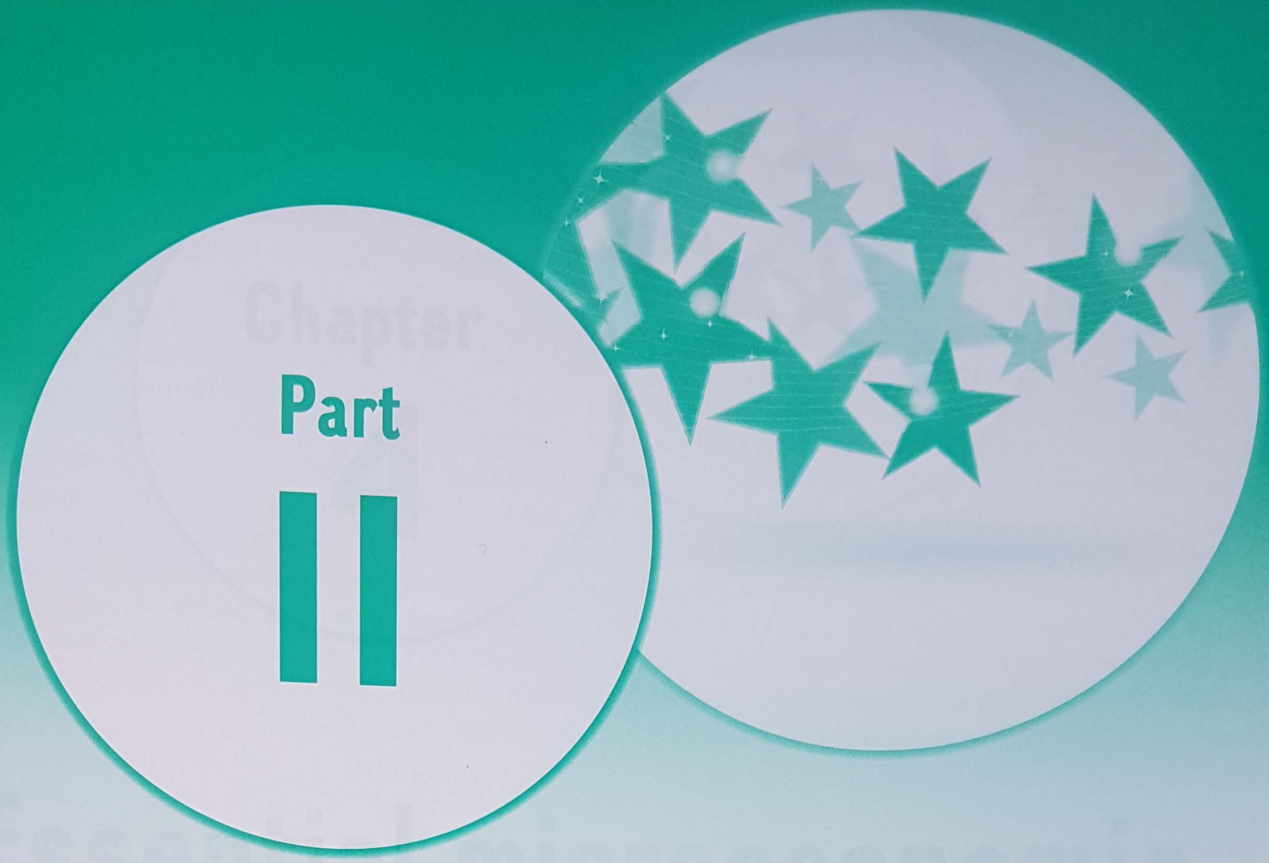

Part

II

The Microeconomics of European Integration

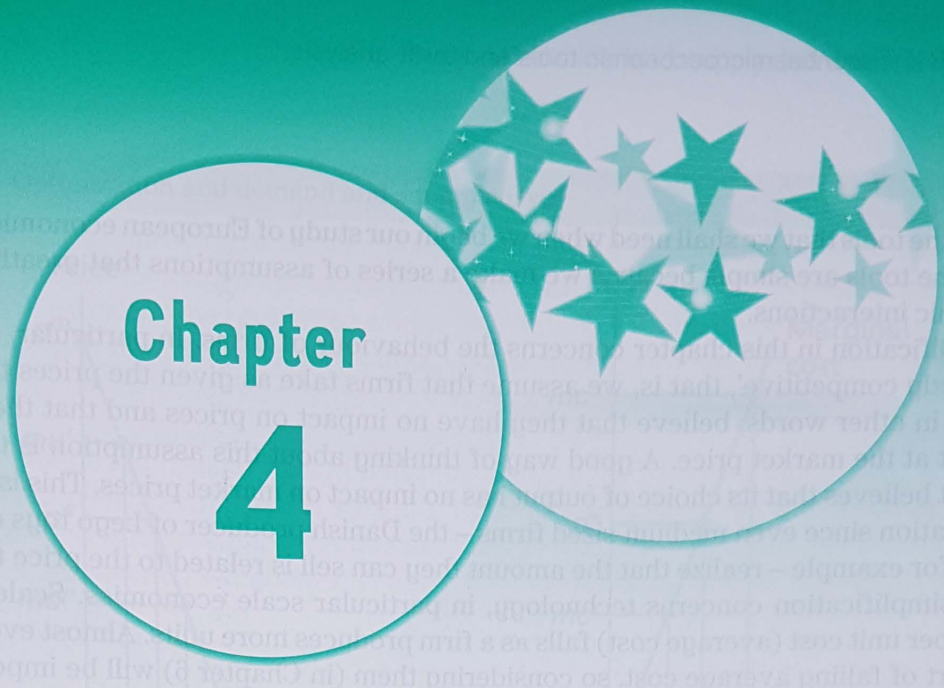

Chapter

4

Essential microeconomic tools and tariff analysis

Chapter Contents

Introduction

This chapter presents the tools that we shall need when we begin our study of European economic integration in the next chapter. The tools are simple because we make a series of assumptions that greatly reduce the complexity of economic interactions.

The primary simplification in this chapter concerns the behaviour of firms. In particular, all firms are assumed to be 'perfectly competitive', that is, we assume that firms take as given the prices they observe in the market. Firms, in other words, believe that they have no impact on prices and that they could sell as much as they want at the market price. A good way of thinking about this assumption is to view each firm as so small that it believes that its choice of output has no impact on market prices. This is obviously a very rough approximation since even medium-sized firms – the Danish producer of Lego toys or the Dutch brewer of Heineken, for example – realize that the amount they can sell is related to the price they charge.

The second key simplification concerns technology, in particular scale economies. Scale economies refer to the way that per unit cost (average cost) falls as a firm produces more units. Almost every industry is subject to some sort of falling average cost, so considering them (in Chapter 6) will be important, but a great deal of simplification can be gained by ignoring them. This simplification, in turn, allows us to master the essentials before adding in more complexity in subsequent chapters.

4.1 Preliminaries I: supply and demand diagrams

Assessing many economic aspects of European integration is made clearer with the help of a simple yet flexible diagram with which to determine the price and volume of imports, as well as the level of domestic consumption and production. The diagram we use – the 'import supply and import demand diagram' – is based on straightforward supply and demand analysis. But to begin from the beginning, we quickly review where demand and supply curves come from. Note that this section assumes that readers have had some exposure to supply and demand analysis; our treatment is intended as a review rather than an introduction. Readers who find it too brief should consult an introductory economics textbook.

Well-prepared readers may want to skip this section, moving straight on to Section 4.2.

4.1.1 Demand curves and marginal utility

A demand curve shows how much consumers would buy of a particular good at any particular price. Generally speaking, consumers strive to spend their money in a way that makes them better off. Their demand curve is thus based on some sort of economic calculation.

To see this, the left-hand panel of Figure 4.1 plots the 'marginal utility' curve for a typical consumer. But what do 'utility' and 'marginal' mean in this context? Utility means nothing more and nothing less than 'happiness', and we measured the happiness in euros. Money sounds like a shallow measure of happiness, but we are talking about the happiness people get from consuming goods, like a cappuccino or a bottle of fresh-squeezed orange juice. For such things, we are weighing the cost of buying the thing versus the money we have to give up to get it, so the money is a natural – if not perfect – measure of the happiness the good gives us. The word 'marginal' is here used to mean nothing more and nothing less than 'one more'. Putting together, 'marginal utility' means the money-value of consuming one more cappuccino.

For example, if we are considering the demand for cups of coffee, the marginal utility curve shows how much extra joy a consumer gets from having one more cup starting from any given number of cups already consumed. Typically the extra joy from an extra cup falls with the number of cups bought per day. For example, if the consumer buys very few cups of coffee today, say c' in the diagram, the gain from buying an extra one is likely to be pretty high, for example mu' in the diagram. If, however, the consumer has already bought lots of cups already, then the gain from one more is likely to be much lower. This is shown by the pair, c'' and mu''.

This marginal utility curve allows us to work out how much the consumer would buy at any given price. Suppose the consumer could buy as many cups as she likes at the price p^*. How many would she buy? If the consumer is wise, and we assume she is, she will buy cups of coffee up to the point where the last one bought is just barely worth the price.

Figure 4.1 Optimization and demand and supply curves

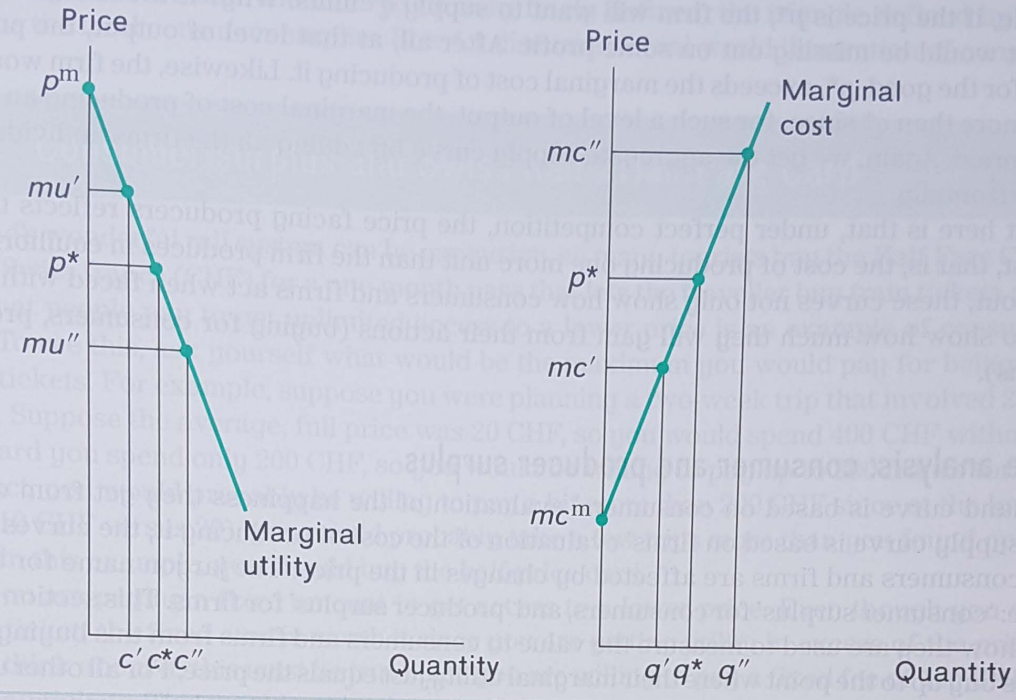

In the diagram, this level of purchase is given by c^* since the extra benefit (marginal utility) from buying an extra cup exceeds the cost of doing so (the price) for all levels of purchase up to c^*. At this point, the consumer finds that additional cups would not be worth the price. For example, the marginal utility from buying c^* plus one cups of coffee would be below p^*.

