- Cox Regression (II)

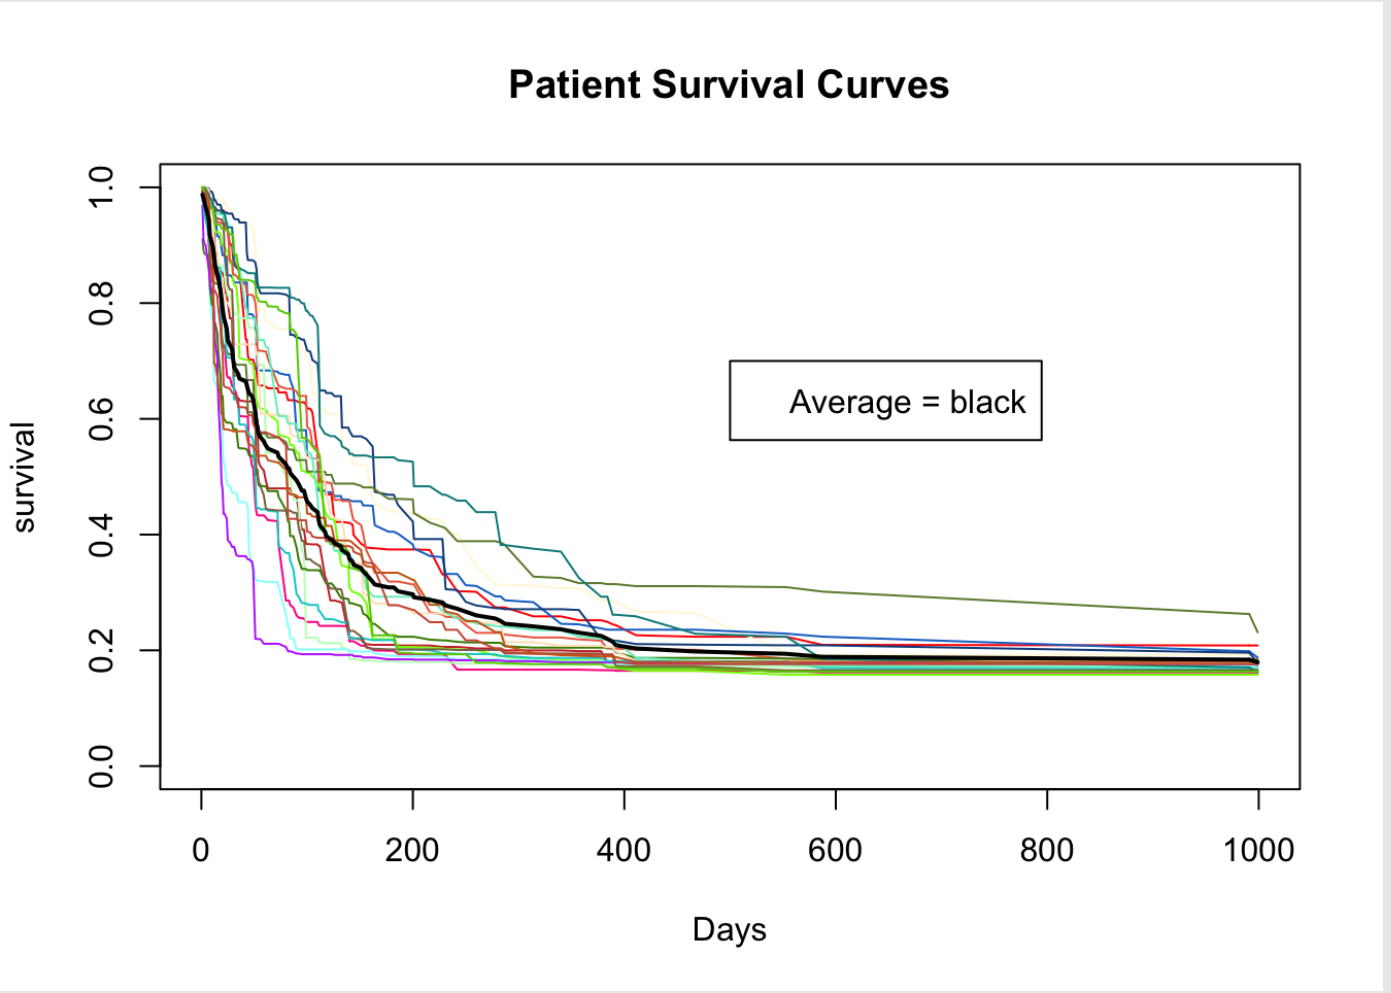

Cox Regression Model

The scale on which ***linearity*** is assumed is the ***log-hazard*** scale:

$$h(t|X) = h_0(t)exp\left(X_1\beta_1 + X_2\beta_2 + X_3\beta_3 + \cdots + X_p\beta_p\right)$$

$$log\left(\frac{h(t|X)}{h_0(t)}\right) = X_1\beta_1 + X_2\beta_2 + X_3\beta_3 + \cdots + X_p\beta_p$$

- $h_0(t)$ is the **baseline hazard function**

- the exponential function represents the effect of the linear combination of the covariates X on the hazard

The aim is to determine the **joint** effect of the covariates on the hazard or to focus on a **specific** effect.

Proportional hazards (PH)

The hazard ***at any given time*** for an individual in one group is proportional to the hazard ***at any given time*** for an individual in the other group. If the hazard functions are proportional **->** survival functions **do not cross** one another…

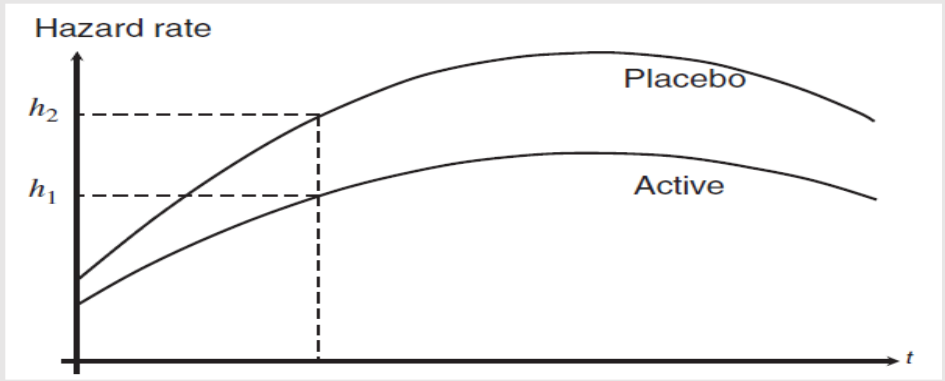

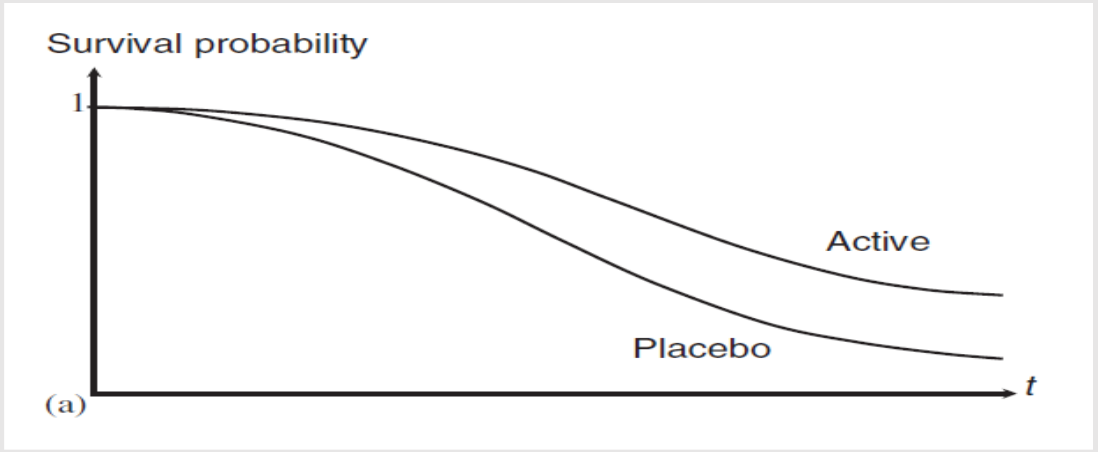

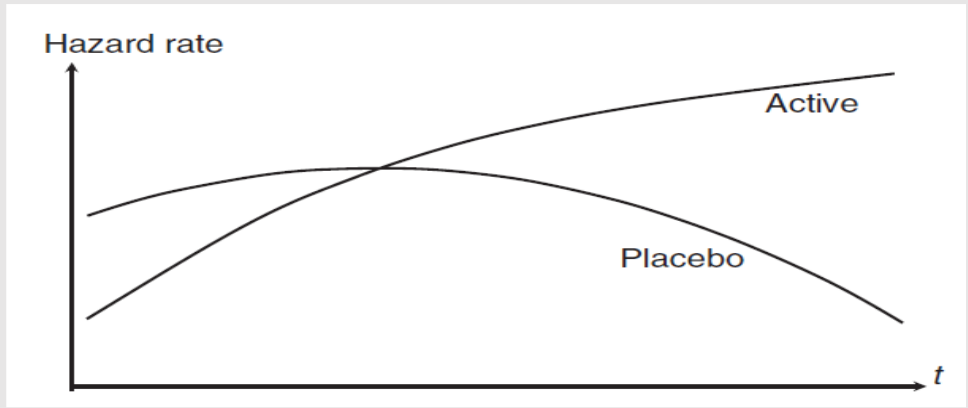

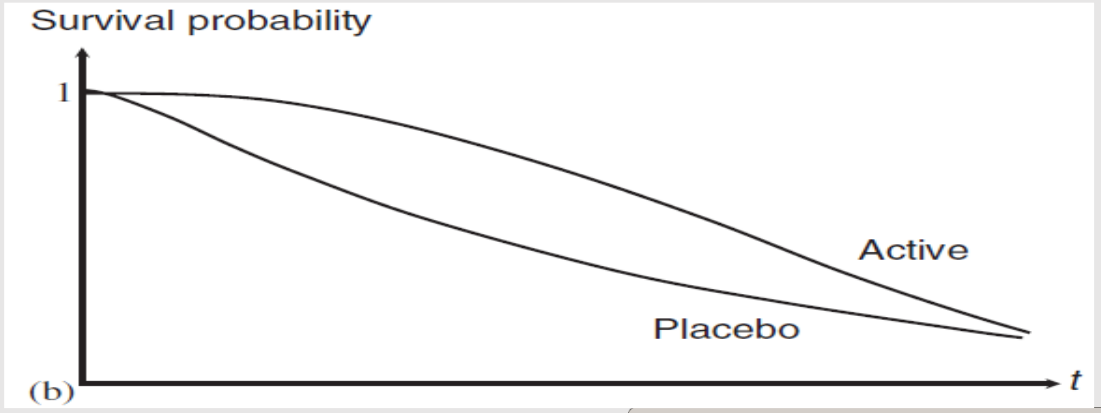

UNITÀ DI BIOSTATISTICA
Dipartimento Universitario Clinico di Scienze Mediche Chirurgiche e della Salute

Cox model assumes **proportional hazards** (PH). Covariates X have always *the same relative effect* along time:

$$h(t|X) = h_0(t)\exp(X_1\beta_1 + X_2\beta_2 + \ldots + X_p\beta_p) = h_0(t)\exp(\mathbf{X\beta})$$

The function $\exp(\mathbf{X\beta})$ does not depend on *t*

Hazard Ratio between two subjects, with covariates X and X* does not depend on *t*:

$$\frac{h_0(t)\exp(\mathbf{X\beta})}{h_0(t)\exp(\mathbf{X^*\beta})} = \exp((\mathbf{X} - \mathbf{X^*})\boldsymbol{\beta})$$

If PH assumption does not hold, *standard* Cox model could be no longer valid **[we could check for this]** [there are extensions]

$$h_i(t|X_i) = h_0(t)exp(X_i\beta)$$

- β_k is the **difference** in the log-hazard function comparing two subpopulations differing in x_k by "1-unit" and that are similar with respect to all other covariates in the model

- the **effect** expressed by β_k is ***adjusted for*** all other covariates in the model, so it has the interpretation of a log-relative hazard associated with a change in x_k, holding other covariates constant ***at some fixed value***

- is it possible to compare ***hypothetical patients*** with different covariates values and check how their ***estimated survival curves*** appear; [remind: the baseline hazard depends on the study cohort…]

- the Cox PH model is a model for the hazard more than a model for survival time, **although they are related one-to-one if no competing risks exists**

Survival function derived from the Cox regression model

(no competing risks, no time-dependent variables)

Once the β are estimated, we can obtain the corresponding survival function:

$$S(t|x) = S_0(t)\exp(\beta x)$$

$\boldsymbol{S_0(t)}$ is derived from an estimate of the **cumulative baseline hazard**
(a derivation in the **non-parametric** form, similar to the Nelson-Aalen formulation)

The estimate of $S_0(t)$ and a fixed set of values for the explanatory variables produce an estimate of the survival function **for a specific person or group**.

