

Chapter 9

adverse selection



9.1 Introduction: Production Game VI

In chapter 7, games of asymmetric information were divided between games with moral hazard, in which agents are identical, and games with adverse selection, in which agents differ. In moral hazard with hidden knowledge and adverse selection, the principal tries to sort out agents of different types. In moral hazard with hidden knowledge, the emphasis is on the agent's action rather than his choice of contract because agents accept contracts before acquiring information. Under adverse selection, the agent has private information about his type or the state of the world before he agrees to a contract, which means that the emphasis is on which contract he will accept.

For comparison with moral hazard, let us consider still another version of the Production Game of chapters 7 and 8.

Production Game VI: Adverse Selection

PLAYERS

The principal and the agent.

THE ORDER OF PLAY

- 0 Nature chooses the agent's ability a , observed by the agent but not by the principal, according to distribution $F(a)$.
- 1 The principal offers the agent one or more wage contracts $w_1(q), w_2(q), \dots$.
- 2 The agent accepts one contract or rejects them all.
- 3 Nature chooses a value for the state of the world, θ , according to distribution $G(\theta)$. Output is then $q = q(a, \theta)$.

PAYOFFS

If the agent rejects all contracts, then $\pi_{agent} = \bar{U}(a)$, which might or might not vary with his type, a ; and $\pi_{principal} = 0$.

Otherwise, $\pi_{agent} = U(w, a)$ and $\pi_{principal} = V(q - w)$.

Under adverse selection, it is not the worker's effort, but his ability, that is noncontractible. Without uncertainty (move [3]), the principal would provide a single contract specifying high wages for high output and low wages for low output, but either high or low output might be observed in equilibrium, unlike under moral hazard – if both types of agents accepted the contract. Also, under adverse selection, unlike moral hazard, offering multiple contracts can be an improvement over offering a single contract. The principal might, for example, provide a flat-wage contract for low-ability agents and an incentive contract for high-ability agents.

Production Game VIa puts specific functional forms into the game to illustrate how to find an equilibrium.

Production Game VIa: Adverse Selection with Particular Parameters

PLAYERS

The principal and the agent.

THE ORDER OF PLAY

- 0 Nature chooses the agent's ability a , unobserved by the principal, according to distribution $F(a)$, which puts probability 0.9 on low ability, $a = 0$, and probability 0.1 on high ability, $a = 10$.
- 1 The principal offers the agent one or more wage contracts
 $W_1 = (w_1(q = 0), w_1(q = 10)), W_2 = (w_2(q = 0), w_2(q = 10)) \dots$
- 2 The agent accepts one contract or rejects them all.
- 3 Nature chooses a value for the state of the world, θ , according to distribution $G(\theta)$, which puts equal weight on 0 and 10. Output is then $q = \text{Min}(a + \theta, 10)$.

PAYOFFS

If the agent rejects all contracts, then depending on his type his reservation payoff is either $\bar{U}_{Low} = 3$ or $\bar{U}_{High} = 2$ and the principal's payoff is $\pi_{principal} = 0$. Otherwise, $U_{agent} = w$ and $V_{principal} = q - w$.

Thus, in Production Game VIa, output is 0 or 10 for the low-ability type of agent, depending on the state of the world, but always 10 for the high-ability agent. The agent types also differ in their reservation payoffs: the low-ability agent would work for an expected wage of 3, but the high-ability agents would require just 2. More realistically the high-ability agent would have a higher reservation wage (his ability might be recognizable in some alternative job), but I have chosen $\bar{U}_{High} = 2$ to illustrate an interesting feature of the equilibrium.

A separating equilibrium is

Principal: Offer $W_1 = \{w_1(q = 0) = 3, w_1(q = 10) = 3\}$,
 $W_2 = \{w_2(q = 0) = 0, w_2(q = 10) = 3\}$.

Low agent: Accept W_1 .

High agent: Accept W_2 .

As usual, this is a weak equilibrium. Both Low and High agents are indifferent about whether they accept or reject their contract. The equilibrium indifference of the agents arises from the open-set problem; if the principal were to specify a wage of 3.01 for W_2 , for example, the high-ability agent would no longer be indifferent about accepting it instead of W_1 .

Let us go through how I came up with the equilibrium contracts above. First, what action does the principal desire from each type of agent? The agents do not choose effort, but they do choose whether or not to work for the principal, and which contract to accept. The low-ability agent's expected output is $0.5(0) + 0.5(10) = 5$, compared to a reservation payoff of 3, so the principal will want to hire the low-ability agent if he can do it at an expected wage of 5 or less. The high-ability agent's expected output is $0.5(10) + 0.5(10) = 10$, compared to a reservation payoff of 2, so the principal will want to hire the high-ability agent if he can do it at an expected wage of 10 or less.

In hidden-action models, the principal tries to construct a contract which will induce the agent to take the single appropriate action. In hidden-knowledge models, the principal tries to make different actions attractive to different types of agent, so the agent's choice depends on the hidden information. The principal's problem, as in Production Game V with its moral hazard and hidden actions, is to maximize his profit subject to

- 1 **Incentive compatibility** (the agent picks the desired contract and actions).
- 2 **Participation** (the agent prefers the contract to his reservation utility).

In a model with hidden knowledge, the incentive compatibility constraint is customarily called the **self-selection constraint**, because it induces the different types of agents to pick different contracts. A big difference from moral hazard is that in a separating equilibrium there will be an entire set of self-selection constraints, one for each type of agent, since the appropriate contract depends on the hidden information. A second big difference is that the incentive compatibility constraint could vanish, instead of multiplying. The principal might decide to give up on separating the types of agent, in which case all he must do is make sure they all participate.

Here, the participation constraints are, letting $\pi_i(W_j)$ denote the expected payoff an agent of type i gets from contract j ,

$$\begin{aligned} \pi_L(W_1) &\geq \bar{U}_{Low}; & 0.5w_1(0) + 0.5w_1(10) &\geq 3, \\ \pi_H(W_2) &\geq \bar{U}_{High}; & 0.5w_2(10) + 0.5w_2(10) &\geq 2. \end{aligned} \tag{9.1}$$

Clearly the contracts in our conjectured equilibrium, $W_1 = (3, 3)$ and $W_2 = (0, 3)$, satisfy the participation constraints. In the equilibrium, the low- and the high-output wages both matter to the low-ability agent, but only the high-output wage matters to the high-ability agent. Both agents, however, end up earning a wage of 3 in each state of the world, the only difference being that contract W_2 would be a very risky contract for the low-ability agent despite being riskless for the high-ability agent. The principal would like to make W_1 risk-free, with the same wage in each state of the world.

In our separating equilibrium, the participation constraint is binding for the “bad” type but not for the “good” type, who would accept a wage as low as 2 if nothing better were available. This is typical of adverse selection models (if there are more than two types it is the participation constraint of the worst type that is binding, and no other). The principal makes the bad type’s contract unattractive for two reasons. First, if he pays less, he keeps more. Second, when the bad type’s contract is less attractive, the good type can be more cheaply lured away to a different contract. The principal allows the good type to earn more than his reservation payoff, on the other hand, because the good type always has the option of lying about his type and choosing the bad type’s contract, and the good type, with his greater skill, could earn a positive payoff from the bad type’s contract. Thus, the principal can never extract all the gains from trade from the good type unless he gives up on making either of his contracts acceptable to the bad type.

Another typical feature of this equilibrium is that the low-ability agent’s contract not only drives him down to his participation constraint, but is riskless. An alternative would be to offer the low-ability agent a contract of the form $W'_1 = (w_l, w_h)$, where it still satisfies the participation constraint because $0.5U(w_l) + 0.5U(w_h) \geq 3$. That is easy enough to do in Production Game VIa, because the agents are risk-neutral, and when $U(w) = w$, the low-ability agent would be as happy with $W'_1 = (0, 6)$ as with $W_1 = (3, 3)$. But W'_1 would create a big problem for self-selection, because the high-ability agent would get an expected payoff of 6 from it, since his output is always high. Also, if the agents were risk-averse, the risky contract would have to have a higher expected wage than W_1 , to make up for the risk, and thus would be more expensive for the principal.

Next, look at the self-selection constraints, which are

$$\begin{aligned} \pi_L(W_1) &\geq \pi_L(W_2); & 0.5w_1(0) + 0.5w_1(10) &\geq 0.5w_2(0) + 0.5w_2(10), \\ \pi_H(W_2) &\geq \pi_H(W_1); & 0.5w_2(10) + 0.5w_2(10) &\geq 0.5w_1(10) + 0.5w_1(10). \end{aligned} \tag{9.2}$$

The first inequality in (9.2) says that the contract W_2 has to have a low enough expected return for the low-ability agent to deter him from accepting it. The second inequality says that the wage contract W_1 must be less attractive than W_2 to the high-ability agent. The conjectured equilibrium contracts $W_1 = (3, 3)$ and $W_2 = (0, 3)$ do this, as can be seen by substituting their values into the constraints:

$$\begin{aligned} \pi_L(W_1) &\geq \pi_L(W_2); & 0.5(3) + 0.5(3) &\geq 0.5(0) + 0.5(3) \\ \pi_H(W_2) &\geq \pi_H(W_1); & 0.5(3) + 0.5(3) &\geq 0.5(3) + 0.5(3). \end{aligned} \tag{9.3}$$

The self-selection constraint is binding for the good type but not for the bad type. This, too, is typical of adverse selection models. The principal wants the good type to reveal his type so he makes the appropriate contract as attractive to the good type as the bad type’s contract. It does not have to be *more* attractive though (here notice the open-set problem), so the principal will minimize his salary expenditures and choose two contracts equally attractive to the good type. In so doing, the principal will have chosen a contract for the good type that is strictly worse for the bad type, who cannot achieve a high output so easily.

It is to show how the participation constraint does not have to be binding for the good type that I assumed $\bar{U}_{High} = 2$ for Production Game VIa. If I had assumed $\bar{U}_{High} = 3$, then we would still have $W_2 = (0, 3)$, but the fact that the “3” came from the self-selection

constraint would be obscured. And although it is typical that the good agent's participation constraint is nonbinding and his incentive compatibility constraint is not, it is by no means necessary. If I had assumed $\bar{U}_{High} = 4$, then we would need $W_2 = (0, 4)$ to satisfy the participation constraint as cheaply as possible, so it would be binding, and then the self-selection constraint would not be binding. Despite all this, modellers most often expect to find the bad type's participation constraint and the good type's self-selection constraint binding in a two-type model, and the worst agent's participation constraint and all other agents' self-selection constraints in a multitype model.

