3 | How markets work (in an imaginary world)

'The supply and demand model, which we introduced in Chapter 3 and have used repeatedly since then, is a model of a perfectly competitive market.' Krugman and Wells (2005: 207)

'Perfect competition is rarely, if ever, found in practice.' Baumol and Blinder (2006: 194)

1 THE STANDARD TEXT

1.1 What is a competitive market?

There is no 'competing' behaviour in a competitive market – no advertising, no price-setting strategies, no rivalry. This is because all buyers and sellers are *price-takers*. This requires large numbers of buyers and sellers, with no one buyer or seller having a significant market share, and all firms producing an identical (or homogeneous) good or service.

1.2 The demand curve

An individual's demand curve describes the relationship between the quantity demanded and the good's own price *ceteris paribus* (holding all other influences constant). These other influences include the individual's preferences, income and the prices of related consumption goods. These may be either complements (such as DVDs and DVD players) or substitutes (such as chicken or beef). Expected future prices may also be important in determining how much is bought currently.

An individual's demand curve is a frontier – it tells us the maximum price he or she is willing to pay to obtain any given quantity. If any of the other influences change, the demand curve *shifts*. To obtain the market demand, we sum the amounts every individual wishes to buy at any given price. Thus the size of the population influences demand.

The shape of the market demand curve shows the responsiveness of quantity demanded to price changes. Normally, as the price increases, quantity demanded decreases, as seen in Figure 3.1.

The responsiveness of quantity demanded to a change in price is measured by the *price elasticity of demand*. It is defined as:

$$e_d = \frac{\%\text{ change in quantity demanded}}{\%\text{ change in price}}$$

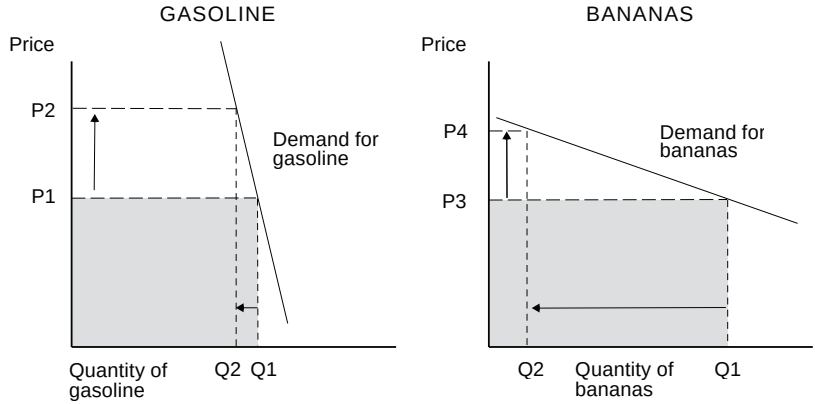

FIGURE 3.1 Inelastic and elastic demand

Suppose a 10 per cent increase in the price of gasoline leads to a 1 per cent decrease in the quantity demanded. The price elasticity of demand for gasoline is 0.1 (1 per cent divided by 10 per cent). In contrast, suppose a 10 per cent increase in the price of bananas leads to a 70 per cent decrease in the quantity demanded. The price elasticity for bananas is 7. Elasticities depend on the availability of substitutes among other things. There are many substitutes for bananas (other fruit), but few substitutes for gasoline (petrol), at least over a short time period.

Total revenue is price multiplied by quantity. Elasticity of demand determines how total revenue changes when price changes. For example, when the price of gasoline is P1, total revenue equals the shaded box in the left-hand diagram of Figure 3.1. When the price goes up to P2 the height of the total revenue box increases by 10 per cent, and its width decreases by 1 per cent. It is clear that the new total revenue box is bigger than before.

Similarly, when the price of bananas is P3, total revenue equals the shaded box in the right-hand diagram in Figure 3.1. When the price goes up to P4 the height of the total revenue box increases by 10 per cent while its width decreases by 70 per cent. Clearly, the new total revenue box is smaller than before.

As an application, suppose the London tube system is losing money. If the objective is solely to increase total revenue, should the tube authority increase or decrease fares? If the demand for tube rides is inelastic, they should increase fares; if it is elastic they should decrease fares. If the elasticity is equal to one (a so-called 'unit-elastic' demand curve) then a fare change would have no effect on total revenue.

1.3 The supply curve

The supply curve describes the relationship between the quantity of a good supplied and its own price, *ceteris paribus*. It too is a frontier, showing the *minimum* price that sellers are willing to accept for any given quantity. Generally speaking, as the price of a product increases, the quantity supplied goes up. The

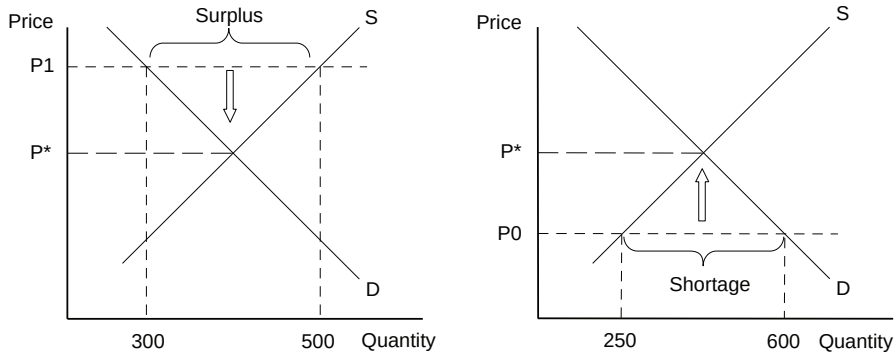

FIGURE 3.2 Movement towards equilibrium

responsiveness of quantity supplied to a change in price is measured by the *price elasticity of supply*.

$$e_s = \frac{\%\text{ change in quantity supply}}{\%\text{ change in price}}$$

The six key shift factors on the supply side are: the weather (especially important for agricultural products); changes in the prices of goods related in production; changes in input prices (or prices of 'factors of production'); changes in technology; changes in the number (and size) of firms in the industry; and changes in expectations about future prices.

Comparing the demand shift factors with the supply shift factors we see only one identical item: expectations of future prices.

1.4 Market equilibrium

When prices are free to fluctuate, market forces move the actual price (and quantity) towards the equilibrium price (and quantity). The left-hand diagram of Figure 3.2 shows that at a price P1, which is above the equilibrium price, P*, there is an excess supply (or surplus) equal to 200 units per period. This creates downward pressure on the price, causing it to fall until equilibrium is restored at P*. The right-hand diagram shows that when the price is below equilibrium there is an excess demand (or shortage). This creates upward pressure on the price, causing the price to increase until equilibrium is restored at P*.

1.5 Comparative static analysis

We simulate change by considering how an exogenous shock would affect the equilibrium position. 'Comparative statics' compares one static equilibrium position with another. The analysis is timeless (we don't know how long anything takes) and ahistorical (it doesn't matter in what order things happen).