The expression for $S(t|x)$ shows that proportional hazard functions dictate that the estimated survival functions ***do not intersect***.

UNITÀ DI BIOSTATISTICA
Dipartimento Universitario Clinico di Scienze Mediche Chirurgiche e della Salute

Block 4.3

Table 3 Univariable and multivariable analysis (primary analysis Cox)

Outcome: Death or CV hosp

Variables	Univariable analysis		Multivariable analysis	
	Hazard ratio (95% CI)	*P*-value	Hazard ratio (95% CI)	*P*-value
Cardiac rehabilitation	0.601 (0.476–0.758)	<0.001	0.578 (0.432–0.773)	<0.001
NSTEMI	1.361 (1.043–1.774)	0.023		
Male	1.168 (0.898–1.517)	0.246		
STEMI	0.908 (0.701–1.176)	0.463		
PCI	1.343 (1.036–1.742)	0.026		
CABG	0.621 (0.465–0.828)	0.001	0.639 (0.466–0.876)	**0.005**
Ejection fraction	0.979 (0.967–0.991)	0.001	0.986 (0.973–0.999)	**0.035**
Diabetes	1.548 (1.219–1.966)	<0.001	1.460 (1.107–1.926)	**0.007**
Hypertension	1.161 (0.887–1.520)	0.276		
Smoking	1.121 (0.860–1.463)	0.398		
Dyslipidaemia	0.897 (0.708–1.138)	0.372		
Beta-blockers	1. 244 (0.910–1.701)	0.171		
ACE-inhibitors/ARBs	1.367 (1.005–1.859)	0.046		
Statins/ezetimibe	0.607 (0.426–0.865)	0.006	0.518 (0.345–0.776)	**0.001**
ASA	0.932 (0.510–1.703)	0.819		
DAPT	1.245 (0.968–1.601)	0.088		
Chronic kidney disease	2.409 (1.823–3.182)	<0.001	2.441 (1.775–3.358)	<0.001
Previous ACS	1.443 (1.111–1.873)	0.006		
Previous PCI	1.718 (1.299–2.272)	<0.001		
Previous CABG	1.884 (1.240–2.861)	0.003		

ACE-inhibitors, angiotensin-converting enzyme inhibitor; ACS, acute coronary syndrome; ARBs, angiotensin receptor blockers; ASA, acetylsalicylic acid; CABG, coronary artery bypass graft; CI, confidence interval; DAPT, dual antiplatelet therapy; NSTEMI, non-ST-elevation myocardial infarction; PCI, percutaneous coronary intervention; STEMI, ST-elevation myocardial infarction.

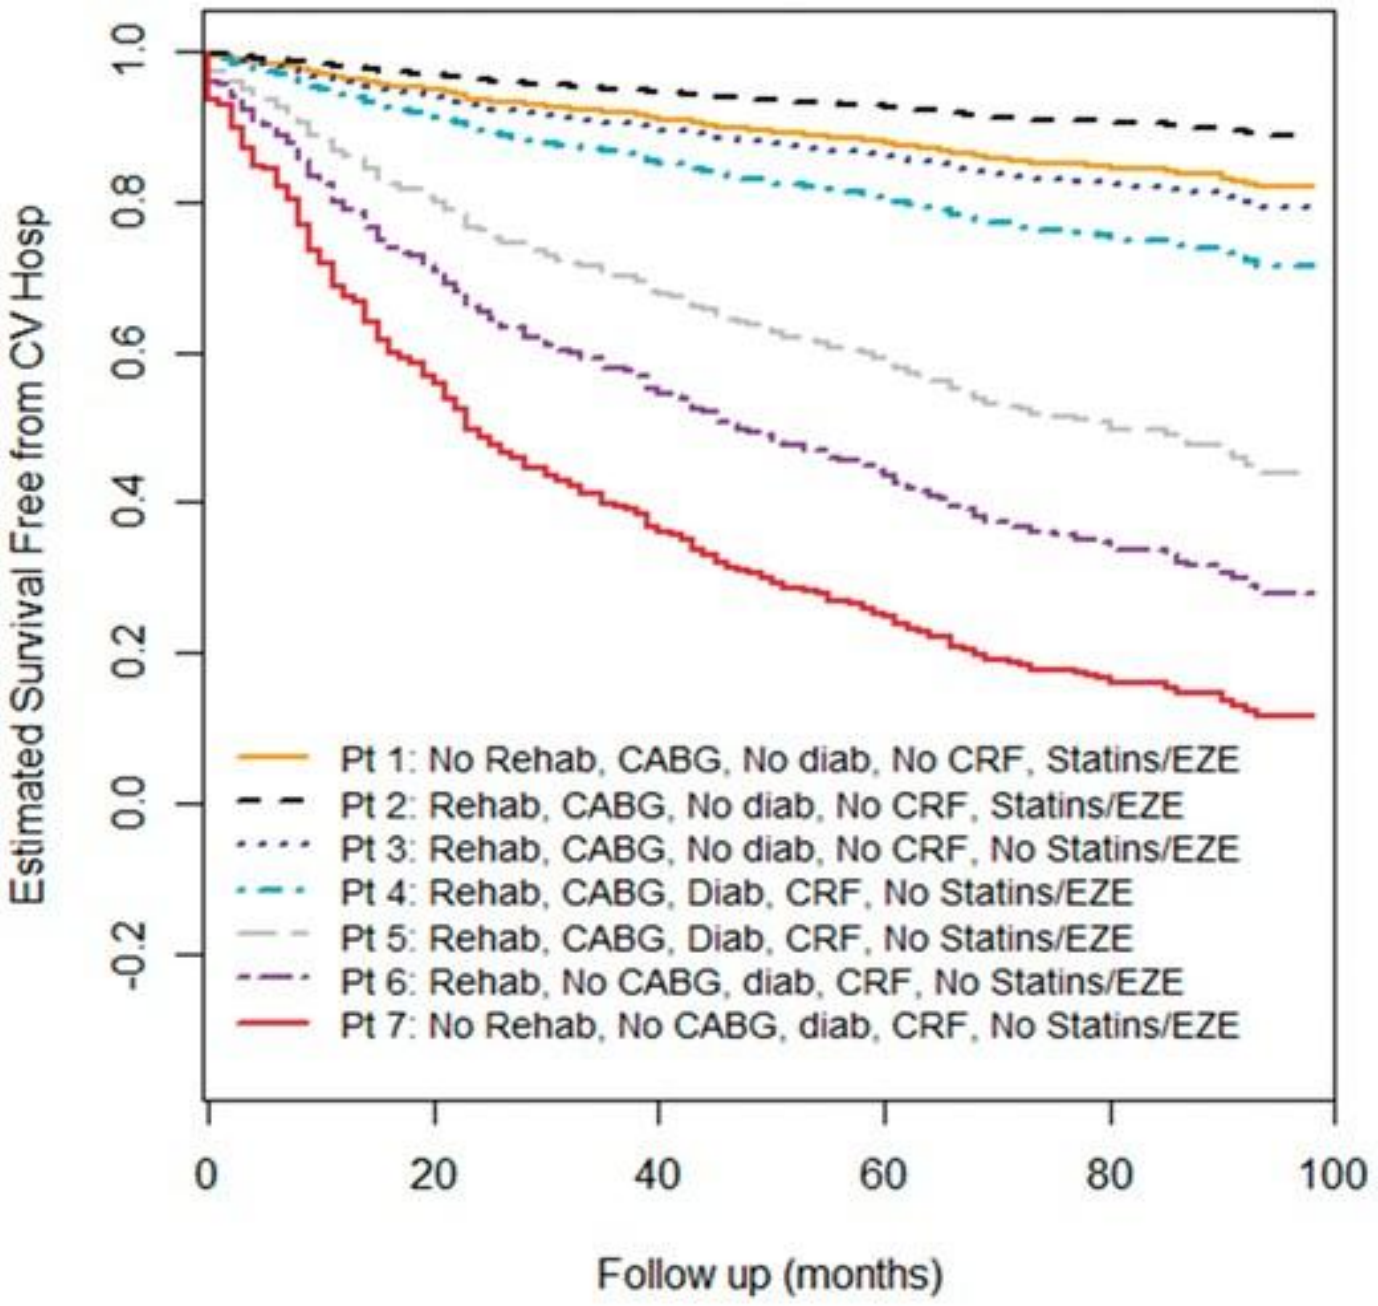

Estimated survival curves from the Cox model.

The curves are estimated for patients having the median ejection fraction (56%) of the population.

CABG : coronary artery bypass graft;
CRF: chronic renal failure;
Diab: diabetes;
EZE: ezetimibe;
Rehab: cardiac rehabilitation.

This study demonstrated the positive effects of CR program in the real world showing a decreased risk of CV hospitalizations and mortality during a long-term follow-up.

UNITÀ DI BIOSTATISTICA
Dipartimento Universitario Clinico di Scienze Mediche Chirurgiche e della Salute

The Cox model assumes that the **hazards are proportional** (PH), which means that the hazard ratio is **constant over time** with different predictor or covariate levels.

This PH assumption in any covariate is quite a strong assumption. Considering the complexity of biological and physiological responses and associations, this assumption has rarely a solid justification.

If PH doesn't exactly hold for a particular covariate but we fit the PH model anyway, then what we are getting is sort of an ***average HR***, averaged over the event times.

The two most common ways to assess the PH assumption are:

- *Visual assessment* by means of *the* ***log-cumulative hazard plot***
- *Testing* of *scaled* ***Schoenfeld residuals***

Eventually, if the non-PH variable is a categorical one, it could make sense using a **stratified** approach

$$h_i(t|X_i) = h_0(t)exp(X_i\beta)$$

$$\int_0^t h_i(u)du = exp(X_i\beta)\int_0^t h_0(u)du$$

$$H_i(t|X_i) = exp(X_i\beta)H_0(t)$$

Cumulative hazard functions

$$log\big(H_i(t|X_i)\big) = X_i\beta + log\big(H_0(t)\big)$$

If the estimated **log-cumulative hazards** for individuals **with different values** of X (categorical) are plotted against time, the curves will be ***parallel*** if the PH assumption is valid.

