

Interpreting and Manipulating Information

THE MARRYING KIND?

True story: Our friend, whom we'll call Sue, was in love. Her beau was an extremely successful executive. He was smart, single, and straight. He professed his love to her. It was a happily-ever-after fairy tale. Well, almost.

The problem was that, at age thirty-seven, Sue wanted to get married and have kids. He was on board with the plan, except that his kids from a previous marriage weren't ready for him to remarry. These things take time, he explained. Sue was willing to wait, so long as she knew that there would be a light at the end of the tunnel. How could she know whether his words were sincere or not? Unfortunately, any public demonstration was out of bounds, as the kids would surely find out.

What she wanted was a credible signal. This is the cousin of a commitment device. In the previous chapter, we emphasized strategies that guaranteed that the person would carry out what he said he'd do. Here, we are looking for something weaker. What Sue wanted was something that would help her understand whether he was truly serious about their relationship.

After much thought, Sue asked him to get a tattoo, a tattoo with her name. A small, discreet tattoo would be just fine. No one else would ever have to see it. If he was in this for the long run, then having Sue's name indelibly inked would be a fitting tribute to their love. But, if commitment wasn't part of his plan, this

would be an embarrassing artifact for his next conquest to discover.

He balked, so Sue left. She found a new love and is now happily married with kids. As for her ex, he is still on the runway, on permanent ground delay.

Tell It Like It Is?

Why can't we just rely on others to tell the truth? The answer is obvious: because it might be against their interests.

Much of the time, people's interests and communications are aligned. When you order a steak medium rare, the waiter can safely assume that you really want the steak medium rare. The waiter is trying to please you and so you do best by telling the truth. Things get a bit trickier when you ask for a recommended entrée or advice on wine. Now the waiter might want to steer you to a more expensive item and thereby increase the likely tip.

The British scientist and novelist C. P. Snow attributes just such strategic insight to the mathematician G. H. Hardy: "If the Archbishop of Canterbury says he believes in God, that's all in the way of business, but if he says he doesn't, one can take it he means what he says."¹ Similarly, when the waiter points you to the less expensive flank steak or bargain Chilean wine, you have every reason to believe him. The waiter might also be right when recommending the expensive entrée, but it is harder to know.

The greater the conflict, the less the message can be trusted. Recall the soccer penalty kicker and the goalie from chapter 5. Suppose that, just as he is getting ready to take his shot, the kicker says: "I am going right." Should the goalie believe him? Of course not. Their interests are totally opposed, and the kicker stands to lose by making his intentions known truthfully in advance. But does this mean that the goalie should assume that the kicker will kick to the left? Again, no. The kicker might be trying a second-level deception—lying by telling the truth. The only rational reaction to an assertion made by another player whose interests are totally opposed to yours is to ignore it completely. Don't assume it to be true, but don't assume its opposite to be true either. (Instead, think about the equilibrium of the actual game ignoring what the other side has said and play accordingly; later in this chapter, we explain just how to do this using the example of bluffing in poker.)

Politicians, advertisers, and children are all players in their own strategic games with their own interests and incentives. And what they are telling us serves their own agendas. How should you interpret information that comes from such sources? And conversely, how can you make your claims credible, knowing

that others will regard what you say with due suspicion? We start our exploration with perhaps the most famous example of divining the truth from interested parties.

KING SOLOMON'S DILEMMA

Two women came before King Solomon, disputing who was the true mother of a child. The Bible takes up the story in 1 Kings (3:24–28):

Then the king said, “Bring me a sword.” So they brought a sword for the king. He then gave an order: “Cut the living child in two and give half to one and half to the other.” The woman whose son was alive was filled with compassion for her son and said to the king, “Please, my lord, give her the living baby! Don’t kill him!” But the other said, “Neither I nor you shall have him. Cut him in two!” Then the king gave his ruling: “Give the living baby to the first woman. Do not kill him; she is his mother.” When all Israel heard the verdict the king had given, they held the king in awe, because they saw that he had wisdom from God to administer justice.

Alas, strategic experts cannot leave a good story alone. Would the king’s device have worked if the second woman, the false claimant, had understood what was going on? No.

The second woman made a strategic blunder. It was her answer in favor of dividing the child that distinguished her from the true mother. She should have simply repeated whatever the first woman said; with both women saying the same thing, the king would not have been able to say which one was the true mother.

The king was more lucky than wise; his strategy worked only because of the second woman’s error. As for what Solomon should have done, we offer that as a case study in chapter 14.

DEVICES FOR MANIPULATING INFORMATION

The kinds of problems faced by Sue and Solomon arise in most strategic interactions. Some players know more than others about something that affects the payoffs for them all. Some who possess extra information are keen to

conceal it (like the false claimant); others are equally keen to reveal the truth (like the true mother). Players with less information (like King Solomon) typically want to elicit it truthfully from those who know.

Pretending wisdom greater than that of Solomon, game theorists have examined several devices that serve these purposes. In this chapter we will illustrate and explain them in simple terms.

The general principle governing all such situations is: Actions (including tattoos) speak louder than words. Players should watch what another player does, not what he or she says. And, knowing that the others will interpret actions in this way, each player should in turn try to manipulate actions for their information content.

Such games of manipulating behavior to manipulate others' inferences, and seeing through others' manipulation of our inferences, go on every day in all of our lives. To borrow and twist a line from *The Love Song of J. Alfred Prufrock*, you must constantly "prepare a face to meet the faces that you meet." If you do not recognize that your "face," or more generally your actions, are being interpreted in this way, you are likely to behave in a way that works to your own disadvantage, often quite seriously so. Therefore the lessons of this chapter are among the most important you will learn in all of game theory.

Strategic game players who possess any special information will try to conceal it if they will be hurt when other players find out the truth. And they will take actions that, when appropriately interpreted, reveal information that works favorably for them. They know that their actions, like their faces, leak information. They will choose actions that promote favorable leakage; such strategies are called *signaling*. They will act in ways that reduce or eliminate unfavorable leakage; this is *signal jamming*. It typically consists of mimicking something that is appropriate under different circumstances than the ones at hand.

If you want to elicit information from someone else, you should set up a situation where that person would find it optimal to take one action if the information was of one kind, and another action if it was of another kind; action (or inaction) then reveals the information.* This strategy is called *screening*. For example, Sue's request for a tattoo was her screening test. We will now illustrate and explain the working of these devices.