In the left-hand diagram of Figure 3.3 we show the effect of an increase in the

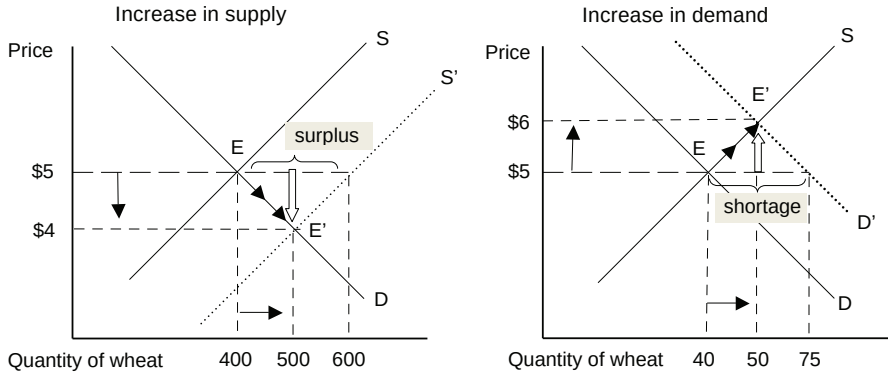

FIGURE 3.3 Comparative static analysis

supply of wheat from S to S' – perhaps caused by a fall in the price of fertilizer. At the original price of $5 there is now a surplus of wheat. This causes the price to fall to $4, eliminating the surplus. In the right-hand diagram of Figure 3.3 we show the effect of an increase in the demand for wheat – perhaps caused by an increase in incomes. The original demand line *shifts rightwards* to D', causing a shortage at the original equilibrium price of $5. This causes the price to increase to $6, at which point the shortage is eliminated.

Note that an expectation of a future price increase causes supply to shift left and demand to shift right. Both these shifts cause prices to increase now: an example of a self-fulfilling prophecy.

1.6 A government-regulated price ceiling: rent controls

Governments often try to control market prices using price ceilings and price floors. Rent control (an example of a price ceiling) is an attempt to help low-income families afford the cost of accommodation. Textbooks emphasize that attempts to overrule market forces always lead to unintended effects that usually hurt the very group the government is intending to help.

Figure 3.4 shows the market for apartments in Montreal assuming all apartments are identical. The going rent is $1,000 a month and 2 million units are rented. When the government imposes a rent ceiling of $800, fewer apartments are offered for rent and more demanded, causing a shortage of 400,000 rental units. The shortage is likely to get worse the longer the rent control is in effect, as apartment buildings are knocked down or converted to condominiums.

Shortages induced by price controls in competitive markets lead to *inefficiency*: missed opportunities to make some people better off at no cost to anyone else.

The first inefficiency is an inappropriate distribution of apartments among renters. For example, 'empty-nesters' want to downsize, while households with new children want something bigger. These moves benefit both parties, but are hampered by the shortage created by the rent control.

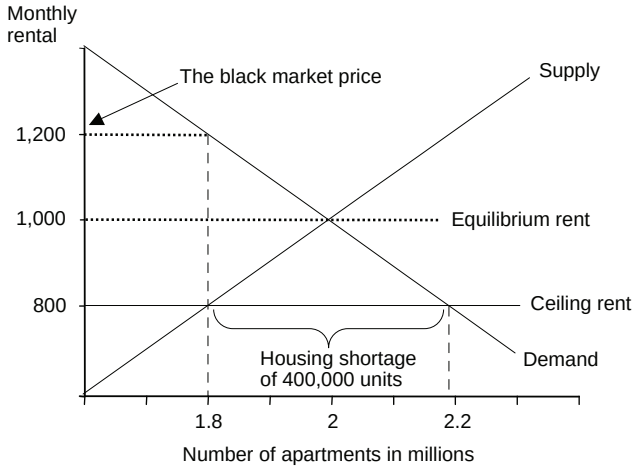

FIGURE 3.4 The effect of rent control

A second inefficiency is the wasted time, energy and money spent searching for an apartment.

A third inefficiency is that the quality of apartments will become undesirably low. Some tenants would be happy to pay for better conditions, and landlords would be happy to provide them for increased rent. This is a missed opportunity.

Finally, price ceilings encourage *illegal activities*, specifically the emergence of black markets – side payments (or bribes) to obtain an apartment. Given the shortage of apartments under rent controls, Figure 3.4 indicates that buyers are willing to pay up to $1,200 a month – $400 more than the legal ceiling. So, we can expect side payments as high as $400 a month.

Our analysis contains five predictions that we will state generically (not tailored to the market for apartments). First, price controls in competitive markets lead to shortages that get worse the longer they are in effect (prediction 1). Next, the fundamental reason shortages are bad is that they are inefficient, and this inefficiency manifests itself in three distinct ways: an inefficient distribution of the good among buyers (prediction 2); wasted resources trying to buy the good (prediction 3); and an inefficiently low quality of the good (prediction 4). Finally, whenever there are unsatisfied wants because of legal restrictions, crime will always arise to profit from them (prediction 5).

1.7 A government-regulated price floor: minimum wages

Figure 3.5 depicts a competitive market for unskilled workers. The equilibrium wage is $9 an hour, and total employment is 15 million workers. Suppose the government decides that $9 an hour is not a living wage, and imposes a minimum wage of $12 an hour. The impact is 3 million fewer jobs, 3 million more people willing to work and unemployment (or a surplus of labour) of 6 million workers.

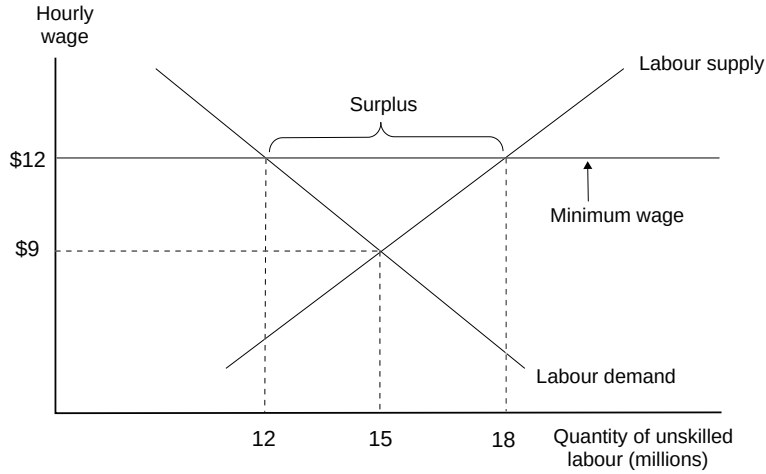

FIGURE 3.5 The effect of a minimum wage

Any minimum wage above the equilibrium has the same qualitative effect; but the higher the minimum wage, the worse it is.

The surplus caused by a price floor creates inefficiencies – missed opportunities – that resemble those created by price ceilings. First, there is an *inefficient allocation* of sales among sellers. With a minimum wage there may be some job seekers who really want to work but cannot find a job, and others who have a job but are almost indifferent as to whether they work or not. Second, sellers (job seekers) waste time and effort searching for a buyer (an employer). Third, suppliers offer an inefficiently high quality to try to attract buyers, who might have preferred the original quality at a lower price. Finally, price ceilings provide an incentive for *illegal activity* – only in this case it is sellers (job seekers) bribing buyers (employers), or employment arrangements out of sight of the law.

1.8 Who bears the cost of sales taxes?

Contrary to popular belief, the person who ends up 'paying' a sales (or excise) tax is not the same person on whom the tax is levied. Rather, the incidence of the tax depends on the relative size of the price elasticities of demand and supply. The texts demonstrate this proposition using demand and supply diagrams.

The left-hand diagram of Figure 3.6 shows supply and demand for parking spaces. We assume the government collects the sales tax from producers of parking spaces. This adds to producers' costs, so a $4 per unit sales tax shifts the supply curve upwards by $4 per unit for each level of output. According to the diagram, the effect is to raise the equilibrium price from $6 to $7. Effectively $1 of the tax has been passed on to consumers in higher prices. The remainder, $3 per unit, is paid by producers. Finally, the tax raises $1,600 in revenue for the government ($4 x 400 units).

