ and both averages in parenthesis are averages of random variables with finite mean under Assumption 6.1.2 (and are thus $O_p(1)$). Indeed, by Hölder's Inequality (B.16)

$$\mathbb{E}(\|\mathbf{x}_i\|^3 |e_i|) \leq \left(\mathbb{E}(\|\mathbf{x}_i\|^3)^{4/3} \right)^{3/4} (\mathbb{E}e_i^4)^{1/4} = \left(\mathbb{E}\|\mathbf{x}_i\|^4 \right)^{3/4} (\mathbb{E}e_i^4)^{1/4} < \infty.$$

We have established (6.29), as desired.

Theorem 6.7.1 Under Assumption 6.1.2, as $n \rightarrow \infty$, $\widehat{\mathbf{\Omega}} \xrightarrow{p} \mathbf{\Omega}$ and $\widehat{\mathbf{V}}_{\boldsymbol{\beta}}^W \xrightarrow{p} \mathbf{V}_{\boldsymbol{\beta}}$.

For an alternative proof of this result, see Section 6.21.

6.8 Summary of Covariance Matrix Notation

The notation we have introduced may be somewhat confusing so it is helpful to write it down in one place. The exact variance of $\hat{\beta}$ and the asymptotic variance of $\sqrt{n}(\hat{\beta} - \beta)$ are

$$\begin{aligned}\mathbf{V}_{\hat{\beta}} &= \text{var}(\hat{\beta} | \mathbf{X}) = (\mathbf{X}'\mathbf{X})^{-1} (\mathbf{X}'\mathbf{D}\mathbf{X}) (\mathbf{X}'\mathbf{X})^{-1} \\ \mathbf{V}_{\beta} &= \text{avar}(\sqrt{n}(\hat{\beta} - \beta)) = \mathbf{Q}_{xx}^{-1} \mathbf{\Omega} \mathbf{Q}_{xx}^{-1}\end{aligned}$$

The White estimates of these two covariance matrices are

$$\begin{aligned}\hat{\mathbf{V}}_{\hat{\beta}}^W &= (\mathbf{X}'\mathbf{X})^{-1} \left(\sum_{i=1}^n x_i x_i' \hat{e}_i^2 \right) (\mathbf{X}'\mathbf{X})^{-1} \\ \hat{\mathbf{V}}_{\beta}^W &= \hat{\mathbf{Q}}_{xx}^{-1} \hat{\mathbf{\Omega}} \hat{\mathbf{Q}}_{xx}^{-1}\end{aligned}$$

and satisfy the simple relationship

$$\hat{\mathbf{V}}_{\beta}^W = n \hat{\mathbf{V}}_{\hat{\beta}}^W.$$

Similarly, under the assumption of homoskedasticity the exact and asymptotic variances simplify to

$$\begin{aligned}\mathbf{V}_{\hat{\beta}}^0 &= (\mathbf{X}'\mathbf{X})^{-1} \sigma^2 \\ \mathbf{V}_{\beta}^0 &= \mathbf{Q}_{xx}^{-1} \sigma^2\end{aligned}$$

and their standard estimators are

$$\begin{aligned}\hat{\mathbf{V}}_{\hat{\beta}}^0 &= (\mathbf{X}'\mathbf{X})^{-1} s^2 \\ \hat{\mathbf{V}}_{\beta}^0 &= \hat{\mathbf{Q}}_{xx}^{-1} s^2\end{aligned}$$

which also satisfy the relationship

$$\hat{\mathbf{V}}_{\beta}^0 = n \hat{\mathbf{V}}_{\hat{\beta}}^0.$$

The exact formula and estimates are useful when constructing test statistics and standard errors, as in practice we are interested in the covariance matrix of the estimates $\hat{\beta}$, not in the asymptotic variance. However, for theoretical purposes the asymptotic formula (variances and their estimates) are more useful, as these retain non-degenerate limits as the sample sizes diverge. That is why both sets of notation are useful.

6.9 Alternative Covariance Matrix Estimators*

In Section 6.7 we introduced $\hat{\mathbf{V}}_{\beta}^W$ as an estimator of \mathbf{V}_{β} . $\hat{\mathbf{V}}_{\beta}^W$ is a scaled version of $\hat{\mathbf{V}}_{\hat{\beta}}^W$ from Section 4.11, where we also introduced the alternative heteroskedasticity-robust covariance matrix estimators $\hat{\mathbf{V}}_{\hat{\beta}}$, $\tilde{\mathbf{V}}_{\hat{\beta}}$ and $\bar{\mathbf{V}}_{\hat{\beta}}$. We now discuss the consistency properties of these estimators.

To do so we introduce their scaled versions, e.g. $\hat{\mathbf{V}}_{\beta} = n \hat{\mathbf{V}}_{\hat{\beta}}$, $\tilde{\mathbf{V}}_{\beta} = n \tilde{\mathbf{V}}_{\hat{\beta}}$, and $\bar{\mathbf{V}}_{\beta} = n \bar{\mathbf{V}}_{\hat{\beta}}$. These are (alternative) estimates of the asymptotic covariance matrix \mathbf{V}_{β} .

First, consider $\hat{\mathbf{V}}_{\beta}$. Notice that $\hat{\mathbf{V}}_{\beta} = n \hat{\mathbf{V}}_{\hat{\beta}} = \frac{n}{n-k} \hat{\mathbf{V}}_{\hat{\beta}}^W$ where $\hat{\mathbf{V}}_{\hat{\beta}}^W$ was defined in (6.27) and shown consistent for \mathbf{V}_{β} in Theorem 6.7.1. If k is fixed as $n \rightarrow \infty$, then $\frac{n}{n-k} \rightarrow 1$ and thus

$$\hat{\mathbf{V}}_{\beta} = (1 + o(1)) \hat{\mathbf{V}}_{\hat{\beta}}^W \xrightarrow{p} \mathbf{V}_{\beta}.$$

Thus $\hat{\mathbf{V}}_{\beta}$ is consistent for \mathbf{V}_{β} .

The alternative estimators $\tilde{\mathbf{V}}_\beta$ and $\bar{\mathbf{V}}_\beta$ take the form (6.27) but with $\hat{\mathbf{\Omega}}$ replaced by

$$\tilde{\mathbf{\Omega}} = \frac{1}{n} \sum_{i=1}^n (1 - h_{ii})^{-2} \mathbf{x}_i \mathbf{x}_i' \hat{e}_i^2$$

and

$$\bar{\mathbf{\Omega}} = \frac{1}{n} \sum_{i=1}^n (1 - h_{ii})^{-1} \mathbf{x}_i \mathbf{x}_i' \hat{e}_i^2,$$

respectively. To show that these estimators also consistent for \mathbf{V}_β , given $\hat{\mathbf{\Omega}} \xrightarrow{p} \mathbf{\Omega}$, it is sufficient to show that the differences $\tilde{\mathbf{\Omega}} - \hat{\mathbf{\Omega}}$ and $\bar{\mathbf{\Omega}} - \hat{\mathbf{\Omega}}$ converge in probability to zero as $n \rightarrow \infty$.

The trick is to use the fact that the leverage values are asymptotically negligible:

$$h_n^* = \max_{1 \leq i \leq n} h_{ii} = o_p(1). \quad (6.33)$$

(See Theorem 6.22.1 in Section 6.22.) Then using the Triangle Inequality

$$\begin{aligned} \|\bar{\mathbf{\Omega}} - \hat{\mathbf{\Omega}}\| &\leq \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i \mathbf{x}_i'\| \hat{e}_i^2 \left| (1 - h_{ii})^{-1} - 1 \right| \\ &\leq \left(\frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i\|^2 \hat{e}_i^2 \right) \left| (1 - h_n^*)^{-1} - 1 \right|. \end{aligned}$$

The sum in parenthesis can be shown to be $O_p(1)$ under Assumption 6.1.2 by the same argument as in the proof of Theorem 6.7.1. (In fact, it can be shown to converge in probability to $\mathbb{E}(\|\mathbf{x}_i\|^2 e_i^2)$.) The term in absolute values is $o_p(1)$ by (6.33). Thus the product is $o_p(1)$, which means that $\bar{\mathbf{\Omega}} = \hat{\mathbf{\Omega}} + o_p(1) \rightarrow \mathbf{\Omega}$.

Similarly,

$$\begin{aligned} \|\tilde{\mathbf{\Omega}} - \hat{\mathbf{\Omega}}\| &\leq \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i \mathbf{x}_i'\| \hat{e}_i^2 \left| (1 - h_{ii})^{-2} - 1 \right| \\ &\leq \left(\frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i\|^2 \hat{e}_i^2 \right) \left| (1 - h_n^*)^{-2} - 1 \right| \\ &= o_p(1). \end{aligned}$$

Theorem 6.9.1 Under Assumption 6.1.2, as $n \rightarrow \infty$, $\tilde{\mathbf{\Omega}} \xrightarrow{p} \mathbf{\Omega}$, $\bar{\mathbf{\Omega}} \xrightarrow{p} \mathbf{\Omega}$, $\hat{\mathbf{V}}_\beta \xrightarrow{p} \mathbf{V}_\beta$, $\tilde{\mathbf{V}}_\beta \xrightarrow{p} \mathbf{V}_\beta$, and $\bar{\mathbf{V}}_\beta \xrightarrow{p} \mathbf{V}_\beta$.

Theorem 6.9.1 shows that the alternative covariance matrix estimators are also consistent for the asymptotic covariance matrix.

6.10 Functions of Parameters

Sometimes we are interested in a transformation of the coefficient vector $\boldsymbol{\beta} = (\beta_1, \dots, \beta_k)$. For example, we may be interested in a single coefficient β_j , or a ratio β_j/β_l . In these cases we can write

the transformation as a function of the coefficients, e.g. $\boldsymbol{\theta} = \mathbf{r}(\boldsymbol{\beta})$ for some function $\mathbf{r} : \mathbb{R}^k \rightarrow \mathbb{R}^q$. The estimate of $\boldsymbol{\theta}$ is

$$\widehat{\boldsymbol{\theta}} = \mathbf{r}(\widehat{\boldsymbol{\beta}}).$$

By the continuous mapping theorem (Theorem 5.9.1) and the fact $\widehat{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\beta}$ we can deduce that $\widehat{\boldsymbol{\theta}}$ is consistent for $\boldsymbol{\theta}$.

Theorem 6.10.1 *Under Assumption 6.1.1, if $\mathbf{r}(\boldsymbol{\beta})$ is continuous at the true value of $\boldsymbol{\beta}$, then as $n \rightarrow \infty$, $\widehat{\boldsymbol{\theta}} \xrightarrow{p} \boldsymbol{\theta}$.*

Furthermore, if the transformation is sufficiently smooth, by the Delta Method (Theorem 5.10.3) we can show that $\widehat{\boldsymbol{\theta}}$ is asymptotically normal.

Assumption 6.10.1 *$\mathbf{r}(\boldsymbol{\beta}) : \mathbb{R}^k \rightarrow \mathbb{R}^q$ is continuously differentiable at the true value of $\boldsymbol{\beta}$ and $\mathbf{R} = \frac{\partial}{\partial \boldsymbol{\beta}} \mathbf{r}(\boldsymbol{\beta})'$ has rank q .*

Theorem 6.10.2 Asymptotic Distribution of Functions of Parameters

Under Assumptions 6.1.2 and 6.10.1, as $n \rightarrow \infty$,

$$\sqrt{n}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \xrightarrow{d} N(\mathbf{0}, \mathbf{V}_{\boldsymbol{\theta}}) \quad (6.34)$$

where

$$\mathbf{V}_{\boldsymbol{\theta}} = \mathbf{R}' \mathbf{V}_{\boldsymbol{\beta}} \mathbf{R} \quad (6.35)$$

In many cases, the function $\mathbf{r}(\boldsymbol{\beta})$ is linear:

$$\mathbf{r}(\boldsymbol{\beta}) = \mathbf{R}' \boldsymbol{\beta}$$

for some $k \times q$ matrix \mathbf{R} . In particular, if \mathbf{R} is a “selector matrix”

$$\mathbf{R} = \begin{pmatrix} \mathbf{I} \\ \mathbf{0} \end{pmatrix} \quad (6.36)$$

then we can conformably partition $\boldsymbol{\beta} = (\boldsymbol{\beta}'_1, \boldsymbol{\beta}'_2)'$ so that $\mathbf{R}' \boldsymbol{\beta} = \boldsymbol{\beta}_1$ for $\boldsymbol{\beta} = (\boldsymbol{\beta}'_1, \boldsymbol{\beta}'_2)'$. Then

$$\mathbf{V}_{\boldsymbol{\theta}} = \begin{pmatrix} \mathbf{I} & \mathbf{0} \end{pmatrix} \mathbf{V}_{\boldsymbol{\beta}} \begin{pmatrix} \mathbf{I} \\ \mathbf{0} \end{pmatrix} = \mathbf{V}_{11},$$

the upper-left sub-matrix of \mathbf{V}_{11} given in (6.21). In this case (6.34) states that

$$\sqrt{n}(\widehat{\boldsymbol{\beta}}_1 - \boldsymbol{\beta}_1) \xrightarrow{d} N(\mathbf{0}, \mathbf{V}_{11}).$$

That is, subsets of $\widehat{\boldsymbol{\beta}}$ are approximately normal with variances given by the conformable subcomponents of \mathbf{V} .

To illustrate the case of a nonlinear transformation, take the example $\theta = \beta_j/\beta_l$ for $j \neq l$. Then

$$\mathbf{R} = \begin{pmatrix} 0 \\ \vdots \\ 1/\beta_l \\ \vdots \\ -\beta_j/\beta_l^2 \\ \vdots \\ 0 \end{pmatrix} \quad (6.37)$$

so

$$\mathbf{V}_\theta = \mathbf{V}_{jj}/\beta_l^2 + \mathbf{V}_{ll}\beta_j^2/\beta_l^4 - 2\mathbf{V}_{jl}\beta_j/\beta_l^3$$

where \mathbf{V}_{ab} denotes the ab 'th element of \mathbf{V}_β .

For inference we need an estimate of the asymptotic variance matrix $\mathbf{V}_\theta = \mathbf{R}'\mathbf{V}_\beta\mathbf{R}$, and for this it is typical to use a plug-in estimator. The natural estimator of \mathbf{R} is the derivative evaluated at the point estimates

$$\widehat{\mathbf{R}} = \frac{\partial}{\partial \boldsymbol{\beta}} \mathbf{r}(\widehat{\boldsymbol{\beta}})'. \quad (6.38)$$

The derivative in (6.38) may be calculated analytically or numerically. By analytically, we mean working out for the formula for the derivative and replacing the unknowns by point estimates. For example, if $\theta = \beta_j/\beta_l$, then $\frac{\partial}{\partial \boldsymbol{\beta}} \mathbf{r}(\boldsymbol{\beta})$ is (6.37). However in some cases the function $\mathbf{r}(\boldsymbol{\beta})$ may be extremely complicated and a formula for the analytic derivative may not be easily available. In this case calculation by numerical differentiation may be preferable. Let $\delta_l = (0 \cdots 1 \cdots 0)'$ be the unit vector with the “1” in the l 'th place. Then the jl 'th element of a numerical derivative $\widehat{\mathbf{R}}$ is

$$\widehat{\mathbf{R}}_{jl} = \frac{\mathbf{r}_j(\widehat{\boldsymbol{\beta}} + \delta_l \varepsilon) - \mathbf{r}_j(\widehat{\boldsymbol{\beta}})}{\varepsilon}$$

for some small ε .

The estimate of \mathbf{V}_θ is

$$\widehat{\mathbf{V}}_\theta = \widehat{\mathbf{R}}' \widehat{\mathbf{V}}_\beta \widehat{\mathbf{R}}. \quad (6.39)$$

Alternatively, $\widehat{\mathbf{V}}_\beta^0$, $\widetilde{\mathbf{V}}_\beta$ or $\overline{\mathbf{V}}_\beta$ may be used in place of $\widehat{\mathbf{V}}_\beta$. For example, the homoskedastic covariance matrix estimator is

$$\widehat{\mathbf{V}}_\theta^0 = \widehat{\mathbf{R}}' \widehat{\mathbf{V}}_\beta^0 \widehat{\mathbf{R}} = \widehat{\mathbf{R}}' \widehat{\mathbf{Q}}_{xx}^{-1} \widehat{\mathbf{R}} s^2 \quad (6.40)$$

Given (6.38), (6.39) and (6.40) are simple to calculate using matrix operations.

As the primary justification for $\widehat{\mathbf{V}}_\theta$ is the asymptotic approximation (6.34), $\widehat{\mathbf{V}}_\theta$ is often called an **asymptotic covariance matrix estimator**.

The estimator $\widehat{\mathbf{V}}_\theta$ is consistent for \mathbf{V}_θ under the conditions of Theorem 6.10.2 since $\widehat{\mathbf{V}}_\beta \xrightarrow{p} \mathbf{V}_\beta$ by Theorem 6.7.1, and

$$\widehat{\mathbf{R}} = \frac{\partial}{\partial \boldsymbol{\beta}} \mathbf{r}(\widehat{\boldsymbol{\beta}})' \xrightarrow{p} \frac{\partial}{\partial \boldsymbol{\beta}} \mathbf{r}(\boldsymbol{\beta})' = \mathbf{R}$$

since $\widehat{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\beta}$ and the function $\frac{\partial}{\partial \boldsymbol{\beta}} \mathbf{r}(\boldsymbol{\beta})'$ is continuous.

Theorem 6.10.3 Under Assumptions 6.1.2 and 6.10.1, as $n \rightarrow \infty$,

$$\widehat{\mathbf{V}}_\theta \xrightarrow{p} \mathbf{V}_\theta.$$

Theorem 6.10.3 shows that $\widehat{\mathbf{V}}_{\boldsymbol{\theta}}$ is consistent for $\mathbf{V}_{\boldsymbol{\theta}}$ and thus may be used for asymptotic inference. In practice, we may set

$$\widehat{\mathbf{V}}_{\boldsymbol{\theta}} = \widehat{\mathbf{R}}' \widehat{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}} \widehat{\mathbf{R}} = n^{-1} \widehat{\mathbf{R}}' \widehat{\mathbf{V}}_{\boldsymbol{\beta}} \widehat{\mathbf{R}} \quad (6.41)$$

as an estimate of the variance of $\widehat{\boldsymbol{\theta}}$, or substitute an alternative covariance estimator such as $\overline{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}}$.

6.11 Asymptotic Standard Errors

As described in Section 4.12, a standard error is an estimate of the standard deviation of the distribution of an estimator. Thus if $\widehat{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}}$ is an estimate of the covariance matrix of $\widehat{\boldsymbol{\beta}}$, then standard errors are the square roots of the diagonal elements of this matrix. These take the form

$$s(\widehat{\beta}_j) = \sqrt{\widehat{\mathbf{V}}_{\widehat{\beta}_j}} = \sqrt{[\widehat{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}}]_{jj}}.$$

Standard errors for $\widehat{\boldsymbol{\theta}}$ are constructed similarly. Supposing that $q = 1$ (so $h(\boldsymbol{\beta})$ is real-valued), then the standard error for $\widehat{\theta}$ is the square root of (6.41)

$$s(\widehat{\theta}) = \sqrt{\widehat{\mathbf{R}}' \widehat{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}} \widehat{\mathbf{R}}} = \sqrt{n^{-1} \widehat{\mathbf{R}}' \widehat{\mathbf{V}}_{\boldsymbol{\beta}} \widehat{\mathbf{R}}}.$$

When the justification is based on asymptotic theory we call $s(\widehat{\beta}_j)$ or $s(\widehat{\theta})$ an **asymptotic standard error** for $\widehat{\beta}_j$ or $\widehat{\theta}$. When reporting your results, it is good practice to report standard errors for each reported estimate, and this includes functions and transformations of your parameter estimates. This helps users of the work (including yourself) assess the estimation precision.

We illustrate using the log wage regression

$$\log(\text{Wage}) = \beta_1 \text{education} + \beta_2 \text{experience} + \beta_3 \text{experience}^2/100 + \beta_4 + e.$$

Consider the following three parameters of interest.

1. Percentage return to education:

$$\theta_1 = 100\beta_1$$

(100 times the partial derivative of the conditional expectation of log wages with respect to *education*.)

2. Percentage return to experience for individuals with 10 years of experience

$$\theta_2 = 100\beta_2 + 20\beta_3$$

(100 times the partial derivative of the conditional expectation of log wages with respect to *experience*, evaluated at *experience* = 10)

3. Experience level which maximizes expected log wages:

$$\theta_3 = -50\beta_2/\beta_3$$

(The level of *experience* at which the partial derivative of the conditional expectation of log wages with respect to *experience* equals 0.)

The 4×1 vector \mathbf{R} for these three examples is

$$\mathbf{R} = \begin{pmatrix} 100 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} 0 \\ 100 \\ 2 \\ 0 \end{pmatrix}, \quad \begin{pmatrix} 0 \\ -50/\beta_3 \\ 50\beta_2/\beta_3^2 \\ 0 \end{pmatrix},$$

respectively.

We use the subsample of married black women (all experience levels), which has 982 observations. The point estimates and standard errors are

$$\log(\widehat{Wage}) = \begin{matrix} 0.118 & \textit{education} + & 0.016 & \textit{experience} - & 0.022 & \textit{experience}^2/100 + & 0.947 \\ (0.008) & & (0.006) & & (0.012) & & (0.157) \end{matrix} . \quad (6.42)$$

The standard errors are the square roots of the Horn-Horn-Duncan covariance matrix estimate

$$\overline{\mathbf{V}}_{\hat{\beta}} = \begin{pmatrix} 0.632 & 0.131 & -0.143 & -11.1 \\ 0.131 & 0.390 & -0.731 & -6.25 \\ -0.143 & -0.731 & 1.48 & 9.43 \\ -11.1 & -6.25 & 9.43 & 246 \end{pmatrix} \times 10^{-4}. \quad (6.43)$$

We calculate that

$$\begin{aligned} \hat{\theta}_1 &= 100\hat{\beta}_1 \\ &= 100 \times 0.118 \\ &= 11.8 \end{aligned}$$

$$\begin{aligned} s(\hat{\theta}_1) &= \sqrt{100^2 \times 0.632 \times 10^{-4}} \\ &= 0.8 \end{aligned}$$

$$\begin{aligned} \hat{\theta}_2 &= 100\hat{\beta}_2 + 20\hat{\beta}_3 \\ &= 100 \times 0.016 - 20 \times 0.022 \\ &= 1.16 \end{aligned}$$

$$\begin{aligned} s(\hat{\theta}_2) &= \sqrt{\begin{pmatrix} 100 & 20 \end{pmatrix} \begin{pmatrix} 0.390 & -0.731 \\ -0.731 & 1.48 \end{pmatrix} \begin{pmatrix} 100 \\ 20 \end{pmatrix} \times 10^{-4}} \\ &= 0.55 \end{aligned}$$

$$\begin{aligned} \hat{\theta}_3 &= -50\hat{\beta}_2/\hat{\beta}_3 \\ &= 50 \times 0.016/0.022 \\ &= 35.2 \end{aligned}$$

$$\begin{aligned} s(\hat{\theta}_3) &= \sqrt{\begin{pmatrix} -50/\hat{\beta}_3 & 50\hat{\beta}_2/\hat{\beta}_3^2 \end{pmatrix} \begin{pmatrix} 0.390 & -0.731 \\ -0.731 & 1.48 \end{pmatrix} \begin{pmatrix} -50/\hat{\beta}_3 \\ 50\hat{\beta}_2/\hat{\beta}_3^2 \end{pmatrix} \times 10^{-4}} \\ &= 7.0 \end{aligned}$$

The calculations show that the estimate of the percentage return to education is about 12% per year, with a standard error of 0.8. The estimate of the percentage return to experience for those with 10 years of experience is 1.2% per year, with a standard error of 0.6. And the estimate of the experience level which maximizes expected log wages is 35 years, with a standard error of 7.

6.12 t statistic

Let $\theta = r(\beta) : \mathbb{R}^k \rightarrow \mathbb{R}$ be any parameter of interest (for example, θ could be a single element of β), $\hat{\theta}$ its estimate and $s(\hat{\theta})$ its asymptotic standard error. Consider the statistic

$$t_n(\theta) = \frac{\hat{\theta} - \theta}{s(\hat{\theta})}. \quad (6.44)$$

Different writers have called (6.44) a **t-statistic**, a **t-ratio**, a **z-statistic** or a **studentized statistic**, sometimes using the different labels to distinguish between finite-sample and asymptotic inference. As the statistics themselves are always (6.44) we won't make such a distinction, and will simply refer to $t_n(\theta)$ as a t-statistic or a t-ratio. We also often suppress the parameter dependence, writing it as t_n . The t-statistic is a simple function of the estimate, its standard error, and the parameter.

By Theorems 6.10.2 and 6.10.3, $\sqrt{n}(\hat{\theta} - \theta) \xrightarrow{d} N(0, V_\theta)$ and $\widehat{V}_\theta \xrightarrow{p} V_\theta$. Thus

$$\begin{aligned} t_n(\theta) &= \frac{\hat{\theta} - \theta}{s(\hat{\theta})} \\ &= \frac{\sqrt{n}(\hat{\theta} - \theta)}{\sqrt{\widehat{V}_\theta}} \\ &\xrightarrow{d} \frac{N(0, V_\theta)}{\sqrt{V_\theta}} \\ &= Z \sim N(0, 1). \end{aligned}$$

The last equality is by the property that linear scales of normal distributions are normal.

Thus the asymptotic distribution of the t-ratio $t_n(\theta)$ is the standard normal. Since this distribution does not depend on the parameters, we say that $t_n(\theta)$ is **asymptotically pivotal**. In special cases (such as the normal regression model, see Section 3.18), the statistic t_n has an exact t distribution, and is therefore exactly free of unknowns. In this case, we say that t_n is **exactly pivotal**. In general, however, pivotal statistics are unavailable and we must rely on asymptotically pivotal statistics.

As we will see in the next section, it is also useful to consider the distribution of the **absolute t-ratio** $|t_n(\theta)|$. Since $t_n(\theta) \xrightarrow{d} Z$, the continuous mapping theorem yields $|t_n(\theta)| \xrightarrow{d} |Z|$. Letting $\Phi(u) = \Pr(Z \leq u)$ denote the standard normal distribution function, we can calculate that the distribution function of $|Z|$ is

$$\begin{aligned} \Pr(|Z| \leq u) &= \Pr(-u \leq Z \leq u) \\ &= \Pr(Z \leq u) - \Pr(Z < -u) \\ &= \Phi(u) - \Phi(-u) \\ &= 2\Phi(u) - 1 \\ &\stackrel{\text{def}}{=} \overline{\Phi}(u). \end{aligned} \quad (6.45)$$

Theorem 6.12.1 Under Assumptions 6.1.2 and 6.10.1, $t_n(\theta) \xrightarrow{d} Z \sim N(0, 1)$ and $|t_n(\theta)| \xrightarrow{d} |Z|$.

The asymptotic normality of Theorem 6.12.1 is used to justify confidence intervals and tests for the parameters.

6.13 Confidence Intervals

The OLS estimate $\hat{\beta}$ is a **point estimate** for β , meaning that $\hat{\beta}$ is a single value in \mathbb{R}^k . A broader concept is a **set estimate** C_n which is a collection of values in \mathbb{R}^k . When the parameter θ is real-valued then it is common to focus on intervals $C_n = [L_n, U_n]$ and which is called an **interval estimate** for θ . The goal of an interval estimate C_n is to contain the true value, e.g. $\theta \in C_n$, with high probability.

The interval estimate C_n is a function of the data and hence is random. The **coverage probability** of the interval $C_n = [L_n, U_n]$ is $\Pr_\theta(\theta \in C_n)$. The randomness comes from C_n as the parameter θ is treated as fixed.

Interval estimates C_n are typically called **confidence intervals** as the goal is typically to set the coverage probability to equal a pre-specified target, typically 90% or 95%. C_n is called a $(1 - \alpha)\%$ confidence interval if $\inf_\theta \Pr_\theta(\theta \in C_n) = 1 - \alpha$.

There is not a unique method to construct confidence intervals. For example, a simple (yet silly) interval is

$$C_n = \begin{cases} \mathbb{R} & \text{with probability } 1 - \alpha \\ \hat{\theta} & \text{with probability } \alpha \end{cases}$$

By construction, if $\hat{\theta}$ has a continuous distribution, $\Pr(\theta \in C_n) = 1 - \alpha$, so this confidence interval has perfect coverage, but C_n is uninformative about θ and is therefore not useful.

When we have an asymptotically normal parameter estimate $\hat{\theta}$ with standard error $s(\hat{\theta})$, the standard confidence interval for θ takes the form

$$C_n = \left[\hat{\theta} - c \cdot s(\hat{\theta}), \quad \hat{\theta} + c \cdot s(\hat{\theta}) \right] \quad (6.46)$$

where $c > 0$ is a pre-specified constant. This confidence interval is symmetric about the point estimate $\hat{\theta}$, and its length is proportional to the standard error $s(\hat{\theta})$.

Equivalently, C_n is the set of parameter values for θ such that the t-statistic $t_n(\theta)$ is smaller (in absolute value) than c , that is

$$C_n = \{\theta : |t_n(\theta)| \leq c\} = \left\{ \theta : -c \leq \frac{\hat{\theta} - \theta}{s(\hat{\theta})} \leq c \right\}.$$

The coverage probability of this confidence interval is

$$\Pr(\theta \in C_n) = \Pr(|t_n(\theta)| \leq c)$$

which is generally unknown. We can approximate the coverage probability by taking the asymptotic limit as $n \rightarrow \infty$. Since $|t_n(\theta)|$ is asymptotically $|Z|$ (Theorem 6.12.1), it follows that as $n \rightarrow \infty$ that

$$\Pr(\theta \in C_n) \rightarrow \Pr(|Z| \leq c) = \bar{\Phi}(c)$$

where $\bar{\Phi}(u)$ is given in (6.45). We call this the **asymptotic coverage probability**. Since the t-ratio is asymptotically pivotal, the asymptotic coverage probability is independent of the parameter θ , and is only a function of c .

As we mentioned before, an ideal confidence interval has a pre-specified probability coverage $1 - \alpha$, typically 90% or 95%. This means selecting the constant c so that

$$\bar{\Phi}(c) = 1 - \alpha.$$

Effectively, this makes c a function of α , and can be backed out of a normal distribution table. For example, $\alpha = 0.05$ (a 95% interval) implies $c = 1.96$ and $\alpha = 0.1$ (a 90% interval) implies $c = 1.645$. Rounding 1.96 to 2, we obtain the most commonly used confidence interval in applied econometric practice

$$C_n = \left[\hat{\theta} - 2s(\hat{\theta}), \quad \hat{\theta} + 2s(\hat{\theta}) \right]. \quad (6.47)$$

This is a useful rule-of thumb. This asymptotic 95% confidence interval C_n is simple to compute and can be roughly calculated from tables of coefficient estimates and standard errors. (Technically, it is an asymptotic 95.4% interval, due to the substitution of 2.0 for 1.96, but this distinction is overly precise.)

Theorem 6.13.1 *Under Assumptions 6.1.2 and 6.10.1, for C_n defined in (6.46), $\Pr(\theta \in C_n) \rightarrow \bar{\Phi}(c)$. For $c = 1.96$, $\Pr(\theta \in C_n) \rightarrow 0.95$.*

Confidence intervals are a simple yet effective tool to assess estimation uncertainty. When reading a set of empirical results, look at the estimated coefficient estimates and the standard errors. For a parameter of interest, compute the confidence interval C_n and consider the meaning of the spread of the suggested values. If the range of values in the confidence interval are too wide to learn about θ , then do not jump to a conclusion about θ based on the point estimate alone.

For illustration, consider the three examples presented in Section 6.11 based on the log wage regression for single asian men.

Percentage return to education. A 95% asymptotic confidence interval is $11.8 \pm 1.96 \times 0.8 = [10.2, 13.3]$.

Percentage return to experience for individuals with 12 years experience. A 90% asymptotic confidence interval is $1.1 \pm 1.645 \times 0.4 = [0.5, 1.8]$.

Experience level which maximizes expected log wages. An 80% asymptotic confidence interval is $35 \pm 1.28 \times 7 = [26, 44]$.

6.14 Regression Intervals

In the linear regression model the conditional mean of y_i given $\mathbf{x}_i = \mathbf{x}$ is

$$m(\mathbf{x}) = \mathbb{E}(y_i | \mathbf{x}_i = \mathbf{x}) = \mathbf{x}'\boldsymbol{\beta}.$$

In some cases, we want to estimate $m(\mathbf{x})$ at a particular point \mathbf{x} . Notice that this is a (linear) function of $\boldsymbol{\beta}$. Letting $r(\boldsymbol{\beta}) = \mathbf{x}'\boldsymbol{\beta}$ and $\theta = r(\boldsymbol{\beta})$, we see that $\hat{m}(\mathbf{x}) = \hat{\theta} = \mathbf{x}'\hat{\boldsymbol{\beta}}$ and $\mathbf{R} = \mathbf{x}$, so $s(\hat{\theta}) = \sqrt{\mathbf{x}'\hat{\mathbf{V}}_{\hat{\boldsymbol{\beta}}}\mathbf{x}}$. Thus an asymptotic 95% confidence interval for $m(\mathbf{x})$ is

$$\left[\mathbf{x}'\hat{\boldsymbol{\beta}} \pm 1.96\sqrt{\mathbf{x}'\hat{\mathbf{V}}_{\hat{\boldsymbol{\beta}}}\mathbf{x}} \right].$$

It is interesting to observe that if this is viewed as a function of \mathbf{x} , the width of the confidence set is dependent on \mathbf{x} .

To illustrate, we return to the log wage regression (3.11) of Section 3.7. The estimated regression equation is

$$\log(\widehat{Wage}) = \mathbf{x}'\hat{\boldsymbol{\beta}} = 0.155x + 0.698$$

where $x = \text{education}$. The covariance matrix estimate from (4.35) is

$$\hat{\mathbf{V}}_{\hat{\boldsymbol{\beta}}} = \begin{pmatrix} 0.001 & -0.015 \\ -0.015 & 0.243 \end{pmatrix}.$$

Thus the 95% confidence interval for the regression takes the form

$$0.155x + 0.698 \pm 1.96\sqrt{0.001x^2 - 0.030x + 0.243}.$$

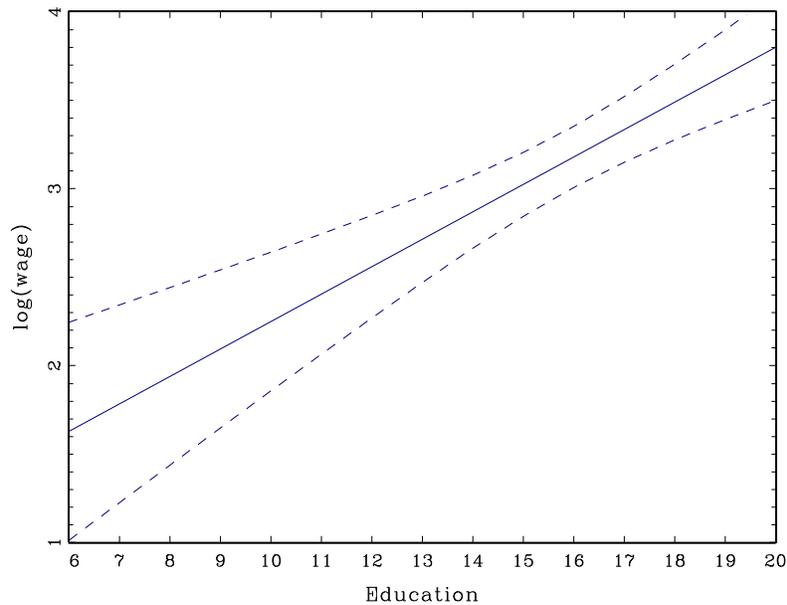


Figure 6.6: Wage on Education Regression Intervals

The estimated regression and 95% intervals are shown in Figure 6.6. Notice that the confidence bands take a hyperbolic shape. This means that the regression line is less precisely estimated for very large and very small values of *education*.

Plots of the estimated regression line and confidence intervals are especially useful when the regression includes nonlinear terms. To illustrate, consider the log wage regression (6.42) which includes experience and its square, with covariance matrix (6.43). We are interested in plotting the regression estimate and regression intervals as a function of *experience*. Since the regression also includes *education*, to plot the estimates in a simple graph we need to fix *education* at a specific value. We select *education*=12. This only affects the level of the estimated regression, since *education* enters without an interaction. Define the points of evaluation

$$\mathbf{z}(x) = \begin{pmatrix} 12 \\ x \\ x^2/100 \\ 1 \end{pmatrix}$$

where $x = \textit{experience}$.

Thus the 95% regression interval for *education*=12, as a function of $x = \textit{experience}$ is

$$\begin{aligned} & 0.118 \times 12 + 0.016 x - 0.022 x^2/100 + 0.947 \\ & \pm 1.96 \sqrt{\mathbf{z}(x)' \begin{pmatrix} 0.632 & 0.131 & -0.143 & -11.1 \\ 0.131 & 0.390 & -0.731 & -6.25 \\ -0.143 & -0.731 & 1.48 & 9.43 \\ -11.1 & -6.25 & 9.43 & 246 \end{pmatrix} \mathbf{z}(x) \times 10^{-4}} \\ & = 0.016 x - .00022 x^2 + 2.36 \\ & \pm 0.0196 \sqrt{70.608 - 9.356 x + 0.54428 x^2 - 0.01462 x^3 + 0.000148 x^4} \end{aligned}$$

The estimated regression and 95% intervals are shown in Figure 6.7. The regression interval widens greatly for small and large values of experience, indicating considerable uncertainty about the effect of experience on mean wages for this population. The confidence bands take a more complicated shape than in Figure 6.6 due to the nonlinear specification.

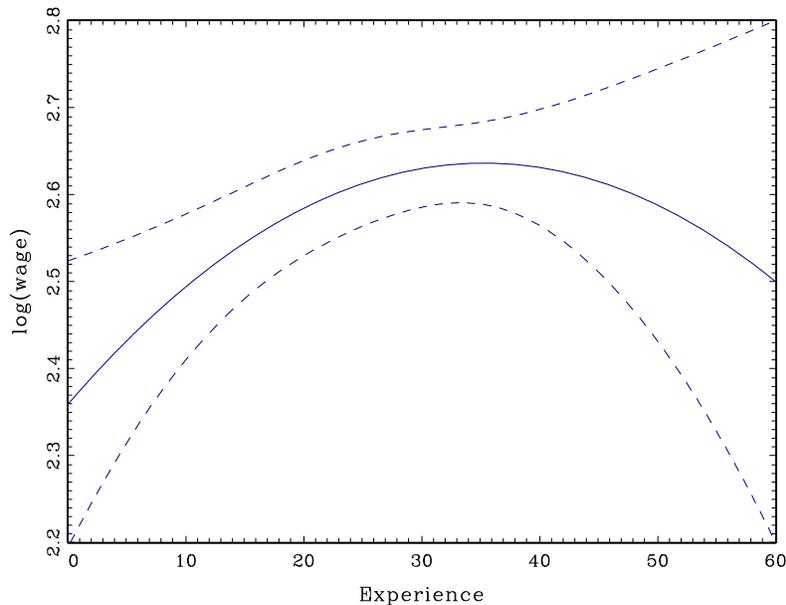


Figure 6.7: Wage on Experience Regression Intervals

6.15 Forecast Intervals

Suppose we are given a value of the regressor vector \mathbf{x}_{n+1} for an individual outside the sample, and we want to forecast (guess) y_{n+1} for this individual. This is equivalent to forecasting y_{n+1} given $\mathbf{x}_{n+1} = \mathbf{x}$, which will generally be a function of \mathbf{x} . A reasonable forecasting rule is the conditional mean $m(\mathbf{x})$ as it is the mean-square-minimizing forecast. A point forecast is the estimated conditional mean $\hat{m}(\mathbf{x}) = \mathbf{x}'\hat{\boldsymbol{\beta}}$. We would also like a measure of uncertainty for the forecast.

The forecast error is $\hat{e}_{n+1} = y_{n+1} - \hat{m}(\mathbf{x}) = e_{n+1} - \mathbf{x}'(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})$. As the out-of-sample error e_{n+1} is independent of the in-sample estimate $\hat{\boldsymbol{\beta}}$, this has variance

$$\begin{aligned}\mathbb{E}\hat{e}_{n+1}^2 &= \mathbb{E}(e_{n+1}^2 | \mathbf{x}_{n+1} = \mathbf{x}) + \mathbf{x}'\mathbb{E}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})' \mathbf{x} \\ &= \sigma^2(\mathbf{x}) + \mathbf{x}'\mathbf{V}_{\hat{\boldsymbol{\beta}}}\mathbf{x}.\end{aligned}$$

Assuming $\mathbb{E}(e_{n+1}^2 | \mathbf{x}_{n+1}) = \sigma^2$, the natural estimate of this variance is $\hat{\sigma}^2 + \mathbf{x}'\hat{\mathbf{V}}_{\hat{\boldsymbol{\beta}}}\mathbf{x}$, so a standard error for the forecast is $\hat{s}(\mathbf{x}) = \sqrt{\hat{\sigma}^2 + \mathbf{x}'\hat{\mathbf{V}}_{\hat{\boldsymbol{\beta}}}\mathbf{x}}$. Notice that this is different from the standard error for the conditional mean.

The conventional 95% forecast interval for y_{n+1} uses a normal approximation and sets

$$\left[\mathbf{x}'\hat{\boldsymbol{\beta}} \pm 2\hat{s}(\mathbf{x}) \right].$$

It is difficult, however, to fully justify this choice. It would be correct if we have a normal approximation to the ratio

$$\frac{e_{n+1} - \mathbf{x}'(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})}{\hat{s}(\mathbf{x})}.$$

The difficulty is that the equation error e_{n+1} is generally non-normal, and asymptotic theory cannot be applied to a single observation. The only special exception is the case where e_{n+1} has the exact distribution $N(0, \sigma^2)$, which is generally invalid.

To get an accurate forecast interval, we need to estimate the conditional distribution of e_{n+1} given $\mathbf{x}_{n+1} = \mathbf{x}$, which is a much more difficult task. Perhaps due to this difficulty, many applied

forecasters use the simple approximate interval $\left[\mathbf{x}'\hat{\boldsymbol{\beta}} \pm 2\hat{s}(\mathbf{x}) \right]$ despite the lack of a convincing justification.

6.16 Wald Statistic

Let $\boldsymbol{\theta} = \mathbf{r}(\boldsymbol{\beta}) : \mathbb{R}^k \rightarrow \mathbb{R}^q$ be any parameter vector of interest, $\hat{\boldsymbol{\theta}}$ its estimate and $\hat{\mathbf{V}}_{\hat{\boldsymbol{\theta}}}$ its covariance matrix estimator. Consider the quadratic form

$$W_n(\boldsymbol{\theta}) = (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})' \hat{\mathbf{V}}_{\hat{\boldsymbol{\theta}}}^{-1} (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) = n (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})' \hat{\mathbf{V}}_{\boldsymbol{\theta}}^{-1} (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}). \quad (6.48)$$

where $\hat{\mathbf{V}}_{\boldsymbol{\theta}} = n \hat{\mathbf{V}}_{\hat{\boldsymbol{\theta}}}$. When $q = 1$, then $W_n(\boldsymbol{\theta}) = t_n(\boldsymbol{\theta})^2$ is the square of the t-ratio. When $q > 1$, $W_n(\boldsymbol{\theta})$ is typically called a **Wald statistic**. We are interested in its sampling distribution.

The asymptotic distribution of $W_n(\boldsymbol{\theta})$ is simple to derive given Theorem 6.10.2 and Theorem 6.10.3, which show that

$$\sqrt{n} (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \xrightarrow{d} \mathbf{Z} \sim \mathbf{N}(\mathbf{0}, \mathbf{V}_{\boldsymbol{\theta}})$$

and

$$\hat{\mathbf{V}}_{\boldsymbol{\theta}} \xrightarrow{p} \mathbf{V}_{\boldsymbol{\theta}}.$$

It follows that

$$W_n(\boldsymbol{\theta}) = \sqrt{n} (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})' \hat{\mathbf{V}}_{\boldsymbol{\theta}}^{-1} \sqrt{n} (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \xrightarrow{d} \mathbf{Z}' \mathbf{V}_{\boldsymbol{\theta}}^{-1} \mathbf{Z} \quad (6.49)$$

a quadratic in the normal random vector \mathbf{Z} . Here we can appeal to a useful result from probability theory. (See Theorem B.9.3 in the Appendix.)

Theorem 6.16.1 *If $\mathbf{Z} \sim \mathbf{N}(\mathbf{0}, \mathbf{A})$ with $\mathbf{A} > 0$, $q \times q$, then $\mathbf{Z}' \mathbf{A}^{-1} \mathbf{Z} \sim \chi_q^2$, a chi-square random variable with q degrees of freedom.*

The asymptotic distribution in (6.49) takes exactly this form. Note that $\mathbf{V}_{\boldsymbol{\theta}} > 0$ since \mathbf{R} is full rank under Assumption 6.10.1 It follows that $W_n(\boldsymbol{\theta})$ converges in distribution to a chi-square random variable.

Theorem 6.16.2 *Under Assumptions 6.1.2 and 6.10.1, as $n \rightarrow \infty$,*

$$W_n(\boldsymbol{\theta}) \xrightarrow{d} \chi_q^2.$$

Theorem 6.16.2 is used to justify multivariate confidence regions and multivariate hypothesis tests.

6.17 Homoskedastic Wald Statistic

Under the conditional homoskedasticity assumption $\mathbb{E}(e_i^2 | \mathbf{x}_i) = \sigma^2$ we can construct the Wald statistic using the homoskedastic covariance matrix estimator $\hat{\mathbf{V}}_{\boldsymbol{\theta}}^0$ defined in (6.40). This yields a homoskedastic Wald statistic

$$W_n^0(\boldsymbol{\theta}) = (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})' (\hat{\mathbf{V}}_{\hat{\boldsymbol{\theta}}}^0)^{-1} (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) = n (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})' (\hat{\mathbf{V}}_{\boldsymbol{\theta}}^0)^{-1} (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}). \quad (6.50)$$

Under the additional assumption of conditional homoskedasticity, it has the same asymptotic distribution as $W_n(\boldsymbol{\theta})$.

Theorem 6.17.1 Under Assumptions 6.1.2 and 6.10.1, and $\mathbb{E}(e_i^2 | \mathbf{x}_i) = \sigma^2$, as $n \rightarrow \infty$,

$$W_n^0(\boldsymbol{\theta}) \xrightarrow{d} \chi_q^2.$$

6.18 Confidence Regions

A confidence region C_n is a set estimator for $\boldsymbol{\theta} \in \mathbb{R}^q$ when $q > 1$. A confidence region C_n is a set in \mathbb{R}^q intended to cover the true parameter value with a pre-selected probability $1 - \alpha$. Thus an ideal confidence region has the coverage probability $\Pr(\boldsymbol{\theta} \in C_n) = 1 - \alpha$. In practice it is typically not possible to construct a region with exact coverage, but we can calculate its asymptotic coverage.

When the parameter estimate satisfies the conditions of Theorem 6.16.2, a good choice for a confidence region is the ellipse

$$C_n = \{\boldsymbol{\theta} : W_n(\boldsymbol{\theta}) \leq c_{1-\alpha}\}.$$

with $c_{1-\alpha}$ the $1 - \alpha$ 'th quantile of the χ_q^2 distribution. (Thus $F_q(c_{1-\alpha}) = 1 - \alpha$.) These quantiles can be found from the χ_q^2 critical value table.

Theorem 6.16.2 implies

$$\Pr(\boldsymbol{\theta} \in C_n) \rightarrow \Pr(\chi_q^2 \leq c_{1-\alpha}) = 1 - \alpha$$

which shows that C_n has asymptotic coverage $(1 - \alpha)\%$.

To illustrate the construction of a confidence region, consider the estimated regression (6.42) of the model

$$\log(\widehat{Wage}) = \beta_1 \text{education} + \beta_2 \text{experience} + \beta_3 \text{experience}^2/100 + \beta_4.$$

Suppose that the two parameters of interest are the percentage return to education $\theta_1 = 100\beta_1$ and the percentage return to experience for individuals with 10 years experience $\theta_2 = 100\beta_2 + 20\beta_3$. These two parameters are a linear transformation of the regression parameters with point estimates

$$\widehat{\boldsymbol{\theta}} = \begin{pmatrix} 0 & 100 & 0 & 0 \\ 0 & 0 & 100 & 20 \end{pmatrix} \widehat{\boldsymbol{\beta}} = \begin{pmatrix} 11.8 \\ 1.2 \end{pmatrix},$$

and have the covariance matrix estimate

$$\begin{aligned} \widehat{\mathbf{V}}_{\widehat{\boldsymbol{\theta}}} &= \begin{pmatrix} 0 & 100 & 0 & 0 \\ 0 & 0 & 100 & 20 \end{pmatrix} \widehat{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}} \begin{pmatrix} 0 & 0 \\ 100 & 0 \\ 0 & 100 \\ 0 & 20 \end{pmatrix} \\ &= \begin{pmatrix} 0.632 & 0.103 \\ 0.103 & 0.157 \end{pmatrix} \end{aligned}$$

with inverse

$$\widehat{\mathbf{V}}_{\widehat{\boldsymbol{\theta}}}^{-1} = \begin{pmatrix} 1.77 & -1.16 \\ -1.16 & 7.13 \end{pmatrix}.$$

Thus the Wald statistic is

$$\begin{aligned} W_n(\boldsymbol{\theta}) &= (\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta})' \widehat{\mathbf{V}}_{\widehat{\boldsymbol{\theta}}}^{-1} (\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \\ &= \begin{pmatrix} 11.8 - \theta_1 \\ 1.2 - \theta_2 \end{pmatrix}' \begin{pmatrix} 1.77 & -1.16 \\ -1.16 & 7.13 \end{pmatrix} \begin{pmatrix} 11.8 - \theta_1 \\ 1.2 - \theta_2 \end{pmatrix} \\ &= 1.77(11.8 - \theta_1)^2 - 2.32(11.8 - \theta_1)(1.2 - \theta_2) + 7.13(1.2 - \theta_2)^2. \end{aligned}$$

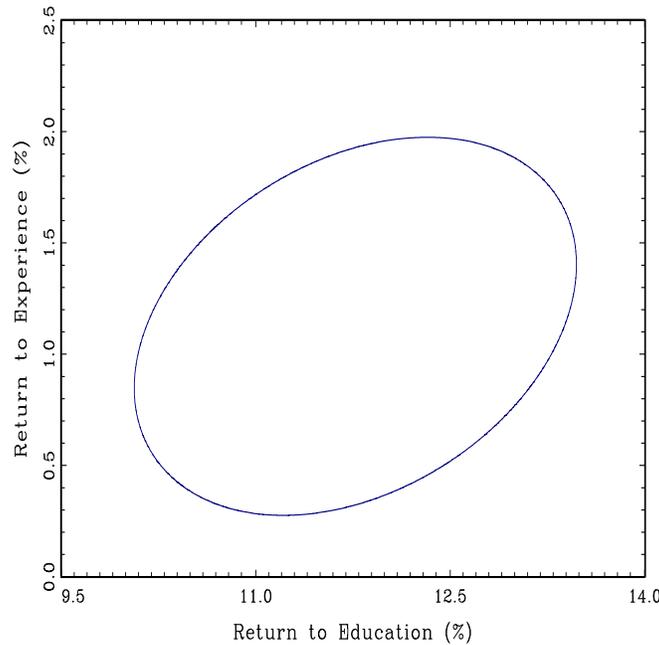


Figure 6.8: Confidence Region for Return to Experience and Return to Education

The 90% quantile of the χ_2^2 distribution is 4.605 (we use the χ_2^2 distribution as the dimension of θ is two), so an asymptotic 90% confidence region for the two parameters is the interior of the ellipse $W_n(\theta) = 4.605$ which is displayed in Figure 6.8. Since the estimated correlation of the two coefficient estimates is modest (about 0.3) the region is modestly elliptical.

6.19 Semiparametric Efficiency in the Projection Model

In Section 4.6 we presented the Gauss-Markov theorem, which stated that in the homoskedastic CEF model, in the class of linear unbiased estimators the one with the smallest variance is least-squares. As we noted in that section, the restriction to linear unbiased estimators is unsatisfactory as it leaves open the possibility that an alternative (non-linear) estimator could have a smaller asymptotic variance. In addition, the restriction to the homoskedastic CEF model is also unsatisfactory as the projection model is more relevant for empirical application. The question remains: what is the most efficient estimator of the projection coefficient β (or functions $\theta = h(\beta)$) in the projection model?

It turns out that it is straightforward to show that the projection model falls in the estimator class considered in Proposition 5.13.2. It follows that the least-squares estimator is semiparametrically efficient in the sense that it has the smallest asymptotic variance in the class of semiparametric estimators of β . This is a more powerful and interesting result than the Gauss-Markov theorem.

To see this, it is worth rephrasing Proposition 5.13.2 with amended notation. Suppose that a parameter of interest is $\theta = g(\mu)$ where $\mu = \mathbb{E}z_i$, for which the moment estimators are $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n z_i$ and $\hat{\theta} = g(\hat{\mu})$. Let $\mathcal{L}_2(g) = \left\{ F : \mathbb{E} \|z\|^2 < \infty, g(u) \text{ is continuously differentiable at } u = \mathbb{E}z \right\}$ be the set of distributions for which $\hat{\theta}$ satisfies the central limit theorem.

Proposition 6.19.1 *In the class of distributions $F \in \mathcal{L}_2(g)$, $\hat{\theta}$ is semiparametrically efficient for θ in the sense that its asymptotic variance equals the semiparametric efficiency bound.*

Proposition 6.19.1 says that under the minimal conditions in which $\hat{\theta}$ is asymptotically normal, then no semiparametric estimator can have a smaller asymptotic variance than $\hat{\theta}$.

To show that an estimator is semiparametrically efficient it is sufficient to show that it falls in the class covered by this Proposition. To show that the projection model falls in this class, we write $\beta = \mathbf{Q}_{xx}^{-1} \mathbf{Q}_{xy} = \mathbf{g}(\mu)$ where $\mu = \mathbb{E}z_i$ and $z_i = (x_i x_i', x_i y_i)$. The class $\mathcal{L}_2(\mathbf{g})$ equals the class of distributions

$$\mathcal{L}_4(\beta) = \left\{ F : \mathbb{E}y^4 < \infty, \mathbb{E}\|x\|^4 < \infty, \mathbb{E}x_i x_i' > 0 \right\}.$$

Proposition 6.19.2 *In the class of distributions $F \in \mathcal{L}_4(\beta)$, the least-squares estimator $\hat{\beta}$ is semiparametrically efficient for β .*

The least-squares estimator is an asymptotically efficient estimator of the projection coefficient because the latter is a smooth function of sample moments and the model implies no further restrictions. However, if the class of permissible distributions is restricted to a strict subset of $\mathcal{L}_4(\beta)$ then least-squares can be inefficient. For example, the linear CEF model with heteroskedastic errors is a strict subset of $\mathcal{L}_4(\beta)$, and the GLS estimator has a smaller asymptotic variance than OLS. In this case, the knowledge that true conditional mean is linear allows for more efficient estimation of the unknown parameter.

From Proposition 6.19.1 we can also deduce that plug-in estimators $\hat{\theta} = \mathbf{h}(\hat{\beta})$ are semiparametrically efficient estimators of $\theta = \mathbf{h}(\beta)$ when \mathbf{h} is continuously differentiable. We can also deduce that other parameters estimators are semiparametrically efficient, such as $\hat{\sigma}^2$ for σ^2 . To see this, note that we can write

$$\begin{aligned} \sigma^2 &= \mathbb{E} (y_i - x_i' \beta)^2 \\ &= \mathbb{E} y_i^2 - 2 \mathbb{E} (y_i x_i') \beta + \beta' \mathbb{E} (x_i x_i') \beta \\ &= Q_{yy} - \mathbf{Q}_{yx} \mathbf{Q}_{xx}^{-1} \mathbf{Q}_{xy} \end{aligned}$$

which is a smooth function of the moments Q_{yy} , Q_{yx} and Q_{xx} . Similarly the estimator $\hat{\sigma}^2$ equals

$$\begin{aligned} \hat{\sigma}^2 &= \frac{1}{n} \sum_{i=1}^n \hat{e}_i^2 \\ &= \hat{Q}_{yy} - \hat{Q}_{yx} \hat{Q}_{xx}^{-1} \hat{Q}_{xy} \end{aligned}$$

Since the variables y_i^2 , $y_i x_i'$ and $x_i x_i'$ all have finite variances when $F \in \mathcal{L}_4(\beta)$, the conditions of Proposition 6.19.1 are satisfied. We conclude:

Proposition 6.19.3 *In the class of distributions $F \in \mathcal{L}_4(\beta)$, $\hat{\sigma}^2$ is semiparametrically efficient for σ^2 .*

6.20 Semiparametric Efficiency in the Homoskedastic Regression Model*

In Section 6.19 we showed that the OLS estimator is semiparametrically efficient in the projection model. What if we restrict attention to the classical homoskedastic regression model? Is OLS still efficient in this class? In this section we derive the asymptotic semiparametric efficiency bound for this model, and show that it is the same as that obtained by the OLS estimator. Therefore it turns out that least-squares is efficient in this class as well.

Recall that in the homoskedastic regression model the asymptotic variance of the OLS estimator $\widehat{\beta}$ for β is $\mathbf{V}_{\beta}^0 = \mathbf{Q}_{xx}^{-1}\sigma^2$. Therefore, as described in Section 5.13, it is sufficient to find a parametric submodel whose Cramer-Rao bound for estimation of β is \mathbf{V}_{β}^0 . This would establish that \mathbf{V}_{β}^0 is the semiparametric variance bound and the OLS estimator $\widehat{\beta}$ is semiparametrically efficient for β .

Let the joint density of y and \mathbf{x} be written as $f(y, \mathbf{x}) = f_1(y | \mathbf{x}) f_2(\mathbf{x})$, the product of the conditional density of y given \mathbf{x} and the marginal density of \mathbf{x} . Now consider the parametric submodel

$$f(y, \mathbf{x} | \boldsymbol{\theta}) = f_1(y | \mathbf{x}) (1 + (y - \mathbf{x}'\boldsymbol{\beta})(\mathbf{x}'\boldsymbol{\theta})/\sigma^2) f_2(\mathbf{x}). \quad (6.51)$$

You can check that in this submodel the marginal density of \mathbf{x} is $f_2(\mathbf{x})$ and the conditional density of y given \mathbf{x} is $f_1(y | \mathbf{x}) (1 + (y - \mathbf{x}'\boldsymbol{\beta})(\mathbf{x}'\boldsymbol{\theta})/\sigma^2)$. To see that the latter is a valid conditional density, observe that the regression assumption implies that $\int y f_1(y | \mathbf{x}) dy = \mathbf{x}'\boldsymbol{\beta}$ and therefore

$$\begin{aligned} & \int f_1(y | \mathbf{x}) (1 + (y - \mathbf{x}'\boldsymbol{\beta})(\mathbf{x}'\boldsymbol{\theta})/\sigma^2) dy \\ &= \int f_1(y | \mathbf{x}) dy + \int f_1(y | \mathbf{x}) (y - \mathbf{x}'\boldsymbol{\beta}) dy (\mathbf{x}'\boldsymbol{\theta})/\sigma^2 \\ &= 1. \end{aligned}$$

In this parametric submodel the conditional mean of y given \mathbf{x} is

$$\begin{aligned} \mathbb{E}_{\boldsymbol{\theta}}(y | \mathbf{x}) &= \int y f_1(y | \mathbf{x}) (1 + (y - \mathbf{x}'\boldsymbol{\beta})(\mathbf{x}'\boldsymbol{\theta})/\sigma^2) dy \\ &= \int y f_1(y | \mathbf{x}) dy + \int y f_1(y | \mathbf{x}) (y - \mathbf{x}'\boldsymbol{\beta})(\mathbf{x}'\boldsymbol{\theta})/\sigma^2 dy \\ &= \int y f_1(y | \mathbf{x}) dy + \int (y - \mathbf{x}'\boldsymbol{\beta})^2 f_1(y | \mathbf{x}) (\mathbf{x}'\boldsymbol{\theta})/\sigma^2 dy \\ &+ \int (y - \mathbf{x}'\boldsymbol{\beta}) f_1(y | \mathbf{x}) dy (\mathbf{x}'\boldsymbol{\beta})(\mathbf{x}'\boldsymbol{\theta})/\sigma^2 \\ &= \mathbf{x}'(\boldsymbol{\beta} + \boldsymbol{\theta}), \end{aligned}$$

using the homoskedasticity assumption $\int (y - \mathbf{x}'\boldsymbol{\beta})^2 f_1(y | \mathbf{x}) dy = \sigma^2$. This means that in this parametric submodel, the conditional mean is linear in \mathbf{x} and the regression coefficient is $\boldsymbol{\beta}(\boldsymbol{\theta}) = \boldsymbol{\beta} + \boldsymbol{\theta}$.

We now calculate the score for estimation of $\boldsymbol{\theta}$. Since

$$\frac{\partial}{\partial \boldsymbol{\theta}} \log f(y, \mathbf{x} | \boldsymbol{\theta}) = \frac{\partial}{\partial \boldsymbol{\theta}} \log (1 + (y - \mathbf{x}'\boldsymbol{\beta})(\mathbf{x}'\boldsymbol{\theta})/\sigma^2) = \frac{\mathbf{x}(y - \mathbf{x}'\boldsymbol{\beta})/\sigma^2}{1 + (y - \mathbf{x}'\boldsymbol{\beta})(\mathbf{x}'\boldsymbol{\theta})/\sigma^2}$$

the score is

$$\mathbf{s} = \frac{\partial}{\partial \boldsymbol{\theta}} \log f(y, \mathbf{x} | \boldsymbol{\theta}_0) = \mathbf{x}e/\sigma^2.$$

The Cramer-Rao bound for estimation of $\boldsymbol{\theta}$ (and therefore $\boldsymbol{\beta}(\boldsymbol{\theta})$ as well) is

$$(\mathbb{E}(\mathbf{s}\mathbf{s}'))^{-1} = (\sigma^{-4}\mathbb{E}((\mathbf{x}e)(\mathbf{x}e)'))^{-1} = \sigma^2\mathbf{Q}_{xx}^{-1} = \mathbf{V}_{\beta}^0.$$

We have shown that there is a parametric submodel (6.51) whose Cramer-Rao bound for estimation of β is identical to the asymptotic variance of the least-squares estimator, which therefore is the semiparametric variance bound.

Theorem 6.20.1 *In the homoskedastic regression model, the semiparametric variance bound for estimation of β is $\mathbf{V}^0 = \sigma^2 \mathbf{Q}_{xx}^{-1}$ and the OLS estimator is semiparametrically efficient.*

This result is similar to the Gauss-Markov theorem, in that it asserts the efficiency of the least-squares estimator in the context of the homoskedastic regression model. The difference is that the Gauss-Markov theorem states that OLS has the smallest variance among the set of unbiased linear estimators, while Theorem 6.20.1 states that OLS has the smallest asymptotic variance among all regular estimators. This is a much more powerful statement.

6.21 Uniformly Consistent Residuals*

It seems natural to view the residuals \hat{e}_i as estimates of the unknown errors e_i . Are they consistent estimates? In this section we develop an appropriate convergence result. This is not a widely-used technique, and can safely be skipped by most readers.

Notice that we can write the residual as

$$\begin{aligned}\hat{e}_i &= y_i - \mathbf{x}'_i \hat{\beta} \\ &= e_i + \mathbf{x}'_i \beta - \mathbf{x}'_i \hat{\beta} \\ &= e_i - \mathbf{x}'_i (\hat{\beta} - \beta).\end{aligned}\tag{6.52}$$

Since $\hat{\beta} - \beta \xrightarrow{p} \mathbf{0}$ it seems reasonable to guess that \hat{e}_i will be close to e_i if n is large.

We can bound the difference in (6.52) using the Schwarz inequality (A.10) to find

$$|\hat{e}_i - e_i| = \left| \mathbf{x}'_i (\hat{\beta} - \beta) \right| \leq \|\mathbf{x}_i\| \|\hat{\beta} - \beta\|.\tag{6.53}$$

To bound (6.53) we can use $\|\hat{\beta} - \beta\| = O_p(n^{-1/2})$ from Theorem 6.3.2, but we also need to bound the random variable $\|\mathbf{x}_i\|$. If the regressor is bounded, that is, $\|\mathbf{x}_i\| \leq B < \infty$, then $|\hat{e}_i - e_i| \leq B \|\hat{\beta} - \beta\| = O_p(n^{-1/2})$. However if the regressor does not have bounded support then we have to be more careful.

The key is Theorem 5.12.1 which shows that $\mathbb{E} \|\mathbf{x}_i\|^r < \infty$ implies $\mathbf{x}_i = o_p(n^{1/r})$ uniformly in i , or

$$n^{-1/r} \max_{1 \leq i \leq n} \|\mathbf{x}_i\| \xrightarrow{p} 0.$$

Applied to (6.53) we obtain

$$\begin{aligned}\max_{1 \leq i \leq n} |\hat{e}_i - e_i| &\leq \max_{1 \leq i \leq n} \|\mathbf{x}_i\| \|\hat{\beta} - \beta\| \\ &= o_p(n^{-1/2+1/r}).\end{aligned}$$

We have shown the following.

Theorem 6.21.1 Under Assumption 6.1.2 and $\mathbb{E} \|\mathbf{x}_i\|^r < \infty$, then uniformly in $1 \leq i \leq n$

$$\hat{e}_i = e_i + o_p(n^{-1/2+1/r}). \quad (6.54)$$

The rate of convergence in (6.54) depends on r . Assumption 6.1.2 requires $r \geq 4$, so the rate of convergence is at least $o_p(n^{-1/4})$. As r increases, the rate improves. As a limiting case, from Theorem 5.12.1 we see that if $\mathbb{E} \exp(\mathbf{t}'\mathbf{x}_i) < \infty$ for all $\|\mathbf{t}\| < \infty$ then $\mathbf{x}_i = o_p(\log n)$ uniformly in i , and thus $\hat{e}_i = e_i + o_p(n^{-1/2} \log n)$.

We mentioned in Section 6.7 that there are multiple ways to prove the consistency of the covariance matrix estimator $\hat{\mathbf{\Omega}}$. We now show that Theorem 6.21.1 provides one simple method to establish (6.32) and thus Theorem 6.7.1. Let $q_n = \max_{1 \leq i \leq n} |\hat{e}_i - e_i| = o_p(n^{-1/4})$. Since

$$\hat{e}_i^2 - e_i^2 = 2e_i(\hat{e}_i - e_i) + (\hat{e}_i - e_i)^2,$$

then

$$\begin{aligned} \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' (\hat{e}_i^2 - e_i^2) \right\| &\leq \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i \mathbf{x}_i'\| |\hat{e}_i^2 - e_i^2| \\ &\leq \frac{2}{n} \sum_{i=1}^n \|\mathbf{x}_i\|^2 |e_i| |\hat{e}_i - e_i| + \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i\|^2 |\hat{e}_i - e_i|^2 \\ &\leq \frac{2}{n} \sum_{i=1}^n \|\mathbf{x}_i\|^2 |e_i| q_n + \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i\|^2 q_n^2 \\ &\leq o_p(n^{-1/4}). \end{aligned}$$

6.22 Asymptotic Leverage*

Recall the definition of leverage from (3.21)

$$h_{ii} = \mathbf{x}_i' (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i.$$

These are the diagonal elements of the projection matrix \mathbf{P} and appear in the formula for leave-one-out prediction errors and several covariance matrix estimators. We can show that under iid sampling the leverage values are uniformly asymptotically small.

Let $\lambda_{\min}(\mathbf{A})$ and $\lambda_{\max}(\mathbf{A})$ denote the smallest and largest eigenvalues of a symmetric square matrix \mathbf{A} , and note that $\lambda_{\max}(\mathbf{A}^{-1}) = (\lambda_{\min}(\mathbf{A}))^{-1}$.

Since $\frac{1}{n} \mathbf{X}'\mathbf{X} \xrightarrow{p} \mathbf{Q}_{xx} > 0$ then by the CMT, $\lambda_{\min}(\frac{1}{n} \mathbf{X}'\mathbf{X}) \xrightarrow{p} \lambda_{\min}(\mathbf{Q}_{xx}) > 0$. (The latter is positive since \mathbf{Q}_{xx} is positive definite and thus all its eigenvalues are positive.) Then by the Trace Inequality (A.13)

$$\begin{aligned} h_{ii} &= \mathbf{x}_i' (\mathbf{X}'\mathbf{X})^{-1} \mathbf{x}_i \\ &= \text{tr} \left(\left(\frac{1}{n} \mathbf{X}'\mathbf{X} \right)^{-1} \frac{1}{n} \mathbf{x}_i \mathbf{x}_i' \right) \\ &\leq \lambda_{\max} \left(\left(\frac{1}{n} \mathbf{X}'\mathbf{X} \right)^{-1} \right) \text{tr} \left(\frac{1}{n} \mathbf{x}_i \mathbf{x}_i' \right) \\ &= \left(\lambda_{\min} \left(\frac{1}{n} \mathbf{X}'\mathbf{X} \right) \right)^{-1} \frac{1}{n} \|\mathbf{x}_i\|^2 \\ &\leq (\lambda_{\min}(\mathbf{Q}_{xx}) + o_p(1))^{-1} \frac{1}{n} \max_{1 \leq i \leq n} \|\mathbf{x}_i\|^2. \end{aligned} \quad (6.55)$$

Theorem 5.12.1 shows that $\mathbb{E} \|\mathbf{x}_i\|^r < \infty$ implies $\max_{1 \leq i \leq n} \|\mathbf{x}_i\| = o_p(n^{1/r})$ and thus (6.55) is $o_p(n^{2/r-1})$.

Theorem 6.22.1 *If \mathbf{x}_i is independent and identically distributed and $\mathbb{E} \|\mathbf{x}_i\|^r < \infty$ for some $r \geq 2$, then uniformly in $1 \leq i \leq n$, $h_{ii} = o_p(n^{2/r-1})$.*

For any $r \geq 2$ then $h_{ii} = o_p(1)$ (uniformly in $i \leq n$). Larger r implies a stronger rate of convergence, for example $r = 4$ implies $h_{ii} = o_p(n^{-1/2})$.

Theorem (6.22.1) implies that under random sampling with finite variances and large samples, no individual observation should have a large leverage value. Consequently individual observations should not be influential, unless one of these conditions is violated.

Exercises

Exercise 6.1 Take the model $y_i = \mathbf{x}'_{1i}\boldsymbol{\beta}_1 + \mathbf{x}'_{2i}\boldsymbol{\beta}_2 + e_i$ with $\mathbb{E}\mathbf{x}_i e_i = \mathbf{0}$. Suppose that $\boldsymbol{\beta}_1$ is estimated by regressing y_i on \mathbf{x}_{1i} only. Find the probability limit of this estimator. In general, is it consistent for $\boldsymbol{\beta}_1$? If not, under what conditions is this estimator consistent for $\boldsymbol{\beta}_1$?

Exercise 6.2 Let \mathbf{y} be $n \times 1$, \mathbf{X} be $n \times k$ (rank k). $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}$ with $\mathbb{E}(\mathbf{x}_i e_i) = 0$. Define the *ridge regression* estimator

$$\widehat{\boldsymbol{\beta}} = \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}'_i + \lambda \mathbf{I}_k \right)^{-1} \left(\sum_{i=1}^n \mathbf{x}_i y_i \right) \quad (6.56)$$

where $\lambda > 0$ is a fixed constant. Find the probability limit of $\widehat{\boldsymbol{\beta}}$ as $n \rightarrow \infty$. Is $\widehat{\boldsymbol{\beta}}$ consistent for $\boldsymbol{\beta}$?

Exercise 6.3 For the ridge regression estimator (6.56), set $\lambda = cn$ where $c > 0$ is fixed as $n \rightarrow \infty$. Find the probability limit of $\widehat{\boldsymbol{\beta}}$ as $n \rightarrow \infty$.

Exercise 6.4 Verify some of the calculations reported in Section 6.4. Specifically, suppose that x_{1i} and x_{2i} only take the values $\{-1, +1\}$, symmetrically, with

$$\begin{aligned} \Pr(x_{1i} = x_{2i} = 1) &= \Pr(x_{1i} = x_{2i} = -1) = 3/8 \\ \Pr(x_{1i} = 1, x_{2i} = -1) &= \Pr(x_{1i} = -1, x_{2i} = 1) = 1/8 \\ \mathbb{E}(e_i^2 \mid x_{1i} = x_{2i}) &= \frac{5}{4} \\ \mathbb{E}(e_i^2 \mid x_{1i} \neq x_{2i}) &= \frac{1}{4}. \end{aligned}$$

Verify the following:

- (a) $\mathbb{E}x_{1i} = 0$
- (b) $\mathbb{E}x_{1i}^2 = 1$
- (c) $\mathbb{E}x_{1i}x_{2i} = \frac{1}{2}$
- (d) $\mathbb{E}(e_i^2) = 1$
- (e) $\mathbb{E}(x_{1i}^2 e_i^2) = 1$
- (f) $\mathbb{E}(x_{1i}x_{2i}e_i^2) = \frac{7}{8}$.

Exercise 6.5 Show (6.20)-(6.23).

Exercise 6.6 The model is

$$\begin{aligned} y_i &= \mathbf{x}'_i \boldsymbol{\beta} + e_i \\ \mathbb{E}(\mathbf{x}_i e_i) &= \mathbf{0} \\ \boldsymbol{\Omega} &= \mathbb{E}(\mathbf{x}_i \mathbf{x}'_i e_i^2). \end{aligned}$$

Find the method of moments estimators $(\widehat{\boldsymbol{\beta}}, \widehat{\boldsymbol{\Omega}})$ for $(\boldsymbol{\beta}, \boldsymbol{\Omega})$.

- (a) In this model, are $(\widehat{\boldsymbol{\beta}}, \widehat{\boldsymbol{\Omega}})$ efficient estimators of $(\boldsymbol{\beta}, \boldsymbol{\Omega})$?
- (b) If so, in what sense are they efficient?

Exercise 6.7 Of the variables $(y_i^*, y_i, \mathbf{x}_i)$ only the pair (y_i, \mathbf{x}_i) are observed. In this case, we say that y_i^* is a *latent* variable. Suppose

$$\begin{aligned}y_i^* &= \mathbf{x}_i' \boldsymbol{\beta} + e_i \\ \mathbb{E}(\mathbf{x}_i e_i) &= \mathbf{0} \\ y_i &= y_i^* + u_i\end{aligned}$$

where u_i is a measurement error satisfying

$$\begin{aligned}\mathbb{E}(\mathbf{x}_i u_i) &= \mathbf{0} \\ \mathbb{E}(y_i^* u_i) &= 0\end{aligned}$$

Let $\widehat{\boldsymbol{\beta}}$ denote the OLS coefficient from the regression of y_i on \mathbf{x}_i .

- Is $\boldsymbol{\beta}$ the coefficient from the linear projection of y_i on \mathbf{x}_i ?
- Is $\widehat{\boldsymbol{\beta}}$ consistent for $\boldsymbol{\beta}$ as $n \rightarrow \infty$?
- Find the asymptotic distribution of $\sqrt{n}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta})$ as $n \rightarrow \infty$.

Exercise 6.8 Find the asymptotic distribution of $\sqrt{n}(\widehat{\sigma}^2 - \sigma^2)$ as $n \rightarrow \infty$.

Exercise 6.9 The model is

$$\begin{aligned}y_i &= x_i \beta + e_i \\ \mathbb{E}(e_i | x_i) &= 0\end{aligned}$$

where $x_i \in \mathbb{R}$. Consider the two estimators

$$\begin{aligned}\widehat{\beta} &= \frac{\sum_{i=1}^n x_i y_i}{\sum_{i=1}^n x_i^2} \\ \widetilde{\beta} &= \frac{1}{n} \sum_{i=1}^n \frac{y_i}{x_i}.\end{aligned}$$

- Under the stated assumptions, are both estimators consistent for β ?
- Are there conditions under which either estimator is efficient?

Exercise 6.10 In the homoskedastic regression model $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}$ with $\mathbb{E}(e_i | \mathbf{x}_i) = 0$ and $\mathbb{E}(e_i^2 | \mathbf{x}_i) = \sigma^2$, suppose $\widehat{\boldsymbol{\beta}}$ is the OLS estimate of $\boldsymbol{\beta}$ with covariance matrix estimate $\widehat{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}}$, based on a sample of size n . Let $\widehat{\sigma}^2$ be the estimate of σ^2 . You wish to forecast an out-of-sample value of y_{n+1} given that $\mathbf{x}_{n+1} = \mathbf{x}$. Thus the available information is the sample (\mathbf{y}, \mathbf{X}) , the estimates $(\widehat{\boldsymbol{\beta}}, \widehat{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}}, \widehat{\sigma}^2)$, the residuals $\widehat{\mathbf{e}}$, and the out-of-sample value of the regressors, \mathbf{x}_{n+1} .

- Find a point forecast of y_{n+1} .
- Find an estimate of the variance of this forecast

Exercise 6.11 As in Exercise 3.21, use the CPS dataset and the subsample of white male Hispanics. Estimate the regression

$$\log(\widehat{Wage}) = \beta_1 \text{education} + \beta_2 \text{experience} + \beta_3 \text{experience}^2/100 + \beta_4.$$

- Report the coefficients and robust standard errors.

- (b) Let θ be the ratio of the return to one year of education to the return to one year of experience. Write θ as a function of the regression coefficients and variables. Compute $\hat{\theta}$ from the estimated model.
- (c) Write out the formula for the asymptotic standard error for $\hat{\theta}$ as a function of the covariance matrix for $\hat{\beta}$. Compute $\hat{s}(\hat{\theta})$ from the estimated model.
- (d) Construct a 90% asymptotic confidence interval for θ from the estimated model.
- (e) Compute the regression function at $edu = 12$ and $experience = 20$. Compute a 95% confidence interval for the regression function at this point.
- (f) Consider an out-of-sample individual with 16 years of education and 5 years experience. Construct an 80% forecast interval for their log wage and wage. [To obtain the forecast interval for the wage, apply the exponential function to both endpoints.]

Chapter 7

Restricted Estimation

7.1 Introduction

In the linear projection model

$$\begin{aligned}y_i &= \mathbf{x}'_i \boldsymbol{\beta} + e_i \\ \mathbb{E}(\mathbf{x}_i e_i) &= 0\end{aligned}$$

a common task is to impose a constraint on the coefficient vector $\boldsymbol{\beta}$. For example, partitioning $\mathbf{x}'_i = (\mathbf{x}'_{1i}, \mathbf{x}'_{2i})$ and $\boldsymbol{\beta}' = (\boldsymbol{\beta}'_1, \boldsymbol{\beta}'_2)$, a typical constraint is an exclusion restriction of the form $\boldsymbol{\beta}_2 = \mathbf{0}$. In this case the constrained model is

$$\begin{aligned}y_i &= \mathbf{x}'_{1i} \boldsymbol{\beta}_1 + e_i \\ \mathbb{E}(\mathbf{x}_i e_i) &= 0\end{aligned}$$

At first glance this appears the same as the linear projection model, but there is one important difference: the error e_i is uncorrelated with the entire regressor vector $\mathbf{x}'_i = (\mathbf{x}'_{1i}, \mathbf{x}'_{2i})$ not just the included regressor \mathbf{x}_{1i} .

In general, a set of q linear constraints on $\boldsymbol{\beta}$ takes the form

$$\mathbf{R}' \boldsymbol{\beta} = \mathbf{c} \tag{7.1}$$

where \mathbf{R} is $k \times q$, $\text{rank}(\mathbf{R}) = q < k$ and \mathbf{c} is $q \times 1$. The assumption that \mathbf{R} is full rank means that the constraints are linearly independent (there are no redundant or contradictory constraints).

The constraint $\boldsymbol{\beta}_2 = \mathbf{0}$ discussed above is a special case of the constraint (7.1) with

$$\mathbf{R} = \begin{pmatrix} \mathbf{0} \\ \mathbf{I} \end{pmatrix}, \tag{7.2}$$

a selector matrix, and $\mathbf{c} = \mathbf{0}$.

Another common restriction is that a set of coefficients sum to a known constant, i.e. $\beta_1 + \beta_2 = 1$. This constraint arises in a constant-return-to-scale production function. Other common restrictions include the equality of coefficients $\beta_1 = \beta_2$, and equal and offsetting coefficients $\beta_1 = -\beta_2$.

A typical reason to impose a constraint is that we believe (or have information) that the constraint is true. By imposing the constraint we hope to improve estimation efficiency. The goal is to obtain consistent estimates with reduced variance relative to the unconstrained estimator.

The questions then arise: How should we estimate the coefficient vector $\boldsymbol{\beta}$ imposing the linear restriction (7.1)? If we impose such constraints, what is the sampling distribution of the resulting estimator? How should we calculate standard errors? These are the questions explored in this chapter.

7.2 Constrained Least Squares

An intuitively appealing method to estimate a constrained linear projection is to minimize the least-squares criterion subject to the constraint $\mathbf{R}'\boldsymbol{\beta} = \mathbf{c}$.

The constrained least-squares estimator is

$$\tilde{\boldsymbol{\beta}}_{\text{cls}} = \underset{\mathbf{R}'\boldsymbol{\beta}=\mathbf{c}}{\operatorname{argmin}} SSE_n(\boldsymbol{\beta}) \quad (7.3)$$

where

$$SSE_n(\boldsymbol{\beta}) = \sum_{i=1}^n (y_i - \mathbf{x}'_i\boldsymbol{\beta})^2 = \mathbf{y}'\mathbf{y} - 2\mathbf{y}'\mathbf{X}\boldsymbol{\beta} + \boldsymbol{\beta}'\mathbf{X}'\mathbf{X}\boldsymbol{\beta}. \quad (7.4)$$

The estimator $\tilde{\boldsymbol{\beta}}_{\text{cls}}$ minimizes the sum of squared errors over all $\boldsymbol{\beta}$ such that $\boldsymbol{\beta} \in \mathbf{B}_{\mathbf{R}}$, or equivalently such that the restriction (7.1) holds. We call $\tilde{\boldsymbol{\beta}}_{\text{cls}}$ the **constrained least-squares** (CLS) estimator. We follow the convention of using a tilde “~” rather than a hat “^” to indicate that $\tilde{\boldsymbol{\beta}}_{\text{cls}}$ is a restricted estimator in contrast to the unrestricted least-squares estimator $\hat{\boldsymbol{\beta}}$, and write it as $\tilde{\boldsymbol{\beta}}_{\text{cls}}$ to be clear that the estimation method is CLS.

One method to find the solution to (7.3) uses the technique of Lagrange multipliers. The problem (7.3) is equivalent to the minimization of the Lagrangian

$$\mathcal{L}(\boldsymbol{\beta}, \boldsymbol{\lambda}) = \frac{1}{2}SSE_n(\boldsymbol{\beta}) + \boldsymbol{\lambda}'(\mathbf{R}'\boldsymbol{\beta} - \mathbf{c}) \quad (7.5)$$

over $(\boldsymbol{\beta}, \boldsymbol{\lambda})$, where $\boldsymbol{\lambda}$ is an $s \times 1$ vector of Lagrange multipliers. The first-order conditions for minimization of (7.5) are

$$\frac{\partial}{\partial \boldsymbol{\beta}} \mathcal{L}(\tilde{\boldsymbol{\beta}}_{\text{cls}}, \tilde{\boldsymbol{\lambda}}_{\text{cls}}) = -\mathbf{X}'\mathbf{y} + \mathbf{X}'\mathbf{X}\tilde{\boldsymbol{\beta}}_{\text{cls}} + \mathbf{R}\tilde{\boldsymbol{\lambda}}_{\text{cls}} = \mathbf{0} \quad (7.6)$$

and

$$\frac{\partial}{\partial \boldsymbol{\lambda}} \mathcal{L}(\tilde{\boldsymbol{\beta}}_{\text{cls}}, \tilde{\boldsymbol{\lambda}}_{\text{cls}}) = \mathbf{R}'\tilde{\boldsymbol{\beta}}_{\text{cls}} - \mathbf{c} = \mathbf{0}. \quad (7.7)$$

Premultiplying (7.6) by $\mathbf{R}'(\mathbf{X}'\mathbf{X})^{-1}$ we obtain

$$-\mathbf{R}'\hat{\boldsymbol{\beta}} + \mathbf{R}'\tilde{\boldsymbol{\beta}}_{\text{cls}} + \mathbf{R}'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{R}\tilde{\boldsymbol{\lambda}}_{\text{cls}} = \mathbf{0} \quad (7.8)$$

where $\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}$ is the unrestricted least-squares estimator. Imposing $\mathbf{R}'\tilde{\boldsymbol{\beta}}_{\text{cls}} - \mathbf{c} = \mathbf{0}$ from (7.7) and solving for $\tilde{\boldsymbol{\lambda}}_{\text{cls}}$ we find

$$\tilde{\boldsymbol{\lambda}}_{\text{cls}} = \left[\mathbf{R}'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{R} \right]^{-1} (\mathbf{R}'\hat{\boldsymbol{\beta}} - \mathbf{c}).$$

Notice that $(\mathbf{X}'\mathbf{X})^{-1} > 0$ and \mathbf{R} full rank imply that $\mathbf{R}'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{R} > 0$ and is hence invertible. (See Section A.8.)

Substituting this expression into (7.6) and solving for $\tilde{\boldsymbol{\beta}}_{\text{cls}}$ we find the solution to the constrained minimization problem (7.3)

$$\tilde{\boldsymbol{\beta}}_{\text{cls}} = \hat{\boldsymbol{\beta}} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{R} \left[\mathbf{R}'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{R} \right]^{-1} (\mathbf{R}'\hat{\boldsymbol{\beta}} - \mathbf{c}). \quad (7.9)$$

(See Exercise 7.4 to verify that (3.4) satisfies (7.1).)

This is a general formula for the CLS estimator. It also can be written as

$$\tilde{\boldsymbol{\beta}}_{\text{cls}} = \hat{\boldsymbol{\beta}} - \hat{\mathbf{Q}}_{xx}^{-1}\mathbf{R} \left(\mathbf{R}'\hat{\mathbf{Q}}_{xx}^{-1}\mathbf{R} \right)^{-1} (\mathbf{R}'\hat{\boldsymbol{\beta}} - \mathbf{c}). \quad (7.10)$$

Given $\tilde{\boldsymbol{\beta}}_{\text{cls}}$ the residuals are

$$\tilde{e}_i = y_i - \mathbf{x}'_i \tilde{\boldsymbol{\beta}}_{\text{cls}}. \quad (7.11)$$

The moment estimator of σ^2 is

$$\tilde{\sigma}_{\text{cls}}^2 = \frac{1}{n} \sum_{i=1}^n \tilde{e}_i^2.$$

A bias-corrected version of $\tilde{\sigma}_{\text{cls}}^2$ is

$$s_{\text{cls}}^2 = \frac{1}{n - k + q} \sum_{i=1}^n \tilde{e}_i^2.$$

You can show (See Exercise 7.6) that in the homoskedastic linear regression model under (7.1),

$$\mathbb{E}(s_{\text{cls}}^2 \mid \mathbf{X}) = \sigma^2 \quad (7.12)$$

so that s_{cls}^2 is unbiased for σ^2 .

7.3 Exclusion Restriction

While (7.9) is a general formula for the CLS estimator, in most cases the estimator can be found by applying least-squares to a reparameterized equation. To illustrate, let us return to the first example presented at the beginning of the chapter – a simple exclusion restriction. Recall the unconstrained model is

$$y_i = \mathbf{x}'_{1i} \boldsymbol{\beta}_1 + \mathbf{x}'_{2i} \boldsymbol{\beta}_2 + e_i \quad (7.13)$$

the exclusion restriction is $\boldsymbol{\beta}_2 = \mathbf{0}$, and the constrained equation is

$$y_i = \mathbf{x}'_{1i} \boldsymbol{\beta}_1 + e_i. \quad (7.14)$$

In this setting the CLS estimator is OLS of y_i on \mathbf{x}_{1i} . (See Exercise 7.1.) We can write this as

$$\tilde{\boldsymbol{\beta}}_1 = \left(\sum_{i=1}^n \mathbf{x}_{1i} \mathbf{x}'_{1i} \right)^{-1} \left(\sum_{i=1}^n \mathbf{x}_{1i} y_i \right). \quad (7.15)$$

The CLS estimator of the entire vector $\boldsymbol{\beta}' = (\boldsymbol{\beta}'_1, \boldsymbol{\beta}'_2)$ is

$$\tilde{\boldsymbol{\beta}} = \begin{pmatrix} \tilde{\boldsymbol{\beta}}_1 \\ \mathbf{0} \end{pmatrix}. \quad (7.16)$$

It is not immediately obvious, but (7.9) and (7.16) are algebraically (and numerically) equivalent. To see this, the first component of (7.9) with (7.2) is

$$\tilde{\boldsymbol{\beta}}_1 = (\mathbf{I} \quad \mathbf{0}) \left[\hat{\boldsymbol{\beta}} - \hat{\mathbf{Q}}_{xx}^{-1} \begin{pmatrix} \mathbf{0} \\ \mathbf{I} \end{pmatrix} \left[(\mathbf{0} \quad \mathbf{I}) \hat{\mathbf{Q}}_{xx}^{-1} \begin{pmatrix} \mathbf{0} \\ \mathbf{I} \end{pmatrix} \right]^{-1} (\mathbf{0} \quad \mathbf{I}) \hat{\boldsymbol{\beta}} \right].$$

Using (3.33) this equals

$$\begin{aligned} \tilde{\boldsymbol{\beta}}_1 &= \hat{\boldsymbol{\beta}}_1 - \hat{\mathbf{Q}}^{12} (\hat{\mathbf{Q}}^{22})^{-1} \hat{\boldsymbol{\beta}}_2 \\ &= \hat{\boldsymbol{\beta}}_1 + \hat{\mathbf{Q}}_{11.2}^{-1} \hat{\mathbf{Q}}_{12} \hat{\mathbf{Q}}_{22}^{-1} \hat{\mathbf{Q}}_{22.1} \hat{\boldsymbol{\beta}}_2 \\ &= \hat{\mathbf{Q}}_{11.2}^{-1} (\hat{\mathbf{Q}}_{1y} - \hat{\mathbf{Q}}_{12} \hat{\mathbf{Q}}_{22}^{-1} \hat{\mathbf{Q}}_{2y}) \\ &\quad + \hat{\mathbf{Q}}_{11.2}^{-1} \hat{\mathbf{Q}}_{12} \hat{\mathbf{Q}}_{22}^{-1} \hat{\mathbf{Q}}_{22.1} \hat{\mathbf{Q}}_{22.1}^{-1} (\hat{\mathbf{Q}}_{2y} - \hat{\mathbf{Q}}_{21} \hat{\mathbf{Q}}_{11}^{-1} \hat{\mathbf{Q}}_{1y}) \\ &= \hat{\mathbf{Q}}_{11.2}^{-1} (\hat{\mathbf{Q}}_{1y} - \hat{\mathbf{Q}}_{12} \hat{\mathbf{Q}}_{22}^{-1} \hat{\mathbf{Q}}_{21} \hat{\mathbf{Q}}_{11}^{-1} \hat{\mathbf{Q}}_{1y}) \\ &= \hat{\mathbf{Q}}_{11.2}^{-1} (\hat{\mathbf{Q}}_{11} - \hat{\mathbf{Q}}_{12} \hat{\mathbf{Q}}_{22}^{-1} \hat{\mathbf{Q}}_{21}) \hat{\mathbf{Q}}_{11}^{-1} \hat{\mathbf{Q}}_{1y} \\ &= \hat{\mathbf{Q}}_{11}^{-1} \hat{\mathbf{Q}}_{1y} \end{aligned}$$

which is (7.16) as originally claimed.

7.4 Minimum Distance

A minimum distance estimator tries to find a parameter value which satisfies the constraint which is as close as possible to the unconstrained estimate. Let $\hat{\beta}$ be the unconstrained least-squares estimator, and for some $k \times k$ positive definite weight matrix $\mathbf{W}_n > 0$ define the quadratic criterion function

$$J_n(\beta) = n(\hat{\beta} - \beta)' \mathbf{W}_n (\hat{\beta} - \beta) \quad (7.17)$$

This is a (squared) weighted Euclidean distance between $\hat{\beta}$ and β . $J_n(\beta)$ is small if β is close to $\hat{\beta}$, and is minimized at zero only if $\beta = \hat{\beta}$. A **minimum distance estimator** $\tilde{\beta}_{\text{md}}$ for β minimizes $J_n(\beta)$ subject to the constraint (7.1), that is,

$$\tilde{\beta}_{\text{md}} = \underset{\mathbf{R}'\beta=c}{\operatorname{argmin}} J_n(\beta). \quad (7.18)$$

The CLS estimator is the special case when $\mathbf{W}_n = \hat{\mathbf{Q}}_{xx}$, and we write this criterion function as

$$J_n^0(\beta) = n(\hat{\beta} - \beta)' \hat{\mathbf{Q}}_{xx} (\hat{\beta} - \beta) \quad (7.19)$$

To see the equality of CLS and minimum distance, rewrite the least-squares criterion as follows. Write the unconstrained least-squares fitted equation as $y_i = \mathbf{x}'_i \hat{\beta} + \hat{\varepsilon}_i$ and substitute this equation into $SSE_n(\beta)$ to obtain

$$\begin{aligned} SSE_n(\beta) &= \sum_{i=1}^n (y_i - \mathbf{x}'_i \beta)^2 \\ &= \sum_{i=1}^n (\mathbf{x}'_i \hat{\beta} + \hat{\varepsilon}_i - \mathbf{x}'_i \beta)^2 \\ &= \sum_{i=1}^n \hat{\varepsilon}_i^2 + (\hat{\beta} - \beta)' \left(\sum_{i=1}^n \mathbf{x}_i \mathbf{x}'_i \right) (\hat{\beta} - \beta) \\ &= n\hat{\sigma}^2 + J_n^0(\beta) \end{aligned} \quad (7.20)$$

where the third equality uses the fact that $\sum_{i=1}^n \mathbf{x}_i \hat{\varepsilon}_i = 0$, and the last line uses $\sum_{i=1}^n \mathbf{x}_i \mathbf{x}'_i = n\hat{\mathbf{Q}}_{xx}$. The expression (7.20) only depends on β through $J_n^0(\beta)$. Thus minimization of $SSE_n(\beta)$ and $J_n^0(\beta)$ are equivalent, and hence $\tilde{\beta}_{\text{md}} = \tilde{\beta}_{\text{cls}}$ when $\mathbf{W}_n = \hat{\mathbf{Q}}_{xx}$.

We can solve for $\tilde{\beta}_{\text{md}}$ explicitly by the method of Lagrange multipliers. The Lagrangian is

$$\mathcal{L}(\beta, \lambda) = \frac{1}{2} J_n(\beta, \mathbf{W}_n) + \lambda' (\mathbf{R}'\beta - c)$$

which is minimized over (β, λ) . The solution is

$$\tilde{\lambda}_{\text{md}} = n(\mathbf{R}'\mathbf{W}_n^{-1}\mathbf{R})^{-1} (\mathbf{R}'\hat{\beta} - c) \quad (7.21)$$

$$\tilde{\beta}_{\text{md}} = \hat{\beta} - \mathbf{W}_n^{-1}\mathbf{R}(\mathbf{R}'\mathbf{W}_n^{-1}\mathbf{R})^{-1} (\mathbf{R}'\hat{\beta} - c). \quad (7.22)$$

(See Exercise 7.7.) Comparing (7.22) with (7.10) we can see that $\tilde{\beta}_{\text{md}}$ specializes to $\tilde{\beta}_{\text{cls}}$ when we set $\mathbf{W}_n = \hat{\mathbf{Q}}_{xx}$.

An obvious question is which weight matrix \mathbf{W}_n is best. We will address this question after we derive the asymptotic distribution for a general weight matrix.

7.5 Asymptotic Distribution

We first show that the class of minimum distance estimators are consistent for the population parameters when the constraints are valid.

Assumption 7.5.1 $\mathbf{R}'\boldsymbol{\beta} = \mathbf{c}$ where \mathbf{R} is $k \times q$ with $\text{rank}(\mathbf{R}) = q$.

Assumption 7.5.2 $\mathbf{W}_n \xrightarrow{p} \mathbf{W} > 0$.

Theorem 7.5.1 Consistency

Under Assumptions 6.1.1, 7.5.1, and 7.5.2, $\tilde{\boldsymbol{\beta}}_{\text{md}} \xrightarrow{p} \boldsymbol{\beta}$ as $n \rightarrow \infty$.

For a proof, see Exercise 7.8.

Theorem 7.5.1 shows that consistency holds for any weight matrix with a positive definite limit, so the result includes the CLS estimator.

Similarly, the constrained estimators are asymptotically normally distributed.

Theorem 7.5.2 Asymptotic Normality

Under Assumptions 6.1.2, 7.5.1, and 7.5.2,

$$\sqrt{n} \left(\tilde{\boldsymbol{\beta}}_{\text{md}} - \boldsymbol{\beta} \right) \xrightarrow{d} \text{N}(\mathbf{0}, \mathbf{V}_{\boldsymbol{\beta}}(\mathbf{W})) \quad (7.23)$$

as $n \rightarrow \infty$, where

$$\begin{aligned} \mathbf{V}_{\boldsymbol{\beta}}(\mathbf{W}) &= \mathbf{V}_{\boldsymbol{\beta}} - \mathbf{W}^{-1} \mathbf{R} (\mathbf{R}' \mathbf{W}^{-1} \mathbf{R})^{-1} \mathbf{R}' \mathbf{V}_{\boldsymbol{\beta}} \\ &\quad - \mathbf{V}_{\boldsymbol{\beta}} \mathbf{R} (\mathbf{R}' \mathbf{W}^{-1} \mathbf{R})^{-1} \mathbf{R}' \mathbf{W}^{-1} \\ &\quad + \mathbf{W}^{-1} \mathbf{R} (\mathbf{R}' \mathbf{W}^{-1} \mathbf{R})^{-1} \mathbf{R}' \mathbf{V}_{\boldsymbol{\beta}} \mathbf{R} (\mathbf{R}' \mathbf{W}^{-1} \mathbf{R})^{-1} \mathbf{R}' \mathbf{W}^{-1} \end{aligned} \quad (7.24)$$

and $\mathbf{V}_{\boldsymbol{\beta}} = \mathbf{Q}_{xx}^{-1} \boldsymbol{\Omega} \mathbf{Q}_{xx}^{-1}$.

For a proof, see Exercise 7.9.

Theorem 7.5.2 shows that the minimum distance estimator is asymptotically normal for all positive definite weight matrices. The asymptotic variance depends on \mathbf{W} . The theorem includes the CLS estimator as a special case by setting $\mathbf{W} = \mathbf{Q}_{xx}$.

Theorem 7.5.3 Asymptotic Distribution of CLS Estimator

Under Assumptions 6.1.2 and 7.5.1, as $n \rightarrow \infty$

$$\sqrt{n} \left(\tilde{\beta}_{\text{cls}} - \beta \right) \xrightarrow{d} N(\mathbf{0}, \mathbf{V}_{\text{cls}})$$

where

$$\begin{aligned} \mathbf{V}_{\text{cls}} &= \mathbf{V}_{\beta} - \mathbf{Q}_{xx}^{-1} \mathbf{R} (\mathbf{R}' \mathbf{Q}_{xx}^{-1} \mathbf{R})^{-1} \mathbf{R}' \mathbf{V}_{\beta} \\ &\quad - \mathbf{V}_{\beta} \mathbf{R} (\mathbf{R}' \mathbf{Q}_{xx}^{-1} \mathbf{R})^{-1} \mathbf{R}' \mathbf{Q}_{xx}^{-1} \\ &\quad + \mathbf{Q}_{xx}^{-1} \mathbf{R} (\mathbf{R}' \mathbf{Q}_{xx}^{-1} \mathbf{R})^{-1} \mathbf{R}' \mathbf{V}_{\beta} \mathbf{R} (\mathbf{R}' \mathbf{Q}_{xx}^{-1} \mathbf{R})^{-1} \mathbf{R}' \mathbf{Q}_{xx}^{-1} \end{aligned}$$

For a proof, see Exercise 7.10.

7.6 Efficient Minimum Distance Estimator

Theorem 7.5.2 shows that the minimum distance estimators, which include CLS as a special case, are asymptotically normal with an asymptotic covariance matrix which depends on the weight matrix \mathbf{W} . The asymptotically optimal weight matrix is the one which minimizes the asymptotic variance $\mathbf{V}_{\beta}(\mathbf{W})$. This turns out to be $\mathbf{W} = \mathbf{V}_{\beta}^{-1}$ as is shown in Theorem 7.6.1 below. Since \mathbf{V}_{β}^{-1} is unknown this weight matrix cannot be used for a feasible estimator, but we can replace \mathbf{V}_{β}^{-1} with a consistent estimate $\hat{\mathbf{V}}_{\beta}^{-1}$ and the asymptotic distribution (and efficiency) are unchanged. We call the minimum distance estimator setting $\mathbf{W}_n = \hat{\mathbf{V}}_{\beta}^{-1}$ the **efficient minimum distance estimator** and takes the form

$$\tilde{\beta}_{\text{emd}} = \hat{\beta} - \hat{\mathbf{V}}_{\beta} \mathbf{R} \left(\mathbf{R}' \hat{\mathbf{V}}_{\beta} \mathbf{R} \right)^{-1} \left(\mathbf{R}' \hat{\beta} - \mathbf{c} \right). \quad (7.25)$$

The asymptotic distribution of (7.25) can be deduced from Theorem 7.5.2. (See Exercises 7.11 and 7.12.)