Once the self-selection and participation constraints are satisfied, weakly or strictly, the agents will not deviate from their equilibrium actions. All that remains to be checked is whether the principal could increase his payoff. He cannot, because he makes a profit from either contract, and having driven the low-ability agent down to his reservation payoff and the high-ability agent down to the minimum payoff needed to achieve separation, he cannot further reduce their pay.

Competition and Pooling

As with hidden actions, if principals compete in offering contracts under hidden information, a **competition constraint** would be added: the equilibrium contract must be as attractive as possible to the agent, since otherwise another principal could profitably lure him away. An equilibrium may also need to satisfy a part of the competition constraint not found in hidden actions models: either a **nonpooling constraint** or a **nonseparating constraint**. If one of several competing principals wishes to construct a pair of separating contracts C_1 and C_2 , he must construct it so that not only do agents choose C_1 and C_2 depending on the state of the world (to satisfy incentive compatibility), but also they prefer (C_1, C_2) to a pooling contract C_3 (to satisfy nonpooling). We only have one principal in Production Game VI, though, so competition constraints are irrelevant.

Although it is true, that the participation constraints must be satisfied for agents who accept the contracts, it is not always the case that they accept different contracts in equilibrium, and if they do not, they do not need to satisfy self-selection constraints.

*If all types of agents choose the same strategy in all states, the equilibrium is **pooling**. Otherwise, it is **separating**.*

The distinction between pooling and separating is different from the distinction between equilibrium concepts. A model might have multiple Nash equilibria, some pooling and some separating. Moreover, a single equilibrium – even a pooling one – can include several contracts, but if it is pooling the agent always uses the same strategy, regardless of type. If the agent's equilibrium strategy is mixed, the equilibrium is pooling if the agent always picks the same mixed strategy, even though the messages and efforts would differ across realizations of the game.

These two terms came up in section 6.2 in the game of PhD Admissions. Neither type of student applied in the pooling equilibrium, but one type did in the separating equilibrium. In a principal–agent model, the principal tries to design the contract to achieve separation unless the incentives turn out to be too costly. In Production Game VI, the equilibrium was separating, since the two types of agents choose different contracts.

A separating contract need not be fully separating. If agents who observe a state variable $\theta \leq 4$ accept contract C_1 but other agents accept C_2 , then the equilibrium is separating but it does not separate out every type. We say that the equilibrium is **fully revealing** if the agent's choice of contract always conveys his private information to the principal. Between pooling and fully revealing equilibria are the **imperfectly separating** equilibria synonymously called **semi-separating**, **partially separating**, **partially revealing**, or **partially pooling** equilibria.

The possibility of a pooling equilibrium reveals one more step we need to take to establish that the proposed separating equilibrium in Production Game VIa is really an equilibrium: would the principal do better by offering a pooling contract instead, or a separating contract under which one type of agent does not participate? All of my derivation above was to show that the agents would not deviate from the proposed equilibrium, but it might still be that the principal would deviate.

First, would the principal prefer pooling? Then all that is necessary is that the contract as cheaply as possible induce both types of agent to participate. Here, that would require that we make the contract barely acceptable to the type with the lowest ability and highest reservation payoff, the low-ability agent. The contract (3, 3) offered by itself would do that, but it would not increase profits over W_1 and W_2 in our equilibrium above. Either pooling or separating would yield profits of $0.9(0.5(0 - 3) + 0.5(10 - 3)) + 0.1(0.5(10 - 3) + 0.5(10 - 3)) = 2.5$.

Second, would the principal prefer a separating contract that "gave up" on one type of agent? The principal would not want to drive away the high-ability agent, of course, though he could do so by offering a high wage for $q = 0$ and a low wage for $q = 10$, because the high-ability agent has both greater output and a lower reservation payoff (if we had $\bar{U}_{High} = 11$ then the outcome would be different). But if the principal did not have to offer a contract that gave the low-ability agent his reservation payoff of 3, he could be more stingy towards the high-ability agent. If there were no low-ability agent, the principal would offer a contract such as (0, 2) to the high-ability agent, driving him down to his reservation payoff and increasing the profits from hiring him. Here, however, there are not enough high-ability agents for that to be a good strategy for the principal. His payoff would decline to $0.9(0) + 0.1(0.5(10 - 2) + 0.5(10 - 2)) = 0.8$, a big decline from 2.5. If 99 percent of the agents were high-ability, instead of 10 percent, things would have turned out differently, but there are too many agents who have low-ability yet can be efficiently hired for the principal to give up on them.

The Production Game is one setting for adverse selection, and is a good foundation for modelling it, but the best-known setting, and one which well illustrates the power of the idea in explaining everyday phenomenon, is the used-car market. We will look at that market in the next few sections. All adverse selection games are games of incomplete information, but they might or might not contain uncertainty, moves by Nature occurring after the agents take their first actions. We will continue using games of certainty in sections 9.2 and 9.3, and wait to look at the effect of uncertainty in section 9.4. The first game will model a used car market in which the quality of the car is known to the seller but not the buyer, and the various versions of the game will differ in the types and numbers of the buyers and sellers. Section 9.4 will return to models with uncertainty, in a model of adverse selection in insurance. One result there will be that a Nash equilibrium in pure strategies fails to exist for certain parameter values. Section 9.5 applies the idea of adverse selection to explain the magnitude of the bid-ask spread in financial markets, and section 9.6 touches on a variety of other applications.

9.2 Adverse Selection under Certainty: Lemons I and II

Akerlof stimulated an entire field of research with his 1970 model of the market for shoddy used cars (“lemons”), in which adverse selection arises because car quality is better known to the seller than to the buyer. In agency terms, the principal contracts to buy from the agent a car whose quality, which might be high or low, is noncontractible despite the lack of uncertainty.

We will spend considerable time adding twists to a model of the market in used-cars. The game will have one buyer and one seller, but this will simulate competition between buyers, as discussed in section 7.2, because the seller moves first. If the model had symmetric information there would be no consumer surplus. It will often be convenient to discuss the game as if it had many sellers, interpreting one seller to whom Nature randomly assigns a type to be a population of sellers of different types, one of whom is drawn by Nature to participate in the game.

The Basic Lemons Model

PLAYERS

A buyer and a seller.

THE ORDER OF PLAY

- 0 Nature chooses quality type θ for the seller according to the distribution $F(\theta)$.
The seller knows θ , but while the buyer knows F , he does not know the θ of the particular seller he faces.
- 1 The buyer offers a price P .
- 2 The seller accepts or rejects.

PAYOFFS

If the buyer rejects the offer, both players receive payoffs of zero.
Otherwise, $\pi_{buyer} = V(\theta) - P$ and $\pi_{seller} = P - U(\theta)$, where V and U will be defined later.

The payoffs of both players are normalized to zero if no transaction takes place. (A normalization is part of the notation of the model rather than a substantive assumption.) The model assigns the players’ utility a base value of zero when no transaction takes place, and the payoff functions show changes from that base. The seller, for instance, gains P if the sale takes place but loses $U(\theta)$ from giving up the car.

The functions $F(\theta)$, $U(\theta)$, and $V(\theta)$ will be specified differently in different versions of the game. We start with identical tastes and two types (Lemons I), and generalize to a continuum of types (Lemons II). Section 9.3 specifies first that the sellers are identical and value cars more than buyers (Lemons III), next that the sellers have heterogeneous tastes (Lemons IV). We will look less formally at other modifications involving risk aversion and the relative numbers of buyers and sellers.

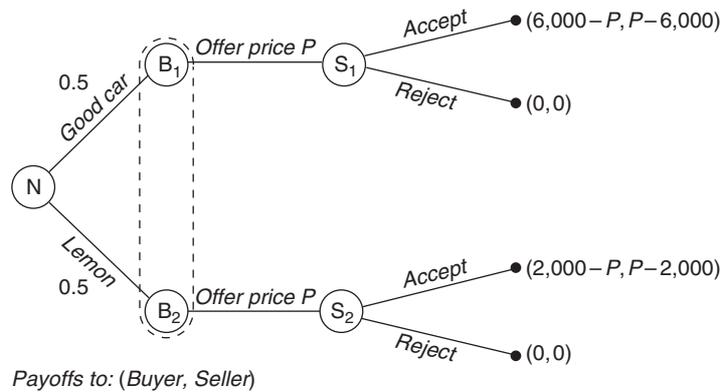


Figure 9.1 An extensive form for Lemons I.

Lemons I: Identical Tastes, Two Types of Sellers

Let good cars have quality 6,000 and bad cars (lemons) quality 2,000, so $\theta \in \{2,000, 6,000\}$, and suppose that half the cars in the world are of the first type and the other half of the second type. A payoff profile of $(0, 0)$ will represent the status quo, in which the buyer has \$50,000 and the seller has the car. Assume that both players are risk-neutral and they value quality at one dollar per unit, so after a trade the payoffs are $\pi_{buyer} = \theta - P$ and $\pi_{seller} = P - \theta$. Figure 9.1 shows the extensive form.

If he could observe quality at the time of his purchase, the buyer would be willing to accept a contract to pay \$6,000 for a good car and \$2,000 for a lemon. He cannot observe quality, however, and we assume that he cannot enforce a contract based on his discovery once the purchase is made. Given these restrictions, if the seller offers \$4,000, a price equal to the average quality, the buyer will deduce that the car is a lemon. The very fact that the car is for sale demonstrates its low quality. Knowing that for \$4,000 he would be sold only lemons, the buyer would refuse to pay more than \$2,000. Let us assume that an indifferent seller sells his car, in which case half of the cars are traded in equilibrium, all of them lemons.

A friendly advisor might suggest to the owner of a good car that he wait until all the lemons have been sold and then sell his own car, since everyone knows that only good cars have remained unsold. But allowing for such behavior changes the model by adding a new action. If it were anticipated, the owners of lemons would also hold back and wait for the price to rise. Such a game could be formally analyzed as a war of attrition (section 3.2).

The outcome that half the cars are held off the market is interesting, though not startling, since half the cars do have genuinely higher quality. It is a formalization of Groucho Marx's wisecrack that he would refuse to join any club that would accept him as a member. Lemons II will have a more dramatic outcome.