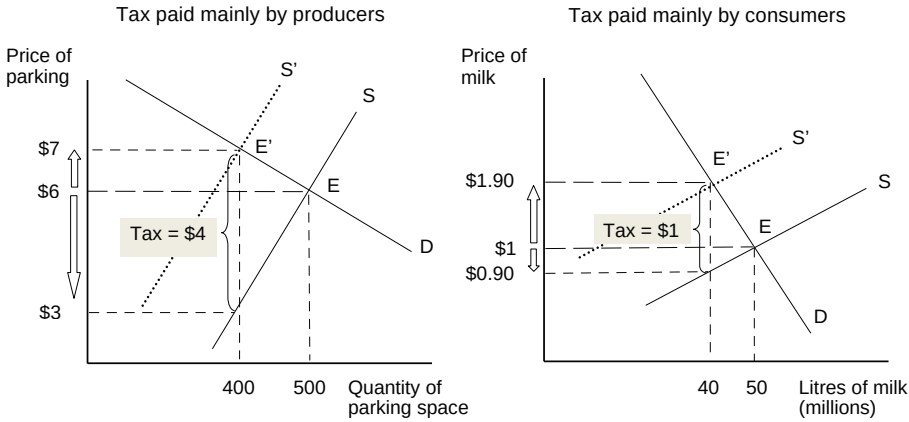

FIGURE 3.6 The incidence of taxation

In the right-hand diagram, we have supply and demand for milk. A sales tax of $1 per litre of milk shifts the supply curve up by $1 per litre. According to the diagram, the tax raises $40 million in government revenue ($1 x 40 million litres), and raises the equilibrium price by 90 cents. In contrast to the previous example, where producers paid most of the tax, here most of the tax is passed on to consumers, and only 10 cents per litre is paid by producers. What causes this difference in who bears the burden of the tax?

It turns out that the incidence of the tax depends on the relative size of the price elasticities of demand and supply at the equilibrium prices and quantities. The actual formula (almost never revealed by textbooks) is:

$$\text{Formula 1: proportion of sales tax borne by consumers} = \frac{e_s}{e_s + e_d}$$

$$\text{Formula 2: proportion of sales tax borne by producers} = \frac{e_d}{e_s + e_d}$$

where e_s and e_d are the elasticities of supply and demand in equilibrium. Note that the proportions sum to one – the total tax is split between buyers and sellers.

Formula 1 shows that the greater the price elasticity of supply, and the smaller the price elasticity of demand, the more the tax is paid by consumers. If we understand elasticity to mean 'responsiveness', this amounts to a claim that if producers are responsive (or flexible, or elastic) to price changes, while consumers are unresponsive (or inflexible or inelastic), the more the tax burden falls on consumers. As in the martial art of t'ai chi, the flexible opponent will always beat the inflexible one. In this case, the more flexible side of the market avoids the larger part of the sales tax, while the inflexible one pays it.

1.9 The costs of taxation

In Figure 3.6 the tax on milk reduced consumption by 10 million litres. This is milk that would have been consumed in the absence of the tax, to the mutual

benefit of both producers and consumers. Nobody would have been worse off. So, an excise tax creates inefficiency. This represents a cost of the tax over and above the money paid to the government in taxes. This extra cost is referred to as the excess burden or deadweight loss of the tax. Economists say that the real cost of a tax is not what people pay but what they don't pay – meaning the mutually beneficial trades that no longer occur because of a tax.

2 THE ANTI-TEXT

2.1 The demand and supply model is sold as a generic tool

The material summarized in the previous section is often called 'How markets work', the terminology of Parkin and Bade (2006) and Mankiw et al. (2006). It follows a discussion of models and methodology that emphasizes the overriding importance of predictive power, and it contains applications to a broad range of labour and product markets. The range of these applications, the position of these chapters near the front of the text, and the immediately preceding methodological discussion that plays down realism of assumptions, *all suggest that the supply and demand framework is a generic tool that can be applied to all markets*. Colander et al. (2006: 72) are explicit here, saying supply and demand provides 'a good off-the-cuff answer for any economic question'.

But the supply and demand framework is actually a simplified representation of a perfectly competitive market structure, one which (according to some textbooks) is so rare as to be hardly ever found in practice. Many textbooks are quite explicit about this, but only much later in the texts, when they discuss perfect competition. For example, Krugman and Wells state in their Chapter 9: 'The supply and demand model, which we introduced in Chapter 3 and have used repeatedly since then, is a model of a perfectly competitive market' (2005: 207).[1]

We can demonstrate the equivalence between supply and demand and perfect competition by showing that perfect competition is the *only* market structure where a supply curve exists. While this demonstration is delayed until Chapter 6, the intuition is straightforward. Unlike the competitive firm, which is a price taker, the non-competitive firm faces a downward-sloping demand curve and has to decide on its best price–output combination. But the best price depends on the position of the demand curve; as a result, there is no unique relationship between price and the quantity supplied, and hence no supply curve.

Since supply curves exist only in perfectly competitive markets, we need to know: first, how many markets are perfectly competitive in the real world? And second, even though the competitive model is not (strictly speaking) applicable to non-competitive markets, can it be usefully applied as an approximation? We address those questions next.

Question for your professor: If an industry is not perfectly competitive, can we still draw the industry supply curve? (Right answer: No.)

2.2 How many markets are perfectly competitive?

To this point we've talked vaguely about 'non-competitive' markets. Let's be more precise. Textbooks categorize markets according to the number of producers and the type of product, as shown in Table 3.1.

Non-competitive markets fall into three types: monopolistic competition, oligopoly and monopoly. Firms in all three non-competitive markets have 'market power', which means they face a downward-sloping demand curve, and so can choose the price of their product rather than simply accept a 'market price' like the perfectly competitive firm. Market power derives either from the firm being large relative to the industry (monopoly), or from having a product that is unique (or differentiated) in some way (monopolistic competition), or for both reasons (oligopoly). The question is: how prevalent is perfect competition relative to the other market structures?

The key requirement for perfect competition is price-taking behaviour. All texts agree that this requires large numbers of buyers and sellers and an identical product. But two other assumptions are often included: perfect information and easy entry by firms into (and exit from) the industry. This last assumption is necessary to show the long-run optimality qualities of perfect competition (discussed in Chapter 6). Easy entry (and exit) is not necessary for price-taking behaviour.[2]

Concerning the information requirements, four of eleven leading US textbooks state that perfect information is required for perfect competition (Beaulier and Mounts 2008).[3] Of the remainder, several state that market participants just need to be 'well informed'. An extreme position is that of Mankiw (2004) – one of the leading US textbooks – who chooses not to mention any information requirements for perfect competition at all.

TABLE 3.1 Types of market structure

How many producers are there?	Are products differentiated? NO	YES
Many	Perfect competition	Monopolistic competition
Few	Oligopoly	Oligopoly
One	Monopoly	Monopoly

Given the amount of research on the effects of imperfect information, this difference of opinion among the textbooks is odd. Joseph Stiglitz received the Nobel Prize for his work on information economics in 2001. In his acceptance speech, published the following year in the flagship journal of the American Economic Association, he explains:

> For more than 100 years, formal modeling in economics has focused on models in which information was assumed to be perfect. Of course, everyone recognized that information was in fact imperfect, but the hope ... was that economies in which information was not too imperfect would look very much like economies in which information was perfect. One of the main results of our research was to show that this was not true; that even a small amount of information imperfection could have a profound effect on the nature of the equilibrium. (2002: 461)

Why could even a small amount of imperfect information have a profound effect? Stiglitz gives the following example:

> Assume for example, as in the standard theory, that all firms were charging the competitive price, but there were an epsilon cost of searching, of going to another store. Then any firm which charged half an epsilon more would lose no customers and thus would choose to increase its price. Similarly, it would pay all other firms to increase their prices. But at the higher price, it would again pay each to increase price, and so on until the price charged at every firm is the monopoly price, even though search costs are small. (Ibid.: 477)

In the above quote, 'epsilon' stands for an arbitrarily small quantity. Just an epsilon of costs of acquiring information could lead otherwise competitive firms to charge the monopoly price. The point is that even slight departures from free, and hence perfect, information have large consequences. Depending on the market, other consequences could be an equilibrium where the market does not clear (the quantity demanded differs from the quantity supplied *in equilibrium*), or even multiple equilibria.