This is the demand curve for one individual. When we want to know how many cups will be bought in a particular market, we add all consumers' individual marginal utility curves horizontally. This is obvious once you think about it. If the price is p^* and there are 100 identical consumers, market demand will be 100 times c^*. And a similar calculation holds for all prices. In particular, at p^m, no one will buy coffee, so individual and group demand is zero.

A key point to retain from this is that the price that consumers face reflects the marginal utility of consuming a little more.

4.1.2 Supply curves and marginal costs

Derivation of the supply curve follows a similar logic, but here the optimization is done by firms. The right-hand panel of Figure 4.1 shows the 'marginal cost' curve facing a typical firm (assume they are all identical for the sake of simplicity). As before, marginal means 'one more' and 'cost' means cost. Thus marginal cost is the extra cost involved in making one more unit of the good.

While the marginal cost of production in the real world often declines with the scale of production, allowing for this involves consideration of scale economies and these, in turn, introduce a whole range of complicating factors that would merely clutter the analysis at this stage. To keep it simple, we assume that firms are operating at a point where the marginal cost is upward sloping; that is, that the cost of producing an extra unit rises as the total number of units produced rises. The curve in the diagram shows, for example, that it costs mc' to produce one more unit when the production level (e.g. the number of cups of coffee per day) is q'. This is less than the cost, mc'', of producing an extra unit when the firm is producing q'' units per year.

Using this curve we can determine the firm's supply behaviour. Presuming that the firm wants to make the most money possible from selling coffee (or as economists put it, they want to maximize profit), the firm will supply goods up to the point where the marginal cost just equals the price. With a little reflection,

you will see this is the correct answer to the question: How many cups of coffee should the firm produce and sell at the price p^* in order to make the most money?

For example, if the price is p^*, the firm will want to supply q^* units. Why? If the firm offered one less than q^* units, it would be missing out on some profit. After all, at that level of output, the price the firm would receive for the good, p^*, exceeds the marginal cost of producing it. Likewise, the firm would not want to supply any more than q^* since, for such a level of output, the marginal cost of producing an extra unit is more than the price. Again, we get the aggregate supply curve by adding all the firms' individual marginal cost curves horizontally.

A key point here is that, under perfect competition, the price facing producers reflects the marginal production cost, that is, the cost of producing one more unit than the firm produces in equilibrium.

As it turns out, these curves not only show how consumers and firms act when faced with a particular price, they also show how much they will gain from their actions (buying for consumers, producing and selling for firms).

4.1.3 Welfare analysis: consumer and producer surplus

Since the demand curve is based on consumers' evaluation of the happiness they get from consuming a good and the supply curve is based on firms' evaluation of the cost of producing it, the curves can be used to show how consumers and firms are affected by changes in the price. The jargon name for the concepts we will use are: 'consumer surplus' for consumers, and 'producer surplus' for firms. This section explains the concepts and how they are used to measure the value to consumers and firms from this buying and selling.

Consumers buy up to the point where their marginal utility just equals the price. For all other units bought, the marginal utility exceeds the price. This means that the consumer gets what is known as 'consumer surplus' from buying c^* units at price p^* (see Figure 4.2). In plain English, this says that consumers get more (in terms of utility) than they pay for. How much more?

For the first unit bought, the marginal unit was mu' but the price paid was only p^*, so the surplus is the area shown by the rectangle 'a'. For the second unit, the marginal utility was somewhat lower

Figure 4.2 Deriving consumer and producer surplus

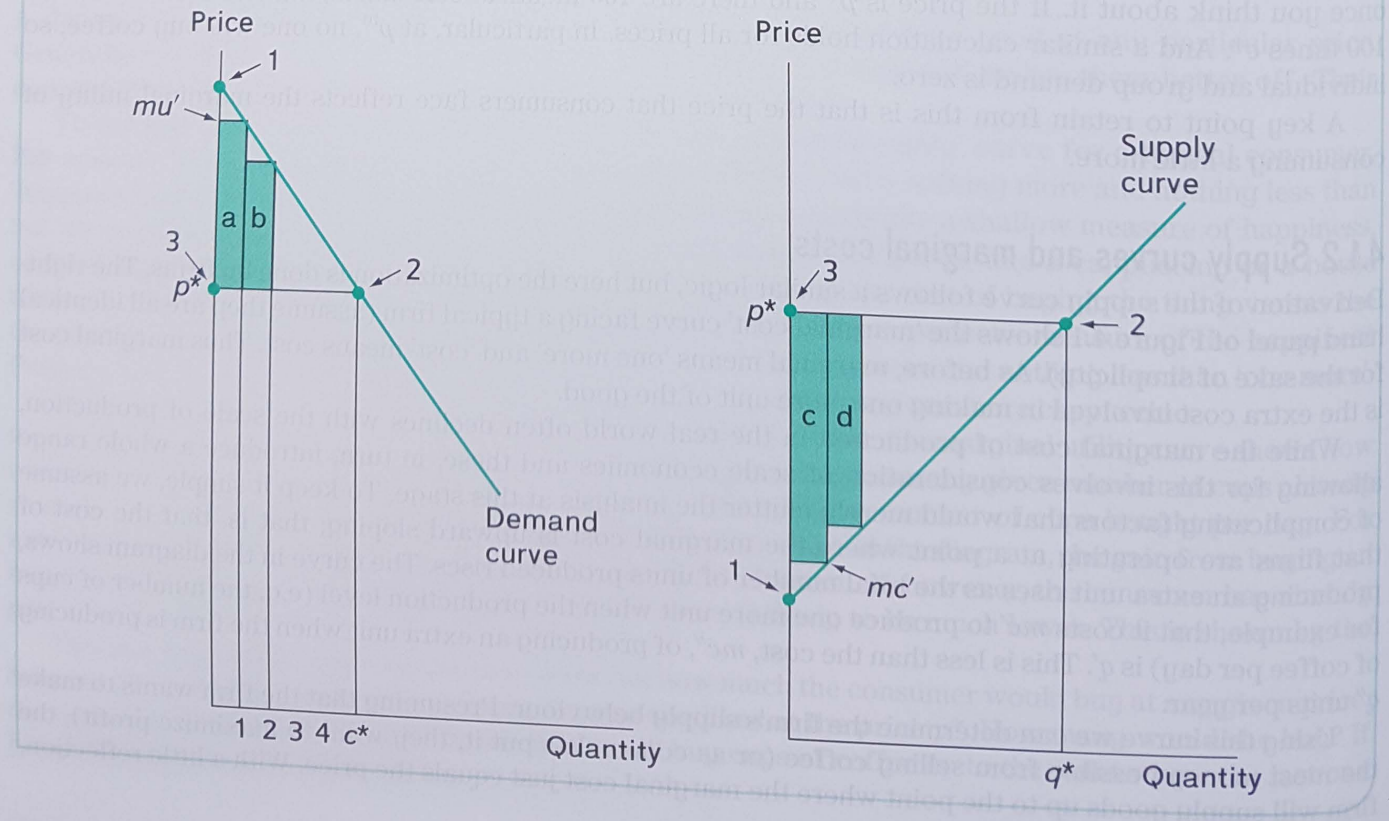

(not shown in the diagram), so the surplus is lower; specifically, it is given by the area 'b'. Doing the same for all units shows that buying c^* units at p^* yields a total consumer surplus equal to the sum of all the resulting rectangles. If we take the units to be very finely defined, the triangle defined by the points 1, 2 and 3 gives us the total consumer surplus. Box 4.1 discusses a real-world illustration of consumer surplus.

Box 4.1 Consumer surplus and Swiss Rail's Half Fare Card

Switzerland's wonderful rail system can be expensive, so many tourists buy the Half Fare Card; in 2018, it cost 120 Swiss francs (CHF) for a one-month pass that lets the traveller buy train tickets at half price. The fact that people pay to get unlimited access to a lower price is an example of consumer surplus in action. To see this, ask yourself what would be the maximum you would pay for being able to buy half-price tickets. For example, suppose you were planning a two-week trip that involved 20 individual train trips. Suppose the average, full price was 20 CHF, so you would spend 400 CHF without the card. With the card you spend only 200 CHF, so you would be willing to pay up to 200 francs for a Half Fare Card. In fact, you would probably be willing to pay a bit more than 200 CHF since at the lower per-trip price (i.e. 10 CHF versus 20), you would probably take a few trips more than you found optimal at the full price. In this example, you should buy the half-price card.

Here you are paying a fixed amount to get access to a lower price. Even though you still have to buy every ticket, the fixed sum is worth it since you are getting utility in excess of the price you pay. When you think about it, this must be true since you are willing to pay a fixed fee of 120 CHF to be able to buy cheap tickets. This would not be the case if there were not consumer surplus.

An analogous line of reasoning shows us that the triangle formed by points 1, 2 and 3 in the right-hand panel gives us a measure of the gain firms get from being able to sell q^* units at a price of p^*. Consider the first unit sold. The marginal cost of producing this unit was mc' but this was sold for p^* so the firm earns a surplus, what we call the 'producer surplus', equal to the rectangle 'c' in the right-hand panel. Doing the same exercise for each unit sold shows that the total producer surplus is equal to the triangle defined by points 1, 2 and 3.

By drawing similar diagrams on your own, you should be able to convince yourself that a price rise increases producer surplus and decreases consumer surplus. A price drop does the opposite.

4.2 Preliminaries II: introduction to open-economy supply and demand analysis

This section introduces the 'workhorse' diagram – the open-economy supply and demand diagram – that is essential to our study of European economic integration. Well-prepared readers may consider skipping, moving straight on to the tariff analysis in Section 4.3. The diagram, however, is used throughout this chapter and the next, so even advanced readers may wish to briefly review the diagram's foundations; if nothing else, such a review will help with the terminology.

4.2.1 The import demand curve

We first look at where the import demand curve comes from. Figure 4.3 facilitates the analysis.