Lemons II: Identical Tastes, a Continuum of Types of Sellers

One might wonder whether the outcome of Lemons I was an artifact of the assumption of just two types. Lemons II generalizes the game by allowing the seller to be any of a continuum of types. We will assume that the quality types are uniformly distributed

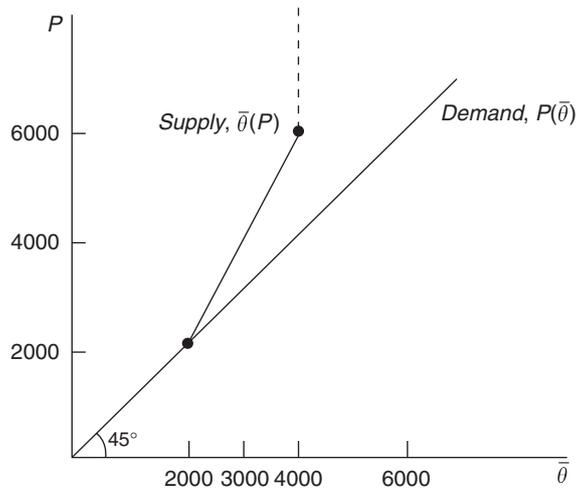


Figure 9.2 Lemons II: identical tastes.

between 2,000 and 6,000. The average quality is $\bar{\theta} = 4,000$, which is therefore the price the buyer would be willing to pay for a car of unknown quality if all cars were on the market. The probability density is zero except on the support $[2,000, 6,000]$, where it is $f(\theta) = 1/(6,000 - 2,000)$, and the cumulative density is

$$\begin{aligned}
 F(\theta) &= \int_{2,000}^{\theta} f(x)dx, \\
 &= \int_{2,000}^{\theta} \frac{1}{4,000} dx = \left|_{x=2,000}^{\theta} \frac{x}{4,000} \right., \\
 &= \frac{\theta}{4,000} - 0.5.
 \end{aligned} \tag{9.4}$$

The payoff functions are the same as in Lemons I.

The equilibrium price must be less than \$4,000 in Lemons II because, as in Lemons I, not all cars are put on the market at that price. Owners are willing to sell only if the quality of their cars is less than 4,000, so while the average quality of all used cars is 4,000, the average quality offered for sale is 3,000. The price cannot be \$4,000 when the average quality is 3,000, so the price must drop at least to \$3,000.

If that happens, the owners of cars with values from 3,000 to 4,000 pull their cars off the market and the average of those remaining is 2,500. The acceptable price falls to \$2,500, and the unravelling continues until the price reaches its equilibrium level of \$2,000. But at $P = 2,000$ the number of cars on the market is infinitesimal. The market has completely collapsed!

Figure 9.2 puts the price of used-cars on one axis and the average quality of cars offered for sale on the other. Each price leads to a different average quality, $\bar{\theta}(P)$, and the slope of $\bar{\theta}(P)$ is greater than one because average quality does not rise proportionately with price. If the price rises, the quality of the *marginal* car offered for sale equals the new price, but

the quality of the *average* car offered for sale is much lower. In equilibrium, the average quality must equal the price, so the equilibrium lies on the 45° line through the origin. That line is a demand schedule of sorts, just as $\bar{\theta}(P)$ is a supply schedule. The only intersection is the point (2,000, 2,000).

9.3 Heterogeneous Tastes: Lemons III and IV

The outcome that no cars are traded is extreme, but there is no efficiency loss in either Lemons I or Lemons II. Since all the players have identical tastes, it does not matter who ends up owning the cars. But the players of the next game, whose tastes differ, have real need of a market.

Lemons III : Buyers Value Cars More than Sellers

Assume that sellers value their cars at exactly their qualities θ but that buyers have valuations 20 percent greater, and, moreover, outnumber the sellers. The payoffs if trade occurs are $\pi_{\text{buyer}} = 1.2\theta - P$ and $\pi_{\text{seller}} = P - \theta$. In equilibrium, the sellers will capture the gains from trade.

In figure 9.3, the curve $\bar{\theta}(P)$ is much the same as in Lemons II, but the equilibrium condition is no longer that price and average quality lie on the 45° line, but that they lie on the demand schedule $P(\bar{\theta})$, which has a slope of 1.2 instead of 1.0. The demand and supply schedules intersect only at $(P = 3,000, \bar{\theta}(P) = 2,500)$. Because buyers are willing to pay a premium, we only see **partial adverse selection**; the equilibrium is partially pooling. The outcome is inefficient, because in a world of perfect information all the cars would be

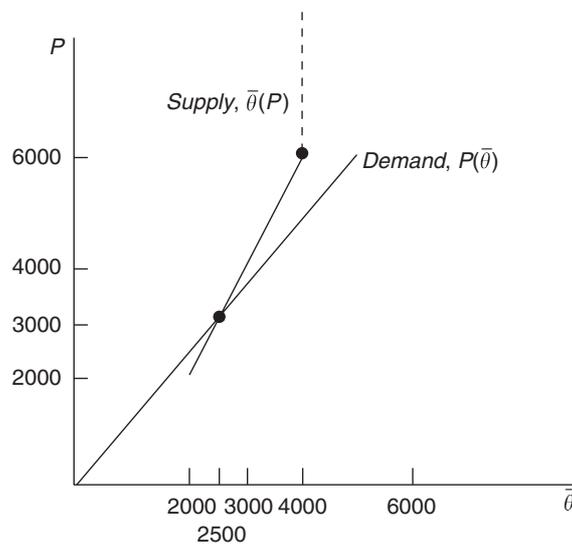


Figure 9.3 Buyers value cars more than sellers: Lemons III.

owned by the “buyers,” who value them more, but under adverse selection they only end up owning the low-quality cars.

Lemons IV: Sellers’ Valuations Differ

In Lemons IV, we dig a little deeper to explain why trade occurs, and we model sellers as consumers whose valuations of quality have changed since they bought their cars. For a particular seller, the valuation of one unit of quality is $(1 + \varepsilon)$, where the random disturbance ε can be either positive or negative and has an expected value of zero. The disturbance could arise because of the seller’s mistake – he did not realize how much he would enjoy driving when he bought the car – or because conditions changed – he switched to a job closer to home. Payoffs if a trade occurs are $\pi_{buyer} = \theta - P$ and $\pi_{seller} = P - (1 + \varepsilon)\theta$.

If $\varepsilon = -0.15$ and $\theta = 2,000$ for a particular seller, then \$1,700 is the lowest price at which he would resell his car. The average quality of cars offered for sale at price P is the expected quality of cars valued by their owners at less than P , that is,

$$\bar{\theta}(P) = E(\theta | (1 + \varepsilon)\theta \leq P). \quad (9.5)$$

Suppose that a large number of new buyers, greater in number than the sellers, appear in the market, and let their valuation of one unit of quality be \$1. The demand schedule, shown in figure 9.4, is the 45° line through the origin. Figure 9.4 shows one possible shape for the supply schedule $\bar{\theta}(P)$, although to specify it precisely we would have to specify the distribution of the disturbances.

In contrast to Lemons I, II, and III, here if $P \geq \$6,000$ some car owners would be reluctant to sell, because they received positive disturbances to their valuations. The average quality of cars on the market is less than 4,000 even at $P = \$6,000$. On the other hand, even if

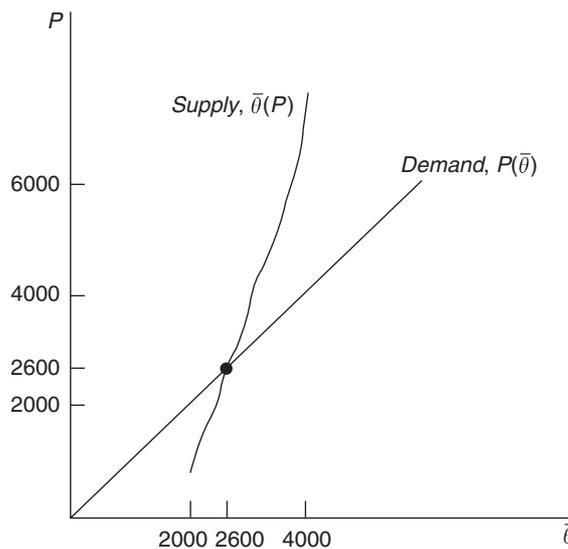


Figure 9.4 Lemons IV: sellers’ valuations differ.

$P = \$2,000$ some sellers with low-quality cars *and* negative realizations of the disturbance do sell, so the average quality remains above 2,000. Under some distributions of ε , a few sellers hate their cars so much they would pay to have them taken away.

The equilibrium drawn in figure 9.4 is ($P = \$2,600$, $\bar{\theta} = 2,600$). Some used cars are sold, but the number is inefficiently low. Some of the sellers have high-quality cars but negative disturbances, and although they would like to sell their cars to someone who values them more, they will not sell at a price of \$2,600.

A theme running through all four Lemons models is that when quality is unknown to the buyer, less trade occurs. Lemons I and II show how trade diminishes, while Lemons III and IV show that the disappearance can be inefficient because some sellers value cars less than some buyers. Next we will use Lemons III, the simplest model with gains from trade, to look at various markets with more sellers than buyers, excess supply, and risk-averse buyers.

More Sellers than Buyers

In analyzing Lemons III we assumed that buyers outnumbered sellers. As a result, the sellers earned producer surplus. In the original equilibrium, all the sellers with quality less than 3,000 offered a price of \$3,000 and earned a surplus of up to \$1,000. There were more buyers than sellers, so every seller who wished to sell was able to do so, but the price equalled the buyers' expected utility, so no buyer who failed to purchase was dissatisfied. The market cleared.

If, instead, sellers outnumber buyers, what price should a seller offer? At \$3,000, not all would-be sellers can find buyers. A seller who proposed a lower price would find willing buyers despite the somewhat lower expected quality. The buyer's trade-off between lower price and lower quality is shown in figure 9.3, in which the expected consumer surplus is the vertical distance between the price (the height of the supply schedule) and the demand schedule. When the price is \$3,000 and the average quality is 2,500, the buyer expects a consumer surplus of zero, which is $\$3,000 - \$1.2 \cdot 2,500$. The combination of price and quality that buyers like best is (\$2,000, 2,000), because if there were enough sellers with quality $\theta = 2,000$ to satisfy the demand, each buyer would pay $P = \$2,000$ for a car worth \$2,400 to him, acquiring a surplus of \$400. If there were fewer sellers, the equilibrium price would be higher and some sellers would receive producer surplus.

Heterogeneous Buyers: Excess Supply

If buyers have different valuations for quality, the market might not clear, as Charles Wilson (1980) points out. Assume that the number of buyers willing to pay \$1.2 per unit of quality exceeds the number of sellers, but that buyer Smith is an eccentric whose demand for high quality is unusually strong. He would pay \$100,000 for a car of quality 5,000 or greater, and \$0 for a car of any lower quality.