If imperfect information undermines the competitive model, is there a better alternative? Stiglitz explains: 'a central consequence of imperfect information is that ... product markets are more aptly described by models of imperfect competition, where ... [firms] perceive themselves facing downward sloping demand schedules' (1985: 34).

Apparently, the prevalence of competition depends on the likelihood of having perfect information. To appreciate how implausible the assumption of perfect information is, it helps to realize that many information asymmetries (some people knowing more than others) are inevitable. Job applicants know more about their ability than prospective employers; workers know more about their work effort than management; management knows more about their firms than potential investors; borrowers know more about their likelihood of default

than lenders; people buying insurance know more about their efforts to avoid risk than insurers. According to Stiglitz, information imperfections are so pervasive 'it is hard to imagine what a world with perfect information would be like' (2002: 469).

Given the prevalence of imperfect information, Stiglitz's argument seems to leave little room to apply the competitive model. This leads him to ask why the competitive paradigm persisted for so long. He says:

> Despite its deficiencies, the competitive paradigm did provide insights into many economic phenomena. There are some markets in which the issues which we have discussed are not important – the market for wheat or corn – though even there, pervasive government interventions make the reigning competitive paradigm of limited relevance. (Ibid.: 488)

So, if not for government interventions, the competitive model would be good enough to describe the markets for wheat and corn. Are there other markets where information is close enough to perfect that the competitive model can be applied? What do the textbooks themselves say on this question?

Opinions differ as to whether perfect competition actually describes many real-world markets. Those textbooks that include *perfect* information as a requirement state that perfect competition has very limited applicability. For example, Baumol and Blinder say that perfect competition is 'rarely, if ever, found in practice' (2006: 194). Those that don't insist on perfect information make stronger claims for the existence of perfectly competitive markets. For instance, Krugman and Wells state: 'important parts of the economy are fairly well described by perfect competition' (2005: 383). The stronger claim was also made by Ragan and Lipsey, who provide examples of perfectly competitive industries. They state (2005: 259): 'Forest and fish products provide many examples. Agriculture also fits fairly well in most ways since individual farmers are clearly price takers. Many basic raw materials, such as iron ore, tin, and copper, are sold on world markets where most individual firms lack significant market power.'

But not all agricultural and raw material products are perfectly competitive. The existence of market power in markets for diamonds, aluminium and oil are well known.[4] Perhaps less well appreciated is the existence of market power in agricultural markets. At least in North America, many farmers and fishers are increasingly squeezed by the market power of the few firms that supply their inputs and the few buyers of their outputs – especially supermarket chains and the fast food industry (Phillips 2003; Lawrence 2004; Schlosser 2001).[5]

But even assuming that the whole of agriculture, forestry, fishing and hunting were perfectly competitive, their combined output is only a very small fraction of total production in the industrialized economies. (It's about 2 per cent of production in Canada in 2006.)[6]

Evidence that price-taking behaviour is rare outside of agriculture is provided

by Blinder et al. (1998). They survey 200 representative firms in the United States, excluding agriculture. They say: 'First of all, we took it for granted that almost all firms in our economy are price-makers rather than price-takers – an assumption amply justified by the survey responses' (ibid.: 12). They find that prices are 'sticky' – set by firms and periodically reviewed; they are not determined instantaneously by supply and demand. They say:

> First, the evidence gathered in this study emphatically supports the mainstream view that sticky prices are the rule, not the exception, in American industry. According to our respondents, the median number of price changes for a typical product in a typical year is just 1.4, and almost half of all prices change no more often than annually. Among firms reporting regular price reviews, annual reviews are by far the most common. At the other end of the spectrum, only about 10 percent of all prices change as often as once a week, and about 7 percent of all firms schedule price reviews at least weekly. (Ibid.: 298)

Competitive firms are price-takers. They never need to review their price schedules. Their prices change continually with shifts in demand and supply. *None* of the firms surveyed by Blinder et al. fell into that category.

Questions for your professor:

1 In the demand and supply model no one is a price-setter. So, who determines what the price will be?
2 Most firms in the real world set their own prices; does the model apply to them?

2.3 Is the competitive model a useful approximation?

No two hairstylists are equally skilled. They sell a differentiated product. Each stylist faces a downward-sloping demand curve, implying that the supply curve for haircuts does not exist, as stated earlier. Furthermore, there is no unique price for haircuts, but instead a range of prices – each price *set* by the hairstylist – depending on the stylist's quality, reputation, location and clientele. This is a non-competitive market.

Nevertheless (we ask, as the devil's advocate), can the competitive model be applied to this market as an *approximation*? Assume away all the complications. Assume all hairstylists are identical. Assume perfect information. Won't the competitive model give us insights into the determinants of the average price of haircuts? Won't the things that cause supply curves to shift left – an increase in the costs of production (shampoo prices go up), or a decrease in the number of firms (hairstylists) – increase the average price of a haircut? If so, the competitive model provides a *useful approximation* even to this non-competitive market.

If this were generally true, the textbook emphasis on competitive markets as a generic tool would be justified. Maybe it's like the law of gravity: strictly speaking it holds only in a vacuum; but it can be usefully applied in everyday life.

This is the position taken by Krugman and Wells, who, after noting that oligopoly is by far the most common market structure, ask, 'Given the prevalence of oligopoly, then, is the analysis ... based on perfect competition still useful?' They argue that it is because '[i]t is also true that predictions from supply and demand analysis are often valid for oligopolies'. Given the complexity of oligopoly models, 'in situations where they do not expect the complications associated with oligopoly to be crucial, economists prefer to adopt the working assumption of perfectly competitive markets' (2005: 383). In other words, the competitive model is simpler and can be applied even to non-competitive markets, because it gives us accurate predictions.[7]

Let us, then, consider the predictive power of the competitive model, focusing on the core applications emphasized in introductory textbooks.

Predictions concerning minimum wages Does the evidence support the predictions of the supply and demand framework concerning the effects of minimum wages? We addressed this question in detail in Chapter 2. In brief, the empirical studies conflict to such an extent that we used it as a case study to illustrate the limitations of hypothesis testing and predictive power as criteria for model selection. The consensus concerning the effects of minimum wages has broken down – though this is not generally reported in the textbooks.

What we have yet to explain is why moderate increases in the minimum wage might *not* reduce employment of low-wage, low-skilled workers. There are several possible explanations, all of which depend on 'frictions' – imperfect information or mobility costs.

One category of explanation is the 'efficiency wage' thesis. If work effort is hard to monitor, workers may shirk. Wage increases make the job more desirable and provide an incentive not to get caught shirking (which might result in getting fired). As a result, workers shirk less and productivity increases. In addition, increased worker morale may reduce labour turnover, which reduces hiring and training costs for the firm. Either way, the higher wages pay for themselves without causing job losses.

An alternative category of explanation is the dynamic monopsony thesis. Here employers are not simply wage-takers: they have some short-run (or temporary or dynamic) power to set wages lower than other firms without losing all their workers. This power may derive from the time and resources necessary for a worker to find a new job, or because taking another job might entail moving home or increased costs of commuting. Either way, moderate minimum wage increases may offset the market power of employers without causing job losses – indeed, they may even cause job gains (as we explain more fully in Chapter 8).