In chapter 1, we argued that poker players should conceal the true strength of their hand by bidding somewhat unpredictably. But the optimal mix of bids is different for hands of different strengths. Therefore, limited information about the probability of a strong hand can be derived from the bids. The same principle holds when someone is trying to convey rather than conceal information:

Actions speak louder than words. To be an effective signal, an action should be incapable of being mimicked by a rational liar: it must be unprofitable when the truth differs from what you want to convey.²

Your personal characteristics—ability, preferences, intentions—constitute the most important information that you have and others lack. They cannot observe these things, but you can take actions that credibly signal the information to them. Likewise, they will attempt to infer your characteristics from your actions. Once you become aware of this, you will start seeing signals everywhere and will scrutinize your own actions for their signal content.

When a law firm recruits summer interns with lavish hospitality, it is saying, “You will be well treated here, because we value you highly. You can believe us because if we valued you less then we would not find it in our interest to spend so much money on you.” In turn, the interns should realize that it doesn’t matter if the food is bad or the entertainment bores them stiff; what’s important is the price.

Many colleges are criticized by their alumni for teaching things that proved of no use in their subsequent careers. But such criticism leaves out the signaling value of education. Skills needed to succeed in particular firms and specialized lines of work are often best learned on the job. What employers cannot easily observe but really need to know is a prospective employee’s general ability to think and learn. A good degree from a good college acts as a signal of such ability. The graduate is in effect saying, “If I were less able, would I have graduated Princeton with honors?”

But such signaling can turn into a rat race. If the more able get only a little more education, the less able might find it profitable to do likewise, be mistaken for the more able, and be given better jobs and wages. Then the truly more able must get even more education to distinguish themselves. Pretty soon, simple clerical jobs require master’s degrees. True abilities remain unchanged; the only people to benefit from the excessive investment in education for signaling are we college professors. Individual workers or firms can do nothing about this wasteful competition; a public policy solution is needed.

IS THE QUALITY GUARANTEED?

Suppose you are in the market to buy a used car. You find two that seem to have the same quality, as far as you can judge. But the first comes with a warranty and the second does not. You surely prefer the first, and are willing to pay more for it. For one thing, you know that if something goes wrong, you will

get it fixed free of charge. However, you will still have to spend a lot of time and suffer a lot of inconvenience, and you are not going to be compensated for these hassles. Here another aspect becomes more relevant. You believe that things are less likely to go wrong with the car under warranty in the first place. Why? To answer that, you have to think about the seller's strategy.

The seller has a much better idea of the quality of the car. If he knows that the car is in good condition and not likely to need costly repairs, offering the warranty is relatively costless to him. However, if he knows that the car is in poor condition, he expects to have to incur a lot of cost to fulfill the warranty. Therefore, even after taking into account the higher price that a car under a warranty may fetch, the worse the quality of the car, the more likely the warranty is to be a losing proposition to the seller.

Therefore the warranty becomes an implied statement by the seller: "I know the quality of the car to be sufficiently good that I can afford to offer the warranty." You could not rely on the mere statement: "I know this car to be of excellent quality." With the warranty, the seller is putting his money where his mouth is. The action of offering the warranty is based on the seller's own gain and loss calculation; therefore it is credible in a way that mere words would not be. Someone who knew his car to be of low quality would not offer the warranty. Therefore the action of offering a warranty serves to separate out sellers who merely "talk the talk" from those who can "walk the walk."

Actions that are intended to convey a player's private information to other players are called *signals*. For a signal to be a credible carrier of a specific item of information, *it must be the case that the action is optimal for the player to take if, but only if, he has that specific information*. Thus we are saying that offering a warranty can be a credible signal of the quality of the car. Of course whether it is credible in a specific instance depends on the kinds of things that are potentially likely to go wrong with that kind of car, the cost of fixing them, and the difference in price between a car under a warranty and a similar-looking car without a warranty. For example, if the expected cost of repairs on a good-quality car is \$500, while that for a poor-quality car is \$2,000, and the price difference with and without a warranty is \$800, then you can infer that a seller offering such a warranty knows his car to be of good quality.

You don't have to wait for the seller to think all this through and offer the warranty if he knows his car to be good. If the facts are as we just stated, you can take the initiative and say: "I will pay you an extra \$800 for the car if you offer me a warranty." This will be a good deal for the seller if, but only if, he knows his car to be of good quality. In fact you could have offered \$600, and he might counter with \$1,800. Any price greater than \$500 and less than \$2,000 for the

warranty will serve to induce sellers of good and bad cars to take different actions and thereby reveal their private information, and the two of you might bargain over this range.

Screening comes into play when the less-informed player requires the more-informed player to take such an information-revealing action. The seller might take the initiative and signal the quality of the car by offering the warranty, or the buyer might take the initiative and screen the seller by asking for a warranty. The two strategies can work in similar ways to reveal private information, although there can be technical gametheoretic differences between the resulting equilibria. When both methods are potentially available, which one is used can depend on the historical, cultural, or institutional context of the transaction.

A credible signal has to be against the interests of an owner who knows his car to be of low quality. To drive home the point, how would you interpret a seller's offer to let you get the car inspected by a mechanic? This is not a credible signal. If the mechanic finds some serious flaw and you walk away, the owner is no worse off than before, regardless of the condition of his car. Therefore the owner of a bad car can make the same offer; the action will not serve to convey the information credibly.*

Warranties are credible signals because they have the crucial cost-difference property. Of course the warranty itself has to be credible in the sense that you can enforce its terms when the need arises. Here we see a big difference between a private seller and a car dealership. Enforcement of a warranty given by a private seller is likely to be much harder. Between the time when the car is sold and when the need for a repair arises, a private seller may move, leaving no forwarding address. Or he may lack the money to pay for the repair, and taking him to court and enforcing a judgment may be too costly to the buyer. A dealership is more likely to be in the business for a longer time and may have a reputation to preserve. Of course a dealer can also try to weasel out of payment by claiming that the problem arose because you did not maintain the car properly or drove it recklessly. But on the whole, revelation of the quality of a car (or other consumer durables) through warranties or other methods is likely to be far more problematic for private transactions than for sales by established dealers.

A similar problem exists for car manufacturers who have not yet established a reputation for high quality. In the late 1990s, Hyundai raised the quality of their cars, but this had not yet been recognized by U.S. consumers. To get its claims of quality across in a dramatic and credible way, in 1999 the company signaled its quality by offering an unprecedented 10-year, 100,000-mile warranty on the power train and 5 years, 50,000 miles on the rest.

A LITTLE HISTORY

George Akerlof chose the used car market as the main example in his classic article showing that information asymmetries can lead to market failures.³ To illustrate the issue in the simplest way, suppose there are just two types of used cars: lemons (bad quality) and peaches (good quality). Suppose that the owner of each lemon is willing to sell it for \$1,000, whereas each potential buyer is willing to pay \$1,500 for a lemon. Suppose the owner of each peach is willing to sell it for \$3,000, whereas each potential buyer is willing to pay \$4,000 for a peach. If the quality of each car were immediately observable to all parties, then the market would work well. All cars would be traded, lemons selling for a price somewhere between \$1,000 and \$1,500, and each peach between \$3,000 and \$4,000.