Theorem 7.6.1 Efficient Minimum Distance Estimator

Under Assumptions 6.1.2 and 7.5.1,

$$\sqrt{n} \left(\tilde{\beta}_{\text{emd}} - \beta \right) \xrightarrow{d} N(\mathbf{0}, \mathbf{V}_{\beta}^*)$$

as $n \rightarrow \infty$, where

$$\mathbf{V}_{\beta}^* = \mathbf{V}_{\beta} - \mathbf{V}_{\beta} \mathbf{R} (\mathbf{R}' \mathbf{V}_{\beta} \mathbf{R})^{-1} \mathbf{R}' \mathbf{V}_{\beta}. \quad (7.26)$$

Since

$$\mathbf{V}_{\beta}^* \leq \mathbf{V}_{\beta} \quad (7.27)$$

the estimator (7.25) has lower asymptotic variance than the unrestricted estimator. Furthermore, for any \mathbf{W} ,

$$\mathbf{V}_{\beta}^* \leq \mathbf{V}_{\beta}(\mathbf{W}) \quad (7.28)$$

so (7.25) is asymptotically efficient in the class of minimum distance estimators.

Theorem 7.6.1 shows that the minimum distance estimator with the smallest asymptotic variance is (7.25). One implication is that the constrained least squares estimator is generally inefficient. The interesting exception is the case of conditional homoskedasticity, in which case the optimal weight matrix is $\mathbf{W} = \mathbf{V}_\beta^{0-1}$ so in this case CLS is an efficient minimum distance estimator. Otherwise when the error is conditionally heteroskedastic, there are asymptotic efficiency gains by using minimum distance rather than least squares.

The fact that CLS is generally inefficient is counter-intuitive and requires some reflection to understand. Standard intuition suggests to apply the same estimation method (least squares) to the unconstrained and constrained models, and this is the most common empirical practice. But Theorem 7.6.1 shows that this is not the efficient estimation method. Instead, the efficient minimum distance estimator has a smaller asymptotic variance. Why? The reason is that the least-squares estimator does not make use of the regressor \mathbf{x}_{2i} . It ignores the information $\mathbb{E}(\mathbf{x}_{2i}e_i) = \mathbf{0}$. This information is relevant when the error is heteroskedastic and the excluded regressors are correlated with the included regressors.

Inequality (7.27) shows that the efficient minimum distance estimator $\tilde{\beta}_{\text{emd}}$ has a smaller asymptotic variance than the unrestricted least squares estimator $\hat{\beta}$. This means that estimation is more efficient by imposing correct restrictions when we use the minimum distance method.

7.7 Exclusion Restriction Revisited

We return to the example of estimation with a simple exclusion restriction. The model is

$$y_i = \mathbf{x}'_{1i}\beta_1 + \mathbf{x}'_{2i}\beta_2 + e_i$$

with the exclusion restriction $\beta_2 = \mathbf{0}$. We have introduced three estimators of β_1 . The first is unconstrained least-squares applied to (7.13), which can be written as

$$\hat{\beta}_1 = \hat{\mathbf{Q}}_{11 \cdot 2}^{-1} \hat{\mathbf{Q}}_{1y \cdot 2}.$$

From Theorem 6.34 and equation (6.21) its asymptotic variance is

$$\text{avar}(\hat{\beta}_1) = \mathbf{Q}_{11 \cdot 2}^{-1} (\boldsymbol{\Omega}_{11} - \mathbf{Q}_{12} \mathbf{Q}_{22}^{-1} \boldsymbol{\Omega}_{21} - \boldsymbol{\Omega}_{12} \mathbf{Q}_{22}^{-1} \mathbf{Q}_{21} + \mathbf{Q}_{12} \mathbf{Q}_{22}^{-1} \boldsymbol{\Omega}_{22} \mathbf{Q}_{22}^{-1} \mathbf{Q}_{21}) \mathbf{Q}_{11 \cdot 2}^{-1}.$$

The second estimator of β_1 is the CLS estimator, which can be written as

$$\tilde{\beta}_{1, \text{cls}} = \hat{\mathbf{Q}}_{11}^{-1} \hat{\mathbf{Q}}_{1y}.$$

Its asymptotic variance can be deduced from Theorem 7.5.3, but it is simpler to apply the CLT directly to show that

$$\text{avar}(\tilde{\beta}_{1, \text{cls}}) = \mathbf{Q}_{11}^{-1} \boldsymbol{\Omega}_{11} \mathbf{Q}_{11}^{-1}. \quad (7.29)$$

The third estimator of β_1 is the efficient minimum distance estimator. Applying (7.25), it equals

$$\tilde{\beta}_{1, \text{md}} = \hat{\beta}_1 - \hat{\mathbf{V}}_{12} \hat{\mathbf{V}}_{22}^{-1} \hat{\beta}_2 \quad (7.30)$$

where we have partitioned

$$\hat{\mathbf{V}}_\beta = \begin{bmatrix} \hat{\mathbf{V}}_{11} & \hat{\mathbf{V}}_{12} \\ \hat{\mathbf{V}}_{21} & \hat{\mathbf{V}}_{22} \end{bmatrix}.$$

From Theorem 7.6.1 its asymptotic variance is

$$\text{avar}(\tilde{\beta}_{1, \text{md}}) = \mathbf{V}_{11} - \mathbf{V}_{12} \mathbf{V}_{22}^{-1} \mathbf{V}_{21}. \quad (7.31)$$

See Exercise 7.13 to verify equations (7.29), (7.30), and (7.31).

In general, the three estimators are different, and they have different asymptotic variances.

It is quite instructive to compare the asymptotic variances of the CLS and unconstrained least-squares estimators to assess whether or not the constrained estimator is necessarily more efficient than the unconstrained estimator.

First, consider the case of conditional homoskedasticity. In this case the two covariance matrices simplify to

$$\text{avar}(\hat{\beta}_1) = \sigma^2 \mathbf{Q}_{11 \cdot 2}^{-1}$$

and

$$\text{avar}(\tilde{\beta}_{1, \text{cls}}) = \sigma^2 \mathbf{Q}_{11}^{-1}.$$

If $\mathbf{Q}_{12} = 0$ (so \mathbf{x}_{1i} and \mathbf{x}_{2i} are orthogonal) then these two variance matrices equal and the two estimators have equal asymptotic efficiency. Otherwise, since $\mathbf{Q}_{12} \mathbf{Q}_{22}^{-1} \mathbf{Q}_{21} \geq 0$, then $\mathbf{Q}_{11} \geq \mathbf{Q}_{11} - \mathbf{Q}_{12} \mathbf{Q}_{22}^{-1} \mathbf{Q}_{21}$, and consequently

$$\mathbf{Q}_{11}^{-1} \sigma^2 \leq (\mathbf{Q}_{11} - \mathbf{Q}_{12} \mathbf{Q}_{22}^{-1} \mathbf{Q}_{21})^{-1} \sigma^2.$$

This means that under conditional homoskedasticity, $\tilde{\beta}_{1, \text{cls}}$ has a lower asymptotic variance matrix than $\hat{\beta}_1$. Therefore in this context, constrained least-squares is more efficient than unconstrained least-squares. This is consistent with our intuition that imposing a correct restriction (excluding an irrelevant regressor) improves estimation efficiency.

However, in the general case of conditional heteroskedasticity this ranking is not guaranteed. In fact what is really amazing is that the variance ranking can be reversed. The CLS estimator can have a larger asymptotic variance than the unconstrained least squares estimator.

To see this let's use the simple heteroskedastic example from Section 6.4. In that example, $Q_{11} = Q_{22} = 1$, $Q_{12} = \frac{1}{2}$, $\Omega_{11} = \Omega_{22} = 1$, and $\Omega_{12} = \frac{7}{8}$. We can calculate (see Exercise 7.14) that $\mathbf{Q}_{11 \cdot 2} = \frac{3}{4}$ and

$$\text{avar}(\hat{\beta}_1) = \frac{2}{3} \tag{7.32}$$

$$\text{avar}(\tilde{\beta}_{1, \text{cls}}) = 1 \tag{7.33}$$

$$\text{avar}(\tilde{\beta}_{1, \text{md}}) = \frac{5}{8}. \tag{7.34}$$

Thus the restricted least-squares estimator $\tilde{\beta}_{1, \text{cls}}$ has a larger variance than the unrestricted least-squares estimator $\hat{\beta}_1$! The minimum distance estimator has the smallest variance of the three, as expected.

What we have found is that when the estimation method is least-squares, deleting the irrelevant variable x_{2i} can actually increase estimation variance, or equivalently, adding an irrelevant variable can actually decrease the estimation variance.

To repeat this unexpected finding, we have shown in a very simple example that it is possible for least-squares applied to the short regression (7.14) to be less efficient for estimation of β_1 than least-squares applied to the long regression (7.13), even though the constraint $\beta_2 = 0$ is valid! This result is strongly counter-intuitive. It seems to contradict our initial motivation for pursuing constrained estimation – to improve estimation efficiency.

It turns out that a more refined answer is appropriate. Constrained estimation is desirable, but not constrained least-squares estimation. While least-squares is asymptotically efficient for estimation of the unconstrained projection model, it is not an efficient estimator of the constrained projection model.

7.8 Variance and Standard Error Estimation

The asymptotic covariance matrix (7.26) may be estimated by replacing \mathbf{V}_β with a consistent estimates such as $\widehat{\mathbf{V}}_\beta$. This variance estimator is

$$\widehat{\mathbf{V}}_\beta^* = \widehat{\mathbf{V}}_\beta - \widehat{\mathbf{V}}_\beta \mathbf{R} \left(\mathbf{R}' \widehat{\mathbf{V}}_\beta \mathbf{R} \right)^{-1} \mathbf{R}' \widehat{\mathbf{V}}_\beta. \quad (7.35)$$

We can calculate standard errors for any linear combination $\mathbf{h}'\tilde{\beta}$ so long as \mathbf{h} does not lie in the range space of \mathbf{R} . A standard error for $\mathbf{h}'\tilde{\beta}$ is

$$s(\mathbf{h}'\tilde{\beta}) = \left(n^{-1} \mathbf{h}' \widehat{\mathbf{V}}_\beta^* \mathbf{h} \right)^{1/2}. \quad (7.36)$$

7.9 Misspecification

What are the consequences for a constrained estimator $\tilde{\beta}$ if the constraint (7.1) is incorrect? To be specific, suppose that

$$\mathbf{R}'\beta = \mathbf{c}^*$$

where \mathbf{c}^* is not necessarily equal to \mathbf{c} .

This situation is a generalization of the analysis of “omitted variable bias” from Section 2.23, where we found that the short regression (e.g. (7.15)) is estimating a different projection coefficient than the long regression (e.g. (7.13)).

One mechanical answer is that we can use the formula (7.22) for the minimum distance estimator to find that

$$\tilde{\beta}_{\text{md}} \xrightarrow{p} \beta_{\text{md}}^* = \beta - \mathbf{W}^{-1} \mathbf{R} \left(\mathbf{R}' \mathbf{W}^{-1} \mathbf{R} \right)^{-1} (\mathbf{c}^* - \mathbf{c}). \quad (7.37)$$

The second term, $\mathbf{W}^{-1} \mathbf{R} \left(\mathbf{R}' \mathbf{W}^{-1} \mathbf{R} \right)^{-1} (\mathbf{c}^* - \mathbf{c})$, shows that imposing an incorrect constraint leads to inconsistency – an asymptotic bias. We can call the limiting value β_{md}^* the minimum-distance projection coefficient or the pseudo-true value implied by the restriction.

However, we can say more.

For example, we can describe some characteristics of the approximating projections. The CLS estimator projection coefficient has the representation

$$\beta_{\text{cls}}^* = \underset{\mathbf{R}'\beta=\mathbf{c}}{\operatorname{argmin}} \mathbb{E} (y_i - \mathbf{x}'_i \beta)^2,$$

the best linear predictor subject to the constraint (7.1). The minimum distance estimator converges to

$$\beta_{\text{md}}^* = \underset{\mathbf{R}'\beta=\mathbf{c}}{\operatorname{argmin}} (\beta - \beta_0)' \mathbf{W} (\beta - \beta_0)$$

where β_0 is the true coefficient. That is, β_{md}^* is the coefficient vector satisfying (7.1) closest to the true value in the weighted Euclidean norm. These calculations show that the constrained estimators are still reasonable in the sense that they produce good approximations to the true coefficient, conditional on being required to satisfy the constraint.

We can also show that $\tilde{\beta}_{\text{md}}$ has an asymptotic normal distribution. The trick is to define the pseudo-true value

$$\beta_n^* = \beta - \mathbf{W}_n^{-1} \mathbf{R} \left(\mathbf{R}' \mathbf{W}_n^{-1} \mathbf{R} \right)^{-1} (\mathbf{c}^* - \mathbf{c}). \quad (7.38)$$

(Note that (7.37) and (7.38) are different!) Then

$$\begin{aligned} \sqrt{n} \left(\tilde{\beta}_{\text{md}} - \beta_n^* \right) &= \sqrt{n} \left(\widehat{\beta} - \beta \right) - \mathbf{W}_n^{-1} \mathbf{R} \left(\mathbf{R}' \mathbf{W}_n^{-1} \mathbf{R} \right)^{-1} \sqrt{n} \left(\mathbf{R}' \widehat{\beta} - \mathbf{c}^* \right) \\ &= \left(\mathbf{I} - \mathbf{W}_n^{-1} \mathbf{R} \left(\mathbf{R}' \mathbf{W}_n^{-1} \mathbf{R} \right)^{-1} \mathbf{R}' \right) \sqrt{n} \left(\widehat{\beta} - \beta \right) \\ &\xrightarrow{d} \left(\mathbf{I} - \mathbf{W}^{-1} \mathbf{R} \left(\mathbf{R}' \mathbf{W}^{-1} \mathbf{R} \right)^{-1} \mathbf{R}' \right) \mathbf{N}(\mathbf{0}, \mathbf{V}_\beta) \\ &= \mathbf{N}(\mathbf{0}, \mathbf{V}_\beta(\mathbf{W})). \end{aligned} \quad (7.39)$$

In particular

$$\sqrt{n} \left(\tilde{\boldsymbol{\beta}}_{\text{emd}} - \boldsymbol{\beta}_n^* \right) \xrightarrow{d} \text{N}(\mathbf{0}, \mathbf{V}_{\boldsymbol{\beta}}^*).$$

This means that even when the constraint (7.1) is misspecified, the conventional covariance matrix estimator (7.35) and standard errors (7.36) are appropriate measures of the sampling variance, though the distributions are centered at the pseudo-true values (or projections) $\boldsymbol{\beta}_n^*$ rather than $\boldsymbol{\beta}$. The fact that the estimators are biased is an unavoidable consequence of misspecification.

An alternative approach to the asymptotic distribution theory under misspecification uses the concept of local alternatives. It is a technical device which might seem a bit artificial, but it is a powerful method to derive useful distributional approximations in a wide variety of contexts. The idea is to index the true coefficient $\boldsymbol{\beta}_n$ by n via the relationship

$$\mathbf{R}'\boldsymbol{\beta}_n = \mathbf{c} + \boldsymbol{\delta}n^{-1/2}. \quad (7.40)$$

Equation (7.40) specifies that $\boldsymbol{\beta}_n$ violates (7.1) and thus the constraint is misspecified. However, the constraint is “close” to correct, as the difference $\mathbf{R}'\boldsymbol{\beta}_n - \mathbf{c} = \boldsymbol{\delta}n^{-1/2}$ is “small” in the sense that it decreases with the sample size n . We call (7.40) **local misspecification**.

The asymptotic theory is then derived as $n \rightarrow \infty$ under the sequence of probability distributions with the coefficients $\boldsymbol{\beta}_n$. The way to think about this is that the true value of the parameter is $\boldsymbol{\beta}_n$, and it is “close” to satisfying (7.1). The reason why the deviation is proportional to $n^{-1/2}$ is because this is the only choice under which the localizing parameter $\boldsymbol{\delta}$ appears in the asymptotic distribution but does not dominate it. The best way to see this is to work through the asymptotic approximation.

Since $\boldsymbol{\beta}_n$ is the true coefficient value, then $y_i = \mathbf{x}_i'\boldsymbol{\beta}_n + e_i$ and we have the standard representation for the unconstrained estimator, namely

$$\begin{aligned} \sqrt{n} \left(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_n \right) &= \left(\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \right)^{-1} \left(\frac{1}{\sqrt{n}} \sum_{i=1}^n \mathbf{x}_i e_i \right) \\ &\xrightarrow{d} \text{N}(\mathbf{0}, \mathbf{V}_{\boldsymbol{\beta}}). \end{aligned} \quad (7.41)$$

There is no difference under fixed (classical) or local asymptotics, since the right-hand-side is independent of the coefficient $\boldsymbol{\beta}_n$.

A difference arises for the constrained estimator. Using (7.40), $\mathbf{c} = \mathbf{R}'\boldsymbol{\beta}_n - \boldsymbol{\delta}n^{-1/2}$, so

$$\mathbf{R}'\hat{\boldsymbol{\beta}} - \mathbf{c} = \mathbf{R}' \left(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_n \right) + \boldsymbol{\delta}n^{-1/2}$$

and

$$\begin{aligned} \tilde{\boldsymbol{\beta}}_{\text{md}} &= \hat{\boldsymbol{\beta}} - \mathbf{W}_n^{-1} \mathbf{R} \left(\mathbf{R}' \mathbf{W}_n^{-1} \mathbf{R} \right)^{-1} \left(\mathbf{R}' \hat{\boldsymbol{\beta}} - \mathbf{c} \right) \\ &= \hat{\boldsymbol{\beta}} - \mathbf{W}_n^{-1} \mathbf{R} \left(\mathbf{R}' \mathbf{W}_n^{-1} \mathbf{R} \right)^{-1} \mathbf{R}' \left(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_n \right) + \mathbf{W}_n^{-1} \mathbf{R} \left(\mathbf{R}' \mathbf{W}_n^{-1} \mathbf{R} \right)^{-1} \boldsymbol{\delta}n^{-1/2}. \end{aligned}$$

It follows that

$$\begin{aligned} \sqrt{n} \left(\tilde{\boldsymbol{\beta}}_{\text{md}} - \boldsymbol{\beta}_n \right) &= \left(\mathbf{I} - \mathbf{W}_n^{-1} \mathbf{R} \left(\mathbf{R}' \mathbf{W}_n^{-1} \mathbf{R} \right)^{-1} \mathbf{R}' \right) \sqrt{n} \left(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_n \right) \\ &\quad + \mathbf{W}_n^{-1} \mathbf{R} \left(\mathbf{R}' \mathbf{W}_n^{-1} \mathbf{R} \right)^{-1} \boldsymbol{\delta}. \end{aligned}$$

The first term is asymptotically normal (from 7.41). The second term converges in probability to a constant. This is because the $n^{-1/2}$ local scaling in (7.40) is exactly balanced by the \sqrt{n} scaling of the estimator. No alternative rate would have produced this result.

Consequently, we find that the asymptotic distribution equals

$$\begin{aligned} \sqrt{n} \left(\tilde{\boldsymbol{\beta}}_{\text{md}} - \boldsymbol{\beta}_n \right) &\xrightarrow{d} \text{N}(\mathbf{0}, \mathbf{V}_{\boldsymbol{\beta}}) + \mathbf{W}^{-1} \mathbf{R} \left(\mathbf{R}' \mathbf{W}^{-1} \mathbf{R} \right)^{-1} \boldsymbol{\delta} \\ &= \text{N}(\boldsymbol{\delta}^*, \mathbf{V}_{\boldsymbol{\beta}}(\mathbf{W})) \end{aligned} \quad (7.42)$$

where

$$\boldsymbol{\delta}^* = \mathbf{W}^{-1} \mathbf{R} (\mathbf{R}' \mathbf{W}^{-1} \mathbf{R})^{-1} \boldsymbol{\delta}.$$

The asymptotic distribution (7.42) is an approximation of the sampling distribution of the restricted estimator under misspecification. The distribution (7.42) contains an asymptotic bias component $\boldsymbol{\delta}^*$. The approximation is not fundamentally different from (7.39) – they both have the same asymptotic variances, and both reflect the bias due to misspecification. The difference is that (7.39) puts the bias on the left-side of the convergence arrow, while (7.42) has the bias on the right-side. There is no substantive difference between the two, but (7.42) is more convenient for some purposes, such as the analysis of the power of tests, as we will explore in the next chapter.

7.10 Nonlinear Constraints

In some cases it is desirable to impose nonlinear constraints on the parameter vector $\boldsymbol{\beta}$. They can be written as

$$\mathbf{r}(\boldsymbol{\beta}) = \mathbf{0} \tag{7.43}$$

where $\mathbf{r} : \mathbb{R}^k \rightarrow \mathbb{R}^q$. This includes the linear constraints (7.1) as a special case. An example of (7.43) which cannot be written as (7.1) is $\beta_1 \beta_2 = 1$, which is (7.43) with $\mathbf{r}(\boldsymbol{\beta}) = \beta_1 \beta_2 - 1$.

The constrained least-squares and minimum distance estimators of $\boldsymbol{\beta}$ subject to (7.43) solve the minimization problems

$$\tilde{\boldsymbol{\beta}}_{\text{cls}} = \underset{\mathbf{r}(\boldsymbol{\beta})=\mathbf{0}}{\operatorname{argmin}} SSE_n(\boldsymbol{\beta}) \tag{7.44}$$

$$\tilde{\boldsymbol{\beta}}_{\text{md}} = \underset{\mathbf{r}(\boldsymbol{\beta})=\mathbf{0}}{\operatorname{argmin}} J_n(\boldsymbol{\beta}) \tag{7.45}$$

where $SSE_n(\boldsymbol{\beta})$ and $J_n(\boldsymbol{\beta})$ are defined in (7.4) and (7.17), respectively. The solutions minimize the Lagrangians

$$\mathcal{L}(\boldsymbol{\beta}, \boldsymbol{\lambda}) = \frac{1}{2} SSE_n(\boldsymbol{\beta}) + \boldsymbol{\lambda}' \mathbf{r}(\boldsymbol{\beta}) \tag{7.46}$$

or

$$\mathcal{L}(\boldsymbol{\beta}, \boldsymbol{\lambda}) = \frac{1}{2} J_n(\boldsymbol{\beta}) + \boldsymbol{\lambda}' \mathbf{r}(\boldsymbol{\beta}) \tag{7.47}$$

over $(\boldsymbol{\beta}, \boldsymbol{\lambda})$.

Computationally, there is in general no explicit expression for the solutions so they must be found numerically. Algorithms to numerically solve (7.44) and (7.45) are known as **constrained optimization** methods, and are available in programming languages including Matlab, Gauss and R.

Assumption 7.10.1 $\mathbf{r}(\boldsymbol{\beta}) = \mathbf{0}$ with $\operatorname{rank}(\mathbf{R}) = q$, where $\mathbf{R} = \frac{\partial}{\partial \boldsymbol{\beta}} \mathbf{r}(\boldsymbol{\beta})'$.

The asymptotic distribution is a simple generalization of the case of a linear constraint, but the proof is more delicate.

Theorem 7.10.1 Under Assumptions 6.1.2, 7.10.1, and 7.5.2, for $\tilde{\boldsymbol{\beta}} = \tilde{\boldsymbol{\beta}}_{\text{md}}$ and $\tilde{\boldsymbol{\beta}} = \tilde{\boldsymbol{\beta}}_{\text{cls}}$ defined in (7.44) and (7.45),

$$\sqrt{n} (\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}) \xrightarrow{d} \mathbf{N}(\mathbf{0}, \mathbf{V}_{\boldsymbol{\beta}}(\mathbf{W}))$$

as $n \rightarrow \infty$, where $\mathbf{V}_{\boldsymbol{\beta}}(\mathbf{W})$ is defined in (7.24). For $\tilde{\boldsymbol{\beta}}_{\text{cls}}$, $\mathbf{W} = \mathbf{Q}_{xx}$ and $\mathbf{V}_{\boldsymbol{\beta}}(\mathbf{W}) = \mathbf{V}_{\text{cls}}$ as defined in Theorem 7.5.3. $\mathbf{V}_{\boldsymbol{\beta}}(\mathbf{W})$ is minimized with $\mathbf{W} = \mathbf{V}_{\boldsymbol{\beta}}^{-1}$, in which case the asymptotic variance is

$$\mathbf{V}_{\boldsymbol{\beta}}^* = \mathbf{V}_{\boldsymbol{\beta}} - \mathbf{V}_{\boldsymbol{\beta}} \mathbf{R} (\mathbf{R}' \mathbf{V}_{\boldsymbol{\beta}} \mathbf{R})^{-1} \mathbf{R}' \mathbf{V}_{\boldsymbol{\beta}}.$$

The asymptotic variance matrix for the efficient minimum distance estimator can be estimated by

$$\widehat{\mathbf{V}}_{\boldsymbol{\beta}}^* = \widehat{\mathbf{V}}_{\boldsymbol{\beta}} - \widehat{\mathbf{V}}_{\boldsymbol{\beta}} \widehat{\mathbf{R}} (\widehat{\mathbf{R}}' \widehat{\mathbf{V}}_{\boldsymbol{\beta}} \widehat{\mathbf{R}})^{-1} \widehat{\mathbf{R}}' \widehat{\mathbf{V}}_{\boldsymbol{\beta}}$$

where

$$\widehat{\mathbf{R}} = \frac{\partial}{\partial \boldsymbol{\beta}} \mathbf{r}(\tilde{\boldsymbol{\beta}}_{\text{md}})'. \quad (7.48)$$

Standard errors for the elements of $\tilde{\boldsymbol{\beta}}_{\text{md}}$ are the square roots of the diagonal elements of $\widehat{\mathbf{V}}_{\boldsymbol{\beta}}^* = n^{-1} \widehat{\mathbf{V}}_{\boldsymbol{\beta}}^*$.

7.11 Inequality Restrictions

Inequality constraints on the parameter vector $\boldsymbol{\beta}$ take the form

$$\mathbf{r}(\boldsymbol{\beta}) \geq \mathbf{0} \quad (7.49)$$

for some function $\mathbf{r} : \mathbb{R}^k \rightarrow \mathbb{R}^q$. The most common example is a non-negative constraint

$$\beta_1 \geq 0.$$

The constrained least-squares and minimum distance estimators can be written as

$$\tilde{\boldsymbol{\beta}}_{\text{cls}} = \underset{\mathbf{r}(\boldsymbol{\beta}) \geq \mathbf{0}}{\operatorname{argmin}} SSE_n(\boldsymbol{\beta}) \quad (7.50)$$

and

$$\tilde{\boldsymbol{\beta}}_{\text{md}} = \underset{\mathbf{r}(\boldsymbol{\beta}) \geq \mathbf{0}}{\operatorname{argmin}} J_n(\boldsymbol{\beta}). \quad (7.51)$$

Except in special cases the constrained estimators do not have simple algebraic solutions. An important exception is when there is a single non-negativity constraint, e.g. $\beta_1 \geq 0$ with $q = 1$. In this case the constrained estimator can be found by two-step approach. First compute the unconstrained estimator $\widehat{\boldsymbol{\beta}}$. If $\widehat{\beta}_1 \geq 0$ then $\tilde{\boldsymbol{\beta}} = \widehat{\boldsymbol{\beta}}$. Second, if $\widehat{\beta}_1 < 0$ then impose $\beta_1 = 0$ (eliminate the regressor X_1) and re-estimate. This yields the constrained least-squares estimator. While this method works when there is a single non-negativity constraint, it does not immediately generalize to other contexts.

The computational problems (7.50) and (7.51) are examples of **quadratic programming** problems. Quick and easy computer algorithms are available in programming languages including Matlab, Gauss and R.

Inference on inequality-constrained estimators is unfortunately quite challenging. The conventional asymptotic theory gives rise to the following dichotomy. If the true parameter satisfies the strict inequality $\mathbf{r}(\boldsymbol{\beta}) > \mathbf{0}$, then asymptotically the estimator is not subject to the constraint and the inequality-constrained estimator has an asymptotic distribution equal to the unconstrained case. However if the true parameter is on the boundary, e.g. $\mathbf{r}(\boldsymbol{\beta}) = \mathbf{0}$, then the estimator has a truncated structure. This is easiest to see in the one-dimensional case. If we have an estimator $\hat{\beta}$ which satisfies $\sqrt{n}(\hat{\beta} - \beta) \xrightarrow{d} Z = N(0, V_\beta)$ and $\beta = 0$, then the constrained estimator $\tilde{\beta} = \max[\hat{\beta}, 0]$ will have the asymptotic distribution $\sqrt{n}\tilde{\beta} \xrightarrow{d} \max[Z, 0]$, a “half-normal” distribution.

7.12 Constrained MLE

Recall that the log-likelihood function (3.44) for the normal regression model is

$$\log L(\boldsymbol{\beta}, \sigma^2) = -\frac{n}{2} \log(2\pi\sigma^2) - \frac{1}{2\sigma^2} SSE_n(\boldsymbol{\beta}).$$

The constrained maximum likelihood estimator (CMLE) $(\hat{\boldsymbol{\beta}}_{\text{cmle}}, \hat{\sigma}_{\text{cmle}}^2)$ maximizes $\log L(\boldsymbol{\beta}, \sigma^2)$ subject to the constraint (7.43). Since $\log L(\boldsymbol{\beta}, \sigma^2)$ is a function of $\boldsymbol{\beta}$ only through the sum of squared errors $SSE_n(\boldsymbol{\beta})$, maximizing the likelihood is identical to minimizing $SSE_n(\boldsymbol{\beta})$. Hence $\hat{\boldsymbol{\beta}}_{\text{cmle}} = \hat{\boldsymbol{\beta}}_{\text{cls}}$ and $\hat{\sigma}_{\text{cmle}}^2 = \hat{\sigma}_{\text{cls}}^2$.

7.13 Technical Proofs*

Proof of Theorem 7.6.1, Equation (7.28). Let \mathbf{R}_\perp be a full rank $k \times (k - q)$ matrix satisfying $\mathbf{R}'_\perp \mathbf{V}_\beta \mathbf{R} = \mathbf{0}$ and then set $\mathbf{C} = [\mathbf{R}, \mathbf{R}_\perp]$ which is full rank and invertible. Then we can calculate that

$$\begin{aligned} \mathbf{C}' \mathbf{V}_\beta^* \mathbf{C} &= \begin{bmatrix} \mathbf{R}' \mathbf{V}_\beta^* \mathbf{R} & \mathbf{R}' \mathbf{V}_\beta^* \mathbf{R}_\perp \\ \mathbf{R}'_\perp \mathbf{V}_\beta^* \mathbf{R} & \mathbf{R}'_\perp \mathbf{V}_\beta^* \mathbf{R}_\perp \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{R}'_\perp \mathbf{V}_\beta \mathbf{R}_\perp \end{bmatrix} \end{aligned}$$

and

$$\begin{aligned} &\mathbf{C}' \mathbf{V}_\beta(\mathbf{W}) \mathbf{C} \\ &= \begin{bmatrix} \mathbf{R}' \mathbf{V}_\beta^*(\mathbf{W}) \mathbf{R} & \mathbf{R}' \mathbf{V}_\beta^*(\mathbf{W}) \mathbf{R}_\perp \\ \mathbf{R}'_\perp \mathbf{V}_\beta^*(\mathbf{W}) \mathbf{R} & \mathbf{R}'_\perp \mathbf{V}_\beta^*(\mathbf{W}) \mathbf{R}_\perp \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{R}'_\perp \mathbf{V}_\beta \mathbf{R}_\perp + \mathbf{R}'_\perp \mathbf{W} \mathbf{R} (\mathbf{R}' \mathbf{W} \mathbf{R})^{-1} \mathbf{R}' \mathbf{V}_\beta \mathbf{R} (\mathbf{R}' \mathbf{W} \mathbf{R})^{-1} \mathbf{R}' \mathbf{W} \mathbf{R}_\perp \end{bmatrix}. \end{aligned}$$

Thus

$$\begin{aligned} &\mathbf{C}' (\mathbf{V}_\beta(\mathbf{W}) - \mathbf{V}_\beta^*) \mathbf{C} \\ &= \mathbf{C}' \mathbf{V}_\beta(\mathbf{W}) \mathbf{C} - \mathbf{C}' \mathbf{V}_\beta^* \mathbf{C} \\ &= \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{R}'_\perp \mathbf{W} \mathbf{R} (\mathbf{R}' \mathbf{W} \mathbf{R})^{-1} \mathbf{R}' \mathbf{V}_\beta \mathbf{R} (\mathbf{R}' \mathbf{W} \mathbf{R})^{-1} \mathbf{R}' \mathbf{W} \mathbf{R}_\perp \end{bmatrix} \\ &\geq 0 \end{aligned}$$

Since \mathbf{C} is invertible it follows that $\mathbf{V}_\beta(\mathbf{W}) - \mathbf{V}_\beta^* \geq 0$ which is (7.28). \blacksquare

Proof of Theorem 7.10.1. We show the result for the minimum distance estimator $\tilde{\boldsymbol{\beta}} = \tilde{\boldsymbol{\beta}}_{\text{md}}$, as the proof for the constrained least-squares estimator is similar. For simplicity we assume that the

constrained estimator is consistent $\tilde{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\beta}$. This can be shown with more effort, but requires a deeper treatment than appropriate for this textbook.

For each element $r_j(\boldsymbol{\beta})$ of the q -vector $\mathbf{r}(\boldsymbol{\beta})$, by the mean value theorem there exists a $\boldsymbol{\beta}_j^*$ on the line segment joining $\tilde{\boldsymbol{\beta}}$ and $\boldsymbol{\beta}$ such that

$$r_j(\tilde{\boldsymbol{\beta}}) = r_j(\boldsymbol{\beta}) + \frac{\partial}{\partial \boldsymbol{\beta}} r_j(\boldsymbol{\beta}_j^*)' (\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}). \quad (7.52)$$

Let \mathbf{R}_n^* be the $k \times q$ matrix

$$\mathbf{R}_n^* = \begin{bmatrix} \frac{\partial}{\partial \boldsymbol{\beta}} r_1(\boldsymbol{\beta}_1^*) & \frac{\partial}{\partial \boldsymbol{\beta}} r_2(\boldsymbol{\beta}_2^*) & \cdots & \frac{\partial}{\partial \boldsymbol{\beta}} r_q(\boldsymbol{\beta}_q^*) \end{bmatrix}.$$

Since $\tilde{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\beta}$ it follows that $\boldsymbol{\beta}_j^* \xrightarrow{p} \boldsymbol{\beta}$, and by the CMT, $\mathbf{R}_n^* \xrightarrow{p} \mathbf{R}$. Stacking the (7.52), we obtain

$$\mathbf{r}(\tilde{\boldsymbol{\beta}}) = \mathbf{r}(\boldsymbol{\beta}) + \mathbf{R}_n^{*'} (\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}).$$

Since $\mathbf{r}(\tilde{\boldsymbol{\beta}}) = \mathbf{0}$ by construction and $\mathbf{r}(\boldsymbol{\beta}) = \mathbf{0}$ by Assumption 7.5.1, this implies

$$\mathbf{0} = \mathbf{R}_n^{*'} (\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}). \quad (7.53)$$

The first-order condition for (7.47) is

$$\mathbf{W}_n (\hat{\boldsymbol{\beta}} - \tilde{\boldsymbol{\beta}}) = \hat{\mathbf{R}} \tilde{\boldsymbol{\lambda}}.$$

where $\hat{\mathbf{R}}$ is defined in (7.48).

Premultiplying by $\mathbf{R}_n^{*'} \mathbf{W}_n^{-1}$, inverting, and using (7.53), we find

$$\tilde{\boldsymbol{\lambda}} = \left(\mathbf{R}_n^{*'} \mathbf{W}_n^{-1} \hat{\mathbf{R}} \right)^{-1} \mathbf{R}_n^{*'} (\hat{\boldsymbol{\beta}} - \tilde{\boldsymbol{\beta}}) = \left(\mathbf{R}_n^{*'} \mathbf{W}_n^{-1} \hat{\mathbf{R}} \right)^{-1} \mathbf{R}_n^{*'} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}).$$

Thus

$$\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta} = \left(\mathbf{I} - \mathbf{W}_n^{-1} \hat{\mathbf{R}} \left(\mathbf{R}_n^{*'} \mathbf{W}_n^{-1} \hat{\mathbf{R}} \right)^{-1} \mathbf{R}_n^{*'} \right) (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}). \quad (7.54)$$

From Theorem 6.3.2 and Theorem 6.7.1 we find

$$\begin{aligned} \sqrt{n} (\tilde{\boldsymbol{\beta}} - \boldsymbol{\beta}) &= \left(\mathbf{I} - \mathbf{W}_n^{-1} \hat{\mathbf{R}} \left(\mathbf{R}_n^{*'} \mathbf{W}_n^{-1} \hat{\mathbf{R}} \right)^{-1} \mathbf{R}_n^{*'} \right) \sqrt{n} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \\ &\xrightarrow{d} \left(\mathbf{I} - \mathbf{W}^{-1} \mathbf{R} (\mathbf{R}' \mathbf{W}^{-1} \mathbf{R})^{-1} \mathbf{R}' \right) \mathbf{N}(\mathbf{0}, \mathbf{V}_\beta) \\ &= \mathbf{N}(\mathbf{0}, \mathbf{V}_\beta(\mathbf{W})). \end{aligned}$$

■

Exercises

Exercise 7.1 In the model $\mathbf{y} = \mathbf{X}_1\boldsymbol{\beta}_1 + \mathbf{X}_2\boldsymbol{\beta}_2 + \mathbf{e}$, show directly from definition (7.3) that the CLS estimate of $\boldsymbol{\beta} = (\boldsymbol{\beta}_1, \boldsymbol{\beta}_2)$ subject to the constraint that $\boldsymbol{\beta}_2 = \mathbf{0}$ is the OLS regression of \mathbf{y} on \mathbf{X}_1 .

Exercise 7.2 In the model $\mathbf{y} = \mathbf{X}_1\boldsymbol{\beta}_1 + \mathbf{X}_2\boldsymbol{\beta}_2 + \mathbf{e}$, show directly from definition (7.3) that the CLS estimate of $\boldsymbol{\beta} = (\boldsymbol{\beta}_1, \boldsymbol{\beta}_2)$, subject to the constraint that $\boldsymbol{\beta}_1 = \mathbf{c}$ (where \mathbf{c} is some given vector) is the OLS regression of $\mathbf{y} - \mathbf{X}_1\mathbf{c}$ on \mathbf{X}_2 .

Exercise 7.3 In the model $\mathbf{y} = \mathbf{X}_1\boldsymbol{\beta}_1 + \mathbf{X}_2\boldsymbol{\beta}_2 + \mathbf{e}$, with \mathbf{X}_1 and \mathbf{X}_2 each $n \times k$, find the CLS estimate of $\boldsymbol{\beta} = (\boldsymbol{\beta}_1, \boldsymbol{\beta}_2)$, subject to the constraint that $\boldsymbol{\beta}_1 = -\boldsymbol{\beta}_2$.

Exercise 7.4 Verify that for $\tilde{\boldsymbol{\beta}}_{\text{cls}}$ defined in (7.9) that $\mathbf{R}'\tilde{\boldsymbol{\beta}}_{\text{cls}} = \mathbf{c}$.

Exercise 7.5 Let $\tilde{\mathbf{e}}$ be the vector of constrained least-squares residuals (7.11). Show that under (7.1),

(a) $\mathbf{R}'\hat{\boldsymbol{\beta}} - \mathbf{c} = \mathbf{R}'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{e}$

(b) $\tilde{\boldsymbol{\beta}}_{\text{cls}} - \boldsymbol{\beta} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{e} - (\mathbf{X}'\mathbf{X})^{-1}\mathbf{R}\left(\mathbf{R}'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{R}\right)^{-1}\mathbf{R}'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{e}$

(c) $\tilde{\mathbf{e}} = (\mathbf{I} - \mathbf{P} + \mathbf{A})\mathbf{e}$ for $\mathbf{P} = \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$ and some matrix \mathbf{A} (find this matrix \mathbf{A}).

(d) Show that \mathbf{A} is symmetric and idempotent, $\text{tr } \mathbf{A} = q$, and $\mathbf{P}\mathbf{A} = \mathbf{A}$.

Exercise 7.6 Show (7.12), that is, $\mathbb{E}(s_{\text{cls}}^2 | \mathbf{X}) = \sigma^2$, under the assumptions of the homoskedastic regression model and (7.1).

Hint: Use the results of Exercise 7.5

Exercise 7.7 Verify (7.21) and (7.22), and that the minimum distance estimator $\tilde{\boldsymbol{\beta}}_{\text{md}}$ with $\mathbf{W}_n = \hat{\mathbf{Q}}_{xx}$ equals the CLS estimator.

Exercise 7.8 Prove Theorem 7.5.1.

Exercise 7.9 Prove Theorem 7.5.2.

Exercise 7.10 Prove Theorem 7.5.3. (Hint: Use that CLS is a special case of Theorem 7.5.2.)

Exercise 7.11 Verify that (7.26) is $\mathbf{V}_{\boldsymbol{\beta}}(\mathbf{W})$ with $\mathbf{W} = \mathbf{V}_{\boldsymbol{\beta}}^{-1}$.

Exercise 7.12 Prove (7.27). Hint: Use (7.26).

Exercise 7.13 Verify (7.29), (7.30) and (7.31)

Exercise 7.14 Verify (7.32), (7.33), and (7.34).

Exercise 7.15 As in Exercise 6.11 and 3.21, use the CPS dataset and the subsample of white male Hispanics.

- (a) Estimate the regression

$$\begin{aligned}\log(\widehat{Wage}) = & \beta_1 \textit{education} + \beta_2 \textit{experience} + \beta_3 \textit{experience}^2/100 + \beta_4 \textit{Married}_1 \\ & + \beta_5 \textit{Married}_2 + \beta_6 \textit{Married}_3 + \beta_7 \textit{Widowed} + \beta_8 \textit{Divorced} + \beta_9 \textit{Separated} + \beta_{10}\end{aligned}$$

where $\textit{Married}_1$, $\textit{Married}_2$, and $\textit{Married}_3$ are the first three marital status codes as listed in Section 3.19.

- (b) Estimate the equation using constrained least-squares, imposing the constraints $\beta_4 = \beta_7$ and $\beta_8 = \beta_9$, and report the estimates and standard errors
- (c) Estimate the equation using efficient minimum distance, imposing the same constraints, and report the estimates and standard errors
- (d) Under what constraint on the coefficients is the wage equation non-decreasing in experience for experience up to 50?
- (e) Estimate the equation imposing $\beta_4 = \beta_7$, $\beta_8 = \beta_9$, and the inequality from part (d).

Chapter 8

Hypothesis Testing

8.1 Hypotheses

In Chapter 7 we discussed estimation subject to restrictions, including linear restrictions (7.1), nonlinear restrictions (7.43), and inequality restrictions (7.49). In this chapter we discuss **tests** of such restrictions.

Hypothesis tests attempt to assess whether there is evidence to contradict a proposed parametric restriction. Let

$$\boldsymbol{\theta} = \boldsymbol{r}(\boldsymbol{\beta})$$

be a $q \times 1$ parameter of interest where $\boldsymbol{r} : \mathbb{R}^k \rightarrow \Theta \subset \mathbb{R}^q$ is some transformation. For example, $\boldsymbol{\theta}$ may be a single coefficient, e.g. $\boldsymbol{\theta} = \beta_j$, the difference between two coefficients, e.g. $\boldsymbol{\theta} = \beta_j - \beta_\ell$, or the ratio of two coefficients, e.g. $\boldsymbol{\theta} = \beta_j / \beta_\ell$.

A point hypothesis concerning $\boldsymbol{\theta}$ is a proposed restriction such as

$$\boldsymbol{\theta} = \boldsymbol{\theta}_0 \tag{8.1}$$

where $\boldsymbol{\theta}_0$ is a hypothesized value.

More generally, letting $\boldsymbol{\beta} \in \boldsymbol{B} \subset \mathbb{R}^k$ be the parameter space, a hypothesis is a restriction $\boldsymbol{\beta} \in \boldsymbol{B}_0$ where \boldsymbol{B}_0 is a proper subset of \boldsymbol{B} . This specializes to (8.1) by setting $\boldsymbol{B}_0 = \{\boldsymbol{\beta} \in \boldsymbol{B} : \boldsymbol{r}(\boldsymbol{\beta}) = \boldsymbol{\theta}_0\}$.

In this chapter we will focus primarily on point hypotheses of the form (8.1) as they are the most common and relatively simple to handle.

The hypothesis to be tested is called the null hypothesis.

Definition 1 The *null hypothesis*, written \mathbb{H}_0 , is the restriction $\boldsymbol{\theta} = \boldsymbol{\theta}_0$ or $\boldsymbol{\beta} \in \boldsymbol{B}_0$.

We often write the null hypothesis as $\mathbb{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$ or $\mathbb{H}_0 : \boldsymbol{r}(\boldsymbol{\beta}) = \boldsymbol{\theta}_0$.

The complement of the null hypothesis (the collection of parameter values which do not satisfy the null hypothesis) is called the alternative hypothesis.

Definition 2 The *alternative hypothesis*, written \mathbb{H}_1 , is the set $\{\boldsymbol{\theta} \in \Theta : \boldsymbol{\theta} \neq \boldsymbol{\theta}_0\}$ or $\{\boldsymbol{\beta} \in \boldsymbol{B} : \boldsymbol{\beta} \notin \boldsymbol{B}_0\}$.

We often write the alternative hypothesis as $\mathbb{H}_1 : \boldsymbol{\theta} \neq \boldsymbol{\theta}_0$ or $\mathbb{H}_0 : \boldsymbol{r}(\boldsymbol{\beta}) \neq \boldsymbol{\theta}_0$. For simplicity, we often refer to the hypotheses as “the null” and “the alternative”.

In hypothesis testing, we assume that there is a true (but unknown) value of $\boldsymbol{\theta}$ and this value either satisfies \mathbb{H}_0 or does not satisfy \mathbb{H}_0 . The goal of hypothesis testing is to assess whether or not \mathbb{H}_0 is true, by asking if \mathbb{H}_0 is consistent with the observed data.

To be specific, take our example of wage determination and consider the question: Does union membership affect wages? We can turn this into a hypothesis test by specifying the null as the restriction that a coefficient on union membership is zero in a wage regression. Consider, for example, the estimates reported in Table 4.1. The coefficient for “Male Union Member” is 0.095 (a wage premium of 9.5%) and the coefficient for “Female Union Member” is 0.022 (a wage premium of 2.2%). These are estimates, not the true values. The question is: Are the estimates close enough to zero that they are compatible with the hypothesis that the true values of the coefficients are zero? Or are the estimates sufficiently different from zero that this hypothesis must be false?

8.2 Acceptance and Rejection

A hypothesis test either accepts the null hypothesis, or rejects the null hypothesis in favor of the alternative hypothesis. We can describe these two decisions as “Accept \mathbb{H}_0 ” and “Reject \mathbb{H}_0 ”. In the example given in the previous section, the decision would be either to accept the hypothesis that union membership does not affect wages, or to reject the hypothesis in favor of the alternative that union membership does affect wages.

The decision is based on the data, and so is a mapping from the sample space to the decision set. This splits the sample space into two regions S_0 and S_1 such that if the observed sample falls into S_0 we accept \mathbb{H}_0 , while if the sample falls into S_1 we reject \mathbb{H}_0 . The set S_0 can be called the acceptance region and the set S_1 the rejection or critical region.

It is convenient to express this mapping as a real-valued function called a **test statistic**

$$T_n = T_n((y_1, \boldsymbol{x}_1), \dots, (y_n, \boldsymbol{x}_n))$$

relative to a **critical value** c . The hypothesis test then consists of the decision rule

1. Accept \mathbb{H}_0 if $T_n \leq c$,
2. Reject \mathbb{H}_0 if $T_n > c$.

A test statistic T_n should be designed so that small values are likely when \mathbb{H}_0 is true and large values are likely when \mathbb{H}_1 is true. There is a well developed statistical theory concerning the design of optimal tests. We will not review that theory here, but instead refer the reader to Lehmann and Romano (2005). In this chapter we will summarize the main approaches to the design of test statistics.

The most commonly used test statistic is the absolute value of the t-statistic

$$t_n = |t_n(\theta_0)| \tag{8.2}$$

where

$$t_n(\theta) = \frac{\hat{\theta} - \theta}{s(\hat{\theta})} \tag{8.3}$$

is the t-statistic from (6.44), where $\hat{\theta}$ is a point estimate and $s(\hat{\theta})$ its standard error. t_n is an appropriate statistic when testing hypotheses on individual coefficients or real-valued parameters $\theta = h(\boldsymbol{\beta})$ and θ_0 is the hypothesized value. Quite typically, $\theta_0 = 0$, as interest focuses on whether or not a coefficient equals zero, but this is not the only possibility. For example, interest may focus on whether an elasticity θ equals 1, in which case we may wish to test $\mathbb{H}_0 : \theta = 1$.

8.3 Type I Error

A false rejection of the null hypothesis \mathbb{H}_0 (rejecting \mathbb{H}_0 when \mathbb{H}_0 is true) is called a Type I error. The probability of a Type I error is

$$\Pr(\text{Reject } \mathbb{H}_0 \mid \mathbb{H}_0 \text{ true}) = \Pr(T_n > c \mid \mathbb{H}_0 \text{ true}). \quad (8.4)$$

The finite sample **size** of the test is defined as the supremum of (8.4) across all data distributions which satisfy \mathbb{H}_0 . A primary goal of test construction is to limit the incidence of Type I error by bounding the size of the test.

For the reasons discussed in Chapter 6, in typical econometric models the exact sampling distributions of estimators and test statistics are unknown and hence we cannot explicitly calculate (8.4). Instead, we typically rely on asymptotic approximations. We start by assuming that the test statistic has an asymptotic distribution under \mathbb{H}_0 . That is, when \mathbb{H}_0 is true

$$T_n \xrightarrow{d} T \quad (8.5)$$

as $n \rightarrow \infty$ for some continuously-distributed random variable T . This is not a substantive restriction, as most conventional econometric tests satisfy (8.5). Let $G(u) = \Pr(T \leq u)$ denote the distribution of T . We call T (or G) the **asymptotic null distribution**.

It is generally desirable to design test statistics T_n whose asymptotic null distribution G is known and does not depend on unknown parameters. In this case we say that the statistic T_n is **asymptotically pivotal**.

For example, if $T_n = t_n$, the absolute t-statistic from (8.2), then we know from Theorem 6.12.1 that if $\theta = \theta_0$ (that is, the null hypothesis holds), then $t_n \xrightarrow{d} |Z|$ as $n \rightarrow \infty$ where $Z \sim N(0, 1)$. This means that $G(u) = \Pr(T \leq |Z|) = \bar{\Phi}(u)$, the symmetrized normal distribution function defined in (6.45). This distribution does not depend on unknowns and is pivotal.

We define the **asymptotic size** of the test as the asymptotic probability of a Type I error:

$$\begin{aligned} \lim_{n \rightarrow \infty} \Pr(T_n > c \mid \mathbb{H}_0 \text{ true}) &= \Pr(T > c) \\ &= 1 - G(c). \end{aligned}$$

We see that the asymptotic size of the test is a simple function of the asymptotic null distribution G and the critical value c . For example, the asymptotic size of a test based on the absolute t-statistic with critical value c is $1 - \bar{\Phi}(c)$.

In the dominant approach to hypothesis testing, the researcher pre-selects a **significance level** $\alpha \in (0, 1)$ and then selects c so that the (asymptotic) size is no larger than α . When the asymptotic null distribution G is pivotal, we can accomplish this by setting c equal to the $(1 - \alpha)^{th}$ quantile of the distribution G . (If the distribution G is not pivotal, more complicated methods must be used, pointing out the great convenience of using asymptotically pivotal test statistics.) We call c the **asymptotic critical value** because it has been selected from the asymptotic null distribution. For example, since $\bar{\Phi}(1.96) = 0.95$, it follows that the 5% asymptotic critical value for the absolute t-statistic is $c = 1.96$.

8.4 t tests

As we mentioned earlier, the most common test of the one-dimensional hypothesis

$$\mathbb{H}_0 : \theta = \theta_0 \quad (8.6)$$

against the alternative

$$\mathbb{H}_1 : \theta \neq \theta_0 \quad (8.7)$$

is the absolute value of the t-statistic (8.3). We now formally state its asymptotic null distribution, which is a simple application of Theorem 6.12.1.

Theorem 8.4.1 Under Assumptions 6.1.2, 6.10.1, and $\mathbb{H}_0 : \theta = \theta_0$,

$$t_n(\theta_0) \xrightarrow{d} Z.$$

For c satisfying $\alpha = 2(1 - \Phi(c))$,

$$\Pr(|t_n(\theta_0)| > c \mid \mathbb{H}_0) \longrightarrow \alpha$$

so the test “Reject \mathbb{H}_0 if $|t_n(\theta_0)| > c$ ” has asymptotic size α .

The theorem shows that asymptotic critical values can be taken from the normal distribution table.

The alternative hypothesis (8.7) is sometimes called a “two-sided” alternative. Sometimes we are interested in testing for one-sided alternatives such as

$$\mathbb{H}_1 : \theta > \theta_0 \tag{8.8}$$

or

$$\mathbb{H}_1 : \theta < \theta_0. \tag{8.9}$$

Tests of (8.6) against (8.8) or (8.9) are based on the signed t-statistic $t_n = t_n(\theta_0)$. The hypothesis (8.6) is rejected in favor of (8.8) if $t_n > c$ where c satisfies $\alpha = 1 - \Phi(c)$. Negative values of t_n are not taken as evidence against \mathbb{H}_0 , as point estimates $\hat{\theta}$ less than θ_0 do not point to (8.8). Since the critical values are taken from the single tail of the normal distribution, they are smaller than for two-sided tests. Specifically, the asymptotic 5% critical value is $\alpha = 1.645$. Thus, we reject (8.6) in favor of (8.8) if $t_n > 1.645$.

Conversely, tests of (8.6) against (8.9) reject \mathbb{H}_0 for negative t-statistics, e.g. if $t_n \leq -c$. For this alternative large positive values of t_n are not evidence against \mathbb{H}_0 . An asymptotic 5% test rejects if $t_n < -1.645$.

There seems to be an ambiguity. Should we use the two-sided critical value 1.96 or the one-sided critical value 1.645? The answer is that we should use one-sided tests and critical values only when the parameter space is known to satisfy a one-sided restriction such as $\theta \geq \theta_0$. This is when the test of (8.6) against (8.8) makes sense. If the restriction $\theta \geq \theta_0$ is not known *a priori*, then imposing this restriction to test (8.6) against (8.8) does not make sense. Since linear regression coefficients typically do not have *a priori* sign restrictions, we conclude that two-sided tests are generally appropriate.

8.5 Type II Error and Power

A false acceptance of the null hypothesis \mathbb{H}_0 (accepting \mathbb{H}_0 when \mathbb{H}_1 is true) is called a **Type II error**. The rejection probability under the alternative hypothesis is called the **power** of the test, and equals 1 minus the probability of a Type II error:

$$\pi_n(\boldsymbol{\theta}) = \Pr(\text{Reject } \mathbb{H}_0 \mid \mathbb{H}_1 \text{ true}) = \Pr(T_n > c \mid \mathbb{H}_1 \text{ true}).$$

We call $\pi_n(\boldsymbol{\theta})$ the **power function** and is written as a function of $\boldsymbol{\theta}$ to indicate its dependence on the true value of the parameter $\boldsymbol{\theta}$.

In the dominant approach to hypothesis testing, the goal of test construction is to have high power, subject to the constraint that the size of the test is lower than the pre-specified significance level. Generally, the power of a test depends on the true value of the parameter θ , and for a well behaved test the power is increasing both as θ moves away from the null hypothesis θ_0 and as the sample size n increases.

Given the two possible states of the world (\mathbb{H}_0 or \mathbb{H}_1) and the two possible decisions (Accept \mathbb{H}_0 or Reject \mathbb{H}_0), there are four possible pairings of states and decisions as is depicted in the following chart.

Hypothesis Testing Decisions

	Accept \mathbb{H}_0	Reject \mathbb{H}_0
\mathbb{H}_0 true	Correct Decision	Type I Error
\mathbb{H}_1 true	Type II Error	Correct Decision

Given a test statistic T_n , increasing the critical value c increases the acceptance region S_0 while decreasing the rejection region S_1 . This decreases the likelihood of a Type I error (decreases the size) but increases the likelihood of a Type II error (decreases the power). Thus the choice of c involves a trade-off between size and the power. This is why the significance level α of the test cannot be set arbitrarily small. (Otherwise the test will not have meaningful power.)

It is important to consider the power of a test when interpreting hypothesis tests, as an overly narrow focus on size can lead to poor decisions. For example, it is trivial to design a test which has perfect size yet has trivial power. Specifically, for any hypothesis we can use the following test: Generate a random variable $U \sim U[0, 1]$ and reject \mathbb{H}_0 if $U < \alpha$. This test has exact size of α . Yet the test also has power precisely equal to α . When the power of a test equals the size, we say that the test has **trivial power**. Nothing is learned from such a test.

8.6 Statistical Significance

Testing requires a pre-selected choice of significance level α , yet there is no objective scientific basis for choice of α . Never-the-less, the common practice is to set $\alpha = 0.05$ (5%). Alternative values are $\alpha = 0.10$ (10%) and $\alpha = 0.01$ (1%). These choices are somewhat the by-product of traditional tables of critical values and statistical software.

The informal reasoning behind the choice of a 5% critical value is to ensure that Type I errors should be relatively unlikely – that the decision “Reject \mathbb{H}_0 ” has scientific strength – yet the test retains power against reasonable alternatives. The decision “Reject \mathbb{H}_0 ” means that the evidence is inconsistent with the null hypothesis, in the sense that it is relatively unlikely (1 in 20) that data generated by the null hypothesis would yield the observed test result.

In contrast, the decision “Accept \mathbb{H}_0 ” is not a strong statement. It does not mean that the evidence supports \mathbb{H}_0 , only that there is insufficient evidence to reject \mathbb{H}_0 . Because of this, it is more accurate to use the label “Do not Reject \mathbb{H}_0 ” instead of “Accept \mathbb{H}_0 ”.

When a test rejects \mathbb{H}_0 at the 5% significance level it is common to say that the statistic is **statistically significant** and if the test accepts \mathbb{H}_0 it is common to say that the statistic is **not statistically significant** or that that it is **statistically insignificant**. It is helpful to remember that this is simply a way of saying “Using the statistic T_n , the hypothesis \mathbb{H}_0 can [cannot] be rejected at the asymptotic 5% level.” When the null hypothesis $\mathbb{H}_0 : \theta = 0$ is rejected it is common to say that the coefficient θ is statistically significant, because the test has rejected the hypothesis that the coefficient is equal to zero.

Let us return to the example about the union wage premium as measured in Table 4.1. The absolute t-statistic for the coefficient on “Male Union Member” is $0.095/0.020 = 4.75$, which is

greater than the 5% asymptotic critical value of 1.96. Therefore we reject the hypothesis that union membership does not affect wages for men. In this case, we can say that union membership is statistically significant for men. However, the absolute t-statistic for the coefficient on “Female Union Member” is $0.022/0.020 = 1.10$, which is less than 1.96 and therefore we do not reject the hypothesis that union membership does not affect wages for women. In this case we find that membership for women is not statistically significant.

When a test accepts a null hypothesis (when a test is not statistically significant), a common misinterpretation is that this is evidence that the null hypothesis is true. This is incorrect. Failure to reject is by itself not evidence. Without an analysis of power, we do not know the likelihood of making a Type II error, and thus are uncertain. In our wage example, it would be a mistake to write that “the regression finds that female union membership has no effect on wages”. This is an incorrect and most unfortunate interpretation. The test has failed to reject the hypothesis that the coefficient is zero, but that does not mean that the coefficient is actually zero.

When a test rejects a null hypothesis (when a test is statistically significant) it does not mean that the null hypothesis is false (as we could be making a Type I error) but we know that this event is unlikely. Thus it is appropriate to interpret the rejection as an evidential statement: The null hypothesis appears to be incorrect. However, while we can conclude that the true value θ is numerically different than the hypothesized value θ_0 , the test alone does not tell us that the true θ value is meaningfully different than θ_0 . That is, whether or not the deviation of θ from θ_0 is meaningful with respect to the interpretation of the coefficient. This is where an examination of confidence intervals can be quite helpful.

8.7 P-Values

Continuing with the wage regression estimates reported in Table 4.1, consider another question: Does marriage status affect wages? To test the hypothesis that marriage status has no effect on wages, we examine the t-statistics for the coefficients on “Married Male” and “Married Female” in Table 4.1, which are $0.180/0.008 = 22.5$ and $0.016/0.008 = 2.0$, respectively. Both exceed the asymptotic 5% critical value of 1.96, so we reject the hypothesis for both men and women. But the statistic for men is exceptionally high, and that for women is only slightly above the critical value. Suppose in contrast that the t-statistic had been 1.9, which is less than the critical value. This would lead to the decision “Accept \mathbb{H}_0 ” rather than “Reject \mathbb{H}_0 ”. Should we really be making a different decision if the t-statistic is 1.9 rather than 2.0? The difference in values is small, shouldn’t the difference in the decision be also small? Thinking through these examples it seems unsatisfactory to simply report “Accept \mathbb{H}_0 ” or “Reject \mathbb{H}_0 ”. These two decisions do not summarize the evidence. Instead, the magnitude of the statistic T_n suggests a “degree of evidence” against \mathbb{H}_0 . How can we take this into account?

The answer is to report what is known as the **asymptotic p-value**

$$p_n = 1 - G(T_n).$$

Since the distribution function G is monotonically increasing, the p-value is a monotonically decreasing function of T_n and is an equivalent test statistic. Instead of rejecting \mathbb{H}_0 at the significance level α if $T_n > c$, we can reject \mathbb{H}_0 if $p_n < \alpha$. Thus it is sufficient to report p_n , and let the reader decide.

In is instructive to interpret p_n as the **marginal significance level**: the largest value of α for which the test T_n “rejects” the null hypothesis. That is, $p_n = 0.11$ means that T_n rejects \mathbb{H}_0 for all significance levels greater than 0.11, but fails to reject \mathbb{H}_0 for significance levels less than 0.11.

Furthermore, the asymptotic p-value has a very convenient asymptotic null distribution. Since

$T_n \xrightarrow{d} T$ under \mathbb{H}_0 , then $p_n = 1 - G(T_n) \xrightarrow{d} 1 - G(T)$, which has the distribution

$$\begin{aligned} \Pr(1 - G(T) \leq u) &= \Pr(1 - u \leq G(T)) \\ &= 1 - \Pr(T \leq G^{-1}(1 - u)) \\ &= 1 - G(G^{-1}(1 - u)) \\ &= 1 - (1 - u) \\ &= u, \end{aligned}$$

which is the uniform distribution on $[0, 1]$. Thus $p_n \xrightarrow{d} U[0, 1]$. This means that the “unusualness” of p_n is easier to interpret than the “unusualness” of T_n .

An important caveat is that the p-value p_n should not be interpreted as the probability that either hypothesis is true. For example, a common mis-interpretation is that p_n is the probability “that the null hypothesis is false.” This is incorrect. Rather, p_n is a measure of the strength of information against the null hypothesis.

Returning to our empirical example, for the test that the coefficient on “Married Male” is zero, the p-value is 0.000. This means that it would be highly unlikely to observe a t-statistic as large as 22.5 when the true value of the coefficient is zero, and thus we can reject that the true value is zero. When presented with such evidence we can say that we “strongly reject” the null hypothesis, that the test is “highly significant”, or that “the test rejects at any conventional critical value”. In contrast, the p-value for the coefficient on “Married Female” is 0.046. In this context it is typical to say that the test is “marginally significant”, meaning that the test statistic is close to the asymptotic 5% critical value.

A related (but somewhat inferior) empirical practice is to append asterisks (*) to coefficient estimates or test statistics to indicate the level of significance. A common practice to append a single asterisk (*) for an estimate or test statistic which exceeds the 10% critical value (i.e., is significant at the 10% level), append a double asterisk (**) for a test which exceeds the 5% critical value, or append a trip asterisk (***) for a test which exceeds the 1% critical value. Such a practice can be better than a table of raw test statistics as the asterisks permit a quick interpretation of significance. On the other hand, asterisks are inferior to p-values, which are also easy and quick to interpret. The goal is essentially the same; it seems wiser to report p-values whenever possible and avoid the use of asterisks.

Our recommendation is that the best empirical practice is to compute and report the asymptotic p-value p_n rather than simply the test statistic T_n , the binary decision Accept/Reject, or appending asterisks. The p-value is a simple statistic, easy to interpret, and contains more information than the other choices.

We now summarize the main features of hypothesis testing.

1. Select a significance level α .
2. Select a test statistic T_n with asymptotic distribution $T_n \xrightarrow{d} T$ under \mathbb{H}_0 .
3. Set the asymptotic critical value c so that $1 - G(c) = \alpha$, where G is the distribution function of T .
4. Calculate the asymptotic p-value $p_n = 1 - G(T_n)$.
5. Reject \mathbb{H}_0 if $T_n > c$, or equivalently $p_n < \alpha$.
6. Accept \mathbb{H}_0 if $T_n \leq c$, or equivalently $p_n \geq \alpha$.
7. Report p_n to summarize the evidence concerning \mathbb{H}_0 versus \mathbb{H}_1 .

8.8 t-ratios and the Abuse of Testing

In Section 4.15, we argued that a good applied practice is to report coefficient estimates $\hat{\theta}$ and standard errors $s(\hat{\theta})$ for all coefficients of interest in estimated models. With $\hat{\theta}$ and $s(\hat{\theta})$ the reader can easily construct confidence intervals $[\hat{\theta} \pm 2s(\hat{\theta})]$ and t-statistics $(\hat{\theta} - \theta_0) / s(\hat{\theta})$ for hypotheses of interest.

Some applied papers (especially older ones) instead report estimates $\hat{\theta}$ and t-ratios $t_n = \hat{\theta} / s(\hat{\theta})$, not standard errors. Reporting t-ratios instead of standard errors is poor econometric practice. While the same information is being reported (you can back out standard errors by division, e.g. $s(\hat{\theta}) = \hat{\theta} / t_n$), standard errors are generally more helpful to readers than t-ratios. Standard errors help the reader focus on the estimation precision and confidence intervals, while t-ratios focus attention on statistical significance. While statistical significance is important, it is less important that the parameter estimates themselves and their confidence intervals. The focus should be on the meaning of the parameter estimates, their magnitudes, and their interpretation, not on listing which variables have significant (e.g. non-zero) coefficients. In many modern applications, sample sizes are very large so standard errors can be very small. Consequently t-ratios can be large even if the coefficient estimates are economically small. In such contexts it may not be interesting to announce “The coefficient is non-zero!” Instead, what is interesting to announce is that “The coefficient estimate is economically interesting!”

In particular, some applied papers report coefficient estimates and t-ratios, and limit their discussion of the results to describing which variables are “significant” (meaning that their t-ratios exceed 2) and the signs of the coefficient estimates. This is very poor empirical work, and should be studiously avoided. It is also a recipe for banishment of your work to lower tier economics journals.

Fundamentally, the common t-ratio is a test for the hypothesis that a coefficient equals zero. This should be reported and discussed when this is an interesting economic hypothesis of interest. But if this is not the case, it is distracting.

In general, when a coefficient θ is of interest, it is constructive to focus on the point estimate, its standard error, and its confidence interval. The point estimate gives our “best guess” for the value. The standard error is a measure of precision. The confidence interval gives us the range of values consistent with the data. If the standard error is large then the point estimate is not a good summary about θ . The endpoints of the confidence interval describe the bounds on the likely possibilities. If the confidence interval embraces too broad a set of values for θ , then the dataset is not sufficiently informative to render useful inferences about θ . On the other hand if the confidence interval is tight, then the data have produced an accurate estimate, and the focus should be on the value and interpretation of this estimate. In contrast, the statement “the t-ratio is highly significant” has little interpretive value.

The above discussion requires that the researcher knows what the coefficient θ means (in terms of the economic problem) and can interpret values and magnitudes, not just signs. This is critical for good applied econometric practice.

For example, consider the question about the effect of marriage status on mean log wages. We had found that the effect is “highly significant” for men and “marginally significant” for women. Now, let’s construct asymptotic confidence intervals for the coefficients. The one for men is [0.16, 0.20] and that for women is [0.00, 0.03]. This shows that average wages for married men are about 16-20% higher than for unmarried men, which is very substantial, while the difference for women is about 0-3%, which is small. These *magnitudes* are more informative than the results of the hypothesis tests.

8.9 Wald Tests

The t-test is appropriate when the null hypothesis is a real-valued restriction. More generally, there may be multiple restrictions on the coefficient vector β . Suppose that we have $q > 1$ restric-

tions which can be written in the form (8.1). It is natural to estimate $\boldsymbol{\theta} = \mathbf{r}(\boldsymbol{\beta})$ by the plug-in estimate $\widehat{\boldsymbol{\theta}} = \mathbf{r}(\widehat{\boldsymbol{\beta}})$. To test $\mathbb{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$ one approach is to measure the magnitude of the discrepancy $\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0$. As this is a vector, there is more than one measure of its length. One simple measure is the weighted quadratic form known as the **Wald statistic**. This is (6.48) evaluated at the null hypothesis

$$W_n = W_n(\boldsymbol{\theta}_0) = (\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)' \widehat{\mathbf{V}}_{\widehat{\boldsymbol{\theta}}}^{-1} (\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \quad (8.10)$$

where $\widehat{\mathbf{V}}_{\widehat{\boldsymbol{\theta}}} = \widehat{\mathbf{R}}' \widehat{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}} \widehat{\mathbf{R}}$ is an estimate of $\mathbf{V}_{\widehat{\boldsymbol{\theta}}}$ and $\widehat{\mathbf{R}} = \frac{\partial}{\partial \boldsymbol{\beta}} \mathbf{r}(\widehat{\boldsymbol{\beta}})'$. Notice that we can write W_n alternatively as

$$W_n = n (\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)' \widehat{\mathbf{V}}_{\boldsymbol{\theta}}^{-1} (\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)$$

using the asymptotic variance estimate $\widehat{\mathbf{V}}_{\boldsymbol{\theta}}$, or we can write it directly as a function of $\widehat{\boldsymbol{\beta}}$ as

$$W_n = (\mathbf{r}(\widehat{\boldsymbol{\beta}}) - \boldsymbol{\theta}_0)' (\widehat{\mathbf{R}}' \widehat{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}} \widehat{\mathbf{R}})^{-1} (\mathbf{r}(\widehat{\boldsymbol{\beta}}) - \boldsymbol{\theta}_0). \quad (8.11)$$

Also, when $\mathbf{r}(\boldsymbol{\beta}) = \mathbf{R}'\boldsymbol{\beta}$ is a linear function of $\boldsymbol{\beta}$, then the Wald statistic simplifies to

$$W_n = (\mathbf{R}'\widehat{\boldsymbol{\beta}} - \boldsymbol{\theta}_0)' (\mathbf{R}'\widehat{\mathbf{V}}_{\widehat{\boldsymbol{\beta}}}\mathbf{R})^{-1} (\mathbf{R}'\widehat{\boldsymbol{\beta}} - \boldsymbol{\theta}_0).$$

The Wald statistic W_n is a weighted Euclidean measure of the length of the vector $\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0$. When $q = 1$ then $W_n = t_n^2$, the square of the t-statistic, so hypothesis tests based on W_n and $|t_n|$ are equivalent. The Wald statistic (8.10) is a generalization of the t-statistic to the case of multiple restrictions.

As shown in Theorem 6.16.2, when $\boldsymbol{\beta}$ satisfies $\mathbf{r}(\boldsymbol{\beta}) = \boldsymbol{\theta}_0$ then $W_n \xrightarrow{d} \chi_q^2$, a chi-square random variable with q degrees of freedom. Let $G_q(u)$ denote the χ_q^2 distribution function. For a given significance level α , the asymptotic critical value c satisfies $\alpha = 1 - G_q(c)$ and can be found from the chi-square distribution table. For example, the 5% critical values for $q = 1$, $q = 2$, and $q = 3$ are 3.84, 5.99, and 7.82, respectively. An asymptotic test rejects \mathbb{H}_0 in favor of \mathbb{H}_1 if $W_n > c$. As with t-tests, it is conventional to describe a Wald test as “significant” if W_n exceeds the 5% asymptotic critical value.

Theorem 8.9.1 *Under Assumptions 6.1.2 and 6.10.1, and $\mathbb{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$, then*

$$W_n \xrightarrow{d} \chi_q^2,$$

and for c satisfying $\alpha = 1 - G_q(c)$,

$$\Pr(W_n > c \mid \mathbb{H}_0) \longrightarrow \alpha$$

so the test “Reject \mathbb{H}_0 if $W_n > c$ ” has asymptotic size α .

Notice that the asymptotic distribution in Theorem 8.9.1 depends solely on q , the number of restrictions being tested. It does not depend on k , the number of parameters estimated.

The asymptotic p-value for W_n is $p_n = 1 - G_q(W_n)$. It is particularly useful to report p-values instead of the Wald statistic. For example, if you write that a Wald test on eight restrictions ($q = 8$) has the value $W_n = 11.2$, it is difficult for a reader to assess the magnitude of this statistic without the time-consuming and cumbersome process of looking up the critical values from a table. Instead, if you write that the p-value is $p_n = 0.19$ (as is the case for $W_n = 11.2$ and $q = 8$) then it is simple for a reader to interpret its magnitude as “insignificant”.

For example, consider the empirical results presented in Table 4.1. The hypothesis “Union membership does not affect wages” is the joint restriction that both coefficients on “Male Union Member” and “Female Union Member” are zero. We calculate the Wald statistic (8.10) for this joint hypothesis and find $W_n = 23.14$ with a p-value of $p_n = 0.000$. Thus we reject the hypothesis in favor of the alternative that at least one of the coefficients is non-zero. This does not mean that both coefficients are non-zero, just that one of the two is non-zero. Therefore examining the joint Wald statistic and the individual t-statistics is useful for interpretation.

The Wald statistic is named after the statistician Abraham Wald, who showed that W_n has optimal weighted average power in certain settings.

8.10 Homoskedastic Wald Tests

If the error is known to be homoskedastic, then it is appropriate to use the homoskedastic Wald statistic (6.50) which replaces $\widehat{\mathbf{V}}_{\hat{\boldsymbol{\theta}}}$ with the homoskedastic estimate $\widehat{\mathbf{V}}_{\hat{\boldsymbol{\theta}}}^0$. This statistic equals

$$\begin{aligned} W_n^0 &= (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)' (\widehat{\mathbf{V}}_{\hat{\boldsymbol{\theta}}}^0)^{-1} (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \\ &= (\mathbf{r}(\hat{\boldsymbol{\beta}}) - \boldsymbol{\theta}_0)' (\widehat{\mathbf{R}}' (\mathbf{X}'\mathbf{X})^{-1} \widehat{\mathbf{R}})^{-1} (\mathbf{r}(\hat{\boldsymbol{\beta}}) - \boldsymbol{\theta}_0) / s^2. \end{aligned} \quad (8.12)$$

We call (8.12) the **homoskedastic Wald statistic** as it is an appropriate test when the errors are conditionally homoskedastic.

As for W_n , when $q = 1$ then $W_n^0 = t_n^2$, the square of the t-statistic, where the latter is computed with a homoskedastic standard error.

In the case of linear hypotheses $\mathbb{H}_0 : \mathbf{R}'\boldsymbol{\beta} = \boldsymbol{\theta}_0$ the homoskedastic Wald statistic equals

$$W_n^0 = (\mathbf{R}'\hat{\boldsymbol{\beta}} - \boldsymbol{\theta}_0)' (\widehat{\mathbf{R}}' (\mathbf{X}'\mathbf{X})^{-1} \widehat{\mathbf{R}})^{-1} (\mathbf{R}'\hat{\boldsymbol{\beta}} - \boldsymbol{\theta}_0) / s^2. \quad (8.13)$$

Theorem 8.10.1 Under Assumptions 6.1.2 and 6.10.1, $\mathbb{E}(e_i^2 | \mathbf{x}_i) = \sigma^2$, and $\mathbb{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$, then

$$W_n^0 \xrightarrow{d} \chi_q^2,$$

and for c satisfying $\alpha = 1 - G_q(c)$,

$$\Pr(W_n^0 > c | \mathbb{H}_0) \longrightarrow \alpha$$

so the test “Reject \mathbb{H}_0 if $W_n^0 > c$ ” has asymptotic size α .

8.11 Criterion-Based Tests

The Wald statistic is based on the length of the vector $\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0$: the discrepancy between the estimate $\hat{\boldsymbol{\theta}} = \mathbf{r}(\hat{\boldsymbol{\beta}})$ and the hypothesized value $\boldsymbol{\theta}_0$. An alternative class of tests is based on the discrepancy between the criterion function minimized with and without the restriction.

Criterion-based testing applies when we have a criterion function, say $J_n(\boldsymbol{\beta})$ with $\boldsymbol{\beta} \in \mathbf{B}$, which is minimized for estimation, and the goal is to test $\mathbb{H}_0 : \boldsymbol{\beta} \in \mathbf{B}_0$ versus $\mathbb{H}_1 : \boldsymbol{\beta} \notin \mathbf{B}_0$

where $\mathbf{B}_0 \subset \mathbf{B}$. Minimizing the criterion function over \mathbf{B} and \mathbf{B}_0 we obtain the unrestricted and restricted estimators

$$\begin{aligned}\widehat{\boldsymbol{\beta}} &= \operatorname{argmin}_{\boldsymbol{\beta} \in \mathbf{B}} J_n(\boldsymbol{\beta}) \\ \widetilde{\boldsymbol{\beta}} &= \operatorname{argmin}_{\boldsymbol{\beta} \in \mathbf{B}_0} J_n(\boldsymbol{\beta}).\end{aligned}$$

The **criterion-based statistic** for \mathbb{H}_0 versus \mathbb{H}_1 is proportional to

$$\begin{aligned}J_n &= \min_{\boldsymbol{\beta} \in \mathbf{B}_0} J_n(\boldsymbol{\beta}) - \min_{\boldsymbol{\beta} \in \mathbf{B}} J_n(\boldsymbol{\beta}) \\ &= J_n(\widetilde{\boldsymbol{\beta}}) - J_n(\widehat{\boldsymbol{\beta}}).\end{aligned}$$

The criterion-based statistic J_n is sometimes called a **distance** statistic, a **minimum-distance** statistic, or a **likelihood-ratio-like** statistic.

Since \mathbf{B}_0 is a subset of \mathbf{B} , $J_n(\widetilde{\boldsymbol{\beta}}) \geq J_n(\widehat{\boldsymbol{\beta}})$ and thus $J_n \geq 0$. The statistic J_n measures the cost (on the criterion) of imposing the null restriction $\boldsymbol{\beta} \in \mathbf{B}_0$.

8.12 Minimum Distance Tests

The minimum distance test is a criterion-based test where $J_n(\boldsymbol{\beta})$ is the minimum distance criterion (7.17)

$$J_n(\boldsymbol{\beta}) = n \left(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta} \right)' \mathbf{W}_n \left(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta} \right) \quad (8.14)$$

with $\widehat{\boldsymbol{\beta}}$ the unrestricted (LS) estimator. The restricted estimator $\widetilde{\boldsymbol{\beta}}_{\text{md}}$ minimizes (8.14) subject to $\boldsymbol{\beta} \in \mathbf{B}_0$. Observing that $J_n(\widehat{\boldsymbol{\beta}}) = 0$, the minimum distance statistic simplifies to

$$J_n = J_n(\widetilde{\boldsymbol{\beta}}_{\text{md}}) = n \left(\widehat{\boldsymbol{\beta}} - \widetilde{\boldsymbol{\beta}}_{\text{md}} \right)' \mathbf{W}_n \left(\widehat{\boldsymbol{\beta}} - \widetilde{\boldsymbol{\beta}}_{\text{md}} \right). \quad (8.15)$$

The efficient minimum distance estimator $\widetilde{\boldsymbol{\beta}}_{\text{emd}}$ is obtained by setting $\mathbf{W}_n = \widehat{\mathbf{V}}_{\boldsymbol{\beta}}^{-1}$ in (8.14) and (8.15). The efficient minimum distance statistic for $\mathbb{H}_0 : \boldsymbol{\beta} \in \mathbf{B}_0$ is therefore

$$J_n^* = n \left(\widehat{\boldsymbol{\beta}} - \widetilde{\boldsymbol{\beta}}_{\text{emd}} \right)' \widehat{\mathbf{V}}_{\boldsymbol{\beta}}^{-1} \left(\widehat{\boldsymbol{\beta}} - \widetilde{\boldsymbol{\beta}}_{\text{emd}} \right). \quad (8.16)$$

Consider the class of linear hypotheses $\mathbb{H}_0 : \mathbf{R}'\boldsymbol{\beta} = \boldsymbol{\theta}_0$. In this case we know from (7.25) that the efficient minimum distance estimator $\widetilde{\boldsymbol{\beta}}_{\text{emd}}$ subject to the constraint $\mathbf{R}'\boldsymbol{\beta} = \boldsymbol{\theta}_0$ is

$$\widetilde{\boldsymbol{\beta}}_{\text{emd}} = \widehat{\boldsymbol{\beta}} - \widehat{\mathbf{V}}_{\boldsymbol{\beta}} \mathbf{R} \left(\mathbf{R}' \widehat{\mathbf{V}}_{\boldsymbol{\beta}} \mathbf{R} \right)^{-1} \left(\mathbf{R}' \widehat{\boldsymbol{\beta}} - \boldsymbol{\theta}_0 \right)$$

and thus

$$\widehat{\boldsymbol{\beta}} - \widetilde{\boldsymbol{\beta}}_{\text{emd}} = \widehat{\mathbf{V}}_{\boldsymbol{\beta}} \mathbf{R} \left(\mathbf{R}' \widehat{\mathbf{V}}_{\boldsymbol{\beta}} \mathbf{R} \right)^{-1} \left(\mathbf{R}' \widehat{\boldsymbol{\beta}} - \boldsymbol{\theta}_0 \right).$$

Substituting into (8.16) we find

$$\begin{aligned}J_n^* &= n \left(\mathbf{R}' \widehat{\boldsymbol{\beta}} - \boldsymbol{\theta}_0 \right)' \left(\mathbf{R}' \widehat{\mathbf{V}}_{\boldsymbol{\beta}} \mathbf{R} \right)^{-1} \mathbf{R}' \widehat{\mathbf{V}}_{\boldsymbol{\beta}} \widehat{\mathbf{V}}_{\boldsymbol{\beta}}^{-1} \widehat{\mathbf{V}}_{\boldsymbol{\beta}} \mathbf{R} \left(\mathbf{R}' \widehat{\mathbf{V}}_{\boldsymbol{\beta}} \mathbf{R} \right)^{-1} \left(\mathbf{R}' \widehat{\boldsymbol{\beta}} - \boldsymbol{\theta}_0 \right) \\ &= n \left(\mathbf{R}' \widehat{\boldsymbol{\beta}} - \boldsymbol{\theta}_0 \right)' \left(\mathbf{R}' \widehat{\mathbf{V}}_{\boldsymbol{\beta}} \mathbf{R} \right)^{-1} \left(\mathbf{R}' \widehat{\boldsymbol{\beta}} - \boldsymbol{\theta}_0 \right) \\ &= W_n,\end{aligned} \quad (8.17)$$

which is the Wald statistic (8.10).

Thus for linear hypotheses $\mathbb{H}_0 : \mathbf{R}'\boldsymbol{\beta} = \boldsymbol{\theta}_0$, the efficient minimum distance statistic J_n^* is identical to the Wald statistic (8.10). For non-linear hypotheses, however, the Wald and minimum distance statistics are different.

Newey and West (1987) established the asymptotic null distribution of J_n^* for linear and non-linear hypotheses.

Theorem 8.12.1 *Under Assumptions 6.1.2 and 6.10.1, and $\mathbb{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$, then $J_n^* \xrightarrow{d} \chi_q^2$.*

Testing using the minimum distance statistic J_n^* is similar to testing using the Wald statistic W_n . Critical values and p-values are computed using the χ_q^2 distribution. \mathbb{H}_0 is rejected in favor of \mathbb{H}_1 if J_n^* exceeds the level α critical value. The asymptotic p-value is $p_n = 1 - F_q(J_n^*)$.

8.13 Minimum Distance Tests Under Homoskedasticity

If we set $\mathbf{W}_n = \widehat{\mathbf{Q}}_{xx}$ in (8.14) we obtain the criterion (7.19)

$$J_n^0(\boldsymbol{\beta}) = n \left(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta} \right)' \widehat{\mathbf{Q}}_{xx} \left(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta} \right).$$

A minimum distance statistic for $\mathbb{H}_0 : \boldsymbol{\beta} \in \mathbf{B}_0$ is

$$J_n^0 = \min_{\boldsymbol{\beta} \in \mathbf{B}_0} J_n^0(\boldsymbol{\beta}) / s^2.$$

Notice that we have scaled the criterion by the unbiased variance estimator s^2 from (4.21) for reasons which will become clear momentarily.

Equation (7.20) showed that

$$SSE_n(\boldsymbol{\beta}) = n\hat{\sigma}^2 + J_n^0(\boldsymbol{\beta})$$

and so the minimizers of $SSE_n(\boldsymbol{\beta})$ and $J_n^0(\boldsymbol{\beta})$ are identical. Thus the constrained minimizer of $J_n^0(\boldsymbol{\beta})$ is constrained least-squares

$$\tilde{\boldsymbol{\beta}}_{\text{cls}} = \underset{\boldsymbol{\beta} \in \mathbf{B}_0}{\operatorname{argmin}} J_n^0(\boldsymbol{\beta}) = \underset{\boldsymbol{\beta} \in \mathbf{B}_0}{\operatorname{argmin}} SSE_n(\boldsymbol{\beta}) \quad (8.18)$$

and therefore

$$\begin{aligned} J_n^0 &= J_n^0(\tilde{\boldsymbol{\beta}}_{\text{cls}}) / s^2 \\ &= n \left(\widehat{\boldsymbol{\beta}} - \tilde{\boldsymbol{\beta}}_{\text{cls}} \right)' \widehat{\mathbf{Q}}_{xx} \left(\widehat{\boldsymbol{\beta}} - \tilde{\boldsymbol{\beta}}_{\text{cls}} \right) / s^2. \end{aligned}$$

In the special case of linear hypotheses $\mathbb{H}_0 : \mathbf{R}'\boldsymbol{\beta} = \boldsymbol{\theta}_0$, the constrained least-squares estimator subject to $\mathbf{R}'\boldsymbol{\beta} = \boldsymbol{\theta}_0$ has the solution (7.10)

$$\tilde{\boldsymbol{\beta}}_{\text{cls}} = \widehat{\boldsymbol{\beta}} - \widehat{\mathbf{Q}}_{xx}^{-1} \mathbf{R} \left(\mathbf{R}' \widehat{\mathbf{Q}}_{xx}^{-1} \mathbf{R} \right)^{-1} \left(\mathbf{R}' \widehat{\boldsymbol{\beta}} - \boldsymbol{\theta}_0 \right)$$

and solving we find

$$J_n^0 = n \left(\mathbf{R}' \widehat{\boldsymbol{\beta}} - \boldsymbol{\theta}_0 \right)' \left(\mathbf{R}' \widehat{\mathbf{Q}}_{xx}^{-1} \mathbf{R} \right)^{-1} \left(\mathbf{R}' \widehat{\boldsymbol{\beta}} - \boldsymbol{\theta}_0 \right) / s^2 = W_n^0. \quad (8.19)$$

This is the homoskedastic Wald statistic (8.13). Thus for testing linear hypotheses, homoskedastic minimum distance and Wald statistics agree.

For nonlinear hypotheses they disagree, but have the same null asymptotic distribution.

Theorem 8.13.1 Under Assumptions 6.1.2 and 6.10.1, $\mathbb{E}(e_i^2 | \mathbf{x}_i) = \sigma^2$, and $\mathbb{H}_0 : \boldsymbol{\theta} = \boldsymbol{\theta}_0$, then $J_n^0 \xrightarrow{d} \chi_q^2$.

8.14 F Tests

The F statistic for testing $\mathbb{H}_0 : \boldsymbol{\beta} \in \mathbf{B}_0$ is

$$F_n = \frac{\left(SSE_n(\tilde{\boldsymbol{\beta}}_{\text{cls}}) - SSE_n(\hat{\boldsymbol{\beta}}) \right) / q}{SSE_n(\hat{\boldsymbol{\beta}}) / (n - k)} \quad (8.20)$$

where

$$SSE_n(\boldsymbol{\beta}) = \sum_{i=1}^n (y_i - \mathbf{x}_i' \boldsymbol{\beta})^2$$

is the sum-of-squared errors, $\tilde{\boldsymbol{\beta}}_{\text{cls}}$ is the constrained least-squares estimator (8.18), $\hat{\boldsymbol{\beta}}$ is the unconstrained least-squares estimator, q is the number of restrictions, and k is the number of unconstrained coefficients.

Noting that $s^2 = SSE_n(\hat{\boldsymbol{\beta}}) / (n - k)$, we can also write (8.20) as

$$F_n = \frac{SSE_n(\tilde{\boldsymbol{\beta}}_{\text{cls}}) - SSE_n(\hat{\boldsymbol{\beta}})}{qs^2}$$

which is a scale of the difference of sum-of-squared errors, and is thus a criterion-based statistic. Using (7.20) we can also write the statistic as

$$F_n = J_n^0 / q,$$

so the F statistic is identical to the homoskedastic minimum distance statistic divided by the number of restrictions q .

Another useful way of writing (8.20) is

$$F_n = \left(\frac{n - k}{q} \right) \frac{(\tilde{\sigma}^2 - \hat{\sigma}^2)}{\hat{\sigma}^2} \quad (8.21)$$

where

$$\hat{\sigma}^2 = \frac{SSE_n(\hat{\boldsymbol{\beta}})}{n} = \frac{1}{n} \sum_{i=1}^n \hat{e}_i^2$$

is the residual variance estimate under \mathbb{H}_1 and

$$\tilde{\sigma}^2 = \frac{SSE_n(\tilde{\boldsymbol{\beta}}_{\text{cls}})}{n} = \frac{1}{n} \sum_{i=1}^n \tilde{e}_i^2$$

with $\tilde{e}_i = y_i - \mathbf{x}_i' \tilde{\boldsymbol{\beta}}_{\text{cls}}$ is the residual variance estimate under \mathbb{H}_0 .

As we discussed in the previous section, in the special case of linear hypotheses $\mathbb{H}_0 : \mathbf{R}'\boldsymbol{\beta} = \boldsymbol{\theta}_0$, $J_n^0 = W_n^0$. It follows that in this case $F_n = W_n^0 / q$. Thus the F_n statistic equals the homoskedastic Wald statistic divided by q . It follows that they are equivalent tests for \mathbb{H}_0 against \mathbb{H}_1 .

In many statistical packages, linear hypothesis tests are reported as F statistics rather than Wald statistics. While they are equivalent, it is important to know which is being reported to know which critical values to use. (If p-values are directly reported this is not an issue.)

When reporting an F_n statistic it is conventional to calculate critical values and p-values using the $F(q, n - k)$ distribution instead of the asymptotic χ_q^2/q distribution. This is a prudent small sample adjustment, as the F distribution is exact when the errors are independent of the regressors and normally distributed. However, when the degrees of freedom $n - k$ are large then the difference is negligible. More relevantly, if $n - k$ is small enough to make a difference, probably we shouldn't be trusting the asymptotic approximation anyway!

An elegant feature about (8.20) or (8.21) is that they are directly computable from the standard output from two simple OLS regressions, as the sum of squared errors (or regression variance) is a typical printed output from statistical packages, and is often reported in applied tables. Thus F_n can be calculated by hand from standard reported statistics even if you don't have the original data (or if you are sitting in a seminar and listening to a presentation!).

If you are presented with an F_n statistic (or a Wald statistic, as you can just divide by q) but don't have access to critical values, a useful rule of thumb is to know that for large n , the 5% asymptotic critical value is decreasing as q increases, and is less than 2 for $q \geq 7$.

In many statistical packages, when an OLS regression is estimated an "F-statistic" is automatically reported, even though hypothesis test is requested. This is the F statistic F_n where \mathbb{H}_0 restricts all coefficients except the intercept to be zero. This was a popular statistic in the early days of econometric reporting, when sample sizes were very small and researchers wanted to know if there was "any explanatory power" to their regression. This is rarely an issue today, as sample sizes are typically sufficiently large that this F statistic is nearly always highly significant. While there are special cases where this F statistic is useful, these cases are atypical. As a general rule, there is no reason to report this F statistic.

The F_n statistic is named after the statistician Ronald Fisher, one of the founders of modern statistical theory.

8.15 Likelihood Ratio Test

For a model with parameter $\boldsymbol{\theta} \in \Theta$ and likelihood function $L_n(\boldsymbol{\theta})$ the likelihood ratio statistic for $\mathbb{H}_0 : \boldsymbol{\theta} \in \Theta_0$ versus $\mathbb{H}_1 : \boldsymbol{\theta} \in \Theta_0^c$ is

$$\begin{aligned} LR_n &= 2 \left(\max_{\boldsymbol{\theta} \in \Theta} \log L_n(\boldsymbol{\theta}) - \max_{\boldsymbol{\theta} \in \Theta_0} \log L_n(\boldsymbol{\theta}) \right) \\ &= 2 \left(\log L_n(\hat{\boldsymbol{\theta}}) - \log L_n(\tilde{\boldsymbol{\theta}}) \right) \end{aligned}$$

where $\hat{\boldsymbol{\theta}}$ and $\tilde{\boldsymbol{\theta}}$ are the unrestricted and constrained MLE.

In the normal linear model the maximized log likelihood (3.45) at the unrestricted and restricted estimates are

$$\log L \left(\hat{\boldsymbol{\beta}}_{\text{mle}}, \hat{\sigma}_{\text{mle}}^2 \right) = -\frac{n}{2} (\log(2\pi) + 1) - \frac{n}{2} \log(\hat{\sigma}^2)$$

and

$$\log L \left(\tilde{\boldsymbol{\beta}}_{\text{cmle}}, \tilde{\sigma}_{\text{cmle}}^2 \right) = -\frac{n}{2} (\log(2\pi) + 1) - \frac{n}{2} \log(\tilde{\sigma}^2)$$

respectively. Thus the LR statistic is

$$\begin{aligned} LR_n &= n \left(\log(\tilde{\sigma}^2) - \log(\hat{\sigma}^2) \right) \\ &= n \log \left(\frac{\tilde{\sigma}^2}{\hat{\sigma}^2} \right) \end{aligned}$$

which is a monotonic function of $\tilde{\sigma}^2/\hat{\sigma}^2$. Recall that the F statistic (8.21) is also a monotonic function of $\tilde{\sigma}^2/\hat{\sigma}^2$. Thus LR_n and F_n are fundamentally the same statistic and have the same information about \mathbb{H}_0 versus \mathbb{H}_1 .

Furthermore, by a first-order Taylor series approximation

$$LR_n/q = \frac{n}{q} \log \left(1 + \frac{\tilde{\sigma}^2}{\hat{\sigma}^2} - 1 \right) \simeq \frac{n}{q} \left(\frac{\tilde{\sigma}^2}{\hat{\sigma}^2} - 1 \right) \simeq F_n.$$

This shows that the two statistics (LR_n and F_n) will be numerically close. It also shows that the F statistic and the homoskedastic Wald statistic for linear hypotheses can also be interpreted as approximate likelihood ratio statistics under normality.

8.16 Problems with Tests of NonLinear Hypotheses

While the t and Wald tests work well when the hypothesis is a linear restriction on β , they can work quite poorly when the restrictions are nonlinear. This can be seen by a simple example introduced by Lafontaine and White (1986). Take the model

$$\begin{aligned} y_i &= \beta + e_i \\ e_i &\sim N(0, \sigma^2) \end{aligned}$$

and consider the hypothesis

$$\mathbb{H}_0 : \beta = 1.$$

Let $\hat{\beta}$ and $\hat{\sigma}^2$ be the sample mean and variance of y_i . The standard Wald test for \mathbb{H}_0 is

$$W_n = n \frac{(\hat{\beta} - 1)^2}{\hat{\sigma}^2}.$$

Now notice that \mathbb{H}_0 is equivalent to the hypothesis

$$\mathbb{H}_0(s) : \beta^s = 1$$

for any positive integer s . Letting $r(\beta) = \beta^s$, and noting $\mathbf{R} = s\beta^{s-1}$, we find that the standard Wald test for $\mathbb{H}_0(s)$ is

$$W_n(s) = n \frac{(\hat{\beta}^s - 1)^2}{\hat{\sigma}^2 s^2 \hat{\beta}^{2s-2}}.$$

While the hypothesis $\beta^s = 1$ is unaffected by the choice of s , the statistic $W_n(s)$ varies with s . This is an unfortunate feature of the Wald statistic.

To demonstrate this effect, we have plotted in Figure 8.1 the Wald statistic $W_n(s)$ as a function of s , setting $n/\hat{\sigma}^2 = 10$. The increasing solid line is for the case $\hat{\beta} = 0.8$. The decreasing dashed line is for the case $\hat{\beta} = 1.6$. It is easy to see that in each case there are values of s for which the test statistic is significant relative to asymptotic critical values, while there are other values of s for which the test statistic is insignificant. This is distressing since the choice of s is arbitrary and irrelevant to the actual hypothesis.

Our first-order asymptotic theory is not useful to help pick s , as $W_n(s) \xrightarrow{d} \chi_1^2$ under \mathbb{H}_0 for any s . This is a context where **Monte Carlo simulation** can be quite useful as a tool to study and compare the exact distributions of statistical procedures in finite samples. The method uses random simulation to create artificial datasets, to which we apply the statistical tools of interest. This produces random draws from the statistic's sampling distribution. Through repetition, features of this distribution can be calculated.

In the present context of the Wald statistic, one feature of importance is the Type I error of the test using the asymptotic 5% critical value 3.84 – the probability of a false rejection, $\Pr(W_n(s) > 3.84 \mid \beta = 1)$. Given the simplicity of the model, this probability depends only on s , n , and σ^2 . In Table 8.1 we report the results of a Monte Carlo simulation where we vary these

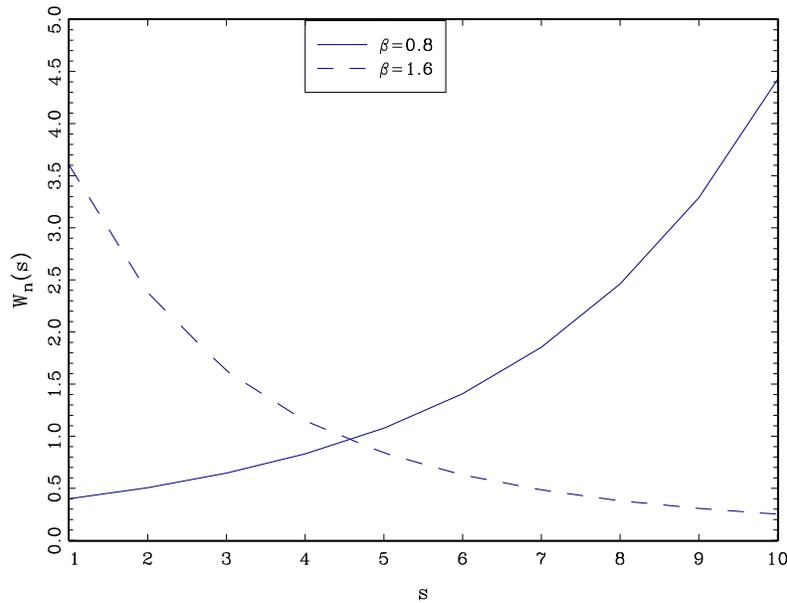


Figure 8.1: Wald Statistic as a function of s

three parameters. The value of s is varied from 1 to 10, n is varied among 20, 100 and 500, and σ is varied among 1 and 3. The Table reports the simulation estimate of the Type I error probability from 50,000 random samples. Each row of the table corresponds to a different value of s – and thus corresponds to a particular choice of test statistic. The second through seventh columns contain the Type I error probabilities for different combinations of n and σ . These probabilities are calculated as the percentage of the 50,000 simulated Wald statistics $W_n(s)$ which are larger than 3.84. The null hypothesis $\beta^s = 1$ is true, so these probabilities are Type I error.

To interpret the table, remember that the ideal Type I error probability is 5% (.05) with deviations indicating distortion. Type I error rates between 3% and 8% are considered reasonable. Error rates above 10% are considered excessive. Rates above 20% are unacceptable. When comparing statistical procedures, we compare the rates row by row, looking for tests for which rejection rates are close to 5% and rarely fall outside of the 3%-8% range. For this particular example the only test which meets this criterion is the conventional $W_n = W_n(1)$ test. Any other choice of s leads to a test with unacceptable Type I error probabilities.