The left-hand panel of the diagram depicts a nation's supply and demand curves for a particular good. As usual, the domestic price is on the vertical axis; quantity is on the horizontal axis. If imports of the good were banned for some reason, the nation would only be able to consume as much as it produced. The result would be a market price of P^* since this is the price where the amount that consumers want to buy just matches the amount that firms want to produce. Plainly, import demand is zero at P^* (for simplicity, we assume that imported and domestic goods are perfect substitutes). This zero-import point is marked in the right-hand panel as point 1; this diagram has the same price on the vertical axis, but plots imports on the horizontal axis.

Figure 4.3 Deriving the import demand curve and welfare changes

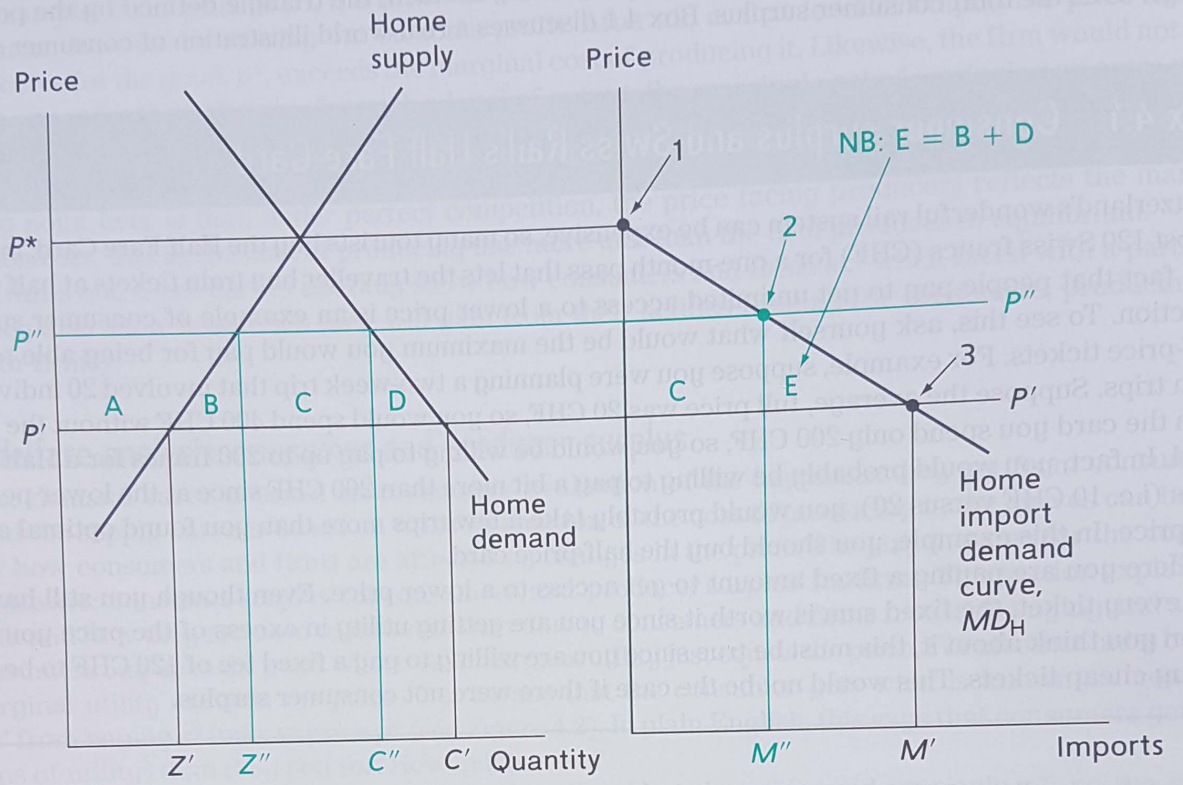

Note: Readers who find these diagrams complicated may benefit from the step-by-step explanations given in the interactive PowerPoint presentation available on the companion website: http://www.mheducation.co.uk/textbooks/baldwin/.

How much would the nation import if the price were lower, say, P'? The first thing to note is that the import price will fix the domestic price. Imports are always available at P', so no consumer would pay more than P'. Of course, domestic producers must match the import competition, so P' becomes the domestic price.

The second thing to note is the impact of P' on consumption, production and imports. Consumption demand would be C' and domestic production would be Z'. As C' exceeds Z', consumers buy more than domestic firms are willing to produce at P'. The 'excess' demand is met by imports. That is to say, imports are the difference between C' and Z' (in symbols, $M' = C' - Z'$).

For convenience, we can show the level of imports that corresponds to P' with a diagram that has imports on the horizontal axis and price on the vertical axis (this is the diagram on the right side of Figure 4.3). In particular, we plot the combination of import demand is M' at the price at P' as point 3 in the right-hand panel of the diagram. Performing the same exercise for P'' yields point 2, and doing the same for every possible import price yields the import demand curve, that is, the amount of imports that the nation wants at any given domestic price. The resulting curve is shown as MD_H in the right-hand panel. (For convenience, we often call the nation under study 'Home' to distinguish it from its trade partner, which we call 'Foreign'.)

Welfare analysis: MD curves as the marginal benefit of imports

When studying European economic integration, a critical question that arises time and again is the extent to which a policy raises or lowers nations' well-being. As before, we will use money to measure well-being – keeping in mind the usual caveats about this being a rather shallow measure of happiness in the broad sense, but a good measure when it comes to pragmatic things like the price of goods.

As it turns out, it is useful to know how to carry out welfare analysis with the right- and left-hand diagrams in Figure 4.3). Consider a rise in the import price from P' to P''. As argued above, the higher import price means the domestic price rises by the same amount. The corresponding level of imports drops to M'', since consumption drops to C'' and production rises to Z''. We can see the welfare analysis in the left-hand

panel using the standard notions of consumer and producer surpluses (see Section 4.1). Specifically, the price rise lowers consumer surplus by A + B + C + D. The same price rise increases producer surplus by A. The right-hand panel shows how this appears in the import demand diagram. From the left-hand panel, the import price rise means a net loss to the country of B + C + D, since the area A cancels out (area A is a gain to Home producers and a loss to Home consumers). In the right-hand panel, these changes are shown as areas C and E; as it turns out, area E equals area B + D.

A powerful perspective: trade volume effects and border price effects

It proves insightful to realize that the MD_H curve shows the marginal benefit of imports to Home. Before explaining why this is true, we show that it is a useful insight. As we saw above, Home loses areas C and E when the price of imports rises from P' to P''. Area C is easy to understand. After the price rise, Home pays more for the units it imported at the old price. Area C is the size of this loss. (Say the price rise was €1.2 per unit and M'' was 100; the loss would be €1.2 times 100; geometrically, this is the area C since a rectangle's area is its height times its base.) Understanding area E is where the insight comes in handy. Home reduces its imports at the new price and area E measures how much it loses from the drop in imports. The marginal value of the first lost unit is the height of the MD_H curve at M''. But since Home had to pay P' for this unit, the net loss is the gap between P' and the MD_H curve. If we add up the gaps for all the extra units imported, we get the area E. The jargon terms for these areas are the 'border price effect' (area C) and the 'import volume effect' (area E).

To understand why MD_H is the marginal benefit of imports, we use three facts and one bit of logic: (1) the MD_H curve is the difference between the domestic demand curve and the domestic supply curve; (2) the domestic supply curve is the domestic marginal cost curve, and the domestic demand curve is the domestic marginal utility curve (see Section 4.1 if these points are unfamiliar); and (3) the difference between domestic marginal utility of consumption and domestic marginal cost of production is the net gain to the nation of producing and consuming one more unit. The logical point is that an extra unit of imports leads to some combination of higher consumption and lower domestic production, and this leads to some combination of higher utility and lower costs; the height of the MD_H curve tells us what that combination is. Or, to put it differently, the nation imports up to the point where the marginal gain from doing so equals the marginal cost. Since the border price is the marginal cost, the border price is also an indication of the marginal benefit of imports.

To see these points in more detail, see the interactive PowerPoint presentations available on this book's Online Learning Centre, http://www.mheducation.co.uk/textbooks/baldwin/.

4.2.2 The export supply curve

Figure 4.4 uses an analogous line of reasoning to derive the import supply schedule. The first thing to keep in mind is that the supply of imports to Home is the supply of exports from foreigners. For simplicity's sake, suppose that there is only one foreign country (simply called 'Foreign' hereafter) and its supply and demand curves look like the left-hand panel of the figure. (Note that the areas in Figure 4.4 are unrelated to the areas in Figure 4.3.)

As with the import demand curve, we start by asking how much Foreign would export for a particular price. For example, how much would it export if the price of its exports was P'? At price P', Foreign firms would produce Z' and Foreign consumers would buy C'. The excess production (equal to $X' = Z' - C'$) would be exported. (Note that, as in the case of import demand, the export price sets the price in Foreign; Foreign firms have no reason to sell for less since they can always export, and competition among Foreign suppliers would prevent any of them from charging Foreign consumers a higher price.) The fact that Foreign would like to export X' when the export price is P' is shown in the right-hand panel at point 2.

As the price for Foreign exports (i.e. Home's import price) rose, Foreign would be willing to supply a higher level of exports for two reasons. The higher price would induce Foreign firms to produce more and Foreign consumers to buy less. For example, the price P'' would bring forth an import supply equal to X'' (this equals $Z'' - C''$); this is shown as point 3 in the right-hand panel. At price P^*, exports are zero. Plotting all such combinations in the right-hand panel produces the export supply curve XS_F. We stress again the simple but critical point that the Foreign export supply is the Home import supply, thus we also label XS_F as MS_H.