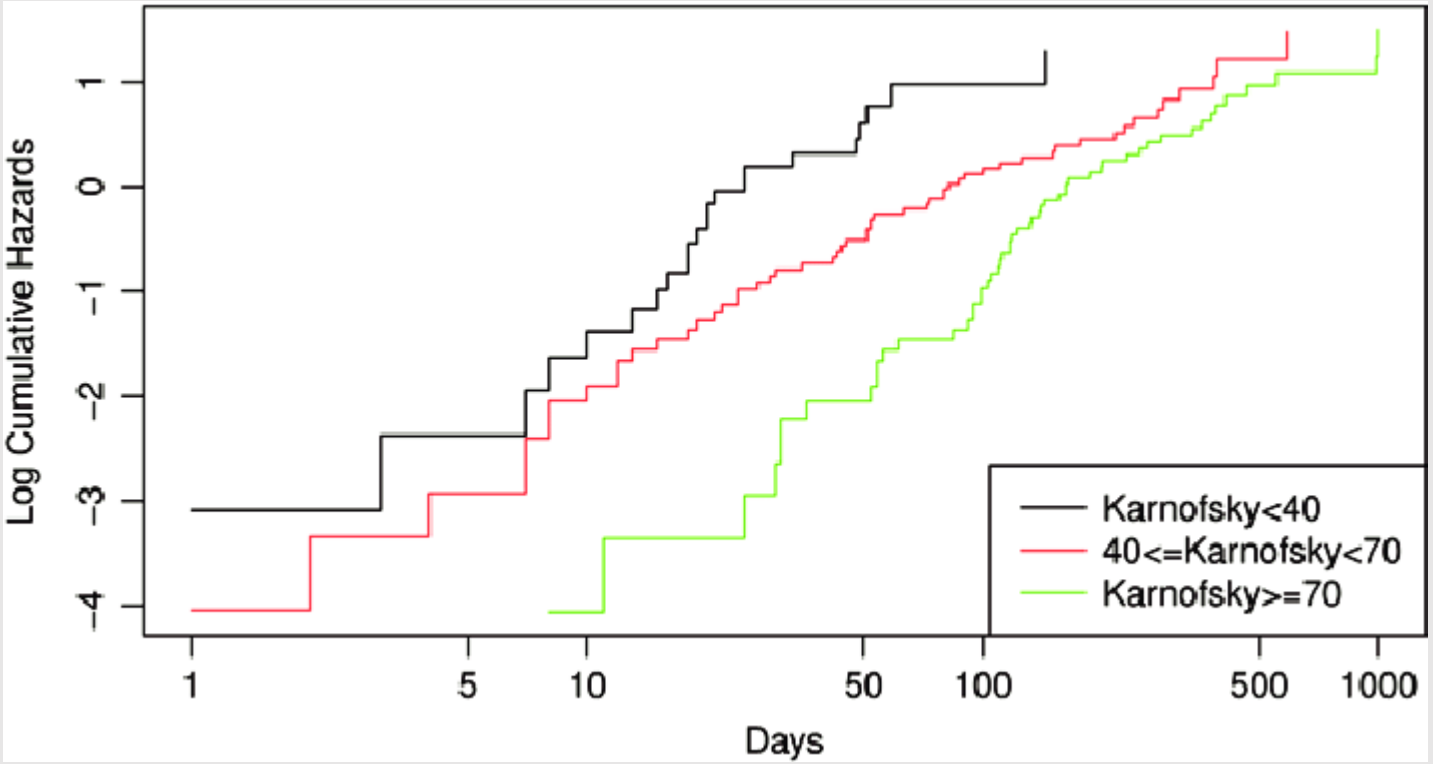

- Values of X need to be **categorical/grouped**
- Just a **visual** appreciation

UNITÀ DI BIOSTATISTICA
Dipartimento Universitario Clinico di Scienze Mediche Chirurgiche e della Salute

Just a note about *Schoenfeld* residuals

Time-varying *residuals* from the model are added to the corresponding **time-invariant** coefficient estimate β and smoothed. The result is a **plot** of an estimate of the regression coefficient for the covariate ***over time***. If the plot is ***reasonably flat*** (there is here a formal test), the PH assumption holds.

$$s_{k,j}$$

Schoenfeld residual for covariate Xj at time tk

$$E(s_{k,j}) + \hat{\beta}_j \approx \beta_j(t_k)$$

The Schoenfeld residuals are the differences between that individual's covariate values at the event time k and the corresponding risk-weighted average of covariate values among all those at risk at that time.

The word "residual" thus makes sense, as it's the difference between an observed covariate value and what you *might have expected* based on all those at risk at that time.

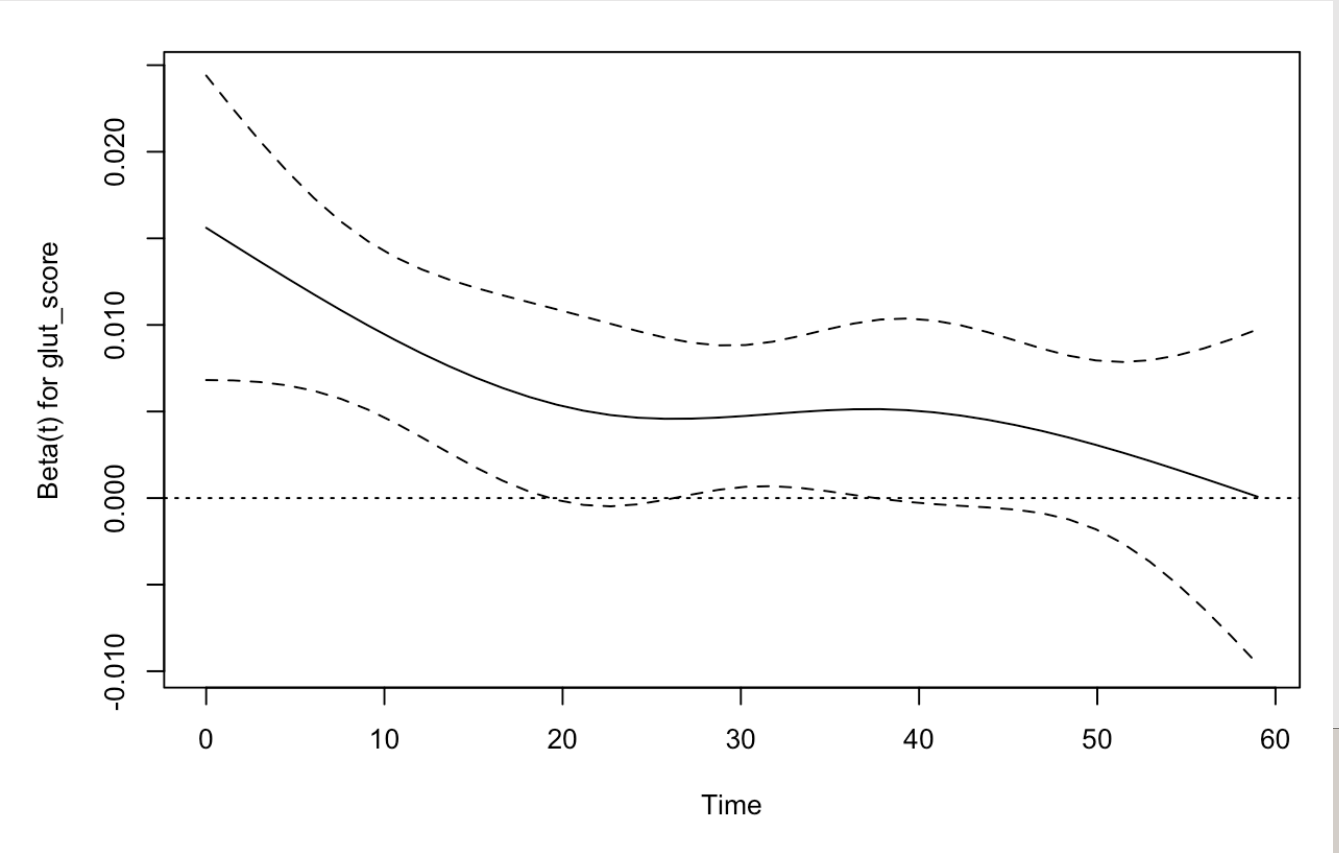

The Stratified Cox Model

Suppose a confounder C has *k* levels on which we would like **to stratify** when comparing *h(t | E)* and *h(t | not E)* where E is an indicator of "exposure".

$$h_i(t|E) = h_{0i}(t)\exp(E\beta)$$
$$i = 1,..,k$$

1. A [non-parametric] baseline hazard is estimated **within** each stratum (solve ev. non PH hazard)

2. If the confounder is controlled using stratification, there is no way to estimate an **hazard ratio** comparing two levels of the confounder.