In Lemons III without Smith, the outcome is a price of \$3,000, an average market quality of 2,500, and a market quality range between 2,000 and 3,000. Smith would be unhappy

with this, since he has zero probability of finding a car he likes. In fact, he would be willing to accept a price of \$6,000, so that all the cars, from quality 2,000 to 6,000, would be offered for sale and the probability that he buys a satisfactory car would rise from 0 to 0.25. But Smith would not want to buy all the cars offered to him, so the equilibrium has two prices, \$3,000 and \$6,000, with excess supply at the higher price.

Strangely enough, Smith's demand function is upward sloping. At a price of \$3,000, he is unwilling to buy; at a price of \$6,000, he is willing, because expected quality rises with price. This does not contradict basic price theory, for the standard assumption of *ceteris paribus* is violated. As the price increases, the quantity demanded would fall if all else stayed the same, but all else does not – quality rises.

Risk Aversion

We have implicitly assumed, by the choice of payoff functions, that the buyers and sellers are both risk-neutral. What happens if they are risk-averse – that is, if the marginal utilities of wealth and car quality are diminishing? Again we will use Lemons III and the assumption of many buyers.

On the seller's side, risk aversion changes nothing. The seller runs no risk because he knows exactly the price he receives and the quality he surrenders. But the buyer does bear risk, because he buys a car of uncertain quality. Although he would pay \$3,600 for a car he knows has quality 3,000, if he is risk-averse he will not pay that much for a car with expected quality 3,000 but actual quality of possibly 2,500 or 3,500: he would obtain less utility from adding 500 quality units than from subtracting 500. The buyer would pay perhaps \$2,900 for a car whose expected quality is 3,000 where the demand schedule is nonlinear, lying everywhere below the demand schedule of the risk-neutral buyer. As a result, the equilibrium has a lower price and average quality.

9.4 Adverse Selection under Uncertainty: Insurance Game III

The term “adverse selection,” like “moral hazard,” comes from insurance. Insurance pays more if there is an accident than otherwise, so it benefits accident-prone customers more than safe ones and a firm's customers are “adversely selected” to be accident-prone. The classic article on adverse selection in insurance markets is Rothschild & Stiglitz (1976), which begins, “Economic theorists traditionally banish discussions of information to footnotes.” How things have changed! Within ten years, information problems came to dominate research in both microeconomics and macroeconomics.

We will follow Rothschild & Stiglitz in using state-space diagrams, and we will use a version of section 8.5's Insurance Game. Under moral hazard, Smith chose whether to be *Careful* or *Careless*. Under adverse selection, Smith cannot affect the probability of a theft, which is chosen by Nature. Rather, Smith is either *Safe* or *Unsafe*, and while he cannot affect the probability that his car will be stolen, he does know what the probability is.

Insurance Game III

PLAYERS

Smith and two insurance companies.

THE ORDER OF PLAY

- 0 Nature chooses Smith to be either *Safe*, with probability 0.6, or *Unsafe*, with probability 0.4. Smith knows his type, but the insurance companies do not.
- 1 Each insurance company offers its own contract (x, y) under which Smith pays premium x unconditionally and receives compensation y if there is a theft.
- 2 Smith picks a contract.
- 3 Nature chooses whether there is a theft, using probability 0.5 if Smith is *Safe* and 0.75 if he is *Unsafe*.

PAYOFFS

Smith's payoff depends on his type and the contract (x, y) that he accepts. Let $U' > 0$ and $U'' < 0$.

$$\pi_{Smith}(Safe) = 0.5U(12 - x) + 0.5U(0 + y - x).$$

$$\pi_{Smith}(Unsafe) = 0.25U(12 - x) + 0.75U(0 + y - x).$$

The companies' payoffs depend on what types of customers accept their contracts, as shown in table 9.1.

Table 9.1 Insurance Game III payoffs

<i>Company payoff</i>	<i>Types of customers</i>
0	No customers
$0.5x + 0.5(x - y)$	Just <i>Safe</i>
$0.25x + 0.75(x - y)$	Just <i>Unsafe</i>
$0.6[0.5x + 0.5(x - y)] + 0.4[0.25x + 0.75(x - y)]$	<i>Unsafe</i> and <i>Safe</i>

Smith is *Safe* with probability 0.6 and *Unsafe* with probability 0.4. Without insurance, Smith's dollar wealth is 12 if there is no theft and 0 if there is, depicted in figure 9.5 as his endowment in state space, $\omega = (12, 0)$. If Smith is *Safe*, a theft occurs with probability 0.5, but if he is *Unsafe* the probability is 0.75. Smith is risk-averse (because $U'' < 0$) and the insurance companies are risk-neutral.

If an insurance company knew that Smith was *Safe*, it could offer him insurance at a premium of 6 with a payout of 12 after a theft, leaving Smith with an allocation of $(6, 6)$. This is the most attractive contract that is not unprofitable, because it fully insures Smith. Whatever the state, his allocation is 6.

Figure 9.5 shows the indifference curves of Smith and an insurance company, with the no-insurance starting point at $\omega = (12, 0)$. A higher insurance premium reduces Smith's

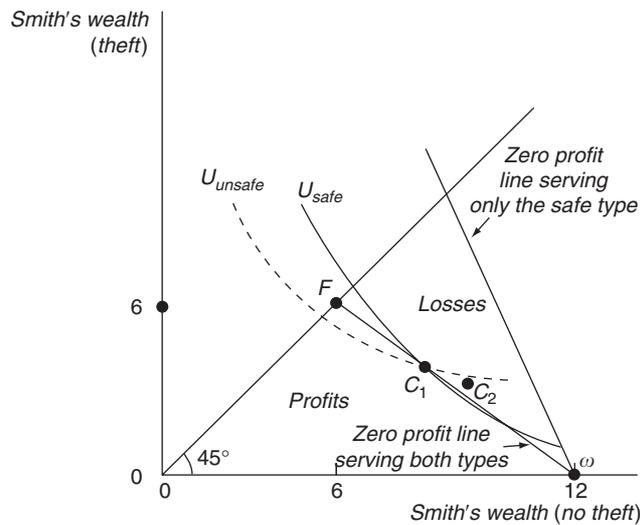


Figure 9.5 Insurance Game III: nonexistence of a pooling equilibrium.

wealth in both states of the world; a higher theft insurance payout increases Smith's wealth in the state of the world in which there is a theft. The insurance company is risk-neutral, so its indifference curve is a straight line with negative slope, since to keep the company's profit constant, the decrease in profit from a rise in Smith's wealth if there is no theft must be balanced by an increase in profit from a fall in Smith's wealth if there is a theft.

If Smith will be a customer regardless of his type, the company's indifference curve based on its expected profits is ωF (although if the company knew that Smith was *Safe*, the indifference curve would be steeper, and if it knew he was *Unsafe*, the curve would be less steep). The insurance company is indifferent between ω and C_1 , at both of which its expected profits are zero. Smith is risk-averse, so his indifference curves are convex, and closest to the origin along the 45° line if the probability of *Theft* is 0.5. He has two sets of indifference curves, solid if he is *Safe* and dotted if he is *Unsafe*.

Figure 9.5 shows why no Nash pooling equilibrium exists. To make zero profits, the equilibrium must lie on the line ωF . It is easiest to think about these problems by imagining an entire population of Smiths, whom we will call "customers." Pick a contract C_1 anywhere on ωF and think about drawing the indifference curves for the *Unsafe* and *Safe* customers that pass through C_1 . *Safe* customers are always willing to trade *Theft* wealth for *No Theft* wealth at a higher rate than *Unsafe* customers. At any point, therefore, the slope of the solid (*Safe*) indifference curve is steeper than that of the dashed (*Unsafe*) curve. Since the slopes of the dashed and solid indifference curves differ, we can insert another contract, C_2 , between them and just barely to the right of ωF . The *Safe* customers prefer contract C_2 to C_1 , but the *Unsafe* customers stay with C_1 , so C_2 is profitable – since C_2 only attracts *Safes*, it need not be to the left of ωF to avoid losses. But then the original contract C_1 was not a Nash equilibrium, and since our argument holds for any pooling contract, no pooling equilibrium exists.

The attraction of the *Safe* customers away from pooling is referred to as **cream skimming**, although profits are still zero when there is competition for the cream. We next consider

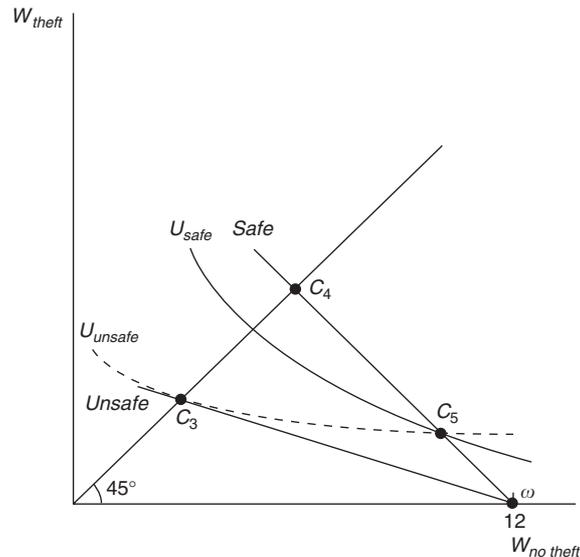


Figure 9.6 A separating equilibrium for Insurance Game III.

whether a separating equilibrium exists, using figure 9.6. The zero-profit condition requires that the *Safe* customers take contracts on ωC_4 and the *Unsafes* on ωC_3 .

The *Unsafes* will be completely insured in any equilibrium, albeit at a high price. On the zero-profit line ωC_3 , the contract they like best is C_3 , which the *Safes* are not tempted to take. The *Safes* would prefer contract C_4 , but C_4 uniformly dominates C_3 , so it would attract *Unsafes* too, and generate losses. To avoid attracting *Unsafes*, the *Safe* contract must be below the *Unsafe* indifference curve. Contract C_5 is the fullest insurance the *Safes* can get without attracting *Unsafes*: it satisfies the self-selection and competition constraints.

Contract C_5 , however, might not be an equilibrium either. Figure 9.7 is the same as figure 9.6 with a few additional points marked. If one firm offered C_6 , it would attract both types, *Unsafe* and *Safe*, away from C_3 and C_5 , because it is to the right of the indifference curves passing through those points. Would C_6 be profitable? That depends on the proportions of the different types. The assumption on which the equilibrium of figure 9.6 is based is that the proportion of *Safes* is 0.6, so the zero-profit line for pooling contracts is ωF and C_6 would be unprofitable. In figure 9.7 it is assumed that the proportion of *Safes* is higher, so the zero-profit line for pooling contracts would be $\omega F'$ and C_6 , lying to its left, is profitable. But we already showed that no pooling contract is Nash, so C_6 cannot be an equilibrium. Since neither a separating pair like (C_3, C_5) nor a pooling contract like C_6 is an equilibrium, no equilibrium whatsoever exists.