Table 8.1
Type I Error Probability of Asymptotic 5% $W_n(s)$ Test

s	$\sigma = 1$			$\sigma = 3$		
	$n = 20$	$n = 100$	$n = 500$	$n = 20$	$n = 100$	$n = 500$
1	.06	.05	.05	.07	.05	.05
2	.08	.06	.05	.15	.08	.06
3	.10	.06	.05	.21	.12	.07
4	.13	.07	.06	.25	.15	.08
5	.15	.08	.06	.28	.18	.10
6	.17	.09	.06	.30	.20	.11
7	.19	.10	.06	.31	.22	.13
8	.20	.12	.07	.33	.24	.14
9	.22	.13	.07	.34	.25	.15
10	.23	.14	.08	.35	.26	.16

Note: Rejection frequencies from 50,000 simulated random samples

In Table 8.1 you can also see the impact of variation in sample size. In each case, the Type I error probability improves towards 5% as the sample size n increases. There is, however, no magic choice of n for which all tests perform uniformly well. Test performance deteriorates as s increases, which is not surprising given the dependence of $W_n(s)$ on s as shown in Figure 8.1.

In this example it is not surprising that the choice $s = 1$ yields the best test statistic. Other choices are arbitrary and would not be used in practice. While this is clear in this particular example, in other examples natural choices are not always obvious and the best choices may in fact appear counter-intuitive at first.

This point can be illustrated through another example which is similar to one developed in Gregory and Veall (1985). Take the model

$$\begin{aligned} y_i &= \beta_0 + x_{1i}\beta_1 + x_{2i}\beta_2 + e_i \\ \mathbb{E}(\mathbf{x}_i e_i) &= \mathbf{0} \end{aligned} \tag{8.22}$$

and the hypothesis

$$\mathbb{H}_0 : \frac{\beta_1}{\beta_2} = \theta_0$$

where θ_0 is a known constant. Equivalently, define $\theta = \beta_1/\beta_2$, so the hypothesis can be stated as $\mathbb{H}_0 : \theta = \theta_0$.

Let $\hat{\boldsymbol{\beta}} = (\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2)$ be the least-squares estimates of (8.22), let $\hat{\mathbf{V}}_{\hat{\boldsymbol{\beta}}}$ be an estimate of the covariance matrix for $\hat{\boldsymbol{\beta}}$ and set $\hat{\theta} = \hat{\beta}_1/\hat{\beta}_2$. Define

$$\hat{\mathbf{R}}_1 = \begin{pmatrix} 0 \\ \frac{1}{\hat{\beta}_2} \\ -\frac{\hat{\beta}_1}{\hat{\beta}_2^2} \end{pmatrix}$$

so that the standard error for $\hat{\theta}$ is $s(\hat{\theta}) = \left(\hat{\mathbf{R}}_1' \hat{\mathbf{V}}_{\hat{\boldsymbol{\beta}}} \hat{\mathbf{R}}_1\right)^{1/2}$. In this case a t-statistic for \mathbb{H}_0 is

$$t_{1n} = \frac{\left(\frac{\hat{\beta}_1}{\hat{\beta}_2} - \theta_0\right)}{s(\hat{\theta})}.$$

An alternative statistic can be constructed through reformulating the null hypothesis as

$$\mathbb{H}_0 : \beta_1 - \theta_0\beta_2 = 0.$$

A t-statistic based on this formulation of the hypothesis is

$$t_{2n} = \frac{\hat{\beta}_1 - \theta_0\hat{\beta}_2}{\left(\mathbf{R}_2' \hat{\mathbf{V}}_{\hat{\boldsymbol{\beta}}} \mathbf{R}_2\right)^{1/2}}.$$

where

$$\mathbf{R}_2 = \begin{pmatrix} 0 \\ 1 \\ -\theta_0 \end{pmatrix}.$$

To compare t_{1n} and t_{2n} we perform another simple Monte Carlo simulation. We let x_{1i} and x_{2i} be mutually independent $N(0, 1)$ variables, e_i be an independent $N(0, \sigma^2)$ draw with $\sigma = 3$, and normalize $\beta_0 = 0$ and $\beta_1 = 1$. This leaves β_2 as a free parameter, along with sample size n . We vary β_2 among .1, .25, .50, .75, and 1.0 and n among 100 and 500.

Table 8.2
Type I Error Probability of Asymptotic 5% t-tests

	$n = 100$				$n = 500$			
	$\Pr(t_n < -1.645)$		$\Pr(t_n > 1.645)$		$\Pr(t_n < -1.645)$		$\Pr(t_n > 1.645)$	
β_2	t_{1n}	t_{2n}	t_{1n}	t_{2n}	t_{1n}	t_{2n}	t_{1n}	t_{2n}
.10	.47	.06	.00	.06	.28	.05	.00	.05
.25	.26	.06	.00	.06	.15	.05	.00	.05
.50	.15	.06	.00	.06	.10	.05	.00	.05
.75	.12	.06	.00	.06	.09	.05	.00	.05
1.00	.10	.06	.00	.06	.07	.05	.02	.05

The one-sided Type I error probabilities $\Pr(t_n < -1.645)$ and $\Pr(t_n > 1.645)$ are calculated from 50,000 simulated samples. The results are presented in Table 8.2. Ideally, the entries in the table should be 0.05. However, the rejection rates for the t_{1n} statistic diverge greatly from this value, especially for small values of β_2 . The left tail probabilities $\Pr(t_{1n} < -1.645)$ greatly exceed 5%, while the right tail probabilities $\Pr(t_{1n} > 1.645)$ are close to zero in most cases. In contrast, the rejection rates for the linear t_{2n} statistic are invariant to the value of β_2 , and are close to the ideal 5% rate for both sample sizes. The implication of Table 4.2 is that the two t-ratios have dramatically different sampling behavior.

The common message from both examples is that Wald statistics are sensitive to the algebraic formulation of the null hypothesis.

A simple solution is to use the minimum distance statistic J_n , which equals W_n with $r = 1$ in the first example, and $|t_{2n}|$ in the second example. The minimum distance statistic is invariant to the algebraic formulation of the null hypothesis, so is immune to this problem. Whenever possible, the Wald statistic should not be used to test nonlinear hypotheses.

8.17 Monte Carlo Simulation

In the Section 8.16 we introduced the method of Monte Carlo simulation to illustrate the small sample problems with tests of nonlinear hypotheses. In this section we describe the method in more detail.

Recall, our data consist of observations (y_i, \mathbf{x}_i) which are random draws from a population distribution F . Let $\boldsymbol{\theta}$ be a parameter and let $T_n = T_n((y_1, \mathbf{x}_1), \dots, (y_n, \mathbf{x}_n), \boldsymbol{\theta})$ be a statistic of interest, for example an estimator $\hat{\theta}$ or a t-statistic $(\hat{\theta} - \theta)/s(\hat{\theta})$. The exact distribution of T_n is

$$G_n(u, F) = \Pr(T_n \leq u \mid F).$$

While the asymptotic distribution of T_n might be known, the exact (finite sample) distribution G_n is generally unknown.

Monte Carlo simulation uses numerical simulation to compute $G_n(u, F)$ for selected choices of F . This is useful to investigate the performance of the statistic T_n in reasonable situations and sample sizes. The basic idea is that for any given F , the distribution function $G_n(u, F)$ can be calculated numerically through simulation. The name Monte Carlo derives from the famous Mediterranean gambling resort where games of chance are played.

The method of Monte Carlo is quite simple to describe. The researcher chooses F (the distribution of the data) and the sample size n . A “true” value of $\boldsymbol{\theta}$ is implied by this choice, or equivalently the value $\boldsymbol{\theta}$ is selected directly by the researcher which implies restrictions on F .

Then the following experiment is conducted by computer simulation:

1. n independent random pairs (y_i^*, \mathbf{x}_i^*) , $i = 1, \dots, n$, are drawn from the distribution F using the computer’s random number generator.
2. The statistic $T_n = T_n((y_1^*, \mathbf{x}_1^*), \dots, (y_n^*, \mathbf{x}_n^*), \boldsymbol{\theta})$ is calculated on this pseudo data.

For step 1, most computer packages have built-in procedures for generating $U[0, 1]$ and $N(0, 1)$ random numbers, and from these most random variables can be constructed. (For example, a chi-square can be generated by sums of squares of normals.)

For step 2, it is important that the statistic be evaluated at the “true” value of θ corresponding to the choice of F .

The above experiment creates one random draw from the distribution $G_n(u, F)$. This is one observation from an unknown distribution. Clearly, from one observation very little can be said. So the researcher repeats the experiment B times, where B is a large number. Typically, we set $B = 1000$ or $B = 5000$. We will discuss this choice later.

Notationally, let the b^{th} experiment result in the draw T_{nb} , $b = 1, \dots, B$. These results are stored. After all B experiments have been calculated, these results constitute a random sample of size B from the distribution of $G_n(u, F) = \Pr(T_{nb} \leq u) = \Pr(T_n \leq u \mid F)$.

From a random sample, we can estimate any feature of interest using (typically) a method of moments estimator. We now describe some specific examples.

Suppose we are interested in the bias, mean-squared error (MSE), and/or variance of the distribution of $\hat{\theta} - \theta$. We then set $T_n = \hat{\theta} - \theta$, run the above experiment, and calculate

$$\begin{aligned}\widehat{Bias}(\hat{\theta}) &= \frac{1}{B} \sum_{b=1}^B T_{nb} = \frac{1}{B} \sum_{b=1}^B \hat{\theta}_b - \theta \\ \widehat{MSE}(\hat{\theta}) &= \frac{1}{B} \sum_{b=1}^B (T_{nb})^2 = \frac{1}{B} \sum_{b=1}^B (\hat{\theta}_b - \theta)^2 \\ \widehat{\text{var}}(\hat{\theta}) &= \widehat{MSE}(\hat{\theta}) - \left(\widehat{Bias}(\hat{\theta}) \right)^2\end{aligned}$$

Suppose we are interested in the Type I error associated with an asymptotic 5% two-sided t-test. We would then set $T_n = \left| \hat{\theta} - \theta \right| / s(\hat{\theta})$ and calculate

$$\hat{P} = \frac{1}{B} \sum_{b=1}^B 1(T_{nb} \geq 1.96), \quad (8.23)$$

the percentage of the simulated t-ratios which exceed the asymptotic 5% critical value.

Suppose we are interested in the 5% and 95% quantile of $T_n = \hat{\theta}$ or $T_n = \left(\hat{\theta} - \theta \right) / s(\hat{\theta})$. We then compute the 5% and 95% sample quantiles of the sample $\{T_{nb}\}$. The $\alpha\%$ sample quantile is a number q_α such that $\alpha\%$ of the sample are less than q_α . A simple way to compute sample quantiles is to sort the sample $\{T_{nb}\}$ from low to high. Then q_α is the N 'th number in this ordered sequence, where $N = (B + 1)\alpha$. It is therefore convenient to pick B so that N is an integer. For example, if we set $B = 999$, then the 5% sample quantile is 50'th sorted value and the 95% sample quantile is the 950'th sorted value.

The typical purpose of a Monte Carlo simulation is to investigate the performance of a statistical procedure (estimator or test) in realistic settings. Generally, the performance will depend on n and F . In many cases, an estimator or test may perform wonderfully for some values, and poorly for others. It is therefore useful to conduct a variety of experiments, for a selection of choices of n and F .

As discussed above, the researcher must select the number of experiments, B . Often this is called the number of **replications**. Quite simply, a larger B results in more precise estimates of the features of interest of G_n , but requires more computational time. In practice, therefore, the choice of B is often guided by the computational demands of the statistical procedure. Since the results of a Monte Carlo experiment are estimates computed from a random sample of size B , it is straightforward to calculate standard errors for any quantity of interest. If the standard error is too large to make a reliable inference, then B will have to be increased.

In particular, it is simple to make inferences about rejection probabilities from statistical tests, such as the percentage estimate reported in (8.23). The random variable $1(T_{nb} \geq 1.96)$ is iid Bernoulli, equalling 1 with probability $p = \mathbb{E}1(T_{nb} \geq 1.96)$. The average (8.23) is therefore an unbiased estimator of p with standard error $s(\hat{p}) = \sqrt{p(1-p)/B}$. As p is unknown, this may be approximated by replacing p with \hat{p} or with an hypothesized value. For example, if we are assessing an asymptotic 5% test, then we can set $s(\hat{p}) = \sqrt{(.05)(.95)/B} \simeq .22/\sqrt{B}$. Hence, standard errors for $B = 100, 1000, \text{ and } 5000$, are, respectively, $s(\hat{p}) = .022, .007, \text{ and } .003$.

Most papers in econometric methods, and some empirical papers, include the results of Monte Carlo simulations to illustrate the performance of their methods. When extending existing results, it is good practice to start by replicating existing (published) results. This is not exactly possible in the case of simulation results, as they are inherently random. For example suppose a paper investigates a statistical test, and reports a simulated rejection probability of 0.07 based on a simulation with $B = 100$ replications. Suppose you attempt to replicate this result, and find a rejection probability of 0.03 (again using $B = 100$ simulation replications). Should you conclude that you have failed in your attempt? Absolutely not! Under the hypothesis that both simulations are identical, you have two independent estimates, $\hat{p}_1 = 0.07$ and $\hat{p}_2 = 0.03$, of a common probability p . The asymptotic (as $B \rightarrow \infty$) distribution of their difference is $\sqrt{B}(\hat{p}_1 - \hat{p}_2) \xrightarrow{d} N(0, 2p(1-p))$, so a standard error for $\hat{p}_1 - \hat{p}_2 = 0.04$ is $\hat{s} = \sqrt{2p(1-p)/B} \simeq 0.03$, using the estimate $p = (\hat{p}_1 + \hat{p}_2)/2$. Since the t-ratio $0.04/0.03 = 1.3$ is not statistically significant, it is incorrect to reject the null hypothesis that the two simulations are identical. The difference between the results $\hat{p}_1 = 0.07$ and $\hat{p}_2 = 0.03$ is consistent with random variation.

What should be done? The first mistake was to copy the previous paper's choice of $B = 100$. Instead, suppose you set $B = 5000$. Suppose you now obtain $\hat{p}_2 = 0.04$. Then $\hat{p}_1 - \hat{p}_2 = 0.03$ and a standard error is $\hat{s} = \sqrt{p(1-p)(1/100 + 1/5000)} \simeq 0.02$. Still we cannot reject the hypothesis that the two simulations are different. Even though the estimates (0.07 and 0.04) appear to be quite different, the difficulty is that the original simulation used a very small number of replications ($B = 100$) so the reported estimate is quite imprecise. In this case, it is appropriate to conclude that your results "replicate" the previous study, as there is no statistical evidence to reject the hypothesis that they are equivalent.

Most journals have policies requiring authors to make available their data sets and computer programs required for empirical results. They do not have similar policies regarding simulations. Never-the-less, it is good professional practice to make your simulations available. The best practice is to post your simulation code on your webpage. This invites others to build on and use your results, leading to possible collaboration, citation, and/or advancement.

8.18 Confidence Intervals by Test Inversion

There is a close relationship between hypothesis tests and confidence intervals. We observed in Section 6.13 that the standard 95% asymptotic confidence interval for a parameter θ is

$$\begin{aligned} C_n &= \left[\hat{\theta} - 1.96 \cdot s(\hat{\theta}), \quad \hat{\theta} + 1.96 \cdot s(\hat{\theta}) \right] \\ &= \{ \theta : |t_n(\theta)| \leq 1.96 \}. \end{aligned} \tag{8.24}$$

That is, we can describe C_n as "The point estimate plus or minus 2 standard errors" or "The set of parameter values not rejected by a two-sided t-test." The second definition, known as "test statistic inversion" is a general method for finding confidence intervals, and typically produces confidence intervals with excellent properties.

Given a test statistic $T_n(\theta)$ and critical value c , the acceptance region "Accept if $T_n(\theta) \leq c$ " is identical to the confidence interval $C_n = \{ \theta : T_n(\theta) \leq c \}$. Since the regions are identical, the probability of coverage $\Pr(\theta \in C_n)$ equals the probability of correct acceptance $\Pr(\text{Accept}|\theta)$ which

is exactly 1 minus the Type I error probability. Thus inverting a test with good Type I error probabilities yields a confidence interval with good coverage probabilities.

Now suppose that the parameter of interest $\theta = r(\boldsymbol{\beta})$ is a nonlinear function of the coefficient vector $\boldsymbol{\beta}$. In this case the standard confidence interval for θ is the set C_n as in (8.24) where $\hat{\theta} = r(\hat{\boldsymbol{\beta}})$ is the point estimate and $s(\hat{\theta}) = \sqrt{\hat{\mathbf{R}}' \hat{\mathbf{V}}_{\hat{\boldsymbol{\beta}}} \hat{\mathbf{R}}}$ is the delta method standard error. This confidence interval is inverting the t-test based on the nonlinear hypothesis $r(\boldsymbol{\beta}) = \theta$. The trouble is that in Section 8.16 we learned that there is no unique t-statistic for tests of nonlinear hypotheses and that the choice of parameterization matters greatly.

For example, if $\theta = \beta_1/\beta_2$ then the coverage probability of the standard interval (8.24) is 1 minus the probability of the Type I error, which as shown in Table 8.2 can be far from the nominal 5%.

In this example a good solution is the same as discussed in Section 8.16 – to rewrite the hypothesis as a linear restriction. The hypothesis $\theta = \beta_1/\beta_2$ is the same as $\theta\beta_2 = \beta_1$. The t-statistic for this restriction is

$$t_n(\theta) = \frac{\hat{\beta}_1 - \hat{\beta}_2\theta}{\left(\mathbf{R}' \hat{\mathbf{V}}_{\hat{\boldsymbol{\beta}}} \mathbf{R}\right)^{1/2}}$$

where

$$\mathbf{R} = \begin{pmatrix} 1 \\ -\theta \end{pmatrix}$$

and $\hat{\mathbf{V}}_{\hat{\boldsymbol{\beta}}}$ is the covariance matrix for $(\hat{\beta}_1 \hat{\beta}_2)$. A 95% confidence interval for $\theta = \beta_1/\beta_2$ is the set of values of θ such that $|t_n(\theta)| \leq 1.96$. Since θ appears in both the numerator and denominator, $t_n(\theta)$ is a non-linear function of θ so the easiest method to find the confidence set is by grid search over θ .

For example, in the wage equation

$$\log(\text{Wage}) = \beta_1 \text{Experience} + \beta_2 \text{Experience}^2/100 + \dots$$

the highest expected wage occurs at $\text{Experience} = -50\beta_1/\beta_2$. From Table 4.1 we have the point estimate $\hat{\theta} = 29.8$ and we can calculate the standard error $s(\hat{\theta}) = 0.022$ for a 95% confidence interval [29.8, 29.9]. However, if we instead invert the linear form of the test we can numerically find the interval [29.1, 30.6] which is much larger. From the evidence presented in Section 8.16 we know the first interval can be quite inaccurate and the second interval is greatly preferred.

8.19 Power and Test Consistency

The **power** of a test is the probability of rejecting \mathbb{H}_0 when \mathbb{H}_1 is true.

For simplicity suppose that y_i is i.i.d. $N(\mu, \sigma^2)$ with σ^2 known, consider the t-statistic $t_n(\mu) = \sqrt{n}(\bar{y} - \mu)/\sigma$, and tests of $\mathbb{H}_0 : \mu = 0$ against $\mathbb{H}_1 : \mu > 0$. We reject \mathbb{H}_0 if $t_n = t_n(0) > c$. Note that

$$t_n = t_n(\mu) + \sqrt{n}\mu/\sigma$$

and $t_n(\mu) = Z$ has an exact $N(0, 1)$ distribution. This is because $t_n(\mu)$ is centered at the true mean μ , while the test statistic $t_n(0)$ is centered at the (false) hypothesized mean of 0.

The power of the test is

$$\Pr(t_n > c \mid \theta) = \Pr(Z + \sqrt{n}\mu/\sigma > c) = 1 - \Phi(c - \sqrt{n}\mu/\sigma).$$

This function is monotonically increasing in μ and n , and decreasing in σ and c .

Notice that for any c and $\mu \neq 0$, the power increases to 1 as $n \rightarrow \infty$. This means that for $\theta \in \mathbb{H}_1$, the test will reject \mathbb{H}_0 with probability approaching 1 as the sample size gets large. We call this property **test consistency**.

Definition 8.19.1 A test of $\mathbb{H}_0 : \boldsymbol{\theta} \in \Theta_0$ is *consistent against fixed alternatives* if for all $\boldsymbol{\theta} \in \Theta_1$, $\Pr(\text{Reject } \mathbb{H}_0 \mid \boldsymbol{\theta}) \rightarrow 1$ as $n \rightarrow \infty$.

For tests of the form “Reject \mathbb{H}_0 if $T_n > c$ ”, a sufficient condition for test consistency is that the T_n diverges to positive infinity with probability one for all $\boldsymbol{\theta} \in \Theta_1$.

Definition 8.19.2 $T_n \xrightarrow{p} \infty$ as $n \rightarrow \infty$ if for all $M < \infty$, $\Pr(T_n \leq M) \rightarrow 0$ as $n \rightarrow \infty$. Similarly, $T_n \xrightarrow{p} -\infty$ as $n \rightarrow \infty$ if for all $M < \infty$, $\Pr(T_n \geq -M) \rightarrow 0$ as $n \rightarrow \infty$.

In general, t-test and Wald tests are consistent against fixed alternatives. Take a t-statistic for a test of $\mathbb{H}_0 : \theta = \theta_0$

$$t_n = \frac{\hat{\theta} - \theta_0}{s(\hat{\theta})}$$

where θ_0 is a known value and $s(\hat{\theta}) = \sqrt{n^{-1}\hat{V}_{\hat{\theta}}}$. Note that

$$t_n = \frac{\hat{\theta} - \theta}{s(\hat{\theta})} + \frac{\sqrt{n}(\theta - \theta_0)}{\sqrt{\hat{V}_{\hat{\theta}}}}.$$

The first term on the right-hand-side converges in distribution to $N(0, 1)$. The second term on the right-hand-side equals zero if $\theta = \theta_0$, converges in probability to $+\infty$ if $\theta > \theta_0$, and converges in probability to $-\infty$ if $\theta < \theta_0$. Thus the two-sided t-test is consistent against $\mathbb{H}_1 : \theta \neq \theta_0$, and one-sided t-tests are consistent against the alternatives for which they are designed.

Theorem 8.19.1 Under Assumptions 6.1.2 and 6.10.1, for $\boldsymbol{\theta} = \mathbf{r}(\boldsymbol{\beta}) \neq \boldsymbol{\theta}_0$ and $q = 1$, then $|t_n| \xrightarrow{p} \infty$, so for any $c < \infty$ the test “Reject \mathbb{H}_0 if $|t_n| > c$ ” is consistent against fixed alternatives.

The Wald statistic for $\mathbb{H}_0 : \boldsymbol{\theta} = \mathbf{r}(\boldsymbol{\beta}) = \boldsymbol{\theta}_0$ against $\mathbb{H}_1 : \boldsymbol{\theta} \neq \boldsymbol{\theta}_0$ is

$$W_n = n \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0 \right)' \hat{\mathbf{V}}_{\hat{\boldsymbol{\theta}}}^{-1} \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0 \right).$$

Under \mathbb{H}_1 , $\hat{\boldsymbol{\theta}} \xrightarrow{p} \boldsymbol{\theta} \neq \boldsymbol{\theta}_0$. Thus $\left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0 \right)' \hat{\mathbf{V}}_{\hat{\boldsymbol{\theta}}}^{-1} \left(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0 \right) \xrightarrow{p} (\boldsymbol{\theta} - \boldsymbol{\theta}_0)' \mathbf{V}_{\boldsymbol{\theta}}^{-1} (\boldsymbol{\theta} - \boldsymbol{\theta}_0) > 0$. Hence under \mathbb{H}_1 , $W_n \xrightarrow{p} \infty$. Again, this implies that Wald tests are consistent tests.

Theorem 8.19.2 Under Assumptions 6.1.2 and 6.10.1, for $\boldsymbol{\theta} = \mathbf{r}(\boldsymbol{\beta}) \neq \boldsymbol{\theta}_0$, then $W_n \xrightarrow{p} \infty$, so for any $c < \infty$ the test “Reject \mathbb{H}_0 if $W_n > c$ ” is consistent against fixed alternatives.

8.20 Asymptotic Local Power

Consistency is a good property for a test, but does not give a useful approximation to the power of a test. To approximate the power function we need a distributional approximation.

The standard asymptotic method for power analysis uses what are called **local alternatives**. This is similar to our analysis of restriction estimation under misspecification (Section 7.9). The technique is to index the parameter by sample size so that the asymptotic distribution of the statistic is continuous in a localizing parameter. In this section we consider t-tests on real-valued parameters and in the next section consider Wald tests. Specifically, we consider parameter vectors β_n which are indexed by sample size n and satisfy the real-valued relationship

$$\theta_n = r(\beta_n) = \theta_0 + n^{-1/2}h \quad (8.25)$$

where the scalar h is called a **localizing parameter**. We index β_n and θ_n by sample size to indicate their dependence on n . The way to think of (8.25) is that the true value of the parameters are β_n and θ_n . The parameter θ_n is close to the hypothesized value θ_0 , with deviation $n^{-1/2}h$.

The specification (8.25) states that for any fixed h , θ_n approaches θ_0 as n gets large. Thus θ_n is “close” or “local” to θ_0 . The concept of a localizing sequence (8.25) might seem odd at first as in the actual world the sample size cannot mechanically affect the value of the parameter. Thus (8.25) should not be interpreted literally. Instead, it should be interpreted as a technical device which allows the asymptotic distribution of the test statistic to be continuous in the alternative hypothesis.

To evaluate the asymptotic distribution of the test statistic we start by examining the scaled estimate centered at the hypothesized value θ_0 . Breaking it into a term centered at the true value θ_n and a remainder we find

$$\begin{aligned} \sqrt{n}(\hat{\theta} - \theta_0) &= \sqrt{n}(\hat{\theta} - \theta_n) + \sqrt{n}(\theta_n - \theta_0) \\ &= \sqrt{n}(\hat{\theta} - \theta_n) + h \end{aligned}$$

where the second equality is (8.25). The first term is asymptotically normal:

$$\sqrt{n}(\hat{\theta} - \theta_n) \xrightarrow{d} \sqrt{V_\theta}Z.$$

where $Z \sim N(0, 1)$. Therefore

$$\sqrt{n}(\hat{\theta} - \theta_0) \xrightarrow{d} \sqrt{V_\theta}Z + h$$

or $N(h, V_\theta)$. This is a continuous asymptotic distribution, and depends continuously on the localizing parameter h .

Applied to the t statistic we find

$$\begin{aligned} t_n &= \frac{\hat{\theta} - \theta_0}{s(\hat{\theta})} \\ &\xrightarrow{d} \frac{\sqrt{V_\theta}Z + h}{\sqrt{V_\theta}} \\ &\sim Z + \delta \end{aligned} \quad (8.26)$$

where $\delta = h/\sqrt{V_\theta}$. This generalizes Theorem 8.4.1 (which assumes \mathbb{H}_0 is true) to allow for local alternatives of the form (8.25).

Consider a t-test of \mathbb{H}_0 against the one-sided alternative $\mathbb{H}_1 : \theta > \theta_0$ which rejects \mathbb{H}_0 for $t_n > c_\alpha$ where $\Phi(c_\alpha) = 1 - \alpha$. The **asymptotic local power** of this test is the limit (as the sample size

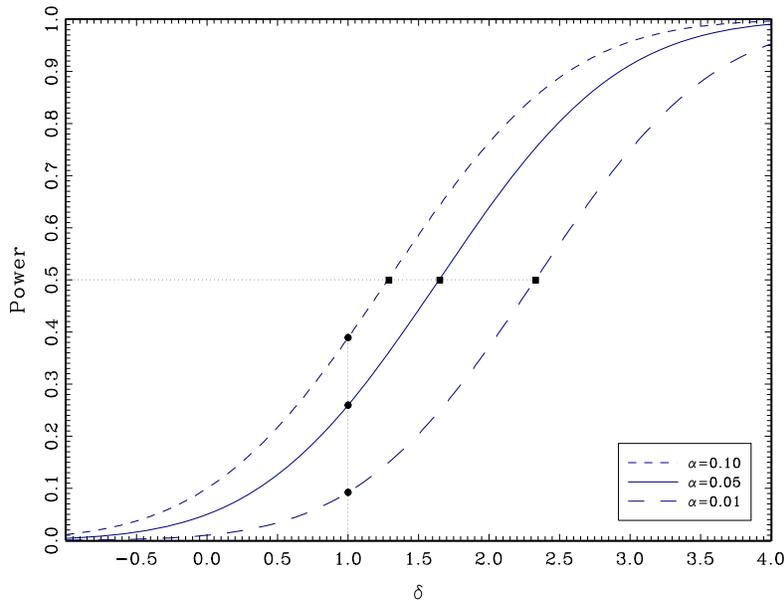


Figure 8.2: Asymptotic Local Power Function of One-Sided t Test

diverges) of the rejection probability under the local alternative (8.25)

$$\begin{aligned}
 \lim_{n \rightarrow \infty} \Pr(\text{Reject } \mathbb{H}_0) &= \lim_{n \rightarrow \infty} \Pr(t_n > c_\alpha) \\
 &= \Pr(Z + \delta > c_\alpha) \\
 &= 1 - \Phi(c_\alpha - \delta) \\
 &= \Phi(\delta - c_\alpha) \\
 &\stackrel{\text{def}}{=} \pi_\alpha(\delta).
 \end{aligned}$$

We call $\pi_\alpha(\delta)$ the **local power function**.

In Figure 8.2 we plot the local power function $\pi_\alpha(\delta)$ as a function of $\delta \in [-1, 4]$ for tests of asymptotic size $\alpha = 0.10$, $\alpha = 0.05$, and $\alpha = 0.01$. $\delta = 0$ corresponds to the null hypothesis so $\pi_\alpha(\delta) = \alpha$. The power functions are monotonically increasing in δ . Note that the power is lower than α for $\delta < 0$ due to the one-sided nature of the test.

We can see that the three power functions are ranked by α so that the test with $\alpha = 0.10$ has higher power than the test with $\alpha = 0.01$. This is the inherent trade-off between size and power. Decreasing size induces a decrease in power, and conversely.

The coefficient δ can be interpreted as the parameter deviation measured as a multiple of the standard error $s(\hat{\theta})$. To see this, recall that $s(\hat{\theta}) = n^{-1/2} \sqrt{\hat{V}_\theta} \simeq n^{-1/2} \sqrt{V_\theta}$ and then note that

$$\delta = \frac{h}{\sqrt{V_\theta}} \simeq \frac{n^{-1/2} h}{s(\hat{\theta})} = \frac{\theta_n - \theta_0}{s(\hat{\theta})}.$$

Thus δ equals the deviation $\theta_n - \theta_0$ expressed as multiples of the standard error $s(\hat{\theta})$. Thus as we examine Figure 8.2, we can interpret the power function at $\delta = 1$ (e.g. 26% for a 5% size test) as the power when the parameter θ_n is one standard error above the hypothesized value. For example, from Table 4.1 the standard error for the coefficient on “Married Female” is 0.008. Thus in this example, $\delta = 1$ corresponds to $\theta_n = 0.008$ or an 0.8% wage premium for married females. Our calculations show that the asymptotic power of a one-sided 5% test against this alternative is about 26%.

The difference between power functions can be measured either vertically or horizontally. For example, in Figure 8.2 there is a vertical dotted line at $\delta = 1$, showing that the asymptotic local power function $\pi_\alpha(\delta)$ equals 39% for $\alpha = 0.10$, equals 26% for $\alpha = 0.05$ and equals 9% for $\alpha = 0.01$. This is the difference in power across tests of differing size, holding fixed the parameter in the alternative

A horizontal comparison can also be illuminating. To illustrate, in Figure 8.2 there is a horizontal dotted line at 50% power. 50% power is a useful benchmark, as it is the point where the test has equal odds of rejection and acceptance. The dotted line crosses the three power curves at $\delta = 1.29$ ($\alpha = 0.10$), $\delta = 1.65$ ($\alpha = 0.05$), and $\delta = 2.33$ ($\alpha = 0.01$). This means that the parameter θ must be at least 1.65 standard errors above the hypothesized value for the one-sided test to have 50% (approximate) power. As discussed above, these values can be interpreted as the multiple of the standard error needed for a coefficient to obtain power equal to 50%.

The ratio of these values (e.g. $1.65/1.29 = 1.28$ for the asymptotic 5% versus 10% tests) measures the relative parameter magnitude needed to achieve the same power. (Thus, for a 5% size test to achieve 50% power, the parameter must be 28% larger than for a 10% size test.) Even more interesting, the square of this ratio (e.g. $(1.65/1.29)^2 = 1.64$) can be interpreted as the increase in sample size needed to achieve the same power under fixed parameters. That is, to achieve 50% power, a 5% size test needs 64% more observations than a 10% size test. This interpretation follows by the following informal argument. By definition and (8.25) $\delta = h/\sqrt{V_\theta} = \sqrt{n}(\theta_n - \theta_0)/\sqrt{V_\theta}$. Thus holding θ and V_θ fixed, we can see that δ^2 is proportional to n .

The analysis of a two-sided t test is similar. (8.26) implies that

$$t_n = \left| \frac{\hat{\theta} - \theta_0}{s(\hat{\theta})} \right| \xrightarrow{d} |Z + \delta|$$

and thus the local power of a two-sided t test is

$$\begin{aligned} \lim_{n \rightarrow \infty} \Pr(\text{Reject } \mathbb{H}_0) &= \lim_{n \rightarrow \infty} \Pr(t_n > c_\alpha) \\ &= \Pr(|Z + \delta| > c_\alpha) \\ &= \Phi(\delta - c_\alpha) - \Phi(-\delta - c_\alpha) \end{aligned}$$

which is monotonically increasing in $|\delta|$.

Theorem 8.20.1 Under Assumptions 6.1.2 and 6.10.1, and $\theta_n = r(\beta_n) = r_0 + n^{-1/2}h$, then

$$t(\theta_0) = \frac{\hat{\theta} - \theta_0}{s(\hat{\theta})} \xrightarrow{d} Z + \delta$$

where $Z \sim N(0, 1)$ and $\delta = h/\sqrt{V_\theta}$. For c_α such that $\Pr(Z > c_\alpha) = \alpha$,

$$\Pr(t(\theta_0) > c_\alpha) \longrightarrow \Phi(\delta - c_\alpha).$$

Furthermore, for c_α such that $\Pr(|Z| > c_\alpha) = \alpha$,

$$\Pr(|t(\theta_0)| > c_\alpha) \longrightarrow \Phi(\delta - c_\alpha) - \Phi(-\delta - c_\alpha).$$

8.21 Asymptotic Local Power, Vector Case

In this section we extend the local power analysis of the previous section to the case of vector-valued alternatives. We generalize (8.25) to allow $\boldsymbol{\theta}_n$ to be vector-valued. The local parameterization takes the form

$$\boldsymbol{\theta}_n = \mathbf{r}(\boldsymbol{\beta}_n) = \boldsymbol{\theta}_0 + n^{-1/2}\mathbf{h} \quad (8.27)$$

where \mathbf{h} is $q \times 1$.

Under (8.27),

$$\begin{aligned} \sqrt{n}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) &= \sqrt{n}(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_n) + \mathbf{h} \\ &\xrightarrow{d} Z_{\mathbf{h}} \sim N(\mathbf{h}, \mathbf{V}_{\boldsymbol{\theta}}), \end{aligned}$$

a normal random vector with mean \mathbf{h} and variance matrix $\mathbf{V}_{\boldsymbol{\theta}}$.

Applied to the Wald statistic we find

$$\begin{aligned} W_n &= n(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)' \widehat{\mathbf{V}}_{\boldsymbol{\theta}}^{-1} (\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \\ &\xrightarrow{d} Z_{\mathbf{h}}' \mathbf{V}_{\boldsymbol{\theta}}^{-1} Z_{\mathbf{h}}. \end{aligned} \quad (8.28)$$

The asymptotic distribution (8.28) is a quadratic in the normal random vector $Z_{\mathbf{h}}$, similar to that found for the asymptotic null distribution of the Wald statistic. The important difference, however, is that $Z_{\mathbf{h}}$ has a mean of \mathbf{h} . The distribution of the quadratic form (8.28) is a close relative of the chi-square distribution.

Theorem 8.21.1 *If $Z_{\mathbf{h}} \sim N(\mathbf{h}, \mathbf{V})$ with $\mathbf{V} > 0$, $q \times q$, then $Z_{\mathbf{h}}' \mathbf{V}^{-1} Z_{\mathbf{h}} \sim \chi_q^2(\lambda)$, a **non-central chi-square random variable** with q degrees of freedom and non-centrality parameter $\lambda = \mathbf{h}' \mathbf{V}^{-1} \mathbf{h}$.*

The convergence (8.28) shows that under the local alternatives (8.27), $W_n \xrightarrow{d} \chi_q^2(\lambda)$. This generalizes the null asymptotic distribution which obtains as the special case $\lambda = 0$. We can use this result to obtain a continuous asymptotic approximation to the power function. For any significance level $\alpha > 0$ set the asymptotic critical value c_{α} so that $\Pr(\chi_q^2 > c_{\alpha}) = \alpha$. Then as $n \rightarrow \infty$,

$$\Pr(W_n > c_{\alpha}) \longrightarrow \Pr(\chi_q^2(\lambda) > c_{\alpha}) \stackrel{def}{=} \pi_{\alpha, q}(\lambda).$$

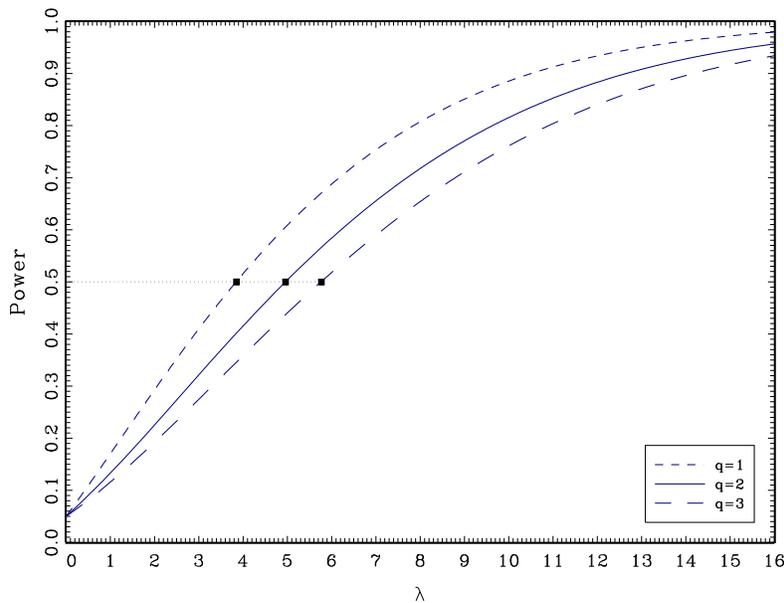
The asymptotic local power function $\pi_{\alpha, q}(\lambda)$ depends only on α , q , and λ .

Theorem 8.21.2 *Under Assumptions 6.1.2 and 6.10.1, and $\boldsymbol{\theta}_n = \mathbf{r}(\boldsymbol{\beta}_n) = \boldsymbol{\theta}_0 + n^{-1/2}\mathbf{h}$, then*

$$W_n \xrightarrow{d} \chi_q^2(\lambda)$$

where $\lambda = \mathbf{h}' \mathbf{V}_{\boldsymbol{\theta}}^{-1} \mathbf{h}$. Furthermore, for c_{α} such that $\Pr(\chi_q^2 > c_{\alpha}) = \alpha$,

$$\Pr(W_n > c_{\alpha}) \longrightarrow \Pr(\chi_q^2(\lambda) > c_{\alpha}).$$

Figure 8.3: Asymptotic Local Power Function, Varying q

The non-central chi-square distribution is a generalization of the chi-square, with $\chi_q^2(\lambda)$ specializing to χ_q^2 when $\lambda = 0$. In the case $q = 1$, $\chi_q^2(\lambda) = |Z + \delta|^2$ with $\lambda = \delta^2$, and thus Theorem 8.21.2 generalizes Theorem 8.20.1 from $q = 1$ to $q \geq 1$.

Figure 8.3 plots $\pi_{0.05,q}(\lambda)$ (the power of asymptotic 5% tests) as a function of λ for $q = 1$, $q = 2$, and $q = 3$. The power functions are monotonically increasing in λ and asymptote to one.

Figure 8.3 also shows the power loss for fixed non-centrality parameter λ as the dimensionality of the test increases. The power curves shift to the right as q increases, resulting in a decrease in power. This is illustrated by the dotted line at 50% power. The dotted line crosses the three power curves at $\lambda = 3.85$ ($q = 1$), $\lambda = 4.96$ ($q = 2$), and $\lambda = 5.77$ ($q = 3$). The ratio of these λ values correspond to the relative sample sizes needed to obtain the same power. Thus increasing the dimension of the test from $q = 1$ to $q = 2$ requires a 28% increase in sample size, or an increase from $q = 1$ to $q = 3$ requires a 50% increase in sample size, to obtain a test with 50% power.

8.22 Technical Proofs*

Proof of Theorem 8.12. The conditions of Theorem 7.10.1 hold, since \mathbb{H}_0 implies Assumption 7.5.1. From (7.54) with $\mathbf{W}_n^{-1} = \widehat{\mathbf{V}}_\beta^{-1}$, we see that

$$\begin{aligned} \sqrt{n} \left(\widehat{\boldsymbol{\beta}} - \widetilde{\boldsymbol{\beta}}_{\text{emd}} \right) &= \widehat{\mathbf{V}}_\beta \widehat{\mathbf{R}} \left(\mathbf{R}_n^{*'} \widehat{\mathbf{V}}_\beta \widehat{\mathbf{R}} \right)^{-1} \mathbf{R}_n^{*'} \sqrt{n} \left(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta} \right) \\ &\xrightarrow{d} \mathbf{V}_\beta \mathbf{R} \left(\mathbf{R}' \mathbf{V}_\beta \mathbf{R} \right)^{-1} \mathbf{R}' \mathbf{N}(\mathbf{0}, \mathbf{V}_\beta) \\ &= \mathbf{V}_\beta \mathbf{R} \mathbf{Z}. \end{aligned}$$

where $\mathbf{Z} \sim \mathbf{N}(\mathbf{0}, (\mathbf{R}' \mathbf{V}_\beta \mathbf{R})^{-1})$. Thus

$$\begin{aligned} J_n^* &= n \left(\widehat{\boldsymbol{\beta}} - \widetilde{\boldsymbol{\beta}}_{\text{emd}} \right)' \widehat{\mathbf{V}}_\beta^{-1} \left(\widehat{\boldsymbol{\beta}} - \widetilde{\boldsymbol{\beta}}_{\text{emd}} \right) \\ &\xrightarrow{d} \mathbf{Z}' \mathbf{R}' \mathbf{V}_\beta \mathbf{V}_\beta^{-1} \mathbf{V}_\beta \mathbf{R} \mathbf{Z} \\ &= \mathbf{Z}' \left(\mathbf{R}' \mathbf{V}_\beta \mathbf{R} \right) \mathbf{Z} \\ &= \chi_q^2. \end{aligned}$$

■

Proof of Theorem 8.21.1. We show that the random variable $Q = Z_h' \mathbf{V}^{-1} Z_h$ depends only on q and λ .

First, let \mathbf{G} be a square root of \mathbf{V} so that $\mathbf{V} = \mathbf{G}\mathbf{G}'$. Define $Z_h^* = \mathbf{G}^{-1} Z_h \sim N(\mathbf{h}^*, \mathbf{I}_q)$ with $\mathbf{h}^* = \mathbf{G}^{-1} \mathbf{h}$. Note that $\lambda = \mathbf{h}' \mathbf{V}^{-1} \mathbf{h} = \mathbf{h}^{*'} \mathbf{h}^*$.

Second, construct an orthogonal $q \times q$ matrix $\mathbf{H} = [\mathbf{H}_1, \mathbf{H}_2]$ whose first column equals $\mathbf{H}_1 = \mathbf{h}^* (\mathbf{h}^{*'} \mathbf{h}^*)^{-1/2}$. Note that $\mathbf{H}_1' \mathbf{h}^* = \lambda^{1/2}$ and $\mathbf{H}_2' \mathbf{h}^* = \mathbf{0}$. Define $Z_h^{**} = \mathbf{H}' Z_h^* \sim N(\mathbf{h}^{**}, \mathbf{I}_q)$ where

$$\mathbf{h}^{**} = \mathbf{H}' \delta^* = \begin{pmatrix} \mathbf{H}_1' \mathbf{h}^* \\ \mathbf{H}_2' \mathbf{h}^* \end{pmatrix} = \begin{pmatrix} \lambda^{1/2} \\ \mathbf{0} \end{pmatrix} \begin{matrix} 1 \\ q-1 \end{matrix}.$$

Note that the distribution of Z_h^{**} is only a function of λ and q .

Finally, observe that

$$\begin{aligned} Q &= Z_h' \mathbf{V}^{-1} Z_h \\ &= Z_h^{*'} Z_h^* \\ &= Z_h^{**'} Z_h^{**} \end{aligned}$$

which is only a function of Z_h^{**} and thus its distribution only depends on λ and q . ■

Exercises

Exercise 8.1 Prove that if an additional regressor \mathbf{X}_{k+1} is added to \mathbf{X} , Theil's adjusted \bar{R}^2 increases if and only if $|t_{k+1}| > 1$, where $t_{k+1} = \hat{\beta}_{k+1}/s(\hat{\beta}_{k+1})$ is the t-ratio for $\hat{\beta}_{k+1}$ and

$$s(\hat{\beta}_{k+1}) = (s^2[(\mathbf{X}'\mathbf{X})^{-1}]_{k+1,k+1})^{1/2}$$

is the homoskedasticity-formula standard error.

Exercise 8.2 You have two independent samples $(\mathbf{y}_1, \mathbf{X}_1)$ and $(\mathbf{y}_2, \mathbf{X}_2)$ which satisfy $\mathbf{y}_1 = \mathbf{X}_1\boldsymbol{\beta}_1 + \mathbf{e}_1$ and $\mathbf{y}_2 = \mathbf{X}_2\boldsymbol{\beta}_2 + \mathbf{e}_2$, where $\mathbb{E}(\mathbf{x}_{1i}e_{1i}) = \mathbf{0}$ and $\mathbb{E}(\mathbf{x}_{2i}e_{2i}) = \mathbf{0}$, and both \mathbf{X}_1 and \mathbf{X}_2 have k columns. Let $\hat{\boldsymbol{\beta}}_1$ and $\hat{\boldsymbol{\beta}}_2$ be the OLS estimates of $\boldsymbol{\beta}_1$ and $\boldsymbol{\beta}_2$. For simplicity, you may assume that both samples have the same number of observations n .

- Find the asymptotic distribution of $\sqrt{n} \left((\hat{\boldsymbol{\beta}}_2 - \hat{\boldsymbol{\beta}}_1) - (\boldsymbol{\beta}_2 - \boldsymbol{\beta}_1) \right)$ as $n \rightarrow \infty$.
- Find an appropriate test statistic for $\mathbb{H}_0 : \boldsymbol{\beta}_2 = \boldsymbol{\beta}_1$.
- Find the asymptotic distribution of this statistic under \mathbb{H}_0 .

Exercise 8.3 The data set `invest.dat` contains data on 565 U.S. firms extracted from Compustat for the year 1987. The variables, in order, are

- I_i Investment to Capital Ratio (multiplied by 100).
- Q_i Total Market Value to Asset Ratio (Tobin's Q).
- C_i Cash Flow to Asset Ratio.
- D_i Long Term Debt to Asset Ratio.

The flow variables are annual sums for 1987. The stock variables are beginning of year.

- Estimate a linear regression of I_i on the other variables. Calculate appropriate standard errors.
- Calculate asymptotic confidence intervals for the coefficients.
- This regression is related to Tobin's q theory of investment, which suggests that investment should be predicted solely by Q_i . Thus the coefficient on Q_i should be positive and the others should be zero. Test the joint hypothesis that the coefficients on C_i and D_i are zero. Test the hypothesis that the coefficient on Q_i is zero. Are the results consistent with the predictions of the theory?
- Now try a non-linear (quadratic) specification. Regress I_i on $Q_i, C_i, D_i, Q_i^2, C_i^2, D_i^2, Q_iC_i, Q_iD_i, C_iD_i$. Test the joint hypothesis that the six interaction and quadratic coefficients are zero.

Exercise 8.4 In a paper in 1963, Marc Nerlove analyzed a cost function for 145 American electric companies. (The problem is discussed in Example 8.3 of Greene, section 1.7 of Hayashi, and the empirical exercise in Chapter 1 of Hayashi). The data file `nerlov.dat` contains his data. The variables are described on page 77 of Hayashi. Nerlov was interested in estimating a *cost function*: $TC = f(Q, PL, PF, PK)$.

- (a) First estimate an unrestricted Cobb-Douglas specification

$$\log TC_i = \beta_1 + \beta_2 \log Q_i + \beta_3 \log PL_i + \beta_4 \log PK_i + \beta_5 \log PF_i + e_i. \quad (8.29)$$

Report parameter estimates and standard errors. You should obtain the same OLS estimates as in Hayashi's equation (1.7.7), but your standard errors may differ.

- (b) What is the economic meaning of the restriction $\mathbb{H}_0 : \beta_3 + \beta_4 + \beta_5 = 1$?
- (c) Estimate (8.29) by constrained least-squares imposing $\beta_3 + \beta_4 + \beta_5 = 1$. Report your parameter estimates and standard errors.
- (d) Estimate (8.29) by efficient minimum distance imposing $\beta_3 + \beta_4 + \beta_5 = 1$. Report your parameter estimates and standard errors.
- (e) Test $\mathbb{H}_0 : \beta_3 + \beta_4 + \beta_5 = 1$ using a Wald statistic
- (f) Test $\mathbb{H}_0 : \beta_3 + \beta_4 + \beta_5 = 1$ using a minimum distance statistic

Chapter 9

Regression Extensions

9.1 NonLinear Least Squares

In some cases we might use a parametric regression function $m(\mathbf{x}, \boldsymbol{\theta}) = \mathbb{E}(y_i | \mathbf{x}_i = \mathbf{x})$ which is a non-linear function of the parameters $\boldsymbol{\theta}$. We describe this setting as **non-linear regression**.

Example 9.1.1 *Exponential Link Regression*

$$m(\mathbf{x}, \boldsymbol{\theta}) = \exp(\mathbf{x}'\boldsymbol{\theta})$$

The exponential link function is strictly positive, so this choice can be useful when it is desired to constrain the mean to be strictly positive.

Example 9.1.2 *Logistic Link Regression*

$$m(\mathbf{x}, \boldsymbol{\theta}) = \Lambda(\mathbf{x}'\boldsymbol{\theta})$$

where

$$\Lambda(u) = (1 + \exp(-u))^{-1} \tag{9.1}$$

is the Logistic distribution function. Since the logistic link function lies in $[0, 1]$, this choice can be useful when the conditional mean is bounded between 0 and 1.

Example 9.1.3 *Exponentially Transformed Regressors*

$$m(x, \boldsymbol{\theta}) = \theta_1 + \theta_2 \exp(\theta_3 x)$$

Example 9.1.4 *Power Transformation*

$$m(x, \boldsymbol{\theta}) = \theta_1 + \theta_2 x^{\theta_3}$$

with $x > 0$.

Example 9.1.5 *Box-Cox Transformed Regressors*

$$m(x, \boldsymbol{\theta}) = \theta_1 + \theta_2 x^{(\theta_3)}$$

where

$$x^{(\lambda)} = \begin{cases} \frac{x^\lambda - 1}{\lambda}, & \text{if } \lambda > 0 \\ \log(x), & \text{if } \lambda = 0 \end{cases} \tag{9.2}$$

and $x > 0$. The function (9.2) is called the Box-Cox Transformation and was introduced by Box and Cox (1964). The function nests linearity ($\lambda = 1$) and logarithmic ($\lambda = 0$) transformations continuously.

Example 9.1.6 *Continuous Threshold Regression*

$$m(x, \boldsymbol{\theta}) = \theta_1 + \theta_2 x + \theta_3 (x - \theta_4) 1(x > \theta_4)$$

Example 9.1.7 *Threshold Regression*

$$m(\mathbf{x}, \boldsymbol{\theta}) = (\boldsymbol{\theta}'_1 \mathbf{x}_1) 1(x_2 < \theta_3) + (\boldsymbol{\theta}'_2 \mathbf{x}_1) 1(x_2 \geq \theta_3)$$

Example 9.1.8 *Smooth Transition*

$$m(\mathbf{x}, \boldsymbol{\theta}) = \boldsymbol{\theta}'_1 \mathbf{x}_1 + (\boldsymbol{\theta}'_2 \mathbf{x}_1) \Lambda\left(\frac{x_2 - \theta_3}{\theta_4}\right)$$

where $\Lambda(u)$ is the logit function (9.1).

What differentiates these examples from the linear regression model is that the conditional mean cannot be written as a linear function of the parameter vector $\boldsymbol{\theta}$.

Nonlinear regression is sometimes adopted because the functional form $m(\mathbf{x}, \boldsymbol{\theta})$ is suggested by an economic model. In other cases, it is adopted as a flexible approximation to an unknown regression function.

The least squares estimator $\hat{\boldsymbol{\theta}}$ minimizes the normalized sum-of-squared-errors

$$S_n(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n (y_i - m(\mathbf{x}_i, \boldsymbol{\theta}))^2.$$

When the regression function is nonlinear, we call this the **nonlinear least squares** (NLLS) estimator. The NLLS residuals are $\hat{e}_i = y_i - m(\mathbf{x}_i, \hat{\boldsymbol{\theta}})$.

One motivation for the choice of NLLS as the estimation method is that the parameter $\boldsymbol{\theta}$ is the solution to the population problem $\min_{\boldsymbol{\theta}} \mathbb{E}(y_i - m(\mathbf{x}_i, \boldsymbol{\theta}))^2$

Since sum-of-squared-errors function $S_n(\boldsymbol{\theta})$ is not quadratic, $\hat{\boldsymbol{\theta}}$ must be found by numerical methods. See Appendix E. When $m(\mathbf{x}, \boldsymbol{\theta})$ is differentiable, then the FOC for minimization are

$$\mathbf{0} = \sum_{i=1}^n \mathbf{m}_{\boldsymbol{\theta}}(\mathbf{x}_i, \hat{\boldsymbol{\theta}}) \hat{e}_i \quad (9.3)$$

where

$$\mathbf{m}_{\boldsymbol{\theta}}(\mathbf{x}, \boldsymbol{\theta}) = \frac{\partial}{\partial \boldsymbol{\theta}} m(\mathbf{x}, \boldsymbol{\theta}).$$

Theorem 9.1.1 Asymptotic Distribution of NLLS Estimator

If the model is identified and $m(\mathbf{x}, \boldsymbol{\theta})$ is differentiable with respect to $\boldsymbol{\theta}$,

$$\sqrt{n}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \xrightarrow{d} N(\mathbf{0}, \mathbf{V}_{\boldsymbol{\theta}})$$

$$\mathbf{V}_{\boldsymbol{\theta}} = (\mathbb{E}(\mathbf{m}_{\boldsymbol{\theta}i} \mathbf{m}'_{\boldsymbol{\theta}i}))^{-1} (\mathbb{E}(\mathbf{m}_{\boldsymbol{\theta}i} \mathbf{m}'_{\boldsymbol{\theta}i} e_i^2)) (\mathbb{E}(\mathbf{m}_{\boldsymbol{\theta}i} \mathbf{m}'_{\boldsymbol{\theta}i}))^{-1}$$

where $\mathbf{m}_{\boldsymbol{\theta}i} = \mathbf{m}_{\boldsymbol{\theta}}(\mathbf{x}_i, \boldsymbol{\theta}_0)$.

Based on Theorem 9.1.1, an estimate of the asymptotic variance \mathbf{V}_θ is

$$\widehat{\mathbf{V}}_\theta = \left(\frac{1}{n} \sum_{i=1}^n \widehat{\mathbf{m}}_{\theta_i} \widehat{\mathbf{m}}'_{\theta_i} \right)^{-1} \left(\frac{1}{n} \sum_{i=1}^n \widehat{\mathbf{m}}_{\theta_i} \widehat{\mathbf{m}}'_{\theta_i} \widehat{e}_i^2 \right) \left(\frac{1}{n} \sum_{i=1}^n \widehat{\mathbf{m}}_{\theta_i} \widehat{\mathbf{m}}'_{\theta_i} \right)^{-1}$$

where $\widehat{\mathbf{m}}_{\theta_i} = \mathbf{m}_\theta(\mathbf{x}_i, \widehat{\boldsymbol{\theta}})$ and $\widehat{e}_i = y_i - m(\mathbf{x}_i, \widehat{\boldsymbol{\theta}})$.

Identification is often tricky in nonlinear regression models. Suppose that

$$m(\mathbf{x}_i, \boldsymbol{\theta}) = \boldsymbol{\beta}'_1 \mathbf{z}_i + \boldsymbol{\beta}'_2 \mathbf{x}_i(\gamma)$$

where $\mathbf{x}_i(\gamma)$ is a function of \mathbf{x}_i and the unknown parameter γ . Examples include $x_i(\gamma) = x_i^\gamma$, $x_i(\gamma) = \exp(\gamma x_i)$, and $x_i(\gamma) = x_i 1(g(x_i) > \gamma)$. The model is linear when $\boldsymbol{\beta}_2 = \mathbf{0}$, and this is often a useful hypothesis (sub-model) to consider. Thus we want to test

$$\mathbb{H}_0 : \boldsymbol{\beta}_2 = \mathbf{0}.$$

However, under \mathbb{H}_0 , the model is

$$y_i = \boldsymbol{\beta}'_1 \mathbf{z}_i + e_i$$

and both $\boldsymbol{\beta}_2$ and γ have dropped out. This means that under \mathbb{H}_0 , γ is not identified. This renders the distribution theory presented in the previous section invalid. Thus when the truth is that $\boldsymbol{\beta}_2 = \mathbf{0}$, the parameter estimates are not asymptotically normally distributed. Furthermore, tests of \mathbb{H}_0 do not have asymptotic normal or chi-square distributions.

The asymptotic theory of such tests have been worked out by Andrews and Ploberger (1994) and B. E. Hansen (1996). In particular, Hansen shows how to use simulation (similar to the bootstrap) to construct the asymptotic critical values (or p-values) in a given application.

Proof of Theorem 9.1.1 (Sketch). NLLS estimation falls in the class of optimization estimators. For this theory, it is useful to denote the true value of the parameter $\boldsymbol{\theta}$ as $\boldsymbol{\theta}_0$.

The first step is to show that $\widehat{\boldsymbol{\theta}} \xrightarrow{p} \boldsymbol{\theta}_0$. Proving that nonlinear estimators are consistent is more challenging than for linear estimators. We sketch the main argument. The idea is that $\widehat{\boldsymbol{\theta}}$ minimizes the sample criterion function $S_n(\boldsymbol{\theta})$, which (for any $\boldsymbol{\theta}$) converges in probability to the mean-squared error function $\mathbb{E}(y_i - m(\mathbf{x}_i, \boldsymbol{\theta}))^2$. Thus it seems reasonable that the minimizer $\widehat{\boldsymbol{\theta}}$ will converge in probability to $\boldsymbol{\theta}_0$, the minimizer of $\mathbb{E}(y_i - m(\mathbf{x}_i, \boldsymbol{\theta}))^2$. It turns out that to show this rigorously, we need to show that $S_n(\boldsymbol{\theta})$ converges *uniformly* to its expectation $\mathbb{E}(y_i - m(\mathbf{x}_i, \boldsymbol{\theta}))^2$, which means that the maximum discrepancy must converge in probability to zero, to exclude the possibility that $S_n(\boldsymbol{\theta})$ is excessively wiggly in $\boldsymbol{\theta}$. Proving uniform convergence is technically challenging, but it can be shown to hold broadly for relevant nonlinear regression models, especially if the regression function $m(\mathbf{x}_i, \boldsymbol{\theta})$ is differentiable in $\boldsymbol{\theta}$. For a complete treatment of the theory of optimization estimators see Newey and McFadden (1994).

Since $\widehat{\boldsymbol{\theta}} \xrightarrow{p} \boldsymbol{\theta}_0$, $\widehat{\boldsymbol{\theta}}$ is close to $\boldsymbol{\theta}_0$ for n large, so the minimization of $S_n(\boldsymbol{\theta})$ only needs to be examined for $\boldsymbol{\theta}$ close to $\boldsymbol{\theta}_0$. Let

$$y_i^0 = e_i + \mathbf{m}'_{\theta_i} \boldsymbol{\theta}_0.$$

For $\boldsymbol{\theta}$ close to the true value $\boldsymbol{\theta}_0$, by a first-order Taylor series approximation,

$$m(\mathbf{x}_i, \boldsymbol{\theta}) \simeq m(\mathbf{x}_i, \boldsymbol{\theta}_0) + \mathbf{m}'_{\theta_i} (\boldsymbol{\theta} - \boldsymbol{\theta}_0).$$

Thus

$$\begin{aligned} y_i - m(\mathbf{x}_i, \boldsymbol{\theta}) &\simeq (e_i + m(\mathbf{x}_i, \boldsymbol{\theta}_0)) - (m(\mathbf{x}_i, \boldsymbol{\theta}_0) + \mathbf{m}'_{\theta_i} (\boldsymbol{\theta} - \boldsymbol{\theta}_0)) \\ &= e_i - \mathbf{m}'_{\theta_i} (\boldsymbol{\theta} - \boldsymbol{\theta}_0) \\ &= y_i^0 - \mathbf{m}'_{\theta_i} \boldsymbol{\theta}. \end{aligned}$$

Hence the sum of squared errors function is

$$S_n(\boldsymbol{\theta}) = \sum_{i=1}^n (y_i - m(\mathbf{x}_i, \boldsymbol{\theta}))^2 \simeq \sum_{i=1}^n (y_i^0 - \mathbf{m}'_{\boldsymbol{\theta}_i} \boldsymbol{\theta})^2$$

and the right-hand-side is the SSE function for a linear regression of y_i^0 on $\mathbf{m}_{\boldsymbol{\theta}_i}$. Thus the NLLS estimator $\hat{\boldsymbol{\theta}}$ has the same asymptotic distribution as the (infeasible) OLS regression of y_i^0 on $\mathbf{m}_{\boldsymbol{\theta}_i}$, which is that stated in the theorem.

9.2 Generalized Least Squares

In the projection model, we know that the least-squares estimator is semi-parametrically efficient for the projection coefficient. However, in the linear regression model

$$\begin{aligned} y_i &= \mathbf{x}'_i \boldsymbol{\beta} + e_i \\ \mathbb{E}(e_i | \mathbf{x}_i) &= 0, \end{aligned}$$

the least-squares estimator is inefficient. The theory of Chamberlain (1987) can be used to show that in this model the semiparametric efficiency bound is obtained by the **Generalized Least Squares** (GLS) estimator (4.13) introduced in Section 4.6.1. The GLS estimator is sometimes called the Aitken estimator. The GLS estimator (9.2) is infeasible since the matrix \mathbf{D} is unknown. A feasible GLS (FGLS) estimator replaces the unknown \mathbf{D} with an estimate $\hat{\mathbf{D}} = \text{diag}\{\hat{\sigma}_1^2, \dots, \hat{\sigma}_n^2\}$. We now discuss this estimation problem.

Suppose that we model the conditional variance using the parametric form

$$\begin{aligned} \sigma_i^2 &= \alpha_0 + \mathbf{z}'_{1i} \boldsymbol{\alpha}_1 \\ &= \boldsymbol{\alpha}' \mathbf{z}_i, \end{aligned}$$

where \mathbf{z}_{1i} is some $q \times 1$ function of \mathbf{x}_i . Typically, \mathbf{z}_{1i} are squares (and perhaps levels) of some (or all) elements of \mathbf{x}_i . Often the functional form is kept simple for parsimony.

Let $\eta_i = e_i^2$. Then

$$\mathbb{E}(\eta_i | \mathbf{x}_i) = \alpha_0 + \mathbf{z}'_{1i} \boldsymbol{\alpha}_1$$

and we have the regression equation

$$\begin{aligned} \eta_i &= \alpha_0 + \mathbf{z}'_{1i} \boldsymbol{\alpha}_1 + \xi_i \\ \mathbb{E}(\xi_i | \mathbf{x}_i) &= 0. \end{aligned} \tag{9.4}$$

This regression error ξ_i is generally heteroskedastic and has the conditional variance

$$\begin{aligned} \text{var}(\xi_i | \mathbf{x}_i) &= \text{var}(e_i^2 | \mathbf{x}_i) \\ &= \mathbb{E}\left((e_i^2 - \mathbb{E}(e_i^2 | \mathbf{x}_i))^2 | \mathbf{x}_i\right) \\ &= \mathbb{E}(e_i^4 | \mathbf{x}_i) - (\mathbb{E}(e_i^2 | \mathbf{x}_i))^2. \end{aligned}$$

Suppose e_i (and thus η_i) were observed. Then we could estimate $\boldsymbol{\alpha}$ by OLS:

$$\hat{\boldsymbol{\alpha}} = (\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'\boldsymbol{\eta} \xrightarrow{p} \boldsymbol{\alpha}$$

and

$$\sqrt{n}(\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}) \xrightarrow{d} \text{N}(\mathbf{0}, \mathbf{V}_{\boldsymbol{\alpha}})$$

where

$$\mathbf{V}_\alpha = (\mathbb{E}(\mathbf{z}_i \mathbf{z}_i'))^{-1} \mathbb{E}(\mathbf{z}_i \mathbf{z}_i' \xi_i^2) (\mathbb{E}(\mathbf{z}_i \mathbf{z}_i'))^{-1}. \quad (9.5)$$

While e_i is not observed, we have the OLS residual $\hat{e}_i = y_i - \mathbf{x}_i' \hat{\boldsymbol{\beta}} = e_i - \mathbf{x}_i' (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})$. Thus

$$\begin{aligned} \phi_i &\equiv \hat{\eta}_i - \eta_i \\ &= \hat{e}_i^2 - e_i^2 \\ &= -2e_i \mathbf{x}_i' (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) + (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})' \mathbf{x}_i \mathbf{x}_i' (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}). \end{aligned}$$

And then

$$\begin{aligned} \frac{1}{\sqrt{n}} \sum_{i=1}^n \mathbf{z}_i \phi_i &= \frac{-2}{n} \sum_{i=1}^n \mathbf{z}_i e_i \mathbf{x}_i' \sqrt{n} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) + \frac{1}{n} \sum_{i=1}^n \mathbf{z}_i (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta})' \mathbf{x}_i \mathbf{x}_i' (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \sqrt{n} \\ &\xrightarrow{p} \mathbf{0} \end{aligned}$$

Let

$$\tilde{\boldsymbol{\alpha}} = (\mathbf{Z}' \mathbf{Z})^{-1} \mathbf{Z}' \hat{\boldsymbol{\eta}} \quad (9.6)$$

be from OLS regression of $\hat{\eta}_i$ on \mathbf{z}_i . Then

$$\begin{aligned} \sqrt{n} (\tilde{\boldsymbol{\alpha}} - \boldsymbol{\alpha}) &= \sqrt{n} (\hat{\boldsymbol{\alpha}} - \boldsymbol{\alpha}) + (n^{-1} \mathbf{Z}' \mathbf{Z})^{-1} n^{-1/2} \mathbf{Z}' \boldsymbol{\phi} \\ &\xrightarrow{d} \mathbf{N}(\mathbf{0}, \mathbf{V}_\alpha) \end{aligned} \quad (9.7)$$

Thus the fact that η_i is replaced with $\hat{\eta}_i$ is asymptotically irrelevant. We call (9.6) the *skedastic* regression, as it is estimating the conditional variance of the regression of y_i on \mathbf{x}_i . We have shown that $\boldsymbol{\alpha}$ is consistently estimated by a simple procedure, and hence we can estimate $\sigma_i^2 = \mathbf{z}_i' \boldsymbol{\alpha}$ by

$$\tilde{\sigma}_i^2 = \tilde{\boldsymbol{\alpha}}' \mathbf{z}_i. \quad (9.8)$$

Suppose that $\tilde{\sigma}_i^2 > 0$ for all i . Then set

$$\tilde{\mathbf{D}} = \text{diag}\{\tilde{\sigma}_1^2, \dots, \tilde{\sigma}_n^2\}$$

and

$$\tilde{\boldsymbol{\beta}} = (\mathbf{X}' \tilde{\mathbf{D}}^{-1} \mathbf{X})^{-1} \mathbf{X}' \tilde{\mathbf{D}}^{-1} \mathbf{y}.$$

This is the feasible GLS, or FGLS, estimator of $\boldsymbol{\beta}$. Since there is not a unique specification for the conditional variance the FGLS estimator is not unique, and will depend on the model (and estimation method) for the skedastic regression.

One typical problem with implementation of FGLS estimation is that in the linear specification (9.4), there is no guarantee that $\tilde{\sigma}_i^2 > 0$ for all i . If $\tilde{\sigma}_i^2 < 0$ for some i , then the FGLS estimator is not well defined. Furthermore, if $\tilde{\sigma}_i^2 \approx 0$ for some i then the FGLS estimator will force the regression equation through the point (y_i, \mathbf{x}_i) , which is undesirable. This suggests that there is a need to bound the estimated variances away from zero. A trimming rule takes the form

$$\bar{\sigma}_i^2 = \max[\tilde{\sigma}_i^2, c\hat{\sigma}^2]$$

for some $c > 0$. For example, setting $c = 1/4$ means that the conditional variance function is constrained to exceed one-fourth of the unconditional variance. As there is no clear method to select c , this introduces a degree of arbitrariness. In this context it is useful to re-estimate the model with several choices for the trimming parameter. If the estimates turn out to be sensitive to its choice, the estimation method should probably be reconsidered.

It is possible to show that if the skedastic regression is correctly specified, then FGLS is asymptotically equivalent to GLS. As the proof is tricky, we just state the result without proof.

Theorem 9.2.1 *If the skedastic regression is correctly specified,*

$$\sqrt{n} \left(\tilde{\boldsymbol{\beta}}_{GLS} - \tilde{\boldsymbol{\beta}}_{FGLS} \right) \xrightarrow{p} \mathbf{0},$$

and thus

$$\sqrt{n} \left(\tilde{\boldsymbol{\beta}}_{FGLS} - \boldsymbol{\beta} \right) \xrightarrow{d} \mathbf{N}(\mathbf{0}, \mathbf{V}_{\boldsymbol{\beta}}),$$

where

$$\mathbf{V}_{\boldsymbol{\beta}} = \left(\mathbb{E} \left(\sigma_i^{-2} \mathbf{x}_i \mathbf{x}_i' \right) \right)^{-1}.$$

Examining the asymptotic distribution of Theorem 9.2.1, the natural estimator of the asymptotic variance of $\tilde{\boldsymbol{\beta}}$ is

$$\tilde{\mathbf{V}}_{\boldsymbol{\beta}}^0 = \left(\frac{1}{n} \sum_{i=1}^n \tilde{\sigma}_i^{-2} \mathbf{x}_i \mathbf{x}_i' \right)^{-1} = \left(\frac{1}{n} \mathbf{X}' \tilde{\mathbf{D}}^{-1} \mathbf{X} \right)^{-1}.$$

which is consistent for $\mathbf{V}_{\boldsymbol{\beta}}$ as $n \rightarrow \infty$. This estimator $\tilde{\mathbf{V}}_{\boldsymbol{\beta}}^0$ is appropriate when the skedastic regression (9.4) is correctly specified.

It may be the case that $\boldsymbol{\alpha}' \mathbf{z}_i$ is only an approximation to the true conditional variance $\sigma_i^2 = \mathbb{E}(e_i^2 | \mathbf{x}_i)$. In this case we interpret $\boldsymbol{\alpha}' \mathbf{z}_i$ as a linear projection of e_i^2 on \mathbf{z}_i . $\tilde{\boldsymbol{\beta}}$ should perhaps be called a quasi-FGLS estimator of $\boldsymbol{\beta}$. Its asymptotic variance is not that given in Theorem 9.2.1. Instead,

$$\mathbf{V}_{\boldsymbol{\beta}} = \left(\mathbb{E} \left((\boldsymbol{\alpha}' \mathbf{z}_i)^{-1} \mathbf{x}_i \mathbf{x}_i' \right) \right)^{-1} \left(\mathbb{E} \left((\boldsymbol{\alpha}' \mathbf{z}_i)^{-2} \sigma_i^2 \mathbf{x}_i \mathbf{x}_i' \right) \right) \left(\mathbb{E} \left((\boldsymbol{\alpha}' \mathbf{z}_i)^{-1} \mathbf{x}_i \mathbf{x}_i' \right) \right)^{-1}.$$

$\mathbf{V}_{\boldsymbol{\beta}}$ takes a sandwich form similar to the covariance matrix of the OLS estimator. Unless $\sigma_i^2 = \boldsymbol{\alpha}' \mathbf{z}_i$, $\tilde{\mathbf{V}}_{\boldsymbol{\beta}}^0$ is inconsistent for $\mathbf{V}_{\boldsymbol{\beta}}$.

An appropriate solution is to use a White-type estimator in place of $\tilde{\mathbf{V}}_{\boldsymbol{\beta}}^0$. This may be written as

$$\begin{aligned} \tilde{\mathbf{V}}_{\boldsymbol{\beta}} &= \left(\frac{1}{n} \sum_{i=1}^n \tilde{\sigma}_i^{-2} \mathbf{x}_i \mathbf{x}_i' \right)^{-1} \left(\frac{1}{n} \sum_{i=1}^n \tilde{\sigma}_i^{-4} \hat{e}_i^2 \mathbf{x}_i \mathbf{x}_i' \right) \left(\frac{1}{n} \sum_{i=1}^n \tilde{\sigma}_i^{-2} \mathbf{x}_i \mathbf{x}_i' \right)^{-1} \\ &= \left(\frac{1}{n} \mathbf{X}' \tilde{\mathbf{D}}^{-1} \mathbf{X} \right)^{-1} \left(\frac{1}{n} \mathbf{X}' \tilde{\mathbf{D}}^{-1} \hat{\mathbf{D}} \tilde{\mathbf{D}}^{-1} \mathbf{X} \right) \left(\frac{1}{n} \mathbf{X}' \tilde{\mathbf{D}}^{-1} \mathbf{X} \right)^{-1} \end{aligned}$$

where $\hat{\mathbf{D}} = \text{diag}\{\hat{e}_1^2, \dots, \hat{e}_n^2\}$. This estimator is robust to misspecification of the conditional variance, and was proposed by Cragg (1992).

In the linear regression model, FGLS is asymptotically superior to OLS. Why then do we not exclusively estimate regression models by FGLS? This is a good question. There are three reasons.

First, FGLS estimation depends on specification and estimation of the skedastic regression. Since the form of the skedastic regression is unknown, and it may be estimated with considerable error, the estimated conditional variances may contain more noise than information about the true conditional variances. In this case, FGLS can do worse than OLS in practice.

Second, individual estimated conditional variances may be negative, and this requires trimming to solve. This introduces an element of arbitrariness which is unsettling to empirical researchers.

Third, and probably most importantly, OLS is a robust estimator of the parameter vector. It is consistent not only in the regression model, but also under the assumptions of linear projection. The GLS and FGLS estimators, on the other hand, require the assumption of a correct conditional mean. If the equation of interest is a linear projection and not a conditional mean, then the OLS

and FGLS estimators will converge in probability to different limits as they will be estimating two different projections. The FGLS probability limit will depend on the particular function selected for the skedastic regression. The point is that the efficiency gains from FGLS are built on the stronger assumption of a correct conditional mean, and the cost is a loss of robustness to misspecification.

9.3 Testing for Heteroskedasticity

The hypothesis of homoskedasticity is that $\mathbb{E}(e_i^2 | \mathbf{x}_i) = \sigma^2$, or equivalently that

$$\mathbb{H}_0 : \boldsymbol{\alpha}_1 = 0$$

in the regression (9.4). We may therefore test this hypothesis by the estimation (9.6) and constructing a Wald statistic. In the classic literature it is typical to impose the stronger assumption that e_i is independent of \mathbf{x}_i , in which case ξ_i is independent of \mathbf{x}_i and the asymptotic variance (9.5) for $\tilde{\boldsymbol{\alpha}}$ simplifies to

$$V_{\boldsymbol{\alpha}} = (\mathbb{E}(\mathbf{z}_i \mathbf{z}_i'))^{-1} \mathbb{E}(\xi_i^2). \quad (9.9)$$

Hence the standard test of \mathbb{H}_0 is a classic F (or Wald) test for exclusion of all regressors from the skedastic regression (9.6). The asymptotic distribution (9.7) and the asymptotic variance (9.9) under independence show that this test has an asymptotic chi-square distribution.

Theorem 9.3.1 *Under \mathbb{H}_0 and e_i independent of \mathbf{x}_i , the Wald test of \mathbb{H}_0 is asymptotically χ_q^2 .*

Most tests for heteroskedasticity take this basic form. The main differences between popular tests are which transformations of \mathbf{x}_i enter \mathbf{z}_i . Motivated by the form of the asymptotic variance of the OLS estimator $\hat{\boldsymbol{\beta}}$, White (1980) proposed that the test for heteroskedasticity be based on setting \mathbf{z}_i to equal all non-redundant elements of \mathbf{x}_i , its squares, and all cross-products. Breusch-Pagan (1979) proposed what might appear to be a distinct test, but the only difference is that they allowed for general choice of \mathbf{z}_i , and replaced $\mathbb{E}(\xi_i^2)$ with $2\sigma^4$ which holds when e_i is $N(\mathbf{0}, \sigma^2)$. If this simplification is replaced by the standard formula (under independence of the error), the two tests coincide.

It is important not to misuse tests for heteroskedasticity. It should not be used to determine whether to estimate a regression equation by OLS or FGLS, nor to determine whether classic or White standard errors should be reported. Hypothesis tests are not designed for these purposes. Rather, tests for heteroskedasticity should be used to answer the scientific question of whether or not the conditional variance is a function of the regressors. If this question is not of economic interest, then there is no value in conducting a test for heteroskedasticity.

9.4 Testing for Omitted NonLinearity

If the goal is to estimate the conditional expectation $\mathbb{E}(y_i | \mathbf{x}_i)$, it is useful to have a general test of the adequacy of the specification.

One simple test for neglected nonlinearity is to add nonlinear functions of the regressors to the regression, and test their significance using a Wald test. Thus, if the model $y_i = \mathbf{x}_i' \boldsymbol{\beta} + \hat{e}_i$ has been fit by OLS, let $\mathbf{z}_i = \mathbf{h}(\mathbf{x}_i)$ denote functions of \mathbf{x}_i which are not linear functions of \mathbf{x}_i (perhaps squares of non-binary regressors) and then fit $y_i = \mathbf{x}_i' \tilde{\boldsymbol{\beta}} + \mathbf{z}_i' \tilde{\boldsymbol{\gamma}} + \tilde{e}_i$ by OLS, and form a Wald statistic for $\boldsymbol{\gamma} = \mathbf{0}$.

Another popular approach is the RESET test proposed by Ramsey (1969). The null model is

$$y_i = \mathbf{x}_i' \boldsymbol{\beta} + e_i$$

which is estimated by OLS, yielding predicted values $\hat{y}_i = \mathbf{x}'_i \hat{\boldsymbol{\beta}}$. Now let

$$\mathbf{z}_i = \begin{pmatrix} \hat{y}_i^2 \\ \vdots \\ \hat{y}_i^m \end{pmatrix}$$

be a $(m - 1)$ -vector of powers of \hat{y}_i . Then run the auxiliary regression

$$y_i = \mathbf{x}'_i \tilde{\boldsymbol{\beta}} + \mathbf{z}'_i \tilde{\boldsymbol{\gamma}} + \tilde{e}_i \quad (9.10)$$

by OLS, and form the Wald statistic W_n for $\boldsymbol{\gamma} = \mathbf{0}$. It is easy (although somewhat tedious) to show that under the null hypothesis, $W_n \xrightarrow{d} \chi_{m-1}^2$. Thus the null is rejected at the $\alpha\%$ level if W_n exceeds the upper $\alpha\%$ tail critical value of the χ_{m-1}^2 distribution.

To implement the test, m must be selected in advance. Typically, small values such as $m = 2, 3$, or 4 seem to work best.

The RESET test appears to work well as a test of functional form against a wide range of smooth alternatives. It is particularly powerful at detecting *single-index* models of the form

$$y_i = G(\mathbf{x}'_i \boldsymbol{\beta}) + e_i$$

where $G(\cdot)$ is a smooth “link” function. To see why this is the case, note that (9.10) may be written as

$$y_i = \mathbf{x}'_i \tilde{\boldsymbol{\beta}} + \left(\mathbf{x}'_i \hat{\boldsymbol{\beta}}\right)^2 \tilde{\gamma}_1 + \left(\mathbf{x}'_i \hat{\boldsymbol{\beta}}\right)^3 \tilde{\gamma}_2 + \cdots + \left(\mathbf{x}'_i \hat{\boldsymbol{\beta}}\right)^m \tilde{\gamma}_{m-1} + \tilde{e}_i$$

which has essentially approximated $G(\cdot)$ by a m 'th order polynomial

9.5 Least Absolute Deviations

We stated that a conventional goal in econometrics is estimation of impact of variation in \mathbf{x}_i on the central tendency of y_i . We have discussed projections and conditional means, but these are not the only measures of central tendency. An alternative good measure is the conditional median.

To recall the definition and properties of the median, let y be a continuous random variable. The median $\theta = \text{med}(y)$ is the value such that $\Pr(y \leq \theta) = \Pr(y \geq \theta) = 0.5$. Two useful facts about the median are that

$$\theta = \underset{\theta}{\text{argmin}} \mathbb{E} |y - \theta| \quad (9.11)$$

and

$$\mathbb{E} \text{sgn}(y - \theta) = 0$$

where

$$\text{sgn}(u) = \begin{cases} 1 & \text{if } u \geq 0 \\ -1 & \text{if } u < 0 \end{cases}$$

is the sign function.

These facts and definitions motivate three estimators of θ . The first definition is the 50th empirical quantile. The second is the value which minimizes $\frac{1}{n} \sum_{i=1}^n |y_i - \theta|$, and the third definition is the solution to the moment equation $\frac{1}{n} \sum_{i=1}^n \text{sgn}(y_i - \theta) = 0$. These distinctions are illusory, however, as these estimators are indeed identical.

Now let's consider the conditional median of y given a random vector \mathbf{x} . Let $m(\mathbf{x}) = \text{med}(y | \mathbf{x})$ denote the conditional median of y given \mathbf{x} . The linear median regression model takes the form

$$y_i = \mathbf{x}'_i \boldsymbol{\beta} + e_i \\ \text{med}(e_i | \mathbf{x}_i) = 0$$

In this model, the linear function $\text{med}(y_i | \mathbf{x}_i = \mathbf{x}) = \mathbf{x}' \boldsymbol{\beta}$ is the conditional median function, and the substantive assumption is that the median function is linear in \mathbf{x} .

Conditional analogs of the facts about the median are

- $\Pr(y_i \leq \mathbf{x}'\boldsymbol{\beta} \mid \mathbf{x}_i = \mathbf{x}) = \Pr(y_i > \mathbf{x}'\boldsymbol{\beta} \mid \mathbf{x}_i = \mathbf{x}) = .5$
- $\mathbb{E}(\text{sgn}(e_i) \mid \mathbf{x}_i) = 0$
- $\mathbb{E}(\mathbf{x}_i \text{sgn}(e_i)) = 0$
- $\boldsymbol{\beta} = \min_{\boldsymbol{\beta}} \mathbb{E}|y_i - \mathbf{x}'_i\boldsymbol{\beta}|$

These facts motivate the following estimator. Let

$$LAD_n(\boldsymbol{\beta}) = \frac{1}{n} \sum_{i=1}^n |y_i - \mathbf{x}'_i\boldsymbol{\beta}|$$

be the average of absolute deviations. The **least absolute deviations** (LAD) estimator of $\boldsymbol{\beta}$ minimizes this function

$$\widehat{\boldsymbol{\beta}} = \underset{\boldsymbol{\beta}}{\text{argmin}} LAD_n(\boldsymbol{\beta})$$

Equivalently, it is a solution to the moment condition

$$\frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \text{sgn}(y_i - \mathbf{x}'_i\widehat{\boldsymbol{\beta}}) = 0. \quad (9.12)$$

The LAD estimator has an asymptotic normal distribution.

Theorem 9.5.1 *Asymptotic Distribution of LAD Estimator*

When the conditional median is linear in \mathbf{x}

$$\sqrt{n}(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \xrightarrow{d} N(\mathbf{0}, \mathbf{V})$$

where

$$\mathbf{V} = \frac{1}{4} (\mathbb{E}(\mathbf{x}_i\mathbf{x}'_i f(0 \mid \mathbf{x}_i)))^{-1} (\mathbb{E}\mathbf{x}_i\mathbf{x}'_i) (\mathbb{E}(\mathbf{x}_i\mathbf{x}'_i f(0 \mid \mathbf{x}_i)))^{-1}$$

and $f(e \mid \mathbf{x})$ is the conditional density of e_i given $\mathbf{x}_i = \mathbf{x}$.

The variance of the asymptotic distribution inversely depends on $f(0 \mid \mathbf{x})$, the conditional density of the error at its median. When $f(0 \mid \mathbf{x})$ is large, then there are many innovations near to the median, and this improves estimation of the median. In the special case where the error is independent of \mathbf{x}_i , then $f(0 \mid \mathbf{x}) = f(0)$ and the asymptotic variance simplifies

$$\mathbf{V} = \frac{(\mathbb{E}\mathbf{x}_i\mathbf{x}'_i)^{-1}}{4f(0)^2} \quad (9.13)$$

This simplification is similar to the simplification of the asymptotic covariance of the OLS estimator under homoskedasticity.

Computation of standard error for LAD estimates typically is based on equation (9.13). The main difficulty is the estimation of $f(0)$, the height of the error density at its median. This can be done with kernel estimation techniques. See Chapter 20. While a complete proof of Theorem 9.5.1 is advanced, we provide a sketch here for completeness.

Proof of Theorem 9.5.1: Similar to NLLS, LAD is an optimization estimator. Let $\boldsymbol{\beta}_0$ denote the true value of $\boldsymbol{\beta}_0$.

The first step is to show that $\hat{\beta} \xrightarrow{p} \beta_0$. The general nature of the proof is similar to that for the NLLS estimator, and is sketched here. For any fixed β , by the WLLN, $LAD_n(\beta) \xrightarrow{p} \mathbb{E} |y_i - \mathbf{x}'_i \beta|$. Furthermore, it can be shown that this convergence is uniform in β . (Proving uniform convergence is more challenging than for the NLLS criterion since the LAD criterion is not differentiable in β .) It follows that $\hat{\beta}$, the minimizer of $LAD_n(\beta)$, converges in probability to β_0 , the minimizer of $\mathbb{E} |y_i - \mathbf{x}'_i \beta|$.

Since $\text{sgn}(a) = 1 - 2 \cdot 1(a \leq 0)$, (9.12) is equivalent to $\bar{\mathbf{g}}_n(\hat{\beta}) = 0$, where $\bar{\mathbf{g}}_n(\beta) = n^{-1} \sum_{i=1}^n \mathbf{g}_i(\beta)$ and $\mathbf{g}_i(\beta) = \mathbf{x}_i (1 - 2 \cdot 1(y_i \leq \mathbf{x}'_i \beta))$. Let $\mathbf{g}(\beta) = \mathbb{E} \mathbf{g}_i(\beta)$. We need three preliminary results. First, by the central limit theorem (Theorem 5.7.1)

$$\sqrt{n}(\bar{\mathbf{g}}_n(\beta_0) - \mathbf{g}(\beta_0)) = -n^{-1/2} \sum_{i=1}^n \mathbf{g}_i(\beta_0) \xrightarrow{d} \mathbf{N}(\mathbf{0}, \mathbb{E} \mathbf{x}_i \mathbf{x}'_i)$$

since $\mathbb{E} \mathbf{g}_i(\beta_0) \mathbf{g}'_i(\beta_0) = \mathbb{E} \mathbf{x}_i \mathbf{x}'_i$. Second using the law of iterated expectations and the chain rule of differentiation,

$$\begin{aligned} \frac{\partial}{\partial \beta'} \mathbf{g}(\beta) &= \frac{\partial}{\partial \beta'} \mathbb{E} \mathbf{x}_i (1 - 2 \cdot 1(y_i \leq \mathbf{x}'_i \beta)) \\ &= -2 \frac{\partial}{\partial \beta'} \mathbb{E} [\mathbf{x}_i \mathbb{E}(1(e_i \leq \mathbf{x}'_i \beta - \mathbf{x}'_i \beta_0) | \mathbf{x}_i)] \\ &= -2 \frac{\partial}{\partial \beta'} \mathbb{E} \left[\mathbf{x}_i \int_{-\infty}^{\mathbf{x}'_i \beta - \mathbf{x}'_i \beta_0} f(e | \mathbf{x}_i) de \right] \\ &= -2 \mathbb{E} [\mathbf{x}_i \mathbf{x}'_i f(\mathbf{x}'_i \beta - \mathbf{x}'_i \beta_0 | \mathbf{x}_i)] \end{aligned}$$

so

$$\frac{\partial}{\partial \beta'} \mathbf{g}(\beta) = -2 \mathbb{E} [\mathbf{x}_i \mathbf{x}'_i f(0 | \mathbf{x}_i)].$$

Third, by a Taylor series expansion and the fact $\mathbf{g}(\beta) = 0$

$$\mathbf{g}(\hat{\beta}) \simeq \frac{\partial}{\partial \beta'} \mathbf{g}(\beta) (\hat{\beta} - \beta).$$

Together

$$\begin{aligned} \sqrt{n}(\hat{\beta} - \beta_0) &\simeq \left(\frac{\partial}{\partial \beta'} \mathbf{g}(\beta_0) \right)^{-1} \sqrt{n} \mathbf{g}(\hat{\beta}) \\ &= (-2 \mathbb{E} [\mathbf{x}_i \mathbf{x}'_i f(0 | \mathbf{x}_i)])^{-1} \sqrt{n} (\mathbf{g}(\hat{\beta}) - \bar{\mathbf{g}}_n(\hat{\beta})) \\ &\simeq \frac{1}{2} (\mathbb{E} [\mathbf{x}_i \mathbf{x}'_i f(0 | \mathbf{x}_i)])^{-1} \sqrt{n} (\bar{\mathbf{g}}_n(\beta_0) - \mathbf{g}(\beta_0)) \\ &\xrightarrow{d} \frac{1}{2} (\mathbb{E} [\mathbf{x}_i \mathbf{x}'_i f(0 | \mathbf{x}_i)])^{-1} \mathbf{N}(\mathbf{0}, \mathbb{E} \mathbf{x}_i \mathbf{x}'_i) \\ &= \mathbf{N}(\mathbf{0}, \mathbf{V}). \end{aligned}$$

The third line follows from an asymptotic empirical process argument and the fact that $\hat{\beta} \xrightarrow{p} \beta_0$.

9.6 Quantile Regression

Quantile regression has become quite popular in recent econometric practice. For $\tau \in [0, 1]$ the τ 'th quantile Q_τ of a random variable with distribution function $F(u)$ is defined as

$$Q_\tau = \inf \{u : F(u) \geq \tau\}$$

When $F(u)$ is continuous and strictly monotonic, then $F(Q_\tau) = \tau$, so you can think of the quantile as the inverse of the distribution function. The quantile Q_τ is the value such that τ (percent) of the mass of the distribution is less than Q_τ . The median is the special case $\tau = .5$.

The following alternative representation is useful. If the random variable U has τ 'th quantile Q_τ , then

$$Q_\tau = \operatorname{argmin}_\theta \mathbb{E} \rho_\tau(U - \theta). \quad (9.14)$$

where $\rho_\tau(q)$ is the piecewise linear function

$$\begin{aligned} \rho_\tau(q) &= \begin{cases} -q(1-\tau) & q < 0 \\ q\tau & q \geq 0 \end{cases} \\ &= q(\tau - 1(q < 0)). \end{aligned} \quad (9.15)$$

This generalizes representation (9.11) for the median to all quantiles.

For the random variables (y_i, \mathbf{x}_i) with conditional distribution function $F(y | \mathbf{x})$ the conditional quantile function $q_\tau(\mathbf{x})$ is

$$Q_\tau(\mathbf{x}) = \inf \{y : F(y | \mathbf{x}) \geq \tau\}.$$

Again, when $F(y | \mathbf{x})$ is continuous and strictly monotonic in y , then $F(Q_\tau(\mathbf{x}) | \mathbf{x}) = \tau$. For fixed τ , the quantile regression function $q_\tau(\mathbf{x})$ describes how the τ 'th quantile of the conditional distribution varies with the regressors.

As functions of \mathbf{x} , the quantile regression functions can take any shape. However for computational convenience it is typical to assume that they are (approximately) linear in \mathbf{x} (after suitable transformations). This linear specification assumes that $Q_\tau(\mathbf{x}) = \boldsymbol{\beta}'_\tau \mathbf{x}$ where the coefficients $\boldsymbol{\beta}_\tau$ vary across the quantiles τ . We then have the linear quantile regression model

$$y_i = \mathbf{x}'_i \boldsymbol{\beta}_\tau + e_i$$

where e_i is the error defined to be the difference between y_i and its τ 'th conditional quantile $\mathbf{x}'_i \boldsymbol{\beta}_\tau$. By construction, the τ 'th conditional quantile of e_i is zero, otherwise its properties are unspecified without further restrictions.

Given the representation (9.14), the quantile regression estimator $\widehat{\boldsymbol{\beta}}_\tau$ for $\boldsymbol{\beta}_\tau$ solves the minimization problem

$$\widehat{\boldsymbol{\beta}}_\tau = \operatorname{argmin}_\boldsymbol{\beta} S_n^\tau(\boldsymbol{\beta})$$

where

$$S_n^\tau(\boldsymbol{\beta}) = \frac{1}{n} \sum_{i=1}^n \rho_\tau(y_i - \mathbf{x}'_i \boldsymbol{\beta})$$

and $\rho_\tau(q)$ is defined in (9.15).

Since the quantile regression criterion function $S_n^\tau(\boldsymbol{\beta})$ does not have an algebraic solution, numerical methods are necessary for its minimization. Furthermore, since it has discontinuous derivatives, conventional Newton-type optimization methods are inappropriate. Fortunately, fast linear programming methods have been developed for this problem, and are widely available.

An asymptotic distribution theory for the quantile regression estimator can be derived using similar arguments as those for the LAD estimator in Theorem 9.5.1.

Theorem 9.6.1 *Asymptotic Distribution of the Quantile Regression Estimator*

When the τ 'th conditional quantile is linear in \mathbf{x}

$$\sqrt{n} \left(\widehat{\boldsymbol{\beta}}_{\tau} - \boldsymbol{\beta}_{\tau} \right) \xrightarrow{d} \mathbf{N}(\mathbf{0}, \mathbf{V}_{\tau}),$$

where

$$\mathbf{V}_{\tau} = \tau(1 - \tau) \left(\mathbb{E}(\mathbf{x}_i \mathbf{x}_i' f(0 | \mathbf{x}_i)) \right)^{-1} \left(\mathbb{E} \mathbf{x}_i \mathbf{x}_i' \right) \left(\mathbb{E}(\mathbf{x}_i \mathbf{x}_i' f(0 | \mathbf{x}_i)) \right)^{-1}$$

and $f(e | \mathbf{x})$ is the conditional density of e_i given $\mathbf{x}_i = \mathbf{x}$.

In general, the asymptotic variance depends on the conditional density of the quantile regression error. When the error e_i is independent of \mathbf{x}_i , then $f(0 | \mathbf{x}_i) = f(0)$, the unconditional density of e_i at 0, and we have the simplification

$$\mathbf{V}_{\tau} = \frac{\tau(1 - \tau)}{f(0)^2} \left(\mathbb{E}(\mathbf{x}_i \mathbf{x}_i') \right)^{-1}.$$

A recent monograph on the details of quantile regression is Koenker (2005).

Exercises

Exercise 9.1 Suppose that $y_i = g(\mathbf{x}_i, \boldsymbol{\theta}) + e_i$ with $\mathbb{E}(e_i | \mathbf{x}_i) = 0$, $\hat{\boldsymbol{\theta}}$ is the NLLS estimator, and $\hat{\mathbf{V}}$ is the estimate of $\text{var}(\hat{\boldsymbol{\theta}})$. You are interested in the conditional mean function $\mathbb{E}(y_i | \mathbf{x}_i = \mathbf{x}) = g(\mathbf{x})$ at some \mathbf{x} . Find an asymptotic 95% confidence interval for $g(\mathbf{x})$.

Exercise 9.2 In Exercise 8.4, you estimated a cost function on a cross-section of electric companies. The equation you estimated was

$$\log TC_i = \beta_1 + \beta_2 \log Q_i + \beta_3 \log PL_i + \beta_4 \log PK_i + \beta_5 \log PF_i + e_i. \quad (9.16)$$

- (a) Following Nerlove, add the variable $(\log Q_i)^2$ to the regression. Do so. Assess the merits of this new specification using a hypothesis test. Do you agree with this modification?
- (b) Now try a non-linear specification. Consider model (9.16) plus the extra term $\beta_6 z_i$, where

$$z_i = \log Q_i (1 + \exp(-(\log Q_i - \beta_7)))^{-1}.$$

In addition, impose the restriction $\beta_3 + \beta_4 + \beta_5 = 1$. This model is called a smooth threshold model. For values of $\log Q_i$ much below β_7 , the variable $\log Q_i$ has a regression slope of β_2 . For values much above β_7 , the regression slope is $\beta_2 + \beta_6$, and the model imposes a smooth transition between these regimes. The model is non-linear because of the parameter β_7 .

The model works best when β_7 is selected so that several values (in this example, at least 10 to 15) of $\log Q_i$ are both below and above β_7 . Examine the data and pick an appropriate range for β_7 .

- (c) Estimate the model by non-linear least squares. I recommend the concentration method: Pick 10 (or more if you like) values of β_7 in this range. For each value of β_7 , calculate z_i and estimate the model by OLS. Record the sum of squared errors, and find the value of β_7 for which the sum of squared errors is minimized.
- (d) Calculate standard errors for all the parameters $(\beta_1, \dots, \beta_7)$.

Exercise 9.3 The data file `cps78.dat` contains 550 observations on 20 variables taken from the May 1978 current population survey. Variables are listed in the file `cps78.pdf`. The goal of the exercise is to estimate a model for the log of earnings (variable `LNWAGE`) as a function of the conditioning variables.

- (a) Start by an OLS regression of `LNWAGE` on the other variables. Report coefficient estimates and standard errors.
- (b) Consider augmenting the model by squares and/or cross-products of the conditioning variables. Estimate your selected model and report the results.
- (c) Are there any variables which seem to be unimportant as a determinant of wages? You may re-estimate the model without these variables, if desired.
- (d) Test whether the error variance is different for men and women. Interpret.
- (e) Test whether the error variance is different for whites and nonwhites. Interpret.
- (f) Construct a model for the conditional variance. Estimate such a model, test for general heteroskedasticity and report the results.
- (g) Using this model for the conditional variance, re-estimate the model from part (c) using FGLS. Report the results.

(h) Do the OLS and FGLS estimates differ greatly? Note any interesting differences.

(i) Compare the estimated standard errors. Note any interesting differences.

Exercise 9.4 For any predictor $g(\mathbf{x}_i)$ for y_i , the mean absolute error (MAE) is

$$\mathbb{E} |y_i - g(\mathbf{x}_i)|.$$

Show that the function $g(\mathbf{x})$ which minimizes the MAE is the conditional median $m(\mathbf{x}) = \text{med}(y_i | \mathbf{x}_i)$.

Exercise 9.5 Define

$$g(u) = \tau - 1(u < 0)$$

where $1(\cdot)$ is the indicator function (takes the value 1 if the argument is true, else equals zero). Let θ satisfy $\mathbb{E}g(y_i - \theta) = 0$. Is θ a quantile of the distribution of y_i ?

Exercise 9.6 Verify equation (9.14).

Chapter 10

The Bootstrap

10.1 Definition of the Bootstrap

Let F denote a distribution function for the population of observations (y_i, \mathbf{x}_i) . Let

$$T_n = T_n((y_1, \mathbf{x}_1), \dots, (y_n, \mathbf{x}_n), F)$$

be a statistic of interest, for example an estimator $\hat{\theta}$ or a t-statistic $(\hat{\theta} - \theta) / s(\hat{\theta})$. Note that we write T_n as possibly a function of F . For example, the t-statistic is a function of the parameter θ which itself is a function of F .

The exact CDF of T_n when the data are sampled from the distribution F is

$$G_n(u, F) = \Pr(T_n \leq u \mid F)$$

In general, $G_n(u, F)$ depends on F , meaning that G changes as F changes.

Ideally, inference would be based on $G_n(u, F)$. This is generally impossible since F is unknown.

Asymptotic inference is based on approximating $G_n(u, F)$ with $G(u, F) = \lim_{n \rightarrow \infty} G_n(u, F)$. When $G(u, F) = G(u)$ does not depend on F , we say that T_n is asymptotically pivotal and use the distribution function $G(u)$ for inferential purposes.

In a seminal contribution, Efron (1979) proposed the bootstrap, which makes a different approximation. The unknown F is replaced by a consistent estimate F_n (one choice is discussed in the next section). Plugged into $G_n(u, F)$ we obtain

$$G_n^*(u) = G_n(u, F_n). \tag{10.1}$$

We call G_n^* the bootstrap distribution. Bootstrap inference is based on $G_n^*(u)$.