3. Stratification generally requires **more data** to obtain the same precision in coefficient estimates

Example of application

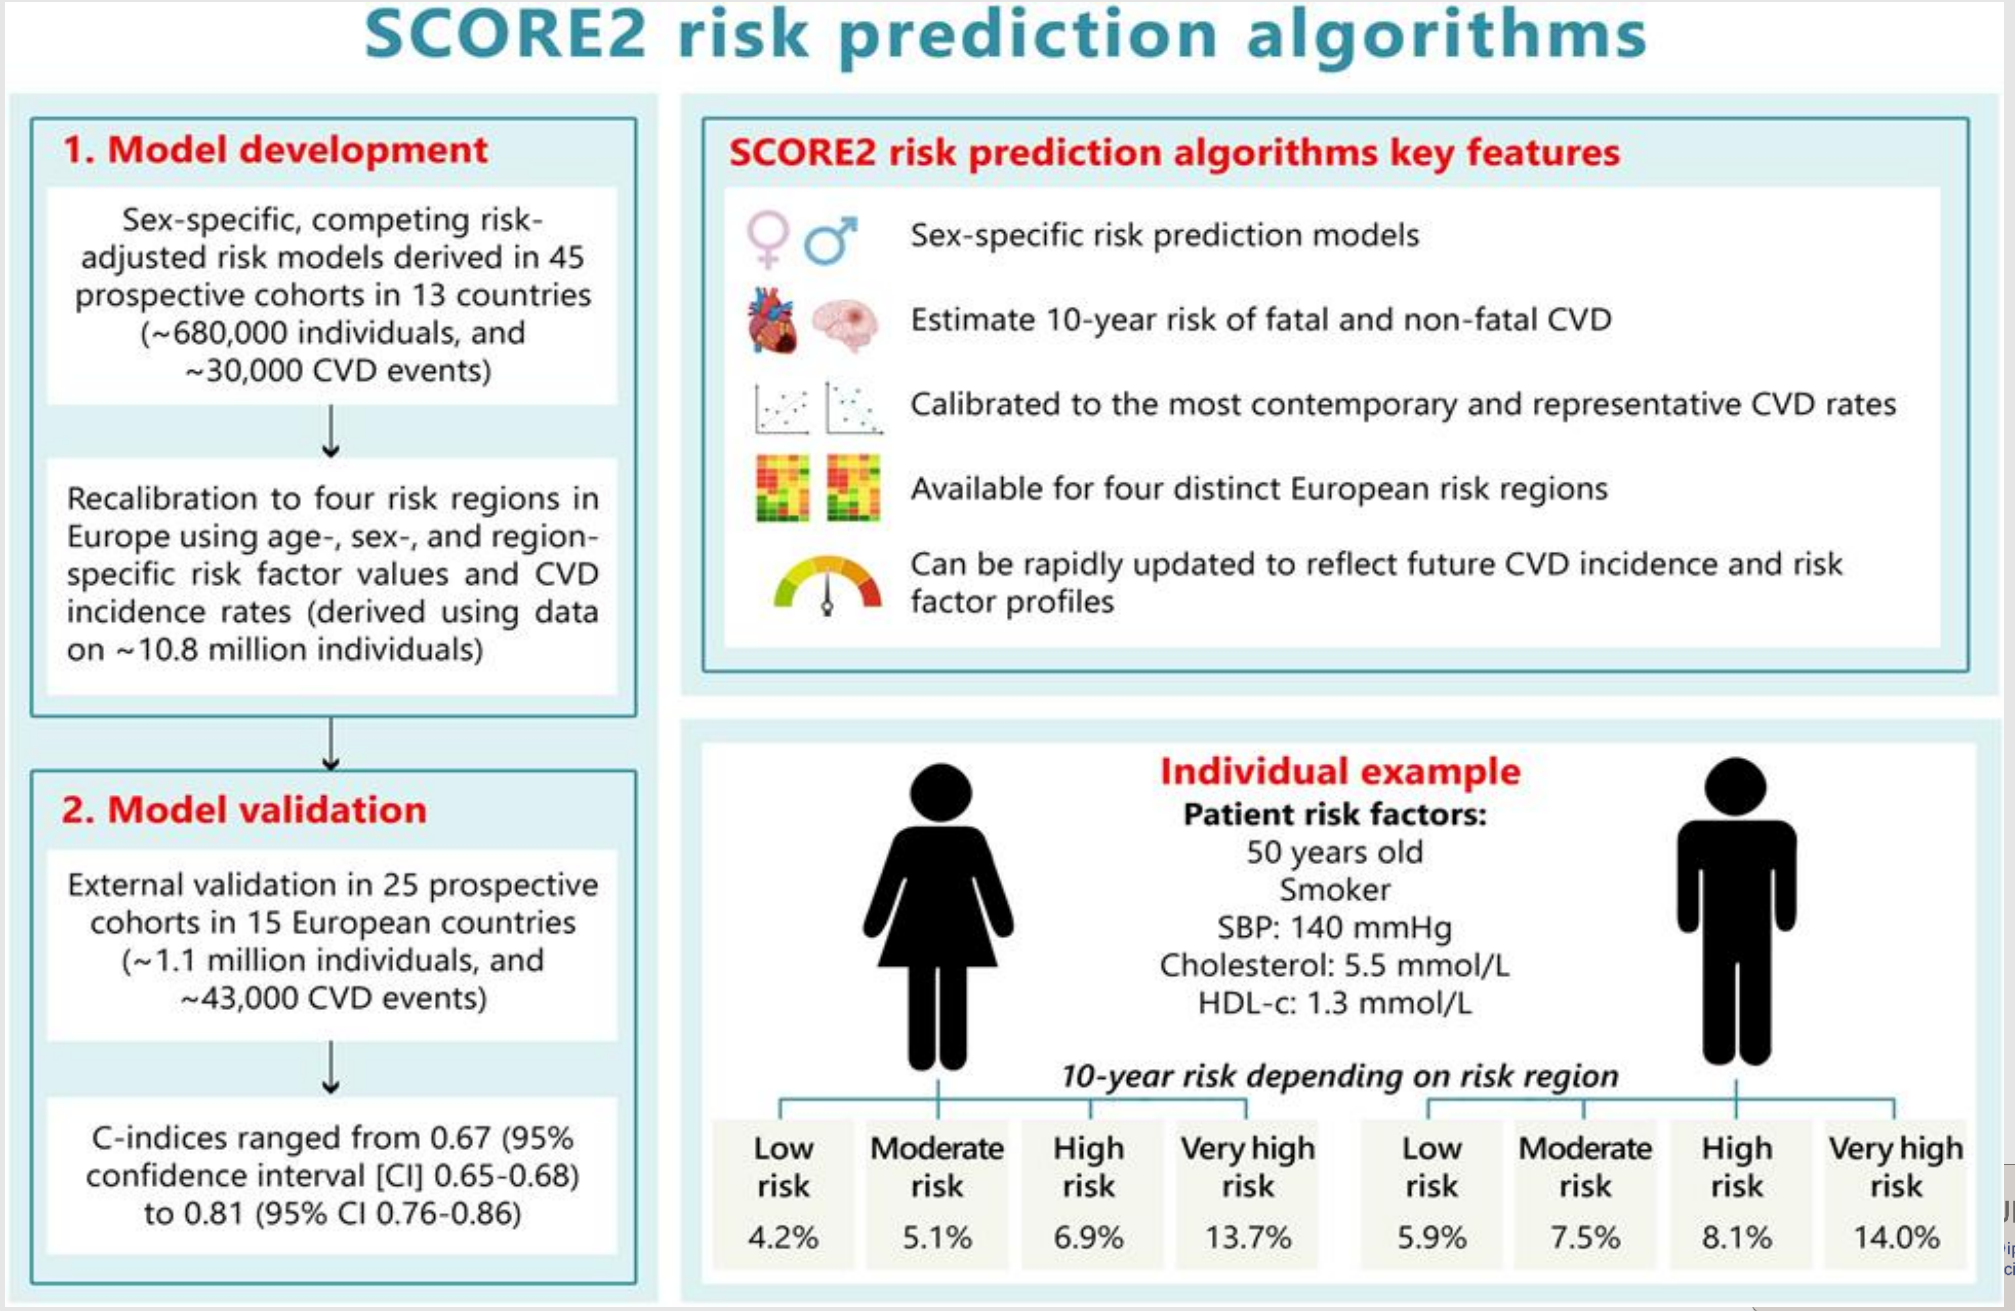

The SCORE2 (OP) algorithms are used to estimate **10-year cardiovascular risk** in individuals aged 40-69 and 70+, respectively. These algorithms, developed by the European Society of Cardiology (ESC), are designed for use across various regions in Europe, including those with low, moderate, high, and very high risk profiles (**different baseline hazard**).

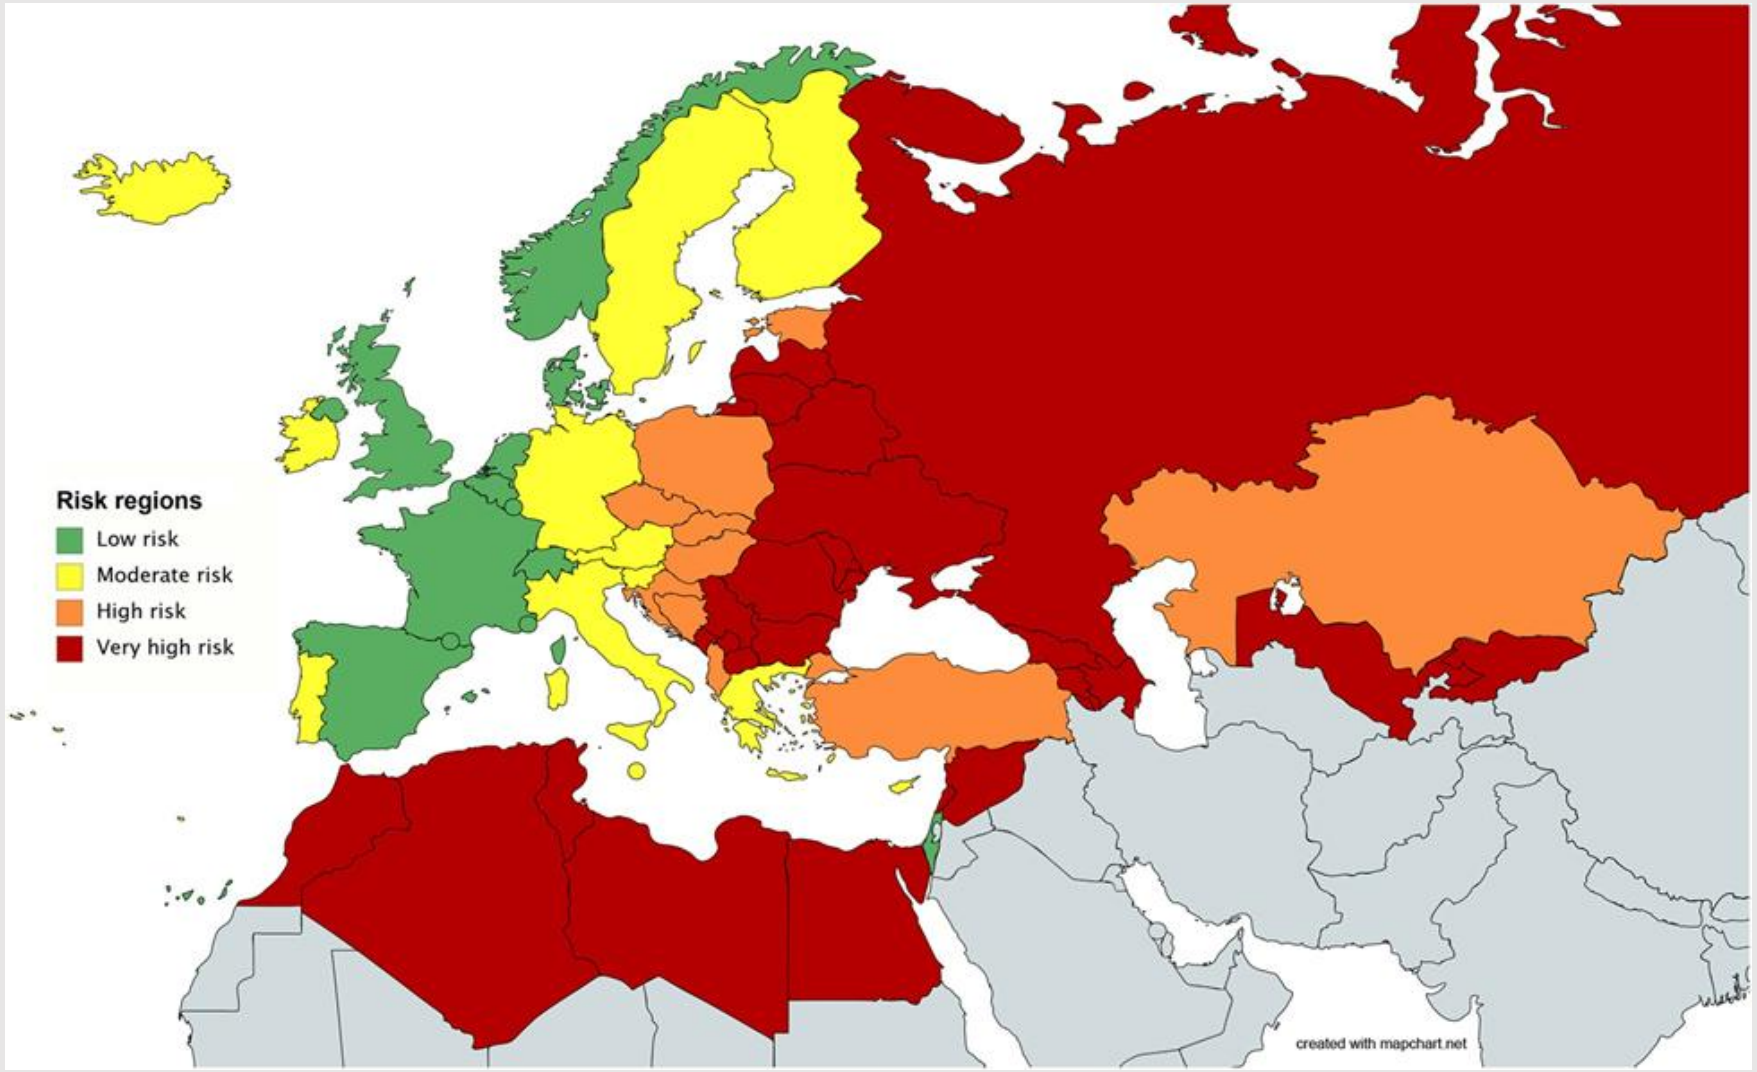

Countries were grouped into **four** risk regions according to their most recently reported WHO **age- and sex-standardized overall CVD mortality rates** per 100,000 population

- **low risk** (<100 CVD deaths per 100,000)
- **moderate risk** (100 to <150 CVD deaths per 100,000)
- **high risk** (150 to <300 CVD deaths per 100,000)
- **very high risk** (≥300 CVD deaths per 100,000)

UNITÀ DI BIOSTATISTICA
Dipartimento Universitario Clinico di Scienze Mediche Chirurgiche e della Salute

