Chapter 34: Asymmetric Information

34.1: Introduction

So far in this course we have assumed that all agents in the market have the same information about the good being traded. In this final chapter we consider the problems that arise if some agents have more information than others. The usual case is that the seller knows more about the product being traded than the buyer - an obvious example is the used car market, which is the example with which we start. However the situation can be the other way round - for example in the insurance market: the buyers of insurance usually know more about their risks than the insurance company which is selling them insurance. This is the second case we consider. A final example, with which you will be familiar, is the labour market - when trying to sell your labour (get a job) you know much more about how good you are than the potential employers. You have to convince them that you are good. I hope that this book helps you to do so.

34.2 The Market for Used Cars

We start with an extremely simple and unrealistic example and then generalise it. We consider the market for used cars (known as the *market for lemons* through the title of an article written by the Nobel Prize winning economist Akerlof). As you know, it is extremely difficult to tell the difference between a used car that is in good condition and one that is in bad condition. In practice there are a variety of devices that are used to help us distinguish the two (such as the appearance of the car and of the people selling it, the existence of otherwise of a guarantee with the car, and so on) but we consider here an extremely simple scenario in which there is absolutely nothing to help a potential buyer decide whether the car is of good quality or not. To keep things really simple, let us assume that there are just two possibilities – the car is either good or it is bad, but the potential buyer can not determine *ex ante* which it is. We assume that the seller knows however, and we investigate in this section the implications of this asymmetry of information.

The willingness of the seller to sell – as expressed in his or her *reservation price* for the car – obviously depends on whether the car is good or bad. Let us assume that the reservation price is £10,000 if the car is good and £5,000 if it is bad. However the seller of a bad car obviously has no incentive to reveal that it is bad – if he or she can pass it off as good and sell it at a higher price as a consequence, so much the better.

Similarly the willingness of a buyer to buy – as expressed in his or her *reservation price* for a car – obviously depends upon whether the car is good or bad. Let us assume that the reservation price is £12,000 if it is good and £6,000 if it is bad. But recall that the buyer does not know whether the car he or she is looking at is good or bad. Suppose, however, that half the sellers are selling good cars and half are selling bad cars – *and that everyone knows that*.

If a buyer is risk neutral then the most that he or she would be willing to pay for a used car is his or her *expected reservation price* – that is £(½ 12,000 + ½ 6,000) = £9,000. But if the buyer knows the reservation prices of the seller then he or she will be able to work out that at that price *only the sellers of bad cars will be willing to sell them* - because £9,000 is below the reservation price of the seller of a good car. So the buyer is able to work out that all the cars on offer must be bad cars – and

is therefore only willing to pay up to £6,000 for a car. In this market only the bad cars will be traded – at a price between £5,000 and £6,000 – and none of the good cars will be exchanged.

This obviously is inefficient – the market has half collapsed. Only the bad cars get traded – the good cars do not get traded. And the reason – the asymmetry of the information. The sellers know the quality of the good being traded but the buyers do not. Notice that it is difficult to find a way round the collapse of the market using just price information. You might argue that the sellers of the good cars just need to ask a price between £10,000 and £12,000 – and that doing so will signal that the cars they are selling are of high quality – but the problem with this is that the sellers of bad cars have even more of an incentive to try and sell their cars between £10,000 and £12,000. So a high price tag on a car does not *guarantee* that the car is good.

This example suggests that the asymmetry of information causes the market to half collapse. In practice we can clearly see some markets which have collapsed in this way, but usually ways are found to get round this asymmetry of information. These ways usually involve the passage of time in some way: by car sellers establishing reputations[1] and offering guarantees of some form[2]. We shall consider other such devices later in the chapter.

In the meantime let us slightly generalise this example. In it there were just two qualities. In practice there is a range of qualities. Let us make the natural assumption that the reservation price of the sellers increases with the quality – so that the supply is an increasing function of the price, and also that the average quality of the goods on offer increases as the prices increases. This two-fold effect of the price on supply and quality means that the effect of price on demand is also two-fold. First we have the usual price effect – as price rises the demand falls. Second we have the quality effect – as the price rises so does the average quality and hence the demand. These two effects counteract each other and the net effect depends upon the relative strengths of the two effects. It seems reasonable that as the price increases the demand increases initially but then falls. The demand and supply schedules might then look like figure 34.1.