But suppose each seller knows the quality of a car, whereas all that buyers know is that half the cars are lemons and half are peaches. If cars are offered for sale in the same proportion, each buyer would be willing to pay at most

$$\frac{1}{2} \times (\$1,500 + \$4,000) = \$2,750.$$

An owner who knows his car to be a peach is not willing to sell at this price.* Therefore only lemons will be offered for sale. Buyers, knowing this, would offer at most \$1,500. The market for peaches would collapse completely, even though buyers are willing to pay a price for provable peaches that sellers are happy to accept. The Panglossian interpretation of markets, namely that they are the best and most efficient institutions for conduct of economic activity, breaks down.

One of us (Dixit) was a graduate student when Akerlof's article first appeared. He and all the other graduate students immediately recognized it as a brilliant and startling idea, the stuff of which scientific revolutions are made. There was just one problem with it: almost all of them drove used cars, most of which they had bought in private deals, and most of which were not lemons. There must be ways in which market participants cope with the information problems that Akerlof had brought to our attention in such a dramatic example.

There are some obvious ways. Some students have a fair bit of mechanical knowledge about cars, and the rest of them can enlist a friend to inspect a car they are thinking of buying. They can get information about the history of the car from networks of mutual friends. And many owners of high-quality cars are forced to sell them at almost any price, because they are moving far away or

even out of the country, or have to switch to bigger cars as their families grow, and so on. Thus there are many practical ways in which markets can mitigate Akerlof's lemons problem.

But we had to wait until Michael Spence's work for the next conceptual breakthrough, namely how strategic actions can communicate information.* He developed the idea of signaling and elucidated the key property—the differences in payoffs from taking an action for players who have different information—that can make signals credible.

The idea of screening evolved from the work of James Mirrlees and William Vickrey but received its clearest statement in the work of Michael Rothschild and Joseph Stiglitz on insurance markets. People have better information about their own risks than do the companies from whom they seek insurance. The companies can require them to take actions, typically to choose from among different plans with different provisions of deductibles and coinsurance. The less risky types will prefer a plan that has a smaller premium but requires them to bear a larger fraction of the risk; this is less attractive to those who know themselves to have higher risk. Thus the choice reveals the insurance applicant's risk type.

This idea of screening by letting people make choices from a suitably designed menu has since become the key to our understanding of many features commonly found in markets, for example, the restrictions on discounted tickets that airlines impose. We will discuss some of these later in this chapter.

The insurance market provided one other input to this topic of information asymmetries. Insurers have long known that their policies selectively attract the worst risks. A life insurance policy that charges the premium of, say, 5 cents for every dollar of coverage will be especially attractive to people whose mortality rate is greater than 5 percent. Of course many people with lower mortality rates will still buy policies, because they need to protect their families, but those at greatest risk will be overrepresented and will buy bigger policies. Raising the price can make matters worse. Now the good risks find the policies too expensive, leaving behind just the worse cases. Once again we have the Groucho Marx effect: anyone willing to buy insurance at those prices is not someone you would want to insure.

In Akerlof's example, potential buyers do not directly know the quality of an individual car and therefore cannot offer different prices for different cars. Thus selling becomes selectively attractive to the owners of lemons. Because the relatively "bad" types are selectively attracted to the transaction, the problem came to be called *adverse selection* in the insurance industry, and the line of research in game theory and economics that deals with problems caused by

information asymmetries has inherited that name.

Just as adverse selection is a problem, sometimes the effect can be turned on its head to create “positive selection.” Starting from its IPO in 1994, Capital One was one of the most successful companies in America. It had a decade of 40 percent compounded growth—and that is excluding mergers and acquisitions. The key to its success was a clever application of selection. Capital One was a new player in the credit card business. Its big innovation was the transfer of balance option, wherein a customer could bring over an outstanding balance from another credit card and get a lower interest rate (at least for some period).

The reason why this was such a profitable offer comes down to positive selection. Roughly speaking, there are three types of credit card customers, what we will call maxpayers, revolvers, and deadbeats. Maxpayers are the folks who pay their bills in full each month and never borrow on the card. Revolvers are the ones who borrow money on the card and pay it back over time. Deadbeats are also borrowers but, unlike revolvers, are going to default on the loan.

From the credit card issuer’s perspective, they obviously lose money on deadbeats. Revolvers are the most profitable of all customers, especially given the high interest rate on credit cards. It may be surprising, but credit card companies also lose money on maxpayers. The reason is that the fees charged to merchants just barely cover the free one-month loan given to these customers. The small profit doesn’t cover billing costs, fraud, and the risk, small but not negligible, that the maxpayer will get divorced (or lose his job) and then default.

Consider who will find the transfer of balance option attractive. Since the maxpayer isn’t borrowing money on the card, there is no reason to switch over to Capital One. The deadbeat is not planning to pay the money back, so here, too, there is little interest in switching. Capital One’s offer is most attractive to the customers who have large amounts outstanding and are planning to pay the loan back. While Capital One may not be able to identify who the profitable customers are, the nature of its offer ends up being attractive just to the profitable type. The offer screens out the unprofitable types. This is the reverse of the Groucho Marx effect. Here, any customer who accepts your offer is one you want to take.

SCREENING AND SIGNALING

You are the chief personnel officer of a company, looking to recruit bright young people who have natural-born talent as managers. Each candidate knows whether he or she has this talent, but you don’t. Even those lacking the talent

look for jobs in your firm, hoping to make a good salary until they are found out. A good manager can generate several million dollars in profits, but a poor one can rack up large losses quickly. Therefore you are on the lookout for evidence of the necessary talent. Unfortunately, such signs are hard to come by. Anyone can come to your interview wearing the right dress and professing the right attitudes; both are widely publicized and easy to imitate. Anyone can get parents, relatives, and friends to write letters attesting to one's leadership skills. You want evidence that is credible and hard to mimic.

What if some candidates can go to a business school and get an MBA? It costs around \$200,000 to get one (when you take into account both tuition and foregone salary). College graduates without an MBA, working in an environment where the specialized managerial talent is irrelevant, can earn \$50,000 per year. Supposing people need to amortize the expense incurred in earning an MBA over five years, you will have to pay at least an extra \$40,000 a year—that is, a total of \$90,000 a year—to a candidate with an MBA.

However, this will make no difference if someone who lacks managerial talent can get an MBA just as easily as someone with this talent. Both types will show up with the certificates, expecting to earn enough to pay off the extra expense and still get more money than they could in other occupations. An MBA will serve to discriminate between the two types only if those with managerial talent somehow find it easier or cheaper to earn this degree.