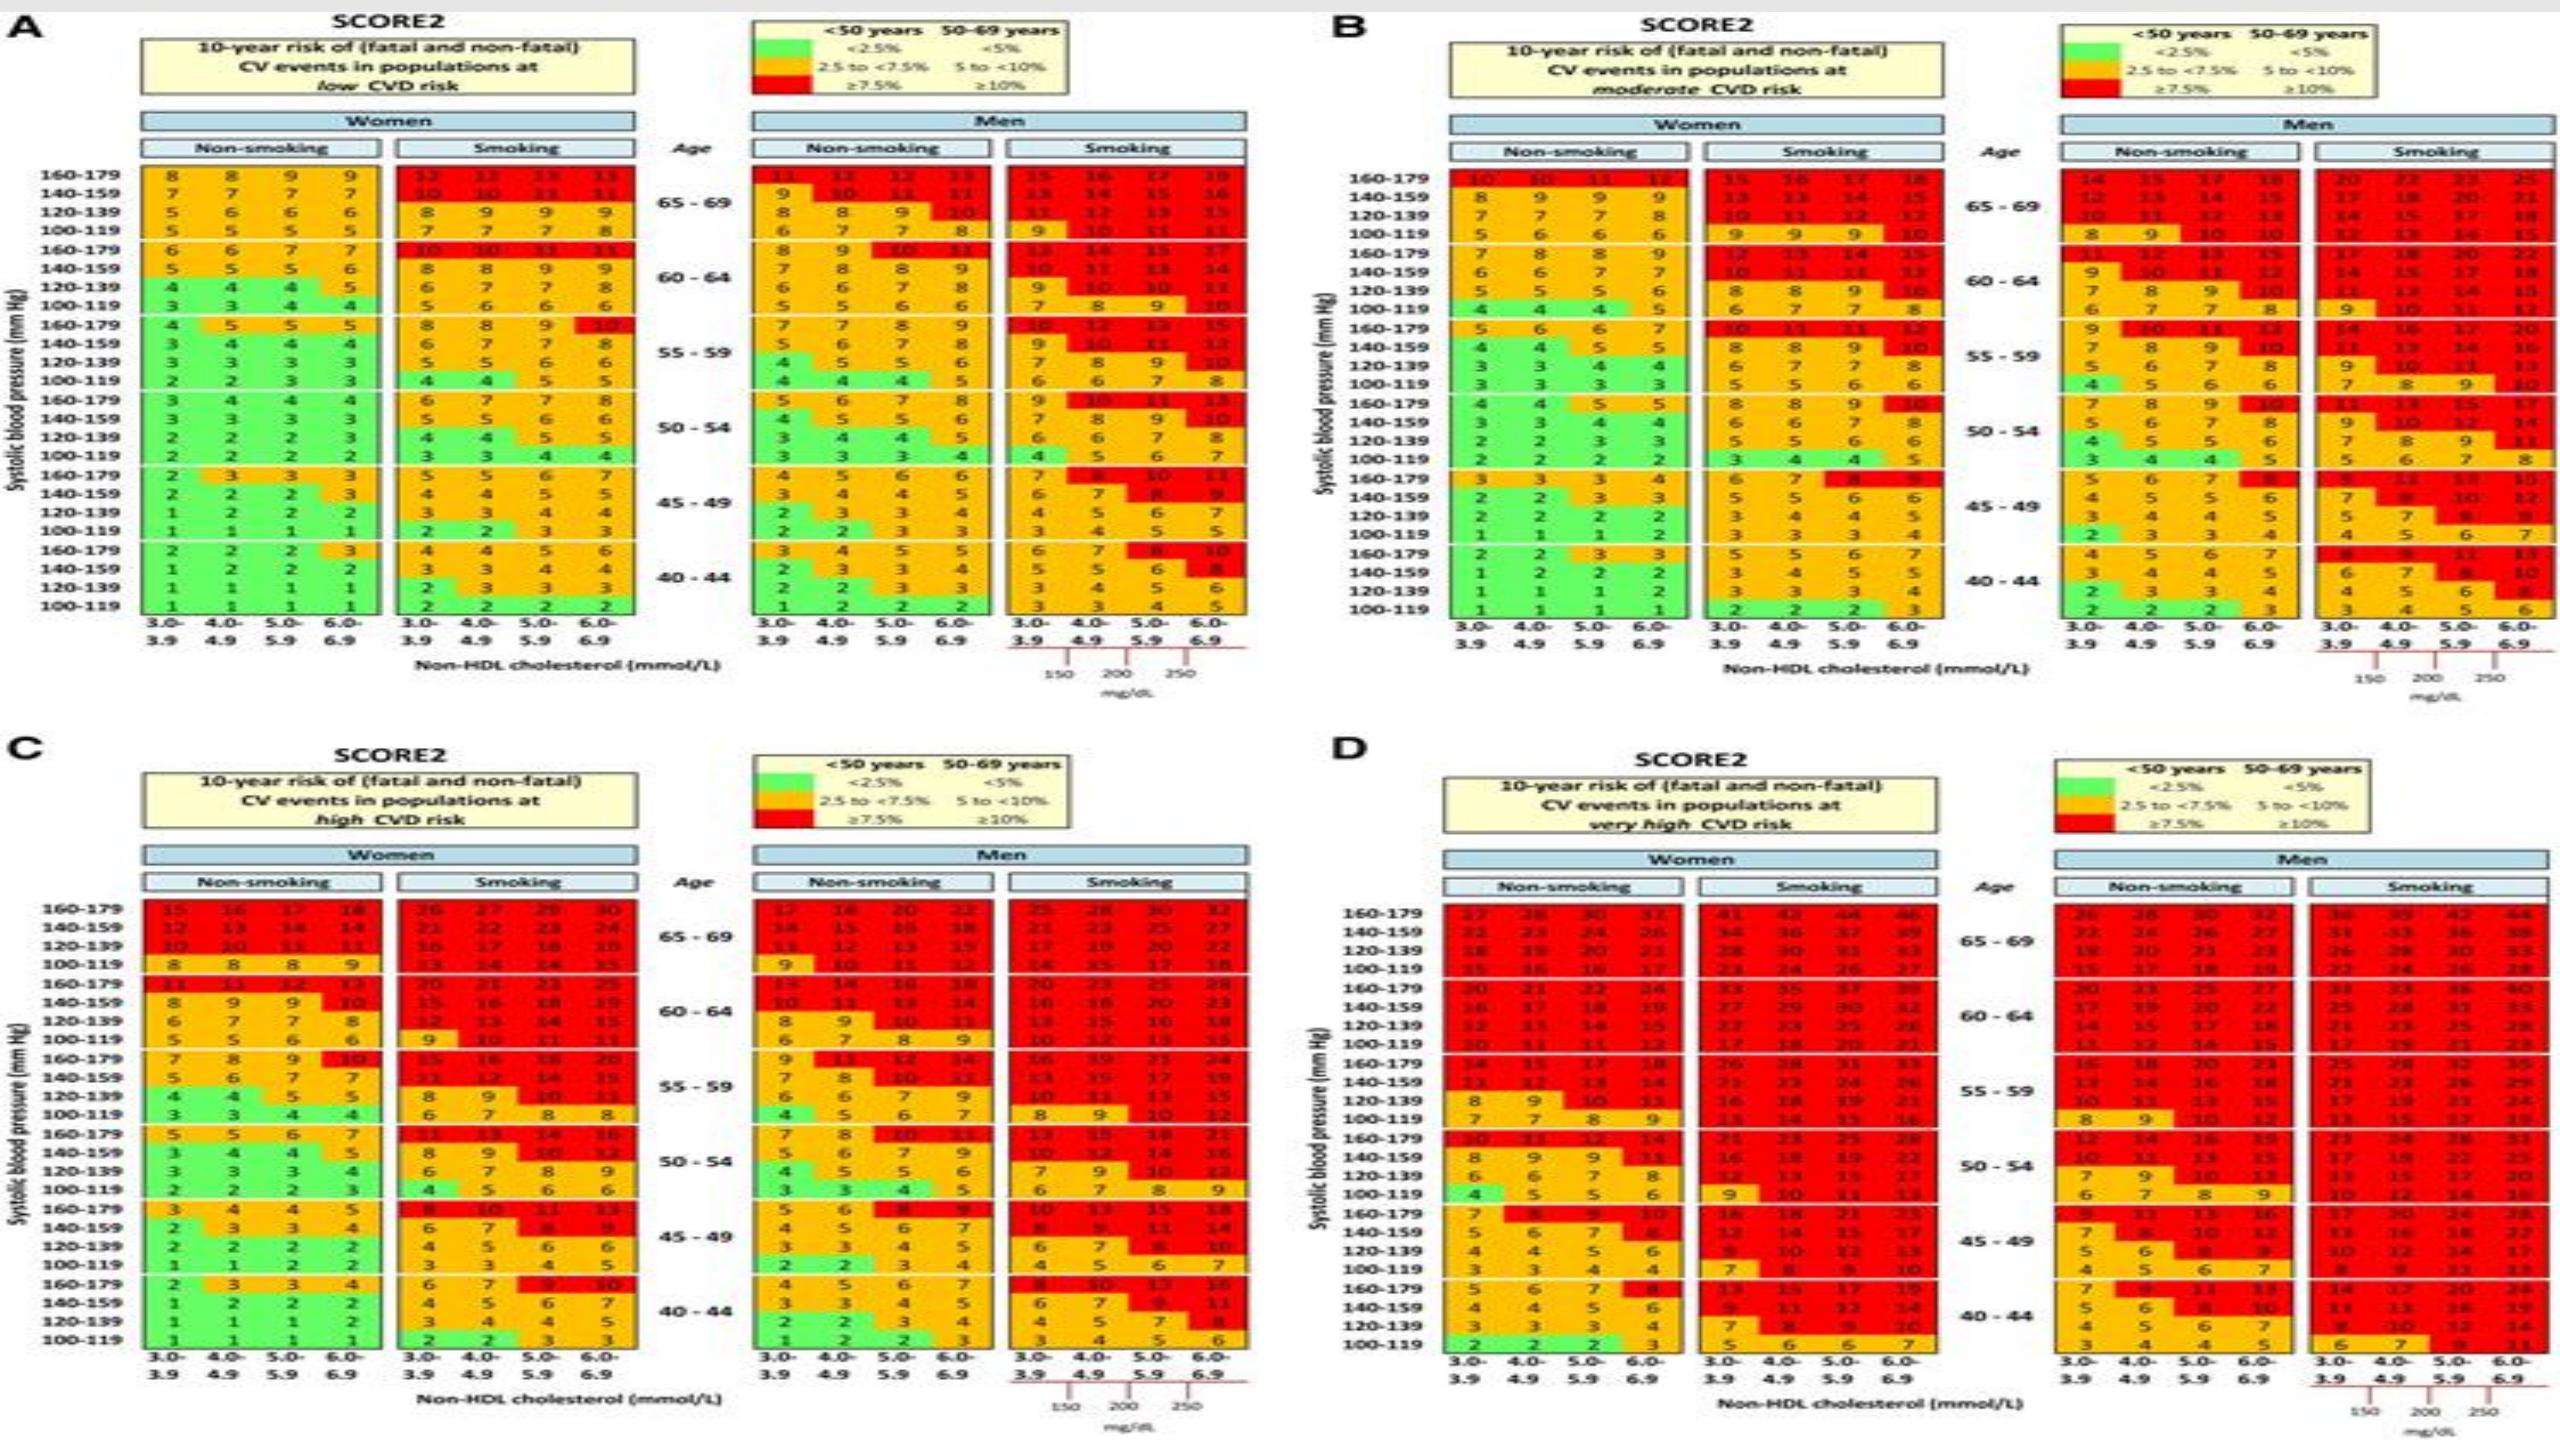
A
SCORE2
10-year risk of (fatal and non-fatal) CV events in populations at low CVD risk
<50 years 50-69 years
<2.5% <5%
2.5 to <7.5% 5 to <10%
≥7.5% ≥10%
Women
Men
Non-smoking
Smoking
Age
65 - 69
60 - 64
55 - 59
50 - 54
45 - 49
40 - 44
Systolic blood pressure (mm Hg)
160-179
140-159
120-139
100-119
3.0-3.9
4.0-4.9
5.0-5.9
6.0-6.9
Non-HDL cholesterol (mmol/L)
150 200 250
mg/dL
B
SCORE2
10-year risk of (fatal and non-fatal) CV events in populations at moderate CVD risk
C
SCORE2
10-year risk of (fatal and non-fatal) CV events in populations at high CVD risk
D
SCORE2
10-year risk of (fatal and non-fatal) CV events in populations at very high CVD risk

Worked Example in R :

- library(ISwR)
- data(melanom)

status: indicator of the patient's status by the end of the study:
1="dead from malignant melanoma"
2= "alive"
3= "dead from other causes"

days: observation time in days
ulc: 1=present (tumor ulcerated) 2 = absent
thick: tumor thickness
sex: 1 for women and 2 for men

	no	status	days	ulc	thick	sex
1	789	3	10	1	676	2
2	13	3	30	2	65	2
3	97	2	35	2	134	2
4	16	3	99	2	290	1
5	21	1	185	1	1208	2
6	469	1	204	1	484	2
7	685	1	210	1	516	2
8	7	1	232	1	1288	2
9	932	3	232	1	322	1
10	944	1	279	1	741	1
11	558	1	295	1	419	1
12	612	3	355	1	16	1
13	2	1	386	1	387	1
14	233	1	426	1	484	2
15	418	1	469	1	242	1
16	765	3	493	1	1256	2
17	777	1	529	1	580	2
18	61	1	621	1	706	2
19	67	1	629	1	548	2
20	819	1	659	1	773	2
21	10	1	667	1	1385	1
22	15	1	718	1	234	2

Consider a model with the single regressor sex:

```
mod.sex <- coxph(Surv(days,status==1)~sex)
summary(mod.sex)
```

```
##       coef exp(coef) se(coef)     z Pr(>|z|)
## sex 0.6622    1.9390   0.2651 2.498   0.0125 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
##     exp(coef) exp(-coef) lower .95 upper .95
## sex     1.939     0.5157     1.153      3.26
##
## Concordance= 0.59  (se = 0.033 )
## Rsquare= 0.03   (max possible= 0.937 )
## Likelihood ratio test= 6.15  on 1 df,   p=0.01314
## Wald test            = 6.24  on 1 df,   p=0.01251
## Score (logrank) test = 6.47  on 1 df,   p=0.01098
```

Males (=2) have an hazard nearly twice than women (=1)

These tests are all equivalent in large samples but may differ somewhat in small-sample cases

R code for the Cox Model

A more elaborate example, involving also a continuous covariate:

```
mod.cov <- coxph(Surv(days,status==1)~sex+log(thick))
summary(mod.cov)
```

'thick' is the tumor thickness; we use logarithm since the distribution is asymmetric:

