

Chapter 17: Asymmetric Information


Asymmetric information occurs when some parties to a particular transaction know more than others. Common examples of asymmetric information include the following:

- You know how productive you are, but the person interviewing you doesn't.
- Your mechanic knows whether your car needs a new fuel filter, but you don't.
- You know how healthy you are, but your insurance company doesn't.

As we will see, asymmetric information often leads to market failure, in the sense that markets will not operate as efficiently as we would like in the absence of perfect information.

The “Lemons” Problem

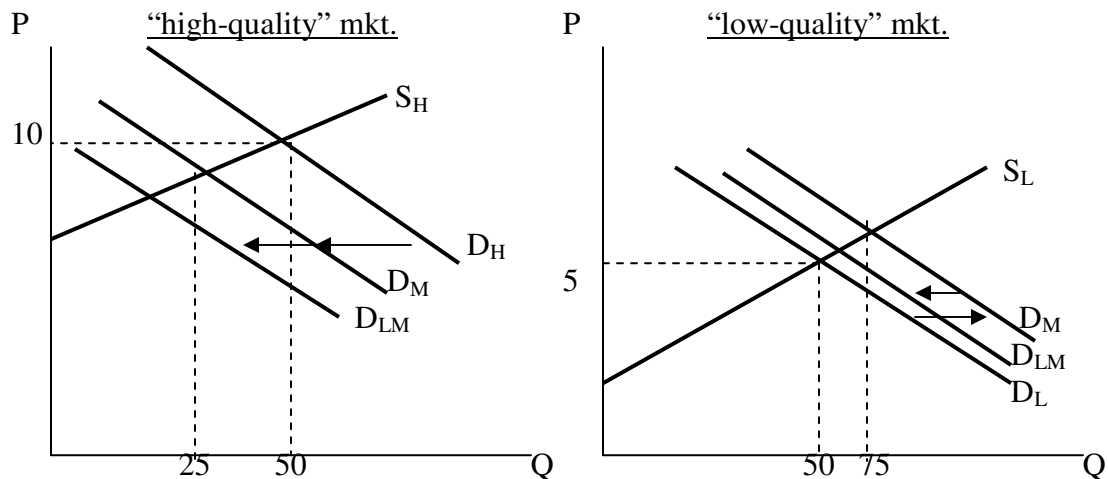
Consider the market for used cars. For simplicity, let's suppose that there are two types of cars – high quality and low quality (“lemons”). Initially, let us also suppose, unrealistically, that when a car is offered for sale, the buyer and seller both know the car's type. In that case, we would expect to see two distinct markets: a market for high quality used cars, and a market for lemons.



Let's suppose that, if there were really two markets, that 50,000 of each type of car would be sold, and that $P_H = \$10,000$ and $P_L = \$5,000$. Notice that S_H is lower than S_L , because people are less willing to sell high-quality, reliable cars and must be paid more for them. Also, D_L is lower than D_H because buyers aren't willing to pay as much for a low-quality car. These two assumptions result in the price of a low-quality car being less than the price of a high-quality car.

Now suppose that there is asymmetric information. Specifically, let's suppose that the seller knows whether the car is a lemon or not, but the buyer does not. After all, the seller has had much more experience with this car than the potential buyer has. How would buyers respond to this lack of information?

At first glance, it might appear that since there are 50,000 of each type of car on the market, that the odds of any particular car being high-quality are 50-50. Thus, demand will change in each "market." In the market for high-quality cars, each consumer perceives a chance that the car he's buying is actually a lemon. Likewise, in the low-quality market, each consumer perceives a chance that the offered car is actually high-quality. (Note that consumers, due to incomplete information, do not actually know what market they are participating in – there will be one demand curve that is the same in both markets). Let's suppose that demand thus becomes D_M ("medium" quality) in both markets.



With demand falling in the high-quality market and rising in the low-quality market, the prices of both types of cars move closer together. However, the quantity of high-quality cars falls while the quantity of low-quality cars rises. Consumers can predict this, so they know that the updated odds of any particular car being high-quality are actually less than 50-50 now. Thus, demand falls as consumers perceive the average car to be of "low-medium" quality (D_{LM}). Note that this shift will be a downward shift in the "high-quality" market *and* in the "low-quality" market since the demand for low-medium cars will always be lower than medium cars.