Suppose that anyone possessing this talent is sure to pass their courses and get an MBA, but anyone without the talent has only a 50 percent chance of success. Now suppose you offer a little more than \$90,000 a year, say \$100,000, to anyone with an MBA. The truly talented find it worthwhile to go and get the degree. What about the untalented? They have a 50 percent chance of making the grade and getting the \$100,000 and a 50 percent chance of failing and having to take another job for the standard \$50,000. With only a 50 percent chance of doubling their salary, an MBA would net them only \$25,000 extra salary on average, so they cannot expect to amortize their MBA expenses over five years. Therefore they will calculate that it is not to their advantage to try for the MBA.

Then you can be assured that anyone with an MBA does have the managerial ability you need; the larger pool of college graduates has sorted itself into two pools in just the right way for you. The MBA serves as a screening device. We emphasize once again that it works because the cost of using the device is less for those you want to attract than for those you want to avoid.

The irony of this is that companies could just as well hire the MBA students on the first day of classes. When the screening device works, only the ones with managerial ability show up. Therefore, firms don't need to wait until the students

graduate to know who's talented and who isn't. Of course, if this practice were to become common, then untalented students would start to show up and be the first in line to drop out. The screening only works so long as people spend the two years to make it through.

Thus this screening device comes at a significant cost. If you could identify the talented directly, you could get them to work for you for just over the \$50,000 that they could have earned elsewhere. Now you have to pay the MBAs more than \$90,000 to make it worth the while of talented students to incur this extra expense in order to identify themselves. The extra \$40,000 per year for five years is the cost of overcoming your informational disadvantage.

The cost can be attributed to the existence of the untalented in the population. If everyone were a good manager, you would not need to do any screening. Thus the untalented, by their mere existence, are inflicting a negative spillover, or a negative externality in the language of economics, on the rest. The talented initially pay the cost, but the company then has to pay them more, so in the end the cost falls on the company. Such "informational externalities" pervade all of the examples below, and you should try to pinpoint them in order to understand exactly what is going on in each.

Is it really worth your while to pay this cost, or would you do better to hire randomly from the whole pool at \$50,000 each and take your chances of hiring some untalented people who will cost you money? The answer depends on what proportion of the population is talented, and the size of the losses that each of them can inflict on your firm. Suppose 25 percent of the population of college graduates lacks managerial talent, and each of them can run up losses of a million dollars before they are found out. Then the random hiring policy will cost you \$250,000 per hire, on average. That exceeds the \$200,000 cost (\$40,000 extra salary over five years) of using the MBA to screen out the untalented. Actually, the proportion with managerial talent is probably much smaller, and the potential loss from poor strategies much larger, so the case for using costly screening devices is much stronger. We like to think that the MBA does teach them a few useful skills, too.

ONE REASON TO GET AN MBA:

A prospective employer may be concerned about hiring and training a young woman only to find that she leaves the labor force to have children. Whether legal or not, such discrimination still arises. How does an MBA help solve the problem?

An MBA serves as a credible signal that the person intends to work for several years. If she was planning to drop out of the labor force in a year, it would not have made sense to have invested the two years in getting an MBA. She would have done much better to have worked for those two years and one more. Practically speaking, it likely takes at least five years to recover the cost of the MBA in terms of tuition and lost salary. Thus you can believe an MBA when she says that she plans to stick around.

Often there are several ways you can identify talent, and you will want to use the cheapest. One way may be to hire people for an in-house training or probationary period. You might let them undertake some small projects under supervision and observe their performance. The cost of this is the salary you have to pay them in the interim, and the risk that the untalented run up some small losses during their probationary period. A second way is to offer contracts with suitably designed backloaded or performance-related compensation. The talented, with confidence in their ability to survive in the firm and generate profits, will be more willing to accept such contracts, while the rest will prefer to take jobs elsewhere that pay a sure \$50,000 a year. A third is to observe the performance of managers in other firms and then try to lure away the proven good ones.

Of course, when all firms are doing this, it alters all their calculations of the costs of hiring apprentices, their salary and performance pay structures, *etc.* Most importantly, competition among firms forces the salaries of the talented above the minimum (for example, \$90,000 with the MBA) needed to attract them. In our example, the salaries could not rise above \$130,000.* If they did, those lacking managerial talent will also find it pays to go for the MBA, and the pool of MBA's will be "contaminated" by the untalented who are lucky enough to pass.

We have thus far looked at the MBA as a screening device—the firm chose it as a condition of hiring and tied the starting pay to the possession of this degree. But it could also work well as a signaling device, initiated by the candidates. Suppose you, the personnel officer, have not thought of this one. You are hiring at random from the pool at \$50,000 a year, and the firm is suffering some losses from the activities of the untalented hires. Someone could come to you with an MBA, explain how it identifies his or her talent, and say: "Knowing that I am a

good manager raises your expectation of the profit the company will make from my services by a million. I will work for you if you will pay me more than \$75,000 a year.” So long as the facts about the ability of the business school to discriminate managerial talent are clear, this will be an attractive proposition for you.

Even though different players initiate the two strategies of screening and signaling, the same principle underlies them both, namely, the action serves to discriminate between the possible types of players or to indicate the specialized information possessed by one of the players.

Signaling via Bureaucracy

In the United States, the government runs a health insurance system called Workers’ Compensation to cover the treatment of work-related injuries or illnesses. The aims are laudable, but the outcomes have problems. It is difficult for those administering the system to know or judge the severity of an injury (or in some cases even its existence) and the cost of treating it. The workers themselves and the doctors treating them have better information but are also subject to severe temptations to overstate the problems and collect larger sums than are warranted. It has been estimated that 20 percent or more of the claims under Workers’ Compensation involve cheating. According to Stan Long, CEO of Oregon’s state-owned Workers’ Compensation insurer, “If you run a system where you give money to everybody who asks, you are going to get a lot of people asking for money.”⁴

The problem can be tackled to some extent using surveillance. The claimants, or at least those suspected of filing false claims, are watched surreptitiously. If they are found doing things incompatible with their claimed injuries—for example, someone with a claim for a severe back injury is seen lifting heavy loads—their claims are denied, and they are prosecuted.

However, surveillance is costly for the scheme, and our analysis of strategies to elicit information suggests some devices to screen those who are truly injured or ill from the false claimants. For example, the claimants could be required to spend a lot of time filling out forms, sitting all day in a bureaucratic office waiting to talk for five minutes to an official, and so on. Those who are actually healthy and can earn good money working all day will have to forgo those earnings and will therefore find this wait too costly. Those who are truly injured and unable to work will be able to spare the time. People often think of bureaucratic delays and inconveniences as proof of the inefficiency of

government, but they may sometimes be valuable strategies to cope with informational problems.