HR of log(thick)=2.18
each 1 point change in log(thick) is associated with a 2.2-fold increase in a patient's risk

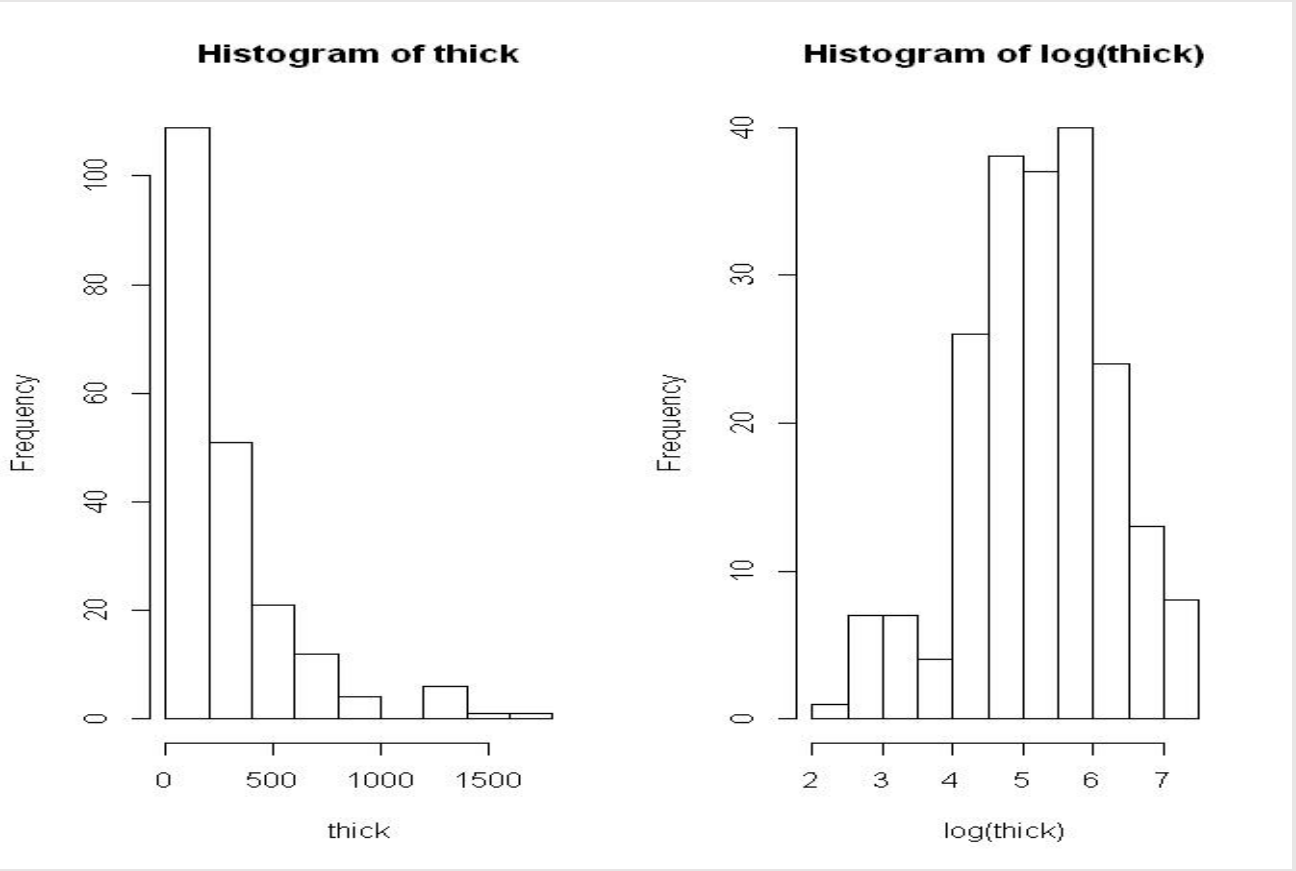


```
##              coef exp(coef) se(coef)     z Pr(>|z|)
## sex         0.4580    1.5809   0.2687 1.705   0.0883 .
## log(thick) 0.7809    2.1834   0.1573 4.963 6.94e-07 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
##             exp(coef) exp(-coef) lower .95 upper .95
## sex             1.581     0.6326    0.9337     2.677
## log(thick)      2.183     0.4580    1.6040     2.972
```

Note that taking into account log(thick) the effect of sex is reduced…

Assessing the PH Assumption (I)

Adjusting for log(thick) does the effect of gender follow a PH model?

If the PH assumption holds, the log cumulative hazards for the two groups, adjusting for log(thick), should be roughly parallel…

Conclusion: not strong evidence of non-PH.

This is a good look at gross departures, but it is far from a formal test…

fit1 <- coxph(Surv(days,status==1) ~ log(thick)+ **strata(sex)**)

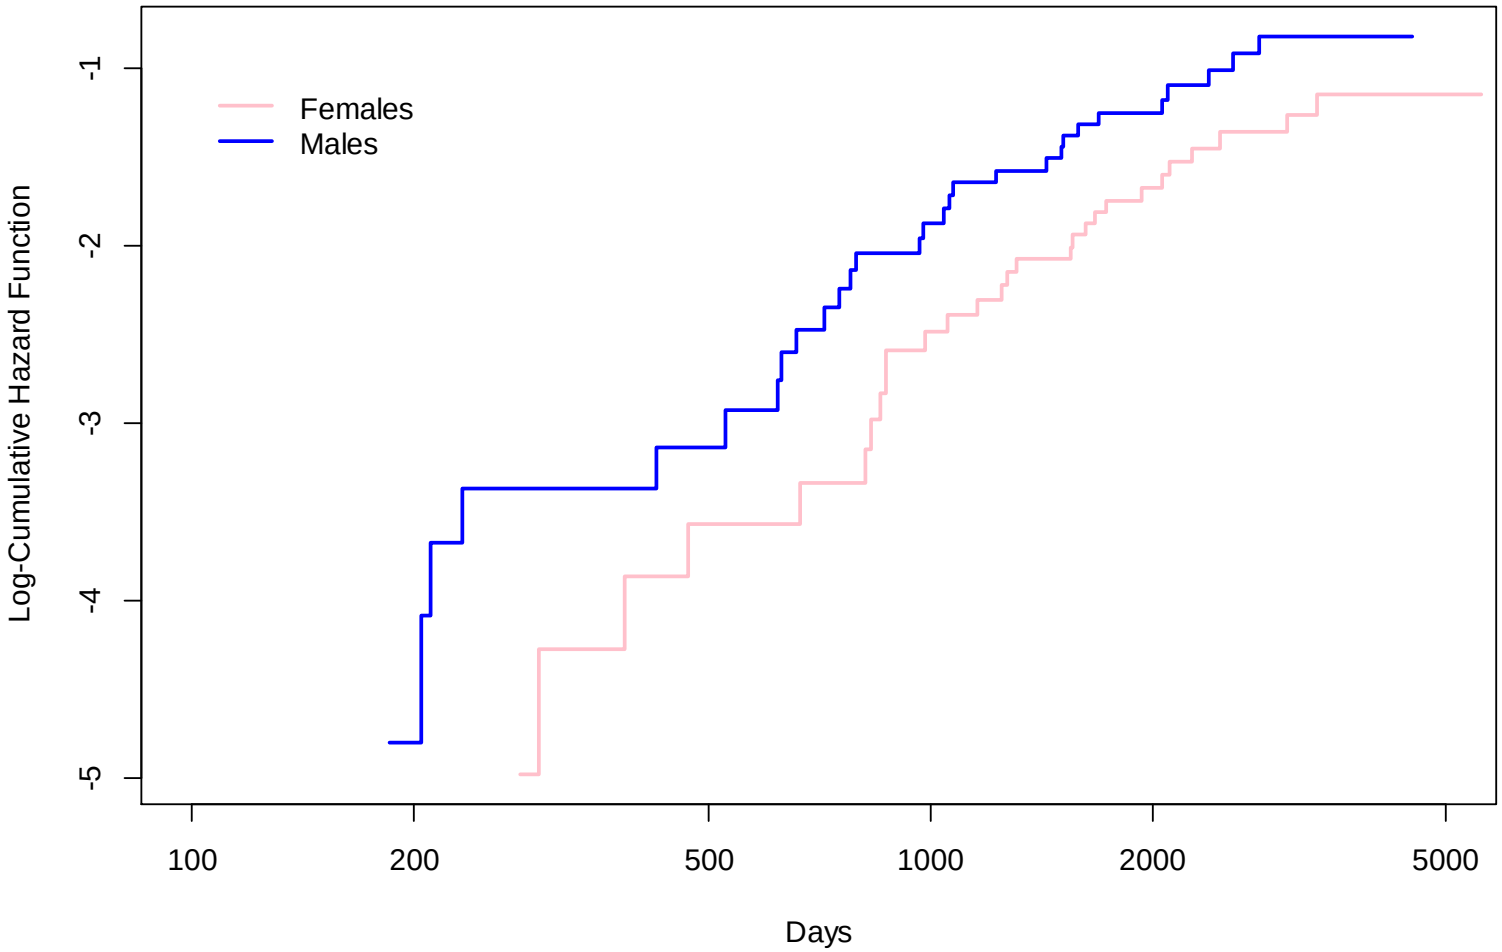

Assessing the PH Assumption (II)

```
check.ph <- cox.zph(mod.cov, transform="km", global=TRUE)
```

Adjusting for gender, does the effect of log(thick) follow a proportional hazards model?

```
##                rho chisq      p
## sex         -0.102 0.587 0.4436
## log(thick) -0.352 5.485 0.0192
## GLOBAL          NA 6.813 0.0332
```

The effect of log(thick) is gradually decreasing with time.

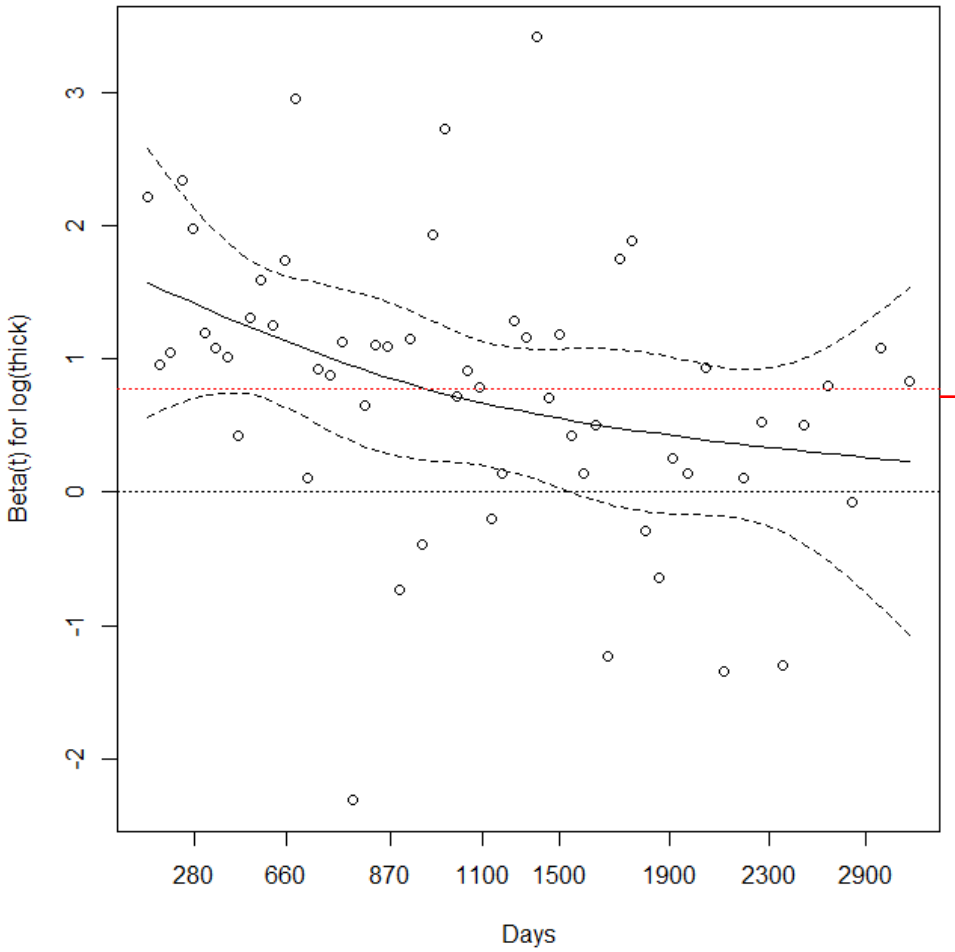

If the PH assumption holds, then the plot of $\beta(t)$ vs time *should be* on a horizontal line.