This shift reduces the proportion of good cars even more, further reducing demand. This process will continue until demand has fallen so low that only low-quality cars are sold on the market. This is an equilibrium because consumers *expect* only lemons, and those expectations are correct, since nobody is willing to sell a high-quality car at such low prices.

This is an example of the lemons problem, showing that low-quality goods tend to drive out high-quality goods when information is asymmetric.

Example: Professional sports and free agency

Similar problems arise when we consider . . .

Adverse Selection

Sometimes a market is inherently set up in such a way that when a seller attempts to sell a product, the only people who buy are the people the seller does *not* want to sell to. This problem is called adverse selection, because sellers will consistently choose the “wrong” customers.

The classic example of adverse selection is the market for health insurance. If I go to buy health insurance on my own, it is extremely expensive. However, if I get health insurance through my employer, it is pretty cheap. Why? The answer is that when health insurance companies offer a policy to an individual, the first people in line to buy the policy are those with health problems. Since I know more about my health than my insurance company does (even after being examined by a doctor), the insurance company knows that it will systematically sell policies to sick people, who will be expensive. Thus, they must charge higher premiums. If I am 28 and healthy, I will be unwilling to pay these higher premiums, so the insurance company sells only to people who are sick and charges very high premiums.

Notice that this is not a problem if I get health insurance through my employer. If *everyone* at SDSU gets health insurance, the insurance company knows that some will be sick but some will be healthy. Thus, it can charge lower premiums. This is the reason why many health-care reform proposals include requirements that everyone be forced to buy health insurance, whether they want it or not, so that adverse selection does not become an issue.

Example: Credit cards

Signaling

One way to get around adverse selection is for customers to send a signal of what “type” of customer they are. Suppose you go to apply for a job. Your potential employer wants to know whether you are a “high-productivity” worker or a “low-productivity” worker. Obviously, you can’t just go in and say “I’m a high productivity worker.” That statement carries no information, because the interviewer knows that it is costless for you to say that. What you need to do is come up with some way to credibly demonstrate that you are a high-productivity worker. There are several ways to try to do this.

For example, you could dress up for the interview. By wearing a nice suit, you send a signal that you care about your job and want to appear professional. This could be interpreted as a signal that you are a “high-productivity” worker. The problem is that anybody can do this. Thus, this does not send a very strong signal.

In order for the signal to convey information, we need two conditions to hold:

1) The signal must be costly. If the signal is cheap or costless (like saying “I’m a good worker”), then *everybody* will send the signal, both good and bad workers, and it will convey no information. The signal needs to be costly so that some workers will opt not to send it.

2) The signal must be less costly for high-productivity workers than low-productivity workers. Presumably, I benefit from sending the signal that I am a high-productivity worker. If it is relatively cheap (the benefits of sending the signal outweigh the costs) for me to send this signal, I will rationally choose to do so. If it is more expensive for low-productivity workers to send the signal (so that the cost of doing so outweighs the benefits), then low-productivity workers will rationally choose not to do so.

Notice that if the signal is such that

$$\text{Cost}_{\text{LOW}} > \text{Benefit} > \text{Cost}_{\text{HIGH}}$$

then only high-quality workers will send the signal. Thus, if the interviewer observes the signal, he will infer that you are high-productivity. If he does not observe the signal, he can infer that you are low-productivity.

Is there a signal for job applicants that satisfies these conditions? Most economists would say yes. Specifically, many labor economists believe that college serves as the signal described here.

1) College obviously is costly. Very costly. Much more costly than just buying a new suit. It’s not just that you pay tens of thousands of dollars in tuition and fees. You also give up the chance to have a real job for the four years you’re in school, costing you perhaps another hundred thousand in forgone earnings.

2) College is less costly for “high-productivity” workers. Everybody pays the same number of dollars, but some people like college while others dislike it. Intelligent people with good work habits and good time-management skills (high productivity) tend to do well in college and enjoy it, or at least dislike it less. People who are less bright, lazy, and unable to manage their time well (low-productivity) often fail in college or greatly dislike it. These people might be unable to send the signal (finish college) even if they wanted to, and it is certainly more costly for them to send the signal since their utility during their college years is much lower than for high-productivity workers.