Let (y_i^*, \mathbf{x}_i^*) denote random variables with the distribution F_n . A random sample from this distribution is called the bootstrap data. The statistic $T_n^* = T_n((y_1^*, \mathbf{x}_1^*), \dots, (y_n^*, \mathbf{x}_n^*), F_n)$ constructed on this sample is a random variable with distribution G_n^* . That is, $\Pr(T_n^* \leq u) = G_n^*(u)$. We call T_n^* the bootstrap statistic. The distribution of T_n^* is identical to that of T_n when the true CDF is F_n rather than F .

The bootstrap distribution is itself random, as it depends on the sample through the estimator F_n .

In the next sections we describe computation of the bootstrap distribution.

10.2 The Empirical Distribution Function

Recall that $F(y, \mathbf{x}) = \Pr(y_i \leq y, \mathbf{x}_i \leq \mathbf{x}) = \mathbb{E}(1(y_i \leq y) 1(\mathbf{x}_i \leq \mathbf{x}))$, where $1(\cdot)$ is the indicator function. This is a population moment. The method of moments estimator is the corresponding

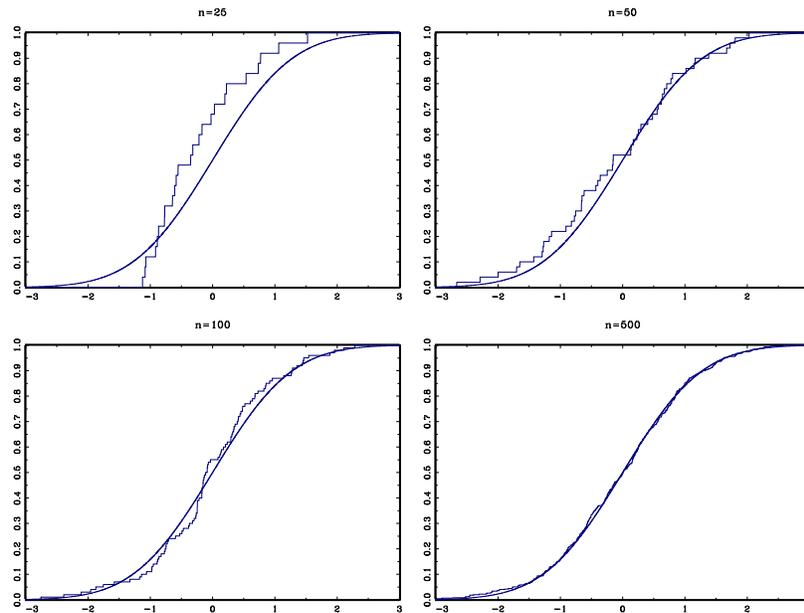


Figure 10.1: Empirical Distribution Functions

sample moment:

$$F_n(y, \mathbf{x}) = \frac{1}{n} \sum_{i=1}^n 1(y_i \leq y) 1(\mathbf{x}_i \leq \mathbf{x}). \quad (10.2)$$

$F_n(y, \mathbf{x})$ is called the empirical distribution function (EDF). F_n is a nonparametric estimate of F . Note that while F may be either discrete or continuous, F_n is by construction a step function.

The EDF is a consistent estimator of the CDF. To see this, note that for any (y, \mathbf{x}) , $1(y_i \leq y) 1(\mathbf{x}_i \leq \mathbf{x})$ is an iid random variable with expectation $F(y, \mathbf{x})$. Thus by the WLLN (Theorem 5.4.2), $F_n(y, \mathbf{x}) \xrightarrow{p} F(y, \mathbf{x})$. Furthermore, by the CLT (Theorem 5.7.1),

$$\sqrt{n}(F_n(y, \mathbf{x}) - F(y, \mathbf{x})) \xrightarrow{d} N(0, F(y, \mathbf{x})(1 - F(y, \mathbf{x}))).$$

To see the effect of sample size on the EDF, in the Figure below, I have plotted the EDF and true CDF for three random samples of size $n = 25, 50, 100$, and 500 . The random draws are from the $N(0, 1)$ distribution. For $n = 25$, the EDF is only a crude approximation to the CDF, but the approximation appears to improve for the large n . In general, as the sample size gets larger, the EDF step function gets uniformly close to the true CDF.

The EDF is a valid discrete probability distribution which puts probability mass $1/n$ at each pair (y_i, \mathbf{x}_i) , $i = 1, \dots, n$. Notationally, it is helpful to think of a random pair (y_i^*, \mathbf{x}_i^*) with the distribution F_n . That is,

$$\Pr(y_i^* \leq y, \mathbf{x}_i^* \leq \mathbf{x}) = F_n(y, \mathbf{x}).$$

We can easily calculate the moments of functions of (y_i^*, \mathbf{x}_i^*) :

$$\begin{aligned} \mathbb{E}h(y_i^*, \mathbf{x}_i^*) &= \int h(y, \mathbf{x}) dF_n(y, \mathbf{x}) \\ &= \sum_{i=1}^n h(y_i, \mathbf{x}_i) \Pr(y_i^* = y_i, \mathbf{x}_i^* = \mathbf{x}_i) \\ &= \frac{1}{n} \sum_{i=1}^n h(y_i, \mathbf{x}_i), \end{aligned}$$

the empirical sample average.

10.3 Nonparametric Bootstrap

The **nonparametric bootstrap** is obtained when the bootstrap distribution (10.1) is defined using the EDF (10.2) as the estimate F_n of F .

Since the EDF F_n is a multinomial (with n support points), in principle the distribution G_n^* could be calculated by direct methods. However, as there are $\binom{2n-1}{n}$ possible samples $\{(y_1^*, \mathbf{x}_1^*), \dots, (y_n^*, \mathbf{x}_n^*)\}$, such a calculation is computationally infeasible. The popular alternative is to use simulation to approximate the distribution. The algorithm is identical to our discussion of Monte Carlo simulation, with the following points of clarification:

- The sample size n used for the simulation is the same as the sample size.
- The random vectors (y_i^*, \mathbf{x}_i^*) are drawn randomly from the empirical distribution. This is equivalent to sampling a pair (y_i, \mathbf{x}_i) randomly from the sample.

The bootstrap statistic $T_n^* = T_n((y_1^*, \mathbf{x}_1^*), \dots, (y_n^*, \mathbf{x}_n^*), F_n)$ is calculated for each bootstrap sample. This is repeated B times. B is known as the number of bootstrap replications. A theory for the determination of the number of bootstrap replications B has been developed by Andrews and Buchinsky (2000). It is desirable for B to be large, so long as the computational costs are reasonable. $B = 1000$ typically suffices.

When the statistic T_n is a function of F , it is typically through dependence on a parameter. For example, the t-ratio $(\hat{\theta} - \theta) / s(\hat{\theta})$ depends on θ . As the bootstrap statistic replaces F with F_n , it similarly replaces θ with θ_n , the value of θ implied by F_n . Typically $\theta_n = \hat{\theta}$, the parameter estimate. (When in doubt use $\hat{\theta}$.)

Sampling from the EDF is particularly easy. Since F_n is a discrete probability distribution putting probability mass $1/n$ at each sample point, sampling from the EDF is equivalent to random sampling a pair (y_i, \mathbf{x}_i) from the observed data **with replacement**. In consequence, a bootstrap sample $\{(y_1^*, \mathbf{x}_1^*), \dots, (y_n^*, \mathbf{x}_n^*)\}$ will necessarily have some ties and multiple values, which is generally not a problem.

10.4 Bootstrap Estimation of Bias and Variance

The bias of $\hat{\theta}$ is $\tau_n = \mathbb{E}(\hat{\theta} - \theta_0)$. Let $T_n(\theta) = \hat{\theta} - \theta$. Then $\tau_n = \mathbb{E}(T_n(\theta_0))$. The bootstrap counterparts are $\hat{\theta}^* = \hat{\theta}((y_1^*, \mathbf{x}_1^*), \dots, (y_n^*, \mathbf{x}_n^*))$ and $T_n^* = \hat{\theta}^* - \theta_n = \hat{\theta}^* - \hat{\theta}$. The bootstrap estimate of τ_n is

$$\tau_n^* = \mathbb{E}(T_n^*).$$

If this is calculated by the simulation described in the previous section, the estimate of τ_n^* is

$$\begin{aligned} \hat{\tau}_n^* &= \frac{1}{B} \sum_{b=1}^B T_{nb}^* \\ &= \frac{1}{B} \sum_{b=1}^B \hat{\theta}_b^* - \hat{\theta} \\ &= \overline{\hat{\theta}^*} - \hat{\theta}. \end{aligned}$$

If $\hat{\theta}$ is biased, it might be desirable to construct a biased-corrected estimator (one with reduced bias). Ideally, this would be

$$\tilde{\theta} = \hat{\theta} - \tau_n,$$

but τ_n is unknown. The (estimated) bootstrap biased-corrected estimator is

$$\begin{aligned}\hat{\theta}^* &= \hat{\theta} - \hat{\tau}_n^* \\ &= \hat{\theta} - (\overline{\hat{\theta}^*} - \hat{\theta}) \\ &= 2\hat{\theta} - \overline{\hat{\theta}^*}.\end{aligned}$$

Note, in particular, that the biased-corrected estimator is *not* $\overline{\hat{\theta}^*}$. Intuitively, the bootstrap makes the following experiment. Suppose that $\hat{\theta}$ is the truth. Then what is the average value of $\hat{\theta}$ calculated from such samples? The answer is $\overline{\hat{\theta}^*}$. If this is lower than $\hat{\theta}$, this suggests that the estimator is *downward-biased*, so a biased-corrected estimator of θ should be *larger* than $\hat{\theta}$, and the best guess is the difference between $\hat{\theta}$ and $\overline{\hat{\theta}^*}$. Similarly if $\overline{\hat{\theta}^*}$ is higher than $\hat{\theta}$, then the estimator is upward-biased and the biased-corrected estimator should be lower than $\hat{\theta}$.

Let $T_n = \hat{\theta}$. The variance of $\hat{\theta}$ is

$$V_n = \mathbb{E}(T_n - \mathbb{E}T_n)^2.$$

Let $T_n^* = \hat{\theta}^*$. It has variance

$$V_n^* = \mathbb{E}(T_n^* - \mathbb{E}T_n^*)^2.$$

The simulation estimate is

$$\hat{V}_n^* = \frac{1}{B} \sum_{b=1}^B (\hat{\theta}_b^* - \overline{\hat{\theta}^*})^2.$$

A bootstrap standard error for $\hat{\theta}$ is the square root of the bootstrap estimate of variance, $s^*(\hat{\theta}) = \sqrt{\hat{V}_n^*}$.

While this standard error may be calculated and reported, it is not clear if it is useful. The primary use of asymptotic standard errors is to construct asymptotic confidence intervals, which are based on the asymptotic normal approximation to the t-ratio. However, the use of the bootstrap presumes that such asymptotic approximations might be poor, in which case the normal approximation is suspected. It appears superior to calculate bootstrap confidence intervals, and we turn to this next.

10.5 Percentile Intervals

For a distribution function $G_n(u, F)$, let $q_n(\alpha, F)$ denote its quantile function. This is the function which solves

$$G_n(q_n(\alpha, F), F) = \alpha.$$

[When $G_n(u, F)$ is discrete, $q_n(\alpha, F)$ may be non-unique, but we will ignore such complications.] Let $q_n(\alpha)$ denote the quantile function of the true sampling distribution, and $q_n^*(\alpha) = q_n(\alpha, F_n)$ denote the quantile function of the bootstrap distribution. Note that this function will change depending on the underlying statistic T_n whose distribution is G_n .

Let $T_n = \hat{\theta}$, an estimate of a parameter of interest. In $(1 - \alpha)\%$ of samples, $\hat{\theta}$ lies in the region $[q_n(\alpha/2), q_n(1 - \alpha/2)]$. This motivates a confidence interval proposed by Efron:

$$C_1 = [q_n^*(\alpha/2), q_n^*(1 - \alpha/2)].$$

This is often called the *percentile confidence interval*.

Computationally, the quantile $q_n^*(\alpha)$ is estimated by $\hat{q}_n^*(\alpha)$, the α 'th sample quantile of the simulated statistics $\{T_{n1}^*, \dots, T_{nB}^*\}$, as discussed in the section on Monte Carlo simulation. The $(1 - \alpha)\%$ Efron percentile interval is then $[\hat{q}_n^*(\alpha/2), \hat{q}_n^*(1 - \alpha/2)]$.

The interval C_1 is a popular bootstrap confidence interval often used in empirical practice. This is because it is easy to compute, simple to motivate, was popularized by Efron early in the history of the bootstrap, and also has the feature that it is translation invariant. That is, if we define $\phi = f(\theta)$ as the parameter of interest for a monotonically increasing function f , then percentile method applied to this problem will produce the confidence interval $[f(q_n^*(\alpha/2)), f(q_n^*(1-\alpha/2))]$, which is a naturally good property.

However, as we show now, C_1 is in a deep sense very poorly motivated.

It will be useful if we introduce an alternative definition of C_1 . Let $T_n(\theta) = \hat{\theta} - \theta$ and let $q_n(\alpha)$ be the quantile function of its distribution. (These are the original quantiles, with θ subtracted.) Then C_1 can alternatively be written as

$$C_1 = [\hat{\theta} + q_n^*(\alpha/2), \hat{\theta} + q_n^*(1-\alpha/2)].$$

This is a bootstrap estimate of the “ideal” confidence interval

$$C_1^0 = [\hat{\theta} + q_n(\alpha/2), \hat{\theta} + q_n(1-\alpha/2)].$$

The latter has coverage probability

$$\begin{aligned} \Pr(\theta_0 \in C_1^0) &= \Pr(\hat{\theta} + q_n(\alpha/2) \leq \theta_0 \leq \hat{\theta} + q_n(1-\alpha/2)) \\ &= \Pr(-q_n(1-\alpha/2) \leq \hat{\theta} - \theta_0 \leq -q_n(\alpha/2)) \\ &= G_n(-q_n(\alpha/2), F_0) - G_n(-q_n(1-\alpha/2), F_0) \end{aligned}$$

which generally is not $1-\alpha$! There is one important exception. If $\hat{\theta} - \theta_0$ has a symmetric distribution about 0, then $G_n(-u, F_0) = 1 - G_n(u, F_0)$, so

$$\begin{aligned} \Pr(\theta_0 \in C_1^0) &= G_n(-q_n(\alpha/2), F_0) - G_n(-q_n(1-\alpha/2), F_0) \\ &= (1 - G_n(q_n(\alpha/2), F_0)) - (1 - G_n(q_n(1-\alpha/2), F_0)) \\ &= \left(1 - \frac{\alpha}{2}\right) - \left(1 - \left(1 - \frac{\alpha}{2}\right)\right) \\ &= 1 - \alpha \end{aligned}$$

and this idealized confidence interval is accurate. Therefore, C_1^0 and C_1 are designed for the case that $\hat{\theta}$ has a symmetric distribution about θ_0 .

When $\hat{\theta}$ does not have a symmetric distribution, C_1 may perform quite poorly.

However, by the translation invariance argument presented above, it also follows that if there exists some monotonically increasing transformation $f(\cdot)$ such that $f(\hat{\theta})$ is symmetrically distributed about $f(\theta_0)$, then the idealized percentile bootstrap method will be accurate.

Based on these arguments, many argue that the percentile interval should not be used unless the sampling distribution is close to unbiased and symmetric.

The problems with the percentile method can be circumvented, at least in principle, by an alternative method.

Let $T_n(\theta) = \hat{\theta} - \theta$. Then

$$\begin{aligned} 1 - \alpha &= \Pr(q_n(\alpha/2) \leq T_n(\theta_0) \leq q_n(1-\alpha/2)) \\ &= \Pr(\hat{\theta} - q_n(1-\alpha/2) \leq \theta_0 \leq \hat{\theta} - q_n(\alpha/2)), \end{aligned}$$

so an exact $(1-\alpha)\%$ confidence interval for θ_0 would be

$$C_2^0 = [\hat{\theta} - q_n(1-\alpha/2), \hat{\theta} - q_n(\alpha/2)].$$

This motivates a bootstrap analog

$$C_2 = [\hat{\theta} - q_n^*(1-\alpha/2), \hat{\theta} - q_n^*(\alpha/2)].$$

Notice that generally this is very different from the Efron interval C_1 ! They coincide in the special case that $G_n^*(u)$ is symmetric about $\hat{\theta}$, but otherwise they differ.

Computationally, this interval can be estimated from a bootstrap simulation by sorting the bootstrap statistics $T_n^* = (\hat{\theta}^* - \hat{\theta})$, which are centered at the sample estimate $\hat{\theta}$. These are sorted to yield the quantile estimates $\hat{q}_n^*(.025)$ and $\hat{q}_n^*(.975)$. The 95% confidence interval is then $[\hat{\theta} - \hat{q}_n^*(.975), \hat{\theta} - \hat{q}_n^*(.025)]$.

This confidence interval is discussed in most theoretical treatments of the bootstrap, but is not widely used in practice.

10.6 Percentile-t Equal-Tailed Interval

Suppose we want to test $\mathbb{H}_0 : \theta = \theta_0$ against $\mathbb{H}_1 : \theta < \theta_0$ at size α . We would set $T_n(\theta) = (\hat{\theta} - \theta) / s(\hat{\theta})$ and reject \mathbb{H}_0 in favor of \mathbb{H}_1 if $T_n(\theta_0) < c$, where c would be selected so that

$$\Pr(T_n(\theta_0) < c) = \alpha.$$

Thus $c = q_n(\alpha)$. Since this is unknown, a bootstrap test replaces $q_n(\alpha)$ with the bootstrap estimate $q_n^*(\alpha)$, and the test rejects if $T_n(\theta_0) < q_n^*(\alpha)$.

Similarly, if the alternative is $\mathbb{H}_1 : \theta > \theta_0$, the bootstrap test rejects if $T_n(\theta_0) > q_n^*(1 - \alpha)$.

Computationally, these critical values can be estimated from a bootstrap simulation by sorting the bootstrap t-statistics $T_n^* = (\hat{\theta}^* - \hat{\theta}) / s(\hat{\theta}^*)$. Note, and this is important, that the bootstrap test statistic is centered at the estimate $\hat{\theta}$, and the standard error $s(\hat{\theta}^*)$ is calculated on the bootstrap sample. These t-statistics are sorted to find the estimated quantiles $\hat{q}_n^*(\alpha)$ and/or $\hat{q}_n^*(1 - \alpha)$.

Let $T_n(\theta) = (\hat{\theta} - \theta) / s(\hat{\theta})$. Then taking the intersection of two one-sided intervals,

$$\begin{aligned} 1 - \alpha &= \Pr(q_n(\alpha/2) \leq T_n(\theta_0) \leq q_n(1 - \alpha/2)) \\ &= \Pr\left(q_n(\alpha/2) \leq (\hat{\theta} - \theta_0) / s(\hat{\theta}) \leq q_n(1 - \alpha/2)\right) \\ &= \Pr\left(\hat{\theta} - s(\hat{\theta})q_n(1 - \alpha/2) \leq \theta_0 \leq \hat{\theta} - s(\hat{\theta})q_n(\alpha/2)\right), \end{aligned}$$

so an exact $(1 - \alpha)\%$ confidence interval for θ_0 would be

$$C_3^0 = [\hat{\theta} - s(\hat{\theta})q_n(1 - \alpha/2), \hat{\theta} - s(\hat{\theta})q_n(\alpha/2)].$$

This motivates a bootstrap analog

$$C_3 = [\hat{\theta} - s(\hat{\theta})q_n^*(1 - \alpha/2), \hat{\theta} - s(\hat{\theta})q_n^*(\alpha/2)].$$

This is often called a *percentile-t confidence interval*. It is *equal-tailed* or *central* since the probability that θ_0 is below the left endpoint approximately equals the probability that θ_0 is above the right endpoint, each $\alpha/2$.

Computationally, this is based on the critical values from the one-sided hypothesis tests, discussed above.

10.7 Symmetric Percentile-t Intervals

Suppose we want to test $\mathbb{H}_0 : \theta = \theta_0$ against $\mathbb{H}_1 : \theta \neq \theta_0$ at size α . We would set $T_n(\theta) = (\hat{\theta} - \theta) / s(\hat{\theta})$ and reject \mathbb{H}_0 in favor of \mathbb{H}_1 if $|T_n(\theta_0)| > c$, where c would be selected so that

$$\Pr(|T_n(\theta_0)| > c) = \alpha.$$

Note that

$$\begin{aligned}\Pr(|T_n(\theta_0)| < c) &= \Pr(-c < T_n(\theta_0) < c) \\ &= G_n(c) - G_n(-c) \\ &\equiv \overline{G}_n(c),\end{aligned}$$

which is a symmetric distribution function. The ideal critical value $c = q_n(\alpha)$ solves the equation

$$\overline{G}_n(q_n(\alpha)) = 1 - \alpha.$$

Equivalently, $q_n(\alpha)$ is the $1 - \alpha$ quantile of the distribution of $|T_n(\theta_0)|$.

The bootstrap estimate is $q_n^*(\alpha)$, the $1 - \alpha$ quantile of the distribution of $|T_n^*|$, or the number which solves the equation

$$\overline{G}_n^*(q_n^*(\alpha)) = G_n^*(q_n^*(\alpha)) - G_n^*(-q_n^*(\alpha)) = 1 - \alpha.$$

Computationally, $q_n^*(\alpha)$ is estimated from a bootstrap simulation by sorting the bootstrap t-statistics $|T_n^*| = |\hat{\theta}^* - \hat{\theta}|/s(\hat{\theta}^*)$, and taking the upper $\alpha\%$ quantile. The bootstrap test rejects if $|T_n(\theta_0)| > q_n^*(\alpha)$.

Let

$$C_4 = [\hat{\theta} - s(\hat{\theta})q_n^*(\alpha), \quad \hat{\theta} + s(\hat{\theta})q_n^*(\alpha)],$$

where $q_n^*(\alpha)$ is the bootstrap critical value for a two-sided hypothesis test. C_4 is called the symmetric percentile-t interval. It is designed to work well since

$$\begin{aligned}\Pr(\theta_0 \in C_4) &= \Pr(\hat{\theta} - s(\hat{\theta})q_n^*(\alpha) \leq \theta_0 \leq \hat{\theta} + s(\hat{\theta})q_n^*(\alpha)) \\ &= \Pr(|T_n(\theta_0)| < q_n^*(\alpha)) \\ &\simeq \Pr(|T_n(\theta_0)| < q_n(\alpha)) \\ &= 1 - \alpha.\end{aligned}$$

If θ is a vector, then to test $\mathbb{H}_0 : \theta = \theta_0$ against $\mathbb{H}_1 : \theta \neq \theta_0$ at size α , we would use a Wald statistic

$$W_n(\theta) = n(\hat{\theta} - \theta)' \hat{V}_\theta^{-1} (\hat{\theta} - \theta)$$

or some other asymptotically chi-square statistic. Thus here $T_n(\theta) = W_n(\theta)$. The ideal test rejects if $W_n \geq q_n(\alpha)$, where $q_n(\alpha)$ is the $(1 - \alpha)\%$ quantile of the distribution of W_n . The bootstrap test rejects if $W_n \geq q_n^*(\alpha)$, where $q_n^*(\alpha)$ is the $(1 - \alpha)\%$ quantile of the distribution of

$$W_n^* = n(\hat{\theta}^* - \hat{\theta})' \hat{V}_\theta^{*-1} (\hat{\theta}^* - \hat{\theta}).$$

Computationally, the critical value $q_n^*(\alpha)$ is found as the quantile from simulated values of W_n^* . Note in the simulation that the Wald statistic is a quadratic form in $(\hat{\theta}^* - \hat{\theta})$, not $(\hat{\theta}^* - \theta_0)$. [This is a typical mistake made by practitioners.]

10.8 Asymptotic Expansions

Let $T_n \in \mathbb{R}$ be a statistic such that

$$T_n \xrightarrow{d} N(0, \sigma^2). \tag{10.3}$$

In some cases, such as when T_n is a t-ratio, then $\sigma^2 = 1$. In other cases σ^2 is unknown. Equivalently, writing $T_n \sim G_n(u, F)$ then for each u and F

$$\lim_{n \rightarrow \infty} G_n(u, F) = \Phi\left(\frac{u}{\sigma}\right),$$

or

$$G_n(u, F) = \Phi\left(\frac{u}{\sigma}\right) + o(1). \quad (10.4)$$

While (10.4) says that G_n converges to $\Phi\left(\frac{u}{\sigma}\right)$ as $n \rightarrow \infty$, it says nothing, however, about the rate of convergence, or the size of the divergence for any particular sample size n . A better asymptotic approximation may be obtained through an *asymptotic expansion*.

The following notation will be helpful. Let a_n be a sequence.

Definition 10.8.1 $a_n = o(1)$ if $a_n \rightarrow 0$ as $n \rightarrow \infty$

Definition 10.8.2 $a_n = O(1)$ if $|a_n|$ is uniformly bounded.

Definition 10.8.3 $a_n = o(n^{-r})$ if $n^r |a_n| \rightarrow 0$ as $n \rightarrow \infty$.

Basically, $a_n = O(n^{-r})$ if it declines to zero like n^{-r} .

We say that a function $g(u)$ is *even* if $g(-u) = g(u)$, and a function $h(u)$ is *odd* if $h(-u) = -h(u)$. The derivative of an even function is odd, and vice-versa.

Theorem 10.8.1 Under regularity conditions and (10.3),

$$G_n(u, F) = \Phi\left(\frac{u}{\sigma}\right) + \frac{1}{n^{1/2}}g_1(u, F) + \frac{1}{n}g_2(u, F) + O(n^{-3/2})$$

uniformly over u , where g_1 is an even function of u , and g_2 is an odd function of u . Moreover, g_1 and g_2 are differentiable functions of u and continuous in F relative to the supremum norm on the space of distribution functions.

The expansion in Theorem 10.8.1 is often called an **Edgeworth expansion**.

We can interpret Theorem 10.8.1 as follows. First, $G_n(u, F)$ converges to the normal limit at rate $n^{1/2}$. To a second order of approximation,

$$G_n(u, F) \approx \Phi\left(\frac{u}{\sigma}\right) + n^{-1/2}g_1(u, F).$$

Since the derivative of g_1 is odd, the density function is skewed. To a third order of approximation,

$$G_n(u, F) \approx \Phi\left(\frac{u}{\sigma}\right) + n^{-1/2}g_1(u, F) + n^{-1}g_2(u, F)$$

which adds a symmetric non-normal component to the approximate density (for example, adding leptokurtosis).

[Side Note: When $T_n = \sqrt{n}(\bar{X}_n - \mu)/\sigma$, a standardized sample mean, then

$$g_1(u) = -\frac{1}{6}\kappa_3(u^2 - 1)\phi(u)$$

$$g_2(u) = -\left(\frac{1}{24}\kappa_4(u^3 - 3u) + \frac{1}{72}\kappa_3^2(u^5 - 10u^3 + 15u)\right)\phi(u)$$

where $\phi(u)$ is the standard normal pdf, and

$$\kappa_3 = \mathbb{E}(X - \mu)^3 / \sigma^3$$

$$\kappa_4 = \mathbb{E}(X - \mu)^4 / \sigma^4 - 3$$

the standardized skewness and excess kurtosis of the distribution of X . Note that when $\kappa_3 = 0$ and $\kappa_4 = 0$, then $g_1 = 0$ and $g_2 = 0$, so the second-order Edgeworth expansion corresponds to the normal distribution.]

Francis Edgeworth

Francis Ysidro Edgeworth (1845-1926) of Ireland, founding editor of the *Economic Journal*, was a profound economic and statistical theorist, developing the theories of indifference curves and asymptotic expansions. He also could be viewed as the first econometrician due to his early use of mathematical statistics in the study of economic data.

10.9 One-Sided Tests

Using the expansion of Theorem 10.8.1, we can assess the accuracy of one-sided hypothesis tests and confidence regions based on an asymptotically normal t-ratio T_n . An asymptotic test is based on $\Phi(u)$.

To the second order, the exact distribution is

$$\Pr(T_n < u) = G_n(u, F_0) = \Phi(u) + \frac{1}{n^{1/2}}g_1(u, F_0) + O(n^{-1})$$

since $\sigma = 1$. The difference is

$$\begin{aligned}\Phi(u) - G_n(u, F_0) &= \frac{1}{n^{1/2}}g_1(u, F_0) + O(n^{-1}) \\ &= O(n^{-1/2}),\end{aligned}$$

so the order of the error is $O(n^{-1/2})$.

A bootstrap test is based on $G_n^*(u)$, which from Theorem 10.8.1 has the expansion

$$G_n^*(u) = G_n(u, F_n) = \Phi(u) + \frac{1}{n^{1/2}}g_1(u, F_n) + O(n^{-1}).$$

Because $\Phi(u)$ appears in both expansions, the difference between the bootstrap distribution and the true distribution is

$$G_n^*(u) - G_n(u, F_0) = \frac{1}{n^{1/2}}(g_1(u, F_n) - g_1(u, F_0)) + O(n^{-1}).$$

Since F_n converges to F at rate \sqrt{n} , and g_1 is continuous with respect to F , the difference $(g_1(u, F_n) - g_1(u, F_0))$ converges to 0 at rate \sqrt{n} . Heuristically,

$$\begin{aligned} g_1(u, F_n) - g_1(u, F_0) &\approx \frac{\partial}{\partial F} g_1(u, F_0) (F_n - F_0) \\ &= O(n^{-1/2}), \end{aligned}$$

The “derivative” $\frac{\partial}{\partial F} g_1(u, F)$ is only heuristic, as F is a function. We conclude that

$$G_n^*(u) - G_n(u, F_0) = O(n^{-1}),$$

or

$$\Pr(T_n^* \leq u) = \Pr(T_n \leq u) + O(n^{-1}),$$

which is an improved rate of convergence over the asymptotic test (which converged at rate $O(n^{-1/2})$). This rate can be used to show that one-tailed bootstrap inference based on the t-ratio achieves a so-called *asymptotic refinement* – the Type I error of the test converges at a faster rate than an analogous asymptotic test.

10.10 Symmetric Two-Sided Tests

If a random variable y has distribution function $H(u) = \Pr(y \leq u)$, then the random variable $|y|$ has distribution function

$$\overline{H}(u) = H(u) - H(-u)$$

since

$$\begin{aligned} \Pr(|y| \leq u) &= \Pr(-u \leq y \leq u) \\ &= \Pr(y \leq u) - \Pr(y \leq -u) \\ &= H(u) - H(-u). \end{aligned}$$

For example, if $Z \sim N(0, 1)$, then $|Z|$ has distribution function

$$\overline{\Phi}(u) = \Phi(u) - \Phi(-u) = 2\Phi(u) - 1.$$

Similarly, if T_n has exact distribution $G_n(u, F)$, then $|T_n|$ has the distribution function

$$\overline{G}_n(u, F) = G_n(u, F) - G_n(-u, F).$$

A two-sided hypothesis test rejects \mathbb{H}_0 for large values of $|T_n|$. Since $T_n \xrightarrow{d} Z$, then $|T_n| \xrightarrow{d} |Z| \sim \overline{\Phi}$. Thus asymptotic critical values are taken from the $\overline{\Phi}$ distribution, and exact critical values are taken from the $\overline{G}_n(u, F_0)$ distribution. From Theorem 10.8.1, we can calculate that

$$\begin{aligned} \overline{G}_n(u, F) &= G_n(u, F) - G_n(-u, F) \\ &= \left(\Phi(u) + \frac{1}{n^{1/2}} g_1(u, F) + \frac{1}{n} g_2(u, F) \right) \\ &\quad - \left(\Phi(-u) + \frac{1}{n^{1/2}} g_1(-u, F) + \frac{1}{n} g_2(-u, F) \right) + O(n^{-3/2}) \\ &= \overline{\Phi}(u) + \frac{2}{n} g_2(u, F) + O(n^{-3/2}), \end{aligned} \tag{10.5}$$

where the simplifications are because g_1 is even and g_2 is odd. Hence the difference between the asymptotic distribution and the exact distribution is

$$\overline{\Phi}(u) - \overline{G}_n(u, F_0) = \frac{2}{n} g_2(u, F_0) + O(n^{-3/2}) = O(n^{-1}).$$

The order of the error is $O(n^{-1})$.

Interestingly, the asymptotic two-sided test has a better coverage rate than the asymptotic one-sided test. This is because the first term in the asymptotic expansion, g_1 , is an even function, meaning that the errors in the two directions exactly cancel out.

Applying (10.5) to the bootstrap distribution, we find

$$\overline{G}_n^*(u) = \overline{G}_n(u, F_n) = \overline{\Phi}(u) + \frac{2}{n}g_2(u, F_n) + O(n^{-3/2}).$$

Thus the difference between the bootstrap and exact distributions is

$$\begin{aligned} \overline{G}_n^*(u) - \overline{G}_n(u, F_0) &= \frac{2}{n}(g_2(u, F_n) - g_2(u, F_0)) + O(n^{-3/2}) \\ &= O(n^{-3/2}), \end{aligned}$$

the last equality because F_n converges to F_0 at rate \sqrt{n} , and g_2 is continuous in F . Another way of writing this is

$$\Pr(|T_n^*| < u) = \Pr(|T_n| < u) + O(n^{-3/2})$$

so the error from using the bootstrap distribution (relative to the true unknown distribution) is $O(n^{-3/2})$. This is in contrast to the use of the asymptotic distribution, whose error is $O(n^{-1})$. Thus a two-sided bootstrap test also achieves an asymptotic refinement, similar to a one-sided test.

A reader might get confused between the two simultaneous effects. Two-sided tests have better rates of convergence than the one-sided tests, and bootstrap tests have better rates of convergence than asymptotic tests.

The analysis shows that there may be a trade-off between one-sided and two-sided tests. Two-sided tests will have more accurate size (Reported Type I error), but one-sided tests might have more power against alternatives of interest. Confidence intervals based on the bootstrap can be asymmetric if based on one-sided tests (equal-tailed intervals) and can therefore be more informative and have smaller length than symmetric intervals. Therefore, the choice between symmetric and equal-tailed confidence intervals is unclear, and needs to be determined on a case-by-case basis.

10.11 Percentile Confidence Intervals

To evaluate the coverage rate of the percentile interval, set $T_n = \sqrt{n}(\hat{\theta} - \theta_0)$. We know that $T_n \xrightarrow{d} N(0, V)$, which is not pivotal, as it depends on the unknown V . Theorem 10.8.1 shows that a first-order approximation

$$G_n(u, F) = \Phi\left(\frac{u}{\sigma}\right) + O(n^{-1/2}),$$

where $\sigma = \sqrt{V}$, and for the bootstrap

$$G_n^*(u) = G_n(u, F_n) = \Phi\left(\frac{u}{\hat{\sigma}}\right) + O(n^{-1/2}),$$

where $\hat{\sigma} = V(F_n)$ is the bootstrap estimate of σ . The difference is

$$\begin{aligned} G_n^*(u) - G_n(u, F_0) &= \Phi\left(\frac{u}{\hat{\sigma}}\right) - \Phi\left(\frac{u}{\sigma}\right) + O(n^{-1/2}) \\ &= -\phi\left(\frac{u}{\sigma}\right) \frac{u}{\sigma} (\hat{\sigma} - \sigma) + O(n^{-1/2}) \\ &= O(n^{-1/2}) \end{aligned}$$

Hence the order of the error is $O(n^{-1/2})$.

The good news is that the percentile-type methods (if appropriately used) can yield \sqrt{n} -convergent asymptotic inference. Yet these methods do not require the calculation of standard

errors! This means that in contexts where standard errors are not available or are difficult to calculate, the percentile bootstrap methods provide an attractive inference method.

The bad news is that the rate of convergence is disappointing. It is no better than the rate obtained from an asymptotic one-sided confidence region. Therefore if standard errors are available, it is unclear if there are any benefits from using the percentile bootstrap over simple asymptotic methods.

Based on these arguments, the theoretical literature (e.g. Hall, 1992, Horowitz, 2001) tends to advocate the use of the percentile-t bootstrap methods rather than percentile methods.

10.12 Bootstrap Methods for Regression Models

The bootstrap methods we have discussed have set $G_n^*(u) = G_n(u, F_n)$, where F_n is the EDF. Any other consistent estimate of F may be used to define a feasible bootstrap estimator. The advantage of the EDF is that it is fully nonparametric, it imposes no conditions, and works in nearly any context. But since it is fully nonparametric, it may be inefficient in contexts where more is known about F . We discuss bootstrap methods appropriate for the linear regression model

$$y_i = \mathbf{x}_i' \boldsymbol{\beta} + e_i$$

$$\mathbb{E}(e_i | \mathbf{x}_i) = 0.$$

The non-parametric bootstrap resamples the observations (y_i^*, \mathbf{x}_i^*) from the EDF, which implies

$$y_i^* = \mathbf{x}_i^{*'} \hat{\boldsymbol{\beta}} + e_i^*$$

$$\mathbb{E}(\mathbf{x}_i^* e_i^*) = \mathbf{0}$$

but generally

$$\mathbb{E}(e_i^* | \mathbf{x}_i^*) \neq 0.$$

The bootstrap distribution does not impose the regression assumption, and is thus an inefficient estimator of the true distribution (when in fact the regression assumption is true.)

One approach to this problem is to impose the very strong assumption that the error ε_i is independent of the regressor \mathbf{x}_i . The advantage is that in this case it is straightforward to construct bootstrap distributions. The disadvantage is that the bootstrap distribution may be a poor approximation when the error is not independent of the regressors.

To impose independence, it is sufficient to sample the \mathbf{x}_i^* and e_i^* independently, and then create $y_i^* = \mathbf{x}_i^{*'} \hat{\boldsymbol{\beta}} + e_i^*$. There are different ways to impose independence. A non-parametric method is to sample the bootstrap errors e_i^* randomly from the OLS residuals $\{\hat{e}_1, \dots, \hat{e}_n\}$. A parametric method is to generate the bootstrap errors e_i^* from a parametric distribution, such as the normal $e_i^* \sim N(0, \hat{\sigma}^2)$.

For the regressors \mathbf{x}_i^* , a nonparametric method is to sample the \mathbf{x}_i^* randomly from the EDF or sample values $\{\mathbf{x}_1, \dots, \mathbf{x}_n\}$. A parametric method is to sample \mathbf{x}_i^* from an estimated parametric distribution. A third approach sets $\mathbf{x}_i^* = \mathbf{x}_i$. This is equivalent to treating the regressors as *fixed in repeated samples*. If this is done, then all inferential statements are made conditionally on the observed values of the regressors, which is a valid statistical approach. It does not really matter, however, whether or not the \mathbf{x}_i are really “fixed” or random.

The methods discussed above are unattractive for most applications in econometrics because they impose the stringent assumption that \mathbf{x}_i and e_i are independent. Typically what is desirable is to impose only the regression condition $\mathbb{E}(e_i | \mathbf{x}_i) = 0$. Unfortunately this is a harder problem.

One proposal which imposes the regression condition without independence is the *Wild Bootstrap*. The idea is to construct a conditional distribution for e_i^* so that

$$\mathbb{E}(e_i^* | \mathbf{x}_i) = 0$$

$$\mathbb{E}(e_i^{*2} | \mathbf{x}_i) = \hat{e}_i^2$$

$$\mathbb{E}(e_i^{*3} | \mathbf{x}_i) = \hat{e}_i^3.$$

A conditional distribution with these features will preserve the main important features of the data. This can be achieved using a two-point distribution of the form

$$\Pr \left(e_i^* = \left(\frac{1 + \sqrt{5}}{2} \right) \hat{e}_i \right) = \frac{\sqrt{5} - 1}{2\sqrt{5}}$$
$$\Pr \left(e_i^* = \left(\frac{1 - \sqrt{5}}{2} \right) \hat{e}_i \right) = \frac{\sqrt{5} + 1}{2\sqrt{5}}$$

For each \mathbf{x}_i , you sample e_i^* using this two-point distribution.

Exercises

Exercise 10.1 Let $F_n(\mathbf{x})$ denote the EDF of a random sample. Show that

$$\sqrt{n}(F_n(\mathbf{x}) - F_0(\mathbf{x})) \xrightarrow{d} N(0, F_0(\mathbf{x})(1 - F_0(\mathbf{x}))).$$

Exercise 10.2 Take a random sample $\{y_1, \dots, y_n\}$ with $\mu = \mathbb{E}y_i$ and $\sigma^2 = \text{var}(y_i)$. Let the statistic of interest be the sample mean $T_n = \bar{y}_n$. Find the population moments $\mathbb{E}T_n$ and $\text{var}(T_n)$. Let $\{y_1^*, \dots, y_n^*\}$ be a random sample from the empirical distribution function and let $T_n^* = \bar{y}_n^*$ be its sample mean. Find the bootstrap moments $\mathbb{E}T_n^*$ and $\text{var}(T_n^*)$.

Exercise 10.3 Consider the following bootstrap procedure for a regression of y_i on \mathbf{x}_i . Let $\hat{\beta}$ denote the OLS estimator from the regression of \mathbf{y} on \mathbf{X} , and $\hat{\mathbf{e}} = \mathbf{y} - \mathbf{X}\hat{\beta}$ the OLS residuals.

(a) Draw a random vector (\mathbf{x}^*, e^*) from the pair $\{(\mathbf{x}_i, \hat{e}_i) : i = 1, \dots, n\}$. That is, draw a random integer i' from $[1, 2, \dots, n]$, and set $\mathbf{x}^* = \mathbf{x}_{i'}$ and $e^* = \hat{e}_{i'}$. Set $y^* = \mathbf{x}^{*\prime}\hat{\beta} + e^*$. Draw (with replacement) n such vectors, creating a random bootstrap data set $(\mathbf{y}^*, \mathbf{X}^*)$.

(b) Regress \mathbf{y}^* on \mathbf{X}^* , yielding OLS estimates $\hat{\beta}^*$ and any other statistic of interest.

Show that this bootstrap procedure is (numerically) *identical* to the non-parametric bootstrap.

Exercise 10.4 Consider the following bootstrap procedure. Using the non-parametric bootstrap, generate bootstrap samples, calculate the estimate $\hat{\theta}^*$ on these samples and then calculate

$$T_n^* = (\hat{\theta}^* - \hat{\theta})/s(\hat{\theta}),$$

where $s(\hat{\theta})$ is the standard error in the original data. Let $q_n^*(.05)$ and $q_n^*(.95)$ denote the 5% and 95% quantiles of T_n^* , and define the bootstrap confidence interval

$$C = \left[\hat{\theta} - s(\hat{\theta})q_n^*(.95), \quad \hat{\theta} - s(\hat{\theta})q_n^*(.05) \right].$$

Show that C exactly equals the Alternative percentile interval (not the percentile-t interval).

Exercise 10.5 You want to test $\mathbb{H}_0 : \theta = 0$ against $\mathbb{H}_1 : \theta > 0$. The test for \mathbb{H}_0 is to reject if $T_n = \hat{\theta}/s(\hat{\theta}) > c$ where c is picked so that Type I error is α . You do this as follows. Using the non-parametric bootstrap, you generate bootstrap samples, calculate the estimates $\hat{\theta}^*$ on these samples and then calculate

$$T_n^* = \hat{\theta}^*/s(\hat{\theta}^*).$$

Let $q_n^*(.95)$ denote the 95% quantile of T_n^* . You replace c with $q_n^*(.95)$, and thus reject \mathbb{H}_0 if $T_n = \hat{\theta}/s(\hat{\theta}) > q_n^*(.95)$. What is wrong with this procedure?

Exercise 10.6 Suppose that in an application, $\hat{\theta} = 1.2$ and $s(\hat{\theta}) = .2$. Using the non-parametric bootstrap, 1000 samples are generated from the bootstrap distribution, and $\hat{\theta}^*$ is calculated on each sample. The $\hat{\theta}^*$ are sorted, and the 2.5% and 97.5% quantiles of the $\hat{\theta}^*$ are .75 and 1.3, respectively.

(a) Report the 95% Efron Percentile interval for θ .

(b) Report the 95% Alternative Percentile interval for θ .

(c) With the given information, can you report the 95% Percentile-t interval for θ ?

Exercise 10.7 The datafile `hprice1.dat` contains data on house prices (sales), with variables listed in the file `hprice1.pdf`. Estimate a linear regression of price on the number of bedrooms, lot size, size of house, and the colonial dummy. Calculate 95% confidence intervals for the regression coefficients using both the asymptotic normal approximation and the percentile-t bootstrap.

Chapter 11

NonParametric Regression

11.1 Introduction

When components of \mathbf{x} are continuously distributed then the conditional expectation function

$$\mathbb{E}(y_i \mid \mathbf{x}_i = \mathbf{x}) = m(\mathbf{x})$$

can take any nonlinear shape. Unless an economic model restricts the form of $m(\mathbf{x})$ to a parametric function, the CEF is inherently **nonparametric**, meaning that the function $m(\mathbf{x})$ is an element of an infinite dimensional class. In this situation, how can we estimate $m(\mathbf{x})$? What is a suitable method, if we acknowledge that $m(\mathbf{x})$ is nonparametric?

There are two main classes of nonparametric regression estimators: kernel estimators, and series estimators. In this chapter we introduce kernel methods.

To get started, suppose that there is a single real-valued regressor x_i . We consider the case of vector-valued regressors later.

11.2 Binned Estimator

For clarity, fix the point x and consider estimation of the single point $m(x)$. This is the mean of y_i for random pairs (y_i, x_i) such that $x_i = x$. If the distribution of x_i were discrete then we could estimate $m(x)$ by taking the average of the sub-sample of observations y_i for which $x_i = x$. But when x_i is continuous then the probability is zero that x_i exactly equals any specific x . So there is no sub-sample of observations with $x_i = x$ and we cannot simply take the average of the corresponding y_i values. However, if the CEF $m(x)$ is continuous, then it should be possible to get a good approximation by taking the average of the observations for which x_i is *close* to x , perhaps for the observations for which $|x_i - x| \leq h$ for some small $h > 0$. Later we will call h a **bandwidth**. This estimator can be written as

$$\hat{m}(x) = \frac{\sum_{i=1}^n 1(|x_i - x| \leq h) y_i}{\sum_{i=1}^n 1(|x_i - x| \leq h)} \quad (11.1)$$

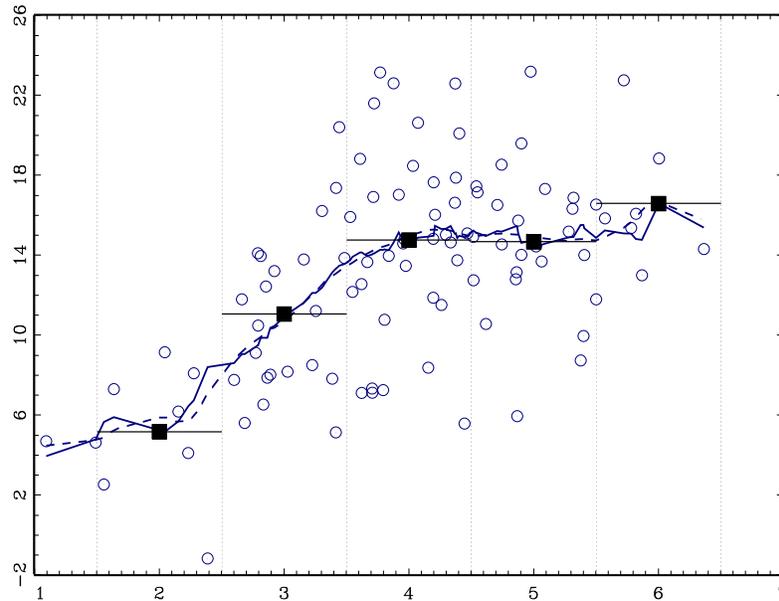
where $1(\cdot)$ is the indicator function. Alternatively, (11.1) can be written as

$$\hat{m}(x) = \sum_{i=1}^n w_i(x) y_i \quad (11.2)$$

where

$$w_i(x) = \frac{1(|x_i - x| \leq h)}{\sum_{j=1}^n 1(|x_j - x| \leq h)}.$$

Notice that $\sum_{i=1}^n w_i(x) = 1$, so (11.2) is a weighted average of the y_i .

Figure 11.1: Scatter of (y_i, x_i) and Nadaraya-Watson regression

It is possible that for some values of x there are no values of x_i such that $|x_i - x| \leq h$, which implies that $\sum_{i=1}^n 1(|x_i - x| \leq h) = 0$. In this case the estimator (11.1) is undefined for those values of x .

To visualize, Figure 11.1 displays a scatter plot of 100 observations on a random pair (y_i, x_i) generated by simulation¹. (The observations are displayed as the open circles.) The estimator (11.1) of the CEF $m(x)$ at $x = 2$ with $h = 1/2$ is the average of the y_i for the observations such that x_i falls in the interval $[1.5, 2.5]$. (Our choice of $h = 1/2$ is somewhat arbitrary. Selection of h will be discussed later.) The estimate is $\hat{m}(2) = 5.16$ and is shown on Figure 11.1 by the first solid square. We repeat the calculation (11.1) for $x = 3, 4, 5$, and 6 , which is equivalent to partitioning the support of x_i into the regions $[1.5, 2.5]$, $[2.5, 3.5]$, $[3.5, 4.5]$, $[4.5, 5.5]$, and $[5.5, 6.5]$. These partitions are shown in Figure 11.1 by the vertical dotted lines, and the estimates (11.1) by the solid squares.

These estimates $\hat{m}(x)$ can be viewed as estimates of the CEF $m(x)$. Sometimes called a binned estimator, this is a step-function approximation to $m(x)$ and is displayed in Figure 11.1 by the horizontal lines passing through the solid squares. This estimate roughly tracks the central tendency of the scatter of the observations (y_i, x_i) . However, the huge jumps in the estimated step function at the edges of the partitions are disconcerting, counter-intuitive, and clearly an artifact of the discrete binning.

If we take another look at the estimation formula (11.1) there is no reason why we need to evaluate (11.1) only on a coarse grid. We can evaluate $\hat{m}(x)$ for any set of values of x . In particular, we can evaluate (11.1) on a fine grid of values of x and thereby obtain a smoother estimate of the CEF. This estimator with $h = 1/2$ is displayed in Figure 11.1 with the solid line. This is a generalization of the binned estimator and by construction passes through the solid squares.

The bandwidth h determines the degree of smoothing. Larger values of h increase the width of the bins in Figure 11.1, thereby increasing the smoothness of the estimate $\hat{m}(x)$ as a function of x . Smaller values of h decrease the width of the bins, resulting in less smooth conditional mean estimates.

¹The distribution is $x_i \sim N(4, 1)$ and $y_i | x_i \sim N(m(x_i), 16)$ with $m(x) = 10 \log(x)$.

11.3 Kernel Regression

One deficiency with the estimator (11.1) is that it is a step function in x , as it is discontinuous at each observation $x = x_i$. That is why its plot in Figure 11.1 is jagged. The source of the discontinuity is that the weights $w_i(x)$ are constructed from indicator functions, which are themselves discontinuous. If instead the weights are constructed from continuous functions then the CEF estimator will also be continuous in x .

To generalize (11.1) it is useful to write the weights $1(|x_i - x| \leq h)$ in terms of the uniform density function on $[-1, 1]$

$$k_0(u) = \frac{1}{2}1(|u| \leq 1).$$

Then

$$1(|x_i - x| \leq h) = 1\left(\left|\frac{x_i - x}{h}\right| \leq 1\right) = 2k_0\left(\frac{x_i - x}{h}\right).$$

and (11.1) can be written as

$$\hat{m}(x) = \frac{\sum_{i=1}^n k_0\left(\frac{x_i - x}{h}\right) y_i}{\sum_{i=1}^n k_0\left(\frac{x_i - x}{h}\right)}. \quad (11.3)$$

The uniform density $k_0(u)$ is a special case of what is known as a **kernel function**.

Definition 11.3.1 A second-order **kernel function** $k(u)$ satisfies $0 \leq k(u) < \infty$, $k(u) = k(-u)$, $\int_{-\infty}^{\infty} k(u) du = 1$ and $\sigma_k^2 = \int_{-\infty}^{\infty} u^2 k(u) du < \infty$.

Essentially, a kernel function is a probability density function which is bounded and symmetric about zero. A generalization of (11.1) is obtained by replacing the uniform kernel with any other kernel function:

$$\hat{m}(x) = \frac{\sum_{i=1}^n k\left(\frac{x_i - x}{h}\right) y_i}{\sum_{i=1}^n k\left(\frac{x_i - x}{h}\right)}. \quad (11.4)$$

The estimator (11.4) also takes the form (11.2) with

$$w_i(x) = \frac{k\left(\frac{x_i - x}{h}\right)}{\sum_{j=1}^n k\left(\frac{x_j - x}{h}\right)}.$$

The estimator (11.4) is known as the **Nadaraya-Watson** estimator, the **kernel regression** estimator, or the **local constant** estimator.

The bandwidth h plays the same role in (11.4) as it does in (11.1). Namely, larger values of h will result in estimates $\hat{m}(x)$ which are smoother in x , and smaller values of h will result in estimates which are more erratic. It might be helpful to consider the two extreme cases $h \rightarrow 0$ and $h \rightarrow \infty$. As $h \rightarrow 0$ we can see that $\hat{m}(x_i) \rightarrow y_i$ (if the values of x_i are unique), so that $\hat{m}(x)$ is simply the scatter of y_i on x_i . In contrast, as $h \rightarrow \infty$ then for all x , $\hat{m}(x) \rightarrow \bar{y}$, the sample mean, so that the nonparametric CEF estimate is a constant function. For intermediate values of h , $\hat{m}(x)$ will lie between these two extreme cases.

The uniform density is not a good kernel choice as it produces discontinuous CEF estimates. To obtain a continuous CEF estimate $\hat{m}(x)$ it is necessary for the kernel $k(u)$ to be continuous. The two most commonly used choices are the **Epanechnikov kernel**

$$k_1(u) = \frac{3}{4} (1 - u^2) 1(|u| \leq 1)$$

and the **normal** or **Gaussian kernel**

$$k_\phi(u) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{u^2}{2}\right).$$

For computation of the CEF estimate (11.4) the scale of the kernel is not important so long as the bandwidth is selected appropriately. That is, for any $b > 0$, $k_b(u) = b^{-1}k\left(\frac{u}{b}\right)$ is a valid kernel function with the identical shape as $k(u)$. Kernel regression with the kernel $k(u)$ and bandwidth h is identical to kernel regression with the kernel $k_b(u)$ and bandwidth h/b .

The estimate (11.4) using the Epanechnikov kernel and $h = 1/2$ is also displayed in Figure 11.1 with the dashed line. As you can see, this estimator appears to be much smoother than that using the uniform kernel.

Two important constants associated with a kernel function $k(u)$ are its variance σ_k^2 and roughness R_k , which are defined as

$$\sigma_k^2 = \int_{-\infty}^{\infty} u^2 k(u) du \quad (11.5)$$

$$R_k = \int_{-\infty}^{\infty} k(u)^2 du. \quad (11.6)$$

Some common kernels and their roughness and variance values are reported in Table 9.1.

Table 9.1: Common Second-Order Kernels

Kernel	Equation	R_k	σ_k^2
Uniform	$k_0(u) = \frac{1}{2} 1(u \leq 1)$	1/2	1/3
Epanechnikov	$k_1(u) = \frac{3}{4} (1 - u^2) 1(u \leq 1)$	3/5	1/5
Biweight	$k_2(u) = \frac{15}{16} (1 - u^2)^2 1(u \leq 1)$	5/7	1/7
Triweight	$k_3(u) = \frac{35}{32} (1 - u^2)^3 1(u \leq 1)$	350/429	1/9
Gaussian	$k_\phi(u) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{u^2}{2}\right)$	$1/(2\sqrt{\pi})$	1

11.4 Local Linear Estimator

The Nadaraya-Watson (NW) estimator is often called a local constant estimator as it locally (about x) approximates the CEF $m(x)$ as a constant function. One way to see this is to observe that $\hat{m}(x)$ solves the minimization problem

$$\hat{m}(x) = \operatorname{argmin}_{\alpha} \sum_{i=1}^n k\left(\frac{x_i - x}{h}\right) (y_i - \alpha)^2.$$

This is a weighted regression of y_i on an intercept only. Without the weights, this estimation problem reduces to the sample mean. The NW estimator generalizes this to a local mean.

This interpretation suggests that we can construct alternative nonparametric estimators of the CEF by alternative local approximations. Many such local approximations are possible. A popular choice is the **local linear** (LL) approximation. Instead of approximating $m(x)$ locally as a constant,

the local linear approximation approximates the CEF locally by a linear function, and estimates this local approximation by locally weighted least squares.

Specifically, for each x we solve the following minimization problem

$$\{\widehat{\alpha}(x), \widehat{\beta}(x)\} = \underset{\alpha, \beta}{\operatorname{argmin}} \sum_{i=1}^n k\left(\frac{x_i - x}{h}\right) (y_i - \alpha - \beta(x_i - x))^2.$$

The local linear estimator of $m(x)$ is the estimated intercept

$$\widehat{m}(x) = \widehat{\alpha}(x)$$

and the local linear estimator of the regression derivative $\nabla m(x)$ is the estimated slope coefficient

$$\widehat{\nabla m}(x) = \widehat{\beta}(x).$$

Computationally, for each x set

$$\mathbf{z}_i(x) = \begin{pmatrix} 1 \\ x_i - x \end{pmatrix}$$

and

$$k_i(x) = k\left(\frac{x_i - x}{h}\right).$$

Then

$$\begin{aligned} \begin{pmatrix} \widehat{\alpha}(x) \\ \widehat{\beta}(x) \end{pmatrix} &= \left(\sum_{i=1}^n k_i(x) \mathbf{z}_i(x) \mathbf{z}_i(x)' \right)^{-1} \sum_{i=1}^n k_i(x) \mathbf{z}_i(x) y_i \\ &= (\mathbf{Z}' \mathbf{K} \mathbf{Z})^{-1} \mathbf{Z}' \mathbf{K} \mathbf{y} \end{aligned}$$

where $\mathbf{K} = \operatorname{diag}\{k_1(x), \dots, k_n(x)\}$.

To visualize, Figure 11.2 displays the scatter plot of the same 100 observations from Figure 11.1, divided into three regions depending on the regressor x_i : $[1, 3]$, $[3, 5]$, $[5, 7]$. A linear regression is fit to the observations in each region, with the observations weighted by the Epanechnikov kernel with $h = 1$. The three fitted regression lines are displayed by the three straight solid lines. The values of these regression lines at $x = 2$, $x = 4$ and $x = 6$, respectively, are the local linear estimates $\widehat{m}(x)$ at $x = 2$, 4, and 6. This estimation is repeated for all x in the support of the regressors, and plotted as the continuous solid line in Figure 11.2.

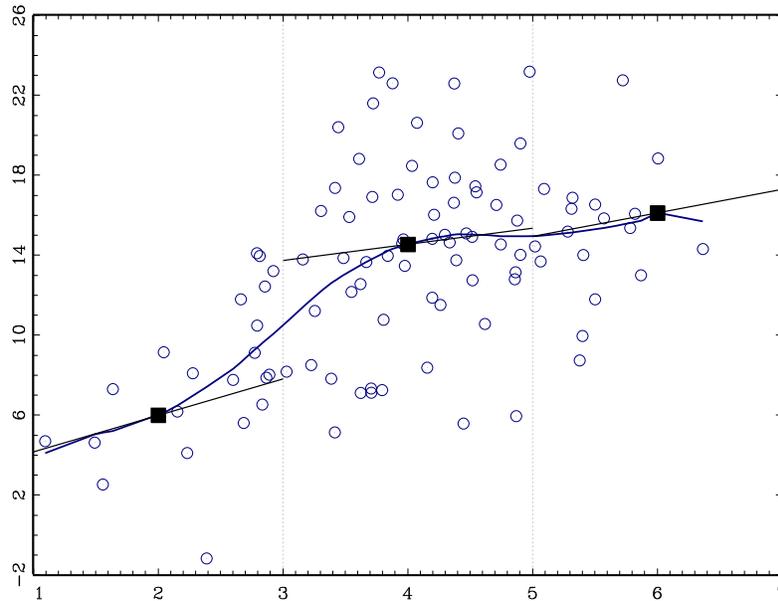
One interesting feature is that as $h \rightarrow \infty$, the LL estimator approaches the full-sample linear least-squares estimator $\widehat{m}(x) \rightarrow \widehat{\alpha} + \widehat{\beta}x$. That is because as $h \rightarrow \infty$ all observations receive equal weight regardless of x . In this sense we can see that the LL estimator is a flexible generalization of the linear OLS estimator.

Which nonparametric estimator should you use in practice: NW or LL? The theoretical literature shows that neither strictly dominates the other, but we can describe contexts where one or the other does better. Roughly speaking, the NW estimator performs better than the LL estimator when $m(x)$ is close to a flat line, but the LL estimator performs better when $m(x)$ is meaningfully non-constant. The LL estimator also performs better for values of x near the boundary of the support of x_i .

11.5 Nonparametric Residuals and Regression Fit

The fitted regression at $x = x_i$ is $\widehat{m}(x_i)$ and the fitted residual is

$$\widehat{e}_i = y_i - \widehat{m}(x_i).$$

Figure 11.2: Scatter of (y_i, x_i) and Local Linear fitted regression

As a general rule, but especially when the bandwidth h is small, it is hard to view $\hat{\epsilon}_i$ as a good measure of the fit of the regression. As $h \rightarrow 0$ then $\hat{m}(x_i) \rightarrow y_i$ and therefore $\hat{\epsilon}_i \rightarrow 0$. This clearly indicates overfitting as the true error is not zero. In general, since $\hat{m}(x_i)$ is a local average which includes y_i , the fitted value will be necessarily close to y_i and the residual $\hat{\epsilon}_i$ small, and the degree of this overfitting increases as h decreases.

A standard solution is to measure the fit of the regression at $x = x_i$ by re-estimating the model excluding the i 'th observation. For Nadaraya-Watson regression, the leave-one-out estimator of $m(x)$ excluding observation i is

$$\tilde{m}_{-i}(x) = \frac{\sum_{j \neq i} k\left(\frac{x_j - x}{h}\right) y_j}{\sum_{j \neq i} k\left(\frac{x_j - x}{h}\right)}.$$

Notationally, the “ $-i$ ” subscript is used to indicate that the i 'th observation is omitted.

The leave-one-out predicted value for y_i at $x = x_i$ equals

$$\tilde{y}_i = \tilde{m}_{-i}(x_i) = \frac{\sum_{j \neq i} k\left(\frac{x_j - x_i}{h}\right) y_j}{\sum_{j \neq i} k\left(\frac{x_j - x_i}{h}\right)}.$$

The leave-one-out residuals (or prediction errors) are the difference between the leave-one-out predicted values and the actual observation

$$\tilde{\epsilon}_i = y_i - \tilde{y}_i.$$

Since \tilde{y}_i is not a function of y_i , there is no tendency for \tilde{y}_i to overfit for small h . Consequently, $\tilde{\epsilon}_i$ is a good measure of the fit of the estimated nonparametric regression.

Similarly, the leave-one-out local-linear residual is $\tilde{\epsilon}_i = y_i - \tilde{\alpha}_i$ with

$$\begin{pmatrix} \tilde{\alpha}_i \\ \tilde{\beta}_i \end{pmatrix} = \left(\sum_{j \neq i} k_{ij} \mathbf{z}_{ij} \mathbf{z}'_{ij} \right)^{-1} \sum_{j \neq i} k_{ij} \mathbf{z}_{ij} y_j,$$

$$z_{ij} = \begin{pmatrix} 1 \\ x_j - x_i \end{pmatrix}$$

and

$$k_{ij} = k\left(\frac{x_j - x_i}{h}\right).$$

11.6 Cross-Validation Bandwidth Selection

As we mentioned before, the choice of bandwidth h is crucial. As h increases, the kernel regression estimators (both NW and LL) become more smooth, ironing out the bumps and wiggles. This reduces estimation variance but at the cost of increased bias and oversmoothing. As h decreases the estimators become more wiggly, erratic, and noisy. It is desirable to select h to trade-off these features. How can this be done systematically?

To be explicit about the dependence of the estimator on the bandwidth, let us write the estimator of $m(x)$ with a given bandwidth h as $\hat{m}(x, h)$, and our discussion will apply equally to the NW and LL estimators.

Ideally, we would like to select h to minimize the mean-squared error (MSE) of $\hat{m}(x, h)$ as a estimate of $m(x)$. For a given value of x the MSE is

$$MSE_n(x, h) = \mathbb{E}(\hat{m}(x, h) - m(x))^2.$$

We are typically interested in estimating $m(x)$ for all values in the support of x . A common measure for the average fit is the integrated MSE

$$\begin{aligned} IMSE_n(h) &= \int MSE_n(x, h) f_x(x) dx \\ &= \int \mathbb{E}(\hat{m}(x, h) - m(x))^2 f_x(x) dx \end{aligned}$$

where $f_x(x)$ is the marginal density of x_i . Notice that we have defined the IMSE as an integral with respect to the density $f_x(x)$. Other weight functions could be used, but it turns out that this is a convenient choice.

The IMSE is closely related with the MSFE of Section 4.9. Let (y_{n+1}, x_{n+1}) be out-of-sample observations (and thus independent of the sample) and consider predicting y_{n+1} given x_{n+1} and the nonparametric estimate $\hat{m}(x, h)$. The natural point estimate for y_{n+1} is $\hat{m}(x_{n+1}, h)$ which has mean-squared forecast error

$$\begin{aligned} MSFE_n(h) &= \mathbb{E}(y_{n+1} - \hat{m}(x_{n+1}, h))^2 \\ &= \mathbb{E}(e_{n+1} + m(x_{n+1}) - \hat{m}(x_{n+1}, h))^2 \\ &= \sigma^2 + \mathbb{E}(m(x_{n+1}) - \hat{m}(x_{n+1}, h))^2 \\ &= \sigma^2 + \int \mathbb{E}(\hat{m}(x, h) - m(x))^2 f_x(x) dx \end{aligned}$$

where the final equality uses the fact that x_{n+1} is independent of $\hat{m}(x, h)$. We thus see that

$$MSFE_n(h) = \sigma^2 + IMSE_n(h).$$

Since σ^2 is a constant independent of the bandwidth h , $MSFE_n(h)$ and $IMSE_n(h)$ are equivalent measures of the fit of the nonparametric regression.

The optimal bandwidth h is the value which minimizes $IMSE_n(h)$ (or equivalently $MSFE_n(h)$). While these functions are unknown, we learned in Theorem 4.9.1 that (at least in the case of linear regression) $MSFE_n$ can be estimated by the sample mean-squared prediction errors. It turns out that this fact extends to nonparametric regression. The nonparametric leave-one-out residuals are

$$\tilde{e}_i(h) = y_i - \tilde{m}_{-i}(x_i, h)$$

where we are being explicit about the dependence on the bandwidth h . The mean squared leave-one-out residuals is

$$CV(h) = \frac{1}{n} \sum_{i=1}^n \tilde{e}_i(h)^2.$$

This function of h is known as the **cross-validation criterion**.

The cross-validation bandwidth \hat{h} is the value which minimizes $CV(h)$

$$\hat{h} = \underset{h \geq h_\ell}{\operatorname{argmin}} CV(h) \quad (11.7)$$

for some $h_\ell > 0$. The restriction $h \geq h_\ell$ is imposed so that $CV(h)$ is not evaluated over unreasonably small bandwidths.

There is not an explicit solution to the minimization problem (11.7), so it must be solved numerically. A typical practical method is to create a grid of values for h , e.g. $[h_1, h_2, \dots, h_J]$, evaluate $CV(h_j)$ for $j = 1, \dots, J$, and set

$$\hat{h} = \underset{h \in [h_1, h_2, \dots, h_J]}{\operatorname{argmin}} CV(h).$$

Evaluation using a coarse grid is typically sufficient for practical application. Plots of $CV(h)$ against h are a useful diagnostic tool to verify that the minimum of $CV(h)$ has been obtained.

We said above that the cross-validation criterion is an estimator of the MSFE. This claim is based on the following result.

Theorem 11.6.1

$$\mathbb{E}(CV(h)) = MSFE_{n-1}(h) = IMSE_{n-1}(h) + \sigma^2 \quad (11.8)$$

Theorem 11.6.1 shows that $CV(h)$ is an unbiased estimator of $IMSE_{n-1}(h) + \sigma^2$. The first term, $IMSE_{n-1}(h)$, is the integrated MSE of the nonparametric estimator using a sample of size $n - 1$. If n is large, $IMSE_{n-1}(h)$ and $IMSE_n(h)$ will be nearly identical, so $CV(h)$ is essentially unbiased as an estimator of $IMSE_n(h) + \sigma^2$. Since the second term (σ^2) is unaffected by the bandwidth h , it is irrelevant for the problem of selection of h . In this sense we can view $CV(h)$ as an estimator of the IMSE, and more importantly we can view the minimizer of $CV(h)$ as an estimate of the minimizer of $IMSE_n(h)$.

To illustrate, Figure 11.3 displays the cross-validation criteria $CV(h)$ for the Nadaraya-Watson and Local Linear estimators using the data from Figure 11.1, both using the Epanechnikov kernel. The CV functions are computed on a grid with intervals 0.01. The CV-minimizing bandwidths are $h = 1.09$ for the Nadaraya-Watson estimator and $h = 1.59$ for the local linear estimator. Figure 11.3 shows the minimizing bandwidths by the arrows. It is not surprising that the CV criteria recommend a larger bandwidth for the LL estimator than for the NW estimator, as the LL employs more smoothing for a given bandwidth.

The CV criterion can also be used to select between different nonparametric estimators. The CV-selected estimator is the one with the lowest minimized CV criterion. For example, in Figure 11.3, the NW estimator has a minimized CV criterion of 16.88, while the LL estimator has a minimized CV criterion of 16.81. Since the LL estimator achieves a lower value of the CV criterion, LL is the CV-selected estimator. The difference (0.07) is small, suggesting that the two estimators are near equivalent in IMSE.

Figure 11.4 displays the fitted CEF estimates (NW and LL) using the bandwidths selected by cross-validation. Also displayed is the true CEF $m(x) = 10 \ln(x)$. Notice that the nonparametric

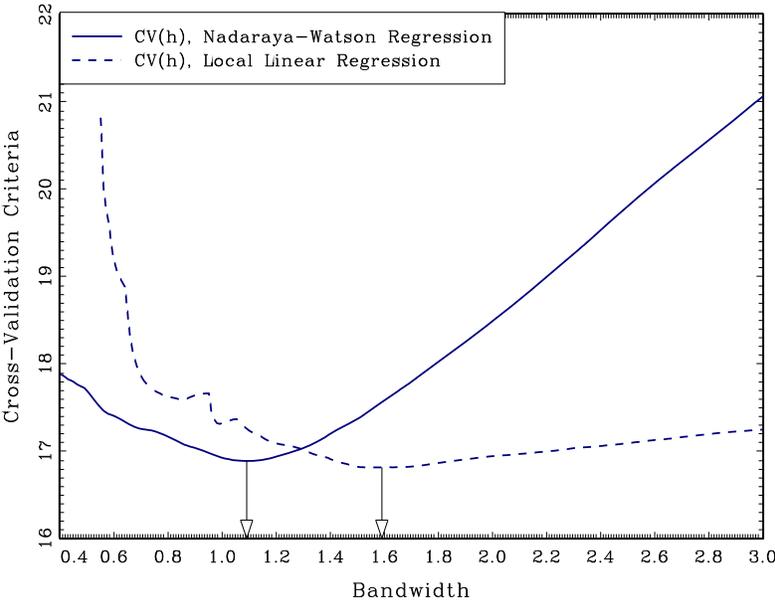


Figure 11.3: Cross-Validation Criteria, Nadaraya-Watson Regression and Local Linear Regression

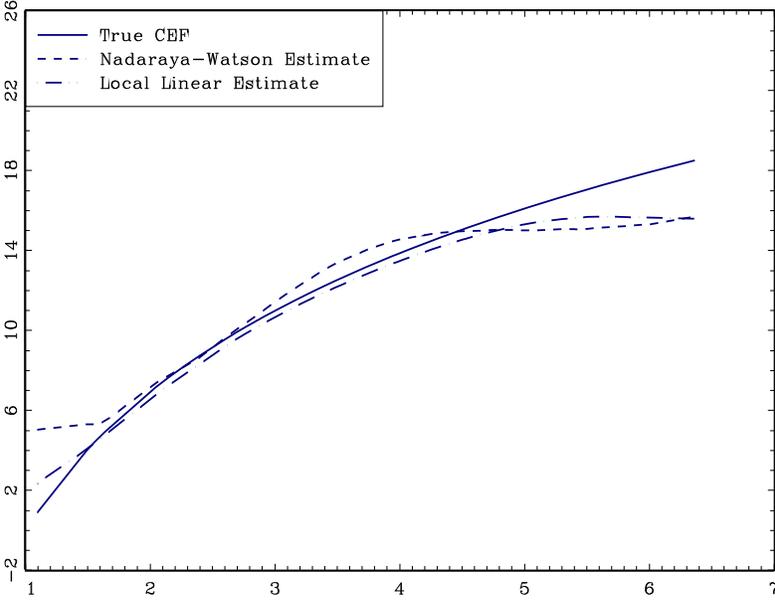


Figure 11.4: Nonparametric Estimates using data-dependent (CV) bandwidths

estimators with the CV-selected bandwidths (and especially the LL estimator) track the true CEF quite well.

Proof of Theorem 11.6.1. Observe that $m(x_i) - \tilde{m}_{-i}(x_i, h)$ is a function only of (x_1, \dots, x_n) and (e_1, \dots, e_n) excluding e_i , and is thus uncorrelated with e_i . Since $\tilde{e}_i(h) = m(x_i) - \tilde{m}_{-i}(x_i, h) + e_i$, then

$$\begin{aligned} \mathbb{E}(CV(h)) &= \mathbb{E}(\tilde{e}_i(h)^2) \\ &= \mathbb{E}(e_i^2) + \mathbb{E}(\tilde{m}_{-i}(x_i, h) - m(x_i))^2 \\ &\quad + 2\mathbb{E}((\tilde{m}_{-i}(x_i, h) - m(x_i))e_i) \\ &= \sigma^2 + \mathbb{E}(\tilde{m}_{-i}(x_i, h) - m(x_i))^2. \end{aligned} \tag{11.9}$$

The second term is an expectation over the random variables x_i and $\tilde{m}_{-i}(x, h)$, which are independent as the second is not a function of the i 'th observation. Thus taking the conditional expectation given the sample excluding the i 'th observation, this is the expectation over x_i only, which is the integral with respect to its density

$$\mathbb{E}_{-i}(\tilde{m}_{-i}(x_i, h) - m(x_i))^2 = \int (\tilde{m}_{-i}(x, h) - m(x))^2 f_x(x) dx.$$

Taking the unconditional expectation yields

$$\begin{aligned} \mathbb{E}(\tilde{m}_{-i}(x_i, h) - m(x_i))^2 &= \mathbb{E} \int (\tilde{m}_{-i}(x, h) - m(x))^2 f_x(x) dx \\ &= IMSE_{n-1}(h) \end{aligned}$$

where this is the IMSE of a sample of size $n - 1$ as the estimator \tilde{m}_{-i} uses $n - 1$ observations. Combined with (11.9) we obtain (11.8), as desired. ■

11.7 Asymptotic Distribution

There is no finite sample distribution theory for kernel estimators, but there is a well developed asymptotic distribution theory. The theory is based on the approximation that the bandwidth h decreases to zero as the sample size n increases. This means that the smoothing is increasingly localized as the sample size increases. So long as the bandwidth does not decrease to zero too quickly, the estimator can be shown to be asymptotically normal, but with a non-trivial bias.

Let $f_x(x)$ denote the marginal density of x_i and $\sigma^2(x) = \mathbb{E}(e_i^2 | x_i = x)$ denote the conditional variance of $e_i = y_i - m(x_i)$.

Theorem 11.7.1 *Let $\widehat{m}(x)$ denote either the Nadarya-Watson or Local Linear estimator of $m(x)$. If x is interior to the support of x_i and $f_x(x) > 0$, then as $n \rightarrow \infty$ and $h \rightarrow 0$ such that $nh \rightarrow \infty$,*

$$\sqrt{nh} (\widehat{m}(x) - m(x) - h^2 \sigma_k^2 B(x)) \xrightarrow{d} \mathbf{N} \left(0, \frac{R_k \sigma^2(x)}{f_x(x)} \right) \quad (11.10)$$

where σ_k^2 and R_k are defined in (11.5) and (11.6). For the Nadaraya-Watson estimator

$$B(x) = \frac{1}{2} m''(x) + f_x(x)^{-1} f'_x(x) m'(x)$$

and for the local linear estimator

$$B(x) = \frac{1}{2} f_x(x) m''(x)$$

There are several interesting features about the asymptotic distribution which are noticeably different than for parametric estimators. First, the estimator converges at the rate \sqrt{nh} , not \sqrt{n} . Since $h \rightarrow 0$, \sqrt{nh} diverges slower than \sqrt{n} , thus the nonparametric estimator converges more slowly than a parametric estimator. Second, the asymptotic distribution contains a non-negligible bias term $h^2 \sigma_k^2 B(x)$. This term asymptotically disappears since $h \rightarrow 0$. Third, the assumptions that $nh \rightarrow \infty$ and $h \rightarrow 0$ mean that the estimator is consistent for the CEF $m(x)$.

The fact that the estimator converges at the rate \sqrt{nh} has led to the interpretation of nh as the “effective sample size”. This is because the number of observations being used to construct $\widehat{m}(x)$ is proportional to nh , not n as for a parametric estimator.

It is helpful to understand that the nonparametric estimator has a reduced convergence rate because the object being estimated – $m(x)$ – is nonparametric. This is harder than estimating a finite dimensional parameter, and thus comes at a cost.

Unlike parametric estimation, the asymptotic distribution of the nonparametric estimator includes a term representing the bias of the estimator. The asymptotic distribution (11.10) shows the form of this bias. Not only is it proportional to the squared bandwidth h^2 (the degree of smoothing), it is proportional to the function $B(x)$ which depends on the slope and curvature of the CEF $m(x)$. Interestingly, when $m(x)$ is constant then $B(x) = 0$ and the kernel estimator has no asymptotic bias. The bias is essentially increasing in the curvature of the CEF function $m(x)$. This is because the local averaging smooths $m(x)$, and the smoothing induces more bias when $m(x)$ is curved.

Theorem 11.7.1 shows that the asymptotic distributions of the NW and LL estimators are similar, with the only difference arising in the bias function $B(x)$. The bias term for the NW estimator has an extra component which depends on the first derivative of the CEF $m(x)$ while the bias term of the LL estimator is invariant to the first derivative. The fact that the bias formula for the LL estimator is simpler and is free of dependence on the first derivative of $m(x)$ suggests that the LL estimator will generally have smaller bias than the NW estimator (but this is not a precise ranking). Since the asymptotic variances in the two distributions are the same, this means that the LL estimator achieves a reduced bias without an effect on asymptotic variance. This analysis has led to the general preference for the LL estimator over the NW estimator in the nonparametrics literature.

One implication of Theorem 11.7.1 is that we can define the asymptotic MSE (AMSE) of $\widehat{m}(x)$ as the squared bias plus the asymptotic variance

$$AMSE(\widehat{m}(x)) = (h^2 \sigma_k^2 B(x))^2 + \frac{R_k \sigma^2(x)}{nh f_x(x)}. \quad (11.11)$$

Focusing on rates, this says

$$AMSE(\widehat{m}(x)) \sim h^4 + \frac{1}{nh} \quad (11.12)$$

which means that the AMSE is dominated by the larger of h^4 and $(nh)^{-1}$. Notice that the bias is increasing in h and the variance is decreasing in h . (More smoothing means more observations are used for local estimation: this increases the bias but decreases estimation variance.) To select h to minimize the AMSE, these two components should balance each other. Setting $h^4 \propto (nh)^{-1}$ means setting $h \propto n^{-1/5}$. Another way to see this is to pick h to minimize the right-hand-side of (11.12). The first-order condition for h is

$$\frac{\partial}{\partial h} \left(h^4 + \frac{1}{nh} \right) = 4h^3 - \frac{1}{nh^2} = 0$$

which when solved for h yields $h = n^{-1/5}$. What this means is that for AMSE-efficient estimation of $m(x)$, the optimal rate for the bandwidth is $h \propto n^{-1/5}$.

Theorem 11.7.2 *The bandwidth which minimizes the AMSE (11.12) is of order $h \propto n^{-1/5}$. With $h \propto n^{-1/5}$ then $AMSE(\widehat{m}(x)) = O(n^{-4/5})$ and $\widehat{m}(x) = m(x) + O_p(n^{-2/5})$.*

This result means that the bandwidth should take the form $h = cn^{-1/5}$. The optimal constant c depends on the kernel k , the bias function $B(x)$ and the marginal density $f_x(x)$. A common misinterpretation is to set $h = n^{-1/5}$, which is equivalent to setting $c = 1$ and is completely arbitrary. Instead, an empirical bandwidth selection rule such as cross-validation should be used in practice.

When $h = cn^{-1/5}$ we can rewrite the asymptotic distribution (11.10) as

$$n^{2/5}(\widehat{m}(x) - m(x)) \xrightarrow{d} N\left(c^2 \sigma_k^2 B(x), \frac{R_k \sigma^2(x)}{c^{1/2} f_x(x)}\right)$$

In this representation, we see that $\widehat{m}(x)$ is asymptotically normal, but with a $n^{2/5}$ rate of convergence and non-zero mean. The asymptotic distribution depends on the constant c through the bias (positively) and the variance (inversely).

The asymptotic distribution in Theorem 11.7.1 allows for the optimal rate $h = cn^{-1/5}$ but this rate is not required. In particular, consider an undersmoothing (smaller than optimal) bandwidth with rate $h = o(n^{-1/5})$. For example, we could specify that $h = cn^{-\alpha}$ for some $c > 0$ and $1/5 < \alpha < 1$. Then $\sqrt{nh}h^2 = O(n^{(1-5\alpha)/2}) = o(1)$ so the bias term in (11.10) is asymptotically negligible so Theorem 11.7.1 implies

$$\sqrt{nh}(\widehat{m}(x) - m(x)) \xrightarrow{d} N\left(0, \frac{R_k \sigma^2(x)}{f_x(x)}\right).$$

That is, the estimator is asymptotically normal without a bias component. Not having an asymptotic bias component is convenient for some theoretical manipulations, so many authors impose the undersmoothing condition $h = o(n^{-1/5})$ to ensure this situation. This convenience comes at a cost. First, the resulting estimator is inefficient as its convergence rate is $O_p(n^{-(1-\alpha)/2}) > O_p(n^{-2/5})$ since $\alpha > 1/5$. Second, the distribution theory is an inherently misleading approximation as it misses a critically key ingredient of nonparametric estimation – the trade-off between bias and variance. The approximation (11.10) is superior precisely because it contains the asymptotic bias component which is a realistic implication of nonparametric estimation. Undersmoothing assumptions should be avoided when possible.

11.8 Conditional Variance Estimation

Let's consider the problem of estimation of the conditional variance

$$\begin{aligned}\sigma^2(x) &= \text{var}(y_i \mid x_i = x) \\ &= \mathbb{E}(e_i^2 \mid x_i = x).\end{aligned}$$

Even if the conditional mean $m(x)$ is parametrically specified, it is natural to view $\sigma^2(x)$ as inherently nonparametric as economic models rarely specify the form of the conditional variance. Thus it is quite appropriate to estimate $\sigma^2(x)$ nonparametrically.

We know that $\sigma^2(x)$ is the CEF of e_i^2 given x_i . Therefore if e_i^2 were observed, $\sigma^2(x)$ could be nonparametrically estimated using NW or LL regression. For example, the ideal NW estimator is

$$\bar{\sigma}^2(x) = \frac{\sum_{i=1}^n k_i(x) e_i^2}{\sum_{i=1}^n k_i(x)}.$$

Since the errors e_i are not observed, we need to replace them with an empirical residual, such as $\hat{e}_i = y_i - \hat{m}(x_i)$ where $\hat{m}(x)$ is the estimated CEF. (The latter could be a nonparametric estimator such as NW or LL, or even a parametric estimator.) Even better, use the leave-one-out prediction errors $\tilde{e}_i = y_i - \hat{m}_{-i}(x_i)$, as these are not subject to overfitting.

With this substitution the NW estimator of the conditional variance is

$$\tilde{\sigma}^2(x) = \frac{\sum_{i=1}^n k_i(x) \tilde{e}_i^2}{\sum_{i=1}^n k_i(x)}. \quad (11.13)$$

This estimator depends on a set of bandwidths h_1, \dots, h_q , but there is no reason for the bandwidths to be the same as those used to estimate the conditional mean. Cross-validation can be used to select the bandwidths for estimation of $\tilde{\sigma}^2(x)$ separately from cross-validation for estimation of $\hat{m}(x)$.

There is one subtle difference between CEF and conditional variance estimation. The conditional variance is inherently non-negative $\sigma^2(x) \geq 0$ and it is desirable for our estimator to satisfy this property. Interestingly, the NW estimator (11.13) is necessarily non-negative, since it is a smoothed average of the non-negative squared residuals, but the LL estimator is not guaranteed to be non-negative for all x . For this reason, the NW estimator is preferred for conditional variance estimation.

Fan and Yao (1998, *Biometrika*) derive the asymptotic distribution of the estimator (11.13). They obtain the surprising result that the asymptotic distribution of this two-step estimator is identical to that of the one-step idealized estimator $\tilde{\sigma}^2(x)$.

11.9 Standard Errors

Theorem 11.7.1 shows the asymptotic variances of both the NW and LL nonparametric regression estimators equal

$$V(x) = \frac{R_k \sigma^2(x)}{f_x(x)}.$$

For standard errors we need an estimate of $V(x)$. A plug-in estimate replaces the unknowns by estimates. The roughness R_k can be found from Table 9.1. The conditional variance can be estimated using (11.13). The density of x_i can be estimated using the methods from Section 20.1. Replacing these estimates into the formula for $V(x)$ we obtain the asymptotic variance estimate

$$\hat{V}(x) = \frac{R_k \tilde{\sigma}^2(x)}{\hat{f}_x(x)}.$$

Then an asymptotic standard error for the kernel estimate $\widehat{m}(\mathbf{x})$ is

$$\widehat{s}(x) = \sqrt{\frac{1}{nh} \widehat{V}(x)}.$$

Plots of the estimated CEF $\widehat{m}(x)$ can be accompanied by confidence intervals $\widehat{m}(x) \pm 2\widehat{s}(x)$. These are known as **pointwise confidence intervals**, as they are designed to have correct coverage at each x , not uniformly in x .

One important caveat about the interpretation of nonparametric confidence intervals is that they are not centered at the true CEF $m(x)$, but rather are centered at the biased or pseudo-true value

$$m^*(x) = m(x) + h^2 \sigma_k^2 B(x).$$

Consequently, a correct statement about the confidence interval $\widehat{m}(x) \pm 2\widehat{s}(x)$ is that it asymptotically contains $m^*(x)$ with probability 95%, not that it asymptotically contains $m(x)$ with probability 95%. The discrepancy is that the confidence interval does not take into account the bias $h^2 \sigma_k^2 B(x)$. Unfortunately, nothing constructive can be done about this. The bias is difficult and noisy to estimate, so making a bias-correction only inflates estimation variance and decreases overall precision. A technical “trick” is to assume undersmoothing $h = o(n^{-1/5})$ but this does not really eliminate the bias, it only assumes it away. The plain fact is that once we honestly acknowledge that the true CEF is nonparametric, it then follows that any finite sample estimate will have finite sample bias, and this bias will be inherently unknown and thus impossible to incorporate into confidence intervals.

11.10 Multiple Regressors

Our analysis has focus on the case of real-valued x_i for simplicity of exposition, but the methods of kernel regression extend easily to the multiple regressor case, at the cost of a reduced rate of convergence. In this section we consider the case of estimation of the conditional expectation function

$$\mathbb{E}(y_i \mid \mathbf{x}_i = \mathbf{x}) = m(\mathbf{x})$$

when

$$\mathbf{x}_i = \begin{pmatrix} x_{1i} \\ \vdots \\ x_{di} \end{pmatrix}$$

is a d -vector.

For any evaluation point \mathbf{x} and observation i , define the kernel weights

$$k_i(\mathbf{x}) = k\left(\frac{x_{1i} - x_1}{h_1}\right) k\left(\frac{x_{2i} - x_2}{h_2}\right) \cdots k\left(\frac{x_{di} - x_d}{h_d}\right),$$

a d -fold product kernel. The kernel weights $k_i(\mathbf{x})$ assess if the regressor vector \mathbf{x}_i is close to the evaluation point \mathbf{x} in the Euclidean space \mathbb{R}^d .

These weights depend on a set of d bandwidths, h_j , one for each regressor. We can group them together into a single vector for notational convenience:

$$\mathbf{h} = \begin{pmatrix} h_1 \\ \vdots \\ h_d \end{pmatrix}.$$

Given these weights, the Nadaraya-Watson estimator takes the form

$$\widehat{m}(\mathbf{x}) = \frac{\sum_{i=1}^n k_i(\mathbf{x}) y_i}{\sum_{i=1}^n k_i(\mathbf{x})}.$$

For the local-linear estimator, define

$$\mathbf{z}_i(\mathbf{x}) = \begin{pmatrix} 1 \\ \mathbf{x}_i - \mathbf{x} \end{pmatrix}$$

and then the local-linear estimator can be written as $\hat{m}(\mathbf{x}) = \hat{\alpha}(\mathbf{x})$ where

$$\begin{aligned} \begin{pmatrix} \hat{\alpha}(\mathbf{x}) \\ \hat{\beta}(\mathbf{x}) \end{pmatrix} &= \left(\sum_{i=1}^n k_i(\mathbf{x}) \mathbf{z}_i(\mathbf{x}) \mathbf{z}_i(\mathbf{x})' \right)^{-1} \sum_{i=1}^n k_i(\mathbf{x}) \mathbf{z}_i(\mathbf{x}) y_i \\ &= (\mathbf{Z}' \mathbf{K} \mathbf{Z})^{-1} \mathbf{Z}' \mathbf{K} \mathbf{y} \end{aligned}$$

where $\mathbf{K} = \text{diag}\{k_1(\mathbf{x}), \dots, k_n(\mathbf{x})\}$.

In multiple regressor kernel regression, cross-validation remains a recommended method for bandwidth selection. The leave-one-out residuals \tilde{e}_i and cross-validation criterion $CV(\mathbf{h})$ are defined identically as in the single regressor case. The only difference is that now the CV criterion is a function over the d -dimensional bandwidth \mathbf{h} . This is a critical practical difference since finding the bandwidth vector $\hat{\mathbf{h}}$ which minimizes $CV(\mathbf{h})$ can be computationally difficult when \mathbf{h} is high dimensional. Grid search is cumbersome and costly, since G gridpoints per dimension imply evaluation of $CV(\mathbf{h})$ at G^d distinct points, which can be a large number. Furthermore, plots of $CV(\mathbf{h})$ against \mathbf{h} are challenging when $d > 2$.

The asymptotic distribution of the estimators in the multiple regressor case is an extension of the single regressor case. Let $f_x(\mathbf{x})$ denote the marginal density of \mathbf{x}_i and $\sigma^2(\mathbf{x}) = \mathbb{E}(e_i^2 \mid \mathbf{x}_i = \mathbf{x})$ the conditional variance of $e_i = y_i - m(\mathbf{x}_i)$. Let $|\mathbf{h}| = h_1 h_2 \cdots h_d$.

Theorem 11.10.1 *Let $\hat{m}(\mathbf{x})$ denote either the Nadaraya-Watson or Local Linear estimator of $m(\mathbf{x})$. If \mathbf{x} is interior to the support of \mathbf{x}_i and $f_x(\mathbf{x}) > 0$, then as $n \rightarrow \infty$ and $h_j \rightarrow 0$ such that $n|\mathbf{h}| \rightarrow \infty$,*

$$\sqrt{n|\mathbf{h}|} \left(\hat{m}(\mathbf{x}) - m(\mathbf{x}) - \sigma_k^2 \sum_{j=1}^d h_j^2 B_{j,j}(\mathbf{x}) \right) \xrightarrow{d} \text{N} \left(0, \frac{R_k^d \sigma^2(\mathbf{x})}{f_x(\mathbf{x})} \right)$$

where for the Nadaraya-Watson estimator

$$B_j(\mathbf{x}) = \frac{1}{2} \frac{\partial^2}{\partial x_j^2} m(\mathbf{x}) + f_x(\mathbf{x})^{-1} \frac{\partial}{\partial x_j} f_x(\mathbf{x}) \frac{\partial}{\partial x_j} m(\mathbf{x})$$

and for the Local Linear estimator

$$B_j(\mathbf{x}) = \frac{1}{2} \frac{\partial^2}{\partial x_j^2} m(\mathbf{x})$$

For notational simplicity consider the case that there is a single common bandwidth h . In this case the AMSE takes the form

$$AMSE(\hat{m}(\mathbf{x})) \sim h^4 + \frac{1}{nh^d}$$

That is, the squared bias is of order h^4 , the same as in the single regressor case, but the variance is of larger order $(nh^d)^{-1}$. Setting h to balance these two components requires setting $h \sim n^{-1/(4+d)}$.

Theorem 11.10.2 *The bandwidth which minimizes the AMSE is of order $h \propto n^{-1/(4+d)}$. With $h \propto n^{-1/(4+d)}$ then $AMSE(\widehat{m}(\mathbf{x})) = O(n^{-4/(4+d)})$ and $\widehat{m}(\mathbf{x}) = m(\mathbf{x}) + O_p(n^{-2/(4+d)})$*

In all estimation problems an increase in the dimension decreases estimation precision. For example, in parametric estimation an increase in dimension typically increases the asymptotic variance. In nonparametric estimation an increase in the dimension typically decreases the convergence rate, which is a more fundamental decrease in precision. For example, in kernel regression the convergence rate $O_p(n^{-2/(4+d)})$ decreases as d increases. The reason is the estimator $\widehat{m}(\mathbf{x})$ is a local average of the y_i for observations such that \mathbf{x}_i is close to \mathbf{x} , and when there are multiple regressors the number of such observations is inherently smaller. This phenomenon – that the rate of convergence of nonparametric estimation decreases as the dimension increases – is called the **curse of dimensionality**.

Chapter 12

Series Estimation

12.1 Approximation by Series

As we mentioned at the beginning of Chapter 11, there are two main methods of nonparametric regression: kernel estimation and series estimation. In this chapter we study series methods.

Series methods approximate an unknown function (e.g. the CEF $m(\mathbf{x})$) with a flexible parametric function, with the number of parameters treated similarly to the bandwidth in kernel regression. A series approximation to $m(\mathbf{x})$ takes the form $m_K(\mathbf{x}) = m_K(\mathbf{x}, \boldsymbol{\beta}_K)$ where $m_K(\mathbf{x}, \boldsymbol{\beta}_K)$ is a known parametric family and $\boldsymbol{\beta}_K$ is an unknown coefficient. The integer K is the dimension of $\boldsymbol{\beta}_K$ and indexes the complexity of the approximation.

A linear series approximation takes the form

$$\begin{aligned} m_K(\mathbf{x}) &= \sum_{j=1}^K z_{jK}(\mathbf{x})\beta_{jK} \\ &= \mathbf{z}_K(\mathbf{x})'\boldsymbol{\beta}_K \end{aligned} \tag{12.1}$$

where $z_{jK}(\mathbf{x})$ are (nonlinear) functions of \mathbf{x} , and are known as **basis functions** or **basis function transformations** of \mathbf{x} .

For real-valued x , a well-known linear series approximation is the p 'th-order **polynomial**

$$m_K(x) = \sum_{j=0}^p x^j \beta_{jK}$$

where $K = p + 1$.

When $\mathbf{x} \in \mathbb{R}^d$ is vector-valued, a p 'th-order polynomial is

$$m_K(\mathbf{x}) = \sum_{j_1=0}^p \cdots \sum_{j_d=0}^p x_1^{j_1} \cdots x_d^{j_d} \beta_{j_1, \dots, j_d K}.$$

This includes all powers and cross-products, and the coefficient vector has dimension $K = (p+1)^d$. In general, a common method to create a series approximation for vector-valued \mathbf{x} is to include all non-redundant cross-products of the basis function transformations of the components of \mathbf{x} .

12.2 Splines

Another common series approximation is a continuous piecewise polynomial function known as a **spline**. While splines can be of any polynomial order (e.g. linear, quadratic, cubic, etc.), a common choice is cubic. To impose smoothness it is common to constrain the spline function to have continuous derivatives up to the order of the spline. Thus a quadratic spline is typically

constrained to have a continuous first derivative, and a cubic spline is typically constrained to have a continuous first and second derivative.

There is more than one way to define a spline series expansion. All are based on the number of **knots** – the join points between the polynomial segments.

To illustrate, a piecewise linear function with two segments and a knot at t is

$$m_K(x) = \begin{cases} m_1(x) = \beta_{00} + \beta_{01}(x - t) & x < t \\ m_2(x) = \beta_{10} + \beta_{11}(x - t) & x \geq t \end{cases}$$

(For convenience we have written the segments functions as polynomials in $x - t$.) The function $m_K(x)$ equals the linear function $m_1(x)$ for $x < t$ and equals $m_2(x)$ for $x > t$. Its left limit at $x = t$ is β_{00} and its right limit is β_{10} , so is continuous if (and only if) $\beta_{00} = \beta_{10}$. Enforcing this constraint is equivalent to writing the function as

$$m_K(x) = \beta_0 + \beta_1(x - t) + \beta_2(x - t) \mathbf{1}(x \geq t)$$

or after transforming coefficients, as

$$m_K(x) = \beta_0 + \beta_1 x + \beta_2(x - t) \mathbf{1}(x \geq t)$$

Notice that this function has $K = 3$ coefficients, the same as a quadratic polynomial.

A piecewise quadratic function with one knot at t is

$$m_K(x) = \begin{cases} m_1(x) = \beta_{00} + \beta_{01}(x - t) + \beta_{02}(x - t)^2 & x < t \\ m_2(x) = \beta_{10} + \beta_{11}(x - t) + \beta_{12}(x - t)^2 & x \geq t \end{cases}$$

This function is continuous at $x = t$ if $\beta_{00} = \beta_{10}$, and has a continuous first derivative if $\beta_{01} = \beta_{11}$. Imposing these constraints and rewriting, we obtain the function

$$m_K(x) = \beta_0 + \beta_1 x + \beta_2 x^2 + \beta_3(x - t)^2 \mathbf{1}(x \geq t).$$

Here, $K = 4$.

Furthermore, a piecewise cubic function with one knot and a continuous second derivative is

$$m_K(x) = \beta_0 + \beta_1 x + \beta_2 x^2 + \beta_3 x^3 + \beta_4(x - t)^3 \mathbf{1}(x \geq t)$$

which has $K = 5$.

The polynomial order p is selected to control the smoothness of the spline, as $m_K(x)$ has continuous derivatives up to $p - 1$.

In general, a p 'th-order spline with N knots at t_1, t_2, \dots, t_N with $t_1 < t_2 < \dots < t_N$ is

$$m_K(x) = \sum_{j=0}^p \beta_j x^j + \sum_{k=1}^N \gamma_k (x - t_k)^p \mathbf{1}(x \geq t_k)$$

which has $K = N + p + 1$ coefficients.

In spline approximation, the typical approach is to treat the polynomial order p as fixed, and select the number of knots N to determine the complexity of the approximation. The knots t_k are typically treated as fixed. A common choice is to set the knots to evenly partition the support \mathcal{X} of \mathbf{x}_i .

12.3 Partially Linear Model

A common use of a series expansion is to allow the CEF to be nonparametric with respect to one variable, yet linear in the other variables. This allows flexibility in a particular variable of interest. A partially linear CEF with vector-valued regressor \mathbf{x}_1 and real-valued continuous x_2 takes the form

$$m(\mathbf{x}_1, x_2) = \mathbf{x}'_1 \boldsymbol{\beta}_1 + m_2(x_2).$$

This model is commonly used when \mathbf{x}_1 are discrete (e.g. binary variables) and x_2 is continuously distributed.

Series methods are particularly convenient for estimation of partially linear models, as we can replace the unknown function $m_2(x_2)$ with a series expansion to obtain

$$\begin{aligned} m(\mathbf{x}) &\simeq m_K(\mathbf{x}) \\ &= \mathbf{x}'_1 \boldsymbol{\beta}_1 + \mathbf{z}'_K \boldsymbol{\beta}_{2K} \\ &= \mathbf{x}'_K \boldsymbol{\beta}_K \end{aligned}$$

where $\mathbf{z}_K = \mathbf{z}_K(x_2)$ are the basis transformations of x_2 (typically polynomials or splines) and $\boldsymbol{\beta}_{2K}$ are coefficients. After transformation the regressors are $\mathbf{x}_K = (\mathbf{x}'_1, \mathbf{z}'_K)$. and the coefficients are $\boldsymbol{\beta}_K = (\boldsymbol{\beta}'_1, \boldsymbol{\beta}'_{2K})'$.

12.4 Additively Separable Models

When \mathbf{x} is multivariate a common simplification is to treat the regression function $m(\mathbf{x})$ as additively separable in the individual regressors, which means that

$$m(\mathbf{x}) = m_1(x_1) + m_2(x_2) + \cdots + m_d(x_d).$$

Series methods are quite convenient for estimation of additively separable models, as we simply apply series expansions (polynomials or splines) separately for each component $m_j(x_j)$. The advantage of additive separability is the reduction in dimensionality. While an unconstrained p 'th order polynomial has $(p+1)^d$ coefficients, an additively separable polynomial model has only $(p+1)d$ coefficients. This can be a major reduction in the number of coefficients. The disadvantage of this simplification is that the interaction effects have been eliminated.