Figure 4.4 Deriving the export supply curve and welfare changes

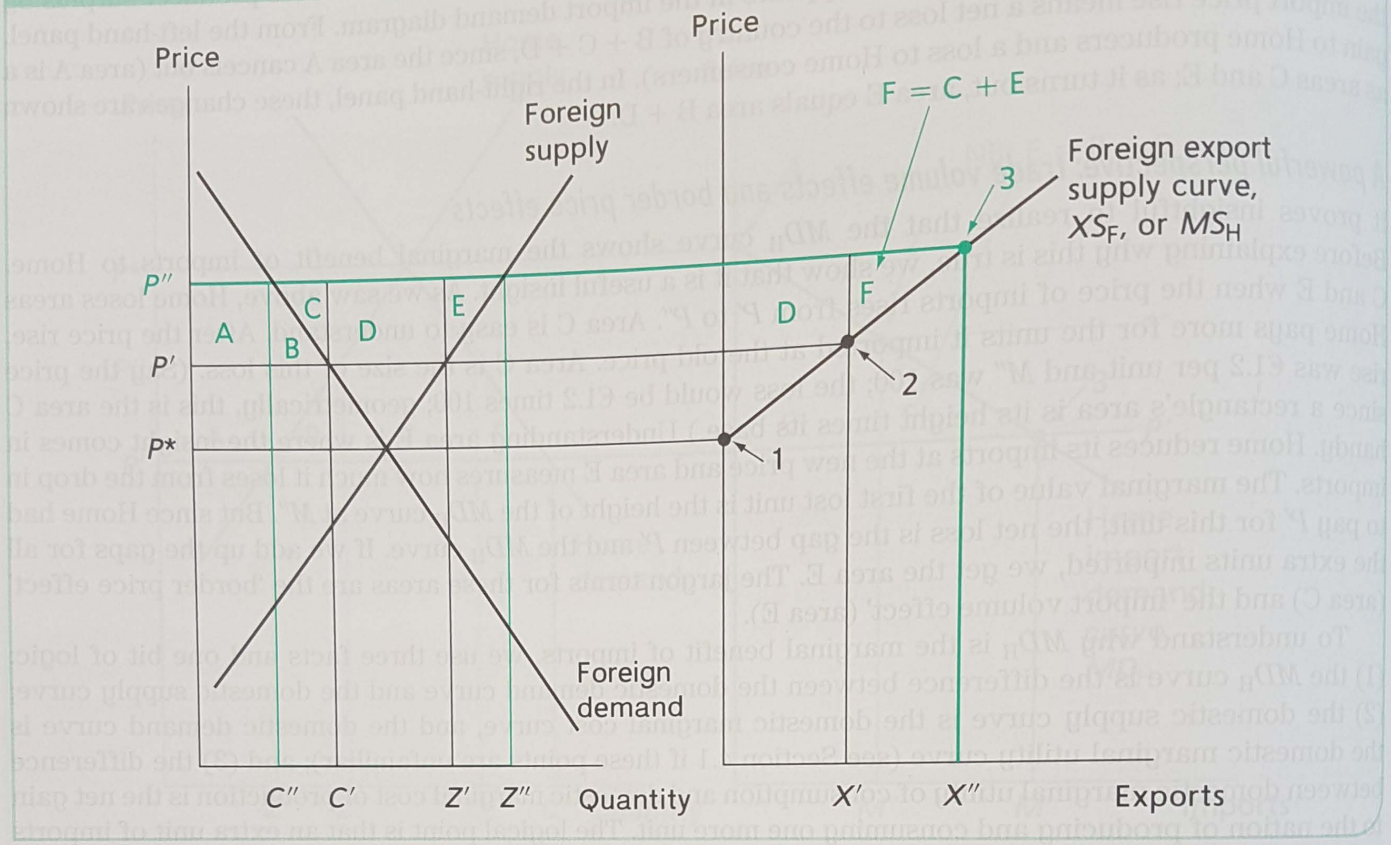

Welfare

The left-hand panel also shows how price changes translate into Foreign welfare changes. If the export price rises from P' to P'', consumers in Foreign lose by A + B (these letters are not related to those in the previous figure), but the Foreign firms gain producer surplus equal to A + B + C + D + E. The net gain is therefore C + D + E. Using the export supply curve XS_F, we can show the same net welfare change in the right-hand panel as the area D + F. Note that the insight from the MD_H curve extends to the XS_F curve, that is, the XS_F curve gives the marginal benefit to Foreign of exporting.

This review of import supply and demand was very rapid – probably too rapid for students who have never used such diagrams and probably too slow for students who have. For those who find themselves in the first category, interactive PowerPoint presentations that go over the diagram in greater detail are available at www.mheducation.co.uk/textbooks/baldwin.

4.2.3 The workhorse diagram: *MD–MS*

The big payoff from having an import supply curve and an import demand curve is that it permits us to find the price and quantity of imports that is produced by the market, in other words the outcome of the interaction of supply and demand in both Home and Foreign. The price is found by putting together import demand and supply as shown in the left-hand panel of Figure 4.5; we drop the 'H' and 'F' subscripts for convenience.

Assuming imported and domestically produced goods are identical, the domestic price is set at the point where the demand and supply of imports meet, namely, P_{FT} (FT stands for free trade). While the import supply and demand diagram, or *MD–MS* diagram for short, is handy for determining the price and volume of imports, it does not permit us to see the impact of price changes on domestic consumers and firms separately. This is where the right-hand panel becomes useful. In particular, we know that the market clears only when the price is P_{FT}, so we know that Home production equals Z and Home consumption equals C. The equilibrium level of imports may be read off either panel. In the left-hand panel, it is shown directly; in the right-hand one, it is the difference between domestic consumption and production.

Figure 4.5 The *MD–MS* and open-economy supply and demand diagrams

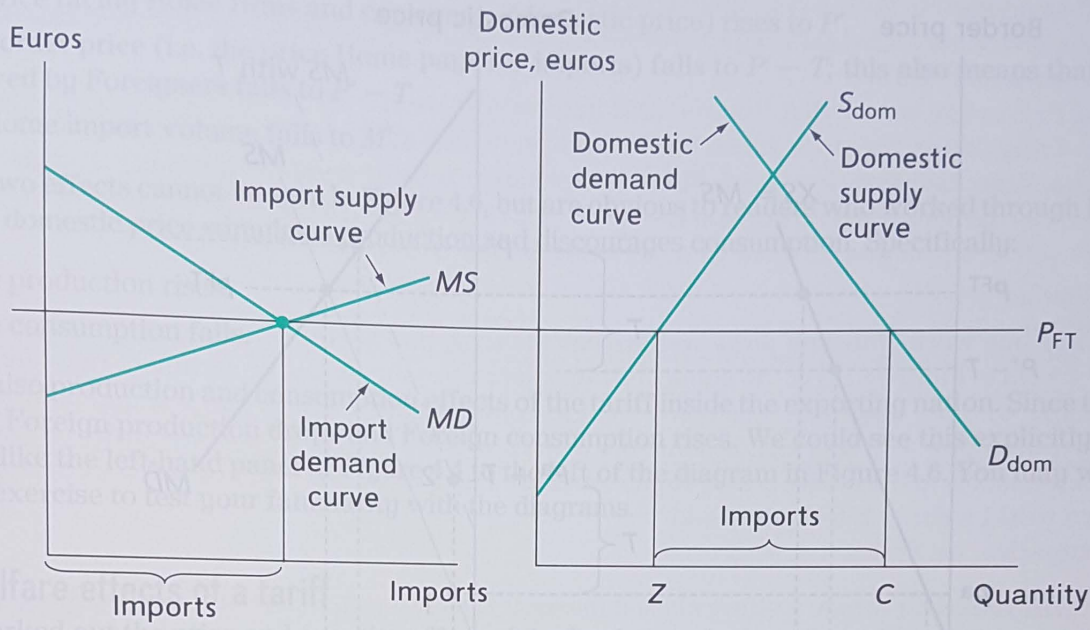

Having explained these basic microeconomic tools, we turn now to using them to study a simple but common real-world problem – the effects of a tax on imports from all nations. Such taxes are called tariffs.

4.3 MFN tariff analysis

To build from simple to complex, we preface the analysis of preferential trade liberalization in Europe with a simpler example, but one that nevertheless is useful for understanding the world. That is, we introduce the basic method of analysis and gain experience in using the diagrams by first studying the impact of removing the simplest type of trade barrier – a tariff that is applied to imports from all trade partners. We call this a non-discriminatory liberalization.

Although this is not what happened when Europe integrated economically, we first look at the non-discriminatory case since it is less complex. An extra benefit of taking this detour is that it helps us understand the effects of the EU lowering its common external tariff – as it does in the context of world trade talks (see Chapter 12). For historical reasons, a non-discriminatory tariff is called a 'most favoured nation' tariff, which provides the handy abbreviation, MFN.

4.3.1 Price and quantity effects of a tariff

The first step is to determine how a tariff changes prices and quantities. To be concrete, suppose that the tariff imposed equals T euros per unit.

The first step in finding the post-tariff price is to work out how the tariff changes the *MD–MS* diagram; here, Figure 4.6 facilitates the analysis. (See Section 4.2 if you are unfamiliar with the *MD–MS* diagram.) The right-hand panel of Figure 4.6 shows the pre-tariff import demand and import supply curves as *MD* and *MS*, respectively. The left-hand panel shows the foreign export supply curve as *XS*. Note that the vertical axis in this right-hand panel shows the domestic price, while the vertical axis in the left-hand panel shows the border price – the difference between the two is simple, but critical (see the note to Figure 4.6).

A tariff shifts up the MS curve

Imposition of a tariff has no effect on the *MD* curve in the right-hand panel since the *MD* curve tells us how much Home would like to import at any given *domestic* price. By contrast, imposing a tariff on imports

Figure 4.6 Price and quantity effects of an MFN tariff

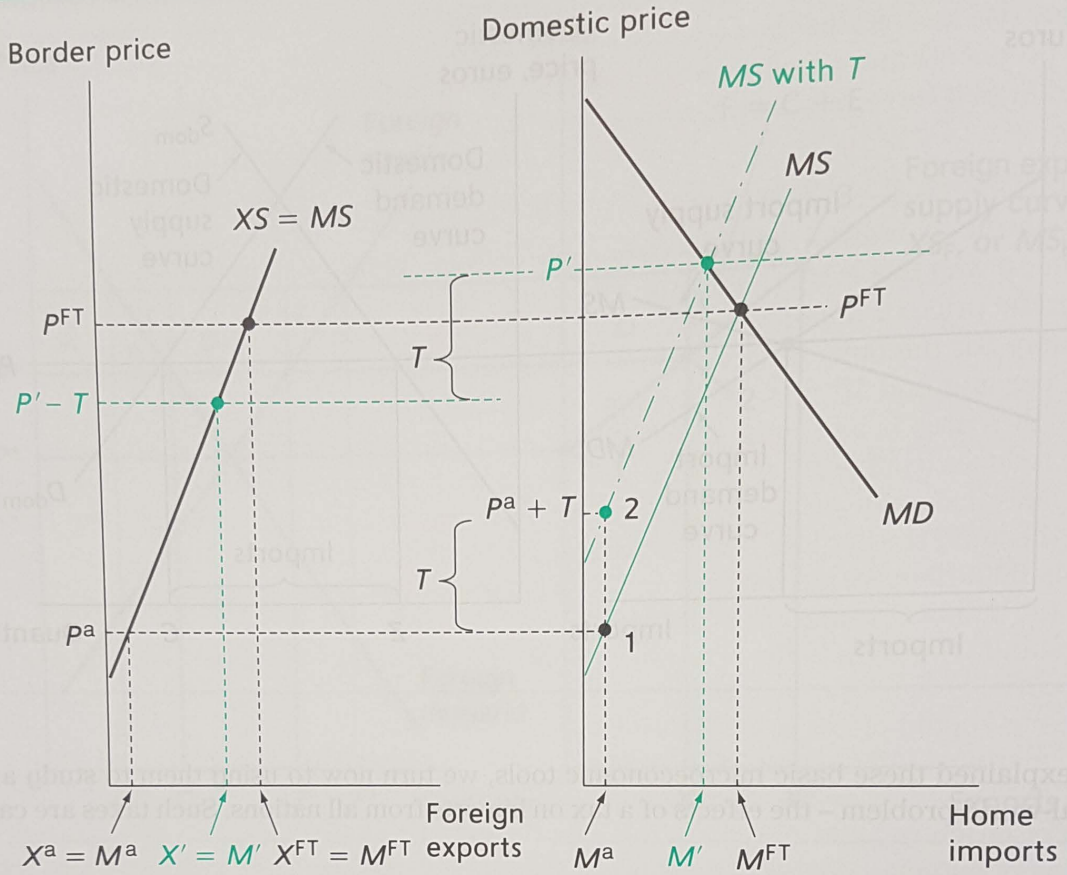

Note: Observe the distinction between the domestic and border prices. The domestic price is the price that domestic consumers pay for the good. The border price is the price foreign producers receive when they sell the good to Home. They can differ because of the tariff (a tariff is nothing more than a tax on imports). When you buy a coffee at a café for, say, 1 euro, the café owner does not get the full euro because the owner has to pay a tax, called the VAT, on your purchase. As a result, the price that the café owner receives is only 80 cents (the VAT is 20 per cent in this example) even though you pay 100 cents. In exactly the same way, foreigners receive a price (the border price) that equals the domestic price minus the tariff.