34.1: the market for used cars

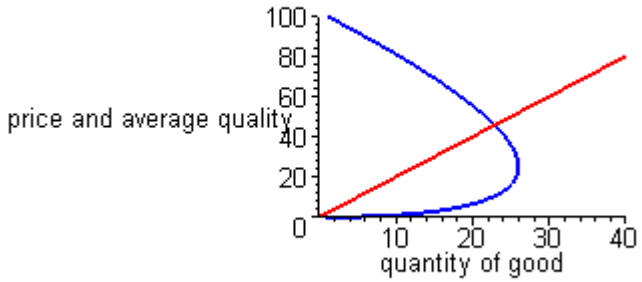

If this happens we have a unique equilibrium where the two curves intersect. We note that, because the quality increases with the price, that the lowest qualities are sold but not the highest qualities. In a sense this is a generalisation of our earlier result – the asymmetry of information could lead to a market in which the high quality goods do not get traded.

[1] Note that 'fly-by-night' sellers usually sell at very low prices, while your established car showroom asks high prices. And what can you infer about the quality of the goods they respectively sell?

[2] It is interesting to note that guarantees are usually not perfect – and that is because the seller is not sure about the quality of the goods that they are selling.

Things could be worse, however. Consider the possibilities below – both of which are consistent with the above story. In figure 34.2 the adverse price effect outweighs the positive quality effect 'too soon' – so that the demand curve fails to intersect the supply curve. In this market there is no equilibrium – no trade takes place. In figure 34.3, the demand curve is such that it intersects twice

34.2: the non-market for used cars

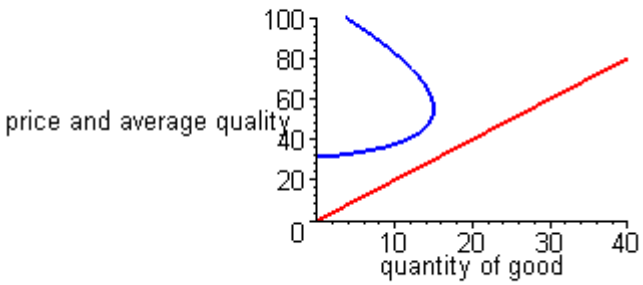

34.3: the multiple-equilibrium-market for used cars

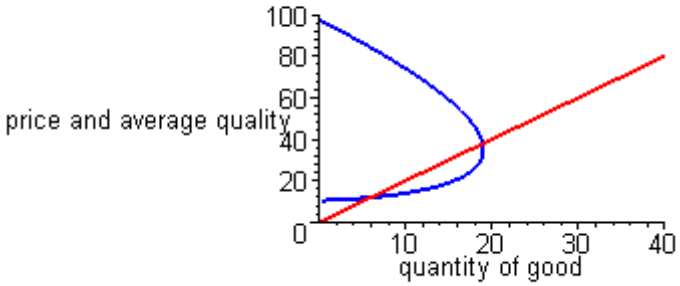

with the supply curve and we have two intersections – one at a low price and one at a higher price. This case is interesting and one might ask whether one of the two equilibria is more 'realistic' than the other. One might be tempted to answer this by adding some kind of *ad hoc* story to explain how the market adjusts when it is not in equilibrium. You should note that we have never really explained how this happens – indeed it is outside the scope of this course, which merely asks whether an equilibrium exists. We have never tried to explain whether and how it is attained[3], and I am a bit reluctant to do so now. But one story which has some plausibility to it is the following: if the price is not an equilibrium price, then if the demand exceeds the supply the price rises, and if the supply exceeds the demand the price falls. This story could be 'justified' by noting that if there is excess demand, then the unsatisfied buyers have an incentive to bid the price up; and if there is excess supply, then the unsatisfied sellers have an incentive to bid the price down. Suppose we accept this story, then can we describe what would happen in the market pictured in figure 34.3? If the price is above the higher-price equilibrium there is excess supply and the price will fall; if the price is between the lower-price equilibrium and the higher-price equilibrium then there is excess demand and the price will rise; if the price is lower than the lower-price equilibrium then there is excess supply and the price will fall. Taken together these suggest that *either* the price will fall to zero (and the market will cease to exist) *or* the price will move towards the higher-price equilibrium. Even in this equilibrium only the lower quality goods are traded. Whatever happens the high quality goods are not traded – at most the low quality goods.

34.3: The Market for Insurance

Let us now consider a market in which it is the buyers who have more information about the quality of the good being traded – the market for insurance. Here when the insurance company sells insurance to an individual, the individual knows more than the company about the risks which he or she is buying insurance against.