The decision to impose additive separability can be based on an economic model which suggests the absence of interaction effects, or can be a model selection decision similar to the selection of the number of series terms. We will discuss model selection methods below.

12.5 Uniform Approximations

A good series approximation $m_K(\mathbf{x})$ will have the property that it gets close to the true CEF $m(\mathbf{x})$ as the complexity K increases. Formal statements can be derived from the theory of functional analysis.

An elegant and famous theorem is the **Stone-Weierstrass theorem**, (Weierstrass, 1885, Stone 1937, 1948) which states that any continuous function can be arbitrarily uniformly well approximated by a polynomial of sufficiently high order. Specifically, the theorem states that for $\mathbf{x} \in \mathbb{R}^d$, if $m(\mathbf{x})$ is continuous on a compact set \mathcal{X} , then for any $\varepsilon > 0$ there exists a polynomial $m_K(\mathbf{x})$ of some order K which is uniformly within ε of $m(\mathbf{x})$:

$$\sup_{\mathbf{x} \in \mathcal{X}} |m_K(\mathbf{x}) - m(\mathbf{x})| \leq \varepsilon. \quad (12.2)$$

Thus the true unknown $m(\mathbf{x})$ can be arbitrarily well approximated by selecting a suitable polynomial.

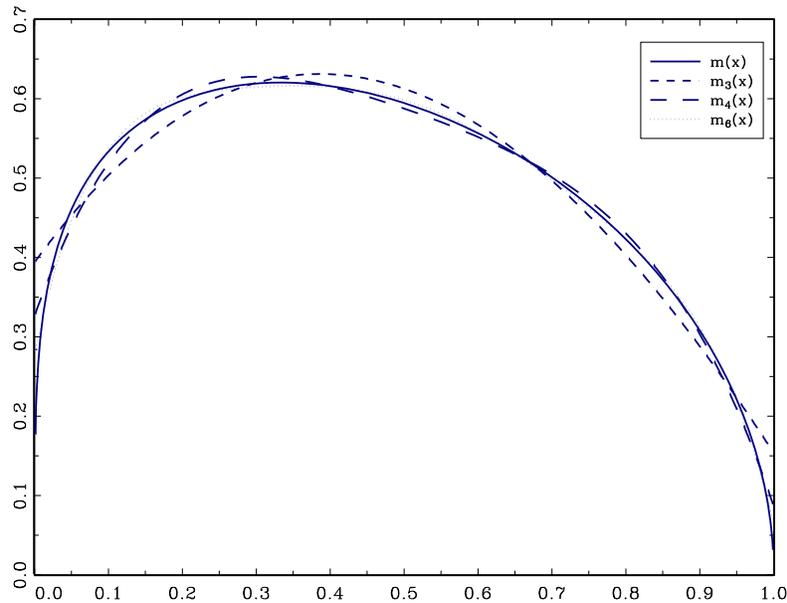


Figure 12.1: True CEF and Best Approximations

The result (12.2) can be strengthened. In particular, if the s^{th} derivative of $m(\mathbf{x})$ is continuous then the uniform approximation error satisfies

$$\sup_{\mathbf{x} \in \mathcal{X}} |m_K(\mathbf{x}) - m(\mathbf{x})| = O(K^{-\alpha}) \quad (12.3)$$

as $K \rightarrow \infty$ where $\alpha = s/d$. This result is more useful than (12.2) because it gives a rate at which the approximation $m_K(\mathbf{x})$ approaches $m(\mathbf{x})$ as K increases.

Both (12.2) and (12.3) hold for spline approximations as well.

Intuitively, the number of derivatives s indexes the smoothness of the function $m(\mathbf{x})$. (12.3) says that the best rate at which a polynomial or spline approximates the CEF $m(\mathbf{x})$ depends on the underlying smoothness of $m(\mathbf{x})$. The more smooth is $m(\mathbf{x})$, the fewer series terms (polynomial order or spline knots) are needed to obtain a good approximation.

To illustrate polynomial approximation, Figure 12.1 displays the CEF $m(x) = x^{1/4}(1-x)^{1/2}$ on $x \in [0, 1]$. In addition, the best approximations using polynomials of order $K = 3$, $K = 4$, and $K = 6$ are displayed. You can see how the approximation with $K = 3$ is fairly crude, but improves with $K = 4$ and especially $K = 6$. Approximations obtained with cubic splines are quite similar so not displayed.

As a series approximation can be written as $m_K(\mathbf{x}) = \mathbf{z}_K(\mathbf{x})'\boldsymbol{\beta}_K$ as in (12.1), then the coefficient of the best uniform approximation (12.3) is then

$$\boldsymbol{\beta}_K^* = \operatorname{argmin}_{\boldsymbol{\beta}_K} \sup_{\mathbf{x} \in \mathcal{X}} |\mathbf{z}_K(\mathbf{x})'\boldsymbol{\beta}_K - m(\mathbf{x})|. \quad (12.4)$$

The approximation error is

$$r_K^*(\mathbf{x}) = m(\mathbf{x}) - \mathbf{z}_K(\mathbf{x})'\boldsymbol{\beta}_K^*.$$

We can write this as

$$m(\mathbf{x}) = \mathbf{z}_K(\mathbf{x})'\boldsymbol{\beta}_K^* + r_K^*(\mathbf{x}) \quad (12.5)$$

to emphasize that the true conditional mean can be written as the linear approximation plus error. A useful consequence of equation (12.3) is

$$\sup_{\mathbf{x} \in \mathcal{X}} |r_K^*(\mathbf{x})| \leq O(K^{-\alpha}). \quad (12.6)$$

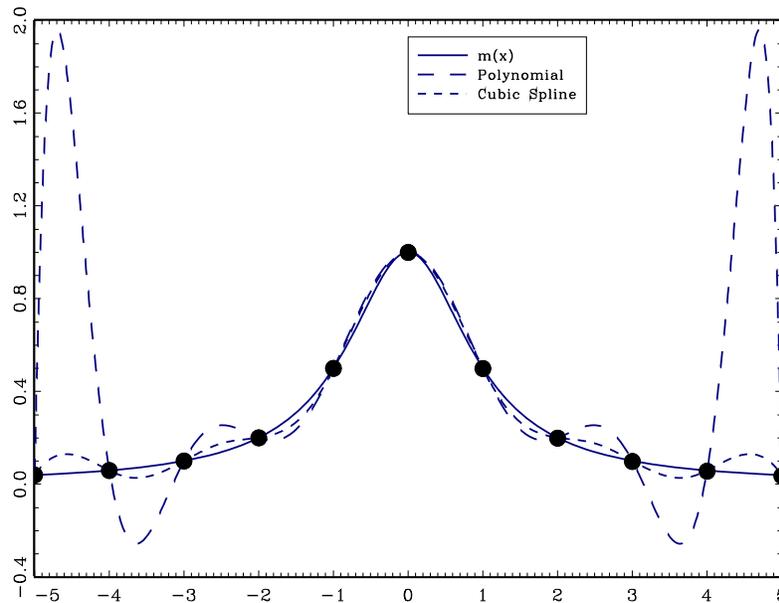


Figure 12.2: True CEF, polynomial interpolation, and spline interpolation

12.6 Runge's Phenomenon

Despite the excellent approximation implied by the Stone-Weierstrass theorem, polynomials have the troubling disadvantage that they are very poor at simple interpolation. The problem is known as **Runge's phenomenon**, and is illustrated in Figure 12.2. The solid line is the CEF $m(x) = (1 + x^2)^{-1}$ displayed on $[-5, 5]$. The circles display the function at the $K = 11$ integers in this interval. The long dashes display the 10th order polynomial fit through these points. Notice that the polynomial approximation is erratic and far from the smooth CEF. This discrepancy gets worse as the number of evaluation points increases, as Runge (1901) showed that the discrepancy increases to infinity with K .

In contrast, splines do not exhibit Runge's phenomenon. In Figure 12.2 the short dashes display a cubic spline with seven knots fit through the same points as the polynomial. While the fitted spline displays some oscillation relative to the true CEF, they are relatively moderate.

Because of Runge's phenomenon, high-order polynomials are not used for interpolation, and are not popular choices for high-order series approximations. Instead, splines are widely used.

12.7 Approximating Regression

For each observation i we observe (y_i, \mathbf{x}_i) and then construct the regressor vector $\mathbf{z}_{Ki} = \mathbf{z}_K(\mathbf{x}_i)$ using the series transformations. Stacking the observations in the matrices \mathbf{y} and \mathbf{Z}_K , the least squares estimate of the coefficient β_K in the series approximation $\mathbf{z}_K(\mathbf{x})'\beta_K$ is

$$\hat{\beta}_K = (\mathbf{Z}'_K \mathbf{Z}_K)^{-1} \mathbf{Z}'_K \mathbf{y},$$

and the least squares estimate of the regression function is

$$\hat{m}_K(\mathbf{x}) = \mathbf{z}_K(\mathbf{x})'\hat{\beta}_K. \quad (12.7)$$

As we learned in Chapter 2, the least-squares coefficient is estimating the best linear predictor of y_i given \mathbf{z}_{Ki} . This is

$$\beta_K = \mathbb{E}(\mathbf{z}_{Ki} \mathbf{z}'_{Ki})^{-1} \mathbb{E}(\mathbf{z}_{Ki} y_i).$$

Given this coefficient, the series approximation is $\mathbf{z}_K(\mathbf{x})'\boldsymbol{\beta}_K$ with approximation error

$$r_K(\mathbf{x}) = m(\mathbf{x}) - \mathbf{z}_K(\mathbf{x})'\boldsymbol{\beta}_K. \quad (12.8)$$

The true CEF equation for y_i is

$$y_i = m(\mathbf{x}_i) + e_i \quad (12.9)$$

with e_i the CEF error. Defining $r_{Ki} = r_K(\mathbf{x}_i)$, we find

$$y_i = \mathbf{z}'_{Ki}\boldsymbol{\beta}_K + e_{Ki}$$

where the equation error is

$$e_{Ki} = r_{Ki} + e_i.$$

Observe that the error e_{Ki} includes the approximation error and thus does not have the properties of a CEF error.

In matrix notation we can write these equations as

$$\begin{aligned} \mathbf{y} &= \mathbf{Z}_K\boldsymbol{\beta}_K + \mathbf{r}_K + \mathbf{e} \\ &= \mathbf{Z}_K\boldsymbol{\beta}_K + \mathbf{e}_K. \end{aligned} \quad (12.10)$$

We now impose some regularity conditions on the regression model to facilitate the theory. Define the $K \times K$ expected design matrix

$$\mathbf{Q}_K = \mathbb{E}(\mathbf{z}_{Ki}\mathbf{z}'_{Ki}),$$

let \mathcal{X} denote the support of \mathbf{x}_i , and define the largest normalized length of the regressor vector in the support of \mathbf{x}_i

$$\zeta_K = \sup_{\mathbf{x} \in \mathcal{X}} (\mathbf{z}_K(\mathbf{x})'\mathbf{Q}_K^{-1}\mathbf{z}_K(\mathbf{x}))^{1/2}. \quad (12.11)$$

ζ_K will increase with K . For example, if the support of the variables $\mathbf{z}_K(\mathbf{x}_i)$ is the unit cube $[0, 1]^K$, then you can compute that $\zeta_K = \sqrt{K}$. As discussed in Newey (1997) and Li and Racine (2007, Corollary 15.1) if the support of \mathbf{x}_i is compact then $\zeta_K = O(K)$ for polynomials and $\zeta_K = O(K^{1/2})$ for splines.

Assumption 12.7.1

1. For some $\alpha > 0$ the series approximation satisfies (12.3).
2. $\mathbb{E}(e_i^2 \mid \mathbf{x}_i) \leq \bar{\sigma}^2 < \infty$.
3. $\lambda_{\min}(\mathbf{Q}_K) \geq \underline{\lambda} > 0$.
4. $K = K(n)$ is a function of n which satisfies $K/n \rightarrow 0$ and $\zeta_K^2 K/n \rightarrow 0$ as $n \rightarrow \infty$.

Assumptions 12.7.1.1 through 12.7.1.3 concern properties of the regression model. Assumption 12.7.1.1 holds with $\alpha = s/d$ if \mathcal{X} is compact and the s 'th derivative of $m(\mathbf{x})$ is continuous. Assumption 12.7.1.2 allows for conditional heteroskedasticity, but requires the conditional variance to be bounded. Assumption 12.7.1.3 excludes near-singular designs. Since estimates of the conditional mean are unchanged if we replace \mathbf{z}_{Ki} with $\mathbf{z}^*_{Ki} = \mathbf{B}_K\mathbf{z}_{Ki}$ for any non-singular \mathbf{B}_K , Assumption 12.7.1.3 can be viewed as holding after transformation by an appropriate non-singular \mathbf{B}_K .

Assumption 12.7.1.4 concerns the choice of the number of series terms, which is under the control of the user. It specifies that K can increase with sample size, but at a controlled rate of growth. Since $\zeta_K = O(K)$ for polynomials and $\zeta_K = O(K^{1/2})$ for splines, Assumption 12.7.1.4 is satisfied if $K^3/n \rightarrow 0$ for polynomials and $K^2/n \rightarrow 0$ for splines. This means that while the number of series terms K can increase with the sample size, K must increase at a much slower rate.

In Section 12.5 we introduced the best uniform approximation, and in this section we introduced the best linear predictor. What is the relationship? They may be similar in practice, but they are not the same and we should be careful to maintain the distinction. Note that from (12.5) we can write $m(\mathbf{x}_i) = \mathbf{z}'_{Ki}\boldsymbol{\beta}_K^* + r_{Ki}^*$ where $r_{Ki}^* = r_K^*(\mathbf{x}_i)$ satisfies $\sup_i |r_{Ki}^*| = O(K^{-\alpha})$ from (12.6). Then the best linear predictor equals

$$\begin{aligned}\boldsymbol{\beta}_K &= \mathbb{E}(\mathbf{z}_{Ki}\mathbf{z}'_{Ki})^{-1} \mathbb{E}(\mathbf{z}_{Ki}y_i) \\ &= \mathbb{E}(\mathbf{z}_{Ki}\mathbf{z}'_{Ki})^{-1} \mathbb{E}(\mathbf{z}_{Ki}m(\mathbf{x}_i)) \\ &= \mathbb{E}(\mathbf{z}_{Ki}\mathbf{z}'_{Ki})^{-1} \mathbb{E}(\mathbf{z}_{Ki}(\mathbf{z}'_{Ki}\boldsymbol{\beta}_K^* + r_{Ki}^*)) \\ &= \boldsymbol{\beta}_K^* + \mathbb{E}(\mathbf{z}_{Ki}\mathbf{z}'_{Ki})^{-1} \mathbb{E}(\mathbf{z}_{Ki}r_{Ki}^*).\end{aligned}$$

Thus the difference between the two approximations is

$$\begin{aligned}r_K(\mathbf{x}) - r_K^*(\mathbf{x}) &= \mathbf{z}_K(\mathbf{x})'(\boldsymbol{\beta}_K^* - \boldsymbol{\beta}_K) \\ &= \mathbf{z}_K(\mathbf{x})'\mathbb{E}(\mathbf{z}_{Ki}\mathbf{z}'_{Ki})^{-1} \mathbb{E}(\mathbf{z}_{Ki}r_{Ki}^*).\end{aligned}\tag{12.12}$$

Observe that by the properties of projection

$$\mathbb{E}(r_{Ki}^{*2}) - \mathbb{E}(\mathbf{r}_{Ki}^*\mathbf{z}_{Ki})'\mathbb{E}(\mathbf{z}_{Ki}\mathbf{z}'_{Ki})^{-1} \mathbb{E}(\mathbf{z}_{Ki}r_{Ki}^*) \geq 0\tag{12.13}$$

and by (12.6)

$$\mathbb{E}(r_{Ki}^{*2}) = \int r_K^*(\mathbf{x})^2 f_x(\mathbf{x})d\mathbf{x} \leq O(K^{-2\alpha}).\tag{12.14}$$

Then applying the Schwarz inequality to (12.12), Definition (12.11), (12.13) and (12.14), we find

$$\begin{aligned}|r_K(\mathbf{x}) - r_K^*(\mathbf{x})| &\leq \left(\mathbf{z}_K(\mathbf{x})'\mathbb{E}(\mathbf{z}_{Ki}\mathbf{z}'_{Ki})^{-1} \mathbf{z}_K(\mathbf{x})\right)^{1/2} \\ &\quad \left(\mathbb{E}(\mathbf{r}_{Ki}^*\mathbf{z}_{Ki})'\mathbb{E}(\mathbf{z}_{Ki}\mathbf{z}'_{Ki})^{-1} \mathbb{E}(\mathbf{z}_{Ki}r_{Ki}^*)\right)^{1/2} \\ &\leq O(\zeta_K K^{-\alpha}).\end{aligned}\tag{12.15}$$

It follows that the best linear predictor approximation error satisfies

$$\sup_{\mathbf{x} \in \mathcal{X}} |r_K(\mathbf{x})| \leq O(\zeta_K K^{-\alpha}).\tag{12.16}$$

The bound (12.16) is probably not the best possible, but it shows that the best linear predictor satisfies a uniform approximation bound. Relative to (12.6), the rate is slower by the factor ζ_K . The bound (12.16) term is $o(1)$ as $K \rightarrow \infty$ if $\zeta_K K^{-\alpha} \rightarrow 0$. A sufficient condition is that $\alpha > 1$ ($s > d$) for polynomials and $\alpha > 1/2$ ($s > d/2$) for splines, where $d = \dim(\mathbf{x})$ and s is the number of continuous derivatives of $m(\mathbf{x})$.

It is also useful to observe that since $\boldsymbol{\beta}_K$ is the best linear approximation to $m(\mathbf{x}_i)$ in mean-square (see Section 2.24), then

$$\begin{aligned}\mathbb{E}r_{Ki}^2 &= \mathbb{E}(m(\mathbf{x}_i) - \mathbf{z}'_{Ki}\boldsymbol{\beta}_K)^2 \\ &\leq \mathbb{E}(m(\mathbf{x}_i) - \mathbf{z}'_{Ki}\boldsymbol{\beta}_K^*)^2 \\ &\leq O(K^{-2\alpha})\end{aligned}\tag{12.17}$$

the final inequality by (12.14).

12.8 Residuals and Regression Fit

The fitted regression at $\mathbf{x} = \mathbf{x}_i$ is $\hat{m}_K(\mathbf{x}_i) = \mathbf{z}'_{Ki} \hat{\boldsymbol{\beta}}_K$ and the fitted residual is

$$\hat{e}_{iK} = y_i - \hat{m}_K(\mathbf{x}_i).$$

The leave-one-out prediction errors are

$$\begin{aligned} \tilde{e}_{iK} &= y_i - \hat{m}_{K,-i}(\mathbf{x}_i) \\ &= y_i - \mathbf{z}'_{Ki} \hat{\boldsymbol{\beta}}_{K,-i} \end{aligned}$$

where $\hat{\boldsymbol{\beta}}_{K,-i}$ is the least-squares coefficient with the i 'th observation omitted. Using (3.38) we can also write

$$\tilde{e}_{iK} = \hat{e}_{iK}(1 - h_{Kii})^{-1}$$

where $h_{Kii} = \mathbf{z}'_{Ki} (\mathbf{Z}'_K \mathbf{Z}_K)^{-1} \mathbf{z}_{Ki}$.

As for kernel regression, the prediction errors \tilde{e}_{iK} are better estimates of the errors than the fitted residuals \hat{e}_{iK} , as they do not have the tendency to “over-fit” when the number of series terms is large.

To assess the fit of the nonparametric regression, the estimate of the mean-square prediction error is

$$\tilde{\sigma}_K^2 = \frac{1}{n} \sum_{i=1}^n \tilde{e}_{iK}^2 = \frac{1}{n} \sum_{i=1}^n \hat{e}_{iK}^2 (1 - h_{Kii})^{-2}$$

and the prediction R^2 is

$$\tilde{R}_K^2 = 1 - \frac{\sum_{i=1}^n \tilde{e}_{iK}^2}{\sum_{i=1}^n (y_i - \bar{y})^2}.$$

12.9 Cross-Validation Model Selection

The cross-validation criterion for selection of the number of series terms is the MSPE

$$CV(K) = \tilde{\sigma}_K^2 = \frac{1}{n} \sum_{i=1}^n \tilde{e}_{iK}^2 (1 - h_{Kii})^{-2}.$$

By selecting the series terms to minimize $CV(K)$, or equivalently maximize \tilde{R}_K^2 , we have a data-dependent rule which is designed to produce estimates with low integrated mean-squared error (IMSE) and mean-squared forecast error (MSFE). As shown in Theorem 11.6.1, $CV(K)$ is an approximately unbiased estimated of the MSFE and IMSE, so finding the model which produces the smallest value of $CV(K)$ is a good indicator that the estimated model has small MSFE and IMSE. The proof of the result is the same for all nonparametric estimators (series as well as kernels) so does not need to be repeated here.

As a practical matter, an estimator corresponds to a set of regressors \mathbf{z}_{Ki} , that is, a set of transformations of the original variables \mathbf{x}_i . For each set of regressions, the regression is estimated and $CV(K)$ calculated, and the estimator is selected which has the smallest value of $CV(K)$. If there are p ordered regressors, then there are p possible estimators. Typically, this calculation is simple even if p is large. However, if the p regressors are unordered (and this is typical) then there are 2^p possible subsets of conceivable models. If p is even moderately large, 2^p can be immensely large so brute-force computation of all models may be computationally demanding.

12.10 Convergence in Mean-Square

The series estimate $\widehat{\beta}_K$ are indexed by K . The point of nonparametric estimation is to let K be flexible so as to incorporate greater complexity when the data are sufficiently informative. This means that K will typically be increasing with sample size n . This invalidates conventional asymptotic distribution theory. However, we can develop extensions which use appropriate matrix norms, and by focusing on real-valued functions of the parameters including the estimated regression function itself.

The asymptotic theory we present in this and the next several sections is largely taken from Newey (1997).

Our first main result shows that the least-squares estimate converges to β_K in mean-square distance.

Theorem 12.10.1 *Under Assumption 12.7.1, as $n \rightarrow \infty$,*

$$\left(\widehat{\beta}_K - \beta_K\right)' \mathbf{Q}_K \left(\widehat{\beta}_K - \beta_K\right) = O_p\left(\frac{K}{n}\right) + o_p\left(K^{-2\alpha}\right) \quad (12.18)$$

The proof of Theorem 12.10.1 is rather technical and deferred to Section 12.16.

The rate of convergence in (12.18) has two terms. The $O_p(K/n)$ term is due to estimation variance. Note in contrast that the corresponding rate would be $O_p(1/n)$ in the parametric case. The difference is that in the parametric case we assume that the number of regressors K is fixed as n increases, while in the nonparametric case we allow the number of regressors K to be flexible. As K increases, the estimation variance increases. The $o_p(K^{-2\alpha})$ term in (12.18) is due to the series approximation error.

Using Theorem 12.10.1 we can establish the following convergence rate for the estimated regression function.

Theorem 12.10.2 *Under Assumption 12.7.1, as $n \rightarrow \infty$,*

$$\int \left(\widehat{m}_K(\mathbf{x}) - m(\mathbf{x})\right)^2 f_x(\mathbf{x}) d\mathbf{x} = O_p\left(\frac{K}{n}\right) + O_p\left(K^{-2\alpha}\right) \quad (12.19)$$

Theorem 12.10.2 shows that the integrated squared difference between the fitted regression and the true CEF converges in probability to zero if $K \rightarrow \infty$ as $n \rightarrow \infty$. The convergence results of Theorem 12.10.2 show that the number of series terms K involves a trade-off similar to the role of the bandwidth h in kernel regression. Larger K implies smaller approximation error but increased estimation variance.

The optimal rate which minimizes the average squared error in (12.19) is $K = O\left(n^{1/(1+2\alpha)}\right)$, yielding an optimal rate of convergence in (12.19) of $O_p\left(n^{-2\alpha/(1+2\alpha)}\right)$. This rate depends on the unknown smoothness α of the true CEF (the number of derivatives s) and so does not directly suggest a practical rule for determining K . Still, the implication is that when the function being estimated is less smooth (α is small) then it is necessary to use a larger number of series terms K to reduce the bias. In contrast, when the function is more smooth then it is better to use a smaller number of series terms K to reduce the variance.

To establish (12.19), using (12.7) and (12.8) we can write

$$\widehat{m}_K(\mathbf{x}) - m(\mathbf{x}) = \mathbf{z}_K(\mathbf{x})' \left(\widehat{\beta}_K - \beta_K\right) - r_K(\mathbf{x}). \quad (12.20)$$

Since e_{Ki} are projection errors, they satisfy $\mathbb{E}(\mathbf{z}_{Ki}e_{Ki}) = 0$ and thus $\mathbb{E}(\mathbf{z}_{Ki}r_{Ki}) = 0$. This means $\int \mathbf{z}_K(\mathbf{x})r_K(\mathbf{x})f_x(\mathbf{x})d\mathbf{x} = 0$. Also observe that $\mathbf{Q}_K = \int \mathbf{z}_K(\mathbf{x})\mathbf{z}_K(\mathbf{x})'f_x(\mathbf{x})d\mathbf{x}$ and $\mathbb{E}r_{Ki}^2 = \int r_K(\mathbf{x})^2f_x(\mathbf{x})d\mathbf{x}$. Then

$$\begin{aligned} & \int (\widehat{m}_K(\mathbf{x}) - m(\mathbf{x}))^2 f_x(\mathbf{x})d\mathbf{x} \\ &= (\widehat{\boldsymbol{\beta}}_K - \boldsymbol{\beta}_K)' \mathbf{Q}_K (\widehat{\boldsymbol{\beta}}_K - \boldsymbol{\beta}_K) + \mathbb{E}r_{Ki}^2 \\ &\leq O_p\left(\frac{K}{n}\right) + O_p(K^{-2\alpha}) \end{aligned}$$

by (12.18) and (12.17), establishing (12.19).

12.11 Uniform Convergence

Theorem 12.10.2 established conditions under which $\widehat{m}_K(\mathbf{x})$ is consistent in a squared error norm. It is also of interest to know the rate at which the largest deviation converges to zero. We have the following rate.

Theorem 12.11.1 *Under Assumption 12.7.1, then as $n \rightarrow \infty$,*

$$\sup_{\mathbf{x} \in \mathcal{X}} |\widehat{m}_K(\mathbf{x}) - m(\mathbf{x})| = O_p\left(\sqrt{\frac{\zeta_K^2 K}{n}}\right) + O_p(\zeta_K K^{-\alpha}). \quad (12.21)$$

Relative to Theorem 12.10.2, the error has been increased multiplicatively by ζ_K . This slower convergence rate is a penalty for the stronger uniform convergence, though it is probably not the best possible rate. Examining the bound in (12.21) notice that the first term is $o_p(1)$ under Assumption 12.7.1.4. The second term is $o_p(1)$ if $\zeta_K K^{-\alpha} \rightarrow 0$, which requires that $K \rightarrow \infty$ and that α be sufficiently large. A sufficient condition is that $s > d$ for polynomials and $s > d/2$ for splines, where $d = \dim(\mathbf{x})$ and s is the number of continuous derivatives of $m(\mathbf{x})$. Thus higher dimensional \mathbf{x} require a smoother CEF $m(\mathbf{x})$ to ensure that the series estimate $\widehat{m}_K(\mathbf{x})$ is uniformly consistent.

The convergence (12.21) is straightforward to show using (12.18). Using (12.20), the Triangle Inequality, the Schwarz inequality (A.10), Definition (12.11), (12.18) and (12.16),

$$\begin{aligned} & \sup_{\mathbf{x} \in \mathcal{X}} |\widehat{m}_K(\mathbf{x}) - m(\mathbf{x})| \\ &\leq \sup_{\mathbf{x} \in \mathcal{X}} \left| \mathbf{z}_K(\mathbf{x})' (\widehat{\boldsymbol{\beta}}_K - \boldsymbol{\beta}_K) \right| + \sup_{\mathbf{x} \in \mathcal{X}} |r_K(\mathbf{x})| \\ &\leq \sup_{\mathbf{x} \in \mathcal{X}} (\mathbf{z}_K(\mathbf{x})' \mathbf{Q}_K^{-1} \mathbf{z}_K(\mathbf{x}))^{1/2} \left((\widehat{\boldsymbol{\beta}}_K - \boldsymbol{\beta}_K)' \mathbf{Q}_K (\widehat{\boldsymbol{\beta}}_K - \boldsymbol{\beta}_K) \right)^{1/2} \\ &\quad + O(\zeta_K K^{-\alpha}) \\ &\leq \zeta_K \left(O_p\left(\frac{K}{n}\right) + O_p(K^{-2\alpha}) \right)^{1/2} + O(\zeta_K K^{-\alpha}), \\ &= O_p\left(\sqrt{\frac{\zeta_K^2 K}{n}}\right) + O_p(\zeta_K K^{-\alpha}). \end{aligned} \quad (12.22)$$

This is (12.21).

12.12 Asymptotic Normality

One advantage of series methods is that the estimators are (in finite samples) equivalent to parametric estimators, so it is easy to calculate covariance matrix estimates. We now show that we can also justify normal asymptotic approximations.

The theory we present in this section will apply to any linear function of the regression function. That is, we allow the parameter of interest to be any non-trivial real-valued linear function of the entire regression function $m(\cdot)$

$$\theta = a(m).$$

This includes the regression function $m(\mathbf{x})$ at a given point \mathbf{x} , derivatives of $m(\mathbf{x})$, and integrals over $m(\mathbf{x})$. Given $\hat{m}_K(\mathbf{x}) = \mathbf{z}_K(\mathbf{x})'\hat{\boldsymbol{\beta}}_K$ as an estimator for $m(\mathbf{x})$, the estimator for θ is

$$\hat{\theta}_K = a(\hat{m}_K) = \mathbf{a}'_K \hat{\boldsymbol{\beta}}_K$$

for some $K \times 1$ vector of constants $\mathbf{a}_K \neq \mathbf{0}$. (The relationship $a(\hat{m}_K) = \mathbf{a}'_K \hat{\boldsymbol{\beta}}_K$ follows since a is linear in m and \hat{m}_K is linear in $\hat{\boldsymbol{\beta}}_K$.)

If K were fixed as $n \rightarrow \infty$, then by standard asymptotic theory we would expect $\hat{\theta}_K$ to be asymptotically normal with variance

$$v_K = \mathbf{a}'_K \mathbf{Q}_K^{-1} \boldsymbol{\Omega}_K \mathbf{Q}_K^{-1} \mathbf{a}_K$$

where

$$\boldsymbol{\Omega}_K = \mathbb{E}(\mathbf{z}_{Ki} \mathbf{z}'_{Ki} e_{Ki}^2).$$

The standard justification, however, is not valid in the nonparametric case, in part because v_K may diverge as $K \rightarrow \infty$, and in part due to the finite sample bias due to the approximation error. Therefore a new theory is required. Interestingly, it turns out that in the nonparametric case $\hat{\theta}_K$ is still asymptotically normal, and v_K is still the appropriate variance for $\hat{\theta}_K$. The proof is different than the parametric case as the dimensions of the matrices are increasing with K , and we need to be attentive to the estimator's bias due to the series approximation.

Theorem 12.12.1 *Under Assumption 12.7.1, if in addition $\mathbb{E}(e_i^4 | \mathbf{x}_i) \leq \kappa_4 < \infty$, $\mathbb{E}(e_i^2 | \mathbf{x}_i) \geq \underline{\sigma}^2 > 0$, and $\zeta_K K^{-\alpha} = O(1)$, then as $n \rightarrow \infty$,*

$$\frac{\sqrt{n} \left(\hat{\theta}_K - \theta + a(r_K) \right)}{v_K^{1/2}} \xrightarrow{d} N(0, 1) \quad (12.23)$$

The proof of Theorem 12.12.1 can be found in Section 12.16.

Theorem 12.12.1 shows that the estimator $\hat{\theta}_K$ is approximately normal with bias $-a(r_K)$ and variance v_K/n . The variance is the same as in the parametric case, but the asymptotic distribution contains an asymptotic bias, similar as is found in kernel regression. We discuss the bias in more detail below.

Notice that Theorem 12.12.1 requires $\zeta_K K^{-\alpha} = O(1)$, which is similar to that found in Theorem 12.11.1 to establish uniform convergence. The bound $\zeta_K K^{-\alpha} = O(1)$ allows K to be constant with n or to increase with n . However, when K is increasing the bound requires that α be sufficient large so that K^α grows faster than ζ_K . A sufficient condition is that $s = d$ for polynomials and $s = d/2$ for splines. The fact that the condition allows for K to be constant means that Theorem 12.12.1 includes parametric least-squares as a special case with explicit attention to estimation bias.

One useful message from Theorem 12.12.1 is that the classic variance formula v_K for $\hat{\theta}_K$ still applies for series regression. Indeed, we can estimate the asymptotic variance using the standard White formula

$$\begin{aligned}\hat{v}_K &= \mathbf{a}'_K \hat{\mathbf{Q}}_K^{-1} \hat{\mathbf{\Omega}}_K \hat{\mathbf{Q}}_K^{-1} \mathbf{a}_K \\ \hat{\mathbf{\Omega}}_K &= \frac{1}{n} \sum_{i=1}^n \mathbf{z}_{Ki} \mathbf{z}'_{Ki} \hat{e}_{iK}^2 \\ \hat{\mathbf{Q}}_K &= \frac{1}{n} \sum_{i=1}^n \mathbf{z}_{Ki} \mathbf{z}'_{Ki}.\end{aligned}$$

Hence a standard error for $\hat{\theta}_K$ is

$$\hat{s}(\theta_K) = \sqrt{\frac{1}{n} \mathbf{a}'_K \hat{\mathbf{Q}}_K^{-1} \hat{\mathbf{\Omega}}_K \hat{\mathbf{Q}}_K^{-1} \mathbf{a}_K}.$$

It can be shown (Newey, 1997) that $\hat{v}_K/v_K \xrightarrow{p} 1$ as $n \rightarrow \infty$ and thus the distribution in (12.23) is unchanged if v_K is replaced with \hat{v}_K .

Theorem 12.12.1 shows that the estimator $\hat{\theta}_K$ has a bias term $a(r_K)$. What is this? It is the same transformation of the function $r_K(\mathbf{x})$ as $\theta = a(m)$ is of the regression function $m(\mathbf{x})$. For example, if $\theta = m(\mathbf{x})$ is the regression at a fixed point \mathbf{x} , then $a(r_K) = r_K(\mathbf{x})$, the approximation error at the same point. If $\theta = \frac{d}{dx}m(x)$ is the regression derivative, then $a(r_K) = \frac{d}{dx}r_K(\mathbf{x})$ is the derivative of the approximation error.

This means that the bias in the estimator $\hat{\theta}_K$ for θ shown in Theorem 12.12.1 is simply the approximation error, transformed by the functional of interest. If we are estimating the regression function then the bias is the error in approximating the regression function; if we are estimating the regression derivative then the bias is the error in the derivative in the approximation error for the regression function.

12.13 Asymptotic Normality with Undersmoothing

An unpleasant aspect about Theorem 12.12.1 is the bias term. An interesting trick is that this bias term can be made asymptotically negligible if we assume that K increases with n at a sufficiently fast rate.

Theorem 12.13.1 *Under Assumption 12.7.1, if in addition $\mathbb{E}(e_i^4|\mathbf{x}_i) \leq \kappa_4 < \infty$, $\mathbb{E}(e_i^2|\mathbf{x}_i) \geq \underline{\sigma}^2 > 0$, $a(r_K^*) \leq O(K^{-\alpha})$, $nK^{-2\alpha} \rightarrow 0$, and $\mathbf{a}'_K \mathbf{Q}_K^{-1} \mathbf{a}_K$ is bounded away from zero, then*

$$\frac{\sqrt{n}(\hat{\theta}_K - \theta)}{v_K^{1/2}} \xrightarrow{d} \text{N}(0, 1). \tag{12.24}$$

The condition $a(r_K^*) \leq O(K^{-\alpha})$ states that the function of interest (for example, the regression function, its derivative, or its integral) applied to the uniform approximation error converges to zero as the number of terms K in the series approximation increases. If $a(m) = m(\mathbf{x})$ then this condition holds by (12.6).

The condition that $\mathbf{a}'_K \mathbf{Q}_K^{-1} \mathbf{a}_K$ is bounded away from zero is simply a technical requirement to exclude degeneracy.

The critical condition is the assumption that $nK^{-2\alpha} \rightarrow 0$. This requires that $K \rightarrow \infty$ at a rate *faster* than $n^{1/2\alpha}$. This is a troubling condition. The optimal rate for estimation of $m(\mathbf{x})$ is $K = O(n^{1/(1+2\alpha)})$. If we set $K = n^{1/(1+2\alpha)}$ by this rule then $nK^{-2\alpha} = n^{1/(1+2\alpha)} \rightarrow \infty$, not zero. Thus this assumption is equivalent to assuming that K is much larger than optimal. The reason why this trick works (that is, why the bias is negligible) is that by increasing K , the asymptotic bias decreases and the asymptotic variance increases and thus the variance dominates. Because K is larger than optimal, we typically say that $\hat{m}_K(\mathbf{x})$ is *undersmoothed* relative to the optimal series estimator.

Many authors like to focus their asymptotic theory on the assumptions in Theorem 12.13.1, as the distribution (12.24) appears cleaner. However, it is a poor use of asymptotic theory. There are three problems with the assumption $nK^{-2\alpha} \rightarrow 0$ and the approximation (12.24). First, it says that if we intentionally pick K to be larger than optimal, we can increase the estimation variance relative to the bias so the variance will dominate the bias. But why would we want to intentionally use an estimator which is sub-optimal? Second, the assumption $nK^{-2\alpha} \rightarrow 0$ does not eliminate the asymptotic bias, it only makes it of lower order than the variance. So the approximation (12.24) is technically valid, but the missing asymptotic bias term is just slightly smaller in asymptotic order, and thus still relevant in finite samples. Third, the condition $nK^{-2\alpha} \rightarrow 0$ is just an assumption, it has nothing to do with actual empirical practice. Thus the difference between (12.23) and (12.24) is in the assumptions, not in the actual reality or in the actual empirical practice. Eliminating a nuisance (the asymptotic bias) through an assumption is a trick, not a substantive use of theory. My strong view is that the result (12.23) is more informative than (12.24). It shows that the asymptotic distribution is normal but has a non-trivial finite sample bias.

12.14 Regression Estimation

A special yet important example of a linear estimator of the regression function is the regression function at a fixed point \mathbf{x} . In the notation of the previous section, $a(m) = m(\mathbf{x})$ and $\mathbf{a}_K = \mathbf{z}_K(\mathbf{x})$. The series estimator of $m(\mathbf{x})$ is $\hat{\theta}_K = \hat{m}_K(\mathbf{x}) = \mathbf{z}_K(\mathbf{x})' \hat{\boldsymbol{\beta}}_K$. As this is a key problem of interest, we restate the asymptotic results of Theorems 12.12.1 and 12.13.1 for this estimator.

Theorem 12.14.1 *Under Assumption 12.7.1, if in addition $\mathbb{E}(e_i^4 | \mathbf{x}_i) \leq \kappa_4 < \infty$, $\mathbb{E}(e_i^2 | \mathbf{x}_i) \geq \underline{\sigma}^2 > 0$, and $\zeta_K K^{-\alpha} = O(1)$, then as $n \rightarrow \infty$,*

$$\frac{\sqrt{n}(\hat{m}_K(\mathbf{x}) - m(\mathbf{x}) + \mathbf{r}_K(\mathbf{x}))}{v_K^{1/2}(\mathbf{x})} \xrightarrow{d} \text{N}(0, 1) \quad (12.25)$$

where

$$v_K(\mathbf{x}) = \mathbf{z}_K(\mathbf{x})' \mathbf{Q}_K^{-1} \boldsymbol{\Omega}_K \mathbf{Q}_K^{-1} \mathbf{z}_K(\mathbf{x}).$$

If $\zeta_K K^{-\alpha} = O(1)$ is replaced by $nK^{-2\alpha} \rightarrow 0$, and $\mathbf{z}_K(\mathbf{x})' \mathbf{Q}_K^{-1} \mathbf{z}_K(\mathbf{x})$ is bounded away from zero, then

$$\frac{\sqrt{n}(\hat{m}_K(\mathbf{x}) - m(\mathbf{x}))}{v_K^{1/2}(\mathbf{x})} \xrightarrow{d} \text{N}(0, 1) \quad (12.26)$$

There are two important features about the asymptotic distribution (12.25).

First, as mentioned in the previous section, it shows how to construct asymptotic standard errors for the CEF $m(\mathbf{x})$. These are

$$\hat{s}(\mathbf{x}) = \sqrt{\frac{1}{n} \mathbf{z}_K(\mathbf{x})' \hat{\mathbf{Q}}_K^{-1} \hat{\boldsymbol{\Omega}}_K \hat{\mathbf{Q}}_K^{-1} \mathbf{z}_K(\mathbf{x})}.$$

Second, (12.25) shows that the estimator has the asymptotic bias component $\mathbf{r}_K(\mathbf{x})$. This is due to the fact that the finite order series is an approximation to the unknown CEF $m(\mathbf{x})$, and this results in finite sample bias.

The asymptotic distribution (12.26) shows that the bias term is negligible if K diverges fast enough so that $nK^{-2\alpha} \rightarrow 0$. As discussed in the previous section, this means that K is larger than optimal.

The assumption that $\mathbf{z}_K(\mathbf{x})' \mathbf{Q}_K^{-1} \mathbf{z}_K(\mathbf{x})$ is bounded away from zero is a technical condition to exclude degenerate cases, and is automatically satisfied if $\mathbf{z}_K(\mathbf{x})$ includes an intercept.

Plots of the CEF estimate $\hat{m}_K(\mathbf{x})$ can be accompanied by 95% confidence intervals $\hat{m}_K(\mathbf{x}) \pm 2\hat{s}(\mathbf{x})$. As we discussed in the chapter on kernel regression, this can be viewed as a confidence interval for the pseudo-true CEF $m_K^*(\mathbf{x}) = m(\mathbf{x}) - \mathbf{r}_K(\mathbf{x})$, not for the true $m(\mathbf{x})$. As for kernel regression, the difference is the unavoidable consequence of nonparametric estimation.

12.15 Kernel Versus Series Regression

In this and the previous chapter we have presented two distinct methods of nonparametric regression based on kernel methods and series methods. Which should be used in practice? Both methods have advantages and disadvantages and there is no clear overall winner.

First, while the asymptotic theory of the two estimators appear quite different, they are actually rather closely related. When the regression function $m(\mathbf{x})$ is twice differentiable ($s = 2$) then the rate of convergence of both the MSE of the kernel regression estimator with optimal bandwidth h and the series estimator with optimal K is $n^{-2/(d+4)}$. There is no difference. If the regression function is smoother than twice differentiable ($s > 2$) then the rate of the convergence of the series estimator improves. This may appear to be an advantage for series methods, but kernel regression can also take advantage of the higher smoothness by using so-called higher-order kernels or local polynomial regression, so perhaps this advantage is not too large.

Both estimators are asymptotically normal and have straightforward asymptotic standard error formulae. The series estimators are a bit more convenient for this purpose, as classic parametric standard error formula work without amendment.

An advantage of kernel methods is that their distributional theory is easier to derive. The theory is all based on local averages which is relatively straightforward. In contrast, series theory is more challenging, dealing with increasing parameter spaces. An important difference in the theory is that for kernel estimators we have explicit representations for the bias while we only have rates for series methods. This means that plug-in methods can be used for bandwidth selection in kernel regression. However, typically we rely on cross-validation, which is equally applicable in both kernel and series regression.

Kernel methods are also relatively easy to implement when the dimension d is large. There is not a major change in the methodology as d increases. In contrast, series methods become quite cumbersome as d increases as the number of cross-terms increases exponentially.

A major advantage of series methods is that it has inherently a high degree of flexibility, and the user is able to implement shape restrictions quite easily. For example, in series estimation it is relatively simple to implement a partial linear CEF, an additively separable CEF, monotonicity, concavity or convexity. These restrictions are harder to implement in kernel regression.

12.16 Technical Proofs

Define $\mathbf{z}_{Ki} = \mathbf{z}_K(\mathbf{x}_i)$ and let $\mathbf{Q}_K^{1/2}$ denote the positive definite square root of \mathbf{Q}_K . As mentioned before Theorem 12.10.1, the regression problem is unchanged if we replace \mathbf{z}_{Ki} with a rotated regressor such as $\mathbf{z}_{Ki}^* = \mathbf{Q}_K^{-1/2} \mathbf{z}_{Ki}$. This is a convenient choice for then $\mathbb{E}(\mathbf{z}_{Ki}^* \mathbf{z}_{Ki}^{*'}) = \mathbf{I}_K$. For notational convenience we will simply write the transformed regressors as \mathbf{z}_{Ki} and set $\mathbf{Q}_K = \mathbf{I}_K$.

We start with some convergence results for the sample design matrix

$$\widehat{\mathbf{Q}}_K = \frac{1}{n} \mathbf{Z}'_K \mathbf{Z}_K = \frac{1}{n} \sum_{i=1}^n \mathbf{z}_{Ki} \mathbf{z}'_{Ki}.$$

Theorem 12.16.1 Under Assumption 12.7.1 and $\mathbf{Q}_K = \mathbf{I}_K$, as $n \rightarrow \infty$,

$$\left\| \widehat{\mathbf{Q}}_K - \mathbf{I}_K \right\| = o_p(1) \quad (12.27)$$

and

$$\lambda_{\min}(\widehat{\mathbf{Q}}_K) \xrightarrow{p} 1. \quad (12.28)$$

Proof. Since

$$\left\| \widehat{\mathbf{Q}}_K - \mathbf{I}_K \right\|^2 = \sum_{j=1}^K \sum_{\ell=1}^K \left(\frac{1}{n} \sum_{i=1}^n (\mathbf{z}_{jKi} \mathbf{z}_{\ell Ki} - \mathbb{E} \mathbf{z}_{jKi} \mathbf{z}_{\ell Ki}) \right)^2$$

then

$$\begin{aligned} \mathbb{E} \left\| \widehat{\mathbf{Q}}_K - \mathbf{I}_K \right\|^2 &= \sum_{j=1}^K \sum_{\ell=1}^K \text{var} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{z}_{jKi} \mathbf{z}_{\ell Ki} \right) \\ &= n^{-1} \sum_{j=1}^K \sum_{\ell=1}^K \text{var} (\mathbf{z}_{jKi} \mathbf{z}_{\ell Ki}) \\ &\leq n^{-1} \mathbb{E} \sum_{j=1}^K \mathbf{z}_{jKi}^2 \sum_{\ell=1}^K \mathbf{z}_{\ell Ki}^2 \\ &= n^{-1} \mathbb{E} (\mathbf{z}'_{Ki} \mathbf{z}_{Ki})^2. \end{aligned} \quad (12.29)$$

Since $\mathbf{z}'_{Ki} \mathbf{z}_{Ki} \leq \zeta_K^2$ by definition (12.11) and using (A.1) we find

$$\mathbb{E} (\mathbf{z}'_{Ki} \mathbf{z}_{Ki}) = \text{tr} (\mathbb{E} \mathbf{z}_{Ki} \mathbf{z}'_{Ki}) = \text{tr} \mathbf{I}_K = K, \quad (12.30)$$

so that

$$\mathbb{E} (\mathbf{z}'_{Ki} \mathbf{z}_{Ki})^2 \leq \zeta_K^2 K \quad (12.31)$$

and hence (12.29) is $o(1)$ under Assumption 12.7.1.4. Theorem 5.11.1 shows that this implies (12.27).

Let $\lambda_1, \lambda_2, \dots, \lambda_K$ be the eigenvalues of $\widehat{\mathbf{Q}}_K - \mathbf{I}_K$ which are real as $\widehat{\mathbf{Q}}_K - \mathbf{I}_K$ is symmetric. Then

$$\left| \lambda_{\min}(\widehat{\mathbf{Q}}_K) - 1 \right| = \left| \lambda_{\min}(\widehat{\mathbf{Q}}_K - \mathbf{I}_K) \right| \leq \left(\sum_{\ell=1}^K \lambda_{\ell}^2 \right)^{1/2} = \left\| \widehat{\mathbf{Q}}_K - \mathbf{I}_K \right\|$$

where the second equality is (A.8). This is $o_p(1)$ by (12.27), establishing (12.28). \blacksquare

Proof of Theorem 12.10.1. As above, assume that the regressors have been transformed so that $\mathbf{Q}_K = \mathbf{I}_K$.

From expression (12.10) we can substitute to find

$$\begin{aligned} \widehat{\boldsymbol{\beta}}_K - \boldsymbol{\beta}_K &= (\mathbf{Z}'_K \mathbf{Z}_K)^{-1} \mathbf{Z}'_K \mathbf{e}_K \\ &= \widehat{\mathbf{Q}}_K^{-1} \left(\frac{1}{n} \mathbf{Z}'_K \mathbf{e}_K \right) \end{aligned} \quad (12.32)$$

Using (12.32) and the Quadratic Inequality (A.14),

$$\begin{aligned}
& \left(\widehat{\boldsymbol{\beta}}_K - \boldsymbol{\beta}_K\right)' \left(\widehat{\boldsymbol{\beta}}_K - \boldsymbol{\beta}_K\right) \\
&= n^{-2} \left(\mathbf{e}'_K \mathbf{Z}_K\right) \widehat{\mathbf{Q}}_K^{-1} \widehat{\mathbf{Q}}_K^{-1} \left(\mathbf{Z}'_K \mathbf{e}_K\right) \\
&\leq \left(\lambda_{\max} \left(\widehat{\mathbf{Q}}_K^{-1}\right)\right)^2 n^{-2} \left(\mathbf{e}'_K \mathbf{Z}_K \mathbf{Z}'_K \mathbf{e}_K\right).
\end{aligned} \tag{12.33}$$

Observe that (12.28) implies

$$\lambda_{\max} \left(\widehat{\mathbf{Q}}_K^{-1}\right) = \left(\lambda_{\max} \left(\widehat{\mathbf{Q}}_K\right)\right)^{-1} = O_p(1). \tag{12.34}$$

Since $e_{Ki} = e_i + r_{Ki}$, and using Assumption 12.7.1.2 and (12.16), then

$$\sup_i \mathbb{E} \left(e_{Ki}^2 | \mathbf{x}_i\right) = \bar{\sigma}^2 + \sup_i r_{Ki}^2 \leq \bar{\sigma}^2 + O\left(\zeta_K^2 K^{-2\alpha}\right). \tag{12.35}$$

As e_{Ki} are projection errors, they satisfy $\mathbb{E}(\mathbf{z}_{Ki} e_{Ki}) = 0$. Since the observations are independent, using (12.30) and (12.35), then

$$\begin{aligned}
n^{-2} \mathbb{E} \left(\mathbf{e}'_K \mathbf{Z}_K \mathbf{Z}'_K \mathbf{e}_K\right) &= n^{-2} \mathbb{E} \left(\sum_{i=1}^n e_{Ki} \mathbf{z}'_{Ki} \sum_{ij=1}^n \mathbf{z}_{Kj} e_{Kj} \right) \\
&= n^{-2} \sum_{i=1}^n \mathbb{E} \left(\mathbf{z}'_{Ki} \mathbf{z}_{Ki} e_{Ki}^2 \right) \\
&\leq n^{-1} \mathbb{E} \left(\mathbf{z}'_{Ki} \mathbf{z}_{Ki} \right) \sup_i \mathbb{E} \left(e_{Ki}^2 | \mathbf{x}_i \right) \\
&\leq \bar{\sigma}^2 \frac{K}{n} + O\left(\frac{\zeta_K^2 K^{1-2\alpha}}{n}\right) \\
&= \bar{\sigma}^2 \frac{K}{n} + o\left(K^{-2\alpha}\right)
\end{aligned} \tag{12.36}$$

since $\zeta_K^2 K/n = o(1)$ by Assumption 12.7.1.4. Theorem 5.11.1 shows that this implies

$$n^{-2} \mathbf{e}'_K \mathbf{Z}_K \mathbf{Z}'_K \mathbf{e}_K = O_p\left(n^{-2}\right) + o_p\left(K^{-2\alpha}\right). \tag{12.37}$$

Together, (12.33), (12.34) and (12.37) imply (12.18). \blacksquare

Proof of Theorem 12.12.1. As above, assume that the regressors have been transformed so that $\mathbf{Q}_K = \mathbf{I}_K$.

Using $m(\mathbf{x}) = \mathbf{z}_K(\mathbf{x})' \boldsymbol{\beta}_K + r_K(\mathbf{x})$ and linearity

$$\begin{aligned}
\theta &= a(m) \\
&= a\left(\mathbf{z}_K(\mathbf{x})' \boldsymbol{\beta}_K\right) + a(r_K) \\
&= \mathbf{a}'_K \boldsymbol{\beta}_K + a(r_K)
\end{aligned}$$

Combined with (12.32) we find

$$\begin{aligned}
\hat{\theta}_K - \theta + a(r_K) &= \mathbf{a}'_K \left(\widehat{\boldsymbol{\beta}}_K - \boldsymbol{\beta}_K\right) \\
&= \frac{1}{n} \mathbf{a}'_K \widehat{\mathbf{Q}}_K^{-1} \mathbf{Z}'_K \mathbf{e}_K
\end{aligned}$$

and thus

$$\begin{aligned} \sqrt{\frac{n}{v_k}} \left(\hat{\theta}_K - \theta_K + a(r_K) \right) &= \sqrt{\frac{n}{v_k}} \mathbf{a}'_K \left(\hat{\beta}_K - \beta_K \right) \\ &= \sqrt{\frac{1}{nv_k}} \mathbf{a}'_K \hat{\mathbf{Q}}_K^{-1} \mathbf{Z}'_K \mathbf{e}_K \\ &= \frac{1}{\sqrt{nv_K}} \mathbf{a}'_K \mathbf{Z}'_K \mathbf{e}_K \end{aligned} \tag{12.38}$$

$$+ \frac{1}{\sqrt{nv_K}} \mathbf{a}'_K \left(\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K \right) \mathbf{Z}'_K \mathbf{e} \tag{12.39}$$

$$+ \frac{1}{\sqrt{nv_K}} \mathbf{a}'_K \left(\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K \right) \mathbf{Z}'_K \mathbf{r}_K. \tag{12.40}$$

where we have used $\mathbf{e}_K = \mathbf{e} + \mathbf{r}_K$. We now take the terms in (12.38)-(12.40) separately.

First, take (12.38). We can write

$$\frac{1}{\sqrt{nv_K}} \mathbf{a}'_K \mathbf{Z}'_K \mathbf{e}_K = \frac{1}{\sqrt{nv_K}} \sum_{i=1}^n \mathbf{a}'_K \mathbf{z}_{Ki} e_{Ki}. \tag{12.41}$$

Observe that $\mathbf{a}'_K \mathbf{z}_{Ki} e_{Ki}$ are independent across i , mean zero, and have variance

$$\mathbb{E} \left(\mathbf{a}'_K \mathbf{z}_{Ki} e_{Ki} \right)^2 = \mathbf{a}'_K \mathbb{E} \left(\mathbf{z}_{Ki} \mathbf{z}'_{Ki} e_{Ki}^2 \right) \mathbf{a}_K = v_K.$$

We will apply the Lindeberg CLT 5.7.2, for which it is sufficient to verify Lyapunov's condition (5.6):

$$\frac{1}{n^2 v_K^2} \sum_{i=1}^n \mathbb{E} \left(\mathbf{a}'_K \mathbf{z}_{Ki} e_{Ki} \right)^4 = \frac{1}{nv_K^2} \mathbb{E} \left(\left(\mathbf{a}'_K \mathbf{z}_{Ki} \right)^4 e_{Ki}^4 \right) \rightarrow 0. \tag{12.42}$$

The assumption that $\zeta_K K^{-\alpha} = O(1)$ means $\zeta_K K^{-\alpha} \leq \kappa_1$ for some $\kappa_1 < \infty$. Then by the c_r inequality and $\mathbb{E} \left(e_i^4 | \mathbf{x}_i \right) \leq \kappa$

$$\sup_i \mathbb{E} \left(e_{Ki}^4 | \mathbf{x}_i \right) \leq 8 \sup_i \left(\mathbb{E} \left(e_i^4 | \mathbf{x}_i \right) + r_{Ki}^4 \right) \leq 8(\kappa + \kappa_1). \tag{12.43}$$

Using (12.43), the Schwarz Inequality, and (12.31)

$$\begin{aligned} \mathbb{E} \left(\left(\mathbf{a}'_K \mathbf{z}_{Ki} \right)^4 e_{Ki}^4 \right) &= \mathbb{E} \left(\left(\mathbf{a}'_K \mathbf{z}_{Ki} \right)^4 \mathbb{E} \left(e_{Ki}^4 | \mathbf{x}_i \right) \right) \\ &\leq 8(\kappa + \kappa_1) \mathbb{E} \left(\mathbf{a}'_K \mathbf{z}_{Ki} \right)^4 \\ &\leq 8(\kappa + \kappa_1) \left(\mathbf{a}'_K \mathbf{a}_K \right)^2 \mathbb{E} \left(\mathbf{z}'_{Ki} \mathbf{z}_{Ki} \right)^2 \\ &= 8(\kappa + \kappa_1) \left(\mathbf{a}'_K \mathbf{a}_K \right)^2 \zeta_K^2 K. \end{aligned} \tag{12.44}$$

Since $\mathbb{E} \left(e_{Ki}^2 | \mathbf{x}_i \right) = \mathbb{E} \left(e_i^2 | \mathbf{x}_i \right) + r_{Ki}^2 \geq \sigma^2$,

$$\begin{aligned} v_K &= \mathbf{a}'_K \mathbb{E} \left(\mathbf{z}_{Ki} \mathbf{z}'_{Ki} e_{Ki}^2 \right) \mathbf{a}_K \\ &\geq \sigma^2 \mathbf{a}'_K \mathbb{E} \left(\mathbf{z}_{Ki} \mathbf{z}'_{Ki} \right) \mathbf{a}_K \\ &= \sigma^2 \mathbf{a}'_K \mathbf{a}_K. \end{aligned} \tag{12.45}$$

Equation (12.44) and (12.45) combine to show that

$$\frac{1}{nv_K^2} \mathbb{E} \left(\left(\mathbf{a}'_K \mathbf{z}_{Ki} \right)^4 e_{Ki}^4 \right) \leq \frac{8(\kappa + \kappa_1) \zeta_K^2 K}{\sigma^4 n} = o(1)$$

under Assumption 12.7.1.4. This establishes Lyapunov's condition (12.42). Hence the Lindeberg CLT applies to (12.41) and we conclude

$$\frac{1}{\sqrt{nv_K}} \mathbf{a}'_K \mathbf{Z}'_K \mathbf{e}_K \xrightarrow{d} N(0, 1). \quad (12.46)$$

Second, take (12.39). Since $\mathbb{E}(\mathbf{e} \mid \mathbf{X}) = 0$, then applying $\mathbb{E}(e_i^2 \mid \mathbf{x}_i) \leq \bar{\sigma}^2$, the Schwarz and Norm Inequalities, (12.45), (12.34) and (12.27),

$$\begin{aligned} & \mathbb{E} \left(\left(\frac{1}{\sqrt{nv_K}} \mathbf{a}'_K (\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K) \mathbf{Z}'_K \mathbf{e} \right)^2 \mid \mathbf{X} \right) \\ &= \frac{1}{nv_K} \mathbf{a}'_K (\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K) \mathbf{Z}'_K \mathbb{E}(\mathbf{e} \mathbf{e}' \mid \mathbf{X}) \mathbf{Z}_K (\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K) \mathbf{a}_K \\ &\leq \frac{\bar{\sigma}^2}{v_K} \mathbf{a}'_K (\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K) \hat{\mathbf{Q}}_K (\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K) \mathbf{a}_K \\ &= \frac{\bar{\sigma}^2}{v_K} \mathbf{a}'_K (\hat{\mathbf{Q}}_K - \mathbf{I}_K) \hat{\mathbf{Q}}_K^{-1} (\hat{\mathbf{Q}}_K - \mathbf{I}_K) \mathbf{a}_K \\ &\leq \frac{\bar{\sigma}^2 \mathbf{a}'_K \mathbf{a}_K}{v_K} \lambda_{\max}(\hat{\mathbf{Q}}_K^{-1}) \left\| \hat{\mathbf{Q}}_K - \mathbf{I}_K \right\|^2 \\ &\leq \frac{\bar{\sigma}^2}{\sigma^2} o_p(1). \end{aligned}$$

This establishes

$$\frac{1}{\sqrt{nv_K}} \mathbf{a}'_K (\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K) \mathbf{Z}'_K \mathbf{e} \xrightarrow{p} 0. \quad (12.47)$$

Third, take (12.40). By the Cauchy-Schwarz inequality, (12.45), and the Quadratic Inequality,

$$\begin{aligned} & \left(\frac{1}{\sqrt{nv_K}} \mathbf{a}'_K (\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K) \mathbf{Z}'_K \mathbf{r}_K \right)^2 \\ &\leq \frac{\mathbf{a}'_K \mathbf{a}_K}{nv_K} \mathbf{r}'_K \mathbf{Z}_K (\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K) (\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K) \mathbf{Z}'_K \mathbf{r}_K \\ &\leq \frac{1}{\sigma^2} \lambda_{\max}(\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K)^2 \frac{1}{n} \mathbf{r}'_K \mathbf{Z}_K \mathbf{Z}'_K \mathbf{r}_K. \end{aligned} \quad (12.48)$$

Observe that since the observations are independent and $\mathbb{E} \mathbf{z}_{Ki} \mathbf{r}_{Ki} = 0$, $\mathbf{z}'_{Ki} \mathbf{z}_{Ki} \leq \zeta_K^2$, and (12.17)

$$\begin{aligned} \mathbb{E} \left(\frac{1}{n} \mathbf{r}'_K \mathbf{Z}_K \mathbf{Z}'_K \mathbf{r}_K \right) &= \mathbb{E} \left(\frac{1}{n} \sum_{i=1}^n r_{Ki} \mathbf{z}'_{Ki} \sum_{ij=1}^n \mathbf{z}_{Kj} r_{Kj} \right) \\ &= \mathbb{E} \left(\frac{1}{n} \sum_{i=1}^n \mathbf{z}'_{Ki} \mathbf{z}_{Ki} r_{Ki}^2 \right) \\ &\leq \zeta_K^2 \mathbb{E}(r_{Ki}^2) \\ &= O(\zeta_K^2 K^{-2\alpha}) \\ &= O(1) \end{aligned}$$

since $\zeta_K K^{-2} = O(1)$. Thus $\frac{1}{n} \mathbf{r}'_K \mathbf{Z}_K \mathbf{Z}'_K \mathbf{r}_K = O_p(1)$. This means that (12.48) is $o_p(1)$ since (12.28) implies

$$\lambda_{\max}(\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K) = \lambda_{\max}(\hat{\mathbf{Q}}_K^{-1}) - 1 = o_p(1). \quad (12.49)$$

Equivalently,

$$\frac{1}{\sqrt{nv_K}} \mathbf{a}'_K (\hat{\mathbf{Q}}_K^{-1} - \mathbf{I}_K) \mathbf{Z}'_K \mathbf{r}_K \xrightarrow{p} 0. \quad (12.50)$$

Equations (12.46), (12.47) and (12.50) applied to (12.38)-(12.40) show that

$$\sqrt{\frac{n}{v_k}} \left(\hat{\theta}_K - \theta_K + a(r_K) \right) \xrightarrow{d} N(0, 1)$$

completing the proof. ■

Proof of Theorem 12.13.1. The assumption that $nK^{-2\alpha} = o(1)$ implies $K^{-\alpha} = o(n^{-1/2})$. Thus

$$\zeta_K K^{-\alpha} \leq o\left(\left(\frac{\zeta_K^2}{n}\right)^{1/2}\right) \leq o\left(\left(\frac{\zeta_K^2 K}{n}\right)^{1/2}\right) = o(1)$$

so the conditions of Theorem 12.12.1 are satisfied. It is thus sufficient to show that

$$\sqrt{\frac{n}{v_k}} a(r_K) = o(1).$$

From (12.12)

$$\begin{aligned} r_K(\mathbf{x}) &= r_K^*(\mathbf{x}) + \mathbf{z}_K(\mathbf{x})' \gamma_K \\ \gamma_K &= \mathbb{E}(\mathbf{z}_{Ki} \mathbf{z}'_{Ki})^{-1} \mathbb{E}(\mathbf{z}_{Ki} r_{Ki}^*). \end{aligned}$$

Thus by linearity, applying (12.45), and the Schwarz inequality

$$\begin{aligned} \sqrt{\frac{n}{v_k}} a(r_K) &= \sqrt{\frac{n}{v_k}} (a(r_K^*) + \mathbf{a}'_K \gamma_K) \\ &\leq \frac{n^{1/2}}{\underline{\sigma}^2 (\mathbf{a}'_K \mathbf{a}_K)^{1/2}} a(r_K^*) \end{aligned} \tag{12.51}$$

$$+ \frac{(n\gamma'_K \gamma_K)^{1/2}}{\underline{\sigma}}. \tag{12.52}$$

By assumption, $n^{1/2} a(r_K^*) = O(n^{1/2} K^{-\alpha}) = o(1)$. By (12.14) and $nK^{-2\alpha} = o(1)$

$$\begin{aligned} n\gamma'_K \gamma_K &= n\mathbb{E}(\mathbf{r}_{Ki}^* \mathbf{z}'_{Ki}) \mathbb{E}(\mathbf{z}_{Ki} \mathbf{z}'_{Ki})^{-1} \mathbb{E}(\mathbf{z}_{Ki} r_{Ki}^*) \\ &\leq nO(K^{-2\alpha}) \\ &= o(1). \end{aligned}$$

Together, both (12.51) and (12.52) are $o(1)$, as required. ■

Chapter 13

Generalized Method of Moments

13.1 Overidentified Linear Model

Consider the linear model

$$\begin{aligned}y_i &= \mathbf{x}'_i \boldsymbol{\beta} + e_i \\ &= \mathbf{x}'_{1i} \boldsymbol{\beta}_1 + \mathbf{x}'_{2i} \boldsymbol{\beta}_2 + e_i \\ \mathbb{E}(\mathbf{x}_i e_i) &= \mathbf{0}\end{aligned}$$

where \mathbf{x}_{1i} is $k \times 1$ and \mathbf{x}_{2i} is $r \times 1$ with $\ell = k + r$. We know that without further restrictions, an asymptotically efficient estimator of $\boldsymbol{\beta}$ is the OLS estimator. Now suppose that we are given the information that $\boldsymbol{\beta}_2 = \mathbf{0}$. Now we can write the model as

$$\begin{aligned}y_i &= \mathbf{x}'_{1i} \boldsymbol{\beta}_1 + e_i \\ \mathbb{E}(\mathbf{x}_i e_i) &= \mathbf{0}.\end{aligned}$$

In this case, how should $\boldsymbol{\beta}_1$ be estimated? One method is OLS regression of y_i on \mathbf{x}_{1i} alone. This method, however, is not necessarily efficient, as there are ℓ restrictions in $\mathbb{E}(\mathbf{x}_i e_i) = \mathbf{0}$, while $\boldsymbol{\beta}_1$ is of dimension $k < \ell$. This situation is called **overidentified**. There are $\ell - k = r$ more moment restrictions than free parameters. We call r the **number of overidentifying restrictions**.

This is a special case of a more general class of moment condition models. Let $\mathbf{g}(y, \mathbf{x}, \mathbf{z}, \boldsymbol{\beta})$ be an $\ell \times 1$ function of a $k \times 1$ parameter $\boldsymbol{\beta}$ with $\ell \geq k$ such that

$$\mathbb{E} \mathbf{g}(y_i, \mathbf{x}_i, \mathbf{z}_i, \boldsymbol{\beta}_0) = \mathbf{0} \tag{13.1}$$

where $\boldsymbol{\beta}_0$ is the true value of $\boldsymbol{\beta}$. In our previous example, $\mathbf{g}(y, \mathbf{z}, \boldsymbol{\beta}) = \mathbf{z} \cdot (y - \mathbf{x}'_1 \boldsymbol{\beta}_1)$. In econometrics, this class of models are called **moment condition models**. In the statistics literature, these are known as **estimating equations**.

As an important special case we will devote special attention to linear moment condition models, which can be written as

$$\begin{aligned}y_i &= \mathbf{x}'_i \boldsymbol{\beta} + e_i \\ \mathbb{E}(\mathbf{z}_i e_i) &= \mathbf{0}.\end{aligned}$$

where the dimensions of \mathbf{x}_i and \mathbf{z}_i are $k \times 1$ and $\ell \times 1$, with $\ell \geq k$. If $k = \ell$ the model is **just identified**, otherwise it is **overidentified**. The variables \mathbf{x}_i may be components and functions of \mathbf{z}_i , but this is not required. This model falls in the class (13.1) by setting

$$\mathbf{g}(y, \mathbf{x}, \mathbf{z}, \boldsymbol{\beta}_0) = \mathbf{z} \cdot (y - \mathbf{x}' \boldsymbol{\beta}) \tag{13.2}$$

13.2 GMM Estimator

Define the sample analog of (13.2)

$$\bar{\mathbf{g}}_n(\boldsymbol{\beta}) = \frac{1}{n} \sum_{i=1}^n \mathbf{g}_i(\boldsymbol{\beta}) = \frac{1}{n} \sum_{i=1}^n \mathbf{z}_i (y_i - \mathbf{x}'_i \boldsymbol{\beta}) = \frac{1}{n} (\mathbf{Z}' \mathbf{y} - \mathbf{Z}' \mathbf{X} \boldsymbol{\beta}). \quad (13.3)$$

The method of moments estimator for $\boldsymbol{\beta}$ is defined as the parameter value which sets $\bar{\mathbf{g}}_n(\boldsymbol{\beta}) = \mathbf{0}$. This is generally not possible when $\ell > k$, as there are more equations than free parameters. The idea of the generalized method of moments (GMM) is to define an estimator which sets $\bar{\mathbf{g}}_n(\boldsymbol{\beta})$ “close” to zero.

For some $\ell \times \ell$ weight matrix $\mathbf{W}_n > 0$, let

$$J_n(\boldsymbol{\beta}) = n \cdot \bar{\mathbf{g}}_n(\boldsymbol{\beta})' \mathbf{W}_n \bar{\mathbf{g}}_n(\boldsymbol{\beta}).$$

This is a non-negative measure of the “length” of the vector $\bar{\mathbf{g}}_n(\boldsymbol{\beta})$. For example, if $\mathbf{W}_n = \mathbf{I}$, then, $J_n(\boldsymbol{\beta}) = n \cdot \bar{\mathbf{g}}_n(\boldsymbol{\beta})' \bar{\mathbf{g}}_n(\boldsymbol{\beta}) = n \cdot \|\bar{\mathbf{g}}_n(\boldsymbol{\beta})\|^2$, the square of the Euclidean length. The GMM estimator minimizes $J_n(\boldsymbol{\beta})$.

Definition 13.2.1 $\hat{\boldsymbol{\beta}}_{GMM} = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} J_n(\boldsymbol{\beta})$.

Note that if $k = \ell$, then $\bar{\mathbf{g}}_n(\hat{\boldsymbol{\beta}}) = \mathbf{0}$, and the GMM estimator is the method of moments estimator. The first order conditions for the GMM estimator are

$$\begin{aligned} \mathbf{0} &= \frac{\partial}{\partial \boldsymbol{\beta}} J_n(\hat{\boldsymbol{\beta}}) \\ &= 2 \frac{\partial}{\partial \boldsymbol{\beta}} \bar{\mathbf{g}}_n(\hat{\boldsymbol{\beta}})' \mathbf{W}_n \bar{\mathbf{g}}_n(\hat{\boldsymbol{\beta}}) \\ &= -2 \left(\frac{1}{n} \mathbf{X}' \mathbf{Z} \right) \mathbf{W}_n \left(\frac{1}{n} \mathbf{Z}' (\mathbf{y} - \mathbf{X} \hat{\boldsymbol{\beta}}) \right) \end{aligned}$$

so

$$2 (\mathbf{X}' \mathbf{Z}) \mathbf{W}_n (\mathbf{Z}' \mathbf{X}) \hat{\boldsymbol{\beta}} = 2 (\mathbf{X}' \mathbf{Z}) \mathbf{W}_n (\mathbf{Z}' \mathbf{y})$$

which establishes the following.

Proposition 13.2.1

$$\hat{\boldsymbol{\beta}}_{GMM} = ((\mathbf{X}' \mathbf{Z}) \mathbf{W}_n (\mathbf{Z}' \mathbf{X}))^{-1} (\mathbf{X}' \mathbf{Z}) \mathbf{W}_n (\mathbf{Z}' \mathbf{y}).$$

While the estimator depends on \mathbf{W}_n , the dependence is only up to scale, for if \mathbf{W}_n is replaced by $c \mathbf{W}_n$ for some $c > 0$, $\hat{\boldsymbol{\beta}}_{GMM}$ does not change.

13.3 Distribution of GMM Estimator

Assume that $\mathbf{W}_n \xrightarrow{p} \mathbf{W} > 0$. Let

$$\mathbf{Q} = \mathbb{E}(\mathbf{z}_i \mathbf{x}'_i)$$

and

$$\mathbf{\Omega} = \mathbb{E}(\mathbf{z}_i \mathbf{z}'_i e_i^2) = \mathbb{E}(\mathbf{g}_i \mathbf{g}'_i),$$

where $\mathbf{g}_i = \mathbf{z}_i e_i$. Then

$$\left(\frac{1}{n} \mathbf{X}' \mathbf{Z}\right) \mathbf{W}_n \left(\frac{1}{n} \mathbf{Z}' \mathbf{X}\right) \xrightarrow{p} \mathbf{Q}' \mathbf{W} \mathbf{Q}$$

and

$$\left(\frac{1}{n} \mathbf{X}' \mathbf{Z}\right) \mathbf{W}_n \left(\frac{1}{\sqrt{n}} \mathbf{Z}' \mathbf{e}\right) \xrightarrow{d} \mathbf{Q}' \mathbf{W} \mathbf{N}(\mathbf{0}, \mathbf{\Omega}).$$

We conclude:

Theorem 13.3.1 *Asymptotic Distribution of GMM Estimator*

$$\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \xrightarrow{d} \mathbf{N}(\mathbf{0}, \mathbf{V}_\beta),$$

where

$$\mathbf{V}_\beta = (\mathbf{Q}' \mathbf{W} \mathbf{Q})^{-1} (\mathbf{Q}' \mathbf{W} \mathbf{\Omega} \mathbf{W} \mathbf{Q}) (\mathbf{Q}' \mathbf{W} \mathbf{Q})^{-1}.$$

In general, GMM estimators are asymptotically normal with “sandwich form” asymptotic variances.

The optimal weight matrix \mathbf{W}_0 is one which minimizes \mathbf{V}_β . This turns out to be $\mathbf{W}_0 = \mathbf{\Omega}^{-1}$. The proof is left as an exercise. This yields the *efficient GMM* estimator:

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}' \mathbf{Z} \mathbf{\Omega}^{-1} \mathbf{Z}' \mathbf{X})^{-1} \mathbf{X}' \mathbf{Z} \mathbf{\Omega}^{-1} \mathbf{Z}' \mathbf{y}.$$

Thus we have

Theorem 13.3.2 *Asymptotic Distribution of Efficient GMM Estimator*

$$\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \xrightarrow{d} \mathbf{N}\left(\mathbf{0}, (\mathbf{Q}' \mathbf{\Omega}^{-1} \mathbf{Q})^{-1}\right).$$

$\mathbf{W}_0 = \mathbf{\Omega}^{-1}$ is not known in practice, but it can be estimated consistently. For any $\mathbf{W}_n \xrightarrow{p} \mathbf{W}_0$, we still call $\hat{\boldsymbol{\beta}}$ the efficient GMM estimator, as it has the same asymptotic distribution.

By “efficient”, we mean that this estimator has the smallest asymptotic variance in the class of GMM estimators with this set of moment conditions. This is a weak concept of optimality, as we are only considering alternative weight matrices \mathbf{W}_n . However, it turns out that the GMM estimator is semiparametrically efficient, as shown by Gary Chamberlain (1987).

If it is known that $\mathbb{E}(\mathbf{g}_i(\boldsymbol{\beta})) = \mathbf{0}$, and this is all that is known, this is a semi-parametric problem, as the distribution of the data is unknown. Chamberlain showed that in this context, no semiparametric estimator (one which is consistent globally for the class of models considered) can have a smaller asymptotic variance than $(\mathbf{G}' \mathbf{\Omega}^{-1} \mathbf{G})^{-1}$ where $\mathbf{G} = \mathbb{E} \frac{\partial}{\partial \boldsymbol{\beta}'} \mathbf{g}_i(\boldsymbol{\beta})$. Since the GMM estimator has this asymptotic variance, it is semiparametrically efficient.

This result shows that in the linear model, no estimator has greater asymptotic efficiency than the efficient linear GMM estimator. No estimator can do better (in this first-order asymptotic sense), without imposing additional assumptions.

13.4 Estimation of the Efficient Weight Matrix

Given any weight matrix $\mathbf{W}_n > 0$, the GMM estimator $\widehat{\boldsymbol{\beta}}$ is consistent yet inefficient. For example, we can set $\mathbf{W}_n = \mathbf{I}_\ell$. In the linear model, a better choice is $\mathbf{W}_n = (\mathbf{Z}'\mathbf{Z})^{-1}$. Given any such first-step estimator, we can define the residuals $\hat{e}_i = y_i - \mathbf{x}'_i \widehat{\boldsymbol{\beta}}$ and moment equations $\hat{\mathbf{g}}_i = \mathbf{z}_i \hat{e}_i = \mathbf{g}(y_i, \mathbf{x}_i, \mathbf{z}_i, \widehat{\boldsymbol{\beta}})$. Construct

$$\begin{aligned}\bar{\mathbf{g}}_n &= \bar{\mathbf{g}}_n(\widehat{\boldsymbol{\beta}}) = \frac{1}{n} \sum_{i=1}^n \hat{\mathbf{g}}_i, \\ \hat{\mathbf{g}}_i^* &= \hat{\mathbf{g}}_i - \bar{\mathbf{g}}_n,\end{aligned}$$

and define

$$\mathbf{W}_n = \left(\frac{1}{n} \sum_{i=1}^n \hat{\mathbf{g}}_i^* \hat{\mathbf{g}}_i^{*'} \right)^{-1} = \left(\frac{1}{n} \sum_{i=1}^n \hat{\mathbf{g}}_i \hat{\mathbf{g}}_i' - \bar{\mathbf{g}}_n \bar{\mathbf{g}}_n' \right)^{-1}. \quad (13.4)$$

Then $\mathbf{W}_n \xrightarrow{p} \boldsymbol{\Omega}^{-1} = \mathbf{W}_0$, and GMM using \mathbf{W}_n as the weight matrix is asymptotically efficient.

A common alternative choice is to set

$$\mathbf{W}_n = \left(\frac{1}{n} \sum_{i=1}^n \hat{\mathbf{g}}_i \hat{\mathbf{g}}_i' \right)^{-1}$$

which uses the uncentered moment conditions. Since $\mathbb{E}\mathbf{g}_i = \mathbf{0}$, these two estimators are asymptotically equivalent under the hypothesis of correct specification. However, Alastair Hall (2000) has shown that the uncentered estimator is a poor choice. When constructing hypothesis tests, under the alternative hypothesis the moment conditions are violated, i.e. $\mathbb{E}\mathbf{g}_i \neq \mathbf{0}$, so the uncentered estimator will contain an undesirable bias term and the power of the test will be adversely affected. A simple solution is to use the centered moment conditions to construct the weight matrix, as in (13.4) above.

Here is a simple way to compute the efficient GMM estimator for the linear model. First, set $\mathbf{W}_n = (\mathbf{Z}'\mathbf{Z})^{-1}$, estimate $\widehat{\boldsymbol{\beta}}$ using this weight matrix, and construct the residual $\hat{e}_i = y_i - \mathbf{x}'_i \widehat{\boldsymbol{\beta}}$. Then set $\hat{\mathbf{g}}_i = \mathbf{z}_i \hat{e}_i$, and let $\hat{\mathbf{g}}$ be the associated $n \times \ell$ matrix. Then the efficient GMM estimator is

$$\widehat{\boldsymbol{\beta}} = \left(\mathbf{X}'\mathbf{Z} (\hat{\mathbf{g}}\hat{\mathbf{g}}' - n\bar{\mathbf{g}}_n\bar{\mathbf{g}}_n')^{-1} \mathbf{Z}'\mathbf{X} \right)^{-1} \mathbf{X}'\mathbf{Z} (\hat{\mathbf{g}}\hat{\mathbf{g}}' - n\bar{\mathbf{g}}_n\bar{\mathbf{g}}_n')^{-1} \mathbf{Z}'\mathbf{y}.$$

In most cases, when we say ‘‘GMM’’, we actually mean ‘‘efficient GMM’’. There is little point in using an inefficient GMM estimator when the efficient estimator is easy to compute.

An estimator of the asymptotic variance of $\widehat{\boldsymbol{\beta}}$ can be seen from the above formula. Set

$$\widehat{\mathbf{V}} = n \left(\mathbf{X}'\mathbf{Z} (\hat{\mathbf{g}}\hat{\mathbf{g}}' - n\bar{\mathbf{g}}_n\bar{\mathbf{g}}_n')^{-1} \mathbf{Z}'\mathbf{X} \right)^{-1}.$$

Asymptotic standard errors are given by the square roots of the diagonal elements of $\frac{1}{n} \widehat{\mathbf{V}}$.

There is an important alternative to the two-step GMM estimator just described. Instead, we can let the weight matrix be considered as a function of $\boldsymbol{\beta}$. The criterion function is then

$$J(\boldsymbol{\beta}) = n \cdot \bar{\mathbf{g}}_n(\boldsymbol{\beta})' \left(\frac{1}{n} \sum_{i=1}^n \mathbf{g}_i^*(\boldsymbol{\beta}) \mathbf{g}_i^{*'}(\boldsymbol{\beta}) \right)^{-1} \bar{\mathbf{g}}_n(\boldsymbol{\beta}).$$

where

$$\mathbf{g}_i^*(\boldsymbol{\beta}) = \mathbf{g}_i(\boldsymbol{\beta}) - \bar{\mathbf{g}}_n(\boldsymbol{\beta})$$

The $\widehat{\boldsymbol{\beta}}$ which minimizes this function is called the **continuously-updated GMM estimator**, and was introduced by L. Hansen, Heaton and Yaron (1996).

The estimator appears to have some better properties than traditional GMM, but can be numerically tricky to obtain in some cases. This is a current area of research in econometrics.

13.5 GMM: The General Case

In its most general form, GMM applies whenever an economic or statistical model implies the $\ell \times 1$ moment condition

$$\mathbb{E}(\mathbf{g}_i(\boldsymbol{\beta})) = \mathbf{0}.$$

Often, this is *all* that is known. Identification requires $l \geq k = \dim(\boldsymbol{\beta})$. The GMM estimator minimizes

$$J(\boldsymbol{\beta}) = n \cdot \bar{\mathbf{g}}_n(\boldsymbol{\beta})' \mathbf{W}_n \bar{\mathbf{g}}_n(\boldsymbol{\beta})$$

where

$$\bar{\mathbf{g}}_n(\boldsymbol{\beta}) = \frac{1}{n} \sum_{i=1}^n \mathbf{g}_i(\boldsymbol{\beta})$$

and

$$\mathbf{W}_n = \left(\frac{1}{n} \sum_{i=1}^n \hat{\mathbf{g}}_i \hat{\mathbf{g}}_i' - \bar{\mathbf{g}}_n \bar{\mathbf{g}}_n' \right)^{-1},$$

with $\hat{\mathbf{g}}_i = \mathbf{g}_i(\tilde{\boldsymbol{\beta}})$ constructed using a preliminary consistent estimator $\tilde{\boldsymbol{\beta}}$, perhaps obtained by first setting $\mathbf{W}_n = \mathbf{I}$. Since the GMM estimator depends upon the first-stage estimator, often the weight matrix \mathbf{W}_n is updated, and then $\tilde{\boldsymbol{\beta}}$ recomputed. This estimator can be iterated if needed.

Theorem 13.5.1 *Distribution of Nonlinear GMM Estimator*

Under general regularity conditions,

$$\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \xrightarrow{d} \mathbf{N}\left(\mathbf{0}, (\mathbf{G}'\boldsymbol{\Omega}^{-1}\mathbf{G})^{-1}\right),$$

where

$$\boldsymbol{\Omega} = \mathbb{E}(\mathbf{g}_i \mathbf{g}_i')$$

and

$$\mathbf{G} = \mathbb{E} \frac{\partial}{\partial \boldsymbol{\beta}'} \mathbf{g}_i(\boldsymbol{\beta}).$$

The variance of $\hat{\boldsymbol{\beta}}$ may be estimated by

$$\hat{\mathbf{V}}_{\boldsymbol{\beta}} = \left(\hat{\mathbf{G}}' \hat{\boldsymbol{\Omega}}^{-1} \hat{\mathbf{G}} \right)^{-1}$$

where

$$\hat{\boldsymbol{\Omega}} = n^{-1} \sum_i \hat{\mathbf{g}}_i^* \hat{\mathbf{g}}_i^{*'}.$$

and

$$\hat{\mathbf{G}} = n^{-1} \sum_i \frac{\partial}{\partial \boldsymbol{\beta}'} \mathbf{g}_i(\hat{\boldsymbol{\beta}}).$$

The general theory of GMM estimation and testing was expounded by L. Hansen (1982).

13.6 Over-Identification Test

Overidentified models ($\ell > k$) are special in the sense that there may not be a parameter value $\boldsymbol{\beta}$ such that the moment condition

$$\mathbb{E}g(y_i, \mathbf{x}_i, \mathbf{z}_i, \boldsymbol{\beta}) = \mathbf{0}$$

holds. Thus the model – the overidentifying restrictions – are testable.

For example, take the linear model $y_i = \boldsymbol{\beta}'_1 \mathbf{x}_{1i} + \boldsymbol{\beta}'_2 \mathbf{x}_{2i} + e_i$ with $\mathbb{E}(\mathbf{x}_{1i}e_i) = \mathbf{0}$ and $\mathbb{E}(\mathbf{x}_{2i}e_i) = \mathbf{0}$. It is possible that $\boldsymbol{\beta}_2 = \mathbf{0}$, so that the linear equation may be written as $y_i = \boldsymbol{\beta}'_1 \mathbf{x}_{1i} + e_i$. However, it is possible that $\boldsymbol{\beta}_2 \neq \mathbf{0}$, and in this case it would be impossible to find a value of $\boldsymbol{\beta}_1$ so that both $\mathbb{E}(\mathbf{x}_{1i}(y_i - \mathbf{x}'_{1i}\boldsymbol{\beta}_1)) = \mathbf{0}$ and $\mathbb{E}(\mathbf{x}_{2i}(y_i - \mathbf{x}'_{1i}\boldsymbol{\beta}_1)) = \mathbf{0}$ hold simultaneously. In this sense an exclusion restriction can be seen as an overidentifying restriction.

Note that $\bar{\mathbf{g}}_n \xrightarrow{p} \mathbb{E}g_i$, and thus $\bar{\mathbf{g}}_n$ can be used to assess whether or not the hypothesis that $\mathbb{E}g_i = \mathbf{0}$ is true or not. The criterion function at the parameter estimates is

$$\begin{aligned} J_n &= n \bar{\mathbf{g}}'_n \mathbf{W}_n \bar{\mathbf{g}}_n \\ &= n^2 \bar{\mathbf{g}}'_n (\hat{\mathbf{g}}' \hat{\mathbf{g}} - n \bar{\mathbf{g}}_n \bar{\mathbf{g}}'_n)^{-1} \bar{\mathbf{g}}_n. \end{aligned}$$

is a quadratic form in $\bar{\mathbf{g}}_n$, and is thus a natural test statistic for $\mathbb{H}_0 : \mathbb{E}g_i = \mathbf{0}$.

Theorem 13.6.1 (*Sargan-Hansen*). *Under the hypothesis of correct specification, and if the weight matrix is asymptotically efficient,*

$$J_n = J_n(\hat{\boldsymbol{\beta}}) \xrightarrow{d} \chi^2_{\ell-k}.$$

The proof of the theorem is left as an exercise. This result was established by Sargan (1958) for a specialized case, and by L. Hansen (1982) for the general case.

The degrees of freedom of the asymptotic distribution are the number of overidentifying restrictions. If the statistic J exceeds the chi-square critical value, we can reject the model. Based on this information alone, it is unclear what is wrong, but it is typically cause for concern. The GMM overidentification test is a very useful by-product of the GMM methodology, and it is advisable to report the statistic J whenever GMM is the estimation method.

When over-identified models are estimated by GMM, it is customary to report the J statistic as a general test of model adequacy.

13.7 Hypothesis Testing: The Distance Statistic

We described before how to construct estimates of the asymptotic covariance matrix of the GMM estimates. These may be used to construct Wald tests of statistical hypotheses.

If the hypothesis is non-linear, a better approach is to directly use the GMM criterion function. This is sometimes called the GMM Distance statistic, and sometimes called a LR-like statistic (the LR is for likelihood-ratio). The idea was first put forward by Newey and West (1987).

For a given weight matrix \mathbf{W}_n , the GMM criterion function is

$$J_n(\boldsymbol{\beta}) = n \cdot \bar{\mathbf{g}}_n(\boldsymbol{\beta})' \mathbf{W}_n \bar{\mathbf{g}}_n(\boldsymbol{\beta})$$

For $\mathbf{h} : \mathbb{R}^k \rightarrow \mathbb{R}^r$, the hypothesis is

$$\mathbb{H}_0 : \mathbf{h}(\boldsymbol{\beta}) = \mathbf{0}.$$

The estimates under \mathbb{H}_1 are

$$\hat{\boldsymbol{\beta}} = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} J_n(\boldsymbol{\beta})$$

and those under \mathbb{H}_0 are

$$\tilde{\boldsymbol{\beta}} = \underset{\mathbf{h}(\boldsymbol{\beta})=\mathbf{0}}{\operatorname{argmin}} J(\boldsymbol{\beta}).$$

The two minimizing criterion functions are $J_n(\hat{\boldsymbol{\beta}})$ and $J_n(\tilde{\boldsymbol{\beta}})$. The GMM distance statistic is the difference

$$D_n = J_n(\tilde{\boldsymbol{\beta}}) - J_n(\hat{\boldsymbol{\beta}}).$$

Proposition 13.7.1 *If the same weight matrix \mathbf{W}_n is used for both null and alternative,*

1. $D \geq 0$
2. $D \xrightarrow{d} \chi_r^2$
3. *If \mathbf{h} is linear in $\boldsymbol{\beta}$, then D equals the Wald statistic.*

If \mathbf{h} is non-linear, the Wald statistic can work quite poorly. In contrast, current evidence suggests that the D_n statistic appears to have quite good sampling properties, and is the preferred test statistic.

Newey and West (1987) suggested to use the same weight matrix \mathbf{W}_n for both null and alternative, as this ensures that $D_n \geq 0$. This reasoning is not compelling, however, and some current research suggests that this restriction is not necessary for good performance of the test.

This test shares the useful feature of LR tests in that it is a natural by-product of the computation of alternative models.

13.8 Conditional Moment Restrictions

In many contexts, the model implies more than an unconditional moment restriction of the form $\mathbb{E}\mathbf{g}_i(\boldsymbol{\beta}) = \mathbf{0}$. It implies a conditional moment restriction of the form

$$\mathbb{E}(e_i(\boldsymbol{\beta}) \mid \mathbf{z}_i) = \mathbf{0}$$

where $e_i(\boldsymbol{\beta})$ is some $s \times 1$ function of the observation and the parameters. In many cases, $s = 1$. The variable \mathbf{z}_i is often called an **instrument**.

It turns out that this conditional moment restriction is much more powerful, and restrictive, than the unconditional moment restriction discussed above.

As discussed later in Chapter 15, the linear model $y_i = \mathbf{x}_i' \boldsymbol{\beta} + e_i$ with instruments \mathbf{z}_i falls into this class under the assumption $\mathbb{E}(e_i \mid \mathbf{z}_i) = 0$. In this case, $e_i(\boldsymbol{\beta}) = y_i - \mathbf{x}_i' \boldsymbol{\beta}$.

It is also helpful to realize that conventional regression models also fall into this class, except that in this case $\mathbf{x}_i = \mathbf{z}_i$. For example, in linear regression, $e_i(\boldsymbol{\beta}) = y_i - \mathbf{x}_i' \boldsymbol{\beta}$, while in a nonlinear regression model $e_i(\boldsymbol{\beta}) = y_i - \mathbf{g}(\mathbf{x}_i, \boldsymbol{\beta})$. In a joint model of the conditional mean $\mathbb{E}(y \mid \mathbf{x}) = \mathbf{x}' \boldsymbol{\beta}$ and variance $\operatorname{var}(y \mid \mathbf{x}) = f(\mathbf{x})' \boldsymbol{\gamma}$, then

$$e_i(\boldsymbol{\beta}, \boldsymbol{\gamma}) = \begin{cases} y_i - \mathbf{x}_i' \boldsymbol{\beta} \\ (y_i - \mathbf{x}_i' \boldsymbol{\beta})^2 - f(\mathbf{x}_i)' \boldsymbol{\gamma} \end{cases}.$$

Here $s = 2$.

Given a conditional moment restriction, an unconditional moment restriction can always be constructed. That is for any $\ell \times 1$ function $\boldsymbol{\phi}(\mathbf{x}_i, \boldsymbol{\beta})$, we can set $\mathbf{g}_i(\boldsymbol{\beta}) = \boldsymbol{\phi}(\mathbf{x}_i, \boldsymbol{\beta}) e_i(\boldsymbol{\beta})$ which

satisfies $\mathbb{E}\mathbf{g}_i(\boldsymbol{\beta}) = \mathbf{0}$ and hence defines a GMM estimator. The obvious problem is that the class of functions ϕ is infinite. Which should be selected?

This is equivalent to the problem of selection of the best instruments. If $x_i \in \mathbb{R}$ is a valid instrument satisfying $\mathbb{E}(e_i | x_i) = 0$, then x_i , x_i^2 , x_i^3 , ..., etc., are all valid instruments. Which should be used?

One solution is to construct an infinite list of potent instruments, and then use the first k instruments. How is k to be determined? This is an area of theory still under development. A recent study of this problem is Donald and Newey (2001).

Another approach is to construct the *optimal instrument*. The form was uncovered by Chamberlain (1987). Take the case $s = 1$. Let

$$\mathbf{R}_i = \mathbb{E} \left(\frac{\partial}{\partial \boldsymbol{\beta}} e_i(\boldsymbol{\beta}) \mid \mathbf{z}_i \right)$$

and

$$\sigma_i^2 = \mathbb{E} (e_i(\boldsymbol{\beta})^2 \mid \mathbf{z}_i).$$

Then the “optimal instrument” is

$$\mathbf{A}_i = -\sigma_i^{-2} \mathbf{R}_i$$

so the optimal moment is

$$\mathbf{g}_i(\boldsymbol{\beta}) = \mathbf{A}_i e_i(\boldsymbol{\beta}).$$

Setting $\mathbf{g}_i(\boldsymbol{\beta})$ to be this choice (which is $k \times 1$, so is just-identified) yields the best GMM estimator possible.

In practice, \mathbf{A}_i is unknown, but its form does help us think about construction of optimal instruments.

In the linear model $e_i(\boldsymbol{\beta}) = y_i - \mathbf{x}_i' \boldsymbol{\beta}$, note that

$$\mathbf{R}_i = -\mathbb{E}(\mathbf{x}_i \mid \mathbf{z}_i)$$

and

$$\sigma_i^2 = \mathbb{E}(e_i^2 \mid \mathbf{z}_i),$$

so

$$\mathbf{A}_i = \sigma_i^{-2} \mathbb{E}(\mathbf{x}_i \mid \mathbf{z}_i).$$

In the case of linear regression, $\mathbf{x}_i = \mathbf{z}_i$, so $\mathbf{A}_i = \sigma_i^{-2} \mathbf{z}_i$. Hence efficient GMM is GLS, as we discussed earlier in the course.

In the case of endogenous variables, note that the efficient instrument \mathbf{A}_i involves the estimation of the conditional mean of \mathbf{x}_i given \mathbf{z}_i . In other words, to get the best instrument for \mathbf{x}_i , we need the best conditional mean model for \mathbf{x}_i given \mathbf{z}_i , not just an arbitrary linear projection. The efficient instrument is also inversely proportional to the conditional variance of e_i . This is the same as the GLS estimator; namely that improved efficiency can be obtained if the observations are weighted inversely to the conditional variance of the errors.

13.9 Bootstrap GMM Inference

Let $\hat{\boldsymbol{\beta}}$ be the 2SLS or GMM estimator of $\boldsymbol{\beta}$. Using the EDF of $(y_i, \mathbf{z}_i, \mathbf{x}_i)$, we can apply the bootstrap methods discussed in Chapter 10 to compute estimates of the bias and variance of $\hat{\boldsymbol{\beta}}$, and construct confidence intervals for $\boldsymbol{\beta}$, identically as in the regression model. However, caution should be applied when interpreting such results.

A straightforward application of the nonparametric bootstrap works in the sense of consistently achieving the first-order asymptotic distribution. This has been shown by Hahn (1996). However, it fails to achieve an asymptotic refinement when the model is over-identified, jeopardizing the

theoretical justification for percentile-t methods. Furthermore, the bootstrap applied J test will yield the wrong answer.

The problem is that in the sample, $\hat{\beta}$ is the “true” value and yet $\bar{g}_n(\hat{\beta}) \neq 0$. Thus according to random variables (y_i^*, z_i^*, x_i^*) drawn from the EDF F_n ,

$$\mathbb{E} \left(g_i \left(\hat{\beta} \right) \right) = \bar{g}_n(\hat{\beta}) \neq \mathbf{0}.$$

This means that (y_i^*, z_i^*, x_i^*) do not satisfy the same moment conditions as the population distribution.

A correction suggested by Hall and Horowitz (1996) can solve the problem. Given the bootstrap sample $(\mathbf{y}^*, \mathbf{Z}^*, \mathbf{X}^*)$, define the bootstrap GMM criterion

$$J_n^*(\beta) = n \cdot \left(\bar{g}_n^*(\beta) - \bar{g}_n(\hat{\beta}) \right)' \mathbf{W}_n^* \left(\bar{g}_n^*(\beta) - \bar{g}_n(\hat{\beta}) \right)$$

where $\bar{g}_n(\hat{\beta})$ is from the in-sample data, not from the bootstrap data.

Let $\hat{\beta}_n^*$ minimize $J_n^*(\beta)$, and define all statistics and tests accordingly. In the linear model, this implies that the bootstrap estimator is

$$\hat{\beta}_n^* = (\mathbf{X}^{*'} \mathbf{Z}^* \mathbf{W}_n^* \mathbf{Z}^{*'} \mathbf{X}^*)^{-1} (\mathbf{X}^{*'} \mathbf{Z}^* \mathbf{W}_n^* (\mathbf{Z}^{*'} \mathbf{y}^* - \mathbf{Z}' \hat{e})).$$

where $\hat{e} = \mathbf{y} - \mathbf{X} \hat{\beta}$ are the in-sample residuals. The bootstrap J statistic is $J_n^*(\hat{\beta}_n^*)$.

Brown and Newey (2002) have an alternative solution. They note that we can sample from the observations with the empirical likelihood probabilities \hat{p}_i described in Chapter 14. Since $\sum_{i=1}^n \hat{p}_i g_i(\hat{\beta}) = \mathbf{0}$, this sampling scheme preserves the moment conditions of the model, so no recentering or adjustments is needed. Brown and Newey argue that this bootstrap procedure will be more efficient than the Hall-Horowitz GMM bootstrap.

Exercises

Exercise 13.1 Take the model

$$\begin{aligned} y_i &= \mathbf{x}'_i \boldsymbol{\beta} + e_i \\ \mathbb{E}(\mathbf{x}_i e_i) &= \mathbf{0} \\ e_i^2 &= \mathbf{z}'_i \boldsymbol{\gamma} + \eta_i \\ \mathbb{E}(\mathbf{z}_i \eta_i) &= \mathbf{0}. \end{aligned}$$

Find the method of moments estimators $(\hat{\boldsymbol{\beta}}, \hat{\boldsymbol{\gamma}})$ for $(\boldsymbol{\beta}, \boldsymbol{\gamma})$.

Exercise 13.2 Take the single equation

$$\begin{aligned} \mathbf{y} &= \mathbf{X}\boldsymbol{\beta} + \mathbf{e} \\ \mathbb{E}(\mathbf{e} \mid \mathbf{Z}) &= \mathbf{0} \end{aligned}$$

Assume $\mathbb{E}(e_i^2 \mid \mathbf{z}_i) = \sigma^2$. Show that if $\hat{\boldsymbol{\beta}}$ is estimated by GMM with weight matrix $\mathbf{W}_n = (\mathbf{Z}'\mathbf{Z})^{-1}$, then

$$\sqrt{n}(\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \xrightarrow{d} \mathbf{N}\left(\mathbf{0}, \sigma^2 (\mathbf{Q}'\mathbf{M}^{-1}\mathbf{Q})^{-1}\right)$$

where $\mathbf{Q} = \mathbb{E}(\mathbf{z}_i \mathbf{x}'_i)$ and $\mathbf{M} = \mathbb{E}(\mathbf{z}_i \mathbf{z}'_i)$.