shifts up the *MS* curve by *T*. The reason is simple. After the tariff is imposed, the domestic price must be higher by *T* to get Foreign to offer the same quantity as it offered before the tariff. Consider an example. How much would Foreign supply before the tariff if the Home domestic price before the tariff were P^a? The answer, which is given by point 1 on the *MS* curve, is M^a. After the tariff, we get a different answer. To get Foreign to offer M^a after the tariff, the domestic price must be $P^a + T$ so that Foreign sees a border price of P^a.

Having shown that the tariff shifts up the *MS* curve, consider next the tariff's impact on equilibrium prices and quantities.

The new equilibrium prices and quantities

Even without a diagram, readers will surely realize that a tariff raises the domestic price and lowers imports. After all, a tariff is a tax on imports and it is intuitively obvious that putting a tax on imports will raise prices somewhat and lower imports somewhat. Why do we need a diagram?

The diagram helps us be more specific about this intuition; this specificity allows us to work out how much the nations gain or lose from the tariff. Returning to our analysis, note that, after the tariff, the old import supply curve is no longer valid. The new import supply curve, labelled *MS with T*, is what matters and the equilibrium price is set at the point where the new import supply curve and the import demand curve cross. As intuition would have it, the new price – marked *P′* in the diagram – is higher than the pre-tariff price

P^{FT} (as already noted, FT stands for free trade). Because of the higher domestic price, Home imports are reduced to M' from M^{FT}. To summarize, there are five price and quantity effects of the tariff:

1 The price facing Home firms and consumers (domestic price) rises to P'.
2 The border price (i.e. the price Home pays for imports) falls to $P' - T$; this also means that the price received by Foreigners falls to $P' - T$.
3 The Home import volume falls to M'.

The other two effects cannot be seen in Figure 4.6, but are obvious to readers who worked through Figure 4.3. The higher domestic price stimulates production and discourages consumption. Specifically:

4 Home production rises.
5 Home consumption falls.

There are also production and consumption effects of the tariff inside the exporting nation. Since the border price falls, Foreign production drops and Foreign consumption rises. We could see this explicitly if we put a diagram like the left-hand panel of Figure 4.4 to the left of the diagram in Figure 4.6. You may want to do this as an exercise to test your familiarity with the diagrams.

4.3.2 Welfare effects of a tariff

Having worked out the price and quantity effects, it is simple to calculate the welfare effects of the tariffs; that is to say, who wins, who loses and by how much. The analysis is really just a combination of what we did in Figures 4.3 and 4.4; this is done in Figure 4.7. The left-hand panel shows Home's supply and demand, the middle panel shows the world market for imports and the right-hand panel shows the Foreign supply and demand. We start with Home.

Figure 4.7 Welfare effects of an MFN tariff

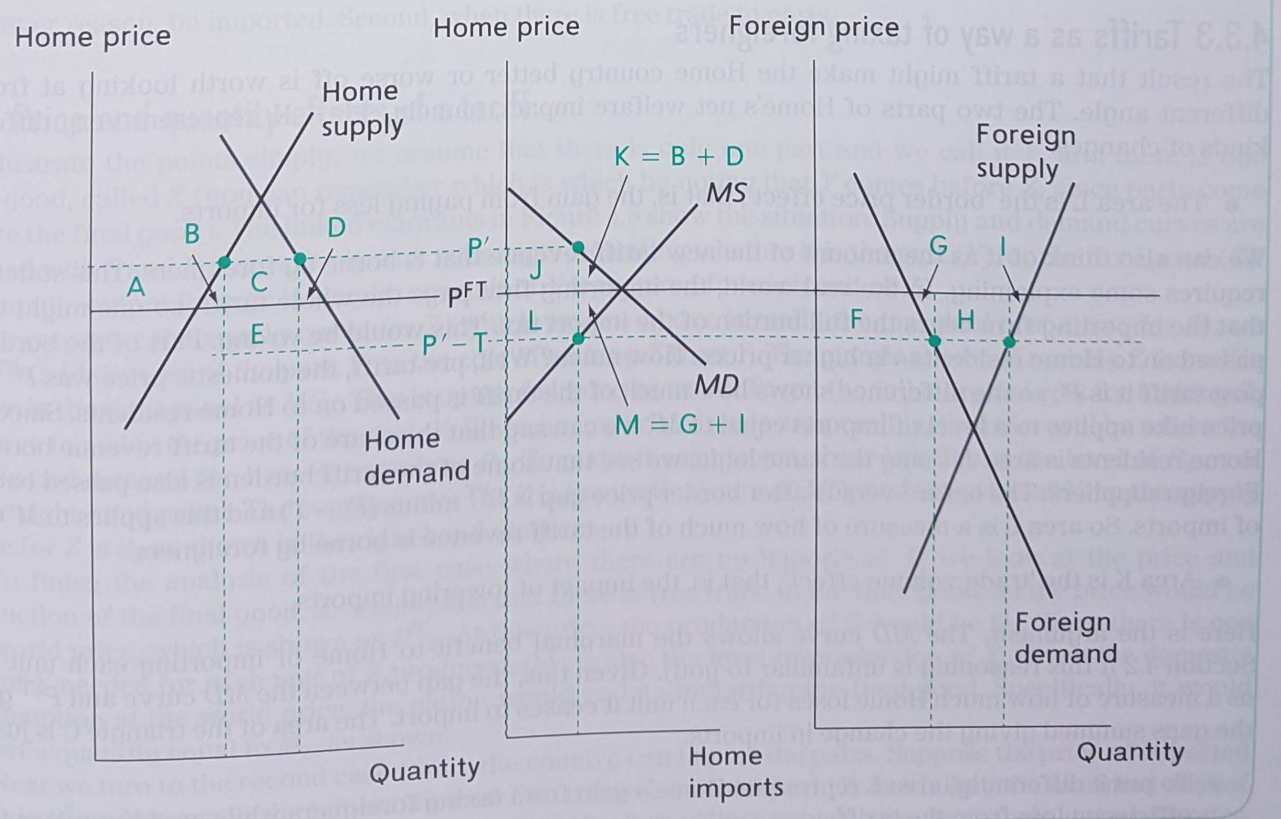

As shown in Figure 4.7, the MFN tariff raises the Home price of the good (to P') while lowering the border price (to $P' - T$). Home consumers lose A + B + C + D, Home producers gain A and Home government gains tariff revenue C + E. The net Home welfare effect is +E − B − D. This can be positive or negative depending upon the size of the tariff (you can show that it will be negative for very large tariffs, but positive for sufficiently small tariffs).

Turning to Foreign, we see that Home's tariff has lowered the border price facing Foreign exporters and this in turn brings down the price faced by Foreign producers and consumers. Foreign consumers gain F while Foreign firms lose F + G + H + I. There is no change in tariff revenue (the tariff is paid to the Home government), so the net impact on Foreign is −G − H − I. This is plainly negative regardless of the tariff's size.

A useful condensation

The first time one works through these welfare calculations, it is useful to consider the full distributional effects as we did in Figure 4.7 (i.e. the impact on consumers, producers and government revenue). Yet, once one is familiar with the diagrams, it is convenient to condense the analysis into a single diagram, like the centre panel in Figure 4.7. This lets us show the overall welfare effects of a Home tariff on both nations. Using the area labels in the centre panel, Home's welfare changes by +L − K, Foreign's welfare changes by −L − M, so world welfare falls by −K − M.

To summarize, we find:

- The tariff reduces Foreign welfare since it means it sells less and receives a lower price.
- The tariff creates private-sector winners and losers (Home firms gain, Home consumers lose), but the losers (consumers) lose more than the gainers (firms) gain.
- Home collects tariff revenue equal to J + L.
- The overall Home welfare change is +L − K; this net effect may be positive or negative; the relative sizes of L and K depend upon the slopes of the *MD* and *MS* curves and on the size of *T*.
- The global impact of the tariff, adding Home and Foreign welfare changes together, is definitely negative.

4.3.3 Tariffs as a way of taxing foreigners

The result that a tariff might make the Home country better or worse off is worth looking at from a different angle. The two parts of Home's net welfare impact, namely, +L − K, represent very different kinds of changes.

- The area L is the 'border price effect', that is, the gain from paying less for imports.

We can also think of it as the amount of the new tariff revenue that is borne by foreigners. This statement requires some explaining. In the real world, the importing firm pays the whole tariff, so one might think that the importing firm bears the full burden of the import tax. This would be wrong. Part of the burden is passed on to Home residents via higher prices. How much? Well, pre-tariff, the domestic price was P^{FT} and post-tariff it is P', so the difference shows how much of the tariff is passed on to Home residents. Since this price hike applies to a level of imports equal to M', we can say that the share of the tariff revenue borne by Home residents is area J. Using the same logic, we see that some of the tariff burden is also passed back to Foreign suppliers. The before-versus-after border price gap is P^{FT} minus $(P' - T)$ and this applies to M' units of imports. So area L is a measure of how much of the tariff revenue is borne by foreigners.

- Area K is the 'trade volume effect', that is, the impact of lowering imports.