The minimum wage application is precisely about whether frictions in the labour market are significant. The easiest way to account for the array of mixed evidence is to concede that frictions are important in certain cases. So, when discussing the effects of minimum wages it is important to contrast the predictions of the competitive model with those from non-competitive models.

A small minority of textbooks do compare the predictions. For example, Ragan and Lipsey (2008: 99) alert the reader that the minimum wage is re-examined in a later chapter on labour markets using a monopsonistic framework.

But the average textbook continues to apply the competitive model *as if it were the only model* relevant to the minimum wage question. Some of these mention the empirical controversy in passing. Others mention the controversy, but dismiss results contrary to the competitive model as wrong (Parkin and Bade 2006: 131). And others pretend there is no empirical controversy and continue to cite results from 'the typical study', which finds that 'a 10 percent increase in the minimum wage depresses teenage employment between 1 and 3 percent' (Mankiw et al. 2006: 125). Krueger (2001: 247) tracks down this 'typical' study to an influential survey paper published in 1982!

A third option is to omit the minimum wage application completely (e.g. Frank et al. 2005; McConnell et al. 2007). Krueger believes this reaction is unfortunate:

> Did astronomy classes stop teaching Newtonian principles once quantum mechanics was discovered? Did physics classes drop lectures on the atom once quarks were discovered? No. Instead, these disciplines explain the limitations of their models, teach the research methods and findings that have been used to establish (and reject) their core principles, and seek to provide students with an understanding of which models work best in which circumstances and why. (2001: 243)

```
Question for your professor: Does the empirical evidence
support the predictions of the supply and demand framework
concerning the effects of minimum wages? (Right answer: It's
very mixed.)
```

Predictions concerning rent controls The main prediction of the competitive model – shortages that get worse the longer the rent control is in effect – depends on the rent ceiling remaining below the equilibrium level: it must be *binding*. If the extent to which it is binding lessens – if the ceiling rent moves towards the equilibrium rent – then we would not expect shortages to worsen. On the contrary, we'd expect them to moderate. But knowing the extent to which the ceiling rent is binding over time is very tricky. It's complicated by the fact that we cannot observe the equilibrium rent.

A second complication is that the type of rent control prevalent nowadays is very different from the type assumed in textbooks – a rigid rent freeze. Controls of this sort were introduced in major US cities during the Second World War, but every city (apart from New York) had abandoned this 'first-generation' rent control by the early 1950s. 'Second-generation' rent control, first introduced in the 1970s, is significantly more flexible. For example, it commonly allows automatic rent increases geared to increasing costs, excludes luxury high-rent buildings and new buildings, restricts conversions, decontrols between tenants, and provides incentives for landlords to maintain or improve quality.

A third complication is that housing units are assets, the desirability of which is impacted by many other factors besides rent control: interest rates, inflation, profit opportunities elsewhere, the local real estate cycle, government housing and tax policies, and current and expected future changes in all relevant variables.

In reviewing the empirical evidence on rent control, Arnott says: 'The impact of these other factors is likely to be significantly greater than any effect due to controls. Trying to discern the effects of rent control in such a situation is akin to trying to hear a whispered conversation across a street of roaring traffic' (1995: 112). He suggests that with the exception of New York City (which retained its first-generation controls) and perhaps Toronto (which had poorly designed second-generation controls) the effects of rent control in North America have been almost imperceptible. This is a dramatic contrast to the treatment in the textbooks. By assuming that the rental housing market is perfectly competitive, and by considering a crude form of price ceiling, most texts suggest that rent controls necessarily have destructive effects.

Why are most textbooks (and most North American economists for that matter) so negative on rent controls? Arnott suggests two reasons: 'The first is ideological. The debate over rent control has been a battleground between those who believe in the free market and those who do not. The echoes of the debate carry over to other policy arenas where its resolution has far more quantitative import. The second is methodological' (ibid.: 117).

The methodological battle is about whether the competitive model is good enough as a generic approximation to most markets. The housing market has many non-competitive elements: apartments are heterogeneous and tastes idiosyncratic, which renders the market thin; search costs are substantial (as evidenced by agents' fees), as are moving costs; and there is a lack of information about who's a good landlord and who's a good tenant. Are these merely details that can be ignored as irrelevant? Most housing economists believe that these are too important to be ignored in practice. Since the mid-1980s most of them have turned their attention to non-competitive models – models that emphasize search costs and the importance of contracts.

Because of this different methodological perspective, they are much less

critical of rent control. Arnott conjectures: 'Perhaps a majority, at least among the younger generation, would agree with the statement that a well-designed rent control program can be beneficial' (ibid.: 99) Yet this research seems to have had no impact on the principles textbooks.[8]

> Question for your professor: Would rent controls necessarily cause shortages if the rental housing market were only imperfectly competitive? (Right answer: No.)

Predictions concerning the incidence of taxation If the evidence presented for the effects of minimum wages and for rent controls is weak, things are even worse when it comes to the incidence of sales taxes: the textbooks present no evidence at all.

This is very strange because the competitive model makes clear predictions: the proportionate burden of a sales tax is determined by the relative elasticities of supply and demand. The texts illustrate this idea using relative slopes of supply and demand in a wide variety of markets. Table 3.2 shows the examples used by ten leading US and Canadian texts. The favourite example is cigarettes (seven cases) – a *highly* oligopolistic industry composed of six US producers; next comes gasoline (five cases) – another non-competitive industry (oligopolistic at the production level, oligopolistic or monopolistically competitive at the retail level); and finally, luxury boats (three cases) – an industry with many sellers but highly differentiated products. None of the examples remotely resembles

TABLE 3.2 Tax incidence applications used in ten major North American textbooks

Text	Example used
Colander (2004, pp. 163–5)	Luxury boats
Gwartney et al. (2006, pp. 94–9)	Gasoline and luxury boats
McConnell and Brue (2005, pp. 589–90)	'A certain domestic wine'
Miller (2004, pp. 125–7, 485–6)	Gasoline and cigarettes
O'Sullivan and Sheffrin (2003, pp. 334–40)	Apartments, cigarettes and luxury boats
Ragan and Lipsey (2008, pp. 84–7)	Cigarettes
Samuelson and Nordhaus (1992, pp. 74–5)	Gasoline, cigarettes; imports; factor inputs
Stiglitz and Walsh (2002, pp. 206–7)	Cigarettes and cheddar cheese
Taylor (2004, pp. 174–6, 348–54)	Gasoline and salt
Tucker (2005, pp. 123–5)	Gasoline, cigarettes and alcoholic beverages

a competitive market. And this explains the complete absence of empirical evidence: while the elasticity of demand could be measured, we cannot measure the elasticity of supply when the supply curve doesn't exist.

Put it this way: it is one thing to generate predictions using hypothetical shifts of a hypothetical supply curve – if it yields accurate predictions this could be a useful approximation. But it is altogether another to test the accuracy of those predictions by *measuring* something that doesn't exist in reality. We suggest this is why *no text* presents *any* corroborating empirical evidence on the ability of the competitive model to predict the incidence of taxation.

> Question for your professor: Can the demand and supply model predict the incidence of taxation in imperfectly competitive markets?

2.4 But don't price floors cause surpluses and price ceilings shortages?

Price ceilings Price ceilings have been imposed on different commodities, in different countries, in different times. During the Second World War, there were price ceilings on many commodities in Britain, Canada and the USA – commodities such as meat, milk, eggs, sugar and gasoline. In every case, shortages developed. Doesn't this confirm the usefulness of the competitive model?

Not really. In Chapter 6 we show that the textbook model of monopoly contains the same prediction: if price ceilings are sufficiently low there will be shortages. Similarly, shortages are also a likely outcome in textbook models of oligopoly and monopolistic competition. The fact that shortages develop in response to price ceilings doesn't demonstrate the superiority of the competitive model.