Benefits in kind have a similar effect. If the government or an insurance company was giving money to the disabled to buy wheelchairs, people might pretend to be disabled. But if it gave wheelchairs directly, the incentive to pretend would be much less, because someone who didn't need a wheelchair would have to make a lot of effort to sell it on the secondhand market and only get a low price for it. Economists usually argue that cash is superior to transfers in kind, because the recipients can make their own optimal decisions to spend cash in the way that best satisfies their preferences, but in the context of asymmetric information, in-kind benefits can be superior because they serve as screening devices.⁵

Signaling by Not Signaling

“Is there any point to which you would wish to draw my attention?”

“To the curious incident of the dog in the nighttime.”

“The dog did nothing in the nighttime.”

“That was the curious incident,” remarked Sherlock Holmes.

In the case of Sherlock Holmes in “Silver Blaze,” the fact that the dog didn't bark meant that the intruder was familiar. In the case where someone doesn't send a signal, that, too, conveys information. Usually it is bad news, but not always.

If the other player knows that you have an opportunity to take an action that will signal something good about yourself, and you fail to take this action, then the other will interpret that as meaning that you do not have that good attribute. You may have innocently overlooked the strategic signaling role of taking or not taking this action, but that will not do you any good.

College students can take many courses for a letter grade (A to F) or on a pass/fail (P or F) basis. Many students think that a P on their transcript will be interpreted as the average passing grade from the letter scale. With grade inflation as it now exists in the United States, this is at least a B+, more likely an A-. Therefore the pass/fail option looks good.

Graduate schools and employers look at transcripts more strategically. They

know that each student has a pretty good estimate of his or her own ability. Those who are so good that they are likely to get an A+ have a strong incentive to signal their ability by taking the course for a letter grade and thereby distinguishing themselves from the average. With many A+ students no longer taking the pass/fail option, the group choosing pass/fail loses much of its upper end. The average grade over this limited pool is no longer an A–, but, say, only a B+. Then those who know they are likely to get an A acquire more of an incentive to distinguish themselves from the herd by taking the course for a letter grade. The pool of pass/fails loses more of its upper end. This process can continue to a point where mostly only those who know they are likely to get a C or worse will choose the pass/fail option. That is how strategic readers of transcripts will interpret a P. Some quite good students who fail to work through this thinking will suffer the consequences of their strategic ignorance.

A friend of ours, John, is brilliant at deal making. He built a worldwide network of classified ad papers through no fewer than 100 acquisitions. When he first sold his company, part of the deal was that he could coinvest with any new acquisition he brought them.* As John explained to the buyer, the fact that he could coinvest would help reassure them that this was a good deal and that they were not overpaying. The buyer understood the reasoning and took it one step further. Did John also understand that if he didn't coinvest, then they would take this as a bad sign and probably wouldn't do the deal? Thus the opportunity to invest would really become a requirement to coinvest. Everything you do sends a signal, including not sending a signal.

Countersignaling

You would think, based on the previous section, that if you have the ability to signal your type, you should. That way, you differentiate yourself from those who can't make the same signal. And yet, some of the people most able to signal refrain from doing so. As Feltovich, Harbaugh, and To explain:

The nouveau riche flaunt their wealth, but the old rich scorn such gauche displays. Minor officials prove their status with petty displays of authority, while the truly powerful show their strength through gestures of magnanimity. People of average education show off the studied regularity of their script, but the well educated often scribble illegibly. Mediocre students answer a teacher's easy questions, but the best students are

embarrassed to prove their knowledge of trivial points. Acquaintances show their good intentions by politely ignoring one's flaws, while close friends show intimacy by teasingly highlighting them. People of moderate ability seek formal credentials to impress employers and society, but the talented often downplay their credentials even if they have bothered to obtain them. A person of average reputation defensively refutes accusations against his character, while a highly respected person finds it demeaning to dignify accusations with a response.⁶

Their insight is that in some circumstances, the best way to signal your ability or type is by not signaling at all, by refusing to play the signaling game. Imagine that there are three types of potential mates: the gold digger, the question mark, and the true love. One partner asks the other to sign a prenuptial with the following argument: I know you say that you love me. Signing the prenup is cheap if you are in this for the love and quite expensive if you are in this relationship for the money.

That is correct. But the partner could well respond: "I know that you can distinguish true loves from gold diggers. It is the question marks that have you confused. You sometimes confuse gold diggers with question marks and other times confuse question marks with true loves. Therefore, if I were to sign the prenup, that would be saying that I felt the need to distinguish myself from the gold diggers. Hence it would be saying that I was a question mark. So I am going to help you realize that I am a true love rather than a question mark by not signing."

Is this really an equilibrium? Imagine that the gold digger and the true love types don't sign and the question marks do sign. As a result, anyone who signs would be viewed as a question mark. This is worse than the position of the true loves. There is no confusion about those who don't sign—the only ones are the gold diggers and true loves, and the partner can tell those apart.

What would happen if the question marks also decided not to sign? Seeing them not sign, their partner would interpret this to mean they must be either a gold digger or a true love. Depending on how likely it is that the question mark will be mistaken for one rather than the other determines whether this would be a good idea or not. If a question mark is more likely to be seen as a gold digger, then not signing is a bad idea.

The larger point is simple. We have ways to figure out people's types besides what they signal. The very fact that they are signaling is a signal that they are trying to differentiate themselves from some other type that can't afford to make

the same signal. In some circumstances, the most powerful signal you can send is that you don't need to signal.*

Sylvia Nasar offers the follow perspective on John Nash: “Fagi Levinson, the [MIT math] department’s den mother, said in 1996: ‘For Nash to deviate from convention is not as shocking as you might think. They were all prima donnas. If a mathematician was mediocre he had to toe the line and be conventional. If he was good, anything went.’”⁷

Prof. Rick Harbaugh, Ph.D., and Ted To did some further investigation into countersignaling. They listened to voicemail messages across the twenty-six University of California and California State University systems, and they found that fewer than 4 percent of economists at schools with a Ph.D. program used a title on their voicemail message, as compared to 27 percent of their colleagues at universities without a doctoral program.⁸ In all cases the faculty had a Ph.D., but reminding the caller of the degree or title suggests that you feel the need of a credential in order to distinguish yourself. The truly impressive faculty could show they were so famous that they didn't need to signal. Hey, just call us Avinash and Barry.

A Quiz: Now you know enough about the manipulation and interpretation of information to take a quiz. We do not call this a Trip to the Gym. It requires no special calculation or math. But we leave it as a quiz instead of offering any discussion of our own, because the correct answers will be highly specific to the situation of each reader. For the same reason, we ask you to grade yourself.