Here is the argument. The *MD* curve shows the marginal benefit to Home of importing each unit (see Section 4.2 if this reasoning is unfamiliar to you). Given this, the gap between the *MD* curve and P^{FT} gives us a measure of how much Home loses for each unit it ceases to import. The area of the triangle C is just all the gaps summed giving the change in imports.

- To put it differently, area L represents Home's gain from taxing foreigners while area K represents an efficiency loss from the tariff.

Given all this, we can say that if T raises Home welfare, then it does so only because the tariff allows the Home government to indirectly tax foreigners enough to offset the tariff's inefficiency effects on the Home economy. That is, T causes economic inefficiency at Home but T is also a way of exploiting foreigners. Since the exploitation gains may outweigh the inefficiency effects, Home may gain from imposing a tariff.

4.3.4 Global welfare effects and retaliation

The global welfare impact is simply a matter of summing up effects; as we saw, it is negative and equal to $-K - M$.

Put in this way, the possibility that Home might gain from a tariff is clearly suspect. For example, if Home and Foreign were symmetric and both imposed tariffs, both would lose the efficiency triangle K and the gain to Home of L on imports would be lost to Home on its exports to Foreign. Home would also lose the deadweight triangle M on exports, so the net loss to each of the symmetric nations would be $-K - M$.

In short, protection by all nations is worse than a zero-sum game. It is exactly this point that underpins the economics of WTO tariff-cutting negotiations. If only one nation liberalizes, it might lose. If, however, the nation's liberalization is coordinated with its trading partners' liberalization, the zero-sum aspect tends to disappear.

4.4 GVC analysis

One of the most important recent developments in trade has been the rise of so-called Global Value Chains (GVCs). These are nothing more than supply chains that cross borders, so, for example, different parts of an Airbus plane are made in different European nations. This section introduces a diagram that allows an analysis of the gains from internationalizing supply chains. It can also be used to understand the cost of disrupting supply chains (as in the case of Brexit).

For most manufactured goods, the production process is more complicated than the one assumed in the previous diagrams. In particular, some parts are imported instead of being made locally. To simplify, we contrast two extreme situations. First, where the parts must be made locally since they cannot, for whatever reason, be imported. Second, when there is free trade in parts.

4.4.1 Price and quantity effects of a tariff

To illustrate the points simply, we assume that there is only one part and we call it Y, and there is one final good, called Z (you can remember which is which by noting that Y comes before Z, since parts come before the final good). The linked diagrams in Figure 4.8 show the situation. Supply and demand curves are marked with S and D respectively with subscripts indicating whether it is the part (Y), or the final good (Z). The input – output linkage is simple; one unit of Y is required for each unit of Z.

To construct the supply curve for Z (which we saw above is really the marginal cost curve for Z), we have to add the cost of the part, Y, to the cost of making Z from Y. The marginal cost of making Z from Y is shown in the right panel as MC_Z. This is upward sloping as usual. To get the full marginal cost for making Z, we have to add in the price of the part, Y. When no imports of parts are possible, the price will be where the supply and demand for Y meet, namely at P_Y. To make it easy to see, there is a line at the level of P_Y over to the right-hand panel. The supply curve for Z is the vertical sum of MC_Z and the price of Y, so the supply curve for Z is S_Z as shown in the right-hand diagram.

To finish the analysis of the first case where there are no imports of Y, we look at the price and production of the final good, Z. We assume that there is free trade in the final good, so the price would be the world price, which is shown as P^W_Z. At this price, the production of Z would be Q_Z. Since there is one unit of Y needed for each unit of Z produced, this is also the level of production of Y. Given the domestic consumption at the world price, the nation would end up importing the final good. Specifically, it would import a quantity equal to M_Z, as shown.

Next we turn to the second case where the country can import the parts. Suppose the price of imported parts is P^P_Y, as shown in the left-hand panel. Since the cost of the part, Y, has fallen, the marginal cost of making Z will also fall. We can see this by drawing the dashed lines between the two panels at the new price

Figure 4.8 The open-economy GVC diagram

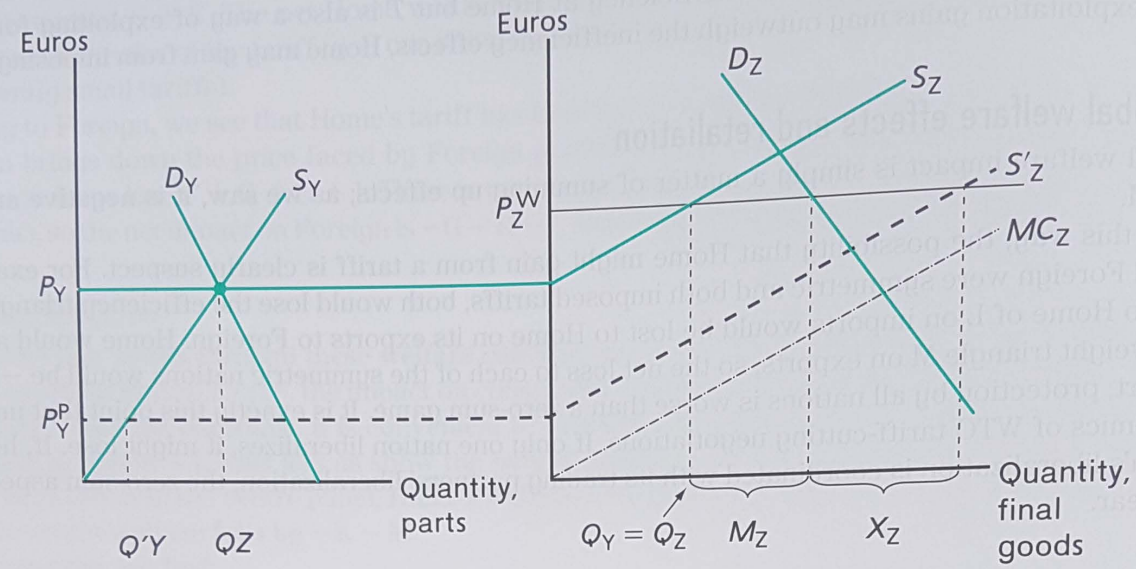

Source: Authors' original elaboration.

for Y, namely P^P_Y. The new supply curve is the vertical sum of the price of Y and the cost of making Z from Y (as before), but because the Y is cheaper, the new supply curve for the final curve is lower at S'_Z.

How does this change trade? The first point is obvious. Since imported parts are cheaper, the nation starts to import parts and this reduces its own Y production to Q'_Y from the previous level of Q_Y. By contrast, the lowering of the supply curve for Y from S_Z to S'_Z increases the nation's production of the final good, Z – so much so that now the nation is an exporter of Z instead of an importer. This makes sense intuitively. The nation was an importer of Z before because it was not very cost-effective in making Z, and part of the inefficiency is that it had to make its own parts. Once it can buy parts from lower-cost foreign producers, it gains cost competitiveness in the final good Y. In this example the gain in cost competitiveness is so great that it switches from being an importer to an exporter.

Interested readers can combine this supply-chain diagram with the tariff analysis to see how changing the tariff on the final good has very different effects than changing the tariff on intermediate goods.

4.5 Types of protection: an economic classification

Tariffs are only one of many types of import barrier that European integration has removed. The first phase of EU integration, 1958–68, focused on tariff removal, but the Single Market Programme that was started in 1986 focused on a much wider range of non-tariff barriers.

While there are several methods of categorizing such barriers, it proves useful to focus on how the barriers affect so-called trade rents. A tariff, for instance, drives a wedge between the Home price and the border price (i.e. the price paid to foreigners). This allows someone (in the tariff case it will be the Home government) to indirectly collect the 'profit' from selling at the high domestic price while buying at the low border price. For historical reasons, economists refer to such profits (area A + B in Figure 4.9) as 'rents'. When it comes to welfare analysis, we must watch the trade rents closely. For some import barriers, Home residents get the rents, but for others no rents are created, or foreigners get them. This distinction is highlighted by distinguishing three categories of trade barrier: domestically captured rent (DCR) barriers; foreign captured rent (FCR) barriers; and 'frictional' barriers. We consider them in turn.

4.5.1 DCR barriers

Tariffs are the classic DCR barrier. Here, the Home government gets the trade rents. From a Home nationwide welfare perspective, however, it does not really matter whether the government, Home firms or

Figure 4.9 Home welfare effects of import protection

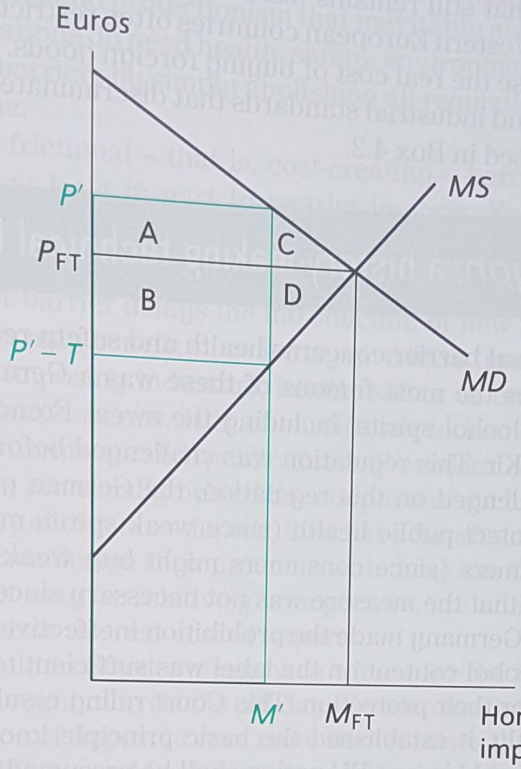

Home consumers earn these rents, as long as the rents are captured domestically. What sorts of barrier other than tariffs would lead to domestically captured rents? Some forms of quotas are DCR barriers. A quota is a quantitative limit on the number of goods that can be imported per year. To control the number of foreign goods entering the country, the government hands out a fixed number of import licences and 'collects' one licence per unit imported. The price and quantity effects of a quota that restricts imports to M' in Figure 4.9 are identical to the effects of a tariff equal to T. The point is that, if imports are limited to M', then the gap between domestic consumption and production can be no more than M', implying that the domestic price must be driven up to P'. Another way to say this is that T is the 'tariff equivalent' of the quota. Now consider the trade rents. With a quota, whoever has the licence can buy the goods at the border price $P' - T$ and resell them in the Home market for P'. This earns the licence holders A + B. If the government gives the licences to Home residents, then the quota is a DCR barrier. If it gives them to foreigners, the quota is an FCR barrier.

4.5.2 FCR barriers

A prime example of an FCR barrier is a 'price undertaking' in the context of an anti-dumping tariff. Under EU law, the Commission can impose a tariff on a non-member nation if that nation's firms are selling goods below cost in the EU market (so-called dumping). In some cases, an anti-dumping tariff is imposed but in other cases no tariff is imposed and the exporting firm promises to raise its price instead. These promises are called 'price undertakings'.