The essence of nonexistence here is that if separating contracts are offered, some company is willing to offer a superior pooling contract; but if a pooling contract is offered, some company is willing to offer a separating contract that makes it unprofitable. A monopoly would have a pure-strategy equilibrium, but in a competitive market only a mixed-strategy Nash equilibrium exists (see Dasgupta & Maskin [1986]).

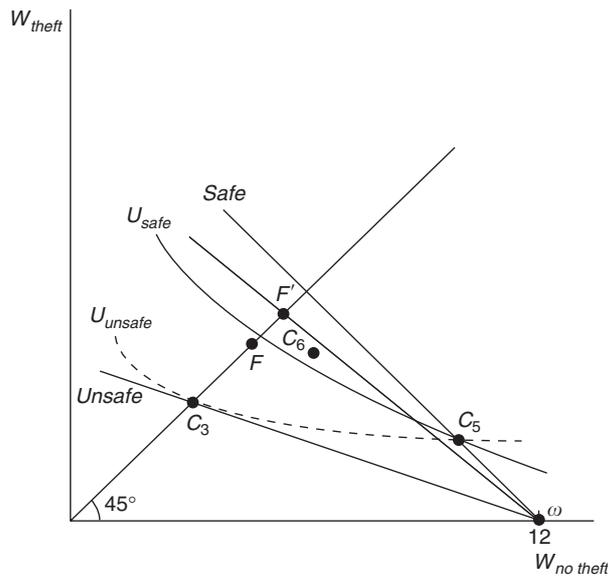


Figure 9.7 Curves for which there is no equilibrium in Insurance Game III.

*9.5 Market Microstructure

The prices of securities such as stocks depend on what investors believe is the value of the assets that underlie them. The values are highly uncertain, and new information about them is constantly being generated. The market microstructure literature is concerned with how new information enters the market. In the paradigmatic situation, an informed trader has private information about the asset value that he hopes to use to make profitable trades, but other traders know that someone might have private information. This is adverse selection, because the informed trader has better information on the value of the stock, and no uninformed trader wants to trade with an informed trader – the informed trade is a “bad type” from the point of view of the other side of the market. An institution that many markets have developed is the “marketmaker” or “specialist,” a trader in a particular stock who is always willing to buy or sell to keep the market going. Other traders feel safer in trading with the marketmaker than with a potentially informed trader, but this just transfers the adverse selection problem to the marketmaker, who always loses when he trades with someone who is informed.

The two models in this section will look at how a marketmaker deals with the problem of informed trading. Both are descendants of the verbal model in Bagehot (1971). (“Bagehot,” pronounced “badget,” is a pseudonym for Jack Treynor. See Glosten & Milgrom [1985] for a formalization.) In the Bagehot model, there may or may not be one or more informed traders, but the informed traders as a group have a trade of fixed size if they are present. The marketmaker must decide how big a bid-ask spread to charge. In the Kyle (1985) model, there is one informed trader, who decides how much to trade. On observing the imbalance of orders, the marketmaker decides what price to offer.

The Bagehot Model

PLAYERS

The informed trader and two competing marketmakers.

THE ORDER OF PLAY

- 0 Nature chooses the asset value v to be either $p_0 - \delta$ or $p_0 + \delta$ with equal probability. The marketmakers never observe the asset value, nor do they observe whether anyone else observes it, but the “informed” trader observes v with probability θ .
- 1 The marketmakers choose their spreads s , offering prices $p_{bid} = p_0 - s/2$ at which they will buy the security and $p_{ask} = p_0 + s/2$ for which they will sell it.
- 2 The informed trader decides whether to buy one unit, sell one unit, or do nothing.
- 3 Noise traders buy n units and sell n units.

PAYOFFS

Everyone is risk-neutral. The informed trader’s payoff is $(v - p_{ask})$ if he buys, $(p_{bid} - v)$ if he sells, and zero if he does nothing. The marketmaker who offers the highest p_{bid} trades with all the customers who wish to sell, and the marketmaker who offers the lowest p_{ask} trades with all the customers who wish to buy. If the marketmakers set equal prices, they split the market evenly. A marketmaker who sells x units gets a payoff of $x(p_{ask} - v)$, and a marketmaker who buys x units gets a payoff of $x(v - p_{bid})$.

Optimal strategies are simple. Competition between the marketmakers will make their prices identical and their profits zero. The informed trader should buy if $v > p_{ask}$ and sell if $v < p_{bid}$. He has no incentive to trade if $v \in [p_{bid}, p_{ask}]$.

A marketmaker will always lose money trading with the informed trader, but if $s > 0$, so $p_{ask} > p_0$ and $p_{bid} < p_0$, he will earn positive expected profits in trading with the noise traders. Since a marketmaker could specialize in either sales or purchases, he must earn zero expected profits overall from either type of trade. Centering the bid-ask spread on the expected value of the stock, p_0 , ensures this. Marketmaker sales will be at the ask price of $(p_0 + s/2)$. With probability 0.5, this is above the true value of the stock, $(p_0 - \delta)$, in which case the informed trader will not buy but the marketmakers will earn a total profit of $n[(p_0 + s/2) - (p_0 - \delta)]$ from the noise traders. With probability 0.5, the ask price of $(p_0 + s/2)$ is below the true value of the stock, $(p_0 + \delta)$, in which case the informed trader will be informed with probability θ and buy one unit and the noise traders will buy n more in any case, so the marketmakers will earn a total expected profit of $(n + \theta)[(p_0 + s/2) - (p_0 + \delta)]$, a negative number. For marketmaker profits from sales at the ask price to be zero overall, this expected profit must be zero:

$$0.5n[(p_0 + s/2) - (p_0 - \delta)] + 0.5(n + \theta)[(p_0 + s/2) - (p_0 + \delta)] = 0. \quad (9.6)$$

Equation (9.6) implies that $n[s/2 + \delta] + (n + \theta)[s/2 - \delta] = 0$, so

$$s^* = \frac{2\delta\theta}{2n + \theta}. \quad (9.7)$$

The profit from marketmaker purchases must similarly equal zero, and will for the same spread s^* , though we will not go through the algebra here.

Equation (9.7) has a number of implications. First, the spread s^* is positive. Even though marketmakers compete and have zero transactions costs, they charge a different price to buy and to sell. They make money dealing with the noise traders but lose money with the informed trader, if he is present. The comparative statics reflect this. The spread s^* rises in δ , the dispersion of the true value, because divergent true values increase losses from trading with the informed trader, and s^* falls in n , which reflects the number of noise traders relative to informed traders, because when there are more noise traders, the profits from trading with them are greater. The spread s^* rises in θ , the probability that the informed trader really has inside information, which is also intuitive but requires a little calculus to demonstrate, using equation (9.7):

$$\frac{\partial s^*}{\partial \theta} = \frac{2\delta}{2n + \theta} - \frac{2\delta\theta}{(2n + \theta)^2} = \left(\frac{1}{(2n + \theta)^2} \right) (4\delta n + 2\delta\theta - 2\delta\theta) > 0. \quad (9.8)$$

The second model of market microstructure, important because it is commonly used as a foundation for more complicated models, is the Kyle model. It focuses on the decision of the informed trader, not the marketmaker. The Kyle model is set up so that marketmaker observes the trade volume before he chooses the price.

The Kyle Model (Kyle [1985])

PLAYERS

The informed trader and two competing marketmakers.

THE ORDER OF PLAY

- 0 Nature chooses the asset value v from a normal distribution with mean p_0 and variance σ_v^2 , observed by the informed trader but not by the marketmakers.
- 1 The informed trader offers a trade of size $x(v)$, which is a purchase if positive and a sale if negative, unobserved by the marketmaker.
- 2 Nature chooses a trade of size u by noise traders, unobserved by the marketmaker, where u is distributed normally with mean zero and variance σ_u^2 .
- 3 The marketmakers observe the total market trade offer $y = x + u$, and choose prices $p(y)$.
- 4 Trades are executed. If y is positive (the market wants to purchase, in net), whichever marketmaker offers the lowest price executes the trades; if y is negative (the market wants to sell, in net), whichever marketmaker offers the highest price executes the trades. The value v is then revealed to everyone.

PAYOFFS

All players are risk-neutral. The informed trader's payoff is $(v - p)x$. The marketmaker's payoff is zero if he does not trade and $(p - v)y$ if he does.

An equilibrium for this game is the strategy profile

$$x(v) = (v - p_0) \left(\frac{\sigma_u}{\sigma_v} \right) \quad (9.9)$$

and

$$p(y) = p_0 + \left(\frac{\sigma_v}{2\sigma_u} \right) y. \quad (9.10)$$

This is reasonable. It says that the informed trader will increase the size of his trade as v gets bigger relative to p_0 (and he will sell, not buy, if $v - p_0 < 0$), and the marketmaker will increase the price he charges if y is bigger and more people want to sell, which is an indicator that the informed trader might be trading heavily. The variances of the asset value (σ_v^2) and the noise trading (σ_u^2) enter as one would expect, and they matter only in their relation to each other. If σ_v^2/σ_u^2 is large, then the asset value fluctuates more than the amount of noise trading, and it is difficult for the informed trader to conceal his trades under the noise. The informed trader will trade less, and a given amount of trading will cause a greater response from the marketmaker. One might say that the market is less “liquid”: a trade of given size will have a greater impact on the price.

I will not (and cannot) prove uniqueness of the equilibrium, since it is very hard to check all possible profiles of nonlinear strategies, but I will show that $\{(9.9), (9.10)\}$ is Nash and is the unique linear equilibrium. To start, hypothesize that the informed trader uses a linear strategy, so

$$x(v) = \alpha + \beta v, \quad (9.11)$$

for some constants α and β . Competition between the marketmakers means that their expected profits will be zero, which requires that the price they offer be the expected value of v . Thus, their equilibrium strategy $p(y)$ will be an unbiased estimate of v given their data y , where they know that y is normally distributed and that

$$\begin{aligned} y &= x + u, \\ &= \alpha + \beta v + u. \end{aligned} \quad (9.12)$$

This means that their best estimate of v given the data y is, following the usual regression rule (which readers unfamiliar with statistics must accept on faith),

$$\begin{aligned} E(v|y) &= E(v) + \left(\frac{\text{cov}(v, y)}{\text{var}(y)} \right) y, \\ &= p_0 + \left(\frac{\beta \sigma_v^2}{\beta^2 \sigma_v^2 + \sigma_u^2} \right) y, \\ &= p_0 + \lambda y, \end{aligned} \quad (9.13)$$

where λ is a new shorthand variable to save writing out the term in parentheses.