A TRIP TO THE BAR

You are on a first date with someone you find attractive. You want to make a good first impression—you won't get a second chance. But you expect your date to be aware that impressions can be faked, so you must devise credible signals of your quality. At the same time, you want to screen your date, to see if your immediate attraction has a more durable basis and decide whether you want to continue the relationship. Find some good strategies for your signaling and screening.

Signal Jamming

If you are buying a used car from the previous owner, you will want to find out how well he cared for it. You might think that its current condition will serve as a signal, that if the car is washed and polished, and its interior is clean and carpets are vacuumed, it is likely to have been well looked after. However, these are signals that even careless owners can mimic when they offer the car for sale. Most importantly, it costs no more for a careless owner than for a careful owner to get the car cleaned. Therefore the signal does not serve to distinguish between the types. As we saw above in the example of the MBA as a signal of managerial talent, this cost difference is essential if the signal is to be effective in making this distinction.

Actually, some small cost differences do exist. Perhaps those who always take good care of their cars take some pride in the fact and may even enjoy washing, polishing, and cleaning the car. Perhaps the careless are very busy and find it hard to spare the time to do these things or get them done. Can small cost differences between the types suffice for the signal to be effective?

The answer depends on the proportions of the two types in the population. To see why, begin by thinking of how prospective buyers will interpret a car's cleanness or dirtiness. If everyone gets the car cleaned prior to putting it up for sale, then a prospective buyer learns nothing from observing its cleanness. When he sees a clean car, he interprets it as nothing other than a random draw from the population of possible owners. A dirty car would be a sure indicator of a careless owner.

Now suppose the proportion of careless owners in the population is quite small. Then a clean car would convey quite a favorable impression: the buyer will think that the probability of the owner being careful is quite high. He will be more likely to buy the car or to pay a higher price for it. For the sake of this benefit, even the careless owners will clean their cars prior to selling. This situation, where all types (or all people possessing different types of information) take the same action, and therefore the action is completely uninformative, is called a *pooling* equilibrium of the signaling game—the different types end up in the same pool of signals. By contrast, the kind of equilibrium where one type signals and the other does not, so that the action accurately identifies or separates the types, is a *separating* equilibrium.

Next suppose the proportion of careless owners is large. Then if everyone cleans his car, a clean car does not convey a favorable impression, and a careless owner does not find it worth his while to incur the cost of cleaning the car. (The

careful owners always have clean cars.) Thus we cannot get a pooling equilibrium. But if no careless owner is cleaning the car, a single one who does so will get mistaken for a careful owner, and will find it worth his while to incur the small cost. Therefore we cannot get a separating equilibrium either. What happens is somewhere in between: each careless owner follows a mixed strategy, cleaning his car with a positive probability but not certainty. The resulting population of clean cars on the market has a mixture of careful and careless owners. The prospective buyers know the mixture and can infer back to the probability that the owner of a particular clean car is careful. Their willingness to pay will depend on this probability. In turn, the willingness to pay should be such that each careless owner is indifferent between cleaning his car at the small cost and leaving it dirty and thereby being identified as a careless owner, saving the cost but getting a lower price for the car. The mathematical calculation of all this gets somewhat intricate.

It requires a formula, known as Bayes' Rule, for inferring the probabilities of types on the basis of observation of their actions. A simple example of using this rule is illustrated below in the context of betting in poker, but the general features are simple to describe. Because the action now conveys only partial information to distinguish the two types, the outcome is called *semi-separating*.

BODYGUARD OF LIES

Espionage in wartime provides particularly good examples of strategies to confuse the signals of the other side. As Churchill famously said (to Stalin at the 1943 Tehran Conference) "In wartime, truth is so precious that she should always be attended by a bodyguard of lies."

There is a story of two rival businessmen who meet in the Warsaw train station. "Where are you going?" says the first. "To Minsk," replies the other. "To Minsk, eh? What a nerve you have! I know that you are telling me that you are going to Minsk because you want me to believe that you are going to Pinsk. But it so happens that I know you really *are* going to Minsk. So why are you lying to me?"⁹

Some of the best lies arise when someone speaks the truth in order not to be believed. On June 27, 2007, Ashraf Marwan died in London after a suspicious fall from the balcony of his fourth-story flat in Mayfair, London. Thus ended the life of a man who was either the best-connected spy for Israel or a brilliant Egyptian double agent.¹⁰

Ashraf Marwan was the son-in-law of Egyptian President Abdel Nasser and

his liaison to the intelligence service. He offered his services to the Israeli Mossad, who determined his goods were real. Marwan was Israel's guide to the Egyptian mindset.

In April 1973, Marwan sent the code "Radish," which meant that a war was imminent. As a result, Israel called up thousands of reservists and wasted tens of millions on what turned out to be a false alarm. Six months later, Marwan signaled "Radish" again. It was October 5. The warning was that Egypt and Syria would simultaneously attack the next day, on the Yom Kippur holiday, at sunset. This time, Marwan's alarm was no longer trusted. The head of military intelligence thought Marwan was a double agent and took his message as evidence that war was not imminent.

The attack came at 2:00 P.M. and almost overran the Israeli army. General Zeira, Israel's intelligence head, lost his job over the fiasco. Whether Marwan was a spy for Israel or a double agent remains uncertain. And if his death wasn't an accident, we don't know if it was the Israelis or the Egyptians who are to blame.

When playing mixed or random strategies, you can't fool the opposition every time. The best you can hope for is to keep them guessing and fool them some of the time. You can know the likelihood of your success but cannot say in advance whether you will succeed on any particular occasion. In this regard, when you know that you are talking to a person who wants to mislead you, it may be best to ignore any statements he makes rather than accept them at face value or to infer that exactly the opposite must be the truth.

Actions do speak a little louder than words. By seeing what your rival does, you can judge the relative likelihood of matters that he wants to conceal from you. It is clear from our examples that you cannot simply take a rival's statements at face value. But that does not mean that you should ignore what he does when trying to discern where his true interests lie. The right proportions to mix one's equilibrium play depend on one's payoffs. Observing a player's move gives some information about the mix being used and is valuable evidence to help infer the rival's payoffs. Betting strategies in poker provide a prime example.

Poker players are well acquainted with the need to mix their plays. John McDonald gives the following advice: "The poker hand must at all times be concealed behind the mask of inconsistency. The good poker player must avoid set practices and act at random, going so far, on occasion, as to violate the elementary principles of correct play."¹¹ A "tight" player who never bluffs seldom wins a large pot; nobody will ever raise him. He may win many small pots, but invariably ends up a loser. A "loose" player who bluffs too often will

always be called, and thus he too goes down to defeat. The best strategy requires a mix of the two.