We quote Krugman and Wells as preferring the competitive model because models of oligopoly (where strategic interaction is the key) are so complex. But the monopoly model is simple – just as simple as the competitive model. Why not use that? Why not champion the usefulness of the monopoly model as a generic tool?

The reason is that such an analysis would tell the wrong 'story'. For all the qualifications that are later tacked on to it, the central textbook story is how the market economy works like an invisible hand, efficiently allocating resources among alternative uses. As we'll see, an economy populated with firms that have market power does not allow a clear-cut story – hence the necessity to study an imaginary economy rather than something resembling the real one.

So, we're not arguing that price ceilings do not cause shortages. The issue is whether a competitive veneer can be smeared over every market as a decent enough approximation. If we accept the official methodology, of hypothesis testing and predictive power, then in each application, in each approximation,

we need to ask which works better: the competitive model or a non-competitive model.

> Questions for your professor: If price ceilings were low enough, would they cause shortages in non-competitive markets too? (Answer: Yes.) So, if price ceilings caused shortages in the Second World War, that can't be taken as empirical support for the demand and supply model, can it?

Price floors Similarly, we are not arguing that price floors don't cause surpluses. All economists would agree that if the minimum wage were raised high enough, a surplus of labour would be created. Where the minimum wage controversy begins is when we ask whether moderate increases have the same effect. As we've explained, the issue revolves around whether labour markets are *perfectly* competitive, or whether there are significant imperfections.

Perhaps one reason for the popularity of the minimum wage application is that there aren't a lot of examples of price floors where governments do not buy up the resulting surplus production. The combination of price floors and a 'government buyer of last resort' has resulted in butter and grain 'mountains' and milk 'lakes' in the European Community. This arrangement most certainly produces surpluses. But there are very few examples of a government imposing a price floor and not buying up the surplus production – besides minimum wages.

Krugman and Wells (2005: 93) use the example of transatlantic airfares. Prior to deregulation of airlines in 1978, airfares were set artificially high by international treaty. Certainly this restricted the quantity demanded, and since airlines couldn't compete in the price dimension, it led to them competing for customers by providing expensive (often unwanted) services. Krugman and Wells argue that it also resulted in surplus production, which manifested itself in empty seats on flights.

But this anecdotal evidence is hardly convincing. There are often empty seats on flights, even without price floors. And with price regulation we would expect airlines to reduce the number of flights to match the limited demand for travel; we would not expect them to increase the number of flights. Yet that's exactly

> Question for your professor: The demand and supply model suggests that suppliers will increase supply when binding price floors are imposed, despite observable surpluses. Isn't this irrational?

what an upward-sloping supply curve says firms would do in the face of a price increase. This is one aspect of a general problem: the competitive framework is based on assumptions that are violated in the context of disequilibrium. We develop this point in the next section.

2.5 What the texts don't tell you about the competitive model

The competitive model is internally inconsistent when not in equilibrium The perfectly competitive demand and supply model seems to make sense in equilibrium. Everyone takes prices as given, which is fine since everyone trades the amount they want, and no one has any incentive to change. But when something happens to disturb equilibrium the story starts to unravel.

Let's go back to the comparative static analysis, explained earlier using Figure 3.3. In the left-hand diagram we assumed a fall in the price of fertilizer shifted the supply curve of wheat to the right. This caused a surplus of wheat at the original equilibrium price. As a result, we are told prices fall. But since no one sets prices, how do they fall?

The lack of an explanation for price movements in the demand and supply model is known as Arrow's Paradox, after the issue raised by Kenneth Arrow (1959): all individuals and firms are assumed to be 'price-takers' and to have no influence over the market price, yet somehow the market price adjusts and reaches the equilibrium value. One 'solution' to this conundrum is to invent an auctioneer, who is 'the visible, if imaginary, embodiment of the invisible hand. He has no economic involvement in the market: no mention is made of his objectives or constraints' (Dixon 1990: 361–2). This fictitious character fills the glaring gap in the demand and supply model to adjust prices in response to excess supply and demand.

If having to invent an auctioneer is bad enough, what's worse is that the auctioneer can't allow any trades to occur until he finds the equilibrium solution. This is because the auctioneer needs eventually to end up at the intersection of the demand and supply curves. If we allow trades before equilibrium is reached, people will have spent some of their budget. As a result, they would not be able to buy what they otherwise would have bought at what would have been the new equilibrium price.

The demand and supply curves are derived assuming market participants can buy or sell all that they wish at the going market price. But they can't do that when there are shortages or surpluses. Out of equilibrium these curves are only 'notional'. They don't tell us how much buyers would try to buy, or sellers would try to sell, if there were a surplus or shortage.

For example, suppose there is a surplus. Do firms ignore this and continue to supply *as if* they could sell all that they wished? If so, the competitive supply curve would tell us what it purports to tell us: the quantity supplied at any given price. But surely it's more likely that firms would notice the surplus and reduce

their production. But if they do, the market supply curve no longer describes the quantity supplied at any given price.

Being unable to sell all they would like at any given price has ramifications for factor markets. Patinkin (1965) argued that excess supply in the goods market 'spills over' to constrain the demand for labour. Instead of the usual labour demand function (where the quantity of labour demanded increases as the wage decreases), sales-constrained firms demand just enough labour to produce the goods they can sell – regardless of how low wages might fall.

Problems also arise with shortages. When demand exceeds supply individual firms can raise their prices without losing all their sales since competitors cannot saturate the market more than they already do. The competitive model assumes that firms pass up this opportunity to exploit their market power.

In sum, when the market is not in equilibrium, the competitive model assumes that market participants continue to act as if it is; they do not exploit all their market opportunities; they do not maximize their profit or utility. This problem becomes more serious the longer the disequilibrium persists. But the model is silent on how fast prices adjust towards equilibrium. The competitive model offers no theory of how prices adjust out of equilibrium. Indeed, there is no theory of price setting in perfect competition at all.

> Question for your professor: If everyone is a price-taker in the competitive demand and supply model, who makes prices fall when there is a surplus?

The requirements for perfect competition are mutually incompatible In 1926, Piero Sraffa, a young Italian economist at Cambridge, made some very inconvenient observations about the supply and demand theory of perfect competition. In particular he argued that the conditions necessary for independence between the supply and demand curves are incompatible with the conditions necessary for large numbers of firms in the industry.

Consider a movement to the right along an industry supply curve. As the industry's output increases, it uses more factors of production. Suppose that this increased usage of factors drives up the price for at least one factor of production – say Factor X. If substitute goods (or complementary goods) also use Factor X, their costs rise and so do their prices. But an increase in the price of a substitute good shifts the demand for the original good (to the right). Thus a movement along the industry supply curve causes a shift in the industry demand curve. Supply and demand would not be independent of each other, and yet they must be if the framework is to provide a clear and determinate result.

We can fix this problem by assuming that perfectly competitive industries do

not influence the prices of any of the inputs they use. This guarantees independence of supply and demand. But this solution opens up a different problem: what is going to limit the size of the firm? If all factors are available at a constant price, why can't firms duplicate plants and grow without limit? If they can, there is nothing to guarantee that there will be large numbers of small-sized firms in the industry – a requirement of perfect competition.

Sraffa's critique led to the development of the model of imperfect competition: many firms, each selling a differentiated product. What limits the size of the firm in this context is that each faces a downward-sloping demand curve.

> Question for your professor: What's your take on Sraffa's (1926) critique of the competitive model?