The function $p(y)$ will be a linear function of y under our assumption that x is a linear function of v . Given that $p(y) = p_0 + \lambda y$, what must next be shown is that x will indeed be a linear function of v . Start by writing the informed trader's expected payoff, which is

$$\begin{aligned} E\pi_i &= E([v - p(y)]x), \\ &= E([v - p_0 - \lambda(x + u)]x), \\ &= [v - p_0 - \lambda(x + 0)]x, \end{aligned} \quad (9.14)$$

since $E(u) = 0$. Maximizing the expected payoff with respect to x yields the first-order condition

$$v - p_0 - 2\lambda x = 0, \quad (9.15)$$

which on rearranging becomes

$$x = -\frac{p_0}{2\lambda} + \left(\frac{1}{2\lambda}\right)v. \quad (9.16)$$

Equation (9.16) establishes that $x(v)$ is linear, given that $p(y)$ is linear. All that is left is to find the value of λ . By comparing (9.16) and (9.11) we can see that $\beta = 1/2\lambda$. Substituting this β into the value of λ from (9.13) leads to

$$\lambda = \frac{\beta\sigma_v^2}{\beta^2\sigma_v^2 + \sigma_u^2} = \frac{\sigma_v^2/2\lambda}{\sigma_v^2/(4\lambda^2) + \sigma_u^2}, \quad (9.17)$$

which upon solving for λ yields $\lambda = \sigma_v/2\sigma_u$. Since $\beta = 1/2\lambda$, it follows that $\beta = \sigma_u/\sigma_v$. These values of λ and β together with equation (9.16) give the strategies asserted at the start in equations (9.9) and (9.10).

The Bagehot model is perhaps a better explanation of why marketmakers might charge a bid/ask spread even under competitive conditions and with zero transactions costs. Its assumption is that the marketmaker cannot change the price depending on volume, but must instead offer a price, and then accept whatever order comes along – a buy order, or a sell order.

The two main divisions of the field of finance are corporate finance and asset pricing. Corporate finance, the study of such things as project choice, capital structure, and mergers has the most obvious applications of game theory, but the Bagehot and Kyle models show that the same techniques are also important in asset pricing. For more, I recommend Harris & Raviv (1995).

*9.6 A Variety of Applications

Price Dispersion

Usually the best model for explaining price dispersion is a search model – Salop & Stiglitz (1977), for example, which is based on buyers whose search costs differ. But although we passed over it quickly in section 9.3, the Lemons model with Smith, the quality-conscious

consumer, generated not only excess supply, but price dispersion as well. Cars of the same average quality were sold for \$3,000 and \$6,000.

Similarly, while the most obvious explanation for why brands of stereo amplifiers sell at different prices is that customers are willing to pay more for higher quality, adverse selection contributes another explanation. Consumers might be willing to pay high prices because they know that high-priced brands could include both high-quality and low-quality amplifiers, whereas low-priced brands are invariably low quality. The low-quality amplifier ends up selling at two prices: a high price in competition with high-quality amplifiers, and, in different stores or under a different name, a low price aimed at customers less willing to trade dollars for quality.

This explanation does depend on sellers of amplifiers incurring a large enough fixed set-up or operating cost. Otherwise, too many low-quality brands would crowd into the market, and the proportion of high-quality brands would be too small for consumers to be willing to pay the high price. The low-quality brands would benefit as a group from entry restrictions: too many of them spoil the market, not through price competition but through degrading the average quality.

Health Insurance

Medical insurance is subject to adverse selection because some people are healthier than others. The variance in health is particularly high among old people, who sometimes have difficulty in obtaining insurance at all. Under basic economic theory this is a puzzle: the price should rise until supply equals demand. The problem is pooling: when the price of insurance is appropriate for the average old person, healthier ones stop buying. The price must rise to keep profits nonnegative, and the market disappears, just as in Lemons II.

If the facts indeed fit this story, adverse selection is an argument for government-enforced pooling. If all old people are required to purchase government insurance, then while the healthier of them may be worse off, the vast majority could be helped.

Using adverse selection to justify medicare, however, points out how dangerous many of the models in this book can be. For policy questions, the best default opinion is that markets are efficient. On closer examination, we have found that many markets are inefficient because of strategic behavior or information asymmetry. It is dangerous, however, to immediately conclude that the government should intervene, because the same arguments applied to government show that the cure might be worse than the disease. The analyst of health-care needs to take seriously the moral hazard and rent seeking that arise from government insurance. Doctors and hospitals will increase the cost and amount of treatment if the government pays for it, and the transfer of wealth from young people to the elderly, which is likely to swamp the gains in efficiency, might distort the shape of the government program from the economist's ideal.

Henry Ford's Five-dollar Day

In 1914 Henry Ford made a much-publicized decision to raise the wage of his auto workers to \$5 a day, considerably above the market wage. This pay hike occurred without pressure from the workers, who were nonunionized. Why did Ford do it?

The pay hike could be explained by either moral hazard or adverse selection. In accordance with the idea of efficiency wages (section 8.1), Ford might have wanted workers who worried about losing their premium job at his factory, because they would work harder and refrain from shirking. Adverse selection could also explain the pay hike: by raising his wage Ford attracted a mixture of low- and high-quality workers, rather than low-quality alone (see Raff & Summers [1987]).

Bank Loans

Suppose that two people come to you for an unsecured loan of \$10,000. One offers to pay an interest rate of 10 percent and the other offers 200 percent. Who do you accept? Like Charles Wilson's car buyer in section 9.3 who chose to buy at a high price, you may choose to lend at a low interest rate.

If a lender raises his interest rate, both his pool of loan applicants and their behavior change because adverse selection and moral hazard contribute to a rise in default rates. Borrowers who expect to default are less concerned about the high interest rate than dependable borrowers, so the number of loans shrinks and the default rate rises (see Stiglitz & Weiss [1981]). In addition, some borrowers shift to higher-risk projects with greater chance of default but higher yields when they are successful. In section 6.6 we went through the model of Diamond (1989) which looks at the evolution of this problem as firms age.

Whether because of moral hazard or adverse selection, asymmetric information can also result in excess demand for bank loans. The savers who own the bank do not save enough at the equilibrium interest rate to provide loans to all the borrowers who want loans. Thus, the bank makes a loan to John Smith, while denying one to Joe, his observationally equivalent twin. Policymakers should carefully consider any laws that rule out arbitrary loan criteria or require banks to treat all customers equally. A bank might wish to restrict its loans to left-handed people, neither from prejudice nor because it is useful to ration loans according to some criterion arbitrary enough to avoid the moral hazard of favoritism by loan officers.

Bernanke (1983) suggests adverse selection in bank loans as an explanation for the Great Depression in the United States. The difficulty in explaining the Depression is not so much the initial stock market crash as the persistence of the unemployment that followed. Bernanke notes that the crash wiped out local banks and dispersed the expertise of the loan officers. After the loss of this expertise, the remaining banks were less willing to lend because of adverse selection, and it was difficult for the economy to recover.

Solutions to Adverse Selection

Even in markets where it apparently does not occur, the threat of adverse selection, like the threat of moral hazard, can be an important influence on market institutions. Adverse selection can be circumvented in a number of ways besides the contractual solutions we have been analyzing. I will mention some of them in the context of the used-car market.

One set of solutions consists of ways to make car quality contractible. Buyers who find that their car is defective may have recourse to the legal system if the sellers were fraudulent, although in the United States the courts are too slow and costly to be fully effective. Other government bodies such as the Federal Trade Commission may do better by

issuing regulations particular to the industry. Even without regulation, private warranties – promises to repair the car if it breaks down – may be easier to enforce than oral claims, by dispelling ambiguity about what level of quality is guaranteed.

Testing (the equivalent of moral hazard's monitoring) is always used to some extent. The prospective driver tries the car on the road, inspects the body, and otherwise tries to reduce information asymmetry. At a cost, he could even reverse the asymmetry by hiring mechanics who can tell him more about the car than the owner himself knows. The rule is not always *caveat emptor*; what should one's response be to an antique dealer who offers to pay \$500 for an apparently worthless old chair?

Reputation can solve adverse selection, just as it can solve moral hazard, but only if the transaction is repeated and the other conditions of the models in chapters 5 and 6 are met. An almost opposite solution is to show that there are innocent motives for a sale; that the owner of the car has gone bankrupt, for example, and his creditor is selling the car cheaply to avoid the holding cost.

Penalties not strictly economic are also important. One example is the social ostracism inflicted by the friend to whom a lemon has been sold; the seller is no longer invited to dinner. Or, the seller might have moral principles that prevent him from defrauding buyers. Such principles, provided they are common knowledge, would help him obtain a higher price in the used-car market. Akerlof himself has worked on the interaction between social custom and markets in his 1980 and 1983 articles. The second of these articles looks directly at the value of inculcating moral principles, using theoretical examples to show that parents might wish to teach their children principles, and that society might wish to give hiring preference to students from elite schools.

It is by violating the assumptions needed for perfect competition that asymmetric information enables government and social institutions to raise efficiency. This points to a major reason for studying asymmetric information: where it is important, noneconomic interference can be helpful instead of harmful. I find the social solutions particularly interesting since, as mentioned earlier in connection with health care, government solutions introduce agency problems as severe as the information problems they solve. Noneconomic behavior is important under adverse selection, in contrast to perfect competition, which allows an "Invisible Hand" to guide the market to efficiency, regardless of the moral beliefs of the traders. If everyone were honest, the lemons problem would disappear because the sellers would truthfully disclose quality. If some fraction of the sellers were honest, but buyers could not distinguish them from the dishonest sellers, the outcome would presumably be somewhere between the outcomes of complete honesty and complete dishonesty. The subject of market ethics is important, and would profit from investigation by scholars trained in economic analysis.

9.7 Adverse Selection and Moral Hazard Combined: Production Game VII

What happens when adverse selection and moral hazard combine, so that agents begin the game with private information but then also take unobserved actions in the course of the game? Production Game VII shows one possibility, in a model that will also be used at the start of chapter 10.

Production Game VII: Adverse Selection and Moral Hazard

PLAYERS

The principal and the agent.

THE ORDER OF PLAY

- 0 Nature chooses the state of the world s , observed by the agent but not by the principal, according to distribution $F(s)$, where the state s is Good with probability 0.5 and Bad with probability 0.5.
- 1 The principal offers the agent a wage contract $w(q)$.
- 2 The agent accepts or rejects the contract.
- 3 The agent chooses effort level e .
- 4 Output is $q = q(e, s)$, where $q(e, \text{good}) = 3e$ and $q(e, \text{bad}) = e$.