Suppose you know that a regular poker rival raises two-thirds of the time and calls one-third of the time when he has a good hand. If he has a poor hand, he folds two-thirds of the time and raises the other third of the time. (In general, it is a bad idea to call when you are bluffing, since you do not expect to have a winning hand.) Then you can construct the following table for the probabilities of his actions.

To avoid possible confusion, we should say that this is not a table of payoffs. The columns do not correspond to the strategies of any player but are the possible workings of chance. The entries in the cells are probabilities, not payoffs.

		Action		
		Raise	Call	Fold
Quality of hand	Good	$2/3$	$1/3$	0
	Poor	$1/3$	0	$2/3$

Suppose that before your rival bids, you believe that good and poor hands are equally likely. Because his mixing probabilities depend on his hand, you get additional information from the bid. If you see him fold, you can be sure he had a poor hand. If he calls, you know his hand is good. But in both these cases, the betting is over. If he raises, the odds are 2:1 that he has a good hand. His bid does not always perfectly reveal his hand, but you know more than when you started. After hearing a raise, you increase the chance that his hand is good from one-half to two-thirds.

The estimation of probabilities conditional on hearing the bid is made using Bayes' Rule. The probability that the other player has a good hand conditional on hearing the bid "X" is the chance that this person would both have a good hand and bid X divided by the chance that he ever bids X. Hearing "fold" implies that his hand must be bad, since a person with a good hand never folds. Hearing "call" implies that his hand must be good, since the only time a player calls is when his hand is good. After hearing "raise," the calculations are only slightly more complicated. The odds that a player both has a good hand and raises is $(1/2)(2/3) = 1/3$, while the chance that the player both has a bad hand and raises—that is, bluffs—is $(1/2)(1/3) = 1/6$. Hence the total chance of hearing

a raise is $1/3 + 1/6 = 1/2$. According to Bayes' Rule, the probability that the hand is good conditional on hearing a raise is the fraction of the total probability of hearing a raise that is due to the times when the player has a strong hand: in this case that fraction is $(1/3)/(1/2) = 2/3$.

PRICE DISCRIMINATION BY SCREENING

The application of the concept of screening that most impinges on your life is price discrimination. For almost any good or service, some people are willing to pay more than others—either because they are richer, more impatient, or just have different tastes. So long as the cost of producing and selling the good to a customer is less than what the customer is willing to pay, the seller would like to serve that customer and get the highest possible price. But that would mean charging different prices to different customers—for example, giving discounts to those who are not willing to pay so much, without giving the same low price to those who would pay more.

That is often difficult. The sellers do not know exactly how much each individual customer is willing to pay. Even if they did, firms would have to try to avoid situations where one customer with a low value buys the item at a low price and then resells it to a high-value customer who was being charged a high price. Here we don't worry about the issue of resale. We focus on the information issue, the fact that firms don't know which customers are which when it comes to who has a high willingness to pay and who doesn't.

To overcome this problem, the trick that sellers commonly use is to create different versions of the same good and price the versions differently. Each customer is free to select any version and pay the price set by the seller for that version, so there is no overt discrimination. But the seller sets the attributes and prices of each version so that different types of customers will choose different versions. These actions implicitly reveal the customers' private information, namely their willingness to pay. The sellers are screening the buyers.

When a new book is published, some people are willing to pay more; these are also likely to be the readers who want to get and read the book immediately, either because they need the information at once or because they want to impress their friends and colleagues with their up-to-date reading. Others are willing to pay less and are content to wait. Publishers take advantage of this inverse relationship between willingness to pay and willingness to wait by publishing the book initially in hardcover at a higher price and then a year or so later issuing a paperback edition at a lower price. The difference in the costs of printing the

two kinds of books is much smaller than the price difference; the “versioning” is just a ploy to screen the buyers. (Question: In what format are you reading this book: hardcover or paperback?)

Producers of computer software often offer a “lite” or “student” version that has fewer features and sells at a substantially lower price. Some users are willing to pay the higher price, perhaps because their employers are the ones paying it. They may also want all the features, or want to have them available just in case they are needed later. Others are willing to pay less and will settle for the basic features. The cost of serving each new customer is very small: just the cost of burning and mailing a CD, or even less in the case of Internet downloads. So the producers would like to cater to those willing to pay less, while charging more to those who are willing to pay more. They do this by offering different versions with different features at different prices. In fact they often produce the lite version by taking the full version and disabling some features. Thus it is somewhat more costly to produce the lite version, even though its price is lower. This seemingly paradoxical situation has to be understood in terms of its purpose, namely to allow the producers to practice price discrimination by screening.

IBM offered two versions of its laser printer. The E version printed at 5 pages per minute, while for \$200 more you could get the fast version that printed at 10 pages per minute. The only difference between the two was that IBM added a chip in the firmware of the E version that added some wait states to slow down the printing.¹² If they hadn’t done this, then they would have had to sell all their printers at one price. But with the slowed down version, they could offer a lower price to home users who were willing to wait longer for their printouts.

The Sharp DVE611 DVD player and their DV740U unit were both made in the same Shanghai plant. The key difference was that the DVE611 lacked the ability to play DVDs formatted to the European standard (called PAL) on television sets that use the American standard (called NTSC). However, it turns out that the functionality was there all along, just hidden from the customer. Sharp had shaved down the system switch button and then covered it with the remote control faceplate. There were some ingenious users who figured this out and shared their discovery on the web. You could restore full functionality simply by punching a hole in the faceplate at the appropriate spot.¹³ Companies often go through great effort to create damaged versions of their goods, and customers often go to great lengths to restore the product.

Airline pricing is probably the example of price discrimination most familiar to readers, so we develop it a little further to give you an idea of the quantitative

aspects of designing such a scheme. For this purpose, we introduce Pie-In-The-Sky (PITS), an airline running a service from Podunk to South Succotash. It carries some business passengers and some tourists; the former type is willing to pay a higher price than the latter. To serve the tourists profitably without giving the same low price to the business travelers, PITS has to develop a way of creating different versions of the same flight and price the versions in such a way that each type will choose a different version. First class and economy class might be one way to do this, and we will take that as our example; another common distinction is that between unrestricted and restricted fares.

Suppose that 30 percent of the customers are businesspeople and 70 percent are tourists; we will do the calculation on the basis of “per 100 customers.” The table shows the maximum price each type is willing to pay for each class of service (technically referred to as the *reservation price*) and the costs of providing the two types of service.