For example, if the agreed level were P' from Figure 4.9, the price undertaking would have the same price and quantity effects as a tariff, T. Importantly, however, the undertaking allows foreign producers, rather than the Home government, to garner the rents A + B. Throughout the industrialized world, and in the EU in particular, it is very common for trade barriers to be arranged so that foreigners earn the rents. One reason is that trade rents are used as a kind of gift to soothe foreign companies and governments that are likely to be angered by the imposition of a trade barrier.

Finally, note that an FCR barrier harms EU welfare more than a DCR barrier. Specifically, the welfare cost of an FCR is always negative (i.e. −A − C), instead of being ambiguous (i.e. B − C). Moreover, the foreign welfare impact is now A − D, so an FCR may end up helping foreigners!

4.5.3 Frictional barriers

An important type of trade barrier that still remains inside the EU consists of what are sometimes called 'technical barriers to trade' (TBTs). Western European countries often restrict imports by subjecting them to a whole range of policies that increase the real cost of buying foreign goods. Some examples are excessive bureaucratic 'red-tape' restrictions and industrial standards that discriminate against foreign goods. One of the most famous examples is discussed in Box 4.2.

Box 4.2 *Cassis de Dijon:* a history-making technical barrier to trade

One very common type of frictional barrier concerns health and safety regulations that have the side effect of hindering trade. Perhaps the most famous of these was a German regulation that forbade the importation of certain low-alcohol spirits, including the sweet French liqueur, Cassis – used in the famous white wine cocktail, Kir. This regulation was challenged before the EU's Court of Justice as a barrier to trade. When challenged on this regulation, the German government argued that the prohibition was necessary to protect public health (since weak spirits more easily promote alcohol tolerance) and to protect consumers (since consumers might buy weak spirits thinking they were strong). In 1979, the Court ruled that the measure was not necessary since widespread availability of low-alcohol drinks (e.g. beer) in Germany made the prohibition ineffective in furthering public health. It also found that putting the alcohol content on the label was sufficient to protect consumers, so the import ban was not necessary for their protection. This Court ruling resulted in the frictional barrier being removed. More importantly, it established the basic principle known as 'mutual recognition' whereby goods that are lawfully sold in one EU nation shall be presumed to be safe for sale in all EU nations. Exceptions to this principle require explicit motivation. By the way, the formal name for this Court case is *Rewe-Zentral AG v Bundesmonopolverwaltung für Branntwein;* no wonder it is called *Cassis de Dijon.*

Since frictional barriers are bad for a nation, one may ask why they are so prevalent. Box 4.3 provides one explanation.

Box 4.3 Why do frictional barriers arise so often?

Government agencies charged with formulating and enforcing standards are often 'captured' by special-interest groups from the regulated industries. Moreover, the Home firms that are to be subjected to the standards often play an important role in setting the standards. For example, when regulating a highly technical field such as elevators, the government (which probably does not employ many full-time elevator experts) naturally asks the opinions of domestic firms that produce elevators. With an eye to their foreign competitors, they quite naturally push for standards that raise the cost of imported goods more than the cost of locally produced goods.

An example can be found in the paper industry. Sweden and Finland produce paper mainly from new trees, while French and German paper producers use a lot of recycled paper and rags. In the early 1990s, the EU was considering a regulation that would require all paper sold in the EU to contain a certain fraction of recycled paper. This sounds like a 'public interest' regulation. However, it also would have had the effect of eliminating the resource-based advantage of Swedish and Finnish firms, much to the joy of French and German firms. In other words, it would have raised the real cost of imports (since the Nordic producers would have had to switch to less efficient techniques). As it turns out, it is not clear which production method is 'greener'. Recycling paper requires lots of chemicals that may be released into the environment, while establishing more tree farms is, well, green – a point that was not raised by French and German paper producers.

As a result of Finland and Sweden joining the EU, the regulation was not adopted, but this shows the subtle mixing of public interest and protectionism that inevitably arises when nations adopt regulations and standards. Of course, nations do need health, safety, environmental and industrial standards, so we cannot eliminate frictional barriers by simply abolishing all regulation. This is one of the issues tackled by the EU's 1992 programme.

One important class of frictional – that is, cost-creating – barriers involves industrial and health standards that are chosen at least in part to restrict imports. For example, some countries refuse to accept safety tests that are performed in foreign countries, even in highly industrialized nations. This forces importers to retest their products in the local country. Beyond raising the real cost of imported goods, this sort of barrier delays the introduction of new products. While this clearly harms consumers, Home producers may benefit since it may give them time to introduce competing varieties. Another example involves imposing industrial, health, safety or environmental standards that differ from internationally recognized norms. It is often difficult to know objectively whether an unusual regulation or standard represents a valid 'public interest' concern or whether it is just a protectionist device. In fact, both motives are usually behind the adoption of such measures.

Regardless of why such policies are adopted, they have the effect of protecting Home producers or service providers. Home firms design their products with these standards in mind, while foreign firms, for whom the Home market may be relatively unimportant, are unlikely to do so. Bringing imported products into conformity raises the real cost of imports.

For example, all cars sold in Sweden must have wipers for the headlights. While this policy may have some merit as a safety regulation (in the old days Sweden had lots of dusty rural roads), it also has the effect of raising the price of imported cars more than it raises the price of Swedish cars. From the drawing board onwards, all Volvo and Saab models – and their production facilities – are designed with these headlight wipers in mind. For other car makers, take Renault as an example, the Swedish market is far too small to really matter. The design of Renaults and Renault's mass production facilities are not optimized for the installation of headlight wipers. Consequently, while it is expensive to put headlight wipers on both Swedish and French cars, it is much more so for French cars. This gives the Swedish car makers an edge in Sweden. Similar sorts of barrier give the French an edge in their domestic market.

4.6 Sources of competitiveness differences

The diagrams in Section 4.3 assumed that the two nations incurred different costs when producing the good whose quantity appears on the horizontal axis (see Figure 4.7, for example). But where do such cost and price differences come from?

A major part of international trade theory is concerned with exactly this question. In that literature it is called the question of 'sources of comparative advantage'. This section introduces some basic notions of comparative advantage theory to help readers understand the real-world sources of these price differences.

4.6.1 Traditional comparative advantage made simple

Comparative advantage analysis starts with a sector-by-sector comparison of the competitiveness of individual nations. To structure our thinking about sectoral competitiveness, it is useful to focus on a simplistic notion of competitiveness – one where cheaper means more competitive. To keep things simple, we brush aside all cost considerations apart from labour productivity and wages. The Home nation's cost of producing a particular good is the number of hours required to produce and sell one unit of the good and Home's wage. In the jargon of trade economists, this is called the 'unit labour cost', which just means the cost of the labour necessary to make one unit of a particular good.

For example, if it takes workers in a UK factory a total of 7 hours to produce an electric fan, and the UK wage is, say, 5 GBP per hour, then the fan costs 35 GBP. Under the assumption of perfect competition, the price on the market would be 35 GBP. Suppose the same electric fan takes 20 hours to make in an Italian factory (since Italian factories are less productive in this example), and the Italian wage is 12 euros an hour, the Italian-made fan would cost 140 euros.

Which fan is cheaper? The answer depends upon the exchange rate, that is, the number of euros per pound. If the exchange rate (EUR/GBP) is 2.00 euros per pound, then converting the GBP price to euros implies that the British-made fan would cost 70 euro – cheaper than the 140 euro price of the Italian-made fan. This is what we mean when we say Britain is more competitive than Italy in terms of fans.

The same sector-by-sector comparison is made for four goods in Table 4.1, namely, electric fans, espresso machines, jet engines and designer silverware. In each case, illustrative hours-per-unit are listed for the two nations. To calculate prices, the table also includes the two nations' wages, and an illustrative exchange rate. The British wages are converted from pounds to euros for the calculation. The second- and third-to-last columns show the calculated prices. The final column shows the ratio of the Italian price to the UK price (both in euros).

Table 4.1 Example of sector-by-sector competitiveness

	Hours needed in:		Wages (local currency)		Exchange rate	Wages in euros		Prices in euros		Price ratio
	UK	Italy	UK (GBP)	Italy (EUR)	EUR per GBP	UK	Italy	UK	Italy	Italy/ UK
Electric fan	7	20	5	12	2	10	7	70	140	2.000
Espresso machine	10	13	5	12	2	10	7	100	91	0.910
Jet engine	1300	3000	5	12	2	10	7	13,000	21,000	1.615
Designer silverware	23	15	5	12	2	10	7	230	105	0.457

Sources of sector-by-sector competitiveness

The relative price is a measure of the UK's competitiveness sector by sector because, when the Italian price is high compared to the British price, it is the UK good that is more competitive. It is instructive to think about the relative price (e.g. 140/70) in a slightly different way. To start with, observe that the relative price depends upon two things: relative labour productivity (i.e. hours needed in Italy over hours needed in the UK); and relative wages measured in a common currency.

Relative productivity is something that changes quite slowly as it depends upon the nations' industrial histories, general level of scientific and technological know-how, management efficiency and experience in making the goods concerned – to mention just a few factors.

The relative wage also moves. But what is the 'right' relative wage? The answer is that the relative wages adjust to ensure that Italy and the UK, or both, are competitive in some sectors but not all. Returning to the illustrative example, Figure 4.10 facilitates the discussion. The dark bars in the chart plot the relative price numbers for the four goods listed in the last column of Table 4.1, when we took the relative Italian-to-UK wage (measured in euros) to be 7/10. The line at 1.0 is relevant since if the bar is above this, the UK is the lower-cost producer (i.e. competitive in the sector). If the bar is below 1.0, Italy is the one that is competitive in the sector (i.e. the low-cost producer). When the relative wage is 7/10 as assumed in Table 4.1, each nation is competitive in two sectors.

Now consider what the situation would look like if the exchange rate were 1.0 euros per GBP. Before turning to calculations, think about what this would do to the relative competitiveness of Italy and the UK. If wages remained constant in local terms (euros and pounds), then the exchange rate change (from 2 euros to 1 euro per point) would raise the relative price of Italian labour. Or, looking at it from the British perspective, it would make UK wages appear to fall – relative to Italian wages – by 50 per cent, namely, from 10 to 5 euros. Plainly, this will tend to improve UK competitiveness in all sectors.