We have already studied the case of the market for insurance with symmetrical information in chapter 25. Let me begin by reminding you of the story. We assume that there are two possible states of the world, state 1 and state 2. We assume initially that the probabilities of these two states are *0.4* and *0.6* respectively *and that both the insurance company and the individual taking out the*

[3] Because it is difficult to formulate a convincing theory of the adjustment of a market out of equilibrium.

insurance know these probabilities. This is the symmetric information case. We assume a perfect (fair) insurance market and a risk-averse individual who has Expected Utility preferences. We assume that the individual has initially an *ex ante* risky income of *10* if state 1 occurs and *70* if state 2 occurs. We then get figure 34.6. Here the point W is the endowment point and the line passing through it with slope equal to *–0.4/0.6* is the fair insurance line[4]. The indifference curves are convex and have a slope equal to *–0.4/0.6* (minus the ratio of the probabilities) along the certainty line because the individual has expected utility preferences. It follows immediately that the optimal point for the individual is the asterisked point: as we already know, the individual chooses to be completely insured and chooses a point on the certainty line.

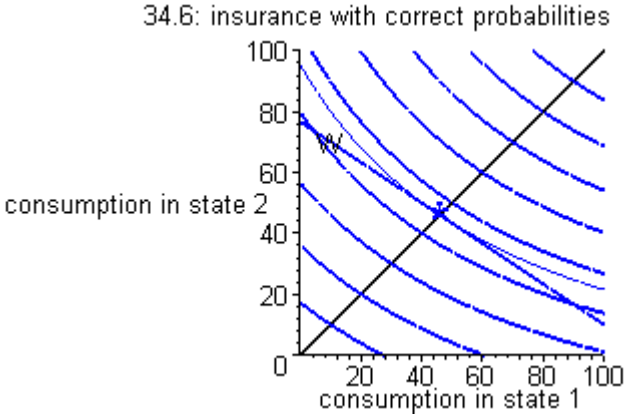

Now there are two possible problems of asymmetric information that may arise in this context. First, with just one individual, the insurance company may not know the true probabilities. Second, with more than one individual who have different probabilities, the insurance company may not be able to distinguish between them. Let us consider these in turn.

The first case may be caused by the presence of the insurance itself. It may be the case that the individual, because he or she is completely covered, is not so worried about whether state 1 happens or not (note that at the chosen point, the individual is insuring *against* state 1). In this case it is possible that the probability of state 1 happening goes up – *unbeknown to the insurance company*. Let us suppose that it rises from *0.4* to *0.5* – but the insurance company continues to think that it is *0.4*. Then we get figure 34.9.

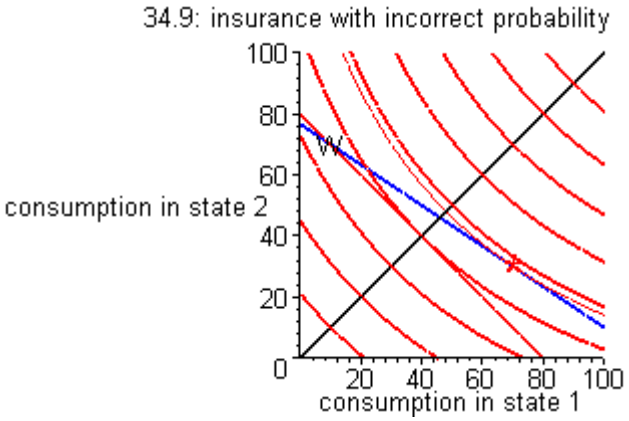

In this figure, point W remains the endowment point and the line passing through that point with a slope equal to *–0.4/0.6* is what the insurance company *thinks* is the fair insurance line – because it is continuing to think that the true probabilities are *0.4* and *0.6*. But the individual knows that the true probabilities are *0.5* and *0.5* and hence his or her indifference curves have a slope equal to *–0.5/0.5*

[4] Recall that the slope of the fair insurance line is minus the ratio of the probability of state 1 to the probability of state 2.

along the certainty line. The individual would choose the asterisked point in the diagram: because the company is offering him more than fair insurance for state 1, he or she now chooses to gamble on state 1 happening. And the insurance company? It loses money – because the fair insurance line is the straight line through W with slope *–0.5/0.5*.