Cox model's estimate for «overall» log thick effect

*P. Grambsch and T. Therneau (1994) Proportional hazards tests and diagnostics based on weighted residuals. *Biometrika,* **81**, 515-26.

Possible **solutions** to non-proportionality (I):

- **Stratification**: covariates with non PH effects may be used as strata
  - no ***direct test*** of association with survival;
  - ok for ***categorical*** covariates, discretization for continuous ones (could be problematic)
  - ***less efficient*** analyses (usually larger sample size needed)

- **Partition** of the time axis: the PH could be *valid* in some time intervals (***landmark* analysis**)

Alternative methods…(II)

- **Cox model with time-varying** coefficients: model the dependence of beta on time (not easy to find the appropriate function… interpretation more complex)

- Use a **different approach**: **Flexible Parametric Survival and Multi-State Models**

Accelerated Failure Time (AFT) Models:

- The *survreg* function in package survival can fit an accelerated failure time model.
- A modified version of *survreg* is implemented in the rms package (*psm* function).
- The eha package also proposes an implementation of the AFT model (function *aftreg*).
- The NADA package proposes the front end of the *survreg* function for left-censored data.
- The simexaft package implements the Simulation-Extrapolation algorithm for the AFT model, that can be used when covariates are subject to measurement error.
- A robust version of the accelerated failure time model can be found in RobustAFT.
- The coarseDataTools package fits AFT models for interval censored data.
- An alternative weighting scheme for parameter estimation in the AFT model is proposed in the imputeYn package.
- The AdapEnetClass package implements elastic net regularisation for the AFT model.

Additive Models:

- Both survival and timereg fit the additive hazards model of Aalen in functions *aareg* and *aalen*, respectively.
- timereg also proposes an implementation of the Cox-Aalen model (that can also be used to perform the Lin, Wei and Ying (1994) goodness-of-fit for Cox regression models) and the partly parametric additive risk model of McKeague and Sasieni.
- A version of the Cox-Aalen model for interval censored data is available in the coxinterval package.
- The uniah package fits shape-restricted additive hazards models.
- The addhazard package contains tools to fit additive hazards model to random sampling, two-phase sampling and two-phase sampling with auxiliary information.

Flexible survival models:

- **flexsurv: Flexible parametric models for time-to-event data**
- **rstpm2: Smooth Survival Models, Including Generalized Survival Models**

A more elaborate example: binary factor + continuous covariate + **stratification** variable:

```
mod.cov.strat <- coxph(Surv(days,status==1)~sex+log(thick)+strata(ulc))
summary(mod.cov.strat)
```

```
##                coef exp(coef) se(coef)     z Pr(>|z|)
## sex          0.3600    1.4333   0.2702 1.332   0.1828
## log(thick) 0.5599    1.7505   0.1784 3.139   0.0017 **
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
##              exp(coef) exp(-coef) lower .95 upper .95
## sex              1.433     0.6977     0.844     2.434
## log(thick)       1.750     0.5713     1.234     2.483
##
## Concordance= 0.673  (se = 0.058 )
## Rsquare= 0.063   (max possible= 0.9 )
## Likelihood ratio test= 13.3  on 2 df,   p=0.001296
## Wald test            = 12.88  on 2 df,   p=0.001598
## Score (logrank) test = 12.98  on 2 df,   p=0.00152
```

Stratifying by the presence or absence of ulcer, significance of the log(thick) has been reduced and sex is no longer significant.

We can plot survival curves estimated for each strata by using survfit on the output of coxph:

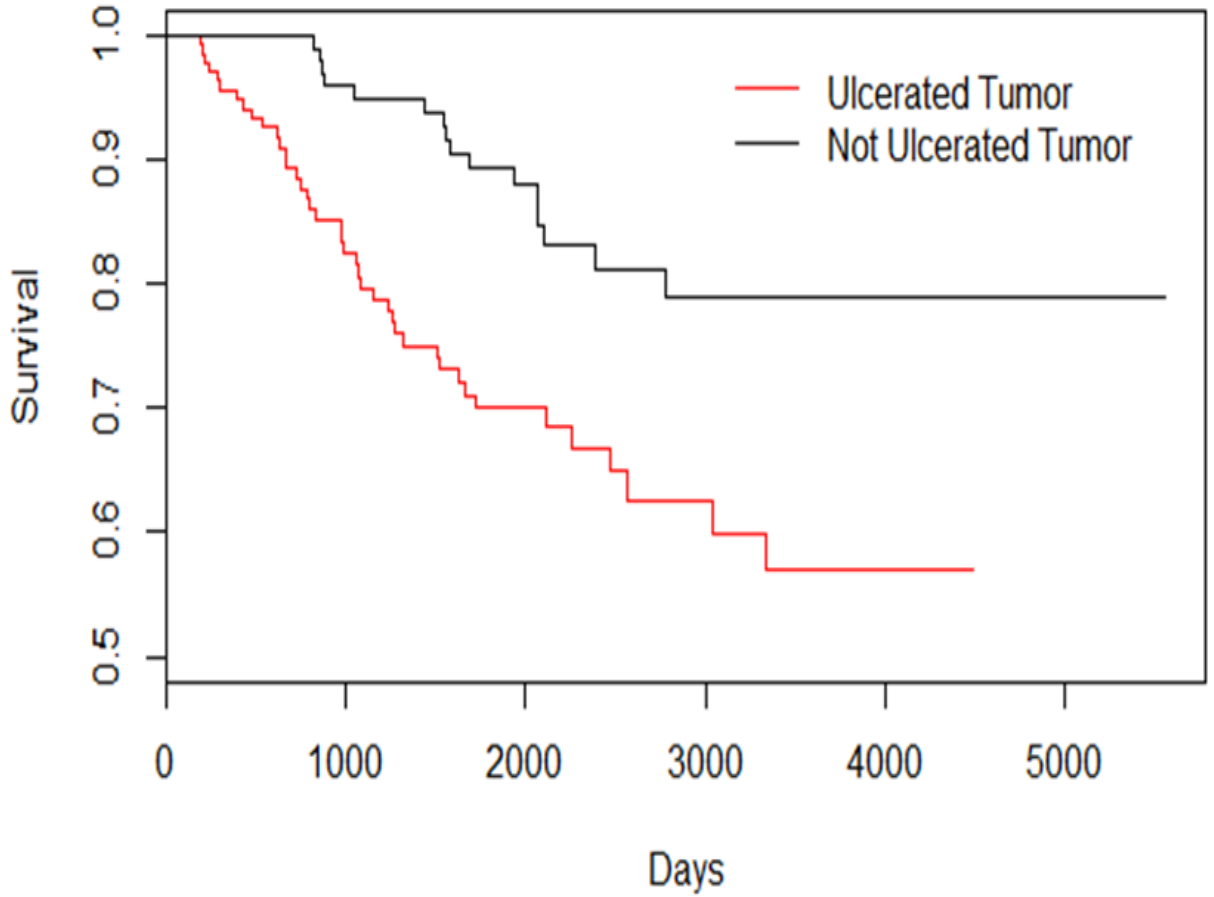

The default for survfit is to generate curves for a *pseudoindividual* for which the covariates are at their mean values.

In the present case, that would correspond to a tumor thickness of 1.86 mm and a gender of 1.39 (!)…

… we have been sloppy in not defining sex as a factor variable, but that would not actually give a different result (HR): coxph subtracts the means of the regressors before fitting, so a 1/2 coding is the same as 0/1, which is what a factor with treatment contrasts usually gives.

[But, defining the factor we can define "hypothetical" pts with certain values for the covariates]

sex.f <- as.factor(sex)