Multiple equilibria Nothing guarantees that the demand and supply curves are linear. They might have backward-bending regions, giving rise to multiple intersection points. For example, Prasch (2008: 88) argues that when needs are an important consideration, the labour supply curve could look like that shown in the left-hand diagram of Figure 3.7.

The standard story (often found in textbooks) describes the section of the curve above W_S. As wages rise, the opportunity cost of leisure increases, causing people to substitute leisure for more work. This is the effect that initially dominates between W_S and W_L. On the other hand, since leisure is a normal good, people want to 'buy' more leisure as their incomes rise. When wages get high enough, this income effect dominates, leading to a backward-bending section above W_L.

Prasch supplements this standard story by considering what happens when

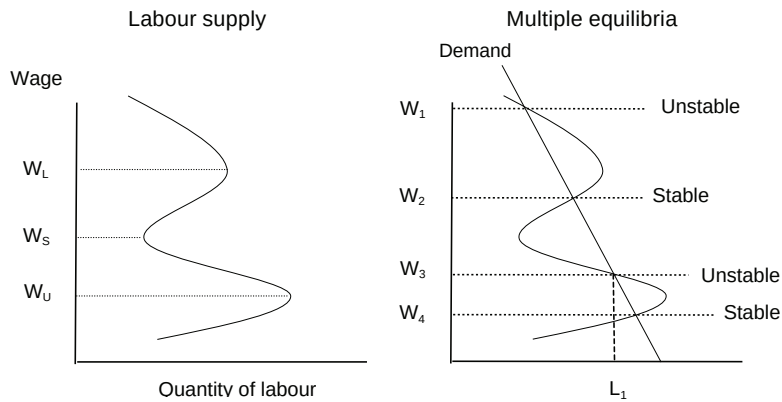

FIGURE 3.7 Multiple equilibria in the labour market

wages fall towards (and even go below) subsistence levels. He argues that when wages go below the level necessary to maintain minimum living standards with normal working hours, households increase their labour supply to abnormal levels. They might hold two jobs or work fourteen-hour shifts. So, below the subsistence wage, which we assume to be W_S, labour supply increases, accounting for another bend in the labour supply curve at W_S.

As wages continue to fall, they will eventually reach the point where the total hours of work required to maintain a socially acceptable standard of living are too long to be sustainable. Below the unsustainable wage, W_U, working hours fall precipitously. Prasch says (ibid.: 88): 'the primary worker and his or her family will be forced by exhaustion, disease, despair, and disrepair to abandon their effort to maintain a standard of living consistent with effective membership in the labour force and, consequently, civil society. They become homeless, petty thieves, or beggars, with strong prospects for a relatively short and miserable life.' This explains the third bend occurring at W_U.

When we confront this labour supply function with a standard downward-sloping labour demand function, we get four possible equilibrium points, as shown in the right-hand diagram of Figure 3.7. Of these, both W_1 and W_3 are unstable. (At a wage slightly below either of these levels, supply exceeds demand, causing wages to continue to fall; similarly, at a wage slightly above either of these wage levels, demand exceeds supply, causing wages to continue to increase.) This leaves two stable equilibria – one of which offers wages quite a bit higher than the subsistence level, W_2; while the other is a poverty trap where wages are substantially below subsistence, W_4.

Prasch uses this construction to show the potential usefulness of minimum wage laws and maximum hours provisions. Either a minimum wage set above W_3, or maximum hours restriction set below L_1, would preclude the poverty trap equilibrium. Interestingly, in this model the legislation pushes the economy to a desirable equilibrium, but once at this equilibrium neither restriction appears 'binding'. That is to say, the equilibrium wage would be above the legal minimum wage and the offered hours would be less than the legal maximum. Prasch notes that this is 'a nice illustration of how market forces can interact with legislation to bring about results that are not immediately evident or expected' (ibid.: 93).

As we shall explain in Chapter 8, there are reasons to think that the demand for labour could also have points where it switches its slope.

Multiple equilibria might also arise out of imperfect information. Stiglitz (2002) argues that if there is a lack of information about quality differences between workers (or goods), then all those workers (or goods) will be lumped into a general category and sell for a wage (or price) that reflects the average quality. Clearly, those selling the better quality have an incentive to try to demonstrate this – to get the information out there – so they can command a premium price.

Conversely, those selling the inferior quality have an incentive to impede the flow of information, to sow confusion and doubt. This leads to the possibility of multiple equilibria, 'one in which information was fully revealed (the market identified the high and low ability people) and another in which it was not (called a pooling equilibrium)' (ibid.: 471).

> Questions for your professor: If markets have multiple equilibria, are some more desirable than others? Is there a role for government in attaining the more desirable ones?

Self-fulfilling prophecies Yet another source of non-uniqueness arises from self-fulfilling prophecies as illustrated in Figure 3.8. Expected future prices influence both the demand and supply curves. Suppose both consumers and producers expect future prices to increase by 10 per cent. Consumers will try to buy more now before prices increase, thus shifting up the demand curve from D1 to D2. Producers will withhold sales now in the expectation of getting higher prices in the future, thus shifting the supply curve left from S1 to S2. It is possible that these shifts will increase the price from P1 to P2 by precisely 10 per cent. If so, there has been a self-fulfilling prophecy: the price is what it is because that's the price we expect. If we had expected a price 40 per cent lower, the price would be 40 per cent lower. Models that embody self-fulfilling prophecies have been used predominantly in macroeconomics to explain instability in aggregate economic activity (Farmer 1993).

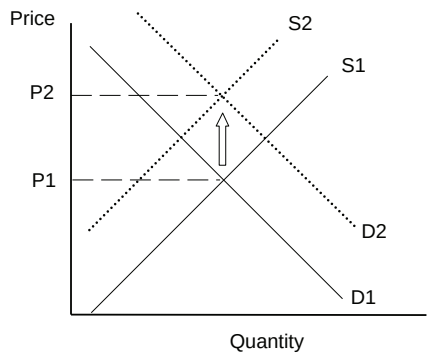

FIGURE 3.8 Self-fulfilling prophecies

> Question for your professor: Changes in expectations about future prices shift both the demand and supply curves. But then what's efficient about prices being at whatever level we expect them to be?

© Andy Singer

Destabilizing speculation and bubbles We've had the Japanese property and stock market bubble (which burst in 1990), the technology stock bubble (which burst in 2001), the Chinese stock market bubble (which burst in 2008) and housing price bubbles in numerous countries which precipitated the financial collapses that began in 2008. Imperfect information is an understatement when it comes to thinking about the future. Yet the textbooks scarcely mention issues of time and uncertainty, the role of speculators or the possibility of price bubbles (i.e. unsustainable price increases driven by expectations that end in a price collapse).

Perhaps some of these issues are beyond the scope of first-year textbooks, but the role of speculators is important enough to warrant consideration. Suppose there are ongoing price fluctuations – for instance, a cycle of boom and bust in commodity prices; what role do the speculators play? Do they make things better or worse? The traditional textbook answer is that if speculators make money, they must buy low and sell high. This extra buying when prices are low, and extra selling when prices are high, implies their activity must act to smooth price fluctuations. Hence, speculators add to the efficiency of markets.

But Mullainathan and Thaler (2004) explain that economists now realize that there are important limits to this argument. First, in the face of irrational traders, the speculator may privately benefit more from trading that helps push prices in the wrong direction than from trading that pushes prices in the right direction. Put another way, it may often pay 'smart money' to follow 'dumb money' rather than to lean against it (Russell and Thaler 1985). For example, if speculators buy when prices are rising, and sell when prices are falling, they could still make money but would add to the amplitude of the price fluctuation.

So, markets per se cannot be relied upon to make rational economic decisions – not even when they are competitively structured.