Exercise 13.3 Take the model $y_i = \mathbf{x}'_i \boldsymbol{\beta} + e_i$ with $\mathbb{E}(\mathbf{z}_i e_i) = \mathbf{0}$. Let $\hat{e}_i = y_i - \mathbf{x}'_i \hat{\boldsymbol{\beta}}$ where $\hat{\boldsymbol{\beta}}$ is consistent for $\boldsymbol{\beta}$ (e.g. a GMM estimator with arbitrary weight matrix). Define the estimate of the optimal GMM weight matrix

$$\mathbf{W}_n = \left(\frac{1}{n} \sum_{i=1}^n \mathbf{z}_i \mathbf{z}'_i \hat{e}_i^2 \right)^{-1}.$$

Show that $\mathbf{W}_n \xrightarrow{p} \boldsymbol{\Omega}^{-1}$ where $\boldsymbol{\Omega} = \mathbb{E}(\mathbf{z}_i \mathbf{z}'_i e_i^2)$.

Exercise 13.4 In the linear model estimated by GMM with general weight matrix \mathbf{W} , the asymptotic variance of $\hat{\boldsymbol{\beta}}_{GMM}$ is

$$\mathbf{V} = (\mathbf{Q}'\mathbf{W}\mathbf{Q})^{-1} \mathbf{Q}'\mathbf{W}\boldsymbol{\Omega}\mathbf{W}\mathbf{Q} (\mathbf{Q}'\mathbf{W}\mathbf{Q})^{-1}$$

- Let \mathbf{V}_0 be this matrix when $\mathbf{W} = \boldsymbol{\Omega}^{-1}$. Show that $\mathbf{V}_0 = (\mathbf{Q}'\boldsymbol{\Omega}^{-1}\mathbf{Q})^{-1}$.
- We want to show that for any \mathbf{W} , $\mathbf{V} - \mathbf{V}_0$ is positive semi-definite (for then \mathbf{V}_0 is the smaller possible covariance matrix and $\mathbf{W} = \boldsymbol{\Omega}^{-1}$ is the efficient weight matrix). To do this, start by finding matrices \mathbf{A} and \mathbf{B} such that $\mathbf{V} = \mathbf{A}'\boldsymbol{\Omega}\mathbf{A}$ and $\mathbf{V}_0 = \mathbf{B}'\boldsymbol{\Omega}\mathbf{B}$.
- Show that $\mathbf{B}'\boldsymbol{\Omega}\mathbf{A} = \mathbf{B}'\boldsymbol{\Omega}\mathbf{B}$ and therefore that $\mathbf{B}'\boldsymbol{\Omega}(\mathbf{A} - \mathbf{B}) = \mathbf{0}$.
- Use the expressions $\mathbf{V} = \mathbf{A}'\boldsymbol{\Omega}\mathbf{A}$, $\mathbf{A} = \mathbf{B} + (\mathbf{A} - \mathbf{B})$, and $\mathbf{B}'\boldsymbol{\Omega}(\mathbf{A} - \mathbf{B}) = \mathbf{0}$ to show that $\mathbf{V} \geq \mathbf{V}_0$.

Exercise 13.5 The equation of interest is

$$\begin{aligned} y_i &= \mathbf{m}(\mathbf{x}_i, \boldsymbol{\beta}) + e_i \\ \mathbb{E}(\mathbf{z}_i e_i) &= \mathbf{0}. \end{aligned}$$

The observed data is $(y_i, \mathbf{z}_i, \mathbf{x}_i)$. \mathbf{z}_i is $\ell \times 1$ and $\boldsymbol{\beta}$ is $k \times 1$, $\ell \geq k$. Show how to construct an efficient GMM estimator for $\boldsymbol{\beta}$.

Exercise 13.6 In the linear model $\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}$ with $\mathbb{E}(\mathbf{x}_i e_i) = \mathbf{0}$, a Generalized Method of Moments (GMM) criterion function for $\boldsymbol{\beta}$ is defined as

$$J_n(\boldsymbol{\beta}) = \frac{1}{n} (\mathbf{y} - \mathbf{X}\boldsymbol{\beta})' \mathbf{X} \widehat{\boldsymbol{\Omega}}^{-1} \mathbf{X}' (\mathbf{y} - \mathbf{X}\boldsymbol{\beta}) \quad (13.5)$$

where $\widehat{\boldsymbol{\Omega}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i' \hat{e}_i^2$, $\hat{e}_i = y_i - \mathbf{x}_i' \hat{\boldsymbol{\beta}}$ are the OLS residuals, and $\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1} \mathbf{X}'\mathbf{y}$ is LS. The GMM estimator of $\boldsymbol{\beta}$, subject to the restriction $\mathbf{h}(\boldsymbol{\beta}) = \mathbf{0}$, is defined as

$$\tilde{\boldsymbol{\beta}} = \underset{\mathbf{h}(\boldsymbol{\beta})=\mathbf{0}}{\operatorname{argmin}} J_n(\boldsymbol{\beta}).$$

The GMM test statistic (the distance statistic) of the hypothesis $\mathbf{h}(\boldsymbol{\beta}) = \mathbf{0}$ is

$$D = J_n(\tilde{\boldsymbol{\beta}}) = \min_{\mathbf{h}(\boldsymbol{\beta})=\mathbf{0}} J_n(\boldsymbol{\beta}). \quad (13.6)$$

(a) Show that you can rewrite $J_n(\boldsymbol{\beta})$ in (13.5) as

$$J_n(\boldsymbol{\beta}) = n (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}})' \hat{\mathbf{V}}_{\boldsymbol{\beta}}^{-1} (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}})$$

thus $\tilde{\boldsymbol{\beta}}$ is the same as the minimum distance estimator.

(b) Show that in this setting, the distance statistic D in (13.6) equals the Wald statistic.

Exercise 13.7 Take the linear model

$$\begin{aligned} y_i &= \mathbf{x}_i' \boldsymbol{\beta} + e_i \\ \mathbb{E}(\mathbf{z}_i e_i) &= \mathbf{0}. \end{aligned}$$

and consider the GMM estimator $\hat{\boldsymbol{\beta}}$ of $\boldsymbol{\beta}$. Let

$$J_n = n \bar{\mathbf{g}}_n(\hat{\boldsymbol{\beta}})' \widehat{\boldsymbol{\Omega}}^{-1} \bar{\mathbf{g}}_n(\hat{\boldsymbol{\beta}})$$

denote the test of overidentifying restrictions. Show that $J_n \xrightarrow{d} \chi_{\ell-k}^2$ as $n \rightarrow \infty$ by demonstrating each of the following:

(a) Since $\boldsymbol{\Omega} > 0$, we can write $\boldsymbol{\Omega}^{-1} = \mathbf{C}\mathbf{C}'$ and $\boldsymbol{\Omega} = \mathbf{C}'^{-1}\mathbf{C}^{-1}$

(b) $J_n = n (\mathbf{C}' \bar{\mathbf{g}}_n(\hat{\boldsymbol{\beta}}))' (\mathbf{C}' \widehat{\boldsymbol{\Omega}} \mathbf{C})^{-1} \mathbf{C}' \bar{\mathbf{g}}_n(\hat{\boldsymbol{\beta}})$

(c) $\mathbf{C}' \bar{\mathbf{g}}_n(\hat{\boldsymbol{\beta}}) = \mathbf{D}_n \mathbf{C}' \bar{\mathbf{g}}_n(\boldsymbol{\beta}_0)$ where

$$\mathbf{D}_n = \mathbf{I}_{\ell} - \mathbf{C}' \left(\frac{1}{n} \mathbf{Z}' \mathbf{X} \right) \left(\left(\frac{1}{n} \mathbf{X}' \mathbf{Z} \right) \widehat{\boldsymbol{\Omega}}^{-1} \left(\frac{1}{n} \mathbf{Z}' \mathbf{X} \right) \right)^{-1} \left(\frac{1}{n} \mathbf{X}' \mathbf{Z} \right) \widehat{\boldsymbol{\Omega}}^{-1} \mathbf{C}'^{-1}$$

$$\bar{\mathbf{g}}_n(\boldsymbol{\beta}_0) = \frac{1}{n} \mathbf{Z}' \mathbf{e}.$$

(d) $\mathbf{D}_n \xrightarrow{p} \mathbf{I}_{\ell} - \mathbf{R}(\mathbf{R}'\mathbf{R})^{-1} \mathbf{R}'$ where $\mathbf{R} = \mathbf{C}' \mathbb{E}(\mathbf{z}_i \mathbf{x}_i')$

(e) $n^{1/2} \mathbf{C}' \bar{\mathbf{g}}_n(\boldsymbol{\beta}_0) \xrightarrow{d} \mathbf{u} \sim \mathbf{N}(\mathbf{0}, \mathbf{I}_{\ell})$

(f) $J_n \xrightarrow{d} \mathbf{u}' (\mathbf{I}_{\ell} - \mathbf{R}(\mathbf{R}'\mathbf{R})^{-1} \mathbf{R}') \mathbf{u}$

(g) $\mathbf{u}' (\mathbf{I}_{\ell} - \mathbf{R}(\mathbf{R}'\mathbf{R})^{-1} \mathbf{R}') \mathbf{u} \sim \chi_{\ell-k}^2$.

Hint: $\mathbf{I}_{\ell} - \mathbf{R}(\mathbf{R}'\mathbf{R})^{-1} \mathbf{R}'$ is a projection matrix.

Chapter 14

Empirical Likelihood

14.1 Non-Parametric Likelihood

An alternative to GMM is **empirical likelihood**. The idea is due to Art Owen (1988, 2001) and has been extended to moment condition models by Qin and Lawless (1994). It is a non-parametric analog of likelihood estimation.

The idea is to construct a multinomial distribution $F(p_1, \dots, p_n)$ which places probability p_i at each observation. To be a valid multinomial distribution, these probabilities must satisfy the requirements that $p_i \geq 0$ and

$$\sum_{i=1}^n p_i = 1. \quad (14.1)$$

Since each observation is observed once in the sample, the log-likelihood function for this multinomial distribution is

$$\log L(p_1, \dots, p_n) = \sum_{i=1}^n \log(p_i). \quad (14.2)$$

First let us consider a just-identified model. In this case the moment condition places no additional restrictions on the multinomial distribution. The maximum likelihood estimators of the probabilities (p_1, \dots, p_n) are those which maximize the log-likelihood subject to the constraint (14.1). This is equivalent to maximizing

$$\sum_{i=1}^n \log(p_i) - \mu \left(\sum_{i=1}^n p_i - 1 \right)$$

where μ is a Lagrange multiplier. The n first order conditions are $0 = p_i^{-1} - \mu$. Combined with the constraint (14.1) we find that the MLE is $p_i = n^{-1}$ yielding the log-likelihood $-n \log(n)$.

Now consider the case of an overidentified model with moment condition

$$\mathbb{E} \mathbf{g}_i(\boldsymbol{\beta}_0) = \mathbf{0}$$

where \mathbf{g} is $\ell \times 1$ and $\boldsymbol{\beta}$ is $k \times 1$ and for simplicity we write $\mathbf{g}_i(\boldsymbol{\beta}) = \mathbf{g}(y_i, \mathbf{z}_i, \mathbf{x}_i, \boldsymbol{\beta})$. The multinomial distribution which places probability p_i at each observation $(y_i, \mathbf{x}_i, \mathbf{z}_i)$ will satisfy this condition if and only if

$$\sum_{i=1}^n p_i \mathbf{g}_i(\boldsymbol{\beta}) = \mathbf{0} \quad (14.3)$$

The **empirical likelihood estimator** is the value of $\boldsymbol{\beta}$ which maximizes the multinomial log-likelihood (14.2) subject to the restrictions (14.1) and (14.3).

The Lagrangian for this maximization problem is

$$\mathcal{L}(\boldsymbol{\beta}, p_1, \dots, p_n, \boldsymbol{\lambda}, \mu) = \sum_{i=1}^n \log(p_i) - \mu \left(\sum_{i=1}^n p_i - 1 \right) - n \boldsymbol{\lambda}' \sum_{i=1}^n p_i \mathbf{g}_i(\boldsymbol{\beta})$$

where $\boldsymbol{\lambda}$ and μ are Lagrange multipliers. The first-order-conditions of \mathcal{L} with respect to p_i , μ , and $\boldsymbol{\lambda}$ are

$$\begin{aligned} \frac{1}{p_i} &= \mu + n \boldsymbol{\lambda}' \mathbf{g}_i(\boldsymbol{\beta}) \\ \sum_{i=1}^n p_i &= 1 \\ \sum_{i=1}^n p_i \mathbf{g}_i(\boldsymbol{\beta}) &= \mathbf{0}. \end{aligned}$$

Multiplying the first equation by p_i , summing over i , and using the second and third equations, we find $\mu = n$ and

$$p_i = \frac{1}{n(1 + \boldsymbol{\lambda}' \mathbf{g}_i(\boldsymbol{\beta}))}.$$

Substituting into \mathcal{L} we find

$$R(\boldsymbol{\beta}, \boldsymbol{\lambda}) = -n \log(n) - \sum_{i=1}^n \log(1 + \boldsymbol{\lambda}' \mathbf{g}_i(\boldsymbol{\beta})). \quad (14.4)$$

For given $\boldsymbol{\beta}$, the Lagrange multiplier $\boldsymbol{\lambda}(\boldsymbol{\beta})$ minimizes $R(\boldsymbol{\beta}, \boldsymbol{\lambda})$:

$$\boldsymbol{\lambda}(\boldsymbol{\beta}) = \underset{\boldsymbol{\lambda}}{\operatorname{argmin}} R(\boldsymbol{\beta}, \boldsymbol{\lambda}). \quad (14.5)$$

This minimization problem is the dual of the constrained maximization problem. The solution (when it exists) is well defined since $R(\boldsymbol{\beta}, \boldsymbol{\lambda})$ is a convex function of $\boldsymbol{\lambda}$. The solution cannot be obtained explicitly, but must be obtained numerically (see section 6.5). This yields the (profile) empirical log-likelihood function for $\boldsymbol{\beta}$.

$$\begin{aligned} R(\boldsymbol{\beta}) &= R(\boldsymbol{\beta}, \boldsymbol{\lambda}(\boldsymbol{\beta})) \\ &= -n \log(n) - \sum_{i=1}^n \log(1 + \boldsymbol{\lambda}(\boldsymbol{\beta})' \mathbf{g}_i(\boldsymbol{\beta})) \end{aligned}$$

The EL estimate $\hat{\boldsymbol{\beta}}$ is the value which maximizes $R(\boldsymbol{\beta})$, or equivalently minimizes its negative

$$\hat{\boldsymbol{\beta}} = \underset{\boldsymbol{\beta}}{\operatorname{argmin}} [-R(\boldsymbol{\beta})] \quad (14.6)$$

Numerical methods are required for calculation of $\hat{\boldsymbol{\beta}}$ (see Section 14.5).

As a by-product of estimation, we also obtain the Lagrange multiplier $\hat{\boldsymbol{\lambda}} = \boldsymbol{\lambda}(\hat{\boldsymbol{\beta}})$, probabilities

$$\hat{p}_i = \frac{1}{n(1 + \hat{\boldsymbol{\lambda}}' \mathbf{g}_i(\hat{\boldsymbol{\beta}}))}$$

and maximized empirical likelihood

$$R(\hat{\boldsymbol{\beta}}) = \sum_{i=1}^n \log(\hat{p}_i). \quad (14.7)$$

14.2 Asymptotic Distribution of EL Estimator

Let β_0 denote the true value of β and define

$$\mathbf{G}_i(\beta) = \frac{\partial}{\partial \beta'} g_i(\beta) \quad (14.8)$$

$$\mathbf{G} = \mathbb{E} \mathbf{G}_i(\beta_0)$$

$$\mathbf{\Omega} = \mathbb{E} (g_i(\beta_0) g_i(\beta_0)')$$

and

$$\mathbf{V} = (\mathbf{G}' \mathbf{\Omega}^{-1} \mathbf{G})^{-1} \quad (14.9)$$

$$\mathbf{V}_\lambda = \mathbf{\Omega} - \mathbf{G} (\mathbf{G}' \mathbf{\Omega}^{-1} \mathbf{G})^{-1} \mathbf{G}' \quad (14.10)$$

For example, in the linear model, $\mathbf{G}_i(\beta) = -z_i x_i'$, $\mathbf{G} = -\mathbb{E}(z_i x_i')$, and $\mathbf{\Omega} = \mathbb{E}(z_i z_i' e_i^2)$.

Theorem 14.2.1 *Under regularity conditions,*

$$\sqrt{n} (\hat{\beta} - \beta_0) \xrightarrow{d} \mathbf{N}(\mathbf{0}, \mathbf{V}_\beta)$$

$$\sqrt{n} \hat{\lambda} \xrightarrow{d} \mathbf{\Omega}^{-1} \mathbf{N}(\mathbf{0}, \mathbf{V}_\lambda)$$

where \mathbf{V} and \mathbf{V}_λ are defined in (14.9) and (14.10), and $\sqrt{n} (\hat{\beta} - \beta_0)$ and $\sqrt{n} \hat{\lambda}$ are asymptotically independent.

The theorem shows that the asymptotic variance \mathbf{V}_β for $\hat{\beta}$ is the same as for efficient GMM. Thus the EL estimator is asymptotically efficient.

Chamberlain (1987) showed that \mathbf{V}_β is the semiparametric efficiency bound for β in the overidentified moment condition model. This means that no consistent estimator for this class of models can have a lower asymptotic variance than \mathbf{V}_β . Since the EL estimator achieves this bound, it is an asymptotically efficient estimator for β .

Proof of Theorem 14.2.1. $(\hat{\beta}, \hat{\lambda})$ jointly solve

$$\mathbf{0} = \frac{\partial}{\partial \lambda} R(\beta, \lambda) = - \sum_{i=1}^n \frac{g_i(\hat{\beta})}{(1 + \hat{\lambda}' g_i(\hat{\beta}))} \quad (14.11)$$

$$\mathbf{0} = \frac{\partial}{\partial \beta} R(\beta, \lambda) = - \sum_{i=1}^n \frac{\mathbf{G}_i(\hat{\beta})' \lambda}{1 + \hat{\lambda}' g_i(\hat{\beta})}. \quad (14.12)$$

Let $\mathbf{G}_n = \frac{1}{n} \sum_{i=1}^n \mathbf{G}_i(\beta_0)$, $\bar{g}_n = \frac{1}{n} \sum_{i=1}^n g_i(\beta_0)$ and $\mathbf{\Omega}_n = \frac{1}{n} \sum_{i=1}^n g_i(\beta_0) g_i(\beta_0)'$.

Expanding (14.12) around $\beta = \beta_0$ and $\lambda = \lambda_0 = \mathbf{0}$ yields

$$\mathbf{0} \simeq \mathbf{G}'_n (\hat{\lambda} - \lambda_0). \quad (14.13)$$

Expanding (14.11) around $\beta = \beta_0$ and $\lambda = \lambda_0 = \mathbf{0}$ yields

$$\mathbf{0} \simeq -\bar{g}_n - \mathbf{G}_n (\hat{\beta} - \beta_0) + \mathbf{\Omega}_n \hat{\lambda} \quad (14.14)$$

Premultiplying by $\mathbf{G}'_n \boldsymbol{\Omega}_n^{-1}$ and using (14.13) yields

$$\begin{aligned} \mathbf{0} &\simeq -\mathbf{G}'_n \boldsymbol{\Omega}_n^{-1} \bar{\mathbf{g}}_n - \mathbf{G}'_n \boldsymbol{\Omega}_n^{-1} \mathbf{G}_n (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) + \mathbf{G}'_n \boldsymbol{\Omega}_n^{-1} \boldsymbol{\Omega}_n \hat{\boldsymbol{\lambda}} \\ &= -\mathbf{G}'_n \boldsymbol{\Omega}_n^{-1} \bar{\mathbf{g}}_n - \mathbf{G}'_n \boldsymbol{\Omega}_n^{-1} \mathbf{G}_n (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) \end{aligned}$$

Solving for $\hat{\boldsymbol{\beta}}$ and using the WLLN and CLT yields

$$\begin{aligned} \sqrt{n} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) &\simeq -(\mathbf{G}'_n \boldsymbol{\Omega}_n^{-1} \mathbf{G}_n)^{-1} \mathbf{G}'_n \boldsymbol{\Omega}_n^{-1} \sqrt{n} \bar{\mathbf{g}}_n \\ &\xrightarrow{d} (\mathbf{G}' \boldsymbol{\Omega}^{-1} \mathbf{G})^{-1} \mathbf{G}' \boldsymbol{\Omega}^{-1} \mathbf{N}(\mathbf{0}, \boldsymbol{\Omega}) \\ &= \mathbf{N}(\mathbf{0}, \mathbf{V}_\beta) \end{aligned} \tag{14.15}$$

Solving (14.14) for $\hat{\boldsymbol{\lambda}}$ and using (14.15) yields

$$\begin{aligned} \sqrt{n} \hat{\boldsymbol{\lambda}} &\simeq \boldsymbol{\Omega}_n^{-1} \left(\mathbf{I} - \mathbf{G}_n (\mathbf{G}'_n \boldsymbol{\Omega}_n^{-1} \mathbf{G}_n)^{-1} \mathbf{G}'_n \boldsymbol{\Omega}_n^{-1} \right) \sqrt{n} \bar{\mathbf{g}}_n \\ &\xrightarrow{d} \boldsymbol{\Omega}^{-1} \left(\mathbf{I} - \mathbf{G} (\mathbf{G}' \boldsymbol{\Omega}^{-1} \mathbf{G})^{-1} \mathbf{G}' \boldsymbol{\Omega}^{-1} \right) \mathbf{N}(\mathbf{0}, \boldsymbol{\Omega}) \\ &= \boldsymbol{\Omega}^{-1} \mathbf{N}(\mathbf{0}, \mathbf{V}_\lambda) \end{aligned} \tag{14.16}$$

Furthermore, since

$$\mathbf{G}' \left(\mathbf{I} - \boldsymbol{\Omega}^{-1} \mathbf{G} (\mathbf{G}' \boldsymbol{\Omega}^{-1} \mathbf{G})^{-1} \mathbf{G}' \right) = \mathbf{0}$$

$\sqrt{n} (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0)$ and $\sqrt{n} \hat{\boldsymbol{\lambda}}$ are asymptotically uncorrelated and hence independent.

14.3 Overidentifying Restrictions

In a parametric likelihood context, tests are based on the difference in the log likelihood functions. The same statistic can be constructed for empirical likelihood. Twice the difference between the unrestricted empirical log-likelihood $-n \log(n)$ and the maximized empirical log-likelihood for the model (14.7) is

$$LR_n = \sum_{i=1}^n 2 \log \left(1 + \hat{\boldsymbol{\lambda}}' \mathbf{g}_i(\hat{\boldsymbol{\beta}}) \right). \tag{14.17}$$

Theorem 14.3.1 *If $\mathbb{E} \mathbf{g}_i(\boldsymbol{\beta}_0) = \mathbf{0}$ then $LR_n \xrightarrow{d} \chi_{\ell-k}^2$.*

The EL overidentification test is similar to the GMM overidentification test. They are asymptotically first-order equivalent, and have the same interpretation. The overidentification test is a very useful by-product of EL estimation, and it is advisable to report the statistic LR_n whenever EL is the estimation method.

Proof of Theorem 14.3.1. First, by a Taylor expansion, (14.15), and (14.16),

$$\begin{aligned} \frac{1}{\sqrt{n}} \sum_{i=1}^n \mathbf{g}_i(\hat{\boldsymbol{\beta}}) &\simeq \sqrt{n} \left(\bar{\mathbf{g}}_n + \mathbf{G}_n (\hat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0) \right) \\ &\simeq \left(\mathbf{I} - \mathbf{G}_n (\mathbf{G}'_n \boldsymbol{\Omega}_n^{-1} \mathbf{G}_n)^{-1} \mathbf{G}'_n \boldsymbol{\Omega}_n^{-1} \right) \sqrt{n} \bar{\mathbf{g}}_n \\ &\simeq \boldsymbol{\Omega}_n \sqrt{n} \hat{\boldsymbol{\lambda}}. \end{aligned}$$

Second, since $\log(1 + u) \simeq u - u^2/2$ for u small,

$$\begin{aligned} LR_n &= \sum_{i=1}^n 2 \log \left(1 + \hat{\lambda}' g_i(\hat{\beta}) \right) \\ &\simeq 2 \hat{\lambda}' \sum_{i=1}^n g_i(\hat{\beta}) - \hat{\lambda}' \sum_{i=1}^n g_i(\hat{\beta}) g_i(\hat{\beta})' \hat{\lambda} \\ &\simeq n \hat{\lambda}' \Omega_n \hat{\lambda} \\ &\xrightarrow{d} N(\mathbf{0}, \mathbf{V}_\lambda)' \Omega^{-1} N(\mathbf{0}, \mathbf{V}_\lambda) \\ &= \chi_{\ell-k}^2 \end{aligned}$$

where the proof of the final equality is left as an exercise.

14.4 Testing

Let the maintained model be

$$\mathbb{E} g_i(\beta) = \mathbf{0} \tag{14.18}$$

where g is $\ell \times 1$ and β is $k \times 1$. By “maintained” we mean that the overidentifying restrictions contained in (14.18) are assumed to hold and are not being challenged (at least for the test discussed in this section). The hypothesis of interest is

$$h(\beta) = \mathbf{0}.$$

where $h : \mathbb{R}^k \rightarrow \mathbb{R}^a$. The restricted EL estimator and likelihood are the values which solve

$$\begin{aligned} \tilde{\beta} &= \operatorname{argmax}_{h(\beta)=\mathbf{0}} R(\beta) \\ R(\tilde{\beta}) &= \max_{h(\beta)=\mathbf{0}} R(\beta). \end{aligned}$$

Fundamentally, the restricted EL estimator $\tilde{\beta}$ is simply an EL estimator with $\ell - k + a$ overidentifying restrictions, so there is no fundamental change in the distribution theory for $\tilde{\beta}$ relative to $\hat{\beta}$. To test the hypothesis $h(\beta)$ while maintaining (14.18), the simple overidentifying restrictions test (14.17) is not appropriate. Instead we use the difference in log-likelihoods:

$$LR_n = 2 \left(R(\hat{\beta}) - R(\tilde{\beta}) \right).$$

This test statistic is a natural analog of the GMM distance statistic.

Theorem 14.4.1 Under (14.18) and $\mathbb{H}_0 : h(\beta) = \mathbf{0}$, $LR_n \xrightarrow{d} \chi_a^2$.

The proof of this result is more challenging and is omitted.

14.5 Numerical Computation

Gauss code which implements the methods discussed below can be found at

<http://www.ssc.wisc.edu/~bhansen/progs/elike.prc>

Derivatives

The numerical calculations depend on derivatives of the dual likelihood function (14.4). Define

$$\mathbf{g}_i^*(\boldsymbol{\beta}, \boldsymbol{\lambda}) = \frac{\mathbf{g}_i(\boldsymbol{\beta})}{(1 + \boldsymbol{\lambda}' \mathbf{g}_i(\boldsymbol{\beta}))}$$

$$\mathbf{G}_i^*(\boldsymbol{\beta}, \boldsymbol{\lambda}) = \frac{\mathbf{G}_i(\boldsymbol{\beta})' \boldsymbol{\lambda}}{1 + \boldsymbol{\lambda}' \mathbf{g}_i(\boldsymbol{\beta})}$$

The first derivatives of (14.4) are

$$\mathbf{R}_\lambda = \frac{\partial}{\partial \boldsymbol{\lambda}} R(\boldsymbol{\beta}, \boldsymbol{\lambda}) = - \sum_{i=1}^n \mathbf{g}_i^*(\boldsymbol{\beta}, \boldsymbol{\lambda})$$

$$\mathbf{R}_\beta = \frac{\partial}{\partial \boldsymbol{\beta}} R(\boldsymbol{\beta}, \boldsymbol{\lambda}) = - \sum_{i=1}^n \mathbf{G}_i^*(\boldsymbol{\beta}, \boldsymbol{\lambda}).$$

The second derivatives are

$$\mathbf{R}_{\lambda\lambda} = \frac{\partial^2}{\partial \boldsymbol{\lambda} \partial \boldsymbol{\lambda}'} R(\boldsymbol{\beta}, \boldsymbol{\lambda}) = \sum_{i=1}^n \mathbf{g}_i^*(\boldsymbol{\beta}, \boldsymbol{\lambda}) \mathbf{g}_i^*(\boldsymbol{\beta}, \boldsymbol{\lambda})'$$

$$\mathbf{R}_{\lambda\beta} = \frac{\partial^2}{\partial \boldsymbol{\lambda} \partial \boldsymbol{\beta}'} R(\boldsymbol{\beta}, \boldsymbol{\lambda}) = \sum_{i=1}^n \left(\mathbf{g}_i^*(\boldsymbol{\beta}, \boldsymbol{\lambda}) \mathbf{G}_i^*(\boldsymbol{\beta}, \boldsymbol{\lambda})' - \frac{\mathbf{G}_i(\boldsymbol{\beta})}{1 + \boldsymbol{\lambda}' \mathbf{g}_i(\boldsymbol{\beta})} \right)$$

$$\mathbf{R}_{\beta\beta} = \frac{\partial^2}{\partial \boldsymbol{\beta} \partial \boldsymbol{\beta}'} R(\boldsymbol{\beta}, \boldsymbol{\lambda}) = \sum_{i=1}^n \left(\mathbf{G}_i^*(\boldsymbol{\beta}, \boldsymbol{\lambda}) \mathbf{G}_i^*(\boldsymbol{\beta}, \boldsymbol{\lambda})' - \frac{\frac{\partial^2}{\partial \boldsymbol{\beta} \partial \boldsymbol{\beta}'} (\mathbf{g}_i(\boldsymbol{\beta})' \boldsymbol{\lambda})}{1 + \boldsymbol{\lambda}' \mathbf{g}_i(\boldsymbol{\beta})} \right)$$

Inner Loop

The so-called “inner loop” solves (14.5) for given $\boldsymbol{\beta}$. The modified Newton method takes a quadratic approximation to $R_n(\boldsymbol{\beta}, \boldsymbol{\lambda})$ yielding the iteration rule

$$\boldsymbol{\lambda}_{j+1} = \boldsymbol{\lambda}_j - \delta (\mathbf{R}_{\lambda\lambda}(\boldsymbol{\beta}, \boldsymbol{\lambda}_j))^{-1} \mathbf{R}_\lambda(\boldsymbol{\beta}, \boldsymbol{\lambda}_j). \quad (14.19)$$

where $\delta > 0$ is a scalar steplength (to be discussed next). The starting value $\boldsymbol{\lambda}_1$ can be set to the zero vector. The iteration (14.19) is continued until the gradient $\mathbf{R}_\lambda(\boldsymbol{\beta}, \boldsymbol{\lambda}_j)$ is smaller than some prespecified tolerance.

Efficient convergence requires a good choice of steplength δ . One method uses the following quadratic approximation. Set $\delta_0 = 0$, $\delta_1 = \frac{1}{2}$ and $\delta_2 = 1$. For $p = 0, 1, 2$, set

$$\boldsymbol{\lambda}_p = \boldsymbol{\lambda}_j - \delta_p (\mathbf{R}_{\lambda\lambda}(\boldsymbol{\beta}, \boldsymbol{\lambda}_j))^{-1} \mathbf{R}_\lambda(\boldsymbol{\beta}, \boldsymbol{\lambda}_j)$$

$$R_p = R(\boldsymbol{\beta}, \boldsymbol{\lambda}_p)$$

A quadratic function can be fit exactly through these three points. The value of δ which minimizes this quadratic is

$$\hat{\delta} = \frac{R_2 + 3R_0 - 4R_1}{4R_2 + 4R_0 - 8R_1}.$$

yielding the steplength to be plugged into (14.19).

A complication is that $\boldsymbol{\lambda}$ must be constrained so that $0 \leq p_i \leq 1$ which holds if

$$n(1 + \boldsymbol{\lambda}' \mathbf{g}_i(\boldsymbol{\beta})) \geq 1 \quad (14.20)$$

for all i . If (14.20) fails, the stepsize δ needs to be decreased.

Outer Loop

The outer loop is the minimization (14.6). This can be done by the modified Newton method described in the previous section. The gradient for (14.6) is

$$\mathbf{R}_\beta = \frac{\partial}{\partial \boldsymbol{\beta}} R(\boldsymbol{\beta}) = \frac{\partial}{\partial \boldsymbol{\beta}} R(\boldsymbol{\beta}, \boldsymbol{\lambda}) = \mathbf{R}_\beta + \boldsymbol{\lambda}'_\beta \mathbf{R}_\lambda = \mathbf{R}_\beta$$

since $\mathbf{R}_\lambda(\boldsymbol{\beta}, \boldsymbol{\lambda}) = 0$ at $\boldsymbol{\lambda} = \boldsymbol{\lambda}(\boldsymbol{\beta})$, where

$$\boldsymbol{\lambda}_\beta = \frac{\partial}{\partial \boldsymbol{\beta}'} \boldsymbol{\lambda}(\boldsymbol{\beta}) = -\mathbf{R}_{\lambda\lambda}^{-1} \mathbf{R}_{\lambda\beta},$$

the second equality following from the implicit function theorem applied to $\mathbf{R}_\lambda(\boldsymbol{\beta}, \boldsymbol{\lambda}(\boldsymbol{\beta})) = 0$.

The Hessian for (14.6) is

$$\begin{aligned} \mathbf{R}_{\beta\beta} &= -\frac{\partial^2}{\partial \boldsymbol{\beta} \partial \boldsymbol{\beta}'} R(\boldsymbol{\beta}) \\ &= -\frac{\partial}{\partial \boldsymbol{\beta}'} [\mathbf{R}_\beta(\boldsymbol{\beta}, \boldsymbol{\lambda}(\boldsymbol{\beta})) + \boldsymbol{\lambda}'_\beta \mathbf{R}_\lambda(\boldsymbol{\beta}, \boldsymbol{\lambda}(\boldsymbol{\beta}))] \\ &= -(\mathbf{R}_{\beta\beta}(\boldsymbol{\beta}, \boldsymbol{\lambda}(\boldsymbol{\beta})) + \mathbf{R}'_{\lambda\beta} \boldsymbol{\lambda}_\beta + \boldsymbol{\lambda}'_\beta \mathbf{R}_{\lambda\beta} + \boldsymbol{\lambda}'_\beta \mathbf{R}_{\lambda\lambda} \boldsymbol{\lambda}_\beta) \\ &= \mathbf{R}'_{\lambda\beta} \mathbf{R}_{\lambda\lambda}^{-1} \mathbf{R}_{\lambda\beta} - \mathbf{R}_{\beta\beta}. \end{aligned}$$

It is not guaranteed that $\mathbf{R}_{\beta\beta} > 0$. If not, the eigenvalues of $\mathbf{R}_{\beta\beta}$ should be adjusted so that all are positive. The Newton iteration rule is

$$\boldsymbol{\beta}_{j+1} = \boldsymbol{\beta}_j - \delta \mathbf{R}_{\beta\beta}^{-1} \mathbf{R}_\beta$$

where δ is a scalar stepsize, and the rule is iterated until convergence.

Chapter 15

Endogeneity

We say that there is endogeneity in the linear model $y = \mathbf{x}'_i \boldsymbol{\beta} + e_i$ if $\boldsymbol{\beta}$ is the parameter of interest and $\mathbb{E}(\mathbf{x}_i e_i) \neq 0$. This cannot happen if $\boldsymbol{\beta}$ is defined by linear projection, so requires a structural interpretation. The coefficient $\boldsymbol{\beta}$ must have meaning separately from the definition of a conditional mean or linear projection.

Example: Measurement error in the regressor. Suppose that (y_i, \mathbf{x}_i^*) are joint random variables, $\mathbb{E}(y_i | \mathbf{x}_i^*) = \mathbf{x}_i^{*'} \boldsymbol{\beta}$ is linear, $\boldsymbol{\beta}$ is the parameter of interest, and \mathbf{x}_i^* is not observed. Instead we observe $\mathbf{x}_i = \mathbf{x}_i^* + \mathbf{u}_i$ where \mathbf{u}_i is an $k \times 1$ measurement error, independent of y_i and \mathbf{x}_i^* . Then

$$\begin{aligned} y_i &= \mathbf{x}_i^{*'} \boldsymbol{\beta} + e_i \\ &= (\mathbf{x}_i - \mathbf{u}_i)' \boldsymbol{\beta} + e_i \\ &= \mathbf{x}_i' \boldsymbol{\beta} + v_i \end{aligned}$$

where

$$v_i = e_i - \mathbf{u}_i' \boldsymbol{\beta}.$$

The problem is that

$$\mathbb{E}(\mathbf{x}_i v_i) = \mathbb{E}[(\mathbf{x}_i^* + \mathbf{u}_i)(e_i - \mathbf{u}_i' \boldsymbol{\beta})] = -\mathbb{E}(\mathbf{u}_i \mathbf{u}_i') \boldsymbol{\beta} \neq 0$$

if $\boldsymbol{\beta} \neq 0$ and $\mathbb{E}(\mathbf{u}_i \mathbf{u}_i') \neq 0$. It follows that if $\hat{\boldsymbol{\beta}}$ is the OLS estimator, then

$$\hat{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\beta}^* = \boldsymbol{\beta} - (\mathbb{E}(\mathbf{x}_i \mathbf{x}_i'))^{-1} \mathbb{E}(\mathbf{u}_i \mathbf{u}_i') \boldsymbol{\beta} \neq \boldsymbol{\beta}.$$

This is called **measurement error bias**.

Example: Supply and Demand. The variables q_i and p_i (quantity and price) are determined jointly by the demand equation

$$q_i = -\beta_1 p_i + e_{1i}$$

and the supply equation

$$q_i = \beta_2 p_i + e_{2i}.$$

Assume that $\mathbf{e}_i = \begin{pmatrix} e_{1i} \\ e_{2i} \end{pmatrix}$ is iid, $\mathbb{E} \mathbf{e}_i = \mathbf{0}$ and $\mathbb{E} \mathbf{e}_i \mathbf{e}_i' = \mathbf{I}_2$ (the latter for simplicity). The question is, if we regress q_i on p_i , what happens?

It is helpful to solve for q_i and p_i in terms of the errors. In matrix notation,

$$\begin{bmatrix} 1 & \beta_1 \\ 1 & -\beta_2 \end{bmatrix} \begin{pmatrix} q_i \\ p_i \end{pmatrix} = \begin{pmatrix} e_{1i} \\ e_{2i} \end{pmatrix}$$

so

$$\begin{aligned} \begin{pmatrix} q_i \\ p_i \end{pmatrix} &= \begin{bmatrix} 1 & \beta_1 \\ 1 & -\beta_2 \end{bmatrix}^{-1} \begin{pmatrix} e_{1i} \\ e_{2i} \end{pmatrix} \\ &= \begin{bmatrix} \beta_2 & \beta_1 \\ 1 & -1 \end{bmatrix} \begin{pmatrix} e_{1i} \\ e_{2i} \end{pmatrix} \\ &= \begin{pmatrix} (\beta_2 e_{1i} + \beta_1 e_{2i}) / (\beta_1 + \beta_2) \\ (e_{1i} - e_{2i}) / (\beta_1 + \beta_2) \end{pmatrix}. \end{aligned}$$

The projection of q_i on p_i yields

$$\begin{aligned} q_i &= \beta^* p_i + \varepsilon_i \\ \mathbb{E}(p_i \varepsilon_i) &= 0 \end{aligned}$$

where

$$\beta^* = \frac{\mathbb{E}(p_i q_i)}{\mathbb{E}(p_i^2)} = \frac{\beta_2 - \beta_1}{2}$$

Hence if it is estimated by OLS, $\hat{\beta} \xrightarrow{p} \beta^*$, which does not equal either β_1 or β_2 . This is called **simultaneous equations bias**.

15.1 Instrumental Variables

Let the equation of interest be

$$y_i = \mathbf{x}'_i \boldsymbol{\beta} + e_i \quad (15.1)$$

where \mathbf{x}_i is $k \times 1$, and assume that $\mathbb{E}(\mathbf{x}_i e_i) \neq \mathbf{0}$ so there is **endogeneity**. We call (15.1) the structural equation. In matrix notation, this can be written as

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}. \quad (15.2)$$

Any solution to the problem of endogeneity requires additional information which we call **instruments**.

Definition 15.1.1 The $\ell \times 1$ random vector \mathbf{z}_i is an **instrumental variable** for (15.1) if $\mathbb{E}(\mathbf{z}_i e_i) = \mathbf{0}$.

In a typical set-up, some regressors in \mathbf{x}_i will be uncorrelated with e_i (for example, at least the intercept). Thus we make the partition

$$\mathbf{x}_i = \begin{pmatrix} \mathbf{x}_{1i} \\ \mathbf{x}_{2i} \end{pmatrix} \begin{matrix} k_1 \\ k_2 \end{matrix} \quad (15.3)$$

where $\mathbb{E}(\mathbf{x}_{1i} e_i) = \mathbf{0}$ yet $\mathbb{E}(\mathbf{x}_{2i} e_i) \neq \mathbf{0}$. We call \mathbf{x}_{1i} exogenous and \mathbf{x}_{2i} endogenous. By the above definition, \mathbf{x}_{1i} is an instrumental variable for (15.1), so should be included in \mathbf{z}_i . So we have the partition

$$\mathbf{z}_i = \begin{pmatrix} \mathbf{x}_{1i} \\ \mathbf{z}_{2i} \end{pmatrix} \begin{matrix} k_1 \\ \ell_2 \end{matrix} \quad (15.4)$$

where $\mathbf{x}_{1i} = \mathbf{z}_{1i}$ are the **included exogenous variables**, and \mathbf{z}_{2i} are the **excluded exogenous variables**. That is \mathbf{z}_{2i} are variables which could be included in the equation for y_i (in the sense that they are uncorrelated with e_i) yet can be *excluded*, as they would have true zero coefficients in the equation.

The model is **just-identified** if $\ell = k$ (i.e., if $\ell_2 = k_2$) and **over-identified** if $\ell > k$ (i.e., if $\ell_2 > k_2$).

We have noted that any solution to the problem of endogeneity requires instruments. This does not mean that valid instruments actually exist.

15.2 Reduced Form

The reduced form relationship between the variables or “regressors” \mathbf{x}_i and the instruments \mathbf{z}_i is found by linear projection. Let

$$\mathbf{\Gamma} = \mathbb{E}(\mathbf{z}_i \mathbf{z}_i')^{-1} \mathbb{E}(\mathbf{z}_i \mathbf{x}_i')$$

be the $\ell \times k$ matrix of coefficients from a projection of \mathbf{x}_i on \mathbf{z}_i , and define

$$\mathbf{u}_i = \mathbf{x}_i - \mathbf{\Gamma}' \mathbf{z}_i$$

as the projection error. Then the reduced form linear relationship between \mathbf{x}_i and \mathbf{z}_i is

$$\mathbf{x}_i = \mathbf{\Gamma}' \mathbf{z}_i + \mathbf{u}_i. \quad (15.5)$$

In matrix notation, we can write (15.5) as

$$\mathbf{X} = \mathbf{Z}\mathbf{\Gamma} + \mathbf{U} \quad (15.6)$$

where \mathbf{U} is $n \times k$.

By construction,

$$\mathbb{E}(\mathbf{z}_i \mathbf{u}_i') = \mathbf{0},$$

so (15.5) is a projection and can be estimated by OLS:

$$\mathbf{x}_i = \hat{\mathbf{\Gamma}}' \mathbf{z}_i + \hat{\mathbf{u}}_i.$$

or

$$\mathbf{X} = \mathbf{Z}\hat{\mathbf{\Gamma}} + \hat{\mathbf{U}}$$

where

$$\hat{\mathbf{\Gamma}} = (\mathbf{Z}'\mathbf{Z})^{-1} (\mathbf{Z}'\mathbf{X}).$$

Substituting (15.6) into (15.2), we find

$$\begin{aligned} \mathbf{y} &= (\mathbf{Z}\mathbf{\Gamma} + \mathbf{U})\boldsymbol{\beta} + \mathbf{e} \\ &= \mathbf{Z}\boldsymbol{\lambda} + \mathbf{v}, \end{aligned} \quad (15.7)$$

where

$$\boldsymbol{\lambda} = \mathbf{\Gamma}\boldsymbol{\beta} \quad (15.8)$$

and

$$\mathbf{v} = \mathbf{U}\boldsymbol{\beta} + \mathbf{e}.$$

Observe that

$$\mathbb{E}(\mathbf{z}_i \mathbf{v}_i) = \mathbb{E}(\mathbf{z}_i \mathbf{u}_i') \boldsymbol{\beta} + \mathbb{E}(\mathbf{z}_i \mathbf{e}_i) = \mathbf{0}.$$

Thus (15.7) is a projection equation and may be estimated by OLS. This is

$$\begin{aligned} \mathbf{y} &= \mathbf{Z}\hat{\boldsymbol{\lambda}} + \hat{\mathbf{v}}, \\ \hat{\boldsymbol{\lambda}} &= (\mathbf{Z}'\mathbf{Z})^{-1} (\mathbf{Z}'\mathbf{y}) \end{aligned}$$

The equation (15.7) is the reduced form for \mathbf{y} . (15.6) and (15.7) together are the **reduced form equations** for the system

$$\begin{aligned} \mathbf{y} &= \mathbf{Z}\boldsymbol{\lambda} + \mathbf{v} \\ \mathbf{X} &= \mathbf{Z}\mathbf{\Gamma} + \mathbf{U}. \end{aligned}$$

As we showed above, OLS yields the reduced-form estimates $(\hat{\boldsymbol{\lambda}}, \hat{\mathbf{\Gamma}})$

15.3 Identification

The structural parameter β relates to (λ, Γ) through (15.8). The parameter β is **identified**, meaning that it can be recovered from the reduced form, if

$$\text{rank}(\Gamma) = k. \quad (15.9)$$

Assume that (15.9) holds. If $\ell = k$, then $\beta = \Gamma^{-1}\lambda$. If $\ell > k$, then for any $\mathbf{W} > 0$, $\beta = (\Gamma' \mathbf{W} \Gamma)^{-1} \Gamma' \mathbf{W} \lambda$.

If (15.9) is not satisfied, then β cannot be recovered from (λ, Γ) . Note that a necessary (although not sufficient) condition for (15.9) is $\ell \geq k$.

Since \mathbf{Z} and \mathbf{X} have the common variables \mathbf{X}_1 , we can rewrite some of the expressions. Using (15.3) and (15.4) to make the matrix partitions $\mathbf{Z} = [\mathbf{Z}_1, \mathbf{Z}_2]$ and $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2]$, we can partition Γ as

$$\begin{aligned} \Gamma &= \begin{bmatrix} \Gamma_{11} & \Gamma_{12} \\ \Gamma_{21} & \Gamma_{22} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{I} & \Gamma_{12} \\ \mathbf{0} & \Gamma_{22} \end{bmatrix} \end{aligned}$$

(15.6) can be rewritten as

$$\begin{aligned} \mathbf{X}_1 &= \mathbf{Z}_1 \\ \mathbf{X}_2 &= \mathbf{Z}_1 \Gamma_{12} + \mathbf{Z}_2 \Gamma_{22} + \mathbf{U}_2. \end{aligned} \quad (15.10)$$

β is identified if $\text{rank}(\Gamma) = k$, which is true if and only if $\text{rank}(\Gamma_{22}) = k_2$ (by the upper-diagonal structure of Γ). Thus the key to identification of the model rests on the $\ell_2 \times k_2$ matrix Γ_{22} in (15.10).

15.4 Estimation

The model can be written as

$$\begin{aligned} y_i &= \mathbf{x}'_i \beta + e_i \\ \mathbb{E}(\mathbf{z}_i e_i) &= \mathbf{0} \end{aligned}$$

or

$$\begin{aligned} \mathbb{E} \mathbf{g}_i(\beta) &= \mathbf{0} \\ \mathbf{g}_i(\beta) &= \mathbf{z}_i (y_i - \mathbf{x}'_i \beta). \end{aligned}$$

This is a moment condition model. Appropriate estimators include GMM and EL. The estimators and distribution theory developed in those Chapter 8 and 9 directly apply. Recall that the GMM estimator, for given weight matrix \mathbf{W}_n , is

$$\hat{\beta} = (\mathbf{X}' \mathbf{Z} \mathbf{W}_n \mathbf{Z}' \mathbf{X})^{-1} \mathbf{X}' \mathbf{Z} \mathbf{W}_n \mathbf{Z}' \mathbf{y}.$$

15.5 Special Cases: IV and 2SLS

If the model is just-identified, so that $k = \ell$, then the formula for GMM simplifies. We find that

$$\begin{aligned} \hat{\beta} &= (\mathbf{X}' \mathbf{Z} \mathbf{W}_n \mathbf{Z}' \mathbf{X})^{-1} \mathbf{X}' \mathbf{Z} \mathbf{W}_n \mathbf{Z}' \mathbf{y} \\ &= (\mathbf{Z}' \mathbf{X})^{-1} \mathbf{W}_n^{-1} (\mathbf{X}' \mathbf{Z})^{-1} \mathbf{X}' \mathbf{Z} \mathbf{W}_n \mathbf{Z}' \mathbf{y} \\ &= (\mathbf{Z}' \mathbf{X})^{-1} \mathbf{Z}' \mathbf{y} \end{aligned}$$

This estimator is often called the **instrumental variables estimator** (IV) of β , where \mathbf{Z} is used as an instrument for \mathbf{X} . Observe that the weight matrix \mathbf{W}_n has disappeared. In the just-identified case, the weight matrix places no role. This is also the method of moments estimator of β , and the EL estimator. Another interpretation stems from the fact that since $\beta = \Gamma^{-1}\lambda$, we can construct the **Indirect Least Squares** (ILS) estimator:

$$\begin{aligned}\widehat{\beta} &= \widehat{\Gamma}^{-1}\widehat{\lambda} \\ &= \left((\mathbf{Z}'\mathbf{Z})^{-1} (\mathbf{Z}'\mathbf{X}) \right)^{-1} \left((\mathbf{Z}'\mathbf{Z})^{-1} (\mathbf{Z}'\mathbf{y}) \right) \\ &= (\mathbf{Z}'\mathbf{X})^{-1} (\mathbf{Z}'\mathbf{Z}) (\mathbf{Z}'\mathbf{Z})^{-1} (\mathbf{Z}'\mathbf{y}) \\ &= (\mathbf{Z}'\mathbf{X})^{-1} (\mathbf{Z}'\mathbf{y}).\end{aligned}$$

which again is the IV estimator.

Recall that the optimal weight matrix is an estimate of the inverse of $\mathbf{\Omega} = \mathbb{E}(\mathbf{z}_i \mathbf{z}_i' e_i^2)$. In the special case that $\mathbb{E}(e_i^2 | \mathbf{z}_i) = \sigma^2$ (homoskedasticity), then $\mathbf{\Omega} = \mathbb{E}(\mathbf{z}_i \mathbf{z}_i') \sigma^2 \propto \mathbb{E}(\mathbf{z}_i \mathbf{z}_i')$ suggesting the weight matrix $\mathbf{W}_n = (\mathbf{Z}'\mathbf{Z})^{-1}$. Using this choice, the GMM estimator equals

$$\widehat{\beta}_{2SLS} = \left(\mathbf{X}'\mathbf{Z} (\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'\mathbf{X} \right)^{-1} \mathbf{X}'\mathbf{Z} (\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'\mathbf{y}$$

This is called the **two-stage-least squares** (2SLS) estimator. It was originally proposed by Theil (1953) and Basmann (1957), and is the classic estimator for linear equations with instruments. Under the homoskedasticity assumption, the 2SLS estimator is efficient GMM, but otherwise it is inefficient.

It is useful to observe that writing

$$\begin{aligned}\mathbf{P} &= \mathbf{Z} (\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}' \\ \widehat{\mathbf{X}} &= \mathbf{P}\mathbf{X} = \mathbf{Z}\widehat{\Gamma}\end{aligned}$$

then the 2SLS estimator is

$$\begin{aligned}\widehat{\beta} &= (\mathbf{X}'\mathbf{P}\mathbf{X})^{-1} \mathbf{X}'\mathbf{P}\mathbf{y} \\ &= (\widehat{\mathbf{X}}'\widehat{\mathbf{X}})^{-1} \widehat{\mathbf{X}}'\mathbf{y}.\end{aligned}$$

The source of the “two-stage” name is since it can be computed as follows

- First regress \mathbf{X} on \mathbf{Z} , vis., $\widehat{\Gamma} = (\mathbf{Z}'\mathbf{Z})^{-1} (\mathbf{Z}'\mathbf{X})$ and $\widehat{\mathbf{X}} = \mathbf{Z}\widehat{\Gamma} = \mathbf{P}\mathbf{X}$.
- Second, regress \mathbf{y} on $\widehat{\mathbf{X}}$, vis., $\widehat{\beta} = (\widehat{\mathbf{X}}'\widehat{\mathbf{X}})^{-1} \widehat{\mathbf{X}}'\mathbf{y}$.

It is useful to scrutinize the projection $\widehat{\mathbf{X}}$. Recall, $\mathbf{X} = [\mathbf{X}_1, \mathbf{X}_2]$ and $\mathbf{Z} = [\mathbf{X}_1, \mathbf{Z}_2]$. Then

$$\begin{aligned}\widehat{\mathbf{X}} &= [\widehat{\mathbf{X}}_1, \widehat{\mathbf{X}}_2] \\ &= [\mathbf{P}\mathbf{X}_1, \mathbf{P}\mathbf{X}_2] \\ &= [\mathbf{X}_1, \mathbf{P}\mathbf{X}_2] \\ &= [\mathbf{X}_1, \widehat{\mathbf{X}}_2],\end{aligned}$$

since \mathbf{X}_1 lies in the span of \mathbf{Z} . Thus in the second stage, we regress \mathbf{y} on \mathbf{X}_1 and $\widehat{\mathbf{X}}_2$. So only the endogenous variables \mathbf{X}_2 are replaced by their fitted values:

$$\widehat{\mathbf{X}}_2 = \mathbf{Z}_1\widehat{\Gamma}_{12} + \mathbf{Z}_2\widehat{\Gamma}_{22}.$$

15.6 Bekker Asymptotics

Bekker (1994) used an alternative asymptotic framework to analyze the finite-sample bias in the 2SLS estimator. Here we present a simplified version of one of his results. In our notation, the model is

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e} \quad (15.11)$$

$$\mathbf{X} = \mathbf{Z}\boldsymbol{\Gamma} + \mathbf{U} \quad (15.12)$$

$$\boldsymbol{\xi} = (\mathbf{e}, \mathbf{U})$$

$$\mathbb{E}(\boldsymbol{\xi} \mid \mathbf{Z}) = \mathbf{0}$$

$$\mathbb{E}(\boldsymbol{\xi}'\boldsymbol{\xi} \mid \mathbf{Z}) = \mathbf{S}$$

As before, \mathbf{Z} is $n \times l$ so there are l instruments.

First, let's analyze the approximate bias of OLS applied to (15.11). Using (15.12),

$$\mathbb{E}\left(\frac{1}{n}\mathbf{X}'\mathbf{e}\right) = \mathbb{E}(\mathbf{x}_i e_i) = \boldsymbol{\Gamma}'\mathbb{E}(\mathbf{z}_i e_i) + \mathbb{E}(\mathbf{u}_i e_i) = \mathbf{s}_{21}$$

and

$$\begin{aligned} \mathbb{E}\left(\frac{1}{n}\mathbf{X}'\mathbf{X}\right) &= \mathbb{E}(\mathbf{x}_i \mathbf{x}_i') \\ &= \boldsymbol{\Gamma}'\mathbb{E}(\mathbf{z}_i \mathbf{z}_i') \boldsymbol{\Gamma} + \mathbb{E}(\mathbf{u}_i \mathbf{z}_i') \boldsymbol{\Gamma} + \boldsymbol{\Gamma}'\mathbb{E}(\mathbf{z}_i \mathbf{u}_i') + \mathbb{E}(\mathbf{u}_i \mathbf{u}_i') \\ &= \boldsymbol{\Gamma}'\mathbf{Q}\boldsymbol{\Gamma} + \mathbf{S}_{22} \end{aligned}$$

where $\mathbf{Q} = \mathbb{E}(\mathbf{z}_i \mathbf{z}_i')$. Hence by a first-order approximation

$$\begin{aligned} \mathbb{E}(\hat{\boldsymbol{\beta}}_{OLS} - \boldsymbol{\beta}) &\approx \left(\mathbb{E}\left(\frac{1}{n}\mathbf{X}'\mathbf{X}\right)\right)^{-1} \mathbb{E}\left(\frac{1}{n}\mathbf{X}'\mathbf{e}\right) \\ &= (\boldsymbol{\Gamma}'\mathbf{Q}\boldsymbol{\Gamma} + \mathbf{S}_{22})^{-1} \mathbf{s}_{21} \end{aligned} \quad (15.13)$$

which is zero only when $\mathbf{s}_{21} = \mathbf{0}$ (when \mathbf{X} is exogenous).

We now derive a similar result for the 2SLS estimator.

$$\hat{\boldsymbol{\beta}}_{2SLS} = (\mathbf{X}'\mathbf{P}\mathbf{X})^{-1} (\mathbf{X}'\mathbf{P}\mathbf{y}).$$

Let $\mathbf{P} = \mathbf{Z}(\mathbf{Z}'\mathbf{Z})^{-1}\mathbf{Z}'$. By the spectral decomposition of an idempotent matrix, $\mathbf{P} = \mathbf{H}\boldsymbol{\Lambda}\mathbf{H}'$ where $\boldsymbol{\Lambda} = \text{diag}(\mathbf{I}_l, \mathbf{0})$. Let $\mathbf{Q} = \mathbf{H}'\boldsymbol{\xi}\mathbf{S}^{-1/2}$ which satisfies $\mathbb{E}\mathbf{Q}'\mathbf{Q} = \mathbf{I}_n$ and partition $\mathbf{Q} = (\mathbf{q}'_1 \mathbf{Q}'_2)$ where \mathbf{q}_1 is $l \times 1$. Hence

$$\begin{aligned} \mathbb{E}\left(\frac{1}{n}\boldsymbol{\xi}'\mathbf{P}\boldsymbol{\xi} \mid \mathbf{Z}\right) &= \frac{1}{n}\mathbf{S}^{1/2'}\mathbb{E}(\mathbf{Q}'\boldsymbol{\Lambda}\mathbf{Q} \mid \mathbf{Z})\mathbf{S}^{1/2} \\ &= \frac{1}{n}\mathbf{S}^{1/2'}\mathbb{E}\left(\frac{1}{n}\mathbf{q}'_1 \mathbf{q}_1\right)\mathbf{S}^{1/2} \\ &= \frac{l}{n}\mathbf{S}^{1/2'}\mathbf{S}^{1/2} \\ &= \alpha\mathbf{S} \end{aligned}$$

where

$$\alpha = \frac{l}{n}.$$

Using (15.12) and this result,

$$\frac{1}{n}\mathbb{E}(\mathbf{X}'\mathbf{P}\mathbf{e}) = \frac{1}{n}\mathbb{E}(\boldsymbol{\Gamma}'\mathbf{Z}'\mathbf{e}) + \frac{1}{n}\mathbb{E}(\mathbf{U}'\mathbf{P}\mathbf{e}) = \alpha\mathbf{s}_{21},$$

and

$$\begin{aligned}\frac{1}{n}\mathbb{E}(\mathbf{X}'\mathbf{P}\mathbf{X}) &= \mathbf{\Gamma}'\mathbb{E}(\mathbf{z}_i\mathbf{z}_i')\mathbf{\Gamma} + \mathbf{\Gamma}'\mathbb{E}(\mathbf{z}_i\mathbf{u}_i) + \mathbb{E}(\mathbf{u}_i\mathbf{z}_i')\mathbf{\Gamma} + \frac{1}{n}\mathbb{E}(\mathbf{U}'\mathbf{P}\mathbf{U}) \\ &= \mathbf{\Gamma}'\mathbf{Q}\mathbf{\Gamma} + \alpha\mathbf{S}_{22}.\end{aligned}$$

Together

$$\begin{aligned}\mathbb{E}(\hat{\boldsymbol{\beta}}_{2SLS} - \boldsymbol{\beta}) &\approx \left(\mathbb{E}\left(\frac{1}{n}\mathbf{X}'\mathbf{P}\mathbf{X}\right)\right)^{-1} \mathbb{E}\left(\frac{1}{n}\mathbf{X}'\mathbf{P}\mathbf{e}\right) \\ &= \alpha(\mathbf{\Gamma}'\mathbf{Q}\mathbf{\Gamma} + \alpha\mathbf{S}_{22})^{-1} \mathbf{s}_{21}.\end{aligned}\tag{15.14}$$

In general this is non-zero, except when $\mathbf{s}_{21} = 0$ (when \mathbf{X} is exogenous). It is also close to zero when $\alpha = 0$. Bekker (1994) pointed out that it also has the reverse implication – that when $\alpha = l/n$ is large, the bias in the 2SLS estimator will be large. Indeed as $\alpha \rightarrow 1$, the expression in (15.14) approaches that in (15.13), indicating that the bias in 2SLS approaches that of OLS as the number of instruments increases.

Bekker (1994) showed further that under the alternative asymptotic approximation that α is fixed as $n \rightarrow \infty$ (so that the number of instruments goes to infinity proportionately with sample size) then the expression in (15.14) is the probability limit of $\hat{\boldsymbol{\beta}}_{2SLS} - \boldsymbol{\beta}$

15.7 Identification Failure

Recall the reduced form equation

$$\mathbf{X}_2 = \mathbf{Z}_1\mathbf{\Gamma}_{12} + \mathbf{Z}_2\mathbf{\Gamma}_{22} + \mathbf{U}_2.$$

The parameter $\boldsymbol{\beta}$ fails to be identified if $\mathbf{\Gamma}_{22}$ has deficient rank. The consequences of identification failure for inference are quite severe.

Take the simplest case where $k = l = 1$ (so there is no \mathbf{Z}_1). Then the model may be written as

$$\begin{aligned}y_i &= x_i\beta + e_i \\ x_i &= z_i\gamma + u_i\end{aligned}$$

and $\mathbf{\Gamma}_{22} = \gamma = \mathbb{E}(z_i x_i) / \mathbb{E}z_i^2$. We see that β is identified if and only if $\gamma \neq 0$, which occurs when $\mathbb{E}(x_i z_i) \neq 0$. Thus identification hinges on the existence of correlation between the excluded exogenous variable and the included endogenous variable.

Suppose this condition fails, so $\mathbb{E}(x_i z_i) = 0$. Then by the CLT

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n z_i e_i \xrightarrow{d} N_1 \sim N(0, \mathbb{E}(z_i^2 e_i^2))\tag{15.15}$$

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n z_i x_i = \frac{1}{\sqrt{n}} \sum_{i=1}^n z_i u_i \xrightarrow{d} N_2 \sim N(0, \mathbb{E}(z_i^2 u_i^2))\tag{15.16}$$

therefore

$$\hat{\beta} - \beta = \frac{\frac{1}{\sqrt{n}} \sum_{i=1}^n z_i e_i}{\frac{1}{\sqrt{n}} \sum_{i=1}^n z_i x_i} \xrightarrow{d} \frac{N_1}{N_2} \sim \text{Cauchy},$$

since the ratio of two normals is Cauchy. This is particularly nasty, as the Cauchy distribution does not have a finite mean. This result carries over to more general settings, and was examined by Phillips (1989) and Choi and Phillips (1992).

Suppose that identification does not completely fail, but is *weak*. This occurs when Γ_{22} is full rank, but *small*. This can be handled in an asymptotic analysis by modeling it as local-to-zero, viz

$$\mathbf{\Gamma}_{22} = n^{-1/2}\mathbf{C},$$

where \mathbf{C} is a full rank matrix. The $n^{-1/2}$ is picked because it provides just the right balancing to allow a rich distribution theory.

To see the consequences, once again take the simple case $k = l = 1$. Here, the instrument x_i is weak for z_i if

$$\gamma = n^{-1/2}c.$$

Then (15.15) is unaffected, but (15.16) instead takes the form

$$\begin{aligned} \frac{1}{\sqrt{n}} \sum_{i=1}^n z_i x_i &= \frac{1}{\sqrt{n}} \sum_{i=1}^n z_i^2 \gamma + \frac{1}{\sqrt{n}} \sum_{i=1}^n z_i u_i \\ &= \frac{1}{n} \sum_{i=1}^n z_i^2 c + \frac{1}{\sqrt{n}} \sum_{i=1}^n z_i u_i \\ &\xrightarrow{d} Qc + N_2 \end{aligned}$$

therefore

$$\hat{\beta} - \beta \xrightarrow{d} \frac{N_1}{Qc + N_2}.$$

As in the case of complete identification failure, we find that $\hat{\beta}$ is inconsistent for β and the asymptotic distribution of $\hat{\beta}$ is non-normal. In addition, standard test statistics have non-standard distributions, meaning that inferences about parameters of interest can be misleading.

The distribution theory for this model was developed by Staiger and Stock (1997) and extended to nonlinear GMM estimation by Stock and Wright (2000). Further results on testing were obtained by Wang and Zivot (1998).

The bottom line is that it is highly desirable to avoid identification failure. Once again, the equation to focus on is the reduced form

$$\mathbf{X}_2 = \mathbf{Z}_1 \mathbf{\Gamma}_{12} + \mathbf{Z}_2 \mathbf{\Gamma}_{22} + \mathbf{U}_2$$

and identification requires $\text{rank}(\mathbf{\Gamma}_{22}) = k_2$. If $k_2 = 1$, this requires $\Gamma_{22} \neq \mathbf{0}$, which is straightforward to assess using a hypothesis test on the reduced form. Therefore in the case of $k_2 = 1$ (one RHS endogenous variable), one constructive recommendation is to explicitly estimate the reduced form equation for \mathbf{X}_2 , construct the test of $\mathbf{\Gamma}_{22} = \mathbf{0}$, and at a minimum check that the test rejects $\mathbb{H}_0 : \mathbf{\Gamma}_{22} = \mathbf{0}$.

When $k_2 > 1$, $\mathbf{\Gamma}_{22} \neq \mathbf{0}$ is not sufficient for identification. It is not even sufficient that each column of $\mathbf{\Gamma}_{22}$ is non-zero (each column corresponds to a distinct endogenous variable in \mathbf{Z}_2). So while a minimal check is to test that each columns of $\mathbf{\Gamma}_{22}$ is non-zero, this cannot be interpreted as definitive proof that $\mathbf{\Gamma}_{22}$ has full rank. Unfortunately, tests of deficient rank are difficult to implement. In any event, it appears reasonable to explicitly estimate and report the reduced form equations for \mathbf{Z}_2 , and attempt to assess the likelihood that $\mathbf{\Gamma}_{22}$ has deficient rank.

Exercises

1. Consider the single equation model

$$y_i = z_i\beta + e_i,$$

where y_i and z_i are both real-valued (1×1). Let $\hat{\beta}$ denote the IV estimator of β using as an instrument a dummy variable d_i (takes only the values 0 and 1). Find a simple expression for the IV estimator in this context.

2. In the linear model

$$y_i = \mathbf{x}'_i\boldsymbol{\beta} + e_i$$

$$\mathbb{E}(e_i | \mathbf{x}_i) = 0$$

suppose $\sigma_i^2 = \mathbb{E}(e_i^2 | x_i)$ is known. Show that the GLS estimator of $\boldsymbol{\beta}$ can be written as an IV estimator using some instrument \mathbf{z}_i . (Find an expression for \mathbf{z}_i .)

3. Take the linear model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}.$$

Let the OLS estimator for $\boldsymbol{\beta}$ be $\hat{\boldsymbol{\beta}}$ and the OLS residual be $\hat{\mathbf{e}} = \mathbf{y} - \mathbf{X}\hat{\boldsymbol{\beta}}$.

Let the IV estimator for $\boldsymbol{\beta}$ using some instrument \mathbf{Z} be $\tilde{\boldsymbol{\beta}}$ and the IV residual be $\tilde{\mathbf{e}} = \mathbf{y} - \mathbf{X}\tilde{\boldsymbol{\beta}}$. If \mathbf{X} is indeed endogenous, will IV “fit” better than OLS, in the sense that $\tilde{\mathbf{e}}'\tilde{\mathbf{e}} < \hat{\mathbf{e}}'\hat{\mathbf{e}}$, at least in large samples?

4. The reduced form between the regressors \mathbf{x}_i and instruments \mathbf{z}_i takes the form

$$\mathbf{x}_i = \boldsymbol{\Gamma}'\mathbf{z}_i + \mathbf{u}_i$$

or

$$\mathbf{X} = \mathbf{Z}\boldsymbol{\Gamma} + \mathbf{U}$$

where \mathbf{x}_i is $k \times 1$, \mathbf{z}_i is $l \times 1$, \mathbf{X} is $n \times k$, \mathbf{Z} is $n \times l$, \mathbf{U} is $n \times k$, and $\boldsymbol{\Gamma}$ is $l \times k$. The parameter $\boldsymbol{\Gamma}$ is defined by the population moment condition

$$\mathbb{E}(\mathbf{z}_i\mathbf{u}'_i) = \mathbf{0}$$

Show that the method of moments estimator for $\boldsymbol{\Gamma}$ is $\hat{\boldsymbol{\Gamma}} = (\mathbf{Z}'\mathbf{Z})^{-1}(\mathbf{Z}'\mathbf{X})$.

5. In the structural model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{e}$$

$$\mathbf{X} = \mathbf{Z}\boldsymbol{\Gamma} + \mathbf{U}$$

with $\boldsymbol{\Gamma}$ $l \times k$, $l \geq k$, we claim that $\boldsymbol{\beta}$ is identified (can be recovered from the reduced form) if $\text{rank}(\boldsymbol{\Gamma}) = k$. Explain why this is true. That is, show that if $\text{rank}(\boldsymbol{\Gamma}) < k$ then $\boldsymbol{\beta}$ cannot be identified.

6. Take the linear model

$$y_i = x_i\beta + e_i$$

$$\mathbb{E}(e_i | x_i) = 0.$$

where x_i and β are 1×1 .

- (a) Show that $\mathbb{E}(x_i e_i) = 0$ and $\mathbb{E}(x_i^2 e_i) = 0$. Is $\mathbf{z}_i = (x_i \quad x_i^2)'$ a valid instrumental variable for estimation of β ?
- (b) Define the 2SLS estimator of β , using \mathbf{z}_i as an instrument for x_i . How does this differ from OLS?
- (c) Find the efficient GMM estimator of β based on the moment condition

$$\mathbb{E}(\mathbf{z}_i (y_i - x_i \beta)) = \mathbf{0}.$$

Does this differ from 2SLS and/or OLS?

7. Suppose that price and quantity are determined by the intersection of the linear demand and supply curves

$$\text{Demand : } Q = a_0 + a_1 P + a_2 Y + e_1$$

$$\text{Supply : } Q = b_0 + b_1 P + b_2 W + e_2$$

where income (Y) and wage (W) are determined outside the market. In this model, are the parameters identified?

8. The data file `card.dat` is taken from Card (1995). There are 2215 observations with 29 variables, listed in `card.pdf`. We want to estimate a wage equation

$$\log(\text{Wage}) = \beta_0 + \beta_1 \text{Educ} + \beta_2 \text{Exper} + \beta_3 \text{Exper}^2 + \beta_4 \text{South} + \beta_5 \text{Black} + e$$

where $\text{Educ} = \text{Education}$ (Years) $\text{Exper} = \text{Experience}$ (Years), and South and Black are regional and racial dummy variables.

- (a) Estimate the model by OLS. Report estimates and standard errors.
- (b) Now treat *Education* as endogenous, and the remaining variables as exogenous. Estimate the model by 2SLS, using the instrument `near4`, a dummy indicating that the observation lives near a 4-year college. Report estimates and standard errors.
- (c) Re-estimate by 2SLS (report estimates and standard errors) adding three additional instruments: `near2` (a dummy indicating that the observation lives near a 2-year college), `fatheduc` (the education, in years, of the father) and `motheduc` (the education, in years, of the mother).
- (d) Re-estimate the model by efficient GMM. I suggest that you use the 2SLS estimates as the first-step to get the weight matrix, and then calculate the GMM estimator from this weight matrix without further iteration. Report the estimates and standard errors.
- (e) Calculate and report the J statistic for overidentification.
- (f) Discuss your findings.

Chapter 16

Univariate Time Series

A time series y_t is a process observed in sequence over time, $t = 1, \dots, T$. To indicate the dependence on time, we adopt new notation, and use the subscript t to denote the individual observation, and T to denote the number of observations.

Because of the sequential nature of time series, we expect that y_t and y_{t-1} are *not* independent, so classical assumptions are not valid.

We can separate time series into two categories: univariate ($y_t \in \mathbb{R}$ is scalar); and multivariate ($y_t \in \mathbb{R}^m$ is vector-valued). The primary model for univariate time series is autoregressions (ARs). The primary model for multivariate time series is vector autoregressions (VARs).

16.1 Stationarity and Ergodicity

Definition 16.1.1 $\{y_t\}$ is *covariance (weakly) stationary* if

$$\mathbb{E}(y_t) = \mu$$

is independent of t , and

$$\text{cov}(y_t, y_{t-k}) = \gamma(k)$$

is independent of t for all k . $\gamma(k)$ is called the **autocovariance function**.

$$\rho(k) = \gamma(k)/\gamma(0) = \text{corr}(y_t, y_{t-k})$$

is the **autocorrelation function**.

Definition 16.1.2 $\{y_t\}$ is **strictly stationary** if the joint distribution of (y_t, \dots, y_{t-k}) is independent of t for all k .

Definition 16.1.3 A stationary time series is **ergodic** if $\gamma(k) \rightarrow 0$ as $k \rightarrow \infty$.

The following two theorems are essential to the analysis of stationary time series. The proofs are rather difficult, however.

Theorem 16.1.1 *If y_t is strictly stationary and ergodic and $x_t = f(y_t, y_{t-1}, \dots)$ is a random variable, then x_t is strictly stationary and ergodic.*

Theorem 16.1.2 (*Ergodic Theorem*). *If y_t is strictly stationary and ergodic and $\mathbb{E}|y_t| < \infty$, then as $T \rightarrow \infty$,*

$$\frac{1}{T} \sum_{t=1}^T y_t \xrightarrow{p} \mathbb{E}(y_t).$$

This allows us to consistently estimate parameters using time-series moments:
The sample mean:

$$\hat{\mu} = \frac{1}{T} \sum_{t=1}^T y_t$$

The sample autocovariance

$$\hat{\gamma}(k) = \frac{1}{T} \sum_{t=1}^T (y_t - \hat{\mu})(y_{t-k} - \hat{\mu}).$$

The sample autocorrelation

$$\hat{\rho}(k) = \frac{\hat{\gamma}(k)}{\hat{\gamma}(0)}.$$

Theorem 16.1.3 *If y_t is strictly stationary and ergodic and $\mathbb{E}y_t^2 < \infty$, then as $T \rightarrow \infty$,*

1. $\hat{\mu} \xrightarrow{p} \mathbb{E}(y_t)$;
2. $\hat{\gamma}(k) \xrightarrow{p} \gamma(k)$;
3. $\hat{\rho}(k) \xrightarrow{p} \rho(k)$.

Proof of Theorem 16.1.3. Part (1) is a direct consequence of the Ergodic theorem. For Part (2), note that

$$\begin{aligned} \hat{\gamma}(k) &= \frac{1}{T} \sum_{t=1}^T (y_t - \hat{\mu})(y_{t-k} - \hat{\mu}) \\ &= \frac{1}{T} \sum_{t=1}^T y_t y_{t-k} - \frac{1}{T} \sum_{t=1}^T y_t \hat{\mu} - \frac{1}{T} \sum_{t=1}^T y_{t-k} \hat{\mu} + \hat{\mu}^2. \end{aligned}$$

By Theorem 16.1.1 above, the sequence $y_t y_{t-k}$ is strictly stationary and ergodic, and it has a finite mean by the assumption that $\mathbb{E}y_t^2 < \infty$. Thus an application of the Ergodic Theorem yields

$$\frac{1}{T} \sum_{t=1}^T y_t y_{t-k} \xrightarrow{p} \mathbb{E}(y_t y_{t-k}).$$

Thus

$$\hat{\gamma}(k) \xrightarrow{p} \mathbb{E}(y_t y_{t-k}) - \mu^2 - \mu^2 + \mu^2 = \mathbb{E}(y_t y_{t-k}) - \mu^2 = \gamma(k).$$

Part (3) follows by the continuous mapping theorem: $\hat{\rho}(k) = \hat{\gamma}(k)/\hat{\gamma}(0) \xrightarrow{p} \gamma(k)/\gamma(0) = \rho(k)$.

16.2 Autoregressions

In time-series, the series $\{\dots, y_1, y_2, \dots, y_T, \dots\}$ are jointly random. We consider the conditional expectation

$$\mathbb{E}(y_t | \mathcal{F}_{t-1})$$

where $\mathcal{F}_{t-1} = \{y_{t-1}, y_{t-2}, \dots\}$ is the past history of the series.

An autoregressive (AR) model specifies that only a finite number of past lags matter:

$$\mathbb{E}(y_t | \mathcal{F}_{t-1}) = \mathbb{E}(y_t | y_{t-1}, \dots, y_{t-k}).$$

A linear AR model (the most common type used in practice) specifies linearity:

$$\mathbb{E}(y_t | \mathcal{F}_{t-1}) = \alpha_0 + \alpha_1 y_{t-1} + \alpha_2 y_{t-2} + \dots + \alpha_k y_{t-k}.$$

Letting

$$e_t = y_t - \mathbb{E}(y_t | \mathcal{F}_{t-1}),$$

then we have the autoregressive model

$$y_t = \alpha_0 + \alpha_1 y_{t-1} + \alpha_2 y_{t-2} + \dots + \alpha_k y_{t-k} + e_t$$

$$\mathbb{E}(e_t | \mathcal{F}_{t-1}) = 0.$$

The last property defines a special time-series process.

Definition 16.2.1 e_t is a *martingale difference sequence (MDS)* if $\mathbb{E}(e_t | \mathcal{F}_{t-1}) = 0$.

Regression errors are naturally a MDS. Some time-series processes may be a MDS as a consequence of optimizing behavior. For example, some versions of the life-cycle hypothesis imply that either changes in consumption, or consumption growth rates, should be a MDS. Most asset pricing models imply that asset returns should be the sum of a constant plus a MDS.

The MDS property for the regression error plays the same role in a time-series regression as does the conditional mean-zero property for the regression error in a cross-section regression. In fact, it is even more important in the time-series context, as it is difficult to derive distribution theories without this property.

A useful property of a MDS is that e_t is uncorrelated with any function of the lagged information \mathcal{F}_{t-1} . Thus for $k > 0$, $\mathbb{E}(y_{t-k} e_t) = 0$.

16.3 Stationarity of AR(1) Process

A mean-zero AR(1) is

$$y_t = \alpha y_{t-1} + e_t.$$

Assume that e_t is iid, $\mathbb{E}(e_t) = 0$ and $\mathbb{E}e_t^2 = \sigma^2 < \infty$.

By back-substitution, we find

$$\begin{aligned} y_t &= e_t + \alpha e_{t-1} + \alpha^2 e_{t-2} + \dots \\ &= \sum_{k=0}^{\infty} \alpha^k e_{t-k}. \end{aligned}$$

Loosely speaking, this series converges if the sequence $\alpha^k e_{t-k}$ gets small as $k \rightarrow \infty$. This occurs when $|\alpha| < 1$.

Theorem 16.3.1 *If and only if $|\alpha| < 1$ then y_t is strictly stationary and ergodic.*

We can compute the moments of y_t using the infinite sum:

$$\begin{aligned} \mathbb{E}y_t &= \sum_{k=0}^{\infty} \alpha^k \mathbb{E}(e_{t-k}) = 0 \\ \text{var}(y_t) &= \sum_{k=0}^{\infty} \alpha^{2k} \text{var}(e_{t-k}) = \frac{\sigma^2}{1 - \alpha^2}. \end{aligned}$$

If the equation for y_t has an intercept, the above results are unchanged, except that the mean of y_t can be computed from the relationship

$$\mathbb{E}y_t = \alpha_0 + \alpha_1 \mathbb{E}y_{t-1},$$

and solving for $\mathbb{E}y_t = \mathbb{E}y_{t-1}$ we find $\mathbb{E}y_t = \alpha_0 / (1 - \alpha_1)$.

16.4 Lag Operator

An algebraic construct which is useful for the analysis of autoregressive models is the lag operator.

Definition 16.4.1 *The lag operator L satisfies $Ly_t = y_{t-1}$.*

Defining $L^2 = LL$, we see that $L^2 y_t = Ly_{t-1} = y_{t-2}$. In general, $L^k y_t = y_{t-k}$. The AR(1) model can be written in the format

$$y_t - \alpha y_{t-1} = e_t$$

or

$$(1 - \alpha L) y_t = e_t.$$

The operator $\alpha(L) = (1 - \alpha L)$ is a polynomial in the operator L . We say that the *root* of the polynomial is $1/\alpha$, since $\rho(z) = 0$ when $z = 1/\alpha$. We call $\alpha(L)$ the autoregressive polynomial of y_t .

From Theorem 16.3.1, an AR(1) is stationary iff $|\alpha| < 1$. Note that an equivalent way to say this is that an AR(1) is stationary iff the root of the autoregressive polynomial is larger than one (in absolute value).

16.5 Stationarity of AR(k)

The AR(k) model is

$$y_t = \alpha_1 y_{t-1} + \alpha_2 y_{t-2} + \cdots + \alpha_k y_{t-k} + e_t.$$

Using the lag operator,

$$y_t - \alpha_1 L y_t - \alpha_2 L^2 y_t - \cdots - \alpha_k L^k y_t = e_t,$$

or

$$\alpha(L) y_t = e_t$$

where

$$\rho(L) = 1 - \alpha_1 L - \alpha_2 L^2 - \cdots - \alpha_k L^k.$$

We call $\alpha(L)$ the autoregressive polynomial of y_t .

The *Fundamental Theorem of Algebra* says that any polynomial can be factored as

$$\alpha(z) = (1 - \lambda_1^{-1} z) (1 - \lambda_2^{-1} z) \cdots (1 - \lambda_k^{-1} z)$$

where the $\lambda_1, \dots, \lambda_k$ are the complex *roots* of $\alpha(z)$, which satisfy $\alpha(\lambda_j) = 0$.

We know that an AR(1) is stationary iff the absolute value of the root of its autoregressive polynomial is larger than one. For an AR(k), the requirement is that all roots are larger than one. Let $|\lambda|$ denote the modulus of a complex number λ .

Theorem 16.5.1 *The AR(k) is strictly stationary and ergodic if and only if $|\lambda_j| > 1$ for all j .*

One way of stating this is that “All roots lie outside the unit circle.”

If one of the roots equals 1, we say that $\alpha(L)$, and hence y_t , “has a unit root”. This is a special case of non-stationarity, and is of great interest in applied time series.

16.6 Estimation

Let

$$\begin{aligned} \mathbf{x}_t &= (1 \quad y_{t-1} \quad y_{t-2} \quad \cdots \quad y_{t-k})' \\ \boldsymbol{\beta} &= (\alpha_0 \quad \alpha_1 \quad \alpha_2 \quad \cdots \quad \alpha_k)'. \end{aligned}$$

Then the model can be written as

$$y_t = \mathbf{x}_t' \boldsymbol{\beta} + e_t.$$

The OLS estimator is

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}' \mathbf{X})^{-1} \mathbf{X}' \mathbf{y}.$$

To study $\hat{\boldsymbol{\beta}}$, it is helpful to define the process $u_t = \mathbf{x}_t e_t$. Note that u_t is a MDS, since

$$\mathbb{E}(u_t | \mathcal{F}_{t-1}) = \mathbb{E}(\mathbf{x}_t e_t | \mathcal{F}_{t-1}) = \mathbf{x}_t \mathbb{E}(e_t | \mathcal{F}_{t-1}) = 0.$$

By Theorem 16.1.1, it is also strictly stationary and ergodic. Thus

$$\frac{1}{T} \sum_{t=1}^T \mathbf{x}_t e_t = \frac{1}{T} \sum_{t=1}^T u_t \xrightarrow{p} \mathbb{E}(u_t) = 0. \quad (16.1)$$

The vector \mathbf{x}_t is strictly stationary and ergodic, and by Theorem 16.1.1, so is $\mathbf{x}_t \mathbf{x}'_t$. Thus by the Ergodic Theorem,

$$\frac{1}{T} \sum_{t=1}^T \mathbf{x}_t \mathbf{x}'_t \xrightarrow{p} \mathbb{E}(\mathbf{x}_t \mathbf{x}'_t) = \mathbf{Q}.$$

Combined with (16.1) and the continuous mapping theorem, we see that

$$\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta} = \left(\frac{1}{T} \sum_{t=1}^T \mathbf{x}_t \mathbf{x}'_t \right)^{-1} \left(\frac{1}{T} \sum_{t=1}^T \mathbf{x}_t e_t \right) \xrightarrow{p} \mathbf{Q}^{-1} \mathbf{0} = \mathbf{0}.$$

We have shown the following:

Theorem 16.6.1 *If the AR(k) process y_t is strictly stationary and ergodic and $\mathbb{E}y_t^2 < \infty$, then $\widehat{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\beta}$ as $T \rightarrow \infty$.*

16.7 Asymptotic Distribution

Theorem 16.7.1 *MDS CLT. If \mathbf{u}_t is a strictly stationary and ergodic MDS and $\mathbb{E}(\mathbf{u}_t \mathbf{u}'_t) = \boldsymbol{\Omega} < \infty$, then as $T \rightarrow \infty$,*

$$\frac{1}{\sqrt{T}} \sum_{t=1}^T \mathbf{u}_t \xrightarrow{d} \mathbf{N}(\mathbf{0}, \boldsymbol{\Omega}).$$

Since $\mathbf{x}_t e_t$ is a MDS, we can apply Theorem 16.7.1 to see that

$$\frac{1}{\sqrt{T}} \sum_{t=1}^T \mathbf{x}_t e_t \xrightarrow{d} \mathbf{N}(\mathbf{0}, \boldsymbol{\Omega}),$$

where

$$\boldsymbol{\Omega} = \mathbb{E}(\mathbf{x}_t \mathbf{x}'_t e_t^2).$$

Theorem 16.7.2 *If the AR(k) process y_t is strictly stationary and ergodic and $\mathbb{E}y_t^4 < \infty$, then as $T \rightarrow \infty$,*

$$\sqrt{T} (\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}) \xrightarrow{d} \mathbf{N}(\mathbf{0}, \mathbf{Q}^{-1} \boldsymbol{\Omega} \mathbf{Q}^{-1}).$$

This is identical in form to the asymptotic distribution of OLS in cross-section regression. The implication is that asymptotic inference is the same. In particular, the asymptotic covariance matrix is estimated just as in the cross-section case.

16.8 Bootstrap for Autoregressions

In the non-parametric bootstrap, we constructed the bootstrap sample by randomly resampling from the data values $\{y_t, \mathbf{x}_t\}$. This creates an iid bootstrap sample. Clearly, this cannot work in a time-series application, as this imposes inappropriate independence.

Briefly, there are two popular methods to implement bootstrap resampling for time-series data.

Method 1: Model-Based (Parametric) Bootstrap.

1. Estimate $\hat{\beta}$ and residuals \hat{e}_t .
2. Fix an initial condition $(y_{-k+1}, y_{-k+2}, \dots, y_0)$.
3. Simulate iid draws e_i^* from the empirical distribution of the residuals $\{\hat{e}_1, \dots, \hat{e}_T\}$.
4. Create the bootstrap series y_t^* by the recursive formula

$$y_t^* = \hat{\alpha}_0 + \hat{\alpha}_1 y_{t-1}^* + \hat{\alpha}_2 y_{t-2}^* + \dots + \hat{\alpha}_k y_{t-k}^* + e_t^*.$$

This construction imposes homoskedasticity on the errors e_i^* , which may be different than the properties of the actual e_i . It also presumes that the AR(k) structure is the truth.

Method 2: Block Resampling

1. Divide the sample into T/m blocks of length m .
2. Resample complete blocks. For each simulated sample, draw T/m blocks.
3. Paste the blocks together to create the bootstrap time-series y_t^* .
4. This allows for arbitrary stationary serial correlation, heteroskedasticity, and for model-misspecification.
5. The results may be sensitive to the block length, and the way that the data are partitioned into blocks.
6. May not work well in small samples.

16.9 Trend Stationarity

$$y_t = \mu_0 + \mu_1 t + S_t \tag{16.2}$$

$$S_t = \rho_1 S_{t-1} + \rho_2 S_{t-2} + \dots + \rho_k S_{t-k} + e_t, \tag{16.3}$$

or

$$y_t = \alpha_0 + \alpha_1 t + \rho_1 y_{t-1} + \rho_2 y_{t-2} + \dots + \rho_k y_{t-k} + e_t. \tag{16.4}$$

There are two essentially equivalent ways to estimate the autoregressive parameters $(\alpha_1, \dots, \alpha_k)$.

- You can estimate (16.4) by OLS.
- You can estimate (16.2)-(16.3) sequentially by OLS. That is, first estimate (16.2), get the residual \hat{S}_t , and then perform regression (16.3) replacing S_t with \hat{S}_t . This procedure is sometimes called *Detrending*.

The reason why these two procedures are (essentially) the same is the Frisch-Waugh-Lovell theorem.

Seasonal Effects

There are three popular methods to deal with seasonal data.

- Include dummy variables for each season. This presumes that “seasonality” does not change over the sample.
- Use “seasonally adjusted” data. The seasonal factor is typically estimated by a two-sided weighted average of the data for that season in neighboring years. Thus the seasonally adjusted data is a “filtered” series. This is a flexible approach which can extract a wide range of seasonal factors. The seasonal adjustment, however, also alters the time-series correlations of the data.
- First apply a seasonal differencing operator. If s is the number of seasons (typically $s = 4$ or $s = 12$),

$$\Delta_s y_t = y_t - y_{t-s},$$

or the season-to-season change. The series $\Delta_s y_t$ is clearly free of seasonality. But the long-run trend is also eliminated, and perhaps this was of relevance.

16.10 Testing for Omitted Serial Correlation

For simplicity, let the null hypothesis be an AR(1):

$$y_t = \alpha_0 + \alpha_1 y_{t-1} + u_t. \quad (16.5)$$

We are interested in the question if the error u_t is serially correlated. We model this as an AR(1):

$$u_t = \theta u_{t-1} + e_t \quad (16.6)$$

with e_t a MDS. The hypothesis of no omitted serial correlation is

$$\begin{aligned} \mathbb{H}_0 : \theta &= 0 \\ \mathbb{H}_1 : \theta &\neq 0. \end{aligned}$$

We want to test \mathbb{H}_0 against \mathbb{H}_1 .

To combine (16.5) and (16.6), we take (16.5) and lag the equation once:

$$y_{t-1} = \alpha_0 + \alpha_1 y_{t-2} + u_{t-1}.$$

We then multiply this by θ and subtract from (16.5), to find

$$y_t - \theta y_{t-1} = \alpha_0 - \theta \alpha_0 + \alpha_1 y_{t-1} - \theta \alpha_1 y_{t-1} + u_t - \theta u_{t-1},$$

or

$$y_t = \alpha_0(1 - \theta) + (\alpha_1 + \theta) y_{t-1} - \theta \alpha_1 y_{t-2} + e_t = AR(2).$$

Thus under \mathbb{H}_0 , y_t is an AR(1), and under \mathbb{H}_1 it is an AR(2). \mathbb{H}_0 may be expressed as the restriction that the coefficient on y_{t-2} is zero.

An appropriate test of \mathbb{H}_0 against \mathbb{H}_1 is therefore a Wald test that the coefficient on y_{t-2} is zero. (A simple exclusion test).

In general, if the null hypothesis is that y_t is an AR(k), and the alternative is that the error is an AR(m), this is the same as saying that under the alternative y_t is an AR($k+m$), and this is equivalent to the restriction that the coefficients on $y_{t-k-1}, \dots, y_{t-k-m}$ are jointly zero. An appropriate test is the Wald test of this restriction.

16.11 Model Selection

What is the appropriate choice of k in practice? This is a problem of model selection.

A good choice is to minimize the AIC information criterion

$$AIC(k) = \log \hat{\sigma}^2(k) + \frac{2k}{T},$$

where $\hat{\sigma}^2(k)$ is the estimated residual variance from an AR(k)

One ambiguity in defining the AIC criterion is that the sample available for estimation changes as k changes. (If you increase k , you need more initial conditions.) This can induce strange behavior in the AIC. The appropriate remedy is to fix an upper value \bar{k} , and then reserve the first \bar{k} as initial conditions, and then estimate the models AR(1), AR(2), ..., AR(\bar{k}) on this (unified) sample.

16.12 Autoregressive Unit Roots

The AR(k) model is

$$\begin{aligned}\alpha(L)y_t &= \alpha_0 + e_t \\ \alpha(L) &= 1 - \alpha_1 L - \dots - \alpha_k L^k.\end{aligned}$$

As we discussed before, y_t has a unit root when $\alpha(1) = 0$, or

$$\alpha_1 + \alpha_2 + \dots + \alpha_k = 1.$$

In this case, y_t is non-stationary. The ergodic theorem and MDS CLT do not apply, and test statistics are asymptotically non-normal.

A helpful way to write the equation is the so-called Dickey-Fuller reparameterization:

$$\Delta y_t = \rho_0 y_{t-1} + \rho_1 \Delta y_{t-1} + \dots + \rho_{k-1} \Delta y_{t-(k-1)} + e_t. \quad (16.7)$$

These models are equivalent linear transformations of one another. The DF parameterization is convenient because the parameter ρ_0 summarizes the information about the unit root, since $\alpha(1) = -\rho_0$. To see this, observe that the lag polynomial for the y_t computed from (16.7) is

$$(1 - L) - \rho_0 L - \rho_1 (L - L^2) - \dots - \rho_{k-1} (L^{k-1} - L^k)$$

But this must equal $\rho(L)$, as the models are equivalent. Thus

$$\alpha(1) = (1 - 1) - \rho_0 - (1 - 1) - \dots - (1 - 1) = -\rho_0.$$

Hence, the hypothesis of a unit root in y_t can be stated as

$$\mathbb{H}_0 : \rho_0 = 0.$$

Note that the model is stationary if $\rho_0 < 0$. So the natural alternative is

$$\mathbb{H}_1 : \rho_0 < 0.$$

Under \mathbb{H}_0 , the model for y_t is

$$\Delta y_t = \mu + \rho_1 \Delta y_{t-1} + \dots + \rho_{k-1} \Delta y_{t-(k-1)} + e_t,$$

which is an AR($k-1$) in the first-difference Δy_t . Thus if y_t has a (single) unit root, then Δy_t is a stationary AR process. Because of this property, we say that if y_t is non-stationary but $\Delta^d y_t$ is stationary, then y_t is “integrated of order d ”, or $I(d)$. Thus a time series with unit root is $I(1)$.

Since α_0 is the parameter of a linear regression, the natural test statistic is the t-statistic for \mathbb{H}_0 from OLS estimation of (16.7). Indeed, this is the most popular unit root test, and is called the Augmented Dickey-Fuller (ADF) test for a unit root.

It would seem natural to assess the significance of the ADF statistic using the normal table. However, under \mathbb{H}_0 , y_t is non-stationary, so conventional normal asymptotics are invalid. An alternative asymptotic framework has been developed to deal with non-stationary data. We do not have the time to develop this theory in detail, but simply assert the main results.

Theorem 16.12.1 Dickey-Fuller Theorem.

If $\rho_0 = 0$ then as $T \rightarrow \infty$,

$$T\hat{\rho}_0 \xrightarrow{d} (1 - \rho_1 - \rho_2 - \cdots - \rho_{k-1}) DF_\alpha$$

$$ADF = \frac{\hat{\rho}_0}{s(\hat{\rho}_0)} \rightarrow DF_t.$$

The limit distributions DF_α and DF_t are non-normal. They are skewed to the left, and have negative means.

The first result states that $\hat{\rho}_0$ converges to its true value (of zero) at rate T , rather than the conventional rate of $T^{1/2}$. This is called a “super-consistent” rate of convergence.

The second result states that the t-statistic for $\hat{\rho}_0$ converges to a limit distribution which is non-normal, but does not depend on the parameters ρ . This distribution has been extensively tabulated, and may be used for testing the hypothesis \mathbb{H}_0 . Note: The standard error $s(\hat{\rho}_0)$ is the conventional (“homoskedastic”) standard error. But the theorem does not require an assumption of homoskedasticity. Thus the Dickey-Fuller test is robust to heteroskedasticity.

Since the alternative hypothesis is one-sided, the ADF test rejects \mathbb{H}_0 in favor of \mathbb{H}_1 when $ADF < c$, where c is the critical value from the ADF table. If the test rejects \mathbb{H}_0 , this means that the evidence points to y_t being stationary. If the test does not reject \mathbb{H}_0 , a common conclusion is that the data suggests that y_t is non-stationary. This is not really a correct conclusion, however. All we can say is that there is insufficient evidence to conclude whether the data are stationary or not.

We have described the test for the setting of with an intercept. Another popular setting includes as well a linear time trend. This model is

$$\Delta y_t = \mu_1 + \mu_2 t + \rho_0 y_{t-1} + \rho_1 \Delta y_{t-1} + \cdots + \rho_{k-1} \Delta y_{t-(k-1)} + e_t. \quad (16.8)$$

This is natural when the alternative hypothesis is that the series is stationary about a linear time trend. If the series has a linear trend (e.g. GDP, Stock Prices), then the series itself is non-stationary, but it may be stationary around the linear time trend. In this context, it is a silly waste of time to fit an AR model to the level of the series without a time trend, as the AR model cannot conceivably describe this data. The natural solution is to include a time trend in the fitted OLS equation. When conducting the ADF test, this means that it is computed as the t-ratio for ρ_0 from OLS estimation of (16.8).

If a time trend is included, the test procedure is the same, but different critical values are required. The ADF test has a different distribution when the time trend has been included, and a different table should be consulted.

Most texts include as well the critical values for the extreme polar case where the intercept has been omitted from the model. These are included for completeness (from a pedagogical perspective) but have no relevance for empirical practice where intercepts are always included.

Chapter 17

Multivariate Time Series

A multivariate time series \mathbf{y}_t is a vector process $m \times 1$. Let $\mathcal{F}_{t-1} = (\mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \dots)$ be all lagged information at time t . The typical goal is to find the conditional expectation $\mathbb{E}(\mathbf{y}_t | \mathcal{F}_{t-1})$. Note that since \mathbf{y}_t is a vector, this conditional expectation is also a vector.

17.1 Vector Autoregressions (VARs)

A VAR model specifies that the conditional mean is a function of only a finite number of lags:

$$\mathbb{E}(\mathbf{y}_t | \mathcal{F}_{t-1}) = \mathbb{E}(\mathbf{y}_t | \mathbf{y}_{t-1}, \dots, \mathbf{y}_{t-k}).$$

A linear VAR specifies that this conditional mean is linear in the arguments:

$$\mathbb{E}(\mathbf{y}_t | \mathbf{y}_{t-1}, \dots, \mathbf{y}_{t-k}) = \mathbf{a}_0 + \mathbf{A}_1 \mathbf{y}_{t-1} + \mathbf{A}_2 \mathbf{y}_{t-2} + \dots + \mathbf{A}_k \mathbf{y}_{t-k}.$$

Observe that \mathbf{a}_0 is $m \times 1$, and each of \mathbf{A}_1 through \mathbf{A}_k are $m \times m$ matrices.

Defining the $m \times 1$ regression error

$$\mathbf{e}_t = \mathbf{y}_t - \mathbb{E}(\mathbf{y}_t | \mathcal{F}_{t-1}),$$

we have the VAR model

$$\begin{aligned} \mathbf{y}_t &= \mathbf{a}_0 + \mathbf{A}_1 \mathbf{y}_{t-1} + \mathbf{A}_2 \mathbf{y}_{t-2} + \dots + \mathbf{A}_k \mathbf{y}_{t-k} + \mathbf{e}_t \\ \mathbb{E}(\mathbf{e}_t | \mathcal{F}_{t-1}) &= \mathbf{0}. \end{aligned}$$

Alternatively, defining the $mk + 1$ vector

$$\mathbf{x}_t = \begin{pmatrix} 1 \\ \mathbf{y}_{t-1} \\ \mathbf{y}_{t-2} \\ \vdots \\ \mathbf{y}_{t-k} \end{pmatrix}$$

and the $m \times (mk + 1)$ matrix

$$\mathbf{A} = \left(\mathbf{a}_0 \quad \mathbf{A}_1 \quad \mathbf{A}_2 \quad \dots \quad \mathbf{A}_k \right),$$

then

$$\mathbf{y}_t = \mathbf{A} \mathbf{x}_t + \mathbf{e}_t.$$

The VAR model is a system of m equations. One way to write this is to let a'_j be the j th row of \mathbf{A} . Then the VAR system can be written as the equations

$$Y_{jt} = a'_j \mathbf{x}_t + e_{jt}.$$

Unrestricted VARs were introduced to econometrics by Sims (1980).

17.2 Estimation

Consider the moment conditions

$$\mathbb{E}(\mathbf{x}_t e_{jt}) = \mathbf{0},$$

$j = 1, \dots, m$. These are implied by the VAR model, either as a regression, or as a linear projection.