Type of service	PITS's cost	Reservation price		PITS's potential profit	
		Tourist	Business	Tourist	Business
Economy	100	140	225	40	125
First	150	175	300	25	150

Begin by setting up a situation that is ideal from PITS's point of view. Suppose it knows the type of each customer, for example, by observing their dress as they come to make their reservations. Also suppose that there are no legal prohibitions or resale possibilities. Then PITS can practice what is called perfect price discrimination. To each businessperson it could sell a first-class ticket at \$300 for a profit of $\$300 - \$150 = \$150$, or an economy ticket at \$225, for a profit of $\$225 - \$100 = \$125$. The former is better for PITS. To each tourist, it could sell a first-class ticket at \$175 for a profit of $\$175 - \$150 = \$25$, or an economy ticket at \$140 for a profit of $\$140 - \$100 = \$40$; the latter is better for PITS. Ideally, PITS would like to sell only first-class tickets to business travelers and only economy-class tickets to tourists, in each case at a price equal to the maximum willingness to pay. PITS's total profit per 100 customers from this strategy will be

$$(\$140 - \$100) \times 70 + (\$300 - \$150) \times 30 = \$40 \times 70 + \$150 \times 30 = \$2800 + \$4500 = \$7300.$$

Now turn to the more realistic scenario where PITS cannot identify the type of each customer, or is not allowed to use the information for purposes of overt discrimination. How can it use the versions to screen the customers?

Most importantly, it cannot charge the business travelers their full willingness to pay for first-class seats. They could buy economy-class seats for \$140 when they are willing to pay \$225; doing so would give them an extra benefit, or “consumer surplus” in the jargon of economics, of \$85. They might use it, for example, for better food or accommodation on their trip. Paying the maximum \$300 that they are willing to pay for a first-class seat would give them no consumer surplus. Therefore they would switch to economy class, and screening would fail.

The maximum that PITS can charge for first class must give business travelers at least as much extra benefit as the \$85 they can get if they buy an economy-class ticket, so the price of first-class tickets can be at most $\$300 - \$85 = \$215$. (Perhaps it should be \$214 to create a definite positive reason for business travelers to choose first class, but we will ignore the trivial difference.) PITS’s profit will be

$$(140 - 100) \times 70 + (215 - 150) \times 30 = 40 \times 70 + 65 \times 30 = 2800 + 1950 = 4750.$$

So, as we see, PITS can successfully screen and separate the two types of travelers based on their self-selection of the two types of services. But PITS must sacrifice some profit to achieve this indirect discrimination. It must charge the business travelers less than their full willingness to pay. As a result, PITS’s profit per 100 passengers drops from the \$7,300 it could achieve if it could discriminate overtly with direct knowledge of each customer’s type to the \$4,750 it achieves from the indirect discrimination based on self-selection. The difference, \$2,550, is precisely 85 times 30, where 85 is the drop in the first-class fare below the business travelers’ full willingness to pay for this service and 30 is the number of business travelers.

PITS has to keep the first-class fare sufficiently low to give the business travelers enough incentive to choose this service and not “defect” to making the choice that PITS intends for the tourists. Such a requirement, or constraint, on the screener’s strategy is called an *incentive compatibility constraint*.

The only way PITS could charge more than \$215 to business travelers without inducing their defection would be to increase the economy-class fare.

For example, if the first-class fare is \$240 and the economy-class fare is \$165, then business travelers get equal extra benefit (consumer surplus) from the two classes: \$300–240 from first class and \$225–165 from economy class, or \$60 from each, so they are (only just) willing to buy first-class tickets.

TRIP TO THE GYM NO. 5

There is also a participation constraint for business travelers and an incentive compatibility constraint for the tourists. Check that these are automatically satisfied at the stated prices.

But at \$140 the economy-class fare is already at the limit of the tourists' willingness to pay. If PITS raised it to even \$141, it would lose these customers altogether. This requirement, namely that the customer type in question remains willing to buy, is called that type's *participation constraint*. PITS's pricing strategy is thus squeezed between the participation constraint of the tourists and the incentive compatibility constraint of the businesspeople. In this situation, the screening strategy above, charging \$215 for first class and \$140 for economy class, is in fact the most profitable for PITS. It takes a little mathematics to prove that rigorously, so we merely assert it.

Whether this strategy is optimal for PITS depends on the specific numbers in the example. Suppose the proportion of business travelers were much higher, say 50 percent. Then the sacrifice of \$85 on each business traveler may be too high to justify keeping the few tourists. PITS may do better not to serve them at all—that is, violate their participation constraint—and raise the price of first-class service for the business travelers. Indeed, the strategy of discrimination by screening with these numbers of travelers yields

$$(140-100) \times 50 + (215-150) \times 50 = 40 \times 50 + 65 \times 50 = 2000 + 3250 = 5250,$$

while the strategy of serving only business travelers in first class at \$300 would yield

$$(300-150) \times 50 = 150 \times 50 = 7500.$$

If there are only a few customers with low willingness to pay, the seller might find it better not to serve them at all than to offer sufficiently low prices to the mass of high-paying customers to prevent their switching to the low-priced version.

Now that you know what to look for, you will see screening for price discrimination everywhere. And if you look in the research literature, you will see analyses of strategies for screening by self-selection equally frequently.¹⁴ Some of these strategies are quite complicated, and the theories need a lot of mathematics. But the basic idea driving all these instances is the interplay between the twin requirements of incentive compatibility and participation.

CASE STUDY: GOING UNDERCOVER

Another friend of ours, Tanya, is an anthropologist. While most anthropologists travel to the ends of the earth to study some unusual tribe, Tanya did her fieldwork in London. Her subject was witches.

Yes, witches. Even in modern-day London there are still a surprisingly large number of people who gather together to trade spells and study witchcraft. Not that being a modern witch is easy; it requires a certain amount of rationalization to be a witch riding the tube. Often anthropologists have trouble gaining their subject's confidence. But Tanya's group was especially welcoming. When she told them she was an anthropologist, they saw this as a clever ruse: she was really a witch with a great cover story.

One of the unusual features of the witches' meetings is that they took place in the nude. Why might that be?

Case Discussion

Any outsider group has to worry that its members will be observers rather than participants. Are you sitting there making fun of the whole process, or are you being a part of it? If you are sitting there in the nude, it is pretty hard to say that you are just watching and making fun of the others. You are well into it.

Thus the nudity is a credible screening device. If you truly believe in the coven, then it is relatively costless to be there in the nude. But if you are a skeptic, then being there in the nude is hard to explain, both to others and to

yourself.* For the same reason, gang initiation rites often involve taking actions that are relatively cheap if you are truly interested in gang life (tattoos, committing crimes) but quite costly if you are an undercover cop trying to infiltrate the gang.

For more cases on interpreting and manipulating information, see “The Other Person’s Envelope Is Always Greener,” “But One Life to Lay Down for Your Country,” “King Solomon’s Dilemma Redux,” and “The King Lear Problem” in chapter 14.