This is the problem known as *moral hazard*. The very existence of the insurance contract makes the individual less careful and changes the probabilities of the two states occurring. The insurance company loses money if it is not aware that this has happened. This is the reason why, in practice, insurance companies take measures to prevent moral hazard: by requiring that householders install burglar alarms and smoke detectors, for example.

Another problem of asymmetric information occurs when there are different potential buyers of insurance – who differ in their riskiness, but who the company finds it difficult to distinguish between. This is called the problem of *adverse selection*, a term which will become clearer shortly.

Suppose that in an insurance market there are two kinds of potential buyers of insurance – which we call *high-risk* and *low-risk*. For the high-risk people the probability the probability of state 1 happening is *0.5*; for the low risk it is *0.4*. We assume that all individuals have the same initial endowment point *(10, 70)* – as point W above. We have the situation pictured in figure 34.10.

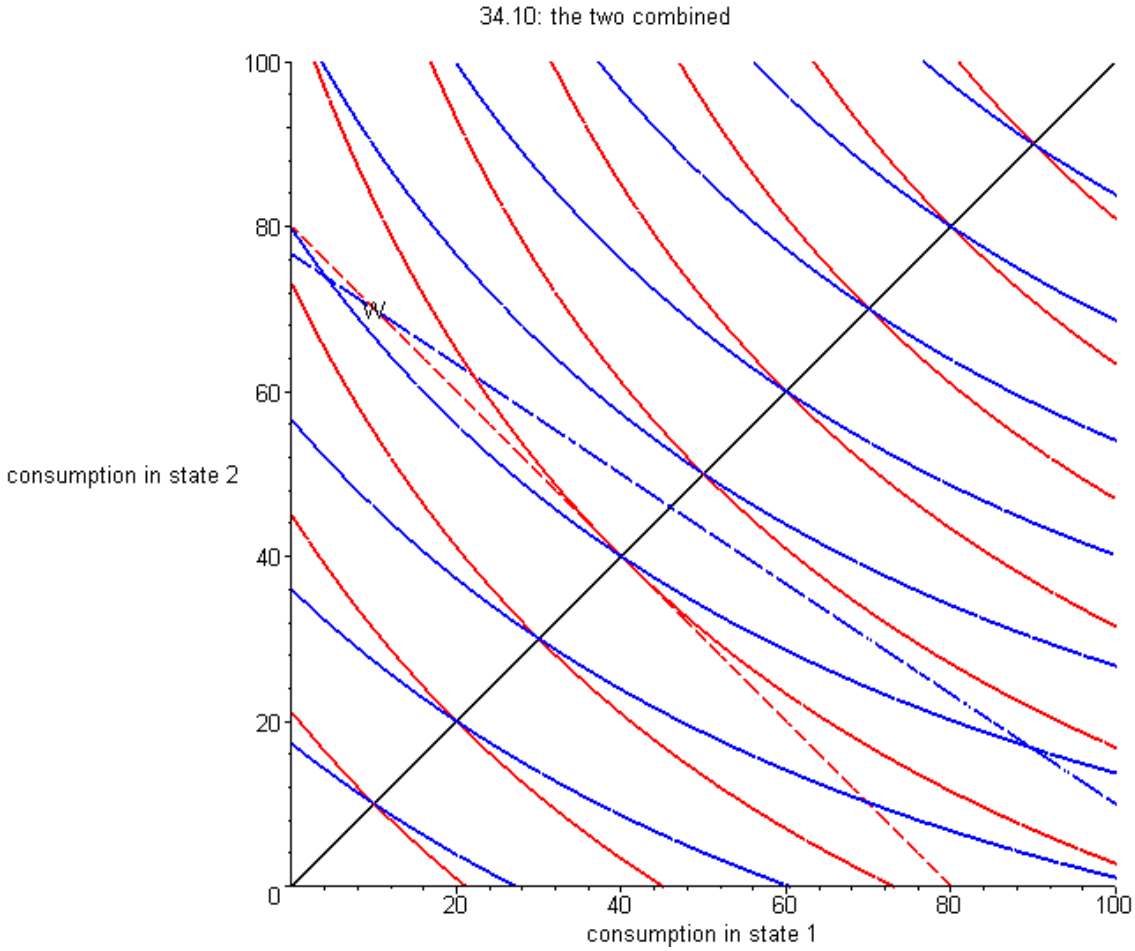

Here the point W is the endowment point for both the high-risk and the low-risk. The straight lines through W are the fair insurance lines for the two groups – the flatter line (with slope *–0.4/0.6*) that appropriate for the low-risk group and the steeper line (with slope *–0.5/0.5*) that appropriate for the high-risk group. The convex curves are the indifference curves of the two groups – those with slope *–0.4/0.6* along the certainty line the indifference curves of the low-risk group and those with slope *–0.5/0.5* along the certainty line the indifference curves of the high-risk group. If the insurance

company can tell the two groups apart, then the solution is simple: the company offers the low-risk budget line to the low-risk group and it offers the high-risk budget line to the high-risk group. Every one gets full insurance. The company breaks even.

But what happens if the company can *not* distinguish between the two groups and has to offer both budget lines to everyone? As we have seen the low-risk group will choose the low-risk line and chose to be fully insured – with this low-risk group the insurance company breaks even. However, the high-risk group will also choose the low-risk line and will choose to bet on state 1 happening (for the reasons we have discussed above). With the high-risk group the insurance company loses money. Unless the company can do something clever, it will decide that, in order not to lose money, that it can only offer the high-risk line. The high-risk group will completely insure – while the low risk group may take out a little insurance (though not complete because the price is unfair to them). This is called a *pooling equilibrium*. Notice that it has the same implications as the market for used cars – the high risk (low quality) drives out the low risk (high quality) participants in the market.

But is there 'something clever' that the insurance company can do? Well, yes, to an extent. It can get the two types to reveal their type. How does it do this? By offering two insurance contracts – one that it knows the low-risk type will accept and the other one that it knows that the high-risk type will accept. The key to this is re-defining a contract. Instead of defining a contract by a budget line, we define it by a point. Consider figure 34.11 and in particular the two asterisked points. These are

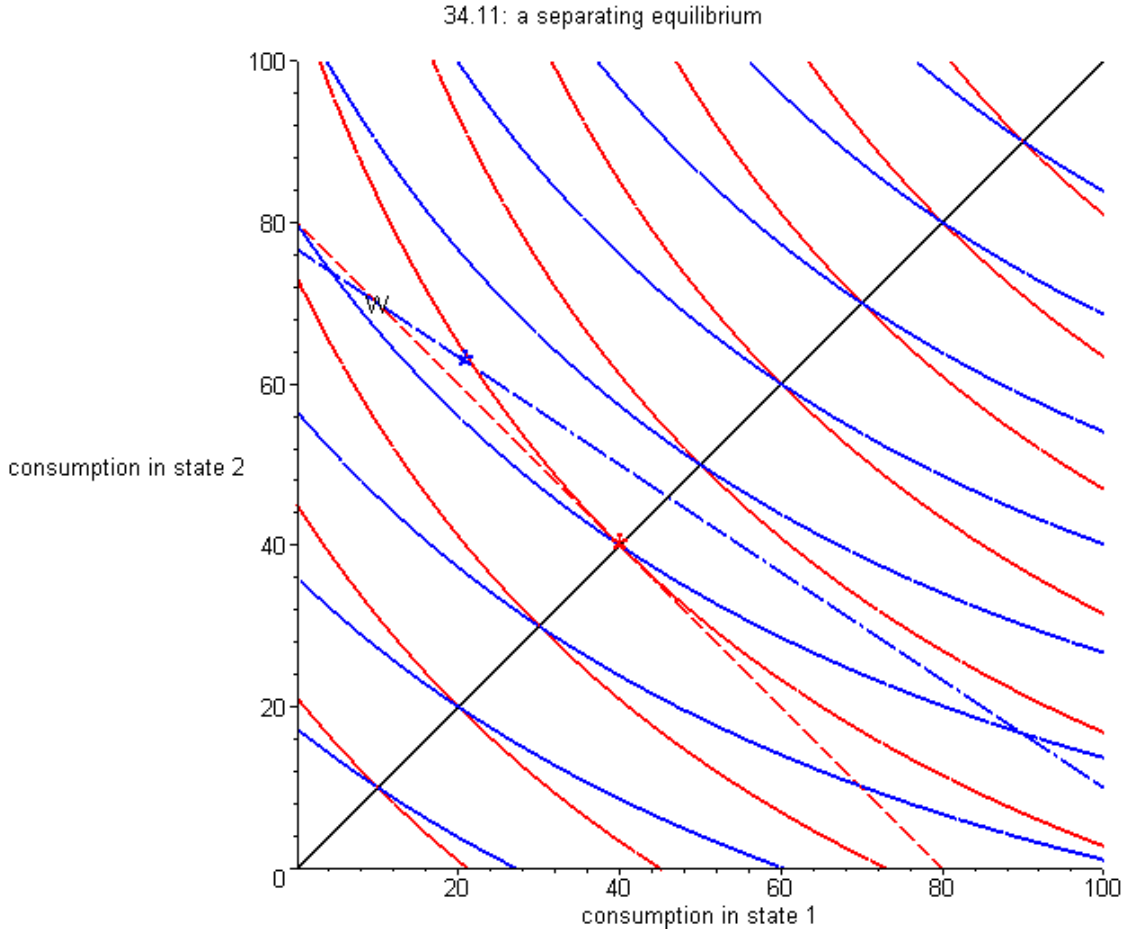

the two contracts the insurance company offers. One of the two contracts is on the certainty line and on the fair insurance line for the high-risk group. This is the contract accepted by the high-risk group. The other contract is on the fair insurance line for the low-risk group and is just to the left of the high-risk indifference curve that passes through the high-risk contract. This is the contract

accepted by the low-risk group. It is not accepted by the high-risk group because is it is "just to the left of the high-risk indifference curve that passes through the high-risk contract". But it is accepted by the low-risk group because it is on a higher low-risk indifference curve than both the point W and the high-risk contract.

This is called a *separating equilibrium*. It separates the two groups. *Ex post*, the company can identify the two groups. The company breaks even because both contracts are on the correct fair insurance line. The only problem is that it does not offer full insurance to the low-risk group. So we still get a residual problem in the presence of asymmetric information.

In practice there are ways round these problems. Many of these ways involve the passage of time and reflect the fact that the contract is repeated many times. You might like to think that if the above story is repeated more than once, then the company can exploit the fact that the choice in the above problem reveals the type – but note, the individuals can anticipate this: if the choice is to be repeated several times it may not be in their interests to reveal their type on the first repetition.

If you look to actual practice, you will see what insurance companies do. Motor insurance is particularly interesting, as it is perhaps here there is considerable uncertainty about how good or bad a driver a particular individual is. In the first place companies try and distinguish between people on the basis of observable characteristics, which are known to be connected with the probability of having an accident – age, sex, etc. In the second place, many companies use No Claim Bonus schemes so that there is an incentive not to have an accident and a reward for not having had one. There are also reductions in premiums for non-smokers and for people who take precautions of various kinds.

34.4: The Labour Market

As you may have already realised the labour market suffers from problems of asymmetric information. When you come to sell your labour you have to convince potential employers that you are better than the next person – and potential employers have to find ways to distinguish between applicants. As far as the second of these is concerned you may be familiar with some of the devices used – interviews, various tests, perhaps even a period of trial employment with some promise of permanent employment.

As to what you can do to convince employees – the usual way is through *signals* of various types. You may include in your CV that you have certain interesting hobbies that make you more employable, that you do certain kinds of voluntary work, and so on. But probably the most important signals that employers look for, and which you can offer, are *qualifications* of various kinds – particularly a degree. There is a high chance that you are doing a degree. Obviously the better the degree that you get the better the signal you put out to potential employers. What is crucial, however, is that employers know that there is value in the signal that you have. They need to know that a first-class mark signals that the student is better than a student with a second-class mark and so on. Good universities make sure that this is the case by making it more difficult to get a first-class mark than a second-class mark. Indeed so it should be. To get a first-class mark you really have to understand the material – not just memorise it – and be good at using it.

I hope that this book helps you understand the material. At the same time, you have to help yourself by *working* through the book – not just reading it. You have to work through the examples I have provided and then make sure you understand the examples and are able to generalise them. Remembering particular examples will not help you – you have to understand the underlying

principles and then be able to apply them in other contexts. The exam questions that I set test these underlying principles – rather than test whether students are able to memorise the text. Often students ask me "where is the answer to this question? We cannot find the answer in any text." I respond by saying "Of course not, if the answer was in the text it would not be a good question. The answer is in your brain." I hope this book helps develop your brain so that you become a good economist – and can put out good signals.

34.5: Summary

We have shown in this chapter that asymmetric information causes problems in markets.

Asymmetric information may cause a market to collapse (disappear completely) or partially collapse (with fewer trades than would be the case with complete information)

We saw that there are some ways around these problems.

There may be the possibility of a separating equilibrium in which information is revealed by the behaviour of agents in the market.

Agents may acquire signals which accurately reveal the information.

In a dynamic world there are other ways round the problems, to which we have alluded but which are too complex to cover in this book.

In practice we see guarantees and reputations which reveal the information.

34.6: Do lemons always have to drive out plums?

This is a simple experiment on the problem of asymmetric information.

You can try it yourself with a group of fellow students. One or more of you should act as 'the experimenter(s)' and the rest of you (make this an even number) as the members of a particular market. The experimenters should organise and implement this asymmetric information experiment a predetermined number of times, which the group as a whole should decide in advance. Each time the following should be implemented.

Half the market members will be designated as (potential) buyers, half as (potential) sellers. Each wants to buy or sell at most one unit. There are two kinds of sellers in the market - but only themselves and the experimenters know which type they are. Indeed the experimenters should find a way of determining (without anyone except the individual seller knowing) whether each seller is one type or the other. The probability of any one seller being one or other of the two types is ½; so the experimenters could (in the vision of only the seller) toss a coin to decide what type he or she is.

The two types are *sellers of a lemon* and *sellers of a plum*. You should first run the experiment with the type of each seller being randomly chosen *each* round of the experiment – so that no-one can infer anything from the fact that a particular seller was the seller of a plum in one period of the experiment as to whether he or she is a seller of a plum or of a lemon in other periods of the experiment.

The following incentive scheme should be put in place by the experimenters, either using real money or hypothetical money (in this latter case, the subjects should imagine that it is real money).

A buyer who buys a plum will be paid £24 but has to pay the agreed price;
A buyer who buys a lemon will be paid £12 but has to pay the agreed price;
A seller of a plum who sells it will be charged £20 but receives the agreed price;
A seller of a lemon who sells it will be charge £10 but receives the agreed price;
Sellers who do not sell will not be charged and buyers who do not buy will not be paid.

Sellers make money by selling at a price above the amount they will be charged whilst buyers make money by buying at a price less than the amount they will be paid. For example, if a buyer buys a plum from a seller at a price of £22.50 then the seller makes a profit of £2.50 from the experiment whilst the buyer makes a profit of £1.50. If a buyer buys a lemon from a seller at a price of £10.50, then the seller makes a profit of £0.50 whilst the buyer makes a profit of £1.50. Of course losses can be made: for example, if a buyer buys a lemon from a seller at a price of £15, then the seller makes a profit of £3 whilst the buyer makes a loss of £5.

The problem is that buyers do not know who the sellers of lemons are, nor who are the sellers of plums. Furthermore only the individual sellers themselves - and the experimenters (WHO WON'T TELL ANYONE) know what type they are.

Now let trade try to take place - by whatever mechanism you as a group have decided: it could be a double auction, it could be shifting bilateral negotiations, etc. It is up to you to decide. Then play the game the predetermined number of times and work out how much each subject has earned. Subjects should, as before, try and earn as much as possible.

After playing the experiment, the group should answer the following questions:

(1) what has this got to do with asymmetric information?
(2) what is the complete information outcome?
(3) did the lemons drive out the plums?
(4) were the sellers of plums able to convince buyers that they *were* the sellers of plums?
(5) were the sellers of lemons able to convince buyers that *they* were the sellers of plums?
(6) did things get better with repetition?

There are various 'variations on a theme' that you could try, and which are informative and interesting. Explicit in the above is that the type of a seller varies randomly from period to period, so that no inferences can be drawn about the type of a seller from his or her type in previous round of the experiment. Clearly, however, in this obviously non-anonymous experiment, people may be able to infer something about the *honesty* of a particular seller by observing his or her behaviour throughout the experiment. A variant of this story – a more realistic variant – is when the seller of a plum remains a seller of a plum throughout all the rounds of the experiment, and a seller of a lemon remains a seller of a lemon throughout all rounds of the experiment. In this case a seller has more to lose by being revealed as dishonest. In this case *reputation effects* become more important, and it may be the case that these effects allow the market to function efficiently – in that plums are sold for higher prices than lemons.

You might then like to explore other devices which help the market operate efficiently – like the offering of *enforceable* guarantees by the sellers or the introduction of *sustainable* legal systems which punish sellers of lemons at high prices. The ultimate objective of the experiment is to see how the problem of asymmetric information may be overcome by sustainable market mechanisms.