The GMM estimator corresponding to these moment conditions is equation-by-equation OLS

$$\hat{\mathbf{a}}_j = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}_j.$$

An alternative way to compute this is as follows. Note that

$$\hat{\mathbf{a}}'_j = \mathbf{y}'_j \mathbf{X} (\mathbf{X}'\mathbf{X})^{-1}.$$

And if we stack these to create the estimate $\hat{\mathbf{A}}$, we find

$$\begin{aligned} \hat{\mathbf{A}} &= \begin{pmatrix} \mathbf{y}'_1 \\ \mathbf{y}'_2 \\ \vdots \\ \mathbf{y}'_{m+1} \end{pmatrix} \mathbf{X} (\mathbf{X}'\mathbf{X})^{-1} \\ &= \mathbf{Y}'\mathbf{X} (\mathbf{X}'\mathbf{X})^{-1}, \end{aligned}$$

where

$$\mathbf{Y} = \begin{pmatrix} \mathbf{y}_1 & \mathbf{y}_2 & \cdots & \mathbf{y}_m \end{pmatrix}$$

the $T \times m$ matrix of the stacked \mathbf{y}'_t .

This (system) estimator is known as the SUR (Seemingly Unrelated Regressions) estimator, and was originally derived by Zellner (1962)

17.3 Restricted VARs

The unrestricted VAR is a system of m equations, each with the same set of regressors. A restricted VAR imposes restrictions on the system. For example, some regressors may be excluded from some of the equations. Restrictions may be imposed on individual equations, or across equations. The GMM framework gives a convenient method to impose such restrictions on estimation.

17.4 Single Equation from a VAR

Often, we are only interested in a single equation out of a VAR system. This takes the form

$$y_{jt} = \mathbf{a}'_j \mathbf{x}_t + e_t,$$

and \mathbf{x}_t consists of lagged values of y_{jt} and the other y'_{lt} s. In this case, it is convenient to re-define the variables. Let $y_t = y_{jt}$, and \mathbf{z}_t be the other variables. Let $e_t = e_{jt}$ and $\beta = a_j$. Then the single equation takes the form

$$y_t = \mathbf{x}'_t \beta + e_t, \tag{17.1}$$

and

$$\mathbf{x}_t = \begin{bmatrix} 1 & \mathbf{y}_{t-1} & \cdots & \mathbf{y}_{t-k} & \mathbf{z}'_{t-1} & \cdots & \mathbf{z}'_{t-k} \end{bmatrix}'.$$

This is just a conventional regression with time series data.

17.5 Testing for Omitted Serial Correlation

Consider the problem of testing for omitted serial correlation in equation (17.1). Suppose that e_t is an AR(1). Then

$$\begin{aligned} y_t &= \mathbf{x}'_t \boldsymbol{\beta} + e_t \\ e_t &= \theta e_{t-1} + u_t \\ \mathbb{E}(u_t \mid \mathcal{F}_{t-1}) &= 0. \end{aligned} \tag{17.2}$$

Then the null and alternative are

$$\mathbb{H}_0 : \theta = 0 \quad \mathbb{H}_1 : \theta \neq 0.$$

Take the equation $y_t = \mathbf{x}'_t \boldsymbol{\beta} + e_t$, and subtract off the equation once lagged multiplied by θ , to get

$$\begin{aligned} y_t - \theta y_{t-1} &= (\mathbf{x}'_t \boldsymbol{\beta} + e_t) - \theta (\mathbf{x}'_{t-1} \boldsymbol{\beta} + e_{t-1}) \\ &= \mathbf{x}'_t \boldsymbol{\beta} - \theta \mathbf{x}'_{t-1} \boldsymbol{\beta} + e_t - \theta e_{t-1}, \end{aligned}$$

or

$$y_t = \theta y_{t-1} + \mathbf{x}'_t \boldsymbol{\beta} + \mathbf{x}'_{t-1} \boldsymbol{\gamma} + u_t, \tag{17.3}$$

which is a valid regression model.

So testing \mathbb{H}_0 versus \mathbb{H}_1 is equivalent to testing for the significance of adding $(y_{t-1}, \mathbf{x}_{t-1})$ to the regression. This can be done by a Wald test. We see that an appropriate, general, and simple way to test for omitted serial correlation is to test the significance of extra lagged values of the dependent variable and regressors.

You may have heard of the Durbin-Watson test for omitted serial correlation, which once was very popular, and is still routinely reported by conventional regression packages. The DW test is appropriate only when regression $y_t = \mathbf{x}'_t \boldsymbol{\beta} + e_t$ is not dynamic (has no lagged values on the RHS), and e_t is iid $N(0, \sigma^2)$. Otherwise it is invalid.

Another interesting fact is that (17.2) is a special case of (17.3), under the restriction $\boldsymbol{\gamma} = -\boldsymbol{\beta}\theta$. This restriction, which is called a common factor restriction, may be tested if desired. If valid, the model (17.2) may be estimated by iterated GLS. (A simple version of this estimator is called Cochrane-Orcutt.) Since the common factor restriction appears arbitrary, and is typically rejected empirically, direct estimation of (17.2) is uncommon in recent applications.

17.6 Selection of Lag Length in an VAR

If you want a data-dependent rule to pick the lag length k in a VAR, you may either use a testing-based approach (using, for example, the Wald statistic), or an information criterion approach. The formula for the AIC and BIC are

$$\begin{aligned} AIC(k) &= \log \det \left(\hat{\boldsymbol{\Omega}}(k) \right) + 2 \frac{p}{T} \\ BIC(k) &= \log \det \left(\hat{\boldsymbol{\Omega}}(k) \right) + \frac{p \log(T)}{T} \\ \hat{\boldsymbol{\Omega}}(k) &= \frac{1}{T} \sum_{t=1}^T \hat{\mathbf{e}}_t(k) \hat{\mathbf{e}}_t(k)' \\ p &= m(km + 1) \end{aligned}$$

where p is the number of parameters in the model, and $\hat{\mathbf{e}}_t(k)$ is the OLS residual vector from the model with k lags. The log determinant is the criterion from the multivariate normal likelihood.

17.7 Granger Causality

Partition the data vector into $(\mathbf{y}_t, \mathbf{z}_t)$. Define the two information sets

$$\begin{aligned}\mathcal{F}_{1t} &= (\mathbf{y}_t, \mathbf{y}_{t-1}, \mathbf{y}_{t-2}, \dots) \\ \mathcal{F}_{2t} &= (\mathbf{y}_t, \mathbf{z}_t, \mathbf{y}_{t-1}, \mathbf{z}_{t-1}, \mathbf{y}_{t-2}, \mathbf{z}_{t-2}, \dots)\end{aligned}$$

The information set \mathcal{F}_{1t} is generated only by the history of \mathbf{y}_t , and the information set \mathcal{F}_{2t} is generated by both \mathbf{y}_t and \mathbf{z}_t . The latter has more information.

We say that \mathbf{z}_t does not *Granger-cause* \mathbf{y}_t if

$$\mathbb{E}(\mathbf{y}_t \mid \mathcal{F}_{1,t-1}) = \mathbb{E}(\mathbf{y}_t \mid \mathcal{F}_{2,t-1}).$$

That is, conditional on information in lagged \mathbf{y}_t , lagged \mathbf{z}_t does not help to forecast \mathbf{y}_t . If this condition does not hold, then we say that \mathbf{z}_t Granger-causes \mathbf{y}_t .

The reason why we call this “Granger Causality” rather than “causality” is because this is not a physical or structure definition of causality. If \mathbf{z}_t is some sort of forecast of the future, such as a futures price, then \mathbf{z}_t may help to forecast \mathbf{y}_t even though it does not “cause” \mathbf{y}_t . This definition of causality was developed by Granger (1969) and Sims (1972).

In a linear VAR, the equation for \mathbf{y}_t is

$$\mathbf{y}_t = \alpha + \rho_1 \mathbf{y}_{t-1} + \dots + \rho_k \mathbf{y}_{t-k} + \mathbf{z}'_{t-1} \boldsymbol{\gamma}_1 + \dots + \mathbf{z}'_{t-k} \boldsymbol{\gamma}_k + e_t.$$

In this equation, \mathbf{z}_t does not Granger-cause \mathbf{y}_t if and only if

$$\mathbb{H}_0 : \boldsymbol{\gamma}_1 = \boldsymbol{\gamma}_2 = \dots = \boldsymbol{\gamma}_k = \mathbf{0}.$$

This may be tested using an exclusion (Wald) test.

This idea can be applied to blocks of variables. That is, \mathbf{y}_t and/or \mathbf{z}_t can be vectors. The hypothesis can be tested by using the appropriate multivariate Wald test.

If it is found that \mathbf{z}_t does not Granger-cause \mathbf{y}_t , then we deduce that our time-series model of $\mathbb{E}(\mathbf{y}_t \mid \mathcal{F}_{t-1})$ does not require the use of \mathbf{z}_t . Note, however, that \mathbf{z}_t may still be useful to explain other features of \mathbf{y}_t , such as the conditional variance.

Clive W. J. Granger

Clive Granger (1934-2009) of England was one of the leading figures in time-series econometrics, and co-winner in 2003 of the Nobel Memorial Prize in Economic Sciences (along with Robert Engle). In addition to formalizing the definition of causality known as Granger causality, he invented the concept of cointegration, introduced spectral methods into econometrics, and formalized methods for the combination of forecasts.

17.8 Cointegration

The idea of cointegration is due to Granger (1981), and was articulated in detail by Engle and Granger (1987).

Definition 17.8.1 The $m \times 1$ series \mathbf{y}_t is *cointegrated* if \mathbf{y}_t is $I(1)$ yet there exists $\boldsymbol{\beta}$, $m \times r$, of rank r , such that $\mathbf{z}_t = \boldsymbol{\beta}'\mathbf{y}_t$ is $I(0)$. The r vectors in $\boldsymbol{\beta}$ are called the *cointegrating vectors*.

If the series \mathbf{y}_t is not cointegrated, then $r = 0$. If $r = m$, then \mathbf{y}_t is $I(0)$. For $0 < r < m$, \mathbf{y}_t is $I(1)$ and cointegrated.

In some cases, it may be believed that $\boldsymbol{\beta}$ is known a priori. Often, $\boldsymbol{\beta} = (1 \quad -1)'$. For example, if \mathbf{y}_t is a pair of interest rates, then $\boldsymbol{\beta} = (1 \quad -1)'$ specifies that the spread (the difference in returns) is stationary. If $\mathbf{y} = (\log(C) \quad \log(I))'$, then $\boldsymbol{\beta} = (1 \quad -1)'$ specifies that $\log(C/I)$ is stationary.

In other cases, $\boldsymbol{\beta}$ may not be known.

If \mathbf{y}_t is cointegrated with a single cointegrating vector ($r = 1$), then it turns out that $\boldsymbol{\beta}$ can be consistently estimated by an OLS regression of one component of \mathbf{y}_t on the others. Thus $\mathbf{y}_t = (Y_{1t}, Y_{2t})$ and $\boldsymbol{\beta} = (\beta_1 \quad \beta_2)$ and normalize $\beta_1 = 1$. Then $\hat{\beta}_2 = (\mathbf{y}'_2 \mathbf{y}_2)^{-1} \mathbf{y}'_2 \mathbf{y}_1 \xrightarrow{p} \beta_2$. Furthermore this estimation is super-consistent: $T(\hat{\beta}_2 - \beta_2) \xrightarrow{d} \text{Limit}$, as first shown by Stock (1987). This is not, in general, a good method to estimate $\boldsymbol{\beta}$, but it is useful in the construction of alternative estimators and tests.

We are often interested in testing the hypothesis of no cointegration:

$$\mathbb{H}_0 : r = 0$$

$$\mathbb{H}_1 : r > 0.$$

Suppose that $\boldsymbol{\beta}$ is known, so $\mathbf{z}_t = \boldsymbol{\beta}'\mathbf{y}_t$ is known. Then under \mathbb{H}_0 \mathbf{z}_t is $I(1)$, yet under \mathbb{H}_1 \mathbf{z}_t is $I(0)$. Thus \mathbb{H}_0 can be tested using a univariate ADF test on \mathbf{z}_t .

When $\boldsymbol{\beta}$ is unknown, Engle and Granger (1987) suggested using an ADF test on the estimated residual $\hat{z}_t = \hat{\boldsymbol{\beta}}'\mathbf{y}_t$, from OLS of y_{1t} on y_{2t} . Their justification was Stock's result that $\hat{\boldsymbol{\beta}}$ is super-consistent under \mathbb{H}_1 . Under \mathbb{H}_0 , however, $\hat{\boldsymbol{\beta}}$ is not consistent, so the ADF critical values are not appropriate. The asymptotic distribution was worked out by Phillips and Ouliaris (1990).

When the data have time trends, it may be necessary to include a time trend in the estimated cointegrating regression. Whether or not the time trend is included, the asymptotic distribution of the test is affected by the presence of the time trend. The asymptotic distribution was worked out in B. Hansen (1992).

17.9 Cointegrated VARs

We can write a VAR as

$$\begin{aligned} \mathbf{A}(L)\mathbf{y}_t &= \mathbf{e}_t \\ \mathbf{A}(L) &= \mathbf{I} - \mathbf{A}_1 L - \mathbf{A}_2 L^2 - \dots - \mathbf{A}_k L^k \end{aligned}$$

or alternatively as

$$\Delta \mathbf{y}_t = \boldsymbol{\Pi} \mathbf{y}_{t-1} + \mathbf{D}(L) \Delta \mathbf{y}_{t-1} + \mathbf{e}_t$$

where

$$\begin{aligned} \boldsymbol{\Pi} &= -\mathbf{A}(1) \\ &= -\mathbf{I} + \mathbf{A}_1 + \mathbf{A}_2 + \dots + \mathbf{A}_k. \end{aligned}$$

Theorem 17.9.1 Granger Representation Theorem

\mathbf{y}_t is cointegrated with $m \times r$ $\boldsymbol{\beta}$ if and only if $\text{rank}(\boldsymbol{\Pi}) = r$ and $\boldsymbol{\Pi} = \boldsymbol{\alpha}\boldsymbol{\beta}'$ where $\boldsymbol{\alpha}$ is $m \times r$, $\text{rank}(\boldsymbol{\alpha}) = r$.

Thus cointegration imposes a restriction upon the parameters of a VAR. The restricted model can be written as

$$\begin{aligned}\Delta \mathbf{y}_t &= \boldsymbol{\alpha}\boldsymbol{\beta}'\mathbf{y}_{t-1} + \mathbf{D}(L)\Delta \mathbf{y}_{t-1} + \mathbf{e}_t \\ \Delta \mathbf{y}_t &= \boldsymbol{\alpha}\mathbf{z}_{t-1} + \mathbf{D}(L)\Delta \mathbf{y}_{t-1} + \mathbf{e}_t.\end{aligned}$$

If $\boldsymbol{\beta}$ is known, this can be estimated by OLS of $\Delta \mathbf{y}_t$ on \mathbf{z}_{t-1} and the lags of $\Delta \mathbf{y}_t$.

If $\boldsymbol{\beta}$ is unknown, then estimation is done by “reduced rank regression”, which is least-squares subject to the stated restriction. Equivalently, this is the MLE of the restricted parameters under the assumption that \mathbf{e}_t is iid $N(\mathbf{0}, \boldsymbol{\Omega})$.

One difficulty is that $\boldsymbol{\beta}$ is not identified without normalization. When $r = 1$, we typically just normalize one element to equal unity. When $r > 1$, this does not work, and different authors have adopted different identification schemes.

In the context of a cointegrated VAR estimated by reduced rank regression, it is simple to test for cointegration by testing the rank of $\boldsymbol{\Pi}$. These tests are constructed as likelihood ratio (LR) tests. As they were discovered by Johansen (1988, 1991, 1995), they are typically called the “Johansen Max and Trace” tests. Their asymptotic distributions are non-standard, and are similar to the Dickey-Fuller distributions.

Chapter 18

Limited Dependent Variables

A “limited dependent variable” y is one which takes a “limited” set of values. The most common cases are

- Binary: $y \in \{0, 1\}$
- Multinomial: $y \in \{0, 1, 2, \dots, k\}$
- Integer: $y \in \{0, 1, 2, \dots\}$
- Censored: $y \in \mathbb{R}^+$

The traditional approach to the estimation of limited dependent variable (LDV) models is parametric maximum likelihood. A parametric model is constructed, allowing the construction of the likelihood function. A more modern approach is semi-parametric, eliminating the dependence on a parametric distributional assumption. We will discuss only the first (parametric) approach, due to time constraints. They still constitute the majority of LDV applications. If, however, you were to write a thesis involving LDV estimation, you would be advised to consider employing a semi-parametric estimation approach.

For the parametric approach, estimation is by MLE. A major practical issue is construction of the likelihood function.

18.1 Binary Choice

The dependent variable $y_i \in \{0, 1\}$. This represents a Yes/No outcome. Given some regressors \mathbf{x}_i , the goal is to describe $\Pr(y_i = 1 \mid \mathbf{x}_i)$, as this is the full conditional distribution.

The linear probability model specifies that

$$\Pr(y_i = 1 \mid \mathbf{x}_i) = \mathbf{x}_i' \boldsymbol{\beta}.$$

As $\Pr(y_i = 1 \mid \mathbf{x}_i) = \mathbb{E}(y_i \mid \mathbf{x}_i)$, this yields the regression: $y_i = \mathbf{x}_i' \boldsymbol{\beta} + e_i$ which can be estimated by OLS. However, the linear probability model does not impose the restriction that $0 \leq \Pr(y_i \mid \mathbf{x}_i) \leq 1$. Even so estimation of a linear probability model is a useful starting point for subsequent analysis.

The standard alternative is to use a function of the form

$$\Pr(y_i = 1 \mid \mathbf{x}_i) = F(\mathbf{x}_i' \boldsymbol{\beta})$$

where $F(\cdot)$ is a known CDF, typically assumed to be symmetric about zero, so that $F(u) = 1 - F(-u)$. The two standard choices for F are

- Logistic: $F(u) = (1 + e^{-u})^{-1}$.

- Normal: $F(u) = \Phi(u)$.

If F is logistic, we call this the *logit* model, and if F is normal, we call this the *probit* model. This model is identical to the latent variable model

$$\begin{aligned} y_i^* &= \mathbf{x}'_i \boldsymbol{\beta} + e_i \\ e_i &\sim F(\cdot) \\ y_i &= \begin{cases} 1 & \text{if } y_i^* > 0 \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

For then

$$\begin{aligned} \Pr(y_i = 1 \mid \mathbf{x}_i) &= \Pr(y_i^* > 0 \mid \mathbf{x}_i) \\ &= \Pr(\mathbf{x}'_i \boldsymbol{\beta} + e_i > 0 \mid \mathbf{x}_i) \\ &= \Pr(e_i > -\mathbf{x}'_i \boldsymbol{\beta} \mid \mathbf{x}_i) \\ &= 1 - F(-\mathbf{x}'_i \boldsymbol{\beta}) \\ &= F(\mathbf{x}'_i \boldsymbol{\beta}). \end{aligned}$$

Estimation is by maximum likelihood. To construct the likelihood, we need the conditional distribution of an individual observation. Recall that if y is Bernoulli, such that $\Pr(y = 1) = p$ and $\Pr(y = 0) = 1 - p$, then we can write the density of y as

$$f(y) = p^y (1 - p)^{1-y}, \quad y = 0, 1.$$

In the Binary choice model, y_i is conditionally Bernoulli with $\Pr(y_i = 1 \mid \mathbf{x}_i) = p_i = F(\mathbf{x}'_i \boldsymbol{\beta})$. Thus the conditional density is

$$\begin{aligned} f(y_i \mid \mathbf{x}_i) &= p_i^{y_i} (1 - p_i)^{1-y_i} \\ &= F(\mathbf{x}'_i \boldsymbol{\beta})^{y_i} (1 - F(\mathbf{x}'_i \boldsymbol{\beta}))^{1-y_i}. \end{aligned}$$

Hence the log-likelihood function is

$$\begin{aligned} \log L(\boldsymbol{\beta}) &= \sum_{i=1}^n \log f(y_i \mid \mathbf{x}_i) \\ &= \sum_{i=1}^n \log (F(\mathbf{x}'_i \boldsymbol{\beta})^{y_i} (1 - F(\mathbf{x}'_i \boldsymbol{\beta}))^{1-y_i}) \\ &= \sum_{i=1}^n [y_i \log F(\mathbf{x}'_i \boldsymbol{\beta}) + (1 - y_i) \log(1 - F(\mathbf{x}'_i \boldsymbol{\beta}))] \\ &= \sum_{y_i=1} \log F(\mathbf{x}'_i \boldsymbol{\beta}) + \sum_{y_i=0} \log(1 - F(\mathbf{x}'_i \boldsymbol{\beta})). \end{aligned}$$

The MLE $\hat{\boldsymbol{\beta}}$ is the value of $\boldsymbol{\beta}$ which maximizes $\log L(\boldsymbol{\beta})$. Standard errors and test statistics are computed by asymptotic approximations. Details of such calculations are left to more advanced courses.

18.2 Count Data

If $y \in \{0, 1, 2, \dots\}$, a typical approach is to employ *Poisson regression*. This model specifies that

$$\begin{aligned} \Pr(y_i = k \mid \mathbf{x}_i) &= \frac{\exp(-\lambda_i) \lambda_i^k}{k!}, \quad k = 0, 1, 2, \dots \\ \lambda_i &= \exp(\mathbf{x}'_i \boldsymbol{\beta}). \end{aligned}$$

The conditional density is the Poisson with parameter λ_i . The functional form for λ_i has been picked to ensure that $\lambda_i > 0$.

The log-likelihood function is

$$\log L(\boldsymbol{\beta}) = \sum_{i=1}^n \log f(y_i | \mathbf{x}_i) = \sum_{i=1}^n (-\exp(\mathbf{x}'_i \boldsymbol{\beta}) + y_i \mathbf{x}'_i \boldsymbol{\beta} - \log(y_i!)).$$

The MLE is the value $\hat{\boldsymbol{\beta}}$ which maximizes $\log L(\boldsymbol{\beta})$.

Since

$$\mathbb{E}(y_i | \mathbf{x}_i) = \lambda_i = \exp(\mathbf{x}'_i \boldsymbol{\beta})$$

is the conditional mean, this motivates the label Poisson “regression.”

Also observe that the model implies that

$$\text{var}(y_i | \mathbf{x}_i) = \lambda_i = \exp(\mathbf{x}'_i \boldsymbol{\beta}),$$

so the model imposes the restriction that the conditional mean and variance of y_i are the same. This may be considered restrictive. A generalization is the negative binomial.

18.3 Censored Data

The idea of “censoring” is that some data above or below a threshold are mis-reported at the threshold. Thus the model is that there is some latent process y_i^* with unbounded support, but we observe only

$$y_i = \begin{cases} y_i^* & \text{if } y_i^* \geq 0 \\ 0 & \text{if } y_i^* < 0 \end{cases} . \quad (18.1)$$

(This is written for the case of the threshold being zero, any known value can substitute.) The observed data y_i therefore come from a mixed continuous/discrete distribution.

Censored models are typically applied when the data set has a meaningful proportion (say 5% or higher) of data at the boundary of the sample support. The censoring process may be explicit in data collection, or it may be a by-product of economic constraints.

An example of a data collection censoring is top-coding of income. In surveys, incomes above a threshold are typically reported at the threshold.

The first censored regression model was developed by Tobin (1958) to explain consumption of durable goods. Tobin observed that for many households, the consumption level (purchases) in a particular period was zero. He proposed the latent variable model

$$y_i^* = \mathbf{x}'_i \boldsymbol{\beta} + e_i \\ e_i \stackrel{iid}{\sim} \text{N}(0, \sigma^2)$$

with the observed variable y_i generated by the censoring equation (18.1). This model (now called the Tobit) specifies that the latent (or ideal) value of consumption may be negative (the household would prefer to sell than buy). All that is reported is that the household purchased zero units of the good.

The naive approach to estimate $\boldsymbol{\beta}$ is to regress y_i on \mathbf{x}_i . This does not work because regression estimates $\mathbb{E}(y_i | \mathbf{x}_i)$, not $\mathbb{E}(y_i^* | \mathbf{x}_i) = \mathbf{x}'_i \boldsymbol{\beta}$, and the latter is of interest. Thus OLS will be biased for the parameter of interest $\boldsymbol{\beta}$.

[Note: it is still possible to estimate $\mathbb{E}(y_i | \mathbf{x}_i)$ by LS techniques. The Tobit framework postulates that this is not inherently interesting, that the parameter of $\boldsymbol{\beta}$ is defined by an alternative statistical structure.]

Consistent estimation will be achieved by the MLE. To construct the likelihood, observe that the probability of being censored is

$$\begin{aligned}\Pr(y_i = 0 \mid \mathbf{x}_i) &= \Pr(y_i^* < 0 \mid \mathbf{x}_i) \\ &= \Pr(\mathbf{x}_i' \boldsymbol{\beta} + e_i < 0 \mid \mathbf{x}_i) \\ &= \Pr\left(\frac{e_i}{\sigma} < -\frac{\mathbf{x}_i' \boldsymbol{\beta}}{\sigma} \mid \mathbf{x}_i\right) \\ &= \Phi\left(-\frac{\mathbf{x}_i' \boldsymbol{\beta}}{\sigma}\right).\end{aligned}$$

The conditional density function above zero is normal:

$$\sigma^{-1} \phi\left(\frac{y - \mathbf{x}_i' \boldsymbol{\beta}}{\sigma}\right), \quad y > 0.$$

Therefore, the density function for $y \geq 0$ can be written as

$$f(y \mid \mathbf{x}_i) = \Phi\left(-\frac{\mathbf{x}_i' \boldsymbol{\beta}}{\sigma}\right)^{1(y=0)} \left[\sigma^{-1} \phi\left(\frac{y - \mathbf{x}_i' \boldsymbol{\beta}}{\sigma}\right)\right]^{1(y>0)},$$

where $1(\cdot)$ is the indicator function.

Hence the log-likelihood is a mixture of the probit and the normal:

$$\begin{aligned}\log L(\boldsymbol{\beta}) &= \sum_{i=1}^n \log f(y_i \mid \mathbf{x}_i) \\ &= \sum_{y_i=0} \log \Phi\left(-\frac{\mathbf{x}_i' \boldsymbol{\beta}}{\sigma}\right) + \sum_{y_i>0} \log \left[\sigma^{-1} \phi\left(\frac{y_i - \mathbf{x}_i' \boldsymbol{\beta}}{\sigma}\right)\right].\end{aligned}$$

The MLE is the value $\hat{\boldsymbol{\beta}}$ which maximizes $\log L(\boldsymbol{\beta})$.

18.4 Sample Selection

The problem of sample selection arises when the sample is a non-random selection of potential observations. This occurs when the observed data is systematically different from the population of interest. For example, if you ask for volunteers for an experiment, and they wish to extrapolate the effects of the experiment on a general population, you should worry that the people who volunteer may be systematically different from the general population. This has great relevance for the evaluation of anti-poverty and job-training programs, where the goal is to assess the effect of “training” on the general population, not just on the volunteers.

A simple sample selection model can be written as the latent model

$$\begin{aligned}y_i &= \mathbf{x}_i' \boldsymbol{\beta} + e_{1i} \\ T_i &= 1(z_i' \boldsymbol{\gamma} + e_{0i} > 0)\end{aligned}$$

where $1(\cdot)$ is the indicator function. The dependent variable y_i is observed if (and only if) $T_i = 1$. Else it is unobserved.

For example, y_i could be a wage, which can be observed only if a person is employed. The equation for T_i is an equation specifying the probability that the person is employed.

The model is often completed by specifying that the errors are jointly normal

$$\begin{pmatrix} e_{0i} \\ e_{1i} \end{pmatrix} \sim N\left(0, \begin{pmatrix} 1 & \rho \\ \rho & \sigma^2 \end{pmatrix}\right).$$

It is presumed that we observe $\{\mathbf{x}_i, \mathbf{z}_i, T_i\}$ for all observations.

Under the normality assumption,

$$e_{1i} = \rho e_{0i} + v_i,$$

where v_i is independent of $e_{0i} \sim N(0, 1)$. A useful fact about the standard normal distribution is that

$$\mathbb{E}(e_{0i} \mid e_{0i} > -x) = \lambda(x) = \frac{\phi(x)}{\Phi(x)},$$

and the function $\lambda(x)$ is called the inverse Mills ratio.

The naive estimator of β is OLS regression of y_i on \mathbf{x}_i for those observations for which y_i is available. The problem is that this is equivalent to conditioning on the event $\{T_i = 1\}$. However,

$$\begin{aligned} \mathbb{E}(e_{1i} \mid T_i = 1, \mathbf{z}_i) &= \mathbb{E}(e_{1i} \mid \{e_{0i} > -\mathbf{z}'_i \gamma\}, \mathbf{z}_i) \\ &= \rho \mathbb{E}(e_{0i} \mid \{e_{0i} > -\mathbf{z}'_i \gamma\}, \mathbf{z}_i) + \mathbb{E}(v_i \mid \{e_{0i} > -\mathbf{z}'_i \gamma\}, \mathbf{z}_i) \\ &= \rho \lambda(\mathbf{z}'_i \gamma), \end{aligned}$$

which is non-zero. Thus

$$e_{1i} = \rho \lambda(\mathbf{z}'_i \gamma) + u_i,$$

where

$$\mathbb{E}(u_i \mid T_i = 1, \mathbf{z}_i) = 0.$$

Hence

$$y_i = \mathbf{x}'_i \beta + \rho \lambda(\mathbf{z}'_i \gamma) + u_i \tag{18.2}$$

is a valid regression equation for the observations for which $T_i = 1$.

Heckman (1979) observed that we could consistently estimate β and ρ from this equation, if γ were known. It is unknown, but also can be consistently estimated by a Probit model for selection. The “Heckit” estimator is thus calculated as follows

- Estimate $\hat{\gamma}$ from a Probit, using regressors \mathbf{z}_i . The binary dependent variable is T_i .
- Estimate $(\hat{\beta}, \hat{\rho})$ from OLS of y_i on \mathbf{x}_i and $\lambda(\mathbf{z}'_i \hat{\gamma})$.
- The OLS standard errors will be incorrect, as this is a two-step estimator. They can be corrected using a more complicated formula. Or, alternatively, by viewing the Probit/OLS estimation equations as a large joint GMM problem.

The Heckit estimator is frequently used to deal with problems of sample selection. However, the estimator is built on the assumption of normality, and the estimator can be quite sensitive to this assumption. Some modern econometric research is exploring how to relax the normality assumption.

The estimator can also work quite poorly if $\lambda(\mathbf{z}'_i \hat{\gamma})$ does not have much in-sample variation. This can happen if the Probit equation does not “explain” much about the selection choice. Another potential problem is that if $\mathbf{z}_i = \mathbf{x}_i$, then $\lambda(\mathbf{z}'_i \hat{\gamma})$ can be highly collinear with \mathbf{x}_i , so the second step OLS estimator will not be able to precisely estimate β . Based this observation, it is typically recommended to find a valid exclusion restriction: a variable should be in \mathbf{z}_i which is not in \mathbf{x}_i . If this is valid, it will ensure that $\lambda(\mathbf{z}'_i \hat{\gamma})$ is not collinear with \mathbf{x}_i , and hence improve the second stage estimator’s precision.

Chapter 19

Panel Data

A panel is a set of observations on individuals, collected over time. An observation is the pair $\{y_{it}, \mathbf{x}_{it}\}$, where the i subscript denotes the individual, and the t subscript denotes time. A panel may be *balanced*:

$$\{y_{it}, \mathbf{x}_{it}\} : t = 1, \dots, T; \quad i = 1, \dots, n,$$

or *unbalanced*:

$$\{y_{it}, \mathbf{x}_{it}\} : \text{For } i = 1, \dots, n, \quad t = \underline{t}_i, \dots, \bar{t}_i.$$

19.1 Individual-Effects Model

The standard panel data specification is that there is an individual-specific effect which enters linearly in the regression

$$y_{it} = \mathbf{x}'_{it}\boldsymbol{\beta} + u_i + e_{it}.$$

The typical maintained assumptions are that the individuals i are mutually independent, that u_i and e_{it} are independent, that e_{it} is iid across individuals and time, and that e_{it} is uncorrelated with \mathbf{x}_{it} .

OLS of y_{it} on \mathbf{x}_{it} is called pooled estimation. It is consistent if

$$\mathbb{E}(\mathbf{x}_{it}u_i) = 0 \tag{19.1}$$

If this condition fails, then OLS is inconsistent. (19.1) fails if the individual-specific unobserved effect u_i is correlated with the observed explanatory variables \mathbf{x}_{it} . This is often believed to be plausible if u_i is an omitted variable.

If (19.1) is true, however, OLS can be improved upon via a GLS technique. In either event, OLS appears a poor estimation choice.

Condition (19.1) is called the *random effects hypothesis*. It is a strong assumption, and most applied researchers try to avoid its use.

19.2 Fixed Effects

This is the most common technique for estimation of non-dynamic linear panel regressions.

The motivation is to allow u_i to be arbitrary, and have arbitrary correlated with \mathbf{x}_i . The goal is to eliminate u_i from the estimator, and thus achieve invariance.

There are several derivations of the estimator.

First, let

$$d_{ij} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{else} \end{cases},$$

and

$$\mathbf{d}_i = \begin{pmatrix} d_{i1} \\ \vdots \\ d_{in} \end{pmatrix},$$

an $n \times 1$ dummy vector with a “1” in the i 'th place. Let

$$\mathbf{u} = \begin{pmatrix} u_1 \\ \vdots \\ u_n \end{pmatrix}.$$

Then note that

$$u_i = \mathbf{d}_i' \mathbf{u},$$

and

$$y_{it} = \mathbf{x}_{it}' \boldsymbol{\beta} + \mathbf{d}_i' \mathbf{u} + e_{it}. \quad (19.2)$$

Observe that

$$\mathbb{E}(e_{it} \mid \mathbf{x}_{it}, \mathbf{d}_i) = 0,$$

so (19.2) is a valid regression, with \mathbf{d}_i as a regressor along with \mathbf{x}_i .

OLS on (19.2) yields estimator $(\hat{\boldsymbol{\beta}}, \hat{\mathbf{u}})$. Conventional inference applies.

Observe that

- This is generally consistent.
- If \mathbf{x}_{it} contains an intercept, it will be collinear with \mathbf{d}_i , so the intercept is typically omitted from \mathbf{x}_{it} .
- Any regressor in \mathbf{x}_{it} which is constant over time for all individuals (e.g., their gender) will be collinear with \mathbf{d}_i , so will have to be omitted.
- There are $n + k$ regression parameters, which is quite large as typically n is very large.

Computationally, you do not want to actually implement conventional OLS estimation, as the parameter space is too large. OLS estimation of $\boldsymbol{\beta}$ proceeds by the FWL theorem. Stacking the observations together:

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \mathbf{D}\mathbf{u} + e,$$

then by the FWL theorem,

$$\begin{aligned} \hat{\boldsymbol{\beta}} &= (\mathbf{X}'(\mathbf{I} - \mathbf{P}_D)\mathbf{X})^{-1} (\mathbf{X}'(\mathbf{I} - \mathbf{P}_D)\mathbf{y}) \\ &= (\mathbf{X}^*{}'\mathbf{X}^*)^{-1} (\mathbf{X}^*{}'\mathbf{y}^*), \end{aligned}$$

where

$$\begin{aligned} \mathbf{y}^* &= \mathbf{y} - \mathbf{D}(\mathbf{D}'\mathbf{D})^{-1}\mathbf{D}'\mathbf{y} \\ \mathbf{X}^* &= \mathbf{X} - \mathbf{D}(\mathbf{D}'\mathbf{D})^{-1}\mathbf{D}'\mathbf{X}. \end{aligned}$$

Since the regression of y_{it} on \mathbf{d}_i is a regression onto individual-specific dummies, the predicted value from these regressions is the individual specific mean \bar{y}_i , and the residual is the demean value

$$y_{it}^* = y_{it} - \bar{y}_i.$$

The fixed effects estimator $\hat{\boldsymbol{\beta}}$ is OLS of y_{it}^* on \mathbf{x}_{it}^* , the dependent variable and regressors in deviation-from-mean form.

Another derivation of the estimator is to take the equation

$$y_{it} = \mathbf{x}'_{it}\boldsymbol{\beta} + u_i + e_{it},$$

and then take individual-specific means by taking the average for the i 'th individual:

$$\frac{1}{T_i} \sum_{t=\underline{t}_i}^{\bar{t}_i} y_{it} = \frac{1}{T_i} \sum_{t=\underline{t}_i}^{\bar{t}_i} \mathbf{x}'_{it}\boldsymbol{\beta} + u_i + \frac{1}{T_i} \sum_{t=\underline{t}_i}^{\bar{t}_i} e_{it}$$

or

$$\bar{y}_i = \bar{\mathbf{x}}'_i\boldsymbol{\beta} + u_i + \bar{e}_i.$$

Subtracting, we find

$$y_{it}^* = \mathbf{x}'_{it}\boldsymbol{\beta} + e_{it}^*,$$

which is free of the individual-effect u_i .

19.3 Dynamic Panel Regression

A dynamic panel regression has a lagged dependent variable

$$y_{it} = \alpha y_{it-1} + \mathbf{x}'_{it}\boldsymbol{\beta} + u_i + e_{it}. \quad (19.3)$$

This is a model suitable for studying dynamic behavior of individual agents.

Unfortunately, the fixed effects estimator is inconsistent, at least if T is held finite as $n \rightarrow \infty$. This is because the sample mean of y_{it-1} is correlated with that of e_{it} .

The standard approach to estimate a dynamic panel is to combine first-differencing with IV or GMM. Taking first-differences of (19.3) eliminates the individual-specific effect:

$$\Delta y_{it} = \alpha \Delta y_{it-1} + \Delta \mathbf{x}'_{it}\boldsymbol{\beta} + \Delta e_{it}. \quad (19.4)$$

However, if e_{it} is iid, then it will be correlated with Δy_{it-1} :

$$\mathbb{E}(\Delta y_{it-1} \Delta e_{it}) = \mathbb{E}((y_{it-1} - y_{it-2})(e_{it} - e_{it-1})) = -\mathbb{E}(y_{it-1} e_{it-1}) = -\sigma_e^2.$$

So OLS on (19.4) will be inconsistent.

But if there are valid instruments, then IV or GMM can be used to estimate the equation. Typically, we use lags of the dependent variable, two periods back, as y_{t-2} is uncorrelated with Δe_{it} . Thus values of y_{it-k} , $k \geq 2$, are valid instruments.

Hence a valid estimator of α and $\boldsymbol{\beta}$ is to estimate (19.4) by IV using y_{t-2} as an instrument for Δy_{t-1} (which is just identified). Alternatively, GMM using y_{t-2} and y_{t-3} as instruments (which is overidentified, but loses a time-series observation).

A more sophisticated GMM estimator recognizes that for time-periods later in the sample, there are more instruments available, so the instrument list should be different for each equation. This is conveniently organized by the GMM principle, as this enables the moments from the different time-periods to be stacked together to create a list of all the moment conditions. A simple application of GMM yields the parameter estimates and standard errors.

Chapter 20

Nonparametric Density Estimation

20.1 Kernel Density Estimation

Let X be a random variable with continuous distribution $F(x)$ and density $f(x) = \frac{d}{dx}F(x)$. The goal is to estimate $f(x)$ from a random sample $\{X_1, \dots, X_n\}$. While $F(x)$ can be estimated by the EDF $\hat{F}(x) = n^{-1} \sum_{i=1}^n 1(X_i \leq x)$, we cannot define $\frac{d}{dx}\hat{F}(x)$ since $\hat{F}(x)$ is a step function. The standard **nonparametric** method to estimate $f(x)$ is based on **smoothing** using a kernel.

While we are typically interested in estimating the entire function $f(x)$, we can simply focus on the problem where x is a specific fixed number, and then see how the method generalizes to estimating the entire function.

Definition 20.1.1 $K(u)$ is a **second-order kernel function** if it is a symmetric zero-mean density function.

Three common choices for kernels include the **Normal**

$$K(u) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{u^2}{2}\right)$$

the **Epanechnikov**

$$K(u) = \begin{cases} \frac{3}{4}(1-u^2), & |u| \leq 1 \\ 0, & |u| > 1 \end{cases}$$

and the **Biweight** or **Quartic**

$$K(u) = \begin{cases} \frac{15}{16}(1-u^2)^2, & |u| \leq 1 \\ 0, & |u| > 1 \end{cases}$$

In practice, the choice between these three rarely makes a meaningful difference in the estimates.

The kernel functions are used to smooth the data. The amount of smoothing is controlled by the **bandwidth** $h > 0$. Let

$$K_h(u) = \frac{1}{h} K\left(\frac{u}{h}\right).$$

be the kernel K rescaled by the bandwidth h . The kernel density estimator of $f(x)$ is

$$\hat{f}(x) = \frac{1}{n} \sum_{i=1}^n K_h(X_i - x).$$

This estimator is the average of a set of weights. If a large number of the observations X_i are near x , then the weights are relatively large and $\hat{f}(x)$ is larger. Conversely, if only a few X_i are near x , then the weights are small and $\hat{f}(x)$ is small. The bandwidth h controls the meaning of “near”.

Interestingly, $\hat{f}(x)$ is a valid density. That is, $\hat{f}(x) \geq 0$ for all x , and

$$\begin{aligned}\int_{-\infty}^{\infty} \hat{f}(x) dx &= \int_{-\infty}^{\infty} \frac{1}{n} \sum_{i=1}^n K_h(X_i - x) dx \\ &= \frac{1}{n} \sum_{i=1}^n \int_{-\infty}^{\infty} K_h(X_i - x) dx \\ &= \frac{1}{n} \sum_{i=1}^n \int_{-\infty}^{\infty} K(u) du = 1\end{aligned}$$

where the second-to-last equality makes the change-of-variables $u = (X_i - x)/h$.

We can also calculate the moments of the density $\hat{f}(x)$. The mean is

$$\begin{aligned}\int_{-\infty}^{\infty} x \hat{f}(x) dx &= \frac{1}{n} \sum_{i=1}^n \int_{-\infty}^{\infty} x K_h(X_i - x) dx \\ &= \frac{1}{n} \sum_{i=1}^n \int_{-\infty}^{\infty} (X_i + uh) K(u) du \\ &= \frac{1}{n} \sum_{i=1}^n X_i \int_{-\infty}^{\infty} K(u) du + \frac{1}{n} \sum_{i=1}^n h \int_{-\infty}^{\infty} u K(u) du \\ &= \frac{1}{n} \sum_{i=1}^n X_i\end{aligned}$$

the sample mean of the X_i , where the second-to-last equality used the change-of-variables $u = (X_i - x)/h$ which has Jacobian h .

The second moment of the estimated density is

$$\begin{aligned}\int_{-\infty}^{\infty} x^2 \hat{f}(x) dx &= \frac{1}{n} \sum_{i=1}^n \int_{-\infty}^{\infty} x^2 K_h(X_i - x) dx \\ &= \frac{1}{n} \sum_{i=1}^n \int_{-\infty}^{\infty} (X_i + uh)^2 K(u) du \\ &= \frac{1}{n} \sum_{i=1}^n X_i^2 + \frac{2}{n} \sum_{i=1}^n X_i h \int_{-\infty}^{\infty} K(u) du + \frac{1}{n} \sum_{i=1}^n h^2 \int_{-\infty}^{\infty} u^2 K(u) du \\ &= \frac{1}{n} \sum_{i=1}^n X_i^2 + h^2 \sigma_K^2\end{aligned}$$

where

$$\sigma_K^2 = \int_{-\infty}^{\infty} u^2 K(u) du$$

is the variance of the kernel. It follows that the variance of the density $\hat{f}(x)$ is

$$\begin{aligned}\int_{-\infty}^{\infty} x^2 \hat{f}(x) dx - \left(\int_{-\infty}^{\infty} x \hat{f}(x) dx \right)^2 &= \frac{1}{n} \sum_{i=1}^n X_i^2 + h^2 \sigma_K^2 - \left(\frac{1}{n} \sum_{i=1}^n X_i \right)^2 \\ &= \hat{\sigma}^2 + h^2 \sigma_K^2\end{aligned}$$

Thus the variance of the estimated density is inflated by the factor $h^2 \sigma_K^2$ relative to the sample moment.

20.2 Asymptotic MSE for Kernel Estimates

For fixed x and bandwidth h observe that

$$\begin{aligned}\mathbb{E}K_h(X - x) &= \int_{-\infty}^{\infty} K_h(z - x) f(z) dz \\ &= \int_{-\infty}^{\infty} K_h(uh) f(x + hu) h du \\ &= \int_{-\infty}^{\infty} K(u) f(x + hu) du\end{aligned}$$

The second equality uses the change-of variables $u = (z - x)/h$. The last expression shows that the expected value is an average of $f(z)$ locally about x .

This integral (typically) is not analytically solvable, so we approximate it using a second order Taylor expansion of $f(x + hu)$ in the argument hu about $hu = 0$, which is valid as $h \rightarrow 0$. Thus

$$f(x + hu) \simeq f(x) + f'(x)hu + \frac{1}{2}f''(x)h^2u^2$$

and therefore

$$\begin{aligned}\mathbb{E}K_h(X - x) &\simeq \int_{-\infty}^{\infty} K(u) \left(f(x) + f'(x)hu + \frac{1}{2}f''(x)h^2u^2 \right) du \\ &= f(x) \int_{-\infty}^{\infty} K(u) du + f'(x)h \int_{-\infty}^{\infty} K(u) u du \\ &\quad + \frac{1}{2}f''(x)h^2 \int_{-\infty}^{\infty} K(u) u^2 du \\ &= f(x) + \frac{1}{2}f''(x)h^2\sigma_K^2.\end{aligned}$$

The bias of $\hat{f}(x)$ is then

$$\text{Bias}(x) = \mathbb{E}\hat{f}(x) - f(x) = \frac{1}{n} \sum_{i=1}^n \mathbb{E}K_h(X_i - x) - f(x) = \frac{1}{2}f''(x)h^2\sigma_K^2.$$

We see that the bias of $\hat{f}(x)$ at x depends on the second derivative $f''(x)$. The sharper the derivative, the greater the bias. Intuitively, the estimator $\hat{f}(x)$ smooths data local to $X_i = x$, so is estimating a smoothed version of $f(x)$. The bias results from this smoothing, and is larger the greater the curvature in $f(x)$.

We now examine the variance of $\hat{f}(x)$. Since it is an average of iid random variables, using first-order Taylor approximations and the fact that n^{-1} is of smaller order than $(nh)^{-1}$

$$\begin{aligned}\text{var}(\hat{f}(x)) &= \frac{1}{n} \text{var}(K_h(X_i - x)) \\ &= \frac{1}{n} \mathbb{E}K_h(X_i - x)^2 - \frac{1}{n} (\mathbb{E}K_h(X_i - x))^2 \\ &\simeq \frac{1}{nh^2} \int_{-\infty}^{\infty} K\left(\frac{z-x}{h}\right)^2 f(z) dz - \frac{1}{n} f(x)^2 \\ &= \frac{1}{nh} \int_{-\infty}^{\infty} K(u)^2 f(x + hu) du \\ &\simeq \frac{f(x)}{nh} \int_{-\infty}^{\infty} K(u)^2 du \\ &= \frac{f(x) R(K)}{nh}.\end{aligned}$$

where $R(K) = \int_{-\infty}^{\infty} K(u)^2 du$ is called the **roughness** of K .

Together, the asymptotic mean-squared error (AMSE) for fixed x is the sum of the approximate squared bias and approximate variance

$$AMSE_h(x) = \frac{1}{4} f''(x)^2 h^4 \sigma_K^4 + \frac{f(x) R(K)}{nh}.$$

A global measure of precision is the asymptotic mean integrated squared error (AMISE)

$$AMISE_h = \int AMSE_h(x) dx = \frac{h^4 \sigma_K^4 R(f'')}{4} + \frac{R(K)}{nh}. \tag{20.1}$$

where $R(f'') = \int (f''(x))^2 dx$ is the roughness of f'' . Notice that the first term (the squared bias) is increasing in h and the second term (the variance) is decreasing in nh . Thus for the AMISE to decline with n , we need $h \rightarrow 0$ but $nh \rightarrow \infty$. That is, h must tend to zero, but at a slower rate than n^{-1} .

Equation (20.1) is an asymptotic approximation to the MSE. We define the asymptotically optimal bandwidth h_0 as the value which minimizes this approximate MSE. That is,

$$h_0 = \underset{h}{\operatorname{argmin}} AMISE_h$$

It can be found by solving the first order condition

$$\frac{d}{dh} AMISE_h = h^3 \sigma_K^4 R(f'') - \frac{R(K)}{nh^2} = 0$$

yielding

$$h_0 = \left(\frac{R(K)}{\sigma_K^4 R(f'')} \right)^{1/5} n^{-1/5}. \tag{20.2}$$

This solution takes the form $h_0 = cn^{-1/5}$ where c is a function of K and f , but not of n . We thus say that the optimal bandwidth is of order $O(n^{-1/5})$. Note that this h declines to zero, but at a very slow rate.

In practice, how should the bandwidth be selected? This is a difficult problem, and there is a large and continuing literature on the subject. The asymptotically optimal choice given in (20.2) depends on $R(K)$, σ_K^2 , and $R(f'')$. The first two are determined by the kernel function. Their values for the three functions introduced in the previous section are given here.

K	$\sigma_K^2 = \int_{-\infty}^{\infty} u^2 K(u) du$	$R(K) = \int_{-\infty}^{\infty} K(u)^2 du$
Gaussian	1	$1/(2\sqrt{\pi})$
Epanechnikov	1/5	1/5
Biweight	1/7	5/7

An obvious difficulty is that $R(f'')$ is unknown. A classic simple solution proposed by Silverman (1986) has come to be known as the **reference bandwidth** or **Silverman's Rule-of-Thumb**. It uses formula (20.2) but replaces $R(f'')$ with $\hat{\sigma}^{-5} R(\phi'')$, where ϕ is the $N(0, 1)$ distribution and $\hat{\sigma}^2$ is an estimate of $\sigma^2 = \operatorname{var}(X)$. This choice for h gives an optimal rule when $f(x)$ is normal, and gives a nearly optimal rule when $f(x)$ is close to normal. The downside is that if the density is very far from normal, the rule-of-thumb h can be quite inefficient. We can calculate that $R(\phi'') = 3/(8\sqrt{\pi})$. Together with the above table, we find the reference rules for the three kernel functions introduced earlier.

- Gaussian Kernel: $h_{rule} = 1.06\hat{\sigma}n^{-1/5}$
- Epanechnikov Kernel: $h_{rule} = 2.34\hat{\sigma}n^{-1/5}$
- Biweight (Quartic) Kernel: $h_{rule} = 2.78\hat{\sigma}n^{-1/5}$

Unless you delve more deeply into kernel estimation methods the rule-of-thumb bandwidth is a good practical bandwidth choice, perhaps adjusted by visual inspection of the resulting estimate $\hat{f}(x)$. There are other approaches, but implementation can be delicate. I now discuss some of these choices. The **plug-in** approach is to estimate $R(f'')$ in a first step, and then plug this estimate into the formula (20.2). This is more treacherous than may first appear, as the optimal h for estimation of the roughness $R(f'')$ is quite different than the optimal h for estimation of $f(x)$. However, there are modern versions of this estimator work well, in particular the iterative method of Sheather and Jones (1991). Another popular choice for selection of h is **cross-validation**. This works by constructing an estimate of the MISE using leave-one-out estimators. There are some desirable properties of cross-validation bandwidths, but they are also known to converge very slowly to the optimal values. They are also quite ill-behaved when the data has some discretization (as is common in economics), in which case the cross-validation rule can sometimes select very small bandwidths leading to dramatically undersmoothed estimates. Fortunately there are remedies, which are known as **smoothed cross-validation** which is a close cousin of the **bootstrap**.

Appendix A

Matrix Algebra

A.1 Notation

A **scalar** a is a single number.

A **vector** \mathbf{a} is a $k \times 1$ list of numbers, typically arranged in a column. We write this as

$$\mathbf{a} = \begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_k \end{pmatrix}$$

Equivalently, a vector \mathbf{a} is an element of Euclidean k space, written as $\mathbf{a} \in \mathbb{R}^k$. If $k = 1$ then \mathbf{a} is a scalar.

A **matrix** \mathbf{A} is a $k \times r$ rectangular array of numbers, written as

$$\mathbf{A} = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1r} \\ a_{21} & a_{22} & \cdots & a_{2r} \\ \vdots & \vdots & & \vdots \\ a_{k1} & a_{k2} & \cdots & a_{kr} \end{bmatrix}$$

By convention a_{ij} refers to the element in the i 'th row and j 'th column of \mathbf{A} . If $r = 1$ then \mathbf{A} is a column vector. If $k = 1$ then \mathbf{A} is a row vector. If $r = k = 1$, then \mathbf{A} is a scalar.

A standard convention (which we will follow in this text whenever possible) is to denote scalars by lower-case italics (a), vectors by lower-case bold italics (\mathbf{a}), and matrices by upper-case bold italics (\mathbf{A}). Sometimes a matrix \mathbf{A} is denoted by the symbol (a_{ij}) .

A matrix can be written as a set of column vectors or as a set of row vectors. That is,

$$\mathbf{A} = [\mathbf{a}_1 \quad \mathbf{a}_2 \quad \cdots \quad \mathbf{a}_r] = \begin{bmatrix} \boldsymbol{\alpha}_1 \\ \boldsymbol{\alpha}_2 \\ \vdots \\ \boldsymbol{\alpha}_k \end{bmatrix}$$

where

$$\mathbf{a}_i = \begin{bmatrix} a_{1i} \\ a_{2i} \\ \vdots \\ a_{ki} \end{bmatrix}$$

are column vectors and

$$\boldsymbol{\alpha}_j = [a_{j1} \quad a_{j2} \quad \cdots \quad a_{jr}]$$

are row vectors.

The **transpose** of a matrix, denoted \mathbf{A}' , is obtained by flipping the matrix on its diagonal. Thus

$$\mathbf{A}' = \begin{bmatrix} a_{11} & a_{21} & \cdots & a_{k1} \\ a_{12} & a_{22} & \cdots & a_{k2} \\ \vdots & \vdots & & \vdots \\ a_{1r} & a_{2r} & \cdots & a_{kr} \end{bmatrix}$$

Alternatively, letting $\mathbf{B} = \mathbf{A}'$, then $b_{ij} = a_{ji}$. Note that if \mathbf{A} is $k \times r$, then \mathbf{A}' is $r \times k$. If \mathbf{a} is a $k \times 1$ vector, then \mathbf{a}' is a $1 \times k$ row vector. An alternative notation for the transpose of \mathbf{A} is \mathbf{A}^\top .

A matrix is **square** if $k = r$. A square matrix is **symmetric** if $\mathbf{A} = \mathbf{A}'$, which requires $a_{ij} = a_{ji}$. A square matrix is **diagonal** if the off-diagonal elements are all zero, so that $a_{ij} = 0$ if $i \neq j$. A square matrix is **upper (lower) diagonal** if all elements below (above) the diagonal equal zero.

An important diagonal matrix is the **identity matrix**, which has ones on the diagonal. The $k \times k$ identity matrix is denoted as

$$\mathbf{I}_k = \begin{bmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & 1 \end{bmatrix}.$$

A **partitioned matrix** takes the form

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} & \cdots & \mathbf{A}_{1r} \\ \mathbf{A}_{21} & \mathbf{A}_{22} & \cdots & \mathbf{A}_{2r} \\ \vdots & \vdots & & \vdots \\ \mathbf{A}_{k1} & \mathbf{A}_{k2} & \cdots & \mathbf{A}_{kr} \end{bmatrix}$$

where the A_{ij} denote matrices, vectors and/or scalars.

A.2 Matrix Addition

If the matrices $\mathbf{A} = (a_{ij})$ and $\mathbf{B} = (b_{ij})$ are of the same order, we define the sum

$$\mathbf{A} + \mathbf{B} = (a_{ij} + b_{ij}).$$

Matrix addition follows the commutative and associative laws:

$$\begin{aligned} \mathbf{A} + \mathbf{B} &= \mathbf{B} + \mathbf{A} \\ \mathbf{A} + (\mathbf{B} + \mathbf{C}) &= (\mathbf{A} + \mathbf{B}) + \mathbf{C}. \end{aligned}$$

A.3 Matrix Multiplication

If \mathbf{A} is $k \times r$ and c is real, we define their product as

$$\mathbf{Ac} = c\mathbf{A} = (a_{ij}c).$$

If \mathbf{a} and \mathbf{b} are both $k \times 1$, then their inner product is

$$\mathbf{a}'\mathbf{b} = a_1b_1 + a_2b_2 + \cdots + a_kb_k = \sum_{j=1}^k a_jb_j.$$

Note that $\mathbf{a}'\mathbf{b} = \mathbf{b}'\mathbf{a}$. We say that two vectors \mathbf{a} and \mathbf{b} are **orthogonal** if $\mathbf{a}'\mathbf{b} = 0$.

If \mathbf{A} is $k \times r$ and \mathbf{B} is $r \times s$, so that the number of columns of \mathbf{A} equals the number of rows of \mathbf{B} , we say that \mathbf{A} and \mathbf{B} are **conformable**. In this event the matrix product \mathbf{AB} is defined. Writing \mathbf{A} as a set of row vectors and \mathbf{B} as a set of column vectors (each of length r), then the matrix product is defined as

$$\begin{aligned} \mathbf{AB} &= \begin{bmatrix} \mathbf{a}'_1 \\ \mathbf{a}'_2 \\ \vdots \\ \mathbf{a}'_k \end{bmatrix} \begin{bmatrix} \mathbf{b}_1 & \mathbf{b}_2 & \cdots & \mathbf{b}_s \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{a}'_1 \mathbf{b}_1 & \mathbf{a}'_1 \mathbf{b}_2 & \cdots & \mathbf{a}'_1 \mathbf{b}_s \\ \mathbf{a}'_2 \mathbf{b}_1 & \mathbf{a}'_2 \mathbf{b}_2 & \cdots & \mathbf{a}'_2 \mathbf{b}_s \\ \vdots & \vdots & & \vdots \\ \mathbf{a}'_k \mathbf{b}_1 & \mathbf{a}'_k \mathbf{b}_2 & \cdots & \mathbf{a}'_k \mathbf{b}_s \end{bmatrix}. \end{aligned}$$

Matrix multiplication is not commutative: in general $\mathbf{AB} \neq \mathbf{BA}$. However, it is associative and distributive:

$$\begin{aligned} \mathbf{A}(\mathbf{BC}) &= (\mathbf{AB})\mathbf{C} \\ \mathbf{A}(\mathbf{B} + \mathbf{C}) &= \mathbf{AB} + \mathbf{AC} \end{aligned}$$

An alternative way to write the matrix product is to use matrix partitions. For example,

$$\begin{aligned} \mathbf{AB} &= \begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{B}_{11} & \mathbf{B}_{12} \\ \mathbf{B}_{21} & \mathbf{B}_{22} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{A}_{11}\mathbf{B}_{11} + \mathbf{A}_{12}\mathbf{B}_{21} & \mathbf{A}_{11}\mathbf{B}_{12} + \mathbf{A}_{12}\mathbf{B}_{22} \\ \mathbf{A}_{21}\mathbf{B}_{11} + \mathbf{A}_{22}\mathbf{B}_{21} & \mathbf{A}_{21}\mathbf{B}_{12} + \mathbf{A}_{22}\mathbf{B}_{22} \end{bmatrix}. \end{aligned}$$

As another example,

$$\begin{aligned} \mathbf{AB} &= \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 & \cdots & \mathbf{A}_r \end{bmatrix} \begin{bmatrix} \mathbf{B}_1 \\ \mathbf{B}_2 \\ \vdots \\ \mathbf{B}_r \end{bmatrix} \\ &= \mathbf{A}_1\mathbf{B}_1 + \mathbf{A}_2\mathbf{B}_2 + \cdots + \mathbf{A}_r\mathbf{B}_r \\ &= \sum_{j=1}^r \mathbf{A}_j\mathbf{B}_j \end{aligned}$$

An important property of the identity matrix is that if \mathbf{A} is $k \times r$, then $\mathbf{AI}_r = \mathbf{A}$ and $\mathbf{I}_k\mathbf{A} = \mathbf{A}$. The $k \times r$ matrix \mathbf{A} , $r \leq k$, is called orthogonal if $\mathbf{A}'\mathbf{A} = \mathbf{I}_r$.

A.4 Trace

The **trace** of a $k \times k$ square matrix \mathbf{A} is the sum of its diagonal elements

$$\text{tr}(\mathbf{A}) = \sum_{i=1}^k a_{ii}.$$

Some straightforward properties for square matrices \mathbf{A} and \mathbf{B} and real c are

$$\begin{aligned}\operatorname{tr}(c\mathbf{A}) &= c \operatorname{tr}(\mathbf{A}) \\ \operatorname{tr}(\mathbf{A}') &= \operatorname{tr}(\mathbf{A}) \\ \operatorname{tr}(\mathbf{A} + \mathbf{B}) &= \operatorname{tr}(\mathbf{A}) + \operatorname{tr}(\mathbf{B}) \\ \operatorname{tr}(\mathbf{I}_k) &= k.\end{aligned}$$

Also, for $k \times r$ \mathbf{A} and $r \times k$ \mathbf{B} we have

$$\operatorname{tr}(\mathbf{AB}) = \operatorname{tr}(\mathbf{BA}). \quad (\text{A.1})$$

Indeed,

$$\begin{aligned}\operatorname{tr}(\mathbf{AB}) &= \operatorname{tr} \begin{bmatrix} \mathbf{a}'_1 \mathbf{b}_1 & \mathbf{a}'_1 \mathbf{b}_2 & \cdots & \mathbf{a}'_1 \mathbf{b}_k \\ \mathbf{a}'_2 \mathbf{b}_1 & \mathbf{a}'_2 \mathbf{b}_2 & \cdots & \mathbf{a}'_2 \mathbf{b}_k \\ \vdots & \vdots & \cdots & \vdots \\ \mathbf{a}'_k \mathbf{b}_1 & \mathbf{a}'_k \mathbf{b}_2 & \cdots & \mathbf{a}'_k \mathbf{b}_k \end{bmatrix} \\ &= \sum_{i=1}^k \mathbf{a}'_i \mathbf{b}_i \\ &= \sum_{i=1}^k \mathbf{b}'_i \mathbf{a}_i \\ &= \operatorname{tr}(\mathbf{BA}).\end{aligned}$$

A.5 Rank and Inverse

The rank of the $k \times r$ matrix ($r \leq k$)

$$\mathbf{A} = [\mathbf{a}_1 \quad \mathbf{a}_2 \quad \cdots \quad \mathbf{a}_r]$$

is the number of linearly independent columns \mathbf{a}_j , and is written as $\operatorname{rank}(\mathbf{A})$. We say that \mathbf{A} has full rank if $\operatorname{rank}(\mathbf{A}) = r$.

A square $k \times k$ matrix \mathbf{A} is said to be **nonsingular** if it has full rank, e.g. $\operatorname{rank}(\mathbf{A}) = k$. This means that there is no $k \times 1$ $\mathbf{c} \neq \mathbf{0}$ such that $\mathbf{Ac} = \mathbf{0}$.

If a square $k \times k$ matrix \mathbf{A} is nonsingular then there exists a unique matrix $k \times k$ matrix \mathbf{A}^{-1} called the **inverse** of \mathbf{A} which satisfies

$$\mathbf{AA}^{-1} = \mathbf{A}^{-1}\mathbf{A} = \mathbf{I}_k.$$

For non-singular \mathbf{A} and \mathbf{C} , some important properties include

$$\begin{aligned}\mathbf{AA}^{-1} &= \mathbf{A}^{-1}\mathbf{A} = \mathbf{I}_k \\ (\mathbf{A}^{-1})' &= (\mathbf{A}')^{-1} \\ (\mathbf{AC})^{-1} &= \mathbf{C}^{-1}\mathbf{A}^{-1} \\ (\mathbf{A} + \mathbf{C})^{-1} &= \mathbf{A}^{-1}(\mathbf{A}^{-1} + \mathbf{C}^{-1})^{-1}\mathbf{C}^{-1} \\ \mathbf{A}^{-1} - (\mathbf{A} + \mathbf{C})^{-1} &= \mathbf{A}^{-1}(\mathbf{A}^{-1} + \mathbf{C}^{-1})^{-1}\mathbf{A}^{-1}\end{aligned}$$

Also, if \mathbf{A} is an orthogonal matrix, then $\mathbf{A}^{-1} = \mathbf{A}'$.

Another useful result for non-singular \mathbf{A} is known as the **Woodbury matrix identity**

$$(\mathbf{A} + \mathbf{BCD})^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1}\mathbf{BC}(\mathbf{C} + \mathbf{CDA}^{-1}\mathbf{BC})^{-1}\mathbf{CDA}^{-1}. \quad (\text{A.2})$$

In particular, for $\mathbf{C} = -1$, $\mathbf{B} = \mathbf{b}$ and $\mathbf{D} = \mathbf{b}'$ for vector \mathbf{b} we find what is known as the **Sherman–Morrison formula**

$$(\mathbf{A} - \mathbf{bb}')^{-1} = \mathbf{A}^{-1} + (1 - \mathbf{b}'\mathbf{A}^{-1}\mathbf{b})^{-1} \mathbf{A}^{-1}\mathbf{bb}'\mathbf{A}^{-1}. \quad (\text{A.3})$$

The following fact about inverting partitioned matrices is quite useful.

$$\begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix}^{-1} = \begin{bmatrix} \mathbf{A}^{11} & \mathbf{A}^{12} \\ \mathbf{A}^{21} & \mathbf{A}^{22} \end{bmatrix} = \begin{bmatrix} \mathbf{A}_{11}^{-1} & -\mathbf{A}_{11}^{-1}\mathbf{A}_{12}\mathbf{A}_{22}^{-1} \\ -\mathbf{A}_{22}^{-1}\mathbf{A}_{21}\mathbf{A}_{11}^{-1} & \mathbf{A}_{22}^{-1} \end{bmatrix} \quad (\text{A.4})$$

where $\mathbf{A}_{11.2} = \mathbf{A}_{11} - \mathbf{A}_{12}\mathbf{A}_{22}^{-1}\mathbf{A}_{21}$ and $\mathbf{A}_{22.1} = \mathbf{A}_{22} - \mathbf{A}_{21}\mathbf{A}_{11}^{-1}\mathbf{A}_{12}$. There are alternative algebraic representations for the components. For example, using the Woodbury matrix identity you can show the following alternative expressions

$$\begin{aligned} \mathbf{A}^{11} &= \mathbf{A}_{11}^{-1} + \mathbf{A}_{11}^{-1}\mathbf{A}_{12}\mathbf{A}_{22.1}^{-1}\mathbf{A}_{21}\mathbf{A}_{11}^{-1} \\ \mathbf{A}^{22} &= \mathbf{A}_{22}^{-1} + \mathbf{A}_{22}^{-1}\mathbf{A}_{21}\mathbf{A}_{11.2}^{-1}\mathbf{A}_{12}\mathbf{A}_{22}^{-1} \\ \mathbf{A}^{12} &= -\mathbf{A}_{11}^{-1}\mathbf{A}_{12}\mathbf{A}_{22.1}^{-1} \\ \mathbf{A}^{21} &= -\mathbf{A}_{22}^{-1}\mathbf{A}_{21}\mathbf{A}_{11.2}^{-1} \end{aligned}$$

Even if a matrix \mathbf{A} does not possess an inverse, we can still define the **Moore–Penrose generalized inverse** \mathbf{A}^- as the matrix which satisfies

$$\begin{aligned} \mathbf{AA}^- \mathbf{A} &= \mathbf{A} \\ \mathbf{A}^- \mathbf{AA}^- &= \mathbf{A}^- \\ \mathbf{AA}^- &\text{ is symmetric} \\ \mathbf{A}^- \mathbf{A} &\text{ is symmetric} \end{aligned}$$

For any matrix \mathbf{A} , the Moore–Penrose generalized inverse \mathbf{A}^- exists and is unique.

For example, if

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{11} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$$

and when \mathbf{A}_{11}^{-1} exists then

$$\mathbf{A}^- = \begin{bmatrix} \mathbf{A}_{11}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}.$$

A.6 Determinant

The **determinant** is a measure of the volume of a square matrix.

While the determinant is widely used, its precise definition is rarely needed. However, we present the definition here for completeness. Let $\mathbf{A} = (a_{ij})$ be a general $k \times k$ matrix. Let $\pi = (j_1, \dots, j_k)$ denote a permutation of $(1, \dots, k)$. There are $k!$ such permutations. There is a unique count of the number of inversions of the indices of such permutations (relative to the natural order $(1, \dots, k)$), and let $\varepsilon_\pi = +1$ if this count is even and $\varepsilon_\pi = -1$ if the count is odd. Then the determinant of \mathbf{A} is defined as

$$\det \mathbf{A} = \sum_{\pi} \varepsilon_\pi a_{1j_1} a_{2j_2} \cdots a_{kj_k}.$$

For example, if \mathbf{A} is 2×2 , then the two permutations of $(1, 2)$ are $(1, 2)$ and $(2, 1)$, for which $\varepsilon_{(1,2)} = 1$ and $\varepsilon_{(2,1)} = -1$. Thus

$$\begin{aligned}\det \mathbf{A} &= \varepsilon_{(1,2)} a_{11} a_{22} + \varepsilon_{(2,1)} a_{21} a_{12} \\ &= a_{11} a_{22} - a_{12} a_{21}.\end{aligned}$$

Some properties include

- $\det(\mathbf{A}) = \det(\mathbf{A}')$
- $\det(c\mathbf{A}) = c^k \det \mathbf{A}$
- $\det(\mathbf{AB}) = (\det \mathbf{A})(\det \mathbf{B})$
- $\det(\mathbf{A}^{-1}) = (\det \mathbf{A})^{-1}$
- $\det \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{bmatrix} = (\det \mathbf{D}) \det(\mathbf{A} - \mathbf{BD}^{-1}\mathbf{C})$ if $\det \mathbf{D} \neq 0$
- $\det \mathbf{A} \neq 0$ if and only if \mathbf{A} is nonsingular
- If \mathbf{A} is triangular (upper or lower), then $\det \mathbf{A} = \prod_{i=1}^k a_{ii}$
- If \mathbf{A} is orthogonal, then $\det \mathbf{A} = \pm 1$

A.7 Eigenvalues

The characteristic equation of a $k \times k$ square matrix \mathbf{A} is

$$\det(\mathbf{A} - \lambda \mathbf{I}_k) = 0.$$

The left side is a polynomial of degree k in λ so it has exactly k roots, which are not necessarily distinct and may be real or complex. They are called the **latent roots** or **characteristic roots** or **eigenvalues** of \mathbf{A} . If λ_i is an eigenvalue of \mathbf{A} , then $\mathbf{A} - \lambda_i \mathbf{I}_k$ is singular so there exists a non-zero vector \mathbf{h}_i such that

$$(\mathbf{A} - \lambda_i \mathbf{I}_k) \mathbf{h}_i = \mathbf{0}.$$

The vector \mathbf{h}_i is called a **latent vector** or **characteristic vector** or **eigenvector** of \mathbf{A} corresponding to λ_i .

We now state some useful properties. Let λ_i and \mathbf{h}_i , $i = 1, \dots, k$ denote the k eigenvalues and eigenvectors of a square matrix \mathbf{A} . Let $\mathbf{\Lambda}$ be a diagonal matrix with the characteristic roots in the diagonal, and let $\mathbf{H} = [\mathbf{h}_1 \cdots \mathbf{h}_k]$.

- $\det(\mathbf{A}) = \prod_{i=1}^k \lambda_i$
- $\text{tr}(\mathbf{A}) = \sum_{i=1}^k \lambda_i$
- \mathbf{A} is non-singular if and only if all its characteristic roots are non-zero.
- If \mathbf{A} has distinct characteristic roots, there exists a nonsingular matrix \mathbf{P} such that $\mathbf{A} = \mathbf{P}^{-1} \mathbf{\Lambda} \mathbf{P}$ and $\mathbf{P} \mathbf{A} \mathbf{P}^{-1} = \mathbf{\Lambda}$.
- If \mathbf{A} is symmetric, then $\mathbf{A} = \mathbf{H} \mathbf{\Lambda} \mathbf{H}'$ and $\mathbf{H}' \mathbf{A} \mathbf{H} = \mathbf{\Lambda}$, and the characteristic roots are all real. $\mathbf{A} = \mathbf{H} \mathbf{\Lambda} \mathbf{H}'$ is called the **spectral decomposition** of a matrix.
- When the eigenvalues of $k \times k$ \mathbf{A} are real they are written in descending order $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_k$. We also write $\lambda_{\min}(\mathbf{A}) = \lambda_k = \min\{\lambda_\ell\}$ and $\lambda_{\max}(\mathbf{A}) = \lambda_1 = \max\{\lambda_\ell\}$.

- The characteristic roots of \mathbf{A}^{-1} are $\lambda_1^{-1}, \lambda_2^{-1}, \dots, \lambda_k^{-1}$.
- The matrix \mathbf{H} has the **orthonormal** properties $\mathbf{H}'\mathbf{H} = \mathbf{I}$ and $\mathbf{H}\mathbf{H}' = \mathbf{I}$.
- $\mathbf{H}^{-1} = \mathbf{H}'$ and $(\mathbf{H}')^{-1} = \mathbf{H}$

A.8 Positive Definiteness

We say that a $k \times k$ symmetric square matrix \mathbf{A} is **positive semi-definite** if for all $\mathbf{c} \neq \mathbf{0}$, $\mathbf{c}'\mathbf{A}\mathbf{c} \geq 0$. This is written as $\mathbf{A} \geq 0$. We say that \mathbf{A} is **positive definite** if for all $\mathbf{c} \neq \mathbf{0}$, $\mathbf{c}'\mathbf{A}\mathbf{c} > 0$. This is written as $\mathbf{A} > 0$.

Some properties include:

- If $\mathbf{A} = \mathbf{G}'\mathbf{B}\mathbf{G}$ with $\mathbf{B} \geq 0$ and some matrix \mathbf{G} , then \mathbf{A} is positive semi-definite. (For any $\mathbf{c} \neq \mathbf{0}$, $\mathbf{c}'\mathbf{A}\mathbf{c} = \alpha'\mathbf{B}\alpha \geq 0$ where $\alpha = \mathbf{G}\mathbf{c}$.) If \mathbf{G} has full column rank and $\mathbf{B} > 0$, then \mathbf{A} is positive definite.
- If \mathbf{A} is positive definite, then \mathbf{A} is non-singular and \mathbf{A}^{-1} exists. Furthermore, $\mathbf{A}^{-1} > 0$.
- $\mathbf{A} > 0$ if and only if it is symmetric and all its characteristic roots are positive.
- By the spectral decomposition, $\mathbf{A} = \mathbf{H}\mathbf{\Lambda}\mathbf{H}'$ where $\mathbf{H}'\mathbf{H} = \mathbf{I}$ and $\mathbf{\Lambda}$ is diagonal with non-negative diagonal elements. All diagonal elements of $\mathbf{\Lambda}$ are strictly positive if (and only if) $\mathbf{A} > 0$.
- If $\mathbf{A} > 0$ then $\mathbf{A}^{-1} = \mathbf{H}\mathbf{\Lambda}^{-1}\mathbf{H}'$.
- If $\mathbf{A} \geq 0$ and $\text{rank}(\mathbf{A}) = r < k$ then $\mathbf{A}^- = \mathbf{H}\mathbf{\Lambda}^-\mathbf{H}'$ where \mathbf{A}^- is the Moore-Penrose generalized inverse, and $\mathbf{\Lambda}^- = \text{diag}(\lambda_1^{-1}, \lambda_2^{-1}, \dots, \lambda_k^{-1}, 0, \dots, 0)$
- If $\mathbf{A} \geq 0$ we can find a matrix \mathbf{B} such that $\mathbf{A} = \mathbf{B}\mathbf{B}'$. We call \mathbf{B} a **matrix square root** of \mathbf{A} . The matrix \mathbf{B} need not be unique. One way to construct \mathbf{B} is to use the spectral decomposition $\mathbf{A} = \mathbf{H}\mathbf{\Lambda}\mathbf{H}'$ where $\mathbf{\Lambda}$ is diagonal, and then set $\mathbf{B} = \mathbf{H}\mathbf{\Lambda}^{1/2}$. There is a unique root \mathbf{B} which is also positive semi-definite $\mathbf{B} \geq 0$.

A square matrix \mathbf{A} is **idempotent** if $\mathbf{A}\mathbf{A} = \mathbf{A}$. If \mathbf{A} is idempotent and symmetric then all its characteristic roots equal either zero or one and is thus positive semi-definite. To see this, note that we can write $\mathbf{A} = \mathbf{H}\mathbf{\Lambda}\mathbf{H}'$ where \mathbf{H} is orthogonal and $\mathbf{\Lambda}$ contains the r (real) characteristic roots. Then

$$\mathbf{A} = \mathbf{A}\mathbf{A} = \mathbf{H}\mathbf{\Lambda}\mathbf{H}'\mathbf{H}\mathbf{\Lambda}\mathbf{H}' = \mathbf{H}\mathbf{\Lambda}^2\mathbf{H}'.$$

By the uniqueness of the characteristic roots, we deduce that $\mathbf{\Lambda}^2 = \mathbf{\Lambda}$ and $\lambda_i^2 = \lambda_i$ for $i = 1, \dots, r$. Hence they must equal either 0 or 1. It follows that the spectral decomposition of idempotent \mathbf{A} takes the form

$$\mathbf{A} = \mathbf{H} \begin{bmatrix} \mathbf{I}_{k-r} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{H}' \quad (\text{A.5})$$

with $\mathbf{H}'\mathbf{H} = \mathbf{I}_k$. Additionally, $\text{tr}(\mathbf{A}) = \text{rank}(\mathbf{A})$.

If \mathbf{A} is **idempotent** then $\mathbf{I} - \mathbf{A}$ is also idempotent.

One useful fact is that \mathbf{A} is idempotent then for any conformable vector \mathbf{c} ,

$$\mathbf{c}'\mathbf{A}\mathbf{c} \leq \mathbf{c}'\mathbf{c} \quad (\text{A.6})$$

$$\mathbf{c}'(\mathbf{I} - \mathbf{A})\mathbf{c} \leq \mathbf{c}'\mathbf{c} \quad (\text{A.7})$$

To see this, note that

$$\mathbf{c}'\mathbf{c} = \mathbf{c}'\mathbf{A}\mathbf{c} + \mathbf{c}'(\mathbf{I} - \mathbf{A})\mathbf{c}.$$

Since \mathbf{A} and $\mathbf{I} - \mathbf{A}$ are idempotent, they are both positive semi-definite, so both $\mathbf{c}'\mathbf{A}\mathbf{c}$ and $\mathbf{c}'(\mathbf{I} - \mathbf{A})\mathbf{c}$ are non-negative. Thus they must satisfy (A.6)-(A.7),.

A.9 Matrix Calculus

Let $\mathbf{x} = (x_1, \dots, x_k)$ be $k \times 1$ and $g(\mathbf{x}) = g(x_1, \dots, x_k) : \mathbb{R}^k \rightarrow \mathbb{R}$. The vector derivative is

$$\frac{\partial}{\partial \mathbf{x}} g(\mathbf{x}) = \begin{pmatrix} \frac{\partial}{\partial x_1} g(\mathbf{x}) \\ \vdots \\ \frac{\partial}{\partial x_k} g(\mathbf{x}) \end{pmatrix}$$

and

$$\frac{\partial}{\partial \mathbf{x}'} g(\mathbf{x}) = \left(\frac{\partial}{\partial x_1} g(\mathbf{x}) \quad \cdots \quad \frac{\partial}{\partial x_k} g(\mathbf{x}) \right).$$

Some properties are now summarized.

- $\frac{\partial}{\partial \mathbf{x}} (\mathbf{a}' \mathbf{x}) = \frac{\partial}{\partial \mathbf{x}} (\mathbf{x}' \mathbf{a}) = \mathbf{a}$
- $\frac{\partial}{\partial \mathbf{x}'} (\mathbf{A} \mathbf{x}) = \mathbf{A}$
- $\frac{\partial}{\partial \mathbf{x}} (\mathbf{x}' \mathbf{A} \mathbf{x}) = (\mathbf{A} + \mathbf{A}') \mathbf{x}$
- $\frac{\partial^2}{\partial \mathbf{x} \partial \mathbf{x}'} (\mathbf{x}' \mathbf{A} \mathbf{x}) = \mathbf{A} + \mathbf{A}'$

A.10 Kronecker Products and the Vec Operator

Let $\mathbf{A} = [\mathbf{a}_1 \ \mathbf{a}_2 \ \cdots \ \mathbf{a}_n]$ be $m \times n$. The **vec** of \mathbf{A} , denoted by $\text{vec}(\mathbf{A})$, is the $mn \times 1$ vector

$$\text{vec}(\mathbf{A}) = \begin{pmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \\ \vdots \\ \mathbf{a}_n \end{pmatrix}.$$

Let $\mathbf{A} = (a_{ij})$ be an $m \times n$ matrix and let \mathbf{B} be any matrix. The **Kronecker product** of \mathbf{A} and \mathbf{B} , denoted $\mathbf{A} \otimes \mathbf{B}$, is the matrix

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{11} \mathbf{B} & a_{12} \mathbf{B} & \cdots & a_{1n} \mathbf{B} \\ a_{21} \mathbf{B} & a_{22} \mathbf{B} & \cdots & a_{2n} \mathbf{B} \\ \vdots & \vdots & \cdots & \vdots \\ a_{m1} \mathbf{B} & a_{m2} \mathbf{B} & \cdots & a_{mn} \mathbf{B} \end{bmatrix}.$$

Some important properties are now summarized. These results hold for matrices for which all matrix multiplications are conformable.

- $(\mathbf{A} + \mathbf{B}) \otimes \mathbf{C} = \mathbf{A} \otimes \mathbf{C} + \mathbf{B} \otimes \mathbf{C}$
- $(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = \mathbf{AC} \otimes \mathbf{BD}$
- $\mathbf{A} \otimes (\mathbf{B} \otimes \mathbf{C}) = (\mathbf{A} \otimes \mathbf{B}) \otimes \mathbf{C}$
- $(\mathbf{A} \otimes \mathbf{B})' = \mathbf{A}' \otimes \mathbf{B}'$
- $\text{tr}(\mathbf{A} \otimes \mathbf{B}) = \text{tr}(\mathbf{A}) \text{tr}(\mathbf{B})$
- If \mathbf{A} is $m \times m$ and \mathbf{B} is $n \times n$, $\det(\mathbf{A} \otimes \mathbf{B}) = (\det(\mathbf{A}))^n (\det(\mathbf{B}))^m$
- $(\mathbf{A} \otimes \mathbf{B})^{-1} = \mathbf{A}^{-1} \otimes \mathbf{B}^{-1}$
- If $\mathbf{A} > 0$ and $\mathbf{B} > 0$ then $\mathbf{A} \otimes \mathbf{B} > 0$
- $\text{vec}(\mathbf{ABC}) = (\mathbf{C}' \otimes \mathbf{A}) \text{vec}(\mathbf{B})$
- $\text{tr}(\mathbf{ABCD}) = \text{vec}(\mathbf{D}')' (\mathbf{C}' \otimes \mathbf{A}) \text{vec}(\mathbf{B})$

A.11 Vector and Matrix Norms

The **Euclidean norm** of an $m \times 1$ vector \mathbf{a} is

$$\begin{aligned}\|\mathbf{a}\| &= (\mathbf{a}'\mathbf{a})^{1/2} \\ &= \left(\sum_{i=1}^m a_i^2 \right)^{1/2}.\end{aligned}$$

The **Frobenius norm** of an $m \times n$ matrix \mathbf{A} is

$$\begin{aligned}\|\mathbf{A}\| &= \|\text{vec}(\mathbf{A})\| \\ &= (\text{tr}(\mathbf{A}'\mathbf{A}))^{1/2} \\ &= \left(\sum_{i=1}^m \sum_{j=1}^n a_{ij}^2 \right)^{1/2}.\end{aligned}$$

If an $m \times m$ matrix \mathbf{A} is symmetric with eigenvalues λ_ℓ , $\ell = 1, \dots, m$, then

$$\|\mathbf{A}\| = \left(\sum_{\ell=1}^m \lambda_\ell^2 \right)^{1/2}.$$

To see this, by the spectral decomposition $\mathbf{A} = \mathbf{H}\mathbf{\Lambda}\mathbf{H}'$ with $\mathbf{H}'\mathbf{H} = \mathbf{I}$ and $\mathbf{\Lambda} = \text{diag}\{\lambda_1, \dots, \lambda_m\}$, so

$$\|\mathbf{A}\| = (\text{tr}(\mathbf{H}\mathbf{\Lambda}\mathbf{H}'\mathbf{H}\mathbf{\Lambda}\mathbf{H}'))^{1/2} = (\text{tr}(\mathbf{\Lambda}\mathbf{\Lambda}))^{1/2} = \left(\sum_{\ell=1}^m \lambda_\ell^2 \right)^{1/2}. \quad (\text{A.8})$$

A useful calculation is for any $m \times 1$ vectors \mathbf{a} and \mathbf{b} , using (A.1),

$$\|\mathbf{ab}'\| = \text{tr}(\mathbf{ba}'\mathbf{ab}')^{1/2} = (\mathbf{b}'\mathbf{ba}'\mathbf{a})^{1/2} = \|\mathbf{a}\| \|\mathbf{b}\|$$

and in particular

$$\|\mathbf{aa}'\| = \|\mathbf{a}\|^2. \quad (\text{A.9})$$

There are other matrix norms. Another norm of frequent use is the spectral norm

$$\|\mathbf{A}\|_S = (\lambda_{\max}(\mathbf{A}'\mathbf{A}))^{1/2}$$

where $\lambda_{\max}(\mathbf{B})$ denotes the largest eigenvalue of the matrix \mathbf{B} .

A.12 Matrix Inequalities

Schwarz Inequality: For any $m \times 1$ vectors \mathbf{a} and \mathbf{b} ,

$$|\mathbf{a}'\mathbf{b}| \leq \|\mathbf{a}\| \|\mathbf{b}\|. \quad (\text{A.10})$$

Schwarz Matrix Inequality: For any $m \times n$ matrices \mathbf{A} and \mathbf{B} ,

$$\|\mathbf{A}'\mathbf{B}\| \leq \|\mathbf{A}\| \|\mathbf{B}\|. \quad (\text{A.11})$$

Triangle Inequality: For any $m \times n$ matrices \mathbf{A} and \mathbf{B} ,

$$\|\mathbf{A} + \mathbf{B}\| \leq \|\mathbf{A}\| + \|\mathbf{B}\|. \quad (\text{A.12})$$

Trace Inequality. For any $m \times m$ matrices \mathbf{A} and \mathbf{B} such that \mathbf{A} is symmetric and $\mathbf{B} \geq 0$

$$\text{tr}(\mathbf{AB}) \leq \lambda_{\max}(\mathbf{A}) \text{tr}(\mathbf{B}) \quad (\text{A.13})$$

where $\lambda_{\max}(\mathbf{A})$ is the largest eigenvalue of \mathbf{A} .

Quadratic Inequality. For any $m \times 1$ \mathbf{b} and $m \times m$ symmetric matrix \mathbf{A}

$$\mathbf{b}'\mathbf{A}\mathbf{b} \leq \lambda_{\max}(\mathbf{A}) \mathbf{b}'\mathbf{b} \quad (\text{A.14})$$

Norm Inequality. For any $m \times m$ matrices \mathbf{A} and \mathbf{B} such that $\mathbf{A} \geq 0$ and $\mathbf{B} \geq 0$

$$\|\mathbf{AB}\| \leq \lambda_{\max}(\mathbf{A}) \|\mathbf{B}\| \quad (\text{A.15})$$

where $\lambda_{\max}(\mathbf{A})$ is the largest eigenvalue of \mathbf{A} .

Eigenvalue Product Inequality. For any $k \times k$ matrices $\mathbf{A} \geq 0$ and $\mathbf{B} \geq 0$, the eigenvalues $\lambda_\ell(\mathbf{AB})$ are real and satisfy

$$\lambda_{\min}(\mathbf{A}) \lambda_{\min}(\mathbf{B}) \leq \lambda_\ell(\mathbf{AB}) = \lambda_\ell(\mathbf{A}^{1/2}\mathbf{B}\mathbf{A}^{1/2}) \leq \lambda_{\max}(\mathbf{A}) \lambda_{\max}(\mathbf{B}) \quad (\text{A.16})$$

(Zhang and Zhang, 2006, Corollary 11)

Jensen's Inequality. If $g(\cdot) : \mathbb{R} \rightarrow \mathbb{R}$ is convex, then for any non-negative weights a_j such that $\sum_{j=1}^m a_j = 1$, and any real numbers x_j

$$g\left(\sum_{j=1}^m a_j x_j\right) \leq \sum_{j=1}^m a_j g(x_j). \quad (\text{A.17})$$

In particular, setting $a_j = 1/m$, then

$$g\left(\frac{1}{m} \sum_{j=1}^m x_j\right) \leq \frac{1}{m} \sum_{j=1}^m g(x_j). \quad (\text{A.18})$$

Loève's c_r Inequality. For $r > 0$,

$$\left|\sum_{j=1}^m a_j\right|^r \leq c_r \sum_{j=1}^m |a_j|^r \quad (\text{A.19})$$

where $c_r = 1$ when $r \leq 1$ and $c_r = m^{r-1}$ when $r \geq 1$.

c_2 Inequality. For any $m \times 1$ vectors \mathbf{a} and \mathbf{b} ,

$$(\mathbf{a} + \mathbf{b})'(\mathbf{a} + \mathbf{b}) \leq 2\mathbf{a}'\mathbf{a} + 2\mathbf{b}'\mathbf{b} \quad (\text{A.20})$$

Proof of Schwarz Inequality: First, suppose that $\|\mathbf{b}\| = 0$. Then $\mathbf{b} = \mathbf{0}$ and both $|\mathbf{a}'\mathbf{b}| = 0$ and $\|\mathbf{a}\| \|\mathbf{b}\| = 0$ so the inequality is true. Second, suppose that $\|\mathbf{b}\| > 0$ and define $\mathbf{c} = \mathbf{a} - \mathbf{b}(\mathbf{b}'\mathbf{b})^{-1} \mathbf{b}'\mathbf{a}$. Since \mathbf{c} is a vector, $\mathbf{c}'\mathbf{c} \geq 0$. Thus

$$0 \leq \mathbf{c}'\mathbf{c} = \mathbf{a}'\mathbf{a} - (\mathbf{a}'\mathbf{b})^2 / (\mathbf{b}'\mathbf{b}).$$

Rearranging, this implies that

$$(\mathbf{a}'\mathbf{b})^2 \leq (\mathbf{a}'\mathbf{a})(\mathbf{b}'\mathbf{b}).$$

Taking the square root of each side yields the result. ■

Proof of Schwarz Matrix Inequality: Partition $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_n]$ and $\mathbf{B} = [\mathbf{b}_1, \dots, \mathbf{b}_n]$. Then by partitioned matrix multiplication, the definition of the matrix Euclidean norm and the Schwarz inequality

$$\begin{aligned} \|\mathbf{A}'\mathbf{B}\| &= \left\| \begin{array}{ccc} \mathbf{a}'_1\mathbf{b}_1 & \mathbf{a}'_1\mathbf{b}_2 & \cdots \\ \mathbf{a}'_2\mathbf{b}_1 & \mathbf{a}'_2\mathbf{b}_2 & \cdots \\ \vdots & \vdots & \ddots \end{array} \right\| \\ &\leq \left\| \begin{array}{ccc} \|\mathbf{a}_1\| \|\mathbf{b}_1\| & \|\mathbf{a}_1\| \|\mathbf{b}_2\| & \cdots \\ \|\mathbf{a}_2\| \|\mathbf{b}_1\| & \|\mathbf{a}_2\| \|\mathbf{b}_2\| & \cdots \\ \vdots & \vdots & \ddots \end{array} \right\| \\ &= \left(\sum_{i=1}^n \sum_{j=1}^n \|\mathbf{a}_i\|^2 \|\mathbf{b}_j\|^2 \right)^{1/2} \\ &= \left(\sum_{i=1}^n \|\mathbf{a}_i\|^2 \right)^{1/2} \left(\sum_{i=1}^n \|\mathbf{b}_i\|^2 \right)^{1/2} \\ &= \left(\sum_{i=1}^n \sum_{j=1}^m \mathbf{a}_{ji}^2 \right)^{1/2} \left(\sum_{i=1}^n \sum_{j=1}^m \|\mathbf{b}_{ji}\|^2 \right)^{1/2} \\ &= \|\mathbf{A}\| \|\mathbf{B}\| \end{aligned}$$

■

Proof of Triangle Inequality: Let $\mathbf{a} = \text{vec}(\mathbf{A})$ and $\mathbf{b} = \text{vec}(\mathbf{B})$. Then by the definition of the matrix norm and the Schwarz Inequality

$$\begin{aligned} \|\mathbf{A} + \mathbf{B}\|^2 &= \|\mathbf{a} + \mathbf{b}\|^2 \\ &= \mathbf{a}'\mathbf{a} + 2\mathbf{a}'\mathbf{b} + \mathbf{b}'\mathbf{b} \\ &\leq \mathbf{a}'\mathbf{a} + 2|\mathbf{a}'\mathbf{b}| + \mathbf{b}'\mathbf{b} \\ &\leq \|\mathbf{a}\|^2 + 2\|\mathbf{a}\| \|\mathbf{b}\| + \|\mathbf{b}\|^2 \\ &= (\|\mathbf{a}\| + \|\mathbf{b}\|)^2 \\ &= (\|\mathbf{A}\| + \|\mathbf{B}\|)^2 \end{aligned}$$

■

Proof of Trace Inequality. By the spectral decomposition for symmetric matrices, $\mathbf{A} = \mathbf{H}\mathbf{\Lambda}\mathbf{H}'$ where $\mathbf{\Lambda}$ has the eigenvalues λ_j of \mathbf{A} on the diagonal and \mathbf{H} is orthonormal. Define $\mathbf{C} = \mathbf{H}'\mathbf{B}\mathbf{H}$ which has non-negative diagonal elements C_{jj} since \mathbf{B} is positive semi-definite. Then

$$\text{tr}(\mathbf{A}\mathbf{B}) = \text{tr}(\mathbf{\Lambda}\mathbf{C}) = \sum_{j=1}^m \lambda_j C_{jj} \leq \max_j \lambda_j \sum_{j=1}^m C_{jj} = \lambda_{\max}(\mathbf{A}) \text{tr}(\mathbf{C})$$

where the inequality uses the fact that $C_{jj} \geq 0$. But note that

$$\text{tr}(\mathbf{C}) = \text{tr}(\mathbf{H}'\mathbf{B}\mathbf{H}) = \text{tr}(\mathbf{H}\mathbf{H}'\mathbf{B}) = \text{tr}(\mathbf{B})$$

since \mathbf{H} is orthonormal. Thus $\text{tr}(\mathbf{AB}) \leq \lambda_{\max}(\mathbf{A}) \text{tr}(\mathbf{B})$ as stated. ■

Proof of Quadratic Inequality: In the Trace Inequality set $\mathbf{B} = \mathbf{b}\mathbf{b}'$ and note $\text{tr}(\mathbf{AB}) = \mathbf{b}'\mathbf{A}\mathbf{b}$ and $\text{tr}(\mathbf{B}) = \mathbf{b}'\mathbf{b}$. ■

Proof of Norm Inequality. Using the Trace Inequality

$$\begin{aligned} \|\mathbf{AB}\| &= (\text{tr}(\mathbf{BAAB}))^{1/2} \\ &= (\text{tr}(\mathbf{AABB}))^{1/2} \\ &\leq (\lambda_{\max}(\mathbf{AA}) \text{tr}(\mathbf{BB}))^{1/2} \\ &= \lambda_{\max}(\mathbf{A}) \|\mathbf{B}\|. \end{aligned}$$

The final equality holds since $\mathbf{A} \geq 0$ implies that $\lambda_{\max}(\mathbf{AA}) = \lambda_{\max}(\mathbf{A})^2$. ■

Proof of Jensen's Inequality (A.17). By the definition of convexity, for any $\lambda \in [0, 1]$

$$g(\lambda x_1 + (1 - \lambda)x_2) \leq \lambda g(x_1) + (1 - \lambda)g(x_2). \quad (\text{A.21})$$

This implies

$$\begin{aligned} g\left(\sum_{j=1}^m a_j x_j\right) &= g\left(a_1 g(x_1) + (1 - a_1) \sum_{j=2}^m \frac{a_j}{1 - a_1} x_j\right) \\ &\leq a_1 g(x_1) + (1 - a_1) g\left(\sum_{j=2}^m b_j x_j\right). \end{aligned}$$

where $b_j = a_j/(1 - a_1)$ and $\sum_{j=2}^m b_j = 1$. By another application of (A.21) this is bounded by

$$a_1 g(x_1) + (1 - a_1) \left(b_2 g(x_2) + (1 - b_2) g\left(\sum_{j=2}^m c_j x_j\right) \right) = a_1 g(x_1) + a_2 g(x_2) + (1 - a_1)(1 - b_2) g\left(\sum_{j=2}^m c_j x_j\right)$$

where $c_j = b_j/(1 - b_2)$. By repeated application of (A.21) we obtain (A.17). ■

Proof of Loève's c_r Inequality. For $r \geq 1$ this is simply a rewriting of the finite form Jensen's inequality (A.18) with $g(u) = u^r$. For $r < 1$, define $b_j = |a_j| / \left(\sum_{j=1}^m |a_j|\right)$. The facts that $0 \leq b_j \leq 1$ and $r < 1$ imply $b_j \leq b_j^r$ and thus

$$1 = \sum_{j=1}^m b_j \leq \sum_{j=1}^m b_j^r$$

which implies

$$\left(\sum_{j=1}^m |a_j|\right)^r \leq \sum_{j=1}^m |a_j|^r.$$

The proof is completed by observing that

$$\left(\sum_{j=1}^m a_j\right)^r \leq \left(\sum_{j=1}^m |a_j|\right)^r.$$

■

Proof of c_2 Inequality. By the c_r inequality, $(a_j + b_j)^2 \leq 2a_j^2 + 2b_j^2$. Thus

$$\begin{aligned}(\mathbf{a} + \mathbf{b})'(\mathbf{a} + \mathbf{b}) &= \sum_{j=1}^m (a_j + b_j)^2 \\ &\leq 2 \sum_{j=1}^m a_j^2 + 2 \sum_{j=1}^m b_j^2 \\ &= 2\mathbf{a}'\mathbf{a} + 2\mathbf{b}'\mathbf{b}\end{aligned}$$

■

Appendix B

Probability

B.1 Foundations

The set S of all possible outcomes of an experiment is called the **sample space** for the experiment. Take the simple example of tossing a coin. There are two outcomes, heads and tails, so we can write $S = \{H, T\}$. If two coins are tossed in sequence, we can write the four outcomes as $S = \{HH, HT, TH, TT\}$.

An **event** A is any collection of possible outcomes of an experiment. An event is a subset of S , including S itself and the null set \emptyset . Continuing the two coin example, one event is $A = \{HH, HT\}$, the event that the first coin is heads. We say that A and B are **disjoint** or **mutually exclusive** if $A \cap B = \emptyset$. For example, the sets $\{HH, HT\}$ and $\{TH\}$ are disjoint. Furthermore, if the sets A_1, A_2, \dots are pairwise disjoint and $\cup_{i=1}^{\infty} A_i = S$, then the collection A_1, A_2, \dots is called a **partition** of S .

The following are elementary set operations:

Union: $A \cup B = \{x : x \in A \text{ or } x \in B\}$.

Intersection: $A \cap B = \{x : x \in A \text{ and } x \in B\}$.

Complement: $A^c = \{x : x \notin A\}$.

The following are useful properties of set operations.

Commutativity: $A \cup B = B \cup A$; $A \cap B = B \cap A$.

Associativity: $A \cup (B \cup C) = (A \cup B) \cup C$; $A \cap (B \cap C) = (A \cap B) \cap C$.

Distributive Laws: $A \cap (B \cup C) = (A \cap B) \cup (A \cap C)$; $A \cup (B \cap C) = (A \cup B) \cap (A \cup C)$.

DeMorgan's Laws: $(A \cup B)^c = A^c \cap B^c$; $(A \cap B)^c = A^c \cup B^c$.

A **probability function** assigns probabilities (numbers between 0 and 1) to events A in S . This is straightforward when S is countable; when S is uncountable we must be somewhat more careful. A set \mathcal{B} is called a **sigma algebra** (or Borel field) if $\emptyset \in \mathcal{B}$, $A \in \mathcal{B}$ implies $A^c \in \mathcal{B}$, and $A_1, A_2, \dots \in \mathcal{B}$ implies $\cup_{i=1}^{\infty} A_i \in \mathcal{B}$. A simple example is $\{\emptyset, S\}$ which is known as the trivial sigma algebra. For any sample space S , let \mathcal{B} be the smallest sigma algebra which contains all of the open sets in S . When S is countable, \mathcal{B} is simply the collection of all subsets of S , including \emptyset and S . When S is the real line, then \mathcal{B} is the collection of all open and closed intervals. We call \mathcal{B} the sigma algebra associated with S . We only define probabilities for events contained in \mathcal{B} .

We now can give the axiomatic definition of probability. Given S and \mathcal{B} , a probability function \Pr satisfies $\Pr(S) = 1$, $\Pr(A) \geq 0$ for all $A \in \mathcal{B}$, and if $A_1, A_2, \dots \in \mathcal{B}$ are pairwise disjoint, then $\Pr(\cup_{i=1}^{\infty} A_i) = \sum_{i=1}^{\infty} \Pr(A_i)$.