Scandals Numerous examples of corporate misbehaviour have been documented, from the accounting scandals at Enron in 2001 and Worldcom in 2002 to the unsupervised creation and trading of financial assets (consisting in part of the now-notorious sub-prime mortgages) so complex that no one really knew what they ultimately consisted of or what they were really worth.

These problems arise where there is imperfect and asymmetric information. Using the competitive model as a generic tool for all markets obscures the importance of information imperfections and the legal and regulatory framework that's necessary to oversee markets and make sure they work as we want them to. More on this in Chapters 6 and 8.

> Question for your professor: The world price of oil hit an all-time high of $147 in July 2008. Many believed that this was in part driven by speculators. How does speculation fit into the demand and supply model?

The legal framework: eviction protection legislation The standard textbook world is implicitly one of perfect information and contracts that are costless to negotiate and enforce. The legal framework within which markets operate gets scarcely a passing mention. These assumptions certainly simplify the discussion, but they also impart a subtle laissez-faire message hiding between the lines of the text itself. It implicitly says that the legal and regulatory framework is (at most) of secondary importance.

The nature of the legal and regulatory framework is crucially important for the efficient functioning of markets, as we'll see repeatedly throughout this *Anti-Textbook*. An example relevant to the rental housing market is the nature of eviction protection legislation. Should tenants be liable for eviction after failing to pay one month's rent? If not, after how many? Should the rule be modified in the depths of winter? Does it matter how high the general level of unemployment is? The wrong balance in eviction protection legislation can create an imbalance in the rental housing market as severe as the first-generation rent freeze did in New York City. The case study here is Paris.

During the severe recession of the early 1980s, many people lost their jobs and became unable to pay their rent. People were being thrown out on to the streets. To prevent that, legislation was passed that gave tenants increased eviction protection that shifted the balance of power between landlord and tenant so much in favour of the tenant that it resulted in only about 70 per cent of all

rents being paid.[9] Landlords had to embark upon months and sometimes years of legal wrangling to evict tenants who defaulted on their rent.

The consequences were in many ways similar to a binding rent freeze: an increase in the quantity of units demanded, a decrease in the quantity supplied, and an excess demand for units. The one difference was that since better-quality tenants were less liable to default on their rent, landlords did have an incentive to upgrade their units (Myatt 2004). Clearly, Paris had the balance wrong in its eviction protection legislation. But what is the right balance?

> Questions for your professor: Is the legal framework within which markets operate important in determining the efficiency of markets? Is this ever going to be discussed?

2.6 Summing up

Using the competitive model as a generic tool applicable to a broad range of markets irrespective of the number of producers, heterogeneity of the product or information imperfections creates an inbuilt bias against government market intervention. It loads the dice against rent controls and minimum wages.

The textbooks justify the generic application of the competitive model because it supposedly gives accurate predictions and because it is simpler than non-competitive alternatives. But the claim about predictive power is backed up by only cursory empirical evidence (minimum wages and rent control), and sometimes by no evidence at all (the proportionate burden of the sales tax). Further, the predictive power of the perfectly competitive model is not compared against that of alternative models. The key issue should be: which model better applies to any given situation? Answering this would require comparing the full array of predictions and a serious look at the evidence.

With regard to the claim that the competitive model is simpler than non-competitive alternatives, no criteria are proposed to evaluate it. It is a subjective judgement call, but not one shared by all members of the profession (Holt 1992). Alone it is not enough to justify the generic use of the competitive model. This is obvious once one considers coming to an alternative judgement call – that the monopoly model is the simplest market structure. Would that then justify applying the monopoly model generically throughout the whole economy? It certainly would not give the required impression of a well-functioning self-regulated market system. Stiglitz concludes that the competitive paradigm has survived so long 'partly because the belief in that paradigm, and the policy prescriptions that were derived from it, has served certain interests' (2002: 488). In other words, it is an 'enabling myth'. Certainly, the overemphasis given to it in the textbooks can hardly be explained in any other way.

Suggestions for further reading

For a critique of the mainstream textbook treatment of rent controls see Arnott (1995). Krueger (2001) is an excellent source of information about how to teach the effect of minimum wages given the mixed empirical evidence. The whole of Prasch's (2008) little book, *How Markets Work: Supply, demand and the 'real world'* (2008), is worth reading. Particular emphasis could be put on Lectures II to VI, pages 29 to 111.

ADDENDUM: THE INDETERMINATE AND UNSTABLE ECONOMY

This brief addendum considers two questions: Is there likely to be a unique (just one) equilibrium for the economy as a whole? And if the economy is not in equilibrium, is there some price adjustment process that will bring it to equilibrium? While these are questions that are normally considered in upper-level courses, the concepts should be understandable even to introductory students.

General equilibrium and partial equilibrium 'General equilibrium' is when all the markets in an economy are in equilibrium. For example, the 'production possibilities frontier' presented in Chapter 2 is a general equilibrium model of a very simple two-good economy. In this construct, impacts on wheat explicitly have implications for cloth. Both markets are simultaneously in equilibrium. 'Partial equilibrium' looks at just one market at a time, as in the supply and demand model of Chapter 3. For the most part, introductory microeconomics courses use a partial equilibrium approach.

Multiple equilibria and why they matter In Figure 3.7 we illustrated a situation in which an individual market had several possible equilibria, two of which were stable. That meant it was not possible to predict where the market price and quantity might end up. It might, perhaps, require knowing where the market price was originally. In such a situation, it might be possible to take action to achieve the most desirable equilibrium.

The same result can hold for the economy as a whole. That is, there could be many possible equilibria in a general equilibrium model, some of which might be stable and others unstable. Indeed, if the whole economy were to consist of only competitive markets, like the ones in Chapter 3, the Sonnenschein-Mantel-Debreu (or SMD) Theorem implies that the simultaneous equilibrium may not be unique. As a result, it's impossible to say which of the possible equilibria is the one at which the economy would settle (Ackerman 2002: 121). The economy is fundamentally indeterminate. Occasionally the SMD Theorem is referred to as the 'Anything Goes Theorem'.

This may sound like an abstract, technical point of no real relevance, but

that's not the case. Economists, when trying to assess the effects of policy changes, sometimes make computer simulation models of the entire economy. Naturally, like any economic model, these general equilibrium computer simulation models are highly simplified descriptions of the real economy. Do these models miss the possibility of multiple equilibria? If so, a researcher could simulate the effect of a policy change (implementing free trade, for example) and reach one conclusion, while perhaps the economy would actually end up in quite a different position. Yet, as Hildenbrand and Kirman observe (1988: 49): 'Almost all of the economic literature, theoretical and applied, turns around models in which the nature of "the equilibrium" is discussed and analysed,' as if that equilibrium were unique.

Does the economy find its way to equilibrium? The question is whether, when the economy is not in equilibrium, some price adjustment process returns the economy to equilibrium. In introductory economics, in the partial equilibrium context, students are told (as in Part One of this chapter) that a market can be brought back into equilibrium by lowering the price when there is excess supply (and raising it if there is excess demand). But Hildenbrand and Kirman (ibid.: 49) note that 'as soon as we leave the two-good case this is no longer true'. To explain, they say: 'Think for a moment of two goods, cars and gasoline. Suppose prices were such that cars were in excess demand and gasoline in excess supply. Normal behaviour ... would be to raise the price of cars and lower that of gasoline' (ibid.: 105). But raising the price of cars also lowers demand for gasoline, increasing the excess supply of gasoline, while lowering gas prices raises the demand for cars, increasing excess demand there. Price adjustments may lead around in circles, with differences between demands and supplies not approaching zero. Ackerman (2002: 122) reviews the issue and the literature.

In the final analysis, the competitive economy is neither determinate nor stable.