PAYOFFS

If the agent rejects all contracts, then $\pi_{\text{agent}} = \bar{U} = 0$ and $\pi_{\text{principal}} = 0$.

Otherwise, $\pi_{\text{agent}} = U(e, w, s) = w - e^2$ and $\pi_{\text{principal}} = V(q - w) = q - w$.

Thus, there is no uncertainty, both principal and agent are risk-neutral in money, and effort is increasingly costly.

In this model, the first-best effort depends on the state of the world. The two social surplus maximization problems are

$$\text{Maximize}_{e_g} 3e_g - e_g^2, \quad (9.18)$$

which is solved by the optimal effort $e_g = 1.5$ (and $q_g = 4.5$) in the good state, and

$$\text{Maximize}_{e_b} e_b - e_b^2, \quad (9.19)$$

which is solved by the optimal effort $e_b = 0.5$ (and $q_b = 0.5$) in the bad state.

The problem is that the principal does not know what level of effort and output are appropriate. He does not want to require high output in both states, because if he does, he will have to pay too high a salary to the agent to compensate for the difficulty of attaining that output in the bad state. Rather, he must solve the following problem:

$$\text{Maximize}_{q_g, q_b, w_g, w_b} [0.5(q_g - w_g) + 0.5(q_b - w_b)], \quad (9.20)$$

where the agent has a choice between two forcing contracts, (q_g, w_g) and (q_b, w_b) , and the contracts must induce participation and self selection.

The self-selection constraints are based on efforts of $e = q/3$ for the good state and $e = q$ for the bad state. In the good state, the agent must choose the good-state contract, so

$$\pi_{\text{agent}}(q_g, w_g | \text{good}) = w_g - \left(\frac{q_g}{3}\right)^2 \geq \pi_{\text{agent}}(q_b, w_b | \text{good}) = w_b - \left(\frac{q_b}{3}\right)^2 \quad (9.21)$$

and in the bad state he must choose the bad-state contract, so

$$\pi_{agent}(q_b, w_b|bad) = w_b - q_b^2 \geq \pi_{agent}(q_g, w_g|bad) = w_g - q_g^2. \quad (9.22)$$

The participation constraints are

$$\pi_{agent}(q_g, w_g|good) = w_g - \left(\frac{q_g}{3}\right)^2 \geq 0 \quad (9.23)$$

and

$$\pi_{agent}(q_b, w_b|bad) = w_b - q_b^2 \geq 0. \quad (9.24)$$

The bad state's participation constraint will be binding, since in the bad state the agent will not be tempted by the good-state contract's higher output and wage. Thus, we can conclude from constraint (9.24) that

$$w_b = q_b^2. \quad (9.25)$$

The good state's participation constraint will not be binding, since there the agent will be left with an informational rent – the principal must leave the agent some surplus to induce him to reveal the good state.

The good state's self-selection constraint will be binding, since in the good state the agent *will* be tempted to take the easier contract appropriate for the bad state. Thus, we can conclude from constraint (9.21) that

$$\begin{aligned} w_g &= \left(\frac{q_g}{3}\right)^2 + w_b - \left(\frac{q_b}{3}\right)^2, \\ &= \left(\frac{q_g}{3}\right)^2 + q_b^2 - \left(\frac{q_b}{3}\right)^2, \end{aligned} \quad (9.26)$$

where the second step substitutes for w_b from equation (9.25). The bad state's self-selection constraint will not be binding, since the agent would then not be tempted to produce a large amount for a large wage.

Now let's return to the principal's maximization problem. Having found expressions for w_b and w_g we can rewrite (9.20) as

$$\underset{q_g, q_b}{\text{Maximize}} \left[0.5 \left(q_g - \left(\frac{q_g}{3}\right)^2 - q_b^2 + \left(\frac{q_b}{3}\right)^2 \right) + 0.5(q_b - q_b^2) \right] \quad (9.27)$$

with no constraints. The first-order conditions are

$$0.5 \left(1 - \frac{2q_g}{9} \right) = 0, \quad (9.28)$$

so $q_g = 4.5$, and

$$0.5 \left(-2q_b + \frac{2q_b}{9} \right) + 0.5(1 - 2q_b) = 0, \quad (9.29)$$

so $q_b \approx 0.26$. We can then find the wages that satisfy the constraints, which are $w_g \approx 2.32$ and $w_b \approx 0.07$.

Thus, in the second-best world of information asymmetry, the effort in the good state remains the first-best effort, but second-best effort in the bad state is lower than first-best. This results from the principal's need to keep the bad-state contract from being too attractive in the good state. Bad-state output and compensation must be suppressed. Good-state output, on the other hand, should be left at the first-best level, since the agent will not be tempted by that contract in the bad state.

Also, observe that in the good state the agent earns an informational rent. As explained earlier, this is because the good-state agent could always earn a positive payoff by pretending the state was bad and taking that contract, so any contract that separates out the good-state agent (while leaving some contract acceptable to the bad-state agent) must also have a positive payoff.

Notes

N9.1 Introduction: Production Game VI

- In moral hazard with hidden knowledge, the contract must ordinarily satisfy only one participation constraint, whereas in adverse selection problems there is a different participation constraint for each type of agent. An exception is if there are constraints limiting how much an agent can be punished in different states of the world. If, for example, there are bankruptcy constraints, then, if the agent has different wealths across the N possible states of the world, there will be N constraints for how negative his wage can be, in addition to the single participation constraint. These can be looked at as **interim** participation constraints, since they represent the idea that the agent wants to get out of the contract once he observes the state of the world midway through the game.
- Gresham's Law ("Bad money drives out good") is a statement of adverse selection. Only debased money will be circulated if the payer knows the quality of his money better than the receiver. The same result occurs if quality is common knowledge, but for legal reasons the receiver is obligated to take the money, whatever its quality. An example of the first is Roman coins with low silver content; and of the second, Zambian currency with an overvalued exchange rate.
- Most adverse selection models have types that could be called "good" and "bad," because one type of agent would like to pool with the other, who would rather be separate. It is also possible to have a model in which both types would rather separate – types of workers who prefer night shifts versus those who prefer day shifts, for example – or two types who both prefer pooling – male and female college students.
- Two curious features of labor markets is that workers of widely differing outputs seem to be paid identical wages and that tests are not used more in hiring decisions. Schmidt and Judiesch (as cited in Seligman [1992], p. 145) have found that in jobs requiring only unskilled and semi-skilled blue-collar workers, the top 1 percent of workers, as defined by performance on ability tests not directly related to output, were 50 percent more productive than the average. In jobs defined as "high complexity" the difference was 127 percent.

At about the same time as Akerlof (1970), another seminal paper appeared on adverse selection, Mirrlees (1971), although the relation only became clear later. Mirrlees looked at optimal taxation and the problem of how the government chooses a tax schedule given that it cannot observe the abilities of its citizens to earn income, and this began the literature on mechanism design. Used cars and income taxes do not appear similar, but in both situations an uninformed player must decide how to behave to another player whose type he does not know. Section 10.4 sets out a descendant of Mirrlees (1971) in a model of government procurement: much of government policy is motivated by the desire to create incentives for efficiency at minimum cost while eliciting information from individuals with superior information.

N9.2 Adverse selection under certainty: Lemons I and II

- Suppose that the cars of Lemons II lasted two periods and did not physically depreciate. A naive economist looking at the market would see new cars selling for \$6,000 (twice \$3,000) and old cars selling for \$2,000 and conclude that the service stream had depreciated by 33 percent. Depreciation and adverse selection are hard to untangle using market data.
- Lemons II uses a uniform distribution. For a general distribution F , the average quality $\bar{\theta}(P)$ of cars with quality P or less is

$$\bar{\theta}(P) = E(\theta|\theta \leq P) = \frac{\int_{-\infty}^P xF'(x)dx}{F(P)}. \quad (9.30)$$

Equation (9.30) also arises in physics (the equation for the center of gravity) and nonlinear econometrics (the likelihood equation). Think of $\bar{\theta}(P)$ as a weighted average of the values of θ up to P , the weights being densities. Having multiplied by all these weights in the numerator, we have to divide by their “sum,” $F(P) = \int_{-\infty}^P F'(x)dx$, in the denominator, giving rise to equation (9.30).

N9.3 Heterogeneous tastes: Lemons III and IV

- You might object to a model in which the buyers of used cars value quality more than the sellers, since the sellers are often richer people. Remember that quality here is “quality of used cars,” which is different from “quality of cars.” The utility functions could be made more complicated without abandoning the basic model. We could specify something like $\pi_{buyer} = \theta + k/\theta - P$, where $\theta^2 > k$. Such a specification implies that the lower is the quality of the car, the greater the difference between the valuations of buyer and seller.
- In the original article, Akerlof (1970), the quality of new cars is uniformly distributed between 0 and 2, and the model is set up differently, with the demand and supply curves offered by different types of traders and net supply and gross supply presented rather confusingly. Usually the best way to model a situation in which traders sell some of their endowment and consume the rest is to use only gross supplies and demands. Each old owner supplies his car to the market, but in equilibrium he might buy it back, having better information about his car than the other consumers. Otherwise, it is easy to count a given unit of demand twice, once in the demand curve and once in the net supply curve.
- See Stiglitz (1987) for a good survey of the relation between price and quality. Leibenstein (1950) uses diagrams to analyze the implications of individual demand being linked to the market price of quantity in markets for “bandwagon,” “snob,” and “Veblen” goods.
- Risk aversion is concerned only with variability of outcomes, not their level. If the quality of used cars ranges from 2,000 to 6,000, buying a used car is risky. If all used cars are of quality 2,000, buying a used car is riskless, because the buyer knows exactly what he is getting.
In Insurance Game III in section 9.4, the separating contract for the *Unsafe* consumer fully insures him: he bears no risk. But in constructing the equilibrium, we had to be very careful to keep the *Unsafes* from being tempted by the risky contract designed for the *Safes*. Risk is a bad thing, but as with old age, the alternative is worse. If Smith were certain his car would be stolen, he would bear no risk, because he would be certain to have low utility.
- To the buyers in Lemons IV, the average quality of cars for a given price is stochastic because they do not know which values of ε were realized. To them, the curve $\bar{\theta}(P)$ is only the *expectation* of the average quality.
- **Lemons III': minimum quality of zero:** If the minimum quality of car in Lemons III were 0, not 2,000, the resulting game (Lemons III') would be close to the original Akerlof (1970) specification. As figure 9.8 shows, the supply schedule and the demand schedule intersect at the origin, so that

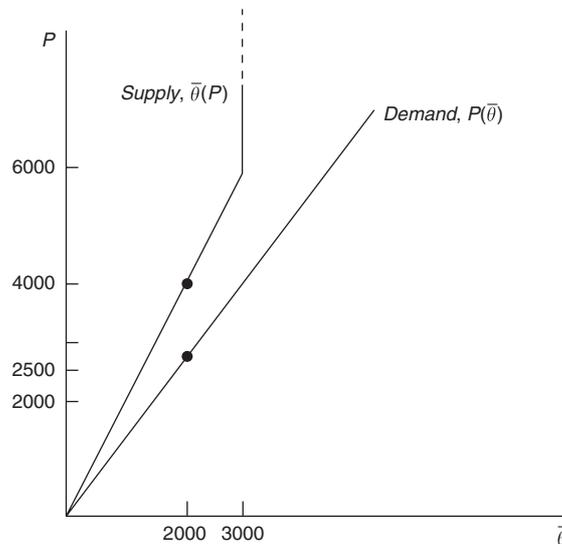


Figure 9.8 Lemons III' when buyers value cars more and the minimum quality is zero.

the equilibrium price is zero and no cars are traded. The market has shut down entirely because of the unravelling effect described in Lemons II. Even though the buyers are willing to accept a quality lower than the dollar price, the price that buyers are willing to pay does not rise with quality as fast as the price needed to extract that average quality from the sellers, and a car of minimum quality is valued exactly the same by buyers and sellers. A 20 percent premium on zero is still zero. The efficiency implications are even stronger than before, because at the optimum all the old cars are sold to new buyers, but in equilibrium, none are.