Converting sex into a factor

mod.cov.strat.f <- coxph(Surv(days,status==1)~sex.f+log(thick)+strata(ulc))
summary(mod.cov.strat.f)

```
mod.cov.strat.f <- coxph(Surv(days,status==1)~sex.f+log(thick)+strata(ulc))
summary(mod.cov.strat.f)
```

```
## Call:
## coxph(formula = Surv(days, status == 1) ~ sex.f + log(thick) +
##     strata(ulc))
##
##   n= 205, number of events= 57
##
##               coef exp(coef) se(coef)     z Pr(>|z|)
## sex.f2     0.3600    1.4333   0.2702 1.332   0.1828
## log(thick) 0.5599    1.7505   0.1784 3.139   0.0017 **
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Now sex.f2 indicates that HR refers to the contrast of level "2" versus level "1" for the factor variable sex, [the same HR value as before]

To estimate survival curves for subjects with *certain* values of the covariates, we could use the option newdata in survfit:

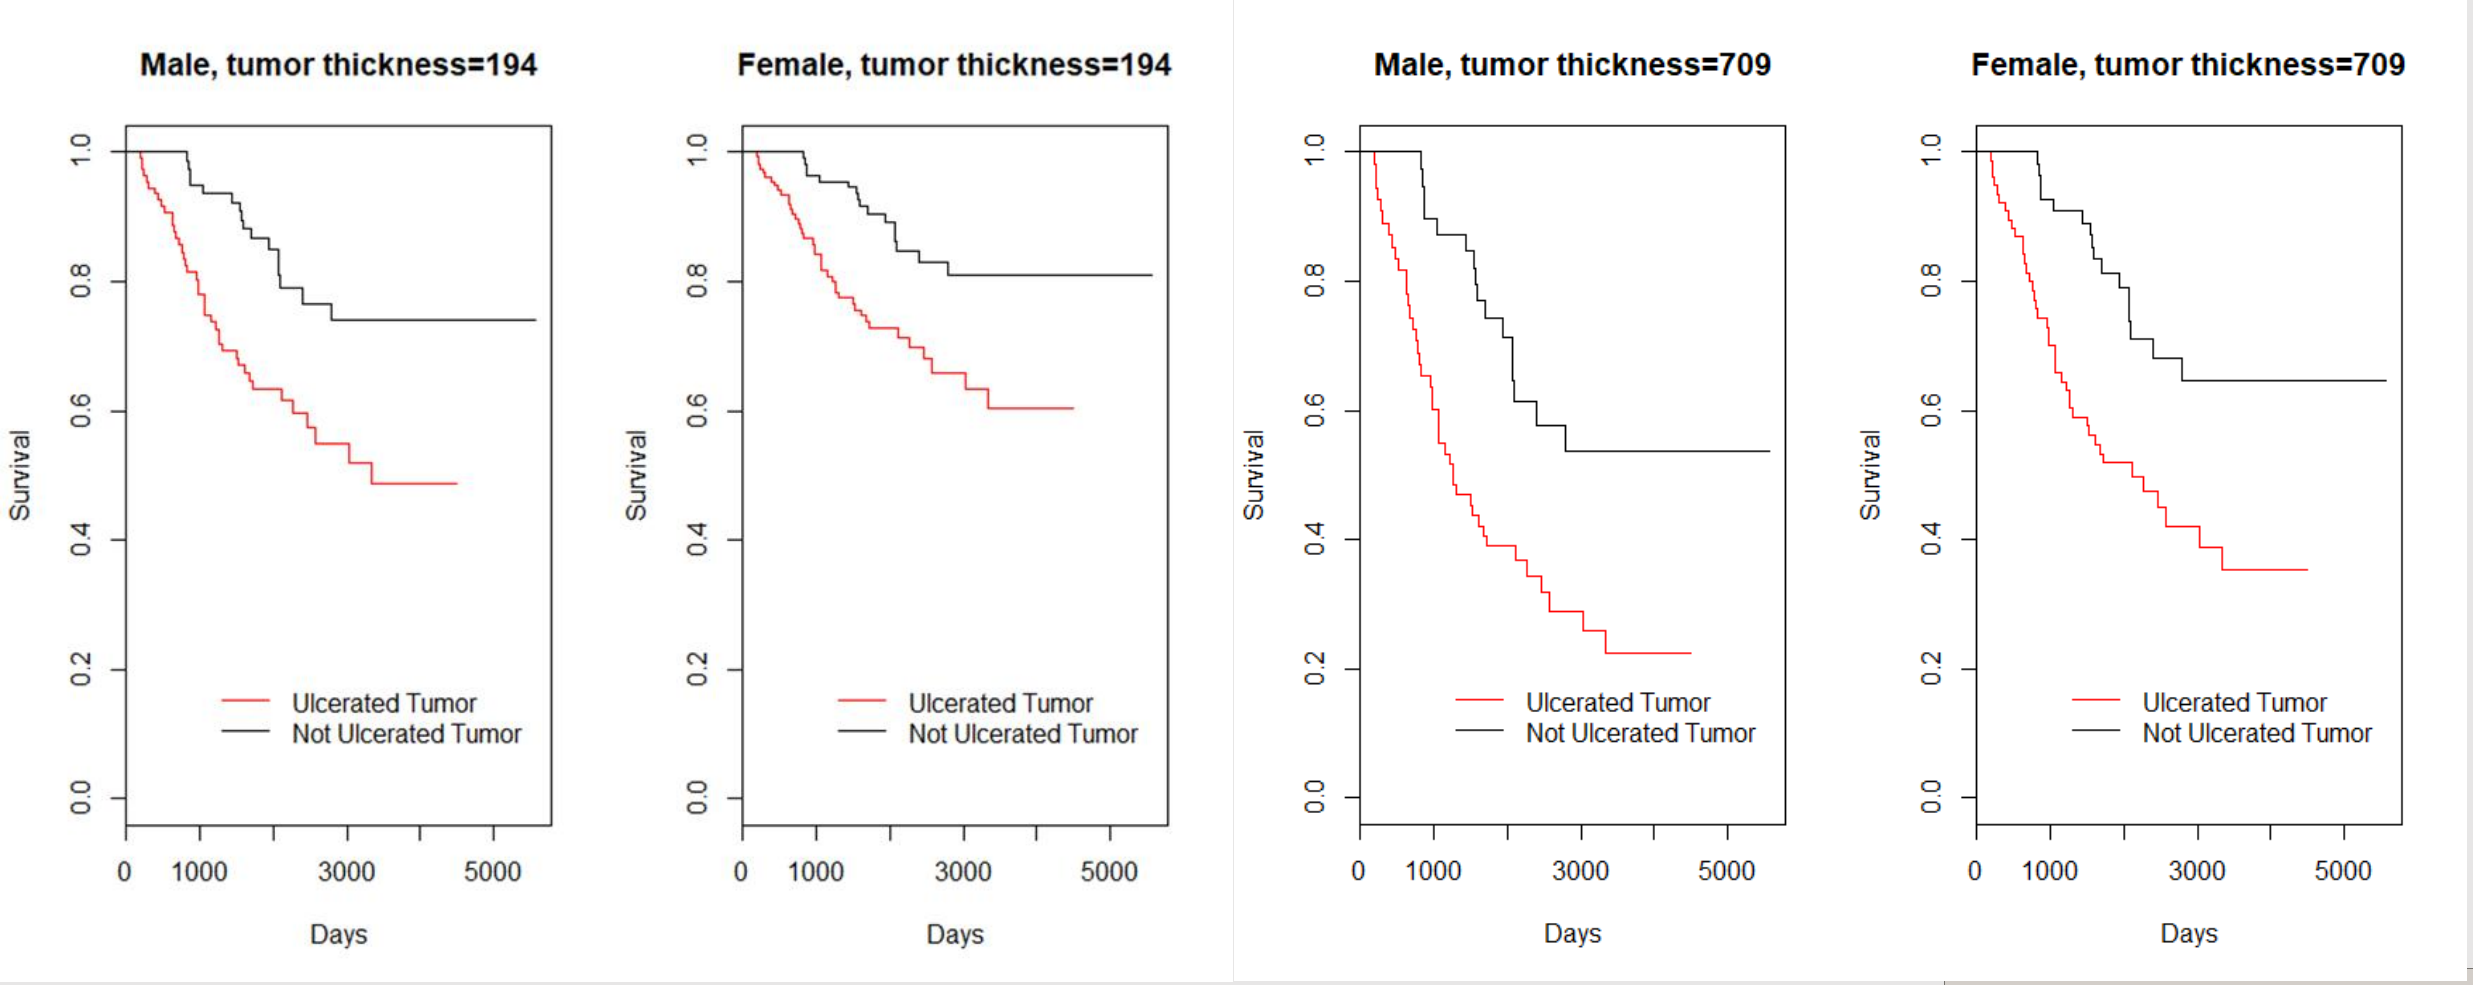

Summary: basic assumptions (all standard methods, KM, log rank & basic Cox):

1. Events of the individuals occur **independently** of one another

Acceptable in «***time to the first event***» analyses

2. Hazard of event *at any given time* for an individual in one group is **proportional** to the hazard *at that time* for an individual in the other group…

hazard functions **do not cross** one another

3. Hazard ratios **are independent of time**

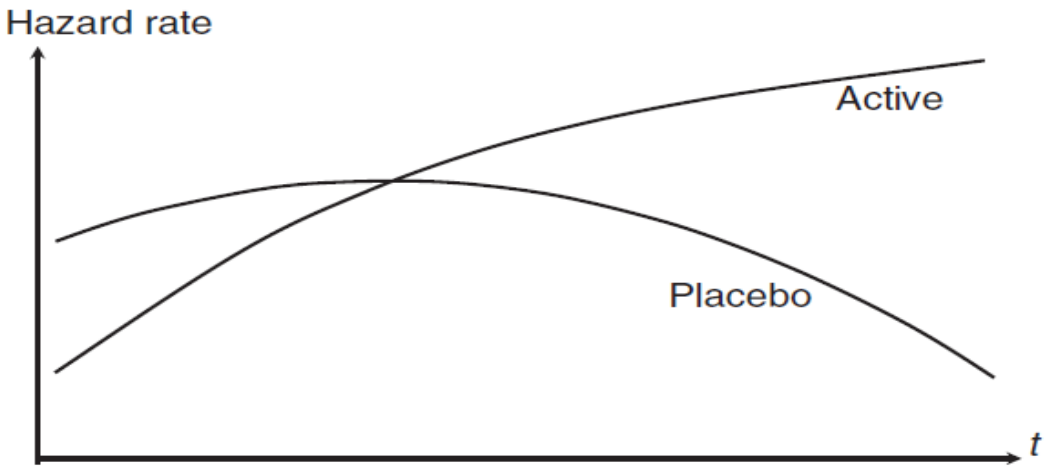

what if the '**treatment' effect changes with time*** ?

***…or we have repeated measures of a covariate ???**

UNITÀ DI BIOSTATISTICA
Dipartimento Universitario Clinico di Scienze Mediche Chirurgiche e della Salute

Last (but not least!):

4. Censoring mechanism is ***independent*** of the event [***conditional*** on covariates in Cox]:

Those still at risk at time *t* are **a random sample** of the population at risk at time *t*, for all *t*…

…is that always true???

Primary and Secondary end point:

Table 2. Efficacy Outcomes.*

Outcome	Apixaban Group (N=9120)		Warfarin Group (N=9081)		Hazard Ratio (95% CI)	P Value
	Patients with Event	Event Rate	Patients with Event	Event Rate		
	no.	%/yr	*no.*	%/yr		
Primary outcome: stroke or systemic embolism	212	1.27	265	1.60	0.79 (0.66–0.95)	0.01
Stroke	199	1.19	250	1.51	0.79 (0.65–0.95)	0.01
Ischemic or uncertain type of stroke **	162	0.97	175	1.05	0.92 (0.74–1.13)	0.42
Hemorrhagic stroke	40	0.24	78	0.47	0.51 (0.35–0.75)	<0.001
Systemic embolism	15	0.09	17	0.10	0.87 (0.44–1.75)	0.70
Key secondary efficacy outcome: death from any cause	603	3.52	669	3.94	0.89 (0.80–0.998)	0.047

Are patients that die **before** experiencing the primary outcome *similar* to the others?

Regression Models and Life-Tables

By D. R. Cox

Imperial College, London

[Read before the Royal Statistical Society, at a meeting organized by the Research Section, on Wednesday, March 8th, 1972, Mr M. J. R. Healy in the Chair]

Summary

The analysis of censored failure times is considered. It is assumed that on each individual are available values of one or more explanatory variables. The hazard function (age-specific failure rate) is taken to be a function of the explanatory variables and unknown regression coefficients multiplied by an arbitrary and unknown function of time. A conditional likelihood is obtained, leading to inferences about the unknown regression coefficients. Some generalizations are outlined.

David Cox (1924-2022**)**

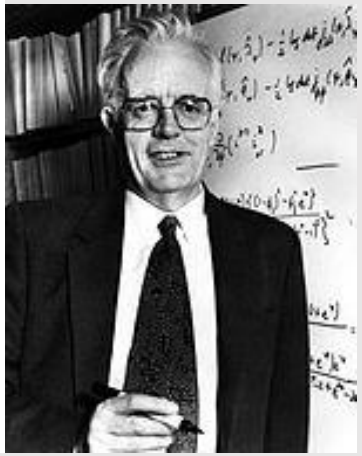

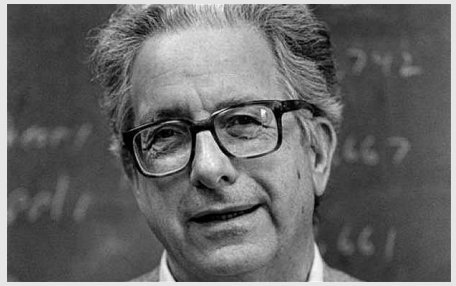

Paul Meier (1924-2011)

Edward L. Kaplan (1920-2006)

NONPARAMETRIC ESTIMATION FROM INCOMPLETE OBSERVATIONS*

E. L. Kaplan

University of California Radiation Laboratory

AND

Paul Meier

University of Chicago

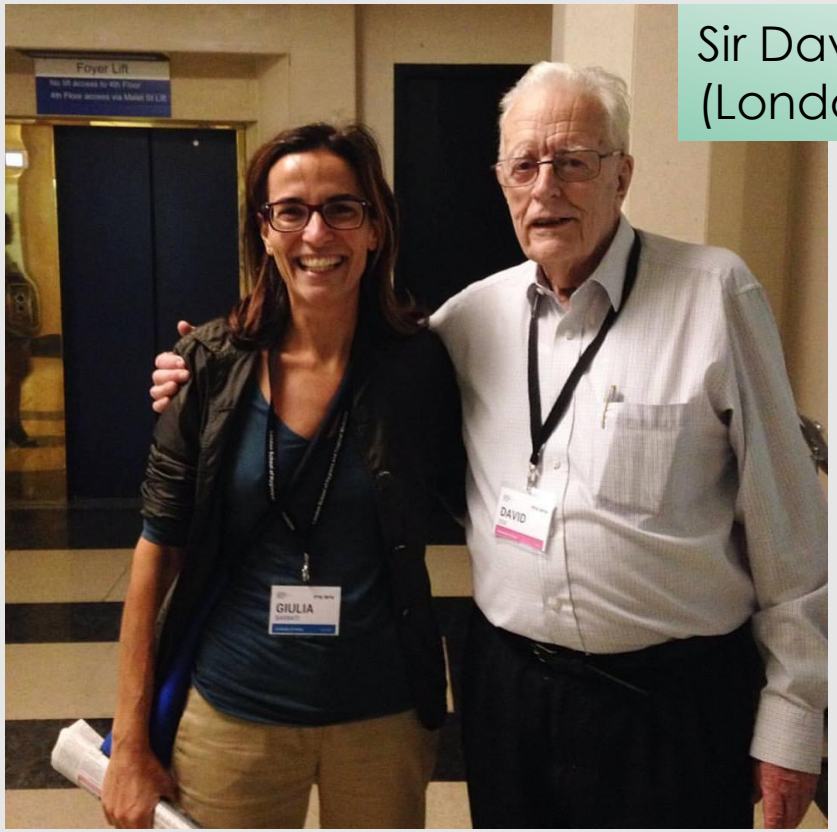

Sir David Cox and me (London, sept. 2016)

Remembering Sir David Cox, 1924–2022

Sir David Cox died on 18 January 2022 at the age of 97. News of his passing was met with an outpouring of tributes. To the Royal Statistical Society, he was “one of the most important statisticians of the past century”. At Nuffield College, Oxford, he was hailed as “a pioneering statistician”. The MRC Biostatistics Unit at Cambridge called him “a giant in the field”, while at St John’s College, Cambridge, he was celebrated as “an inspiring scholar”. In this special collection of articles, friends and colleagues remember Sir David in their own way, while also reflecting on his immense contributions to statistics