FIVE

The Inside Story

Jerome K. Jerome's nineteenth-century comic travelogue *Three Men in a Boat* begins with Mr. Jerome flipping through a medical dictionary in the British Museum:

> I came to typhoid fever—read the symptoms—discovered that I had typhoid fever, must have had it for months without knowing it—wondered what else I had got; turned up St Vitus's Dance—found, as I expected, that I had that too—began to get interested in my case, and determined to sift it to the bottom, and so started alphabetically—read up ague, and learnt that I was sickening for it, and that the acute stage would commence in about another fortnight. Bright's disease, I was relieved to find, I had only in a modified form, and, so far as that was concerned, I might live for years. Cholera I had, with severe complications; and diphtheria I seemed to have been born with. I plodded conscientiously though the twenty-six letters, and the only malady I could conclude I had not got was housemaid's knee.

If you were Mr. Jerome, what would you do? He decided to go for a trip up the river, but then he was not an economist. My advice in such a situation would be to pick up the phone and buy yourself some really generous medical insurance. After all, since

you know for sure that you will be making expensive claims, why not pay the premium for the best possible care?

Yet this raises a question. If people like Mr. Jerome could rush to buy health insurance if and only if they knew they were ill, who would want to insure them?

Inside information

This is more than an idle question. Economists have known for a while that if one party to a deal has inside information and the other does not, then markets may not work as well as we would hope. That makes intuitive sense. But it wasn't until an American economist named George Akerlof published a revolutionary paper in 1970 that the profession realized quite how profound and dramatic the problem might be.

Akerlof chose as his example the market for used cars and showed that even if the market is highly competitive, it simply cannot work if sellers know a lot about the quality of their cars and buyers do not. To take a stark example, let's say that half the used cars on sale are "peaches," and half are "lemons." The peaches are worth more to prospective buyers than to sellers—otherwise the buyers wouldn't be buyers—say, $5,000 to prospective buyers and $4,000 to sellers. The lemons are worthless pieces of junk. Sellers know whether the car they're selling is a lemon or a peach. Buyers have to guess.

A buyer who doesn't mind taking a fair gamble might think that anything between $2,000 and $2,500 would be a reasonable price for a car that has a 50-percent chance of being a peach, and a 50-percent chance of being a lemon. The seller would also think this was a fair deal for a 50/50 prospect, but the seller doesn't face a 50/50 prospect: the seller knows for certain whether the car is a peach or a lemon. The problem is that a seller with a lemon would bite your hand off if offered $2,500, but a seller with a peach would find it rather insulting. Wandering around offering $2,500 for a car, you quickly discover that only lemons are for sale at that price. Of course, if you offered $4,001 you

would also see the peaches on the market—but the lemons won't go away, and $4,001 is not an attractive price for a car that only has a 50-percent chance of running properly.

This isn't just a trivial problem around the fringes of the market. In this scenario, there *is no market*. No seller is willing to sell a peach for less than $4,000, but no buyer is willing to offer that much on a car that has a 50-percent chance of being a lemon. Sellers don't offer peaches for sale, buyers know that they won't, and in the end the only cars that are traded are worthless lemons, which get passed around for next to nothing. Less extreme assumptions about the problem lead to less extreme breakdowns of the market, but the conclusions are similar: if some people know more than others about the quality of a product, then some high-quality products may not be traded at all, or not be traded very much.

Anyone who has ever tried to buy a secondhand car will appreciate that Akerlof was on to something. The market doesn't work nearly as well as it should; secondhand cars tend to be cheap and of poor quality. Sellers with good cars want to hold out for a good price, but because they cannot *prove* that a good car really is a peach, they cannot get that price and prefer to keep the car for themselves. You might expect that the sellers would benefit from their inside information, but in fact there are no winners: smart buyers simply don't show up to play a rigged game.

Let's be clear about how dramatic and how worrying this problem is. What Akerlof described is not a market where some people get ripped off; it's much more serious than that. He described a market that should exist and simply doesn't because of the corrosive force of inside information. Sellers with good cars should be trading with buyers—each trade produces $1,000 of value, because that's the difference between the value to the seller and the value to the buyer. If the price is close to $4,000, the buyer gets more of that value, and if it is close to $5,000, the seller gets more. But Akerlof showed that none of those value-creating trades happen because the buyers will not buy without proof, and the sellers cannot offer proof.