Some economists have gone so far as to argue that what you learn in college is irrelevant. Economically, college is just a way to sort out the men from the boys. While many would argue that college *does* serve this role to some extent, it might appear that there is more than signaling going on. For instance, why do philosophy majors make less than business majors? Philosophy is a much more difficult discipline than business, so one would think, according to signaling theory, that a philosophy major sends a stronger signal than a business major. Maybe majoring in business sends a signal that you really like that kind of career and were willing to spend four years of your youth studying it. In any case, the role of signaling in higher education is certainly one to think about.

Example: Warranties

Moral Hazard

Another information problem that arises is when a party is able to take an unobserved action that affects another party. This is often a problem in the insurance industry. Moral hazard refers to a situation in which an insured party can take unobserved actions that affect the probability of an insurance claim.

e.g. I could buy a home security system. But this costs me money, and my belongings are insured. Therefore, I choose not to buy the security system. But this increases the likelihood of an insurance claim. My insurance company would want me to install the security system, but because it cannot observe my choice, it is stuck.

Notice that this can cause real problems in insurance markets. Consider the case of auto insurance. If I had no insurance, I would drive very carefully. My insurance company would be happy to sell me insurance if I drive carefully. But my insurance company cannot observe my driving habits. Once it sells me insurance, I no longer have to worry about damage to my car, so I can drive recklessly. But my insurance company isn't stupid. It knows that once it sells me insurance, I will drive recklessly. Thus, it charges higher premiums to compensate for my reckless driving.

Both my insurance company and I would be better off if I could somehow credibly promise to keep driving safely after getting insurance. My insurer would face fewer claims, and I would therefore be able to pay lower premiums. But there is no way for me to credibly promise to drive carefully after buying insurance. Hence, moral hazard drives up insurance premiums for everybody, making insurance less affordable.

You can see the basic problem here. After buying insurance, my incentives are not aligned with the incentives of my insurer. My insurer wants me to drive safely, but I have little incentive to do so. We can attempt to minimize the moral hazard problem in several ways:

- 1) Changing premiums based on observable information. If you get frequent speeding tickets, your insurance company knows you are driving recklessly, so you get charged higher premiums. If your driving record is always clean, that shows that you are probably driving safely, so your premiums are lower. Note that in this case, your insurer can at least imperfectly observe your driving habits and set premiums accordingly.
- 2) Co-payments. Some types of insurance require that I pay a percentage of any claim. For instance, most health insurance policies require that I pay 10% of any medical bill. This gives me a financial incentive to stay in good health, since I will pay *something* if medical treatment is necessary. One could do the same with auto insurance.
- 3) Deductibles. If I have an auto accident, I will have to pay the first \$500 of any damage. This at least gives me some incentive to drive carefully, so as to avoid having to pay this deductible. However, note that my incentive to drive carefully is much weaker if I face only a \$500 bill from an accident than if I face a \$25,000 bill for a new car. Or a \$250,000 personal injury judgment.

The Principal-Agent Problem

The last type of asymmetric information we will consider is the principal-agent problem. In this situation, one party (the principal) delegates decision making authority to another party (the agent). The problem occurs when the agent's actions are not directly observed by the principal, and the agent has an incentive to behave in ways not in the principal's interest.

The classic example of a principal-agent problem occurs in the management of a corporation. As you know, a corporation is owned by its stockholders. However, these stockholders do not actually manage the corporation. They hire managers to run the corporation for them. The problem is that stockholders cannot easily observe all of management's actions, and the goals of management may not be the same as the stockholder's goals. Thus the principal-agent problem, with stockholders being the principal and management being the agent.

For instance, stockholders might want the firm to maximize profits. Part of maximizing profit is minimizing cost, which involves keeping managerial salaries down, or at least in line with what comparable managers are paid elsewhere. However, managers know that if they maximize the company's sales, or cash flow, that they will have the resources to pay themselves inflated salaries, or provide nice perks like country club memberships. Granted, it is hypothetically possible for stockholders to find out about these shenanigans, but it is quite difficult.