Some important properties of the probability function include the following

- $\Pr(\emptyset) = 0$
- $\Pr(A) \leq 1$
- $\Pr(A^c) = 1 - \Pr(A)$

- $\Pr(B \cap A^c) = \Pr(B) - \Pr(A \cap B)$
- $\Pr(A \cup B) = \Pr(A) + \Pr(B) - \Pr(A \cap B)$
- If $A \subset B$ then $\Pr(A) \leq \Pr(B)$
- Bonferroni's Inequality: $\Pr(A \cap B) \geq \Pr(A) + \Pr(B) - 1$
- Boole's Inequality: $\Pr(A \cup B) \leq \Pr(A) + \Pr(B)$

For some elementary probability models, it is useful to have simple rules to count the number of objects in a set. These counting rules are facilitated by using the binomial coefficients which are defined for nonnegative integers n and r , $n \geq r$, as

$$\binom{n}{r} = \frac{n!}{r!(n-r)!}.$$

When counting the number of objects in a set, there are two important distinctions. Counting may be **with replacement** or **without replacement**. Counting may be **ordered** or **unordered**. For example, consider a lottery where you pick six numbers from the set $1, 2, \dots, 49$. This selection is without replacement if you are not allowed to select the same number twice, and is with replacement if this is allowed. Counting is ordered or not depending on whether the sequential order of the numbers is relevant to winning the lottery. Depending on these two distinctions, we have four expressions for the number of objects (possible arrangements) of size r from n objects.

	Without Replacement	With Replacement
Ordered	$\frac{n!}{(n-r)!}$	n^r
Unordered	$\binom{n}{r}$	$\binom{n+r-1}{r}$

In the lottery example, if counting is unordered and without replacement, the number of potential combinations is $\binom{49}{6} = 13,983,816$.

If $\Pr(B) > 0$ the **conditional probability** of the event A given the event B is

$$\Pr(A | B) = \frac{\Pr(A \cap B)}{\Pr(B)}.$$

For any B , the conditional probability function is a valid probability function where S has been replaced by B . Rearranging the definition, we can write

$$\Pr(A \cap B) = \Pr(A | B) \Pr(B)$$

which is often quite useful. We can say that the occurrence of B has no information about the likelihood of event A when $\Pr(A | B) = \Pr(A)$, in which case we find

$$\Pr(A \cap B) = \Pr(A) \Pr(B) \tag{B.1}$$

We say that the events A and B are **statistically independent** when (B.1) holds. Furthermore, we say that the collection of events A_1, \dots, A_k are **mutually independent** when for any subset $\{A_i : i \in I\}$,

$$\Pr\left(\bigcap_{i \in I} A_i\right) = \prod_{i \in I} \Pr(A_i).$$

Theorem 3 (*Bayes' Rule*). For any set B and any partition A_1, A_2, \dots of the sample space, then for each $i = 1, 2, \dots$

$$\Pr(A_i | B) = \frac{\Pr(B | A_i) \Pr(A_i)}{\sum_{j=1}^{\infty} \Pr(B | A_j) \Pr(A_j)}$$

B.2 Random Variables

A **random variable** X is a function from a sample space S into the real line. This induces a new sample space – the real line – and a new probability function on the real line. Typically, we denote random variables by uppercase letters such as X , and use lower case letters such as x for potential values and realized values. (This is in contrast to the notation adopted for most of the textbook.) For a random variable X we define its **cumulative distribution function** (CDF) as

$$F(x) = \Pr(X \leq x). \quad (\text{B.2})$$

Sometimes we write this as $F_X(x)$ to denote that it is the CDF of X . A function $F(x)$ is a CDF if and only if the following three properties hold:

1. $\lim_{x \rightarrow -\infty} F(x) = 0$ and $\lim_{x \rightarrow \infty} F(x) = 1$
2. $F(x)$ is nondecreasing in x
3. $F(x)$ is right-continuous

We say that the random variable X is **discrete** if $F(x)$ is a step function. In the latter case, the range of X consists of a countable set of real numbers τ_1, \dots, τ_r . The probability function for X takes the form

$$\Pr(X = \tau_j) = \pi_j, \quad j = 1, \dots, r \quad (\text{B.3})$$

where $0 \leq \pi_j \leq 1$ and $\sum_{j=1}^r \pi_j = 1$.

We say that the random variable X is **continuous** if $F(x)$ is continuous in x . In this case $\Pr(X = \tau) = 0$ for all $\tau \in R$ so the representation (B.3) is unavailable. Instead, we represent the relative probabilities by the **probability density function** (PDF)

$$f(x) = \frac{d}{dx}F(x)$$

so that

$$F(x) = \int_{-\infty}^x f(u)du$$

and

$$\Pr(a \leq X \leq b) = \int_a^b f(u)du.$$

These expressions only make sense if $F(x)$ is differentiable. While there are examples of continuous random variables which do not possess a PDF, these cases are unusual and are typically ignored.

A function $f(x)$ is a PDF if and only if $f(x) \geq 0$ for all $x \in R$ and $\int_{-\infty}^{\infty} f(x)dx = 1$.

B.3 Expectation

For any measurable real function g , we define the **mean** or **expectation** $\mathbb{E}g(X)$ as follows. If X is discrete,

$$\mathbb{E}g(X) = \sum_{j=1}^r g(\tau_j)\pi_j,$$

and if X is continuous

$$\mathbb{E}g(X) = \int_{-\infty}^{\infty} g(x)f(x)dx.$$

The latter is well defined and finite if

$$\int_{-\infty}^{\infty} |g(x)| f(x)dx < \infty. \quad (\text{B.4})$$

If (B.4) does not hold, evaluate

$$I_1 = \int_{g(x)>0} g(x)f(x)dx$$

$$I_2 = - \int_{g(x)<0} g(x)f(x)dx$$

If $I_1 = \infty$ and $I_2 < \infty$ then we define $\mathbb{E}g(X) = \infty$. If $I_1 < \infty$ and $I_2 = \infty$ then we define $\mathbb{E}g(X) = -\infty$. If both $I_1 = \infty$ and $I_2 = \infty$ then $\mathbb{E}g(X)$ is undefined.

Since $\mathbb{E}(a + bX) = a + b\mathbb{E}X$, we say that expectation is a linear operator.

For $m > 0$, we define the m 'th **moment** of X as $\mathbb{E}X^m$ and the m 'th **central moment** as $\mathbb{E}(X - \mathbb{E}X)^m$.

Two special moments are the **mean** $\mu = \mathbb{E}X$ and **variance** $\sigma^2 = \mathbb{E}(X - \mu)^2 = \mathbb{E}X^2 - \mu^2$. We call $\sigma = \sqrt{\sigma^2}$ the **standard deviation** of X . We can also write $\sigma^2 = \text{var}(X)$. For example, this allows the convenient expression $\text{var}(a + bX) = b^2 \text{var}(X)$.

The **moment generating function** (MGF) of X is

$$M(\lambda) = \mathbb{E} \exp(\lambda X).$$

The MGF does not necessarily exist. However, when it does and $\mathbb{E}|X|^m < \infty$ then

$$\left. \frac{d^m}{d\lambda^m} M(\lambda) \right|_{\lambda=0} = \mathbb{E}(X^m)$$

which is why it is called the moment generating function.

More generally, the **characteristic function** (CF) of X is

$$C(\lambda) = \mathbb{E} \exp(i\lambda X)$$

where $i = \sqrt{-1}$ is the imaginary unit. The CF always exists, and when $\mathbb{E}|X|^m < \infty$

$$\left. \frac{d^m}{d\lambda^m} C(\lambda) \right|_{\lambda=0} = i^m \mathbb{E}(X^m).$$

The L^p **norm**, $p \geq 1$, of the random variable X is

$$\|X\|_p = (\mathbb{E}|X|^p)^{1/p}.$$

B.4 Gamma Function

The gamma function is defined for $\alpha > 0$ as

$$\Gamma(\alpha) = \int_0^\infty x^{\alpha-1} \exp(-x) dx.$$

It satisfies the property

$$\Gamma(1 + \alpha) = \Gamma(\alpha)\alpha$$

so for positive integers n ,

$$\Gamma(n) = (n - 1)!$$

Special values include

$$\Gamma(1) = 1$$

and

$$\Gamma\left(\frac{1}{2}\right) = \pi^{1/2}.$$

Sterling's formula is an expansion for the its logarithm

$$\log \Gamma(\alpha) = \frac{1}{2} \log(2\pi) + \left(\alpha - \frac{1}{2}\right) \log \alpha - z + \frac{1}{12\alpha} - \frac{1}{360\alpha^3} + \frac{1}{1260\alpha^5} + \dots$$

B.5 Common Distributions

For reference, we now list some important discrete distribution function.

Bernoulli

$$\begin{aligned}\Pr(X = x) &= p^x(1-p)^{1-x}, & x = 0, 1; & & 0 \leq p \leq 1 \\ \mathbb{E}X &= p \\ \text{var}(X) &= p(1-p)\end{aligned}$$

Binomial

$$\begin{aligned}\Pr(X = x) &= \binom{n}{x} p^x (1-p)^{n-x}, & x = 0, 1, \dots, n; & & 0 \leq p \leq 1 \\ \mathbb{E}X &= np \\ \text{var}(X) &= np(1-p)\end{aligned}$$

Geometric

$$\begin{aligned}\Pr(X = x) &= p(1-p)^{x-1}, & x = 1, 2, \dots; & & 0 \leq p \leq 1 \\ \mathbb{E}X &= \frac{1}{p} \\ \text{var}(X) &= \frac{1-p}{p^2}\end{aligned}$$

Multinomial

$$\begin{aligned}\Pr(X_1 = x_1, X_2 = x_2, \dots, X_m = x_m) &= \frac{n!}{x_1! x_2! \cdots x_m!} p_1^{x_1} p_2^{x_2} \cdots p_m^{x_m}, \\ x_1 + \cdots + x_m &= n; \\ p_1 + \cdots + p_m &= 1 \\ \mathbb{E}X_i &= p_i \\ \text{var}(X_i) &= np_i(1-p_i) \\ \text{cov}(X_i, X_j) &= -np_i p_j\end{aligned}$$

Negative Binomial

$$\begin{aligned}\Pr(X = x) &= \frac{\Gamma(r+x)}{x! \Gamma(r)} p^r (1-p)^{x-1}, & x = 0, 1, 2, \dots; & & 0 \leq p \leq 1 \\ \mathbb{E}X &= \frac{r(1-p)}{p} \\ \text{var}(X) &= \frac{r(1-p)}{p^2}\end{aligned}$$

Poisson

$$\begin{aligned}\Pr(X = x) &= \frac{\exp(-\lambda) \lambda^x}{x!}, & x = 0, 1, 2, \dots, & & \lambda > 0 \\ \mathbb{E}X &= \lambda \\ \text{var}(X) &= \lambda\end{aligned}$$

We now list some important continuous distributions.

Beta

$$f(x) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1}(1-x)^{\beta-1}, \quad 0 \leq x \leq 1; \quad \alpha > 0, \beta > 0$$

$$\mu = \frac{\alpha}{\alpha + \beta}$$

$$\text{var}(X) = \frac{\alpha\beta}{(\alpha + \beta + 1)(\alpha + \beta)^2}$$

Cauchy

$$f(x) = \frac{1}{\pi(1+x^2)}, \quad -\infty < x < \infty$$

$$\mathbb{E}X = \infty$$

$$\text{var}(X) = \infty$$

Exponential

$$f(x) = \frac{1}{\theta} \exp\left(-\frac{x}{\theta}\right), \quad 0 \leq x < \infty; \quad \theta > 0$$

$$\mathbb{E}X = \theta$$

$$\text{var}(X) = \theta^2$$

Logistic

$$f(x) = \frac{\exp(-x)}{(1 + \exp(-x))^2}, \quad -\infty < x < \infty;$$

$$\mathbb{E}X = 0$$

$$\text{var}(X) = \frac{\pi^2}{3}$$

Lognormal

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma x} \exp\left(-\frac{(\log x - \mu)^2}{2\sigma^2}\right), \quad 0 < x < \infty; \quad \sigma > 0$$

$$\mathbb{E}X = \exp(\mu + \sigma^2/2)$$

$$\text{var}(X) = \exp(2\mu + 2\sigma^2) - \exp(2\mu + \sigma^2)$$

Pareto

$$f(x) = \frac{\beta\alpha^\beta}{x^{\beta+1}}, \quad \alpha \leq x < \infty, \quad \alpha > 0, \quad \beta > 0$$

$$\mathbb{E}X = \frac{\beta\alpha}{\beta - 1}, \quad \beta > 1$$

$$\text{var}(X) = \frac{\beta\alpha^2}{(\beta - 1)^2(\beta - 2)}, \quad \beta > 2$$

Uniform

$$f(x) = \frac{1}{b-a}, \quad a \leq x \leq b$$

$$\mathbb{E}X = \frac{a+b}{2}$$

$$\text{var}(X) = \frac{(b-a)^2}{12}$$

Weibull

$$f(x) = \frac{\gamma}{\beta} x^{\gamma-1} \exp\left(-\frac{x^\gamma}{\beta}\right), \quad 0 \leq x < \infty; \quad \gamma > 0, \beta > 0$$

$$\mathbb{E}X = \beta^{1/\gamma} \Gamma\left(1 + \frac{1}{\gamma}\right)$$

$$\text{var}(X) = \beta^{2/\gamma} \left(\Gamma\left(1 + \frac{2}{\gamma}\right) - \Gamma^2\left(1 + \frac{1}{\gamma}\right) \right)$$

Gamma

$$f(x) = \frac{1}{\Gamma(\alpha)\theta^\alpha} x^{\alpha-1} \exp\left(-\frac{x}{\theta}\right), \quad 0 \leq x < \infty; \quad \alpha > 0, \theta > 0$$

$$\mathbb{E}X = \alpha\theta$$

$$\text{var}(X) = \alpha\theta^2$$

Chi-Square

$$f(x) = \frac{1}{\Gamma(r/2)2^{r/2}} x^{r/2-1} \exp\left(-\frac{x}{2}\right), \quad 0 \leq x < \infty; \quad r > 0$$

$$\mathbb{E}X = r$$

$$\text{var}(X) = 2r$$

Normal

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), \quad -\infty < x < \infty; \quad -\infty < \mu < \infty, \sigma^2 > 0$$

$$\mathbb{E}X = \mu$$

$$\text{var}(X) = \sigma^2$$

Student t

$$f(x) = \frac{\Gamma\left(\frac{r+1}{2}\right)}{\sqrt{r\pi}\Gamma\left(\frac{r}{2}\right)} \left(1 + \frac{x^2}{r}\right)^{-\left(\frac{r+1}{2}\right)}, \quad -\infty < x < \infty; \quad r > 0$$

$$\mathbb{E}X = 0 \text{ if } r > 1$$

$$\text{var}(X) = \frac{r}{r-2} \text{ if } r > 2$$

B.6 Multivariate Random Variables

A pair of bivariate random variables (X, Y) is a function from the sample space into \mathbb{R}^2 . The joint CDF of (X, Y) is

$$F(x, y) = \Pr(X \leq x, Y \leq y).$$

If F is continuous, the joint probability density function is

$$f(x, y) = \frac{\partial^2}{\partial x \partial y} F(x, y).$$

For a Borel measurable set $A \in \mathcal{R}^2$,

$$\Pr((X, Y) \in A) = \int \int_A f(x, y) dx dy$$

For any measurable function $g(x, y)$,

$$\mathbb{E}g(X, Y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} g(x, y) f(x, y) dx dy.$$

The **marginal distribution** of X is

$$\begin{aligned} F_X(x) &= \Pr(X \leq x) \\ &= \lim_{y \rightarrow \infty} F(x, y) \\ &= \int_{-\infty}^x \int_{-\infty}^{\infty} f(x, y) dy dx \end{aligned}$$

so the **marginal density** of X is

$$f_X(x) = \frac{d}{dx} F_X(x) = \int_{-\infty}^{\infty} f(x, y) dy.$$

Similarly, the marginal density of Y is

$$f_Y(y) = \int_{-\infty}^{\infty} f(x, y) dx.$$

The random variables X and Y are defined to be **independent** if $f(x, y) = f_X(x)f_Y(y)$. Furthermore, X and Y are independent if and only if there exist functions $g(x)$ and $h(y)$ such that $f(x, y) = g(x)h(y)$.

If X and Y are independent, then

$$\begin{aligned} \mathbb{E}(g(X)h(Y)) &= \int \int g(x)h(y)f(y, x) dy dx \\ &= \int \int g(x)h(y)f_Y(y)f_X(x) dy dx \\ &= \int g(x)f_X(x) dx \int h(y)f_Y(y) dy \\ &= \mathbb{E}g(X) \mathbb{E}h(Y). \end{aligned} \tag{B.5}$$

if the expectations exist. For example, if X and Y are independent then

$$\mathbb{E}(XY) = \mathbb{E}X\mathbb{E}Y.$$

Another implication of (B.5) is that if X and Y are independent and $Z = X + Y$, then

$$\begin{aligned} M_Z(\lambda) &= \mathbb{E} \exp(\lambda(X + Y)) \\ &= \mathbb{E}(\exp(\lambda X) \exp(\lambda Y)) \\ &= \mathbb{E} \exp(\lambda' X) \mathbb{E} \exp(\lambda' Y) \\ &= M_X(\lambda) M_Y(\lambda). \end{aligned} \tag{B.6}$$

The covariance between X and Y is

$$\text{cov}(X, Y) = \sigma_{XY} = \mathbb{E}((X - \mathbb{E}X)(Y - \mathbb{E}Y)) = \mathbb{E}XY - \mathbb{E}X\mathbb{E}Y.$$

The correlation between X and Y is

$$\text{corr}(X, Y) = \rho_{XY} = \frac{\sigma_{XY}}{\sigma_X \sigma_Y}.$$

The Cauchy-Schwarz Inequality implies that

$$|\rho_{XY}| \leq 1. \quad (\text{B.7})$$

The correlation is a measure of linear dependence, free of units of measurement.

If X and Y are independent, then $\sigma_{XY} = 0$ and $\rho_{XY} = 0$. The reverse, however, is not true. For example, if $\mathbb{E}X = 0$ and $\mathbb{E}X^3 = 0$, then $\text{cov}(X, X^2) = 0$.

A useful fact is that

$$\text{var}(X + Y) = \text{var}(X) + \text{var}(Y) + 2\text{cov}(X, Y).$$

An implication is that if X and Y are independent, then

$$\text{var}(X + Y) = \text{var}(X) + \text{var}(Y),$$

the variance of the sum is the sum of the variances.

A $k \times 1$ random vector $\mathbf{X} = (X_1, \dots, X_k)'$ is a function from S to \mathbb{R}^k . Let $\mathbf{x} = (x_1, \dots, x_k)'$ denote a vector in \mathbb{R}^k . (In this Appendix, we use bold to denote vectors. Bold capitals \mathbf{X} are random vectors and bold lower case \mathbf{x} are nonrandom vectors. Again, this is in distinction to the notation used in the bulk of the text) The vector \mathbf{X} has the distribution and density functions

$$F(\mathbf{x}) = \Pr(\mathbf{X} \leq \mathbf{x})$$

$$f(\mathbf{x}) = \frac{\partial^k}{\partial x_1 \cdots \partial x_k} F(\mathbf{x}).$$

For a measurable function $\mathbf{g}: \mathbb{R}^k \rightarrow \mathbb{R}^s$, we define the expectation

$$\mathbb{E}\mathbf{g}(\mathbf{X}) = \int_{\mathbb{R}^k} \mathbf{g}(\mathbf{x})f(\mathbf{x})d\mathbf{x}$$

where the symbol $d\mathbf{x}$ denotes $dx_1 \cdots dx_k$. In particular, we have the $k \times 1$ multivariate mean

$$\boldsymbol{\mu} = \mathbb{E}\mathbf{X}$$

and $k \times k$ covariance matrix

$$\begin{aligned} \boldsymbol{\Sigma} &= \mathbb{E}((\mathbf{X} - \boldsymbol{\mu})(\mathbf{X} - \boldsymbol{\mu})') \\ &= \mathbb{E}\mathbf{X}\mathbf{X}' - \boldsymbol{\mu}\boldsymbol{\mu}' \end{aligned}$$

If the elements of \mathbf{X} are mutually independent, then $\boldsymbol{\Sigma}$ is a diagonal matrix and

$$\text{var}\left(\sum_{i=1}^k \mathbf{X}_i\right) = \sum_{i=1}^k \text{var}(\mathbf{X}_i)$$

B.7 Conditional Distributions and Expectation

The **conditional density** of Y given $\mathbf{X} = \mathbf{x}$ is defined as

$$f_{Y|\mathbf{X}}(y | \mathbf{x}) = \frac{f(\mathbf{x}, y)}{f_{\mathbf{X}}(\mathbf{x})}$$

if $f_{\mathbf{X}}(\mathbf{x}) > 0$. One way to derive this expression from the definition of conditional probability is

$$\begin{aligned}
 f_{Y|X}(y | \mathbf{x}) &= \frac{\partial}{\partial y} \lim_{\varepsilon \rightarrow 0} \Pr(Y \leq y | \mathbf{x} \leq \mathbf{X} \leq \mathbf{x} + \varepsilon) \\
 &= \frac{\partial}{\partial y} \lim_{\varepsilon \rightarrow 0} \frac{\Pr(\{Y \leq y\} \cap \{\mathbf{x} \leq \mathbf{X} \leq \mathbf{x} + \varepsilon\})}{\Pr(\mathbf{x} \leq \mathbf{X} \leq \mathbf{x} + \varepsilon)} \\
 &= \frac{\partial}{\partial y} \lim_{\varepsilon \rightarrow 0} \frac{F(\mathbf{x} + \varepsilon, y) - F(\mathbf{x}, y)}{F_{\mathbf{X}}(\mathbf{x} + \varepsilon) - F_{\mathbf{X}}(\mathbf{x})} \\
 &= \frac{\partial}{\partial y} \lim_{\varepsilon \rightarrow 0} \frac{\frac{\partial}{\partial x} F(\mathbf{x} + \varepsilon, y)}{f_{\mathbf{X}}(\mathbf{x} + \varepsilon)} \\
 &= \frac{\frac{\partial^2}{\partial x \partial y} F(\mathbf{x}, y)}{f_{\mathbf{X}}(\mathbf{x})} \\
 &= \frac{f(\mathbf{x}, y)}{f_{\mathbf{X}}(\mathbf{x})}.
 \end{aligned}$$

The **conditional mean** or **conditional expectation** is the function

$$m(\mathbf{x}) = \mathbb{E}(Y | \mathbf{X} = \mathbf{x}) = \int_{-\infty}^{\infty} y f_{Y|X}(y | \mathbf{x}) dy.$$

The conditional mean $m(\mathbf{x})$ is a function, meaning that when \mathbf{X} equals \mathbf{x} , then the expected value of Y is $m(\mathbf{x})$.

Similarly, we define the conditional variance of Y given $\mathbf{X} = \mathbf{x}$ as

$$\begin{aligned}
 \sigma^2(\mathbf{x}) &= \text{var}(Y | \mathbf{X} = \mathbf{x}) \\
 &= \mathbb{E}\left((Y - m(\mathbf{x}))^2 | \mathbf{X} = \mathbf{x}\right) \\
 &= \mathbb{E}(Y^2 | X = \mathbf{x}) - m(\mathbf{x})^2.
 \end{aligned}$$

Evaluated at $\mathbf{x} = \mathbf{X}$, the conditional mean $m(\mathbf{X})$ and conditional variance $\sigma^2(\mathbf{X})$ are random variables, functions of \mathbf{X} . We write this as $\mathbb{E}(Y | \mathbf{X}) = m(\mathbf{X})$ and $\text{var}(Y | \mathbf{X}) = \sigma^2(\mathbf{X})$. For example, if $\mathbb{E}(Y | \mathbf{X} = \mathbf{x}) = \alpha + \beta' \mathbf{x}$, then $\mathbb{E}(Y | \mathbf{X}) = \alpha + \beta' \mathbf{X}$, a transformation of \mathbf{X} .

The following are important facts about conditional expectations.

Simple Law of Iterated Expectations:

$$\mathbb{E}(\mathbb{E}(Y | \mathbf{X})) = \mathbb{E}(Y) \tag{B.8}$$

Proof:

$$\begin{aligned}
 \mathbb{E}(\mathbb{E}(Y | \mathbf{X})) &= \mathbb{E}(m(\mathbf{X})) \\
 &= \int_{-\infty}^{\infty} m(\mathbf{x}) f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} \\
 &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} y f_{Y|X}(y | \mathbf{x}) f_{\mathbf{X}}(\mathbf{x}) dy d\mathbf{x} \\
 &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} y f(y, \mathbf{x}) dy d\mathbf{x} \\
 &= \mathbb{E}(Y).
 \end{aligned}$$

Law of Iterated Expectations:

$$\mathbb{E}(\mathbb{E}(Y | \mathbf{X}, \mathbf{Z}) | \mathbf{X}) = \mathbb{E}(Y | \mathbf{X}) \tag{B.9}$$

Conditioning Theorem. For any function $g(\mathbf{x})$,

$$\mathbb{E}(g(\mathbf{X})Y \mid \mathbf{X}) = g(\mathbf{X}) \mathbb{E}(Y \mid \mathbf{X}) \quad (\text{B.10})$$

Proof: Let

$$\begin{aligned} h(\mathbf{x}) &= \mathbb{E}(g(\mathbf{X})Y \mid \mathbf{X} = \mathbf{x}) \\ &= \int_{-\infty}^{\infty} g(\mathbf{x})y f_{Y|\mathbf{X}}(y \mid \mathbf{x}) dy \\ &= g(\mathbf{x}) \int_{-\infty}^{\infty} y f_{Y|\mathbf{X}}(y \mid \mathbf{x}) dy \\ &= g(\mathbf{x})m(\mathbf{x}) \end{aligned}$$

where $m(\mathbf{x}) = \mathbb{E}(Y \mid \mathbf{X} = \mathbf{x})$. Thus $h(\mathbf{X}) = g(\mathbf{X})m(\mathbf{X})$, which is the same as $\mathbb{E}(g(\mathbf{X})Y \mid \mathbf{X}) = g(\mathbf{X}) \mathbb{E}(Y \mid \mathbf{X})$.

B.8 Transformations

Suppose that $\mathbf{X} \in \mathbb{R}^k$ with continuous distribution function $F_{\mathbf{X}}(\mathbf{x})$ and density $f_{\mathbf{X}}(\mathbf{x})$. Let $\mathbf{Y} = \mathbf{g}(\mathbf{X})$ where $\mathbf{g}(\mathbf{x}) : \mathbb{R}^k \rightarrow \mathbb{R}^k$ is one-to-one, differentiable, and invertible. Let $\mathbf{h}(\mathbf{y})$ denote the inverse of $\mathbf{g}(\mathbf{x})$. The **Jacobian** is

$$J(\mathbf{y}) = \det \left(\frac{\partial}{\partial \mathbf{y}'} \mathbf{h}(\mathbf{y}) \right).$$

Consider the univariate case $k = 1$. If $g(x)$ is an increasing function, then $g(X) \leq Y$ if and only if $X \leq h(Y)$, so the distribution function of Y is

$$\begin{aligned} F_Y(y) &= \Pr(g(X) \leq y) \\ &= \Pr(X \leq h(Y)) \\ &= F_X(h(Y)). \end{aligned}$$

Taking the derivative, the density of Y is

$$f_Y(y) = \frac{d}{dy} F_Y(y) = f_X(h(Y)) \frac{d}{dy} h(y).$$

If $g(x)$ is a decreasing function, then $g(X) \leq Y$ if and only if $X \geq h(Y)$, so

$$\begin{aligned} F_Y(y) &= \Pr(g(X) \leq y) \\ &= 1 - \Pr(X \geq h(Y)) \\ &= 1 - F_X(h(Y)) \end{aligned}$$

and the density of Y is

$$f_Y(y) = -f_X(h(Y)) \frac{d}{dy} h(y).$$

We can write these two cases jointly as

$$f_Y(y) = f_X(h(Y)) |J(y)|. \quad (\text{B.11})$$

This is known as the **change-of-variables** formula. This same formula (B.11) holds for $k > 1$, but its justification requires deeper results from analysis.

As one example, take the case $X \sim U[0, 1]$ and $Y = -\log(X)$. Here, $g(x) = -\log(x)$ and $h(y) = \exp(-y)$ so the Jacobian is $J(y) = -\exp(y)$. As the range of X is $[0, 1]$, that for Y is $[0, \infty)$. Since $f_X(x) = 1$ for $0 \leq x \leq 1$ (B.11) shows that

$$f_Y(y) = \exp(-y), \quad 0 \leq y \leq \infty,$$

an exponential density.

B.9 Normal and Related Distributions

The **standard normal** density is

$$\phi(x) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right), \quad -\infty < x < \infty.$$

It is conventional to write $X \sim N(0, 1)$, and to denote the standard normal density function by $\phi(x)$ and its distribution function by $\Phi(x)$. The latter has no closed-form solution. The normal density has all moments finite. Since it is symmetric about zero all odd moments are zero. By iterated integration by parts, we can also show that $\mathbb{E}X^2 = 1$ and $\mathbb{E}X^4 = 3$. In fact, for any positive integer m , $\mathbb{E}X^{2m} = (2m - 1)!! = (2m - 1) \cdot (2m - 3) \cdots 1$. Thus $\mathbb{E}X^4 = 3$, $\mathbb{E}X^6 = 15$, $\mathbb{E}X^8 = 105$, and $\mathbb{E}X^{10} = 945$.

If Z is standard normal and $X = \mu + \sigma Z$, then using the change-of-variables formula, X has density

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(x - \mu)^2}{2\sigma^2}\right), \quad -\infty < x < \infty.$$

which is the **univariate normal density**. The mean and variance of the distribution are μ and σ^2 , and it is conventional to write $X \sim N(\mu, \sigma^2)$.

For $\mathbf{x} \in \mathbb{R}^k$, the **multivariate normal density** is

$$f(\mathbf{x}) = \frac{1}{(2\pi)^{k/2} \det(\boldsymbol{\Sigma})^{1/2}} \exp\left(-\frac{(\mathbf{x} - \boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu})}{2}\right), \quad \mathbf{x} \in \mathbb{R}^k.$$

The mean and covariance matrix of the distribution are $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, and it is conventional to write $X \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$.

The MGF and CF of the multivariate normal are $\exp(\boldsymbol{\lambda}'\boldsymbol{\mu} + \boldsymbol{\lambda}'\boldsymbol{\Sigma}\boldsymbol{\lambda}/2)$ and $\exp(i\boldsymbol{\lambda}'\boldsymbol{\mu} - \boldsymbol{\lambda}'\boldsymbol{\Sigma}\boldsymbol{\lambda}/2)$, respectively.

If $\mathbf{X} \in \mathbb{R}^k$ is multivariate normal and the elements of \mathbf{X} are mutually uncorrelated, then $\boldsymbol{\Sigma} = \text{diag}\{\sigma_j^2\}$ is a diagonal matrix. In this case the density function can be written as

$$\begin{aligned} f(\mathbf{x}) &= \frac{1}{(2\pi)^{k/2} \sigma_1 \cdots \sigma_k} \exp\left(-\left(\frac{(x_1 - \mu_1)^2/\sigma_1^2 + \cdots + (x_k - \mu_k)^2/\sigma_k^2}{2}\right)\right) \\ &= \prod_{j=1}^k \frac{1}{(2\pi)^{1/2} \sigma_j} \exp\left(-\frac{(x_j - \mu_j)^2}{2\sigma_j^2}\right) \end{aligned}$$

which is the product of marginal univariate normal densities. This shows that if \mathbf{X} is multivariate normal with uncorrelated elements, then they are mutually independent.

Theorem B.9.1 *If $\mathbf{X} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and $\mathbf{Y} = \mathbf{a} + \mathbf{B}\mathbf{X}$ with \mathbf{B} an invertible matrix, then $\mathbf{Y} \sim N(\mathbf{a} + \mathbf{B}\boldsymbol{\mu}, \mathbf{B}\boldsymbol{\Sigma}\mathbf{B}')$.*

Theorem B.9.2 *Let $\mathbf{X} \sim N(\mathbf{0}, \mathbf{I}_r)$. Then $Q = \mathbf{X}'\mathbf{X}$ is distributed chi-square with r degrees of freedom, written χ_r^2 .*

Theorem B.9.3 *If $\mathbf{Z} \sim N(\mathbf{0}, \mathbf{A})$ with $\mathbf{A} > 0$, $q \times q$, then $\mathbf{Z}'\mathbf{A}^{-1}\mathbf{Z} \sim \chi_q^2$.*

Theorem B.9.4 *Let $Z \sim N(0, 1)$ and $Q \sim \chi_r^2$ be independent. Then $T_r = Z/\sqrt{Q/r}$ is distributed as student's t with r degrees of freedom.*

Proof of Theorem B.9.1. By the change-of-variables formula, the density of $\mathbf{Y} = \mathbf{a} + \mathbf{B}\mathbf{X}$ is

$$f(\mathbf{y}) = \frac{1}{(2\pi)^{k/2} \det(\boldsymbol{\Sigma}_Y)^{1/2}} \exp\left(-\frac{(\mathbf{y} - \boldsymbol{\mu}_Y)' \boldsymbol{\Sigma}_Y^{-1} (\mathbf{y} - \boldsymbol{\mu}_Y)}{2}\right), \quad \mathbf{y} \in \mathbb{R}^k.$$

where $\boldsymbol{\mu}_Y = \mathbf{a} + \mathbf{B}\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}_Y = \mathbf{B}\boldsymbol{\Sigma}\mathbf{B}'$, where we used the fact that $\det(\mathbf{B}\boldsymbol{\Sigma}\mathbf{B}')^{1/2} = \det(\boldsymbol{\Sigma})^{1/2} \det(\mathbf{B})$. ■

Proof of Theorem B.9.2. First, suppose a random variable Q is distributed chi-square with r degrees of freedom. It has the MGF

$$\mathbb{E} \exp(tQ) = \int_0^\infty \frac{1}{\Gamma(\frac{r}{2}) 2^{r/2}} x^{r/2-1} \exp(tx) \exp(-x/2) dy = (1-2t)^{-r/2}$$

where the second equality uses the fact that $\int_0^\infty y^{a-1} \exp(-by) dy = b^{-a} \Gamma(a)$, which can be found by applying change-of-variables to the gamma function. Our goal is to calculate the MGF of $Q = \mathbf{X}'\mathbf{X}$ and show that it equals $(1-2t)^{-r/2}$, which will establish that $Q \sim \chi_r^2$.

Note that we can write $Q = \mathbf{X}'\mathbf{X} = \sum_{j=1}^r Z_j^2$ where the Z_j are independent $N(0, 1)$. The distribution of each of the Z_j^2 is

$$\begin{aligned} \Pr(Z_j^2 \leq y) &= 2 \Pr(0 \leq Z_j \leq \sqrt{y}) \\ &= 2 \int_0^{\sqrt{y}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{x^2}{2}\right) dx \\ &= \int_0^y \frac{1}{\Gamma(\frac{1}{2}) 2^{1/2}} s^{-1/2} \exp\left(-\frac{s}{2}\right) ds \end{aligned}$$

using the change-of-variables $s = x^2$ and the fact $\Gamma(\frac{1}{2}) = \sqrt{\pi}$. Thus the density of Z_j^2 is

$$f_1(x) = \frac{1}{\Gamma(\frac{1}{2}) 2^{1/2}} x^{-1/2} \exp\left(-\frac{x}{2}\right)$$

which is the χ_1^2 and by our above calculation has the MGF of $\mathbb{E} \exp(tZ_j^2) = (1-2t)^{-1/2}$.

Since the Z_j^2 are mutually independent, (B.6) implies that the MGF of $Q = \sum_{j=1}^r Z_j^2$ is $[(1-2t)^{-1/2}]^r = (1-2t)^{-r/2}$, which is the MGF of the χ_r^2 density as desired. ■

Proof of Theorem B.9.3. The fact that $\mathbf{A} > 0$ means that we can write $\mathbf{A} = \mathbf{C}\mathbf{C}'$ where \mathbf{C} is non-singular. Then $\mathbf{A}^{-1} = \mathbf{C}^{-1'}\mathbf{C}^{-1}$ and

$$\mathbf{C}^{-1}\mathbf{Z} \sim N(\mathbf{0}, \mathbf{C}^{-1}\mathbf{A}\mathbf{C}^{-1'}) = N(\mathbf{0}, \mathbf{C}^{-1}\mathbf{C}\mathbf{C}'\mathbf{C}^{-1'}) = N(\mathbf{0}, \mathbf{I}_q).$$

Thus

$$\mathbf{Z}'\mathbf{A}^{-1}\mathbf{Z} = \mathbf{Z}'\mathbf{C}^{-1'}\mathbf{C}^{-1}\mathbf{Z} = (\mathbf{C}^{-1}\mathbf{Z})'(\mathbf{C}^{-1}\mathbf{Z}) \sim \chi_q^2.$$

■

Proof of Theorem B.9.4. Using the simple law of iterated expectations, T_r has distribution

function

$$\begin{aligned}
F(x) &= \Pr\left(\frac{Z}{\sqrt{Q/r}} \leq x\right) \\
&= \mathbb{E}\left\{Z \leq x\sqrt{\frac{Q}{r}}\right\} \\
&= \mathbb{E}\left[\Pr\left(Z \leq x\sqrt{\frac{Q}{r}} \mid Q\right)\right] \\
&= \mathbb{E}\Phi\left(x\sqrt{\frac{Q}{r}}\right)
\end{aligned}$$

Thus its density is

$$\begin{aligned}
f(x) &= \mathbb{E}\frac{d}{dx}\Phi\left(x\sqrt{\frac{Q}{r}}\right) \\
&= \mathbb{E}\left(\phi\left(x\sqrt{\frac{Q}{r}}\right)\sqrt{\frac{Q}{r}}\right) \\
&= \int_0^\infty \left(\frac{1}{\sqrt{2\pi}}\exp\left(-\frac{qx^2}{2r}\right)\right)\sqrt{\frac{q}{r}}\left(\frac{1}{\Gamma\left(\frac{r}{2}\right)2^{r/2}}q^{r/2-1}\exp(-q/2)\right)dq \\
&= \frac{\Gamma\left(\frac{r+1}{2}\right)}{\sqrt{r\pi}\Gamma\left(\frac{r}{2}\right)}\left(1+\frac{x^2}{r}\right)^{-\left(\frac{r+1}{2}\right)}
\end{aligned}$$

which is that of the student t with r degrees of freedom. ■

B.10 Inequalities

Jensen's Inequality. If $g(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}$ is convex, then for any random vector \mathbf{x} for which $\mathbb{E}\|\mathbf{x}\| < \infty$ and $\mathbb{E}|g(\mathbf{x})| < \infty$,

$$g(\mathbb{E}(\mathbf{x})) \leq \mathbb{E}(g(\mathbf{x})). \quad (\text{B.12})$$

Conditional Jensen's Inequality. If $g(\cdot) : \mathbb{R}^m \rightarrow \mathbb{R}$ is convex, then for any random vectors (\mathbf{y}, \mathbf{x}) for which $\mathbb{E}\|\mathbf{y}\| < \infty$ and $\mathbb{E}\|g(\mathbf{y})\| < \infty$,

$$g(\mathbb{E}(\mathbf{y} \mid \mathbf{x})) \leq \mathbb{E}(g(\mathbf{y}) \mid \mathbf{x}). \quad (\text{B.13})$$

Conditional Expectation Inequality. For any $r \geq 1$ such that $\mathbb{E}|y|^r < \infty$, then

$$\mathbb{E}|\mathbb{E}(y \mid \mathbf{x})|^r \leq \mathbb{E}|y|^r < \infty. \quad (\text{B.14})$$

Expectation Inequality. For any random matrix \mathbf{Y} for which $\mathbb{E}\|\mathbf{Y}\| < \infty$,

$$\|\mathbb{E}(\mathbf{Y})\| \leq \mathbb{E}\|\mathbf{Y}\|. \quad (\text{B.15})$$

Hölder's Inequality. If $p > 1$ and $q > 1$ and $\frac{1}{p} + \frac{1}{q} = 1$, then for any random $m \times n$ matrices \mathbf{X} and \mathbf{Y} ,

$$\mathbb{E}\|\mathbf{X}'\mathbf{Y}\| \leq (\mathbb{E}\|\mathbf{X}\|^p)^{1/p} (\mathbb{E}\|\mathbf{Y}\|^q)^{1/q}. \quad (\text{B.16})$$

Cauchy-Schwarz Inequality. For any random $m \times n$ matrices \mathbf{X} and \mathbf{Y} ,

$$\mathbb{E} \|\mathbf{X}'\mathbf{Y}\| \leq \left(\mathbb{E} \|\mathbf{X}\|^2\right)^{1/2} \left(\mathbb{E} \|\mathbf{Y}\|^2\right)^{1/2}. \quad (\text{B.17})$$

Matrix Cauchy-Schwarz Inequality. Tripathi (1999). For any random $\mathbf{x} \in \mathbb{R}^m$ and $\mathbf{y} \in \mathbb{R}^\ell$,

$$\mathbb{E} \mathbf{y} \mathbf{x}' (\mathbb{E} \mathbf{x} \mathbf{x}')^{-1} \mathbb{E} \mathbf{x} \mathbf{y}' \leq \mathbb{E} \mathbf{y} \mathbf{y}' \quad (\text{B.18})$$

Minkowski's Inequality. For any random $m \times n$ matrices \mathbf{X} and \mathbf{Y} ,

$$(\mathbb{E} \|\mathbf{X} + \mathbf{Y}\|^p)^{1/p} \leq (\mathbb{E} \|\mathbf{X}\|^p)^{1/p} + (\mathbb{E} \|\mathbf{Y}\|^p)^{1/p} \quad (\text{B.19})$$

Liapunov's Inequality. For any random $m \times n$ matrix \mathbf{X} and $1 \leq r \leq p$,

$$(\mathbb{E} \|\mathbf{X}\|^r)^{1/r} \leq (\mathbb{E} \|\mathbf{X}\|^p)^{1/p} \quad (\text{B.20})$$

Markov's Inequality (standard form). For any random vector \mathbf{x} and non-negative function $g(\mathbf{x}) \geq 0$,

$$\Pr(g(\mathbf{x}) > \alpha) \leq \alpha^{-1} \mathbb{E} g(\mathbf{x}). \quad (\text{B.21})$$

Markov's Inequality (strong form). For any random vector \mathbf{x} and non-negative function $g(\mathbf{x}) \geq 0$,

$$\Pr(g(\mathbf{x}) > \alpha) \leq \alpha^{-1} \mathbb{E} (g(\mathbf{x}) \mathbf{1}(g(\mathbf{x}) > \alpha)). \quad (\text{B.22})$$

Chebyshev's Inequality. For any random variable x ,

$$\Pr(|x - \mathbb{E}x| > \alpha) \leq \frac{\text{var}(x)}{\alpha^2}. \quad (\text{B.23})$$

Proof of Jensen's Inequality (B.12). Since $g(\mathbf{u})$ is convex, at any point \mathbf{u} there is a nonempty set of subderivatives (linear surfaces touching $g(\mathbf{u})$ at \mathbf{u} but lying below $g(\mathbf{u})$ for all \mathbf{u}). Let $a + \mathbf{b}'\mathbf{u}$ be a subderivative of $g(\mathbf{u})$ at $\mathbf{u} = \mathbb{E}\mathbf{x}$. Then for all \mathbf{u} , $g(\mathbf{u}) \geq a + \mathbf{b}'\mathbf{u}$ yet $g(\mathbb{E}\mathbf{x}) = a + \mathbf{b}'\mathbb{E}\mathbf{x}$. Applying expectations, $\mathbb{E}g(\mathbf{x}) \geq a + \mathbf{b}'\mathbb{E}\mathbf{x} = g(\mathbb{E}\mathbf{x})$, as stated. ■

Proof of Conditional Jensen's Inequality. The same as the proof of (B.12), but using conditional expectations. The conditional expectations exist since $\mathbb{E} \|\mathbf{y}\| < \infty$ and $\mathbb{E} \|g(\mathbf{y})\| < \infty$. ■

Proof of Conditional Expectation Inequality. As the function $|u|^r$ is convex for $r \geq 1$, the Conditional Jensen's inequality implies

$$|\mathbb{E}(y \mid \mathbf{x})|^r \leq \mathbb{E}(|y|^r \mid \mathbf{x}).$$

Taking unconditional expectations and the law of iterated expectations, we obtain

$$\mathbb{E} |\mathbb{E}(y \mid \mathbf{x})|^r \leq \mathbb{E} \mathbb{E}(|y|^r \mid \mathbf{x}) = \mathbb{E} |y|^r < \infty$$

as required. ■

Proof of Expectation Inequality. By the Triangle inequality, for $\lambda \in [0, 1]$,

$$\|\lambda \mathbf{U}_1 + (1 - \lambda) \mathbf{U}_2\| \leq \lambda \|\mathbf{U}_1\| + (1 - \lambda) \|\mathbf{U}_2\|$$

which shows that the matrix norm $g(\mathbf{U}) = \|\mathbf{U}\|$ is convex. Applying Jensen's Inequality (B.12) we find (B.15). ■

Proof of Hölder's Inequality. Since $\frac{1}{p} + \frac{1}{q} = 1$ an application of Jensen's Inequality (A.17) shows that for any real a and b

$$\exp\left[\frac{1}{p}a + \frac{1}{q}b\right] \leq \frac{1}{p}\exp(a) + \frac{1}{q}\exp(b).$$

Setting $u = \exp(a)$ and $v = \exp(b)$ this implies

$$u^{1/p}v^{1/q} \leq \frac{u}{p} + \frac{v}{q}$$

and this inequality holds for any $u > 0$ and $v > 0$.

Set $u = \|\mathbf{X}\|^p / \mathbb{E}\|\mathbf{X}\|^p$ and $v = \|\mathbf{Y}\|^q / \mathbb{E}\|\mathbf{Y}\|^q$. Note that $\mathbb{E}u = \mathbb{E}v = 1$. By the matrix Schwarz Inequality (A.11), $\|\mathbf{X}'\mathbf{Y}\| \leq \|\mathbf{X}\| \|\mathbf{Y}\|$. Thus

$$\begin{aligned} \frac{\mathbb{E}\|\mathbf{X}'\mathbf{Y}\|}{(\mathbb{E}\|\mathbf{X}\|^p)^{1/p}(\mathbb{E}\|\mathbf{Y}\|^q)^{1/q}} &\leq \frac{\mathbb{E}(\|\mathbf{X}\| \|\mathbf{Y}\|)}{(\mathbb{E}\|\mathbf{X}\|^p)^{1/p}(\mathbb{E}\|\mathbf{Y}\|^q)^{1/q}} \\ &= \mathbb{E}\left(u^{1/p}v^{1/q}\right) \\ &\leq \mathbb{E}\left(\frac{u}{p} + \frac{v}{q}\right) \\ &= \frac{1}{p} + \frac{1}{q} \\ &= 1, \end{aligned}$$

which is (B.16). ■

Proof of Cauchy-Schwarz Inequality. Special case of Hölder's with $p = q = 2$.

Proof of Matrix Cauchy-Schwarz Inequality. Define $e = \mathbf{y} - (\mathbb{E}\mathbf{y}\mathbf{x}')(\mathbb{E}\mathbf{x}\mathbf{x}')^{-1}\mathbf{x}$. Note that $\mathbb{E}ee' \geq 0$ is positive semi-definite. We can calculate that

$$\mathbb{E}ee' = \mathbb{E}\mathbf{y}\mathbf{y}' - (\mathbb{E}\mathbf{y}\mathbf{x}')(\mathbb{E}\mathbf{x}\mathbf{x}')^{-1}\mathbb{E}\mathbf{x}\mathbf{y}'.$$

Since the left-hand-side is positive semi-definite, so is the right-hand-side, which means $\mathbb{E}\mathbf{y}\mathbf{y}' \geq (\mathbb{E}\mathbf{y}\mathbf{x}')(\mathbb{E}\mathbf{x}\mathbf{x}')^{-1}\mathbb{E}\mathbf{x}\mathbf{y}'$ as stated. ■

Proof of Liapunov's Inequality. The function $g(u) = u^{p/r}$ is convex for $u > 0$ since $p \geq r$. Set $u = \|\mathbf{X}\|^r$. By Jensen's inequality, $g(\mathbb{E}u) \leq \mathbb{E}g(u)$ or

$$(\mathbb{E}\|\mathbf{X}\|^r)^{p/r} \leq \mathbb{E}(\|\mathbf{X}\|^r)^{p/r} = \mathbb{E}\|\mathbf{X}\|^p.$$

Raising both sides to the power $1/p$ yields $(\mathbb{E}\|\mathbf{X}\|^r)^{1/r} \leq (\mathbb{E}\|\mathbf{X}\|^p)^{1/p}$ as claimed. ■

Proof of Minkowski's Inequality. Note that by rewriting, using the triangle inequality (A.12), and then Hölder's Inequality to the two expectations

$$\begin{aligned} \mathbb{E}\|\mathbf{X} + \mathbf{Y}\|^p &= \mathbb{E}\left(\|\mathbf{X} + \mathbf{Y}\| \|\mathbf{X} + \mathbf{Y}\|^{p-1}\right) \\ &\leq \mathbb{E}\left(\|\mathbf{X}\| \|\mathbf{X} + \mathbf{Y}\|^{p-1}\right) + \mathbb{E}\left(\|\mathbf{Y}\| \|\mathbf{X} + \mathbf{Y}\|^{p-1}\right) \\ &\leq (\mathbb{E}\|\mathbf{X}\|^p)^{1/p} \mathbb{E}\left(\|\mathbf{X} + \mathbf{Y}\|^{q(p-1)}\right)^{1/q} \\ &\quad + (\mathbb{E}\|\mathbf{Y}\|^p)^{1/p} \mathbb{E}\left(\|\mathbf{X} + \mathbf{Y}\|^{q(p-1)}\right)^{1/q} \\ &= \left((\mathbb{E}\|\mathbf{X}\|^p)^{1/p} + (\mathbb{E}\|\mathbf{Y}\|^p)^{1/p}\right) \mathbb{E}\left(\|\mathbf{X} + \mathbf{Y}\|^p\right)^{(p-1)/p} \end{aligned}$$

where the second equality picks q to satisfy $1/p+1/q = 1$, and the final equality uses this fact to make the substitution $q = p/(p-1)$ and then collects terms. Dividing both sides by $\mathbb{E}(\|\mathbf{X} + \mathbf{Y}\|^p)^{(p-1)/p}$, we obtain (B.19). ■

Proof of Markov's Inequality. Let F denote the distribution function of \mathbf{x} . Then

$$\begin{aligned} \Pr(g(\mathbf{x}) \geq \alpha) &= \int_{\{g(\mathbf{u}) \geq \alpha\}} dF(\mathbf{u}) \\ &\leq \int_{\{g(\mathbf{u}) \geq \alpha\}} \frac{g(\mathbf{u})}{\alpha} dF(\mathbf{u}) \\ &= \alpha^{-1} \int 1(g(\mathbf{u}) > \alpha) g(\mathbf{u}) dF(\mathbf{u}) \\ &= \alpha^{-1} \mathbb{E}(g(\mathbf{x}) 1(g(\mathbf{x}) > \alpha)) \end{aligned}$$

the inequality using the region of integration $\{g(\mathbf{u}) > \alpha\}$. This establishes the strong form (B.22). Since $1(g(\mathbf{x}) > \alpha) \leq 1$, the final expression is less than $\alpha^{-1} \mathbb{E}(g(\mathbf{x}))$, establishing the standard form (B.21). ■

Proof of Chebyshev's Inequality. Define $y = (x - \mathbb{E}x)^2$ and note that $\mathbb{E}y = \text{var}(x)$. The events $\{|x - \mathbb{E}x| > \alpha\}$ and $\{y > \alpha^2\}$ are equal, so by an application Markov's inequality we find

$$\Pr(|x - \mathbb{E}x| > \alpha) = \Pr(y > \alpha^2) \leq \alpha^{-2} \mathbb{E}(y) = \alpha^{-2} \text{var}(x)$$

as stated. ■

B.11 Maximum Likelihood

In this section we provide a brief review of the asymptotic theory of maximum likelihood estimation.

When the density of \mathbf{y}_i is $f(\mathbf{y} | \boldsymbol{\theta})$ where F is a known distribution function and $\boldsymbol{\theta} \in \Theta$ is an unknown $m \times 1$ vector, we say that the distribution is **parametric** and that $\boldsymbol{\theta}$ is the **parameter** of the distribution F . The space Θ is the set of permissible value for $\boldsymbol{\theta}$. In this setting the **method of maximum likelihood** is an appropriate technique for estimation and inference on $\boldsymbol{\theta}$. We let $\boldsymbol{\theta}$ denote a generic value of the parameter and let $\boldsymbol{\theta}_0$ denote its true value.

The joint density of a random sample $(\mathbf{y}_1, \dots, \mathbf{y}_n)$ is

$$f_n(\mathbf{y}_1, \dots, \mathbf{y}_n | \boldsymbol{\theta}) = \prod_{i=1}^n f(\mathbf{y}_i | \boldsymbol{\theta}).$$

The **likelihood** of the sample is this joint density evaluated at the observed sample values, viewed as a function of $\boldsymbol{\theta}$. The **log-likelihood** function is its natural logarithm

$$\log L(\boldsymbol{\theta}) = \sum_{i=1}^n \log f(\mathbf{y}_i | \boldsymbol{\theta}).$$

The **likelihood score** is the derivative of the log-likelihood, evaluated at the true parameter value.

$$\mathbf{S}_i = \frac{\partial}{\partial \boldsymbol{\theta}} \log f(\mathbf{y}_i | \boldsymbol{\theta}_0).$$

We also define the **Hessian**

$$\boldsymbol{\mathcal{H}} = -\mathbb{E} \frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} \log f(\mathbf{y}_i | \boldsymbol{\theta}_0) \tag{B.24}$$

and the **outer product** matrix

$$\boldsymbol{\Omega} = \mathbb{E} (\mathbf{S}_i \mathbf{S}_i'). \quad (\text{B.25})$$

We now present three important features of the likelihood.

Theorem B.11.1

$$\left. \frac{\partial}{\partial \boldsymbol{\theta}} \mathbb{E} \log f(\mathbf{y} | \boldsymbol{\theta}) \right|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} = \mathbf{0} \quad (\text{B.26})$$

$$\mathbb{E} \mathbf{S}_i = \mathbf{0} \quad (\text{B.27})$$

and

$$\boldsymbol{\mathcal{H}} = \boldsymbol{\Omega} \equiv \boldsymbol{\mathcal{I}} \quad (\text{B.28})$$

The matrix $\boldsymbol{\mathcal{I}}$ is called the **information**, and the equality (B.28) is called the **information matrix equality**.

The **maximum likelihood estimator (MLE)** $\hat{\boldsymbol{\theta}}$ is the parameter value which maximizes the likelihood (equivalently, which maximizes the log-likelihood). We can write this as

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta} \in \Theta}{\operatorname{argmax}} \log L(\boldsymbol{\theta}). \quad (\text{B.29})$$

In some simple cases, we can find an explicit expression for $\hat{\boldsymbol{\theta}}$ as a function of the data, but these cases are rare. More typically, the MLE $\hat{\boldsymbol{\theta}}$ must be found by numerical methods.

To understand why the MLE $\hat{\boldsymbol{\theta}}$ is a natural estimator for the parameter $\boldsymbol{\theta}$ observe that the standardized log-likelihood is a sample average and an estimator of $\mathbb{E} \log f(\mathbf{y}_i | \boldsymbol{\theta})$:

$$\frac{1}{n} \log L(\boldsymbol{\theta}) = \frac{1}{n} \sum_{i=1}^n \log f(\mathbf{y}_i | \boldsymbol{\theta}) \xrightarrow{p} \mathbb{E} \log f(\mathbf{y}_i | \boldsymbol{\theta}).$$

As the MLE $\hat{\boldsymbol{\theta}}$ maximizes the left-hand-side, we can see that it is an estimator of the maximizer of the right-hand-side. The first-order condition for the latter problem is

$$0 = \frac{\partial}{\partial \boldsymbol{\theta}} \mathbb{E} \log f(\mathbf{y}_i | \boldsymbol{\theta})$$

which holds at $\boldsymbol{\theta} = \boldsymbol{\theta}_0$ by (B.26). This suggests that $\hat{\boldsymbol{\theta}}$ is an estimator of $\boldsymbol{\theta}_0$. In fact, under conventional regularity conditions, $\hat{\boldsymbol{\theta}}$ is consistent, $\hat{\boldsymbol{\theta}} \xrightarrow{p} \boldsymbol{\theta}_0$ as $n \rightarrow \infty$. Furthermore, we can derive its asymptotic distribution.

Theorem B.11.2 Under regularity conditions, $\sqrt{n}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \xrightarrow{d} \mathbf{N}(\mathbf{0}, \boldsymbol{\mathcal{I}}^{-1})$.

We omit the regularity conditions for Theorem B.11.2, but the result holds quite broadly for models which are smooth functions of the parameters. Theorem B.11.2 gives the general form for the asymptotic distribution of the MLE. A famous result shows that the asymptotic variance is the smallest possible.

Theorem B.11.3 Cramer-Rao Lower Bound. *If $\tilde{\boldsymbol{\theta}}$ is an unbiased regular estimator of $\boldsymbol{\theta}$, then $\text{var}(\tilde{\boldsymbol{\theta}}) \geq (n\boldsymbol{\mathcal{I}})^{-}$.*

The Cramer-Rao Theorem shows that the finite sample variance of an unbiased estimator is bounded below by $(n\boldsymbol{\mathcal{I}})^{-1}$. This means that the asymptotic variance of the standardized estimator $\sqrt{n}(\tilde{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)$ is bounded below by $\boldsymbol{\mathcal{I}}^{-1}$. In other words, the best possible asymptotic variance among all (regular) estimators is $\boldsymbol{\mathcal{I}}^{-1}$. An estimator is called asymptotically efficient if its asymptotic variance equals this lower bound. Theorem B.11.2 shows that the MLE has this asymptotic variance, and is thus asymptotically efficient.

Theorem B.11.4 *The MLE is asymptotically efficient in the sense that its asymptotic variance equals the Cramer-Rao Lower Bound.*

Theorem B.11.4 gives a strong endorsement for the MLE in parametric models.

Finally, consider functions of parameters. If $\boldsymbol{\psi} = \mathbf{g}(\boldsymbol{\theta})$ then the MLE of $\boldsymbol{\psi}$ is $\hat{\boldsymbol{\psi}} = \mathbf{g}(\hat{\boldsymbol{\theta}})$. This is because maximization (e.g. (B.29)) is unaffected by parameterization and transformation. Applying the Delta Method to Theorem B.11.2 we conclude that

$$\sqrt{n}(\hat{\boldsymbol{\psi}} - \boldsymbol{\psi}) \simeq \mathbf{G}'\sqrt{n}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \xrightarrow{d} \text{N}(\mathbf{0}, \mathbf{G}'\boldsymbol{\mathcal{I}}^{-1}\mathbf{G}) \quad (\text{B.30})$$

where $\mathbf{G} = \frac{\partial}{\partial \boldsymbol{\theta}} \mathbf{g}(\boldsymbol{\theta}_0)'$. By Theorem B.11.4, $\hat{\boldsymbol{\psi}}$ is an asymptotically efficient estimator for $\boldsymbol{\psi}$ since it is the MLE. The asymptotic variance $\mathbf{G}'\boldsymbol{\mathcal{I}}^{-1}\mathbf{G}$ is the Cramer-Rao lower bound for estimation of $\boldsymbol{\psi}$.

Theorem B.11.5 *The Cramer-Rao lower bound for $\boldsymbol{\psi} = \mathbf{g}(\boldsymbol{\theta})$ is $\mathbf{G}'\boldsymbol{\mathcal{I}}^{-1}\mathbf{G}$, and the MLE $\hat{\boldsymbol{\psi}} = \mathbf{g}(\hat{\boldsymbol{\theta}})$ is asymptotically efficient.*

Proof of Theorem B.11.1. To see (B.26),

$$\begin{aligned} \left. \frac{\partial}{\partial \boldsymbol{\theta}} \mathbb{E} \log f(\mathbf{y} | \boldsymbol{\theta}) \right|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} &= \left. \frac{\partial}{\partial \boldsymbol{\theta}} \int \log f(\mathbf{y} | \boldsymbol{\theta}) f(\mathbf{y} | \boldsymbol{\theta}_0) d\mathbf{y} \right|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} \\ &= \left. \int \frac{\partial}{\partial \boldsymbol{\theta}} f(\mathbf{y} | \boldsymbol{\theta}) \frac{f(\mathbf{y} | \boldsymbol{\theta}_0)}{f(\mathbf{y} | \boldsymbol{\theta})} d\mathbf{y} \right|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} \\ &= \left. \frac{\partial}{\partial \boldsymbol{\theta}} \int f(\mathbf{y} | \boldsymbol{\theta}) d\mathbf{y} \right|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} \\ &= \left. \frac{\partial}{\partial \boldsymbol{\theta}} 1 \right|_{\boldsymbol{\theta}=\boldsymbol{\theta}_0} = \mathbf{0}. \end{aligned}$$

Equation (B.27) follows by exchanging integration and differentiation

$$\mathbb{E} \frac{\partial}{\partial \boldsymbol{\theta}} \log f(\mathbf{y} | \boldsymbol{\theta}_0) = \frac{\partial}{\partial \boldsymbol{\theta}} \mathbb{E} \log f(\mathbf{y} | \boldsymbol{\theta}_0) = \mathbf{0}.$$

Similarly, we can show that

$$\mathbb{E} \left(\frac{\frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} f(\mathbf{y} | \boldsymbol{\theta}_0)}{f(\mathbf{y} | \boldsymbol{\theta}_0)} \right) = \mathbf{0}.$$

By direction computation,

$$\begin{aligned} \frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} \log f(\mathbf{y} | \boldsymbol{\theta}_0) &= \frac{\frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} f(\mathbf{y} | \boldsymbol{\theta}_0)}{f(\mathbf{y} | \boldsymbol{\theta}_0)} - \frac{\frac{\partial}{\partial \boldsymbol{\theta}} f(\mathbf{y} | \boldsymbol{\theta}_0) \frac{\partial}{\partial \boldsymbol{\theta}} f(\mathbf{y} | \boldsymbol{\theta}_0)'}{f(\mathbf{y} | \boldsymbol{\theta}_0)^2} \\ &= \frac{\frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} f(\mathbf{y} | \boldsymbol{\theta}_0)}{f(\mathbf{y} | \boldsymbol{\theta}_0)} - \frac{\partial}{\partial \boldsymbol{\theta}} \log f(\mathbf{y} | \boldsymbol{\theta}_0) \frac{\partial}{\partial \boldsymbol{\theta}} \log f(\mathbf{y} | \boldsymbol{\theta}_0)'. \end{aligned}$$

Taking expectations yields (B.28). \blacksquare

Proof of Theorem B.11.2 Taking the first-order condition for maximization of $\log L(\boldsymbol{\theta})$, and making a first-order Taylor series expansion,

$$\begin{aligned} 0 &= \left. \frac{\partial}{\partial \boldsymbol{\theta}} \log L(\boldsymbol{\theta}) \right|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}} \\ &= \sum_{i=1}^n \frac{\partial}{\partial \boldsymbol{\theta}} \log f(\mathbf{y}_i | \hat{\boldsymbol{\theta}}) \\ &= \sum_{i=1}^n \frac{\partial}{\partial \boldsymbol{\theta}} \log f(\mathbf{y}_i | \boldsymbol{\theta}_0) + \sum_{i=1}^n \frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} \log f(\mathbf{y}_i | \boldsymbol{\theta}_n) (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0), \end{aligned}$$

where $\boldsymbol{\theta}_n$ lies on a line segment joining $\hat{\boldsymbol{\theta}}$ and $\boldsymbol{\theta}_0$. (Technically, the specific value of $\boldsymbol{\theta}_n$ varies by row in this expansion.) Rewriting this equation, we find

$$(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) = \left(- \sum_{i=1}^n \frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} \log f(\mathbf{y}_i | \boldsymbol{\theta}_n) \right)^{-1} \left(\sum_{i=1}^n \mathbf{S}_i \right)$$

where \mathbf{S}_i are the likelihood scores. Since the score \mathbf{S}_i is mean-zero (B.27) with covariance matrix $\boldsymbol{\Omega}$ (equation B.25) an application of the CLT yields

$$\frac{1}{\sqrt{n}} \sum_{i=1}^n \mathbf{S}_i \xrightarrow{d} \mathbf{N}(\mathbf{0}, \boldsymbol{\Omega}).$$

The analysis of the sample Hessian is somewhat more complicated due to the presence of $\boldsymbol{\theta}_n$. Let $\boldsymbol{\mathcal{H}}(\boldsymbol{\theta}) = -\frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} \log f(\mathbf{y}_i, \boldsymbol{\theta})$. If it is continuous in $\boldsymbol{\theta}$, then since $\boldsymbol{\theta}_n \xrightarrow{p} \boldsymbol{\theta}_0$ it follows that $\boldsymbol{\mathcal{H}}(\boldsymbol{\theta}_n) \xrightarrow{p} \boldsymbol{\mathcal{H}}$ and so

$$\begin{aligned} -\frac{1}{n} \sum_{i=1}^n \frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} \log f(\mathbf{y}_i, \boldsymbol{\theta}_n) &= \frac{1}{n} \sum_{i=1}^n \left(-\frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} \log f(\mathbf{y}_i, \boldsymbol{\theta}_n) - \boldsymbol{\mathcal{H}}(\boldsymbol{\theta}_n) \right) + \boldsymbol{\mathcal{H}}(\boldsymbol{\theta}_n) \\ &\xrightarrow{p} \boldsymbol{\mathcal{H}} \end{aligned}$$

by an application of a uniform WLLN. (By uniform, we mean that the WLLN holds uniformly over the parameter value. This requires the second derivative to be a smooth function of the parameter.)

Together,

$$\sqrt{n} (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \xrightarrow{d} \boldsymbol{\mathcal{H}}^{-1} \mathbf{N}(\mathbf{0}, \boldsymbol{\Omega}) = \mathbf{N}(\mathbf{0}, \boldsymbol{\mathcal{H}}^{-1} \boldsymbol{\Omega} \boldsymbol{\mathcal{H}}^{-1}) = \mathbf{N}(\mathbf{0}, \boldsymbol{\mathcal{I}}^{-1}),$$

the final equality using Theorem B.11.1. \blacksquare

Proof of Theorem B.11.3. Let $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)$ be the sample, and set

$$\mathbf{S} = \frac{\partial}{\partial \boldsymbol{\theta}} \log f_n(\mathbf{Y}, \boldsymbol{\theta}_0) = \sum_{i=1}^n \mathbf{S}_i$$

which by Theorem (B.11.1) has mean zero and variance $n\mathcal{I}$. Write the estimator $\tilde{\boldsymbol{\theta}} = \tilde{\boldsymbol{\theta}}(\mathbf{Y})$ as a function of the data. Since $\tilde{\boldsymbol{\theta}}$ is unbiased for any $\boldsymbol{\theta}$,

$$\boldsymbol{\theta} = \mathbb{E}\tilde{\boldsymbol{\theta}} = \int \tilde{\boldsymbol{\theta}}(\mathbf{Y}) f(\mathbf{Y}, \boldsymbol{\theta}) d\mathbf{Y}.$$

Differentiating with respect to $\boldsymbol{\theta}$ and evaluating at $\boldsymbol{\theta}_0$ yields

$$\begin{aligned} \mathbf{I} &= \int \tilde{\boldsymbol{\theta}}(\mathbf{Y}) \frac{\partial}{\partial \boldsymbol{\theta}'} f(\mathbf{Y}, \boldsymbol{\theta}) d\mathbf{Y} \\ &= \int \tilde{\boldsymbol{\theta}}(\mathbf{Y}) \frac{\partial}{\partial \boldsymbol{\theta}'} \log f(\mathbf{Y}, \boldsymbol{\theta}) f(\mathbf{Y}, \boldsymbol{\theta}_0) d\mathbf{Y} \\ &= \mathbb{E}(\tilde{\boldsymbol{\theta}} \mathbf{S}') \\ &= \mathbb{E}((\tilde{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \mathbf{S}') \end{aligned}$$

the final equality since $\mathbb{E}(\tilde{\boldsymbol{\theta}}) = \boldsymbol{\theta}_0$

By the matrix Cauchy-Schwarz inequality (B.18), $\mathbb{E}((\tilde{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \mathbf{S}') = \mathbf{I}$, and $\text{var}(\mathbf{S}) = \mathbb{E}(\mathbf{S}\mathbf{S}') = n\mathcal{I}$,

$$\begin{aligned} \text{var}(\tilde{\boldsymbol{\theta}}) &= \mathbb{E}((\tilde{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) (\tilde{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)') \\ &\geq \mathbb{E}((\tilde{\boldsymbol{\theta}} - \boldsymbol{\theta}_0) \mathbf{S}') \mathbb{E}(\mathbf{S}\mathbf{S}')^{-1} \mathbb{E}(\mathbf{S} (\tilde{\boldsymbol{\theta}} - \boldsymbol{\theta}_0)') \\ &= \mathbb{E}(\mathbf{S}\mathbf{S}')^{-1} \\ &= (n\mathcal{I})^{-1} \end{aligned}$$

as stated. ■

Appendix C

Numerical Optimization

Many econometric estimators are defined by an optimization problem of the form

$$\hat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta} \in \Theta}{\operatorname{argmin}} Q(\boldsymbol{\theta}) \quad (\text{C.1})$$

where the parameter is $\boldsymbol{\theta} \in \Theta \subset \mathbb{R}^m$ and the criterion function is $Q(\boldsymbol{\theta}) : \Theta \rightarrow \mathbb{R}$. For example NLLS, GLS, MLE and GMM estimators take this form. In most cases, $Q(\boldsymbol{\theta})$ can be computed for given $\boldsymbol{\theta}$, but $\hat{\boldsymbol{\theta}}$ is not available in closed form. In this case, numerical methods are required to obtain $\hat{\boldsymbol{\theta}}$.

C.1 Grid Search

Many optimization problems are either one dimensional ($m = 1$) or involve one-dimensional optimization as a sub-problem (for example, a line search). In this context grid search may be employed.

Grid Search. Let $\Theta = [a, b]$ be an interval. Pick some $\varepsilon > 0$ and set $G = (b - a)/\varepsilon$ to be the number of gridpoints. Construct an equally spaced grid on the region $[a, b]$ with G gridpoints, which is $\{\boldsymbol{\theta}(j) = a + j(b - a)/G : j = 0, \dots, G\}$. At each point evaluate the criterion function and find the gridpoint which yields the smallest value of the criterion, which is $\boldsymbol{\theta}(\hat{j})$ where $\hat{j} = \operatorname{argmin}_{0 \leq j \leq G} Q(\boldsymbol{\theta}(j))$. This value $\boldsymbol{\theta}(\hat{j})$ is the gridpoint estimate of $\boldsymbol{\theta}$. If the grid is sufficiently fine to capture small oscillations in $Q(\boldsymbol{\theta})$, the approximation error is bounded by ε , that is, $|\boldsymbol{\theta}(\hat{j}) - \hat{\boldsymbol{\theta}}| \leq \varepsilon$. Plots of $Q(\boldsymbol{\theta}(j))$ against $\boldsymbol{\theta}(j)$ can help diagnose errors in grid selection. This method is quite robust but potentially costly.

Two-Step Grid Search. The gridsearch method can be refined by a two-step execution. For an error bound of ε pick G so that $G^2 = (b - a)/\varepsilon$. For the first step define an equally spaced grid on the region $[a, b]$ with G gridpoints, which is $\{\boldsymbol{\theta}(j) = a + j(b - a)/G : j = 0, \dots, G\}$. At each point evaluate the criterion function and let $\hat{j} = \operatorname{argmin}_{0 \leq j \leq G} Q(\boldsymbol{\theta}(j))$. For the second step define an equally spaced grid on $[\boldsymbol{\theta}(\hat{j} - 1), \boldsymbol{\theta}(\hat{j} + 1)]$ with G gridpoints, which is $\{\boldsymbol{\theta}'(k) = \boldsymbol{\theta}(\hat{j} - 1) + 2k(b - a)/G^2 : k = 0, \dots, G\}$. Let $\hat{k} = \operatorname{argmin}_{0 \leq k \leq G} Q(\boldsymbol{\theta}'(k))$. The estimate of $\hat{\boldsymbol{\theta}}$ is $\boldsymbol{\theta}(\hat{k})$. The advantage of the two-step method over a one-step grid search is that the number of function evaluations has been reduced from $(b - a)/\varepsilon$ to $2\sqrt{(b - a)/\varepsilon}$ which can be substantial. The disadvantage is that if the function $Q(\boldsymbol{\theta})$ is irregular, the first-step grid may not bracket $\hat{\boldsymbol{\theta}}$ which thus would be missed.

C.2 Gradient Methods

Gradient Methods are iterative methods which produce a sequence $\boldsymbol{\theta}_i : i = 1, 2, \dots$ which are designed to converge to $\hat{\boldsymbol{\theta}}$. All require the choice of a starting value $\boldsymbol{\theta}_1$, and all require the

computation of the **gradient** of $Q(\boldsymbol{\theta})$

$$\mathbf{g}(\boldsymbol{\theta}) = \frac{\partial}{\partial \boldsymbol{\theta}} Q(\boldsymbol{\theta})$$

and some require the **Hessian**

$$\mathcal{H}(\boldsymbol{\theta}) = \frac{\partial^2}{\partial \boldsymbol{\theta} \partial \boldsymbol{\theta}'} Q(\boldsymbol{\theta}).$$

If the functions $\mathbf{g}(\boldsymbol{\theta})$ and $\mathcal{H}(\boldsymbol{\theta})$ are not analytically available, they can be calculated numerically. Take the j 'th element of $\mathbf{g}(\boldsymbol{\theta})$. Let δ_j be the j 'th unit vector (zeros everywhere except for a one in the j 'th row). Then for ε small

$$g_j(\boldsymbol{\theta}) \simeq \frac{Q(\boldsymbol{\theta} + \delta_j \varepsilon) - Q(\boldsymbol{\theta})}{\varepsilon}.$$

Similarly,

$$g_{jk}(\boldsymbol{\theta}) \simeq \frac{Q(\boldsymbol{\theta} + \delta_j \varepsilon + \delta_k \varepsilon) - Q(\boldsymbol{\theta} + \delta_k \varepsilon) - Q(\boldsymbol{\theta} + \delta_j \varepsilon) + Q(\boldsymbol{\theta})}{\varepsilon^2}$$

In many cases, numerical derivatives can work well but can be computationally costly relative to analytic derivatives. In some cases, however, numerical derivatives can be quite unstable.

Most gradient methods are a variant of **Newton's method** which is based on a quadratic approximation. By a Taylor's expansion for $\boldsymbol{\theta}$ close to $\hat{\boldsymbol{\theta}}$

$$0 = \mathbf{g}(\hat{\boldsymbol{\theta}}) \simeq \mathbf{g}(\boldsymbol{\theta}) + \mathcal{H}(\boldsymbol{\theta}) (\hat{\boldsymbol{\theta}} - \boldsymbol{\theta})$$

which implies

$$\hat{\boldsymbol{\theta}} = \boldsymbol{\theta} - \mathcal{H}(\boldsymbol{\theta})^{-1} \mathbf{g}(\boldsymbol{\theta}).$$

This suggests the iteration rule

$$\hat{\boldsymbol{\theta}}_{i+1} = \boldsymbol{\theta}_i - \mathcal{H}(\boldsymbol{\theta}_i)^{-1} \mathbf{g}(\boldsymbol{\theta}_i).$$

where

One problem with Newton's method is that it will send the iterations in the wrong direction if $\mathcal{H}(\boldsymbol{\theta}_i)$ is not positive definite. One modification to prevent this possibility is quadratic hill-climbing which sets

$$\hat{\boldsymbol{\theta}}_{i+1} = \boldsymbol{\theta}_i - (\mathcal{H}(\boldsymbol{\theta}_i) + \alpha_i \mathbf{I}_m)^{-1} \mathbf{g}(\boldsymbol{\theta}_i).$$

where α_i is set just above the smallest eigenvalue of $\mathcal{H}(\boldsymbol{\theta}_i)$ if $\mathcal{H}(\boldsymbol{\theta})$ is not positive definite.

Another productive modification is to add a scalar **steplength** λ_i . In this case the iteration rule takes the form

$$\boldsymbol{\theta}_{i+1} = \boldsymbol{\theta}_i - \mathbf{D}_i \mathbf{g}_i \lambda_i \tag{C.2}$$

where $\mathbf{g}_i = \mathbf{g}(\boldsymbol{\theta}_i)$ and $\mathbf{D}_i = \mathcal{H}(\boldsymbol{\theta}_i)^{-1}$ for Newton's method and $\mathbf{D}_i = (\mathcal{H}(\boldsymbol{\theta}_i) + \alpha_i \mathbf{I}_m)^{-1}$ for quadratic hill-climbing.

Allowing the steplength to be a free parameter allows for a line search, a one-dimensional optimization. To pick λ_i write the criterion function as a function of λ

$$Q(\lambda) = Q(\boldsymbol{\theta}_i + \mathbf{D}_i \mathbf{g}_i \lambda)$$

a one-dimensional optimization problem. There are two common methods to perform a line search. A **quadratic approximation** evaluates the first and second derivatives of $Q(\lambda)$ with respect to λ , and picks λ_i as the value minimizing this approximation. The **half-step** method considers the sequence $\lambda = 1, 1/2, 1/4, 1/8, \dots$. Each value in the sequence is considered and the criterion $Q(\boldsymbol{\theta}_i + \mathbf{D}_i \mathbf{g}_i \lambda)$ evaluated. If the criterion has improved over $Q(\boldsymbol{\theta}_i)$, use this value, otherwise move to the next element in the sequence.

Newton's method does not perform well if $Q(\boldsymbol{\theta})$ is irregular, and it can be quite computationally costly if $\mathbf{H}(\boldsymbol{\theta})$ is not analytically available. These problems have motivated alternative choices for the weight matrix D_i . These methods are called **Quasi-Newton** methods. Two popular methods are do to Davidson-Fletcher-Powell (DFP) and Broyden-Fletcher-Goldfarb-Shanno (BFGS).

Let

$$\begin{aligned}\Delta \mathbf{g}_i &= \mathbf{g}_i - \mathbf{g}_{i-1} \\ \Delta \boldsymbol{\theta}_i &= \boldsymbol{\theta}_i - \boldsymbol{\theta}_{i-1}\end{aligned}$$

and . The DFP method sets

$$D_i = D_{i-1} + \frac{\Delta \boldsymbol{\theta}_i \Delta \boldsymbol{\theta}'_i}{\Delta \boldsymbol{\theta}'_i \Delta \mathbf{g}_i} + \frac{D_{i-1} \Delta \mathbf{g}_i \Delta \mathbf{g}'_i D_{i-1}}{\Delta \mathbf{g}'_i D_{i-1} \Delta \mathbf{g}_i}.$$

The BFGS methods sets

$$D_i = D_{i-1} + \frac{\Delta \boldsymbol{\theta}_i \Delta \boldsymbol{\theta}'_i}{\Delta \boldsymbol{\theta}'_i \Delta \mathbf{g}_i} - \frac{\Delta \boldsymbol{\theta}_i \Delta \boldsymbol{\theta}'_i}{(\Delta \boldsymbol{\theta}'_i \Delta \mathbf{g}_i)^2} \Delta \mathbf{g}'_i D_{i-1} \Delta \mathbf{g}_i + \frac{\Delta \boldsymbol{\theta}_i \Delta \mathbf{g}'_i D_{i-1}}{\Delta \boldsymbol{\theta}'_i \Delta \mathbf{g}_i} + \frac{D_{i-1} \Delta \mathbf{g}_i \Delta \boldsymbol{\theta}'_i}{\Delta \boldsymbol{\theta}'_i \Delta \mathbf{g}_i}.$$

For any of the gradient methods, the iterations continue until the sequence has converged in some sense. This can be defined by examining whether $|\boldsymbol{\theta}_i - \boldsymbol{\theta}_{i-1}|$, $|Q(\boldsymbol{\theta}_i) - Q(\boldsymbol{\theta}_{i-1})|$ or $|g(\boldsymbol{\theta}_i)|$ has become small.

C.3 Derivative-Free Methods

All gradient methods can be quite poor in locating the global minimum when $Q(\boldsymbol{\theta})$ has several local minima. Furthermore, the methods are not well defined when $Q(\boldsymbol{\theta})$ is non-differentiable. In these cases, alternative optimization methods are required. One example is the **simplex method** of Nelder-Mead (1965).

A more recent innovation is the method of **simulated annealing (SA)**. For a review see Goffe, Ferrier, and Rodgers (1994). The SA method is a sophisticated random search. Like the gradient methods, it relies on an iterative sequence. At each iteration, a random variable is drawn and added to the current value of the parameter. If the resulting criterion is decreased, this new value is accepted. If the criterion is increased, it may still be accepted depending on the extent of the increase and another randomization. The latter property is needed to keep the algorithm from selecting a local minimum. As the iterations continue, the variance of the random innovations is shrunk. The SA algorithm stops when a large number of iterations is unable to improve the criterion. The SA method has been found to be successful at locating global minima. The downside is that it can take considerable computer time to execute.

Bibliography

- [1] Abadir, Karim M. and Jan R. Magnus (2005): *Matrix Algebra*, Cambridge University Press.
- [2] Aitken, A.C. (1935): "On least squares and linear combinations of observations," *Proceedings of the Royal Statistical Society*, 55, 42-48.
- [3] Akaike, H. (1973): "Information theory and an extension of the maximum likelihood principle." In B. Petroc and F. Csake, eds., *Second International Symposium on Information Theory*.
- [4] Anderson, T.W. and H. Rubin (1949): "Estimation of the parameters of a single equation in a complete system of stochastic equations," *The Annals of Mathematical Statistics*, 20, 46-63.
- [5] Andrews, Donald W. K. (1988): "Laws of large numbers for dependent non-identically distributed random variables," *Econometric Theory*, 4, 458-467.
- [6] Andrews, Donald W. K. (1991), "Asymptotic normality of series estimators for nonparametric and semiparametric regression models," *Econometrica*, 59, 307-345.
- [7] Andrews, Donald W. K. (1993), "Tests for parameter instability and structural change with unknown change point," *Econometrica*, 61, 821-8516.
- [8] Andrews, Donald W. K. and Moshe Buchinsky: (2000): "A three-step method for choosing the number of bootstrap replications," *Econometrica*, 68, 23-51.
- [9] Andrews, Donald W. K. and Werner Ploberger (1994): "Optimal tests when a nuisance parameter is present only under the alternative," *Econometrica*, 62, 1383-1414.
- [10] Ash, Robert B. (1972): *Real Analysis and Probability*, Academic Press.
- [11] Basmann, R. L. (1957): "A generalized classical method of linear estimation of coefficients in a structural equation," *Econometrica*, 25, 77-83.
- [12] Bekker, P.A. (1994): "Alternative approximations to the distributions of instrumental variable estimators," *Econometrica*, 62, 657-681.
- [13] Billingsley, Patrick (1968): *Convergence of Probability Measures*. New York: Wiley.
- [14] Billingsley, Patrick (1995): *Probability and Measure*, 3rd Edition, New York: Wiley.
- [15] Bose, A. (1988): "Edgeworth correction by bootstrap in autoregressions," *Annals of Statistics*, 16, 1709-1722.
- [16] Box, George E. P. and Dennis R. Cox, (1964). "An analysis of transformations," *Journal of the Royal Statistical Society, Series B*, 26, 211-252.
- [17] Breusch, T.S. and A.R. Pagan (1979): "The Lagrange multiplier test and its application to model specification in econometrics," *Review of Economic Studies*, 47, 239-253.

- [18] Brown, B. W. and Whitney K. Newey (2002): "GMM, efficient bootstrapping, and improved inference," *Journal of Business and Economic Statistics*.
- [19] Card, David (1995): "Using geographic variation in college proximity to estimate the return to schooling," in *Aspects of Labor Market Behavior: Essays in Honour of John Vanderkamp*, L.N. Christofides, E.K. Grant, and R. Swidinsky, editors. Toronto: University of Toronto Press.
- [20] Carlstein, E. (1986): "The use of subseries methods for estimating the variance of a general statistic from a stationary time series," *Annals of Statistics*, 14, 1171-1179.
- [21] Casella, George and Roger L. Berger (2002): *Statistical Inference*, 2nd Edition, Duxbury Press.
- [22] Chamberlain, Gary (1987): "Asymptotic efficiency in estimation with conditional moment restrictions," *Journal of Econometrics*, 34, 305-334.
- [23] Choi, In and Peter C.B. Phillips (1992): "Asymptotic and finite sample distribution theory for IV estimators and tests in partially identified structural equations," *Journal of Econometrics*, 51, 113-150.
- [24] Chow, G.C. (1960): "Tests of equality between sets of coefficients in two linear regressions," *Econometrica*, 28, 591-603.
- [25] Cragg, John (1992): "Quasi-Aitken Estimation for Heterskedasticity of Unknown Form" *Journal of Econometrics*, 54, 179-201.
- [26] Davidson, James (1994): *Stochastic Limit Theory: An Introduction for Econometricians*. Oxford: Oxford University Press.
- [27] Davison, A.C. and D.V. Hinkley (1997): *Bootstrap Methods and their Application*. Cambridge University Press.
- [28] Dickey, D.A. and W.A. Fuller (1979): "Distribution of the estimators for autoregressive time series with a unit root," *Journal of the American Statistical Association*, 74, 427-431.
- [29] Donald Stephen G. and Whitney K. Newey (2001): "Choosing the number of instruments," *Econometrica*, 69, 1161-1191.
- [30] Dufour, J.M. (1997): "Some impossibility theorems in econometrics with applications to structural and dynamic models," *Econometrica*, 65, 1365-1387.
- [31] Efron, Bradley (1979): "Bootstrap methods: Another look at the jackknife," *Annals of Statistics*, 7, 1-26.
- [32] Efron, Bradley (1982): *The Jackknife, the Bootstrap, and Other Resampling Plans*. Society for Industrial and Applied Mathematics.
- [33] Efron, Bradley and R.J. Tibshirani (1993): *An Introduction to the Bootstrap*, New York: Chapman-Hall.
- [34] Eicker, F. (1963): "Asymptotic normality and consistency of the least squares estimators for families of linear regressions," *Annals of Mathematical Statistics*, 34, 447-456.
- [35] Engle, Robert F. and Clive W. J. Granger (1987): "Co-integration and error correction: Representation, estimation and testing," *Econometrica*, 55, 251-276.
- [36] Frisch, Ragnar (1933): "Editorial," *Econometrica*, 1, 1-4.

- [37] Frisch, Ragnar and F. Waugh (1933): "Partial time regressions as compared with individual trends," *Econometrica*, 1, 387-401.
- [38] Gallant, A. Ronald and D.W. Nychka (1987): "Seminonparametric maximum likelihood estimation," *Econometrica*, 55, 363-390.
- [39] Gallant, A. Ronald and Halbert White (1988): *A Unified Theory of Estimation and Inference for Nonlinear Dynamic Models*. New York: Basil Blackwell.
- [40] Galton, Francis (1886): "Regression Towards Mediocrity in Hereditary Stature," *The Journal of the Anthropological Institute of Great Britain and Ireland*, 15, 246-263.
- [41] Goldberger, Arthur S. (1964): *Econometric Theory*, Wiley.
- [42] Goldberger, Arthur S. (1968): *Topics in Regression Analysis*, Macmillan
- [43] Goldberger, Arthur S. (1991): *A Course in Econometrics*. Cambridge: Harvard University Press.
- [44] Goffe, W.L., G.D. Ferrier and J. Rogers (1994): "Global optimization of statistical functions with simulated annealing," *Journal of Econometrics*, 60, 65-99.
- [45] Gosset, William S. (a.k.a. "Student") (1908): "The probable error of a mean," *Biometrika*, 6, 1-25.
- [46] Gauss, K.F. (1809): "Theoria motus corporum coelestium," in *Werke*, Vol. VII, 240-254.
- [47] Granger, Clive W. J. (1969): "Investigating causal relations by econometric models and cross-spectral methods," *Econometrica*, 37, 424-438.
- [48] Granger, Clive W. J. (1981): "Some properties of time series data and their use in econometric specification," *Journal of Econometrics*, 16, 121-130.
- [49] Granger, Clive W. J. and Timo Teräsvirta (1993): *Modelling Nonlinear Economic Relationships*, Oxford University Press, Oxford.
- [50] Gregory, A. and M. Veall (1985): "On formulating Wald tests of nonlinear restrictions," *Econometrica*, 53, 1465-1468,
- [51] Haavelmo, T. (1944): "The probability approach in econometrics," *Econometrica*, supplement, 12.
- [52] Hall, A. R. (2000): "Covariance matrix estimation and the power of the overidentifying restrictions test," *Econometrica*, 68, 1517-1527,
- [53] Hall, P. (1992): *The Bootstrap and Edgeworth Expansion*, New York: Springer-Verlag.
- [54] Hall, P. (1994): "Methodology and theory for the bootstrap," *Handbook of Econometrics, Vol. IV*, eds. R.F. Engle and D.L. McFadden. New York: Elsevier Science.
- [55] Hall, P. and J.L. Horowitz (1996): "Bootstrap critical values for tests based on Generalized-Method-of-Moments estimation," *Econometrica*, 64, 891-916.
- [56] Hahn, J. (1996): "A note on bootstrapping generalized method of moments estimators," *Econometric Theory*, 12, 187-197.
- [57] Hamilton, James D. (1994) *Time Series Analysis*.

- [58] Hansen, Bruce E. (1992): "Efficient estimation and testing of cointegrating vectors in the presence of deterministic trends," *Journal of Econometrics*, 53, 87-121.
- [59] Hansen, Bruce E. (1996): "Inference when a nuisance parameter is not identified under the null hypothesis," *Econometrica*, 64, 413-430.
- [60] Hansen, Bruce E. (2006): "Edgeworth expansions for the Wald and GMM statistics for non-linear restrictions," *Econometric Theory and Practice: Frontiers of Analysis and Applied Research*, edited by Dean Corbae, Steven N. Durlauf and Bruce E. Hansen. Cambridge University Press.
- [61] Hansen, Lars Peter (1982): "Large sample properties of generalized method of moments estimators," *Econometrica*, 50, 1029-1054.
- [62] Hansen, Lars Peter, John Heaton, and A. Yaron (1996): "Finite sample properties of some alternative GMM estimators," *Journal of Business and Economic Statistics*, 14, 262-280.
- [63] Hausman, J.A. (1978): "Specification tests in econometrics," *Econometrica*, 46, 1251-1271.
- [64] Heckman, J. (1979): "Sample selection bias as a specification error," *Econometrica*, 47, 153-161.
- [65] Horn, S.D., R.A. Horn, and D.B. Duncan. (1975) "Estimating heteroscedastic variances in linear model," *Journal of the American Statistical Association*, 70, 380-385.
- [66] Horowitz, Joel (2001): "The Bootstrap," *Handbook of Econometrics, Vol. 5*, J.J. Heckman and E.E. Leamer, eds., Elsevier Science, 3159-3228.
- [67] Imbens, G.W. (1997): "One step estimators for over-identified generalized method of moments models," *Review of Economic Studies*, 64, 359-383.
- [68] Imbens, G.W., R.H. Spady and P. Johnson (1998): "Information theoretic approaches to inference in moment condition models," *Econometrica*, 66, 333-357.
- [69] Jarque, C.M. and A.K. Bera (1980): "Efficient tests for normality, homoskedasticity and serial independence of regression residuals," *Economic Letters*, 6, 255-259.
- [70] Johansen, S. (1988): "Statistical analysis of cointegrating vectors," *Journal of Economic Dynamics and Control*, 12, 231-254.
- [71] Johansen, S. (1991): "Estimation and hypothesis testing of cointegration vectors in the presence of linear trend," *Econometrica*, 59, 1551-1580.
- [72] Johansen, S. (1995): *Likelihood-Based Inference in Cointegrated Vector Auto-Regressive Models*, Oxford University Press.
- [73] Johansen, S. and K. Juselius (1992): "Testing structural hypotheses in a multivariate cointegration analysis of the PPP and the UIP for the UK," *Journal of Econometrics*, 53, 211-244.
- [74] Kitamura, Y. (2001): "Asymptotic optimality and empirical likelihood for testing moment restrictions," *Econometrica*, 69, 1661-1672.
- [75] Kitamura, Y. and M. Stutzer (1997): "An information-theoretic alternative to generalized method of moments," *Econometrica*, 65, 861-874..
- [76] Koenker, Roger (2005): *Quantile Regression*. Cambridge University Press.
- [77] Kunsch, H.R. (1989): "The jackknife and the bootstrap for general stationary observations," *Annals of Statistics*, 17, 1217-1241.

- [78] Kwiatkowski, D., P.C.B. Phillips, P. Schmidt, and Y. Shin (1992): "Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root?" *Journal of Econometrics*, 54, 159-178.
- [79] Lafontaine, F. and K.J. White (1986): "Obtaining any Wald statistic you want," *Economics Letters*, 21, 35-40.
- [80] Lehmann, E.L. and George Casella (1998): *Theory of Point Estimation*, 2nd Edition, Springer.
- [81] Lehmann, E.L. and Joseph P. Romano (2005): *Testing Statistical Hypotheses*, 3rd Edition, Springer.
- [82] Lindeberg, Jarl Waldemar, (1922): "Eine neue Herleitung des Exponentialgesetzes in der Wahrscheinlichkeitsrechnung," *Mathematische Zeitschrift*, 15, 211-225.
- [83] Li, Qi and Jeffrey Racine (2007) *Nonparametric Econometrics*.
- [84] Lovell, M.C. (1963): "Seasonal adjustment of economic time series," *Journal of the American Statistical Association*, 58, 993-1010.
- [85] MacKinnon, James G. (1990): "Critical values for cointegration," in Engle, R.F. and C.W. Granger (eds.) *Long-Run Economic Relationships: Readings in Cointegration*, Oxford, Oxford University Press.
- [86] MacKinnon, James G. and Halbert White (1985): "Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties," *Journal of Econometrics*, 29, 305-325.
- [87] Magnus, J. R., and H. Neudecker (1988): *Matrix Differential Calculus with Applications in Statistics and Econometrics*, New York: John Wiley and Sons.
- [88] Mann, H.B. and A. Wald (1943). "On stochastic limit and order relationships," *The Annals of Mathematical Statistics* 14, 217-226.
- [89] Muirhead, R.J. (1982): *Aspects of Multivariate Statistical Theory*. New York: Wiley.
- [90] Nelder, J. and R. Mead (1965): "A simplex method for function minimization," *Computer Journal*, 7, 308-313.
- [91] Nerlove, Marc (1963): "Returns to Scale in Electricity Supply," Chapter 7 of *Measurement in Economics* (C. Christ, et al, eds.). Stanford: Stanford University Press, 167-198.
- [92] Newey, Whitney K. (1990): "Semiparametric efficiency bounds," *Journal of Applied Econometrics*, 5, 99-135.
- [93] Newey, Whitney K. (1997): "Convergence rates and asymptotic normality for series estimators," *Journal of Econometrics*, 79, 147-168.
- [94] Newey, Whitney K. and Daniel L. McFadden (1994): "Large Sample Estimation and Hypothesis Testing," in Robert Engle and Daniel McFadden, (eds.) *Handbook of Econometrics*, vol. IV, 2111-2245, North Holland: Amsterdam.
- [95] Newey, Whitney K. and Kenneth D. West (1987): "Hypothesis testing with efficient method of moments estimation," *International Economic Review*, 28, 777-787.
- [96] Owen, Art B. (1988): "Empirical likelihood ratio confidence intervals for a single functional," *Biometrika*, 75, 237-249.

- [97] Owen, Art B. (2001): *Empirical Likelihood*. New York: Chapman & Hall.
- [98] Park, Joon Y. and Peter C. B. Phillips (1988): "On the formulation of Wald tests of nonlinear restrictions," *Econometrica*, 56, 1065-1083,
- [99] Phillips, Peter C.B. (1989): "Partially identified econometric models," *Econometric Theory*, 5, 181-240.
- [100] Phillips, Peter C.B. and Sam Ouliaris (1990): "Asymptotic properties of residual based tests for cointegration," *Econometrica*, 58, 165-193.
- [101] Politis, D.N. and J.P. Romano (1996): "The stationary bootstrap," *Journal of the American Statistical Association*, 89, 1303-1313.
- [102] Potscher, B.M. (1991): "Effects of model selection on inference," *Econometric Theory*, 7, 163-185.
- [103] Qin, J. and J. Lawless (1994): "Empirical likelihood and general estimating equations," *The Annals of Statistics*, 22, 300-325.
- [104] Ramsey, J. B. (1969): "Tests for specification errors in classical linear least-squares regression analysis," *Journal of the Royal Statistical Society, Series B*, 31, 350-371.
- [105] Rudin, W. (1987): *Real and Complex Analysis*, 3rd edition. New York: McGraw-Hill.
- [106] Runge, Carl (1901): "Über empirische Funktionen und die Interpolation zwischen äquidistanten Ordinaten," *Zeitschrift für Mathematik und Physik*, 46, 224-243.
- [107] Said, S.E. and D.A. Dickey (1984): "Testing for unit roots in autoregressive-moving average models of unknown order," *Biometrika*, 71, 599-608.
- [108] Secrist, Horace (1933): *The Triumph of Mediocrity in Business*. Evanston: Northwestern University.
- [109] Shao, J. and D. Tu (1995): *The Jackknife and Bootstrap*. NY: Springer.
- [110] Sargan, J.D. (1958): "The estimation of economic relationships using instrumental variables," *Econometrica*, 26, 393-415.
- [111] Shao, Jun (2003): *Mathematical Statistics*, 2nd edition, Springer.
- [112] Sheather, S.J. and M.C. Jones (1991): "A reliable data-based bandwidth selection method for kernel density estimation," *Journal of the Royal Statistical Society, Series B*, 53, 683-690.
- [113] Shin, Y. (1994): "A residual-based test of the null of cointegration against the alternative of no cointegration," *Econometric Theory*, 10, 91-115.
- [114] Silverman, B.W. (1986): *Density Estimation for Statistics and Data Analysis*. London: Chapman and Hall.
- [115] Sims, C.A. (1972): "Money, income and causality," *American Economic Review*, 62, 540-552.
- [116] Sims, C.A. (1980): "Macroeconomics and reality," *Econometrica*, 48, 1-48.
- [117] Staiger, D. and James H. Stock (1997): "Instrumental variables regression with weak instruments," *Econometrica*, 65, 557-586.
- [118] Stock, James H. (1987): "Asymptotic properties of least squares estimators of cointegrating vectors," *Econometrica*, 55, 1035-1056.

- [119] Stock, James H. (1991): "Confidence intervals for the largest autoregressive root in U.S. macroeconomic time series," *Journal of Monetary Economics*, 28, 435-460.
- [120] Stock, James H. and Jonathan H. Wright (2000): "GMM with weak identification," *Econometrica*, 68, 1055-1096.
- [121] Stock, James H. and Mark W. Watson (2010): *Introduction to Econometrics*, 3rd edition, Addison-Wesley.
- [122] Stone, Marshall H. (1937): "Applications of the Theory of Boolean Rings to General Topology," *Transactions of the American Mathematical Society*, 41, 375-481.
- [123] Stone, Marshall H. (1948): "The Generalized Weierstrass Approximation Theorem," *Mathematics Magazine*, 21, 167-184.
- [124] Theil, Henri. (1953): "Repeated least squares applied to complete equation systems," The Hague, Central Planning Bureau, mimeo.
- [125] Theil, Henri (1961): *Economic Forecasts and Policy*. Amsterdam: North Holland.
- [126] Theil, Henri. (1971): *Principles of Econometrics*, New York: Wiley.
- [127] Tobin, James (1958): "Estimation of relationships for limited dependent variables," *Econometrica*, 26, 24-36.
- [128] Tripathi, Gautam (1999): "A matrix extension of the Cauchy-Schwarz inequality," *Economics Letters*, 63, 1-3.
- [129] van der Vaart, A.W. (1998): *Asymptotic Statistics*, Cambridge University Press.
- [130] Wald, A. (1943): "Tests of statistical hypotheses concerning several parameters when the number of observations is large," *Transactions of the American Mathematical Society*, 54, 426-482.
- [131] Wang, J. and E. Zivot (1998): "Inference on structural parameters in instrumental variables regression with weak instruments," *Econometrica*, 66, 1389-1404.
- [132] Weierstrass, K. (1885): "Über die analytische Darstellbarkeit sogenannter willkürlicher Functionen einer reellen Veränderlichen," *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin*, 1885.
- [133] White, Halbert (1980): "A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity," *Econometrica*, 48, 817-838.
- [134] White, Halbert (1984): *Asymptotic Theory for Econometricians*, Academic Press.
- [135] Wooldridge, Jeffrey M. (2010) *Econometric Analysis of Cross Section and Panel Data*, 2nd edition, MIT Press.
- [136] Wooldridge, Jeffrey M. (2009) *Introductory Econometrics: A Modern Approach*, 4th edition, Southwestern
- [137] Zellner, Arnold. (1962): "An efficient method of estimating seemingly unrelated regressions, and tests for aggregation bias," *Journal of the American Statistical Association*, 57, 348-368.
- [138] Zhang, Fuzhen and Qingling Zhang (2006): "Eigenvalue inequalities for matrix product," *IEEE Transactions on Automatic Control*, 51, 1506-1509.