Figure 4.10 Illustration of the relative wage that balances competitiveness

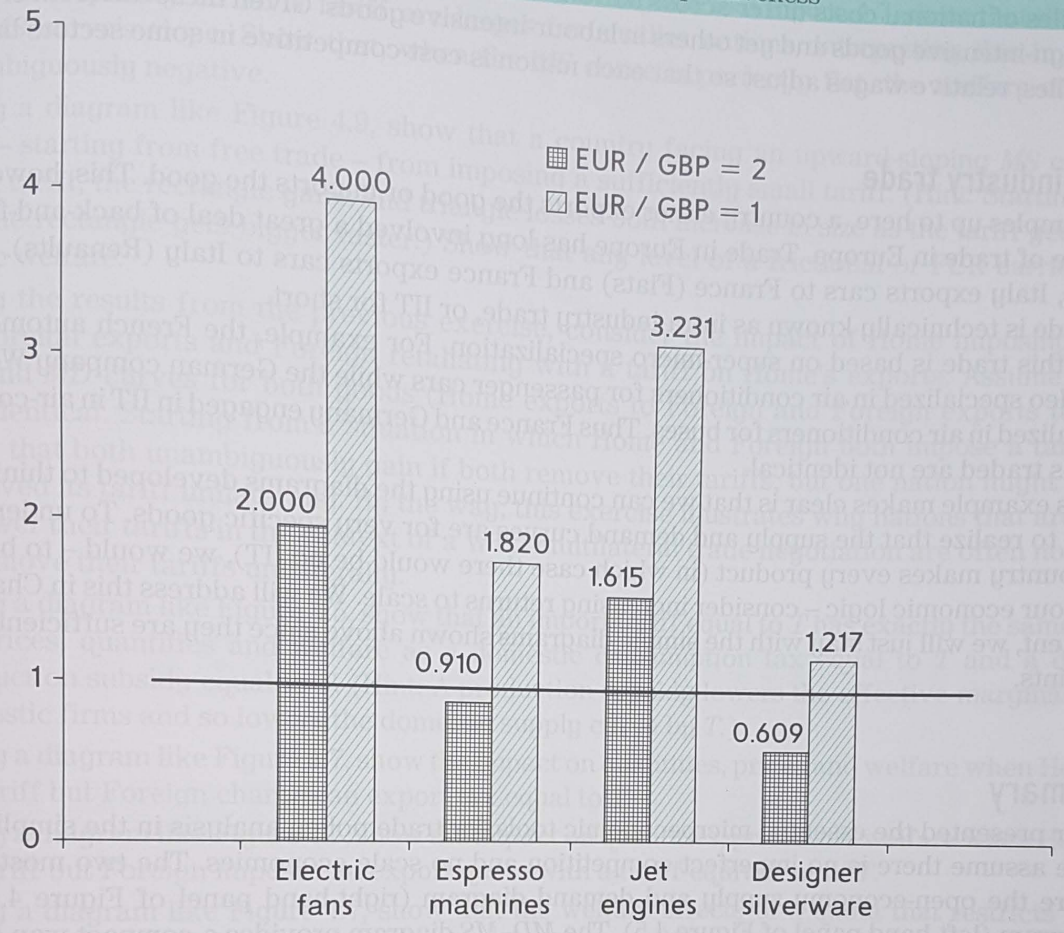

To identify the degree of change, we can redo the price calculations using this simple formula:

$$\frac{\text{Italian price}}{\text{UK price}} = \frac{\text{Italian hours needed}}{\text{UK hours needed}} \times \frac{\text{Italian wage (€)}}{\text{UK wage (£)}} \times \frac{1}{\text{€s per £}}$$

The answers are listed above the light bars in Figure 4.10 (diligent readers should work this out for themselves). The results show that Italy would be uncompetitive in all sectors at this 1.0 exchange rate. Surely this is not an equilibrium because it would mean that Italy imports everything and exports nothing.

The real-world mechanisms for arriving at the equilibrium relative wage are complex. Explaining them would fill a few more chapters. But even without a full understanding, our simple 'thought experiment' serves to elucidate the basic considerations. If Italian workers are too expensive relative to UK workers – taking into account their relative productivity – then Italy would have no exports to pay for its imports. In equilibrium, the relative wage will adjust so that each nation is competitive in some sectors. When exchange rates can move – as is the case for the EUR/GBP rate – then some of the adjustment can come from changes in the number of euros to the pound. When exchange rates are locked in – as they are among the Eurozone nations – the only way to adjust national competitiveness is to change wages directly. Doing so can be painful, as the recent Eurozone crisis has shown.

The simple illustration above looks only at labour productivity and the price of labour (i.e. the wage). In reality, many other things affect the competitiveness of goods – product quality and reliability, for example. These help explain why most Europeans will pay more for a Volkswagen than they will for a similar car made by Renault. Quality and reliability themselves are the outcome of complex factors and interactions among these factors. Since quality is hard to measure and hard to change, simple microeconomics typically ignores it by assuming that only prices matter.

The bottom-line insight of this discussion of comparative advantage is simple. For many reasons, the sectoral profiles of national costs differ across nations. Some nations are really good at engineering goods, others in design-intensive goods and yet others in labour-intensive goods. Given these different cross-sector, national profiles, relative wages adjust so that each nation is cost-competitive in some sectors but not all.

4.6.2 Intra-industry trade

In all the examples up to here, a country either exports the good or imports the good. This, however, is not the main type of trade in Europe. Trade in Europe has long involved a great deal of back-and-forth trade. For example, Italy exports cars to France (Fiats) and France exports cars to Italy (Renaults). This back-and-forth trade is technically known as intra-industry trade, or IIT for short.

Much of this trade is based on super-micro specialization. For example, the French automotive parts company Valeo specialized in air conditioners for passenger cars while the German company Webasto Bus GmbH specialized in air conditioners for buses. Thus France and Germany engaged in IIT in air-conditioners, but the goods traded are not identical.

What this example makes clear is that we can continue using the diagrams developed to think about IIT but we have to realize that the supply and demand curves are for very specific goods. To understand why not every country makes every product (in which case there would be no IIT), we would – to be perfectly complete in our economic logic – consider increasing returns to scale. We will address this in Chapter 6, but for the moment, we will just stay with the simple diagrams shown above since they are sufficient to explain the main points.

4.7 Summary

This chapter presented the essential microeconomic tools for trade policy analysis in the simplified world in which we assume there is no imperfect competition and no scale economies. The two most important diagrams are the open-economy supply and demand diagram (right-hand panel of Figure 4.5) and the *MD–MS* diagram (left-hand panel of Figure 4.5). The *MD–MS* diagram provides a compact way of working out the impact of import protection on prices, quantities and overall Home and Foreign welfare. The open-economy supply and demand diagram allowed us to consider the distributional impact of import protection, that is, to separate the overall effect into its component effects on Home consumers, Home producers and Home revenue. The chapter also covered the open-economy GVC diagram, which showed how to take into consideration the fact that many inputs into the production of goods are imported.

The chapter also discussed types of trade barrier in Europe and classified them according to what happens to the trade 'rents'. Under the first type, DCR barriers, the rents go to domestic residents; with FCR barriers, the rents go to foreigners; and with frictional barriers, the rents disappear. European integration consisted primarily of removing DCR barriers up until the mid-1970s. Subsequent goods-market liberalization has focused on frictional barriers.

The final topic was a quick introduction to the intuition behind the sources of comparative advantage, that is, the reasons why nations are competitive in some but not all sectors.

Self-assessment questions

1 Using a diagram like Figure 4.9, show the full Foreign welfare effects of imposing a Home tariff equal to *T*, i.e. show the impact on Foreign producers and Foreign consumers separately.

2 The exit of Britain from the EU is likely to disrupt supply chains for things like Airbus planes and automobiles like the Mini. Go online to find some information about the supply chains of Airbus and use the Supply-Chain diagram to discuss the impact on the cost of Airbus planes.

3 One way to think about the slope of the *MS* curve is in terms of the 'size' of the home nation. The idea is that the demand from a very small nation has a very small impact on the world price. For example, Switzerland could probably increase its oil imports by 10 per cent without

having any impact on the world oil price. Using a diagram like Figure 4.7, show that the welfare costs of imposing an MFN tariff are larger for smaller nations, interpreting this in terms of the *MS* curve's slope. Show that, when the *MS* curve is perfectly flat, the welfare effects are unambiguously negative.

4 Using a diagram like Figure 4.9, show that a country facing an upward-sloping *MS* curve can gain – starting from free trade – from imposing a sufficiently small tariff. (Hint: Starting from a small tariff, the rectangle gains and triangle losses both increase in size as the tariff gets bigger, but the rectangle gets bigger faster.) Show that any level of a frictional or FCR barrier lowers Home welfare.

5 Using the results from the previous exercise, consider the impact of Home imposing a tariff on Foreign exports and Foreign retaliating with a tariff on Home's exports. Assume that the *MS* and *MD* curves for both goods (Home exports to Foreign and Foreign exports to Home) are identical. Starting from a situation in which Home and Foreign both impose a tariff of T, show that both unambiguously gain if both remove their tariffs, but one nation might lose if it removed its tariff unilaterally. By the way, this exercise illustrates why nations that are willing to lower their tariffs in the context of a WTO multilateral trade negotiation are often not willing to remove their tariffs unilaterally.

6 Using a diagram like Figure 4.5, show that an import tariff equal to T has exactly the same impact on prices, quantities and welfare as a domestic consumption tax equal to T and a domestic production subsidy equal to T. (Hint: A production subsidy lowers the effective marginal cost of domestic firms and so lowers the domestic supply curve by T.)

7 Using a diagram like Figure 4.7, show the impact on quantities, prices and welfare when Home has no tariff but Foreign charges an export tax equal to T.

8 Using a diagram like Figure 4.5, show the impact on quantities, prices and welfare when Home has no tariff but Foreign imposes an export quota with a tariff-equivalent of T.

9 Using a diagram like Figure 4.7, show that the welfare effects of a quota that restricts imports to M' are exactly the same as a tariff equal to T; assume that each quota licence (i.e. the right to import one unit) is sold by the government to the highest bidder.

References and further reading

Further reading: the aficionado's corner

Mankiw, G. (2011) *Principles of Economics*, 6th edition (or earlier), South Western Publishing, New York.

Mathieu, M., T. Spencer and O. Sartor (2014) 'Economic analysis of the US unconventional oil and gas revolution', http://www.voxeu.org/article/limited-economic-impact-us-shale-gas-boom.

Every undergraduate textbook on international economics has a chapter on tariff analysis that covers the same material as this chapter. One particularly accessible treatment can be found in:

Krugman, P. and M. Obstfeld (2005) *International Economics*, 7th edition (or earlier), HarperCollins, New York.

For much more on the economics of trade protection, see:

Vousden, N. (1990) *The Economics of Trade Protection*, Cambridge University Press, Cambridge.

Useful websites

The World Bank's website provides extensive research on trade policy analysis. This includes many papers on non-discriminatory trade policy but also a very large section on preferential trade arrangements under the heading of 'regionalism'. See www.worldbank.org.

The Commission's website on trade issues can be found at http://ec.europa.eu/trade/. It has lots of information on the latest changes to EU trade policy.