N9.4 Adverse selection under uncertainty: Insurance Game III

- Markets with two types of customers are very common in insurance, because it is easy to distinguish male from female, both those types are numerous, and the difference between them is important. Males under age 25 pay almost twice the auto insurance premiums of females, and females pay 10–30 percent less for life insurance. The difference goes both ways, however: Aetna charges a 35-year-old woman 30–50 percent more than a man for medical insurance. One market in which rates do not differ much is disability insurance. Women do make more claims, but the rates are the same because relatively few women buy the product (*Wall Street Journal*, August 27, 1987, p. 21).

N9.6 A variety of applications

- Economics professors sometimes make use of self-selection for student exams. One of my colleagues put the following instructions on an MBA exam, after stating that either Question 5 or 6 must be answered. “The value of Question 5 is less than that of Question 6. Question 5, however, is straightforward and the average student may expect to answer it correctly. Question 6 is more tricky: only those who have understood and absorbed the content of the course well will be able to answer it correctly For a candidate to earn a final course grade of A or higher,

it will be *necessary* for him to answer Question 6 successfully.” Making the question even more self-referential, he asked the students for an explanation of its purpose.

Another of my colleagues tried asking who in his class would be willing to skip the exam and settle for an A-. Those students who were willing received an A-. The others got A's. But nobody had to take the exam. (This method did upset a few people.) More formally, Guasch & Weiss (1980) have looked at adverse selection and the willingness of workers with different abilities to take tests.

- Nalebuff & Scharfstein (1987) have written on testing, generalizing Mirrlees (1974), who showed how a forcing contract in which output is costlessly observed might attain efficiency by punishing only for very low output. In Nalebuff & Scharfstein, testing is costly and agents are risk-averse. They develop an equilibrium in which the employer tests workers with small probability, using high-quality tests and heavy punishments to attain almost the first-best. Under a condition which implies that large expenditures on each test can eliminate false accusations, they show that the principal will test workers with small probability, but use expensive, accurate tests when he does test a worker, and impose a heavy punishment for lying.

Problems

9.1: Insurance with equations and diagrams (easy)

The text analyzes Insurance Game III using diagrams. Here, let us use equations too. Let $U(t) = \log(t)$.

- Give the numeric values (x, y) for the full-information separating contracts C_3 and C_4 from figure 9.6. What are the coordinates for C_3 and C_4 ?
- Why is it not necessary to use the $U(t) = \log(t)$ function to find the values?
- At the separating contract under incomplete information, C_5 , $x = 2.01$. What is y ? Justify the value 2.01 for x . What are the coordinates of C_5 ?
- What is a contract C_6 that might be profitable and that would lure both types away from C_3 and C_5 ?

9.2: Testing and commitment (medium)

Fraction β of workers are talented, with output $a_t = 5$, and fraction $(1 - \beta)$ are untalented, with output $a_u = 0$. Both types have a reservation wage of 1 and are risk-neutral. At a cost of 2 to itself and 1 to the job applicant, employer Apex can test a job applicant and discover his true ability with probability θ , which takes a value of something over 0.5. There is just one period of work. Let $\beta = 0.001$. Suppose that Apex can commit itself to a wage schedule before the workers take the test, and that Apex must test all applicants and pay all the workers it hires the same wage, to avoid grumbling among workers and corruption in the personnel division.

- What is the lowest wage, w_t , that will induce talented workers to apply? What is the lowest wage, w_u , that will induce untalented workers to apply? Which is greater?
- What is the minimum accuracy value θ that will induce Apex to use the test? What are the firm's expected profits per worker who applies?
- Now suppose that Apex can pay w_p to workers who pass the test and w_f to workers who flunk. What are w_p and w_f ? What is the minimum accuracy value θ that will induce Apex to use the test? What are the firm's expected profits per worker who applies?

- (d) What happens if Apex cannot commit to paying the advertised wage, and can decide each applicant's wage individually?
- (e) If Apex cannot commit to testing every applicant, why is there no equilibrium in which either untalented workers do not apply or the firm tests every applicant?

9.3: Finding the mixed-strategy equilibrium in a testing game (medium)

Half of high school graduates are talented, producing output $a = x$, and half are untalented, producing output $a = 0$. Both types have a reservation wage of 1 and are risk-neutral. At a cost of 2 to himself and 1 to the job applicant, an employer can test a graduate and discover his true ability. Employers compete with each other in offering wages but they cooperate in revealing test results, so an employer knows if an applicant has already been tested and failed. There is just one period of work. The employer cannot commit to testing every applicant or any fixed percentage of them.

- (a) Why is there no equilibrium in which either untalented workers do not apply or the employer tests every applicant?
- (b) In equilibrium, the employer tests workers with probability γ and pays those who pass the test w , the talented workers all present themselves for testing, and the untalented workers present themselves with probability α , where possibly $\gamma = 1$ or $\alpha = 1$. Find an expression for the equilibrium value of α in terms of w . Explain why α is not directly a function of x in this expression, even though the employer's main concern is that some workers have a productivity advantage of x .
- (c) If $x = 9$, what are the equilibrium values of α , γ , and w ?
- (d) If $x = 8$, what are the equilibrium values of α , γ , and w ?

9.4: Two-time losers (easy)

Some people are strictly principled and will commit no robberies, even if there is no penalty. Others are incorrigible criminals and will commit two robberies, regardless of the penalty. Society wishes to inflict a particular penalty on criminals as retribution. Retribution requires an expected penalty of 15 per crime (15 if detection is sure, 150 if it has probability 0.1, etc.). Innocent people are sometimes falsely convicted, as shown in table 9.2.

Table 9.2 Two-time losers

Robberies	Convictions		
	0	1	2
0	0.81	0.18	0.01
1	0.60	0.34	0.06
2	0.49	0.42	0.09

Two systems are proposed: (1) a penalty of X for each conviction, and (2) a penalty of 0 for the first conviction, and some amount P for the second conviction.

- (a) What must X and P be to achieve the desired amount of retribution?
- (b) Which system inflicts the smaller cost on innocent people? How much is the cost in each case?
- (c) Compare this with problem 8.2. How are they different?

9.5: Insurance and state-space diagrams (medium)

Two types of risk-averse people, clean-living and dissolute, would like to buy health insurance. Clean-living people become sick with probability 0.3, and dissolute people with probability 0.9. In state-space diagrams with the person's wealth if he is healthy on the vertical axis and if he is sick on the horizontal, every person's initial endowment is (5, 10), because his initial wealth is 10 and the cost of medical treatment is 5.

- (a) What is the expected wealth of each type of person?
- (b) Draw a state-space diagram with the indifference curves for a risk-neutral insurance company that insures each type of person separately. Draw in the post-insurance allocations C_1 for the dissolute and C_2 for the clean-living under the assumption that a person's type is contractible.
- (c) Draw a new state-space diagram with the initial endowment and the indifference curves for the two types of people that go through that point.
- (d) Explain why, under asymmetric information, no pooling contract C_3 can be part of a Nash equilibrium.
- (e) If the insurance company is a monopoly, can a pooling contract be part of a Nash equilibrium?

Adverse Selection in Stock Sales: A Classroom Game for Chapter 9

This is a game modelling a financial market in which sellers of stock know the value of the stock better than buyers do, but buyers get more utility than sellers from a given asset.

There are two kinds of people: buyers and sellers. Each buyer starts with a checking account with \$200. Each seller starts with 4 stock certificates from 4 different companies. Stock face values are \$90, \$70, \$30, and \$10 in equal proportions.

Buyers value stocks more than cash, because they wish to save. Each buyer has the following payoff function, where payoff is measured in utils.

$$\pi(\text{buyer}) = \text{cash} + 1.5 * (\text{stock value}). \quad (9.31)$$

The initial payoff of a buyer is 200 utils from his \$200 in cash, but it would rise to 300 utils if the buyer could convert his \$200 in cash to \$200 in stocks.

Sellers value cash more than stocks, because they wish to consume now. Each seller has payoff function

$$\pi(\text{seller}) = 1.5 * \text{cash} + (\text{stock value}). \quad (9.32)$$

In the game, buyers and sellers buy and sell stock. The buyers pay with checks registered on the blackboard. If buyer Smith buys a stock certificate for \$40 from Jones, he writes “Smith” on the board, and under it writes “Owes \$40 to Jones.” If Smith then buys a certificate for \$90 from Lee, he writes “Owes \$90 to Lee” under the first entry.

The game is repeated with different features. Each time, sellers start with 4 new shares, and buyers start with \$200. Players do *not* keep their earnings from previous rounds. (Think of the 1993 movie, *Groundhog Day*, <http://www.imdb.com/title/tt0107048>.)

- 1 Symmetric information: The instructor writes the value of each company on the board.
- 2 Asymmetric information – decentralized: The instructor tells the sellers, but not the buyers, the value of each company.
- 3 The same as round 2, but starting over with different shares.
- 4 Asymmetric information – centralized: Partial regulation: Sellers must disclose if their stock is worth less than 50. The instructor tells the sellers, but not the buyers, the value of each company.
- 5 Asymmetric information – centralized: Truth regulation: Sellers *may* guarantee their stock value. The instructor tells the sellers, but not the buyers, the value of each company.