The market for used cars is not the only one affected by inside information. Think about furniture in a rented apartment—why is it never built to last? Akerlof's model provides an answer. Apartments have many noticeable, even obvious, attributes that can influence our decision whether or not to rent—size, location, interior design, and so on. But there are also qualities that are difficult to observe—for instance, whether the furniture is durable. The landlord has little incentive to provide expensive, hard-wearing furniture, since this is not one of the features that potential tenants can recognize before they move in, and so not one for which they are willing to pay. (Of course, the landlord might also go for cheap, fragile furniture because he expects the tenants to trash whatever he puts in the flat. But that fear is equally an argument to get in more durable stuff.) As a result, there is a market for rental apartments with flimsy furniture but not a market for rental apartments with durable furniture.

Inside information also means that you can't get a decent meal in a tourist trap like London's Leicester Square, Times Square in Manhattan, or the Plaka in Athens. With few exceptions, the hungry visitor will pay a lot for mediocre cuisine. Tourists are willing to pay high prices because they have no sense of where better alternatives, even just a few streets away, might be found. But the tourist-trap phenomenon is not just about high prices. If it was, we would see a wide range of restaurants, charming little bistros, and downscale pasta or burger joints, all kinds of food from superb to disastrous, all charging a premium. Instead, we see a truncated market—high-quality places, whether the good food is fried chicken or fine dining, are simply not to be found. The reason is simple enough; tourists will only be making a single visit and will find it hard to pick out the great food from the bad. Good restaurants all locate where they are more likely to be appreciated by more informed locals. The bad ones remain . . . the "lemons" of the restaurant trade.

It's worth emphasizing that Akerlof is not describing universal ignorance but a situation where one side knows more than the

other. If buyers and sellers were both ignorant about whether a car was a lemon or a peach there would be no problem: buyers would be willing to pay up to $2,500 for a car that had a 50/50 chance of being a peach; sellers, equally ignorant, would be willing to accept any offer over $2,000. Of course they will strike a deal. It's only when one negotiator knows too much and the other too little that agreement becomes impossible. Because the problem is caused by an uneven grasp on the facts, economists tend to call it "asymmetric information." Ripping the "world of truth" apart, this imbalance in information can completely destroy perfect markets.

Inside information and health insurance

Akerlof's "lemons" problem would be troublesome enough if it applied only to secondhand cars, furnished accommodations, and dubious restaurants in the world's beauty spots. Unfortunately, it also damages the market for more important goods—most notably, health insurance.

Health insurance is important because illnesses are extremely unpredictable and sometimes cost a lot to treat. Not only can some medical treatment be very expensive, it is often impossible to postpone it until a more convenient moment. It can also coincide with periods of low income because people are more likely to need medical care after retirement, and because those who need medical care may also be too ill to work.

So health insurance is a valuable product. If the health insurance market doesn't work well, the results will be excessively high premiums and a large number of uninsured people. This will sound very familiar to readers in the United States, where markets do not, in fact, do a good job of providing medical insurance, precisely because of Akerlof's "lemons" problem.

Let's say that people who are likely prone to sickness are "lemons"; people who are likely to stay healthy are "peaches." If, like Jerome K. Jerome, I suspect myself to be a lemon, I'd be well

advised to buy all the medical insurance going. On the other hand, if you feel fine and all your ancestors lived to be a hundred, then perhaps you will buy medical insurance only if it is extremely cheap. After all, you hardly expect to need it.

Thanks to Akerlof's proof that markets whose players have asymmetric information are doomed, we know that the insurance market may disappear just as the market for good-quality used cars did. You, whose body is a succulent peach, will not find the typical insurance package a good deal; while Mr. Jerome and I, whose bodies are bitter little lemons, will embrace the typical insurance package with open arms. The result is that the insurance company only sells insurance to people who are confident they will use it. As a result, the insurer loses clients who are unlikely to make claims and acquires the unwanted clients who are likely to make costly claims, and then the insurer has to cut back on benefits and raise premiums. People of middling health now find the insurance is too expensive and cancel it, forcing the insurance company to raise premiums even higher to stay in business. More and more people cancel their policies, and in the end only the most sickly of the lemons will buy insurance and at a price that will be nearly impossible to afford.

The insurance companies will, of course, try to repair the insurance market by finding out more information about their customers. Do they smoke? How old are they? Did their parents die of hereditary diseases at the age of thirty-five, or in sports car accidents at the age of a hundred? With more genetic information becoming available, insurance companies will be able to gain an increasingly accurate picture of the costs of providing medical insurance to particular individuals. Previously the insurance market was constrained by the presence of inside information: insurers knew less than those they were insuring. But if insurance companies can continue to close that information gap, they will be willing to provide insurance to more people.

This might sound like the price targeting used by Starbucks and Wholefoods in chapter 2, but in fact it's a different game.

When Starbucks tries to price-target, it knows its own costs and is simply trying to find out whether it can get away with a higher price for some customers. The health insurance companies face a more fundamental task: they don't know how much it will cost to cover the claims of each customer, and if they can't work this out with more accuracy than the customer can, they will simply go bust under the burden of claims. The effects are different, too: price-targeting is a way of getting a larger slice of the pie by squeezing more money from customers, while finding out about insurance customers can create a new pie by making trades possible when previously there were none.

Unfortunately, the insurance market would be saved at some cost: what we would find is that lemons, like Mr. Jerome and me, would be able to buy insurance only at enormous rates. Peaches like you could pick it up for a nominal sum. Both premiums would reflect an actuarially fair rate, which means a rate that covers no more and no less than the likely medical expenses. If companies had really accurate information, perhaps obtainable from the genetic tests of the future, then someone who was likely to get ill would pay hundreds of thousands of dollars in premiums; but this would hardly be insurance at all.

By assessing our individual backgrounds and predicting the cost of providing benefits to each of us, the insurance industry manages to stay afloat; if companies did not raise prices for lemons like me and Mr. Jerome, they would soon go out of business. The problem is that people who expect to have expensive medical needs—the elderly and the chronically ill, for example—will find that their insurance company does not really give them much insurance at all. Because their premiums are adjusted to take these expenses into account, they will pay more for insurance than they would for the out-of-pocket medical costs they would face without insurance.

The curious conclusion, which is obvious in retrospect, is that an insurance policy depends on mutual ignorance. An insurance company can only insure me against an event like a burglary, a fire, or a medical bill if neither of us has any idea whether it will

happen. If we could predict the future, insurance would be meaningless. If my insurance company could predict fires much better than I could, it would sell me insurance only if I didn't need it. And if I knew that my house would burn down, the insurance company should be calling the police rather than selling me fire insurance. Since insurance depends on mutual ignorance, then any advance in medical science, which pushes back the boundaries of ignorance—whether for the insurers, the insured, or both, will weaken the basis of insurance. The more we know, the less we can insure. This is a worrying prospect if we want to give people a chance to protect themselves against the high costs of bad luck.

Making lemonade

A slightly vexing aphorism recommends, "If life deals you lemons, make lemonade." How can we make lemonade out of Akerlof's lemons? To return to Akerlof's original example of the market for secondhand cars, both buyers and sellers have an incentive to try to fix the problem: sellers want to get a decent price for their peaches, and buyers want to buy peaches. If inside information is wrecking the chance of a mutually beneficial deal, both sides will want to find a way to bridge the information gap.

Akerlof won the Nobel Prize in 2001 for his work on the problem of asymmetric information; he shared it with two economists who proposed partial solutions. Michael Spence argued that the person with the information might be able to communicate it in a way that the person without the information could trust. Joe Stiglitz looked at the problem in reverse and explored ways in which the person without the information might uncover it.

Spence realized that it wasn't enough for a seller of peaches simply to say, "All my cars are peaches," because talk is cheap. A seller of lemons can also say, "All my cars are peaches." The buyer wouldn't know who was telling the truth, so the claim itself doesn't carry any information. Spence realized that a real signal of quality would be one that a lemon-seller could not make, or at least, could not afford to make.

An example would be buying an expensive car showroom, an investment affordable only by a businessman who plans to stick around for the long term. A peach-seller expects satisfied customers to return, and to tell their friends about his reliable, trustworthy cars. Over the years the sales would pay for the showroom. A lemon-seller couldn't operate like that; instead, he would sell a few overhyped lemons and then have to move someplace where his reputation for dishonesty could not follow.

It's for this reason that banks always used to build such impressive buildings. In the days before government oversight, who knew whether they were depositing their money with a fly-by-night operation? Customers realize that crooks planning to run off with the money do not first clad their branches with bronze and marble. This is one reason, too, why you will pay more at an established store than at a market stall if you buy a product about which you lack inside information about quality and durability. The established store will still be there to refund your money in the case of a complaint, and that very possibility gives you an assurance that a complaint is less likely to be necessary.

Other economists have used Spence's theory to explain enormously expensive advertising campaigns with no informational content. What, after all, is the information contained within a soft-drink advertisement? "Coca-Cola. Real." Pardon? The only information that potential customers can glean from such an advertisement is that it was expensive to make, and that therefore the Coca-Cola company plans to stick around with the same commitment to high-quality products that it always had.

Spence himself first used his insight to show why students might choose to pursue a degree in philosophy, which is difficult but does not lead to specific career opportunities, like an economics degree or a marketing degree. Assume that employers would like to hire smart, diligent workers but can't tell from an interview who is smart or diligent. Assume also that everyone has to work hard to obtain a philosophy degree, but lazy, dumb people find it particularly troublesome.

Spence then shows that smart, diligent people can prove they're smart and diligent by going to the trouble of getting a philosophy degree. It's not that lazy, dumb people can't get that degree but that they wouldn't want to: employers will pay philosophy graduates enough to compensate them for the trouble but not enough to persuade lazy, dumb people to bother. The employers are willing to do this despite the fact that the philosophy degree itself does not improve the candidate's productivity at all. It is merely a credible signal, because a philosophy degree is too much trouble for lazy, dumb people to acquire. Since Spence himself majored in philosophy at Princeton, perhaps there is something in his idea.

Spence proved that one way to bridge the information gap in markets previously hindered by inside information is for trustworthy vendors to find ways to signal their reliability. High-quality job applicants, banks, used-car salesmen, and soft-drink manufacturers may find it worthwhile to spend excessive amounts of time and money (by pursuing a degree that does not really add to one's qualifications, paying for lavish decorations, buildings, and advertising) simply to distinguish themselves from low-quality job applicants, banks, used-car salesmen, or soft-drink manufacturers.

Spence's ideas suggest that the lemons problem is not insoluble, but they are not particularly reassuring. In some variations of Spence's model, everyone would be better off if the wasteful signal were impossible. If studying philosophy was banned, employers would be unable to distinguish between lazy workers and smart workers and would pay both the same wages, an average of their expected productivity. The lazy workers would be better off; the smart workers might also be better off if the new, lower, wage is more than the old wage minus the cost of getting a philosophy degree. And the employers don't mind; they employ worse workers on average but also have to pay less for them. Akerlof showed that inside information could reduce people's ability to find trades that made both sides better off. Spence showed a way of making those trades possible but also found that the social cost of doing so can be too high.

While Spence asked what the informed side could do to credibly signal information, Stiglitz studied what the uninformed side could do to uncover it. He explicitly considered markets for insurance and concluded that the uninformed insurer was not completely helpless in the face of customers who could predict the likelihood of needing to file insurance claims. The insurer could offer different deals, for instance, reducing the premium but increasing the deductible. This has the effect of reducing the level of insurance: the low premium makes the insurance cheaper, but the higher deductible, which is the amount by which any claim is reduced, means that any claim would pay out less. Low-risk customers would be attracted by that kind of deal, because the insurance is cheaper and they don't expect to claim very often anyway, but high-risk customers would rather pay the higher premium because they expect to claim frequently, so a high deductible will cost them a lot. So insurers could persuade different types of customer to reveal their inside information. This is a little like the self-targeting strategy used by coffee bars in chapter 2. Starbucks offers frills like whipped cream and flavored syrup to persuade customers to reveal whether or not they are price conscious. Aetna Insurance offers four different packages to individuals, with deductibles ranging from $500 to $5,000, to entice policy buyers to reveal their predictions for how many insurance claims they will file. But again, Stiglitz did not conclude that Akerlof's lemons problem could be costlessly solved. On the contrary, he showed that in response to inside information, banks might deny loans to whole sections of society, firms might prefer to pay high wages to privileged insiders rather than lower wages to more workers, and insurance companies would prefer to exclude high-risk individuals. Spence and Stiglitz both showed that you can make lemonade from Akerlof's lemons—but that you cannot get rid of the bitter aftertaste.

Lemons, health care, and the United States

The difficulty of solving the lemons problem may explain why the American health-care system is malfunctioning so badly. The

United States relies upon private health insurance to provide much of the financing for medical costs. This is unusual: in Britain, Canada, and Spain, for example, health-care costs are largely paid for by the government. In Austria, Belgium, France, Germany, and the Netherlands, medical costs are paid for by a system of "social insurance": it is compulsory for most people to buy insurance, but insurance premiums are tied by law to income rather than to the risk of a claim.

The United States system makes it voluntary to buy insurance, and premiums are linked to risk, not to income. But these market-based principles, beloved of many Americans, do not seem to be delivering health care that makes them happy. A recent survey revealed that only 17 percent of respondents in the United States were content with their health-care system and thought no substantial reforms were necessary. Why the discontent?

The superficial reasons are simple enough to describe: the system is hugely expensive, very bureaucratic, and extremely patchy. The expense first: US health care costs a third more, per person, than that of the closest rival, super-rich Switzerland, and twice what many European countries spend. The United States government alone spends more per person than the combination of public and private expenditure in Britain, despite the fact that the British government provides free health care for all residents, while the American government spending program covers only the elderly (Medicare) and some of the marginalized (Medicaid). Most Americans worry about health-care costs and would be stunned to discover that the British government spends less per person than the American government but still manages to provide free health care for everyone. In fact, if you figure in the costs of providing health insurance to government employees and providing tax breaks to encourage private health care, the US government spending on health care, per person, is the highest in the world.

Bureaucracy next. Researchers at the Harvard Medical School found that the administrative costs of the US system, public and private, exceed $1,000 per person. In other words, when you count

all the taxes, premiums, and out-of-pocket expenses, the typical American spends as much on doctor's receptionists and the like as citizens of Singapore and the Czech Republic spend on their entire medical care. Both places are countries with health outcomes very similar to those in the United States: life expectancy and "healthy life" expectancy (a statistic that distinguishes a long healthy life and a long life plagued by years of severe disability) are a shade lower in the Czech Republic than in the United States; and in Singapore they are a little higher than in the United States. The cost of US bureaucracy is also more than three times the $307 cost per person for the administration of the Canadian health system, which produces noticeably superior health outcomes.

Then there is the patchy coverage of the system. Health insurance is usually packaged together with a job, which reduces the efficiency of the labor market; workers are hesitant to quit their jobs without lining another job up first for fear of being uninsured. Worse, 15 percent of citizens have no insurance coverage of any kind—which should be a stunning statistic for the world's richest economy, but probably isn't because it has been lamented for so many years. Compare it to Germany, where 0.2 percent of the population has no coverage, or to Canada or Britain, where everyone is provided for by the government.

Given what we have learned from George Akerlof and his lemons, the troubles of the US health-care system should be no surprise. We should expect a voluntary private insurance system to be patchy. A few people who have more pressing costs than health insurance (for example, the young poor, who have little money and rightly expect that they are unlikely to become seriously ill) will drop out of the system. As a result, health insurance companies, needing to cover their costs, will raise the premiums for the average client, driving out more and more people. Unlike the very stark lemons model, the market does not completely collapse; this is partly because many people find that the risks of having to pay for medical treatment are so worrying that they're willing to pay substantially more than an actuarially fair premium.

The US spends most on medical care

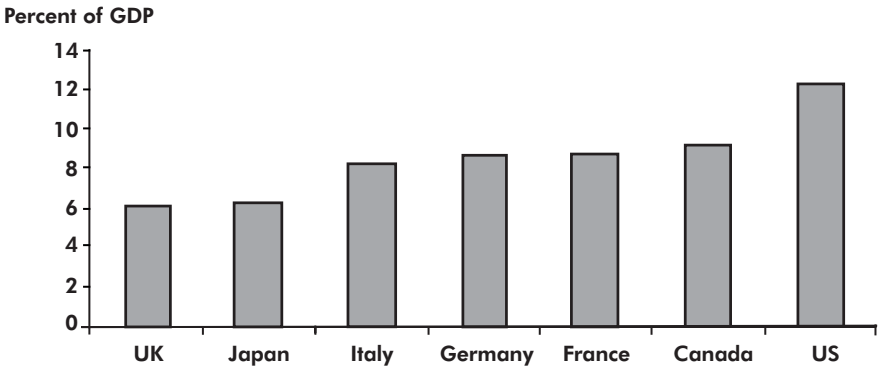

Source: Cutler 2002.

As a result, the process of unraveling stops, but not before many people have been excluded from the system.

Thanks to Spence and Stiglitz, we should also expect insurance companies to devise ways to get around this lemons problem, but that although the solutions may be effective they will probably also be wasteful. The huge bureaucratic burden of the US system is one of the results, as insurance companies struggle to monitor the risks, behavior, and expenses of their customers. The clunky linkage of health insurance with jobs is another result: at first sight, there is no reason why a job should come with health insurance, any more than it should come with a house or free food. Employees are frequently forced to buy the health insurance that is packaged with their job. This packaging compels the healthiest members of society to buy insurance packages and so helps prevent the unraveling of the market. But this solution doesn't come cheap: health-care plans are not chosen by their beneficiaries, who would aim to get the ideal coverage for the right price, but by human resource managers with other priorities, such as making their own lives easy with a "one-size-fits-all" bulk purchase. The result is likely to be further wasteful spending.

Not every drawback of the US health-care system should be blamed on Akerlof's lemons problem. Even without the diffi-

culty of inside information, the system of insurance is problematic, because patients are not always able to choose their treatment. With the insurance company picking up the bill, choosing the appropriate treatment is always going to be something of a matter of negotiation. When you ask somebody else to pay for your health care, don't be surprised if you don't get exactly what you would have chosen yourself.

Nevertheless, it is striking that partial coverage, inefficiency, and high costs are not only the defining characteristics of private health insurance, they are also exactly what we would have predicted armed only with the theoretical models of Akerlof, Spence, and Stiglitz.

Imperfect information—the whole story

The lemons problem ("adverse selection" in the economists' jargon), when inside information guts a market because ignorant buyers are unwilling to pay for quality they cannot observe, is one example of the broader problem of inside information ("asymmetric information" in the jargon). Inside information also produces an obstacle called "moral hazard." The concept is simple: if you compensate people when bad things happen to them, they may get careless.

If my car is insured against theft, I will park wherever I find a space, even on a deserted street that doesn't seem entirely safe. If my insurance doesn't cover theft I may choose to pay a little extra to park in a lot with an attendant. If I lose my job and the government pays unemployment benefit, I may not hurry to find a new job quite as quickly as I would if I had no income whatsoever. If the money in my checking account is insured against a bank failure, why bother to check that the bank is financially sound?

Moral hazard is an inevitable problem in the real economy. While it is impossible for insurance companies (or anyone else)

to avoid moral hazard altogether, they can take steps to reduce it. For example, they do not offer insurance against being fired or becoming pregnant, which is a shame, because it would be great to have that kind of insurance. The reason is easy to see: it is easy to arrange to be fired or to get pregnant. There are many people who would like to leave their jobs and many others who would like to have children, and such people would be particularly eager to buy an insurance policy that would pay them handsomely for putting their plans into action. As a result, moral hazard destroys the market for private unemployment insurance.

On the other hand, public unemployment insurance still exists, in spite of moral hazard. It is not polite to say so, but it is obvious that paying people to be unemployed encourages unemployment. Yet if a government scrapped unemployment benefit, there would still be jobless people, and supporting the jobless is something that every civilized society should do. The truth is that we have a trade-off: it is bad to encourage unemployment but good to support those without incomes.

Both governments and private insurers will try to protect themselves against moral hazard. One of the most common ways insurance companies do this is by modifying the insurance policy to provide incomplete insurance, in the form of a deductible. If my car insurance deductible is $200, the fear of losing that money probably won't persuade me to take extravagant safety precautions, but it should be enough to make me check that my car is locked.

Another way insurance companies can fight moral hazard is by gaining access to the inside information. Health insurers will want to know whether I am a smoker before they set my premiums. Of course, I could lie, but it wouldn't be too hard for insurers to expose my lie; a simple medical examination would reveal that I smoke. When most governments pay unemployment benefits, they do so on the condition that recipients are actively looking for work. Because the government cannot monitor people's job searches perfectly, it pays out only meager unemployment benefits. Yet if the government could really tell how hard unem-

ployed people were looking for jobs, then it would be possible to pay more generous benefits to genuinely deserving recipients.

The problems of imperfect information include adverse selection (lemons) and moral hazard, but there are other, broader, and vaguer issues. For example, my boss would like to pay me extra if I try harder to do a good job, but because he has only a vague idea how hard I am trying, my performance bonus is only a small part of my salary. If my boss could observe my skill and effort perfectly, he could make my entire salary performance-related. Another example: let's say I would like to eat at the best-value restaurant in town; I don't know what it is, so I look for a familiar brand name, where I know I can't go wrong. Knowing that customers won't bother trying to find the cheapest place around, established restaurants are able to charge more than they should.

Do these information problems destroy markets completely? They certainly don't help, but it would be wrong to exaggerate the problems. In spite of asymmetric information, markets do often work well, because people produce ingenious solutions to improve the quality of information—or to reduce the damage caused by imperfect information.

When I buy complicated equipment like a camera, I talk to friends and consult websites and consumer magazines, which I hope will give me useful information about the products I am trying to choose between. Expert reviews providing "inside information" can be particularly helpful when we are ignorant about what we're buying. I rely on them all the time in another market, which suffers from severe information problems: the market for vacations. I'm the kind of person who likes to visit new places, but often I have no idea where to go or what will be fun, where is tacky and who offers a good deal, where is beautiful and where is dangerous. If the problem were insoluble we wouldn't bother to take vacations at all. (Or we might demand that governments provided them, which in my mind conjures up images of being on a waiting list for years before enjoying organized team games

and artificial cheer in an overcast concrete resort.) Instead, we simply buy decent guidebooks and try to learn more on our own.

Health-care provision, yet again, provides a particularly acute example of the problem. It's one thing to check a guidebook to find out what to do on vacation. It's quite another to consult a guidebook to find a heart surgeon. Yet the basic information problem is the same as the one faced by vacationers. Heart surgery patients try to learn more about which doctors have good reputations, which procedures have the highest success rate, and which hospitals provide the best recovery care. Still, most patients would admit that they really don't have much idea just how good their doctor really is.

Market failure versus government failure

This is not exactly comforting. A health-care system based on private insurance will be, as we have discovered, patchy, costly, and bureaucratic. What's more, it will offer patients choices, such as the choice of heart surgeon, that they are not very well qualified to make. So could the government do better? After all, every chapter of this book so far—with the exception of chapter 3—has lamented the causes and costs of market failures. It's tempting to look to the government to sort things out.

Unfortunately, while markets can fail, governments can fail too. Politicians and bureaucrats have their own motivations. Scarcity power, externalities, and imperfect information do not magically disappear when the economy is run or regulated by governments. When market failure and government failure are both present, the choice is often between the lesser of two evils.

An intriguing case in point is Britain's National Health Service (NHS), which offers health care to all citizens. It is almost completely free, although people with jobs need to pay a token amount for prescription drugs. It provides universal coverage: if you walk into any doctor's surgery or any hospital in the country, you will be treated free of charge.

As you would expect, the system gets overcrowded, people often have to wait, and patient choice is not a major feature of the system: you accept whatever treatment the doctor says is appropriate, or nothing. Overall, the medical outcomes are not bad, but the waiting lines for treatment have been a major bone of contention for many years. The same survey, which found 17 percent of American citizens approved of the US health-care system, reports that only 25 percent of the British are happy with their own system—better, but hardly a resounding endorsement.

If you were going blind in Britain, you would be well aware of a recent example of the difficulties faced by such a system. The Royal National Institute of the Blind, along with other organizations representing people with vision problems, has been campaigning vigorously against a ruling by the National Institute for Clinical Excellence (NICE), an agency that evaluates treatments and decides whether the National Health Service should pay for them or not. Heart surgery is on the approved list; nose jobs are not.

The controversy stems from NICE's half-hearted endorsement of a new treatment called photodynamic therapy. The therapy uses a drug called "Visudyne" or verteporfin, combined with a low-intensity laser treatment, to destroy lesions under the surface of the eye's retina, usually without damaging the retina itself. If the lesions are not treated they can irreversibly damage the center of the retina, called the macula. The resultant condition, age-related macular degeneration (ARMD), destroys central vision so that the victim cannot recognize faces, read, or drive. It is the leading cause of blindness in the United Kingdom.

In 2002 NICE filed a report recommending photodynamic therapy only in more extreme cases, only when both eyes are affected, and only in the eye that is less seriously damaged. The implication is that even treated patients will lose their sight in one eye, while others whose sight might be improved are denied treatment altogether.

It is easy to condemn NICE without appreciating its methods and the situation it is in. The basic challenge confronting the

National Health Service is that it has a limited amount of money to spend and an unlimited number of ways of spending it. It is no good asking patients, who pay little or nothing for treatment and will as a result demand more of everything. So NICE must resolve the inevitable dilemmas, determining who will get what type of health care, and who will be left to fend for himself.

How can medical spending be decided under such conditions? What would *you* do if you were in charge of NICE? It's a near-impossible task, but you would probably work out the costs and the effects of each treatment, and then you'd compare them to each other. Sometimes this is quite simple: a treatment with a 20-percent chance of preventing another heart attack is better than a treatment with a 10-percent chance of preventing another heart attack. Under pressure to make decisions, you might go farther and say the first treatment is twice as good, and should be used if it is less than twice as expensive. Even to go that far would be a stretch. How, then, would you compare a treatment that increases the chance of walking again after an accident, with a treatment that reduces the likelihood of going blind? Impossible! But if you were running NICE you would have to try.

The way NICE does it is to calculate the impact of each treatment in "Quality-Adjusted Life Years," or QUALYs. A treatment that saves ten years of life is better than a treatment that saves five years of life; a treatment that gives somebody ten years of able-bodied life is better than a treatment that gives somebody ten years alive but in a coma. Evidently, the value-judgments involved are extraordinarily difficult. Yet they must be made in a system that provides health care free of charge.

As an example, think of the problems involved in judging the QUALY impact of photodynamic therapy, which reduces the chance of blindness. The best way for the Royal National Institute of the Blind to get a higher priority for photodynamic therapy is to argue that a year alive but blind is worth much less than a year alive and fully sighted. If NICE accepts this view, treatments curing blindness will become very valuable on the QUALY mea-

sure, which places years alive and fully sighted high above years alive but blind.

But hang on. The strict logic of the view that "it's bad to be blind" would suggest that while that claim places a high priority on vision treatments, it places a low priority on treating people for *other* illnesses if they are already blind. If two people, one blind and one sighted, simultaneously turn up in a hospital in coronary arrest and there is time to treat only one, the QUALY methodology offers a truly unpalatable conclusion; that it is more worthwhile to help the person who can see rather than the blind patient.

We could backtrack and argue that in fact there is no difference in the value of life for the blind compared with the sighted. That is certainly more comfortable. Unfortunately, in conjunction with the QUALY methodology, this produces the conclusion that there is no point in spending anything at all on photodynamic therapy—or indeed on a pair of glasses. If treatments do not improve the value of people's lives then they are not worth spending money on, particularly when there are many causes, such as treatment for cancer. which certainly do improve the value of people's lives.

It is no wonder that the Royal National Institute of the Blind steers well clear of even mentioning QUALYs. It simply argues that photodynamic therapy is proven to improve vision and so should be comprehensively available. I do not blame the RNIB. But given the problem NICE is trying to solve—allocating finite resources among an unlimited range of medical treatments—it is easy to understand the position NICE takes: in particular the apparently heartless ruling that treatment should be applied in only one eye, leaving the other to go blind. The dispassionate perspective of QUALY analysis argues that the difference between having two good eyes and one good eye is less significant than the difference between having one good eye and none at all. Small wonder that the calculations tend to churn out embarrassing recommendations. But a free service will always be in demand, and it is hard to see a better way of rationing it.

Fixing health care with keyhole economics

Keyhole surgery techniques allow surgeons to operate without making large incisions, minimizing the risk of complications and side effects. Economists often advocate a similar strategy when trying to fix a policy problem: target the problem as closely as possible rather than attempting something a little more drastic.

How, then, can we fix health care? The insurance-based market solution is misfiring badly in the United States, in large part because of the problem identified by Akerlof's lemons model. The result: expensive, bureaucratic care . . . and even that, only for some.

The British approach has been to sweep away the market completely and replace it with a system governed by the decisions of bureaucrats like NICE rather than directed by market prices, as though part of the old Soviet Union had been transplanted into the hospitals and surgeries of the English shires. Fortunately, political and bureaucratic decisions are much more accountable in the United Kingdom than in the USSR, so the system works fairly well. But this is a colossal and wide-reaching response to the serious but rather specific problem of inside information. We owe it to ourselves to ask: is there a "keyhole" solution, which could fix health care without sacrificing the ability of patients to decide how much they value their own eyes?

Keyhole economics would first identify the specific market failures, which fall into three categories: scarcity power, externalities, and imperfect information, plus the issue of fairness. Scarcity power is a potential problem, but for most treatments not a significant one. In the United Kingdom, for example, there are roughly fifteen hundred patients per general practitioner (the doctor who is the first port of call for most patients using the National Health Service). So, a small town of nine thousand people can support six doctors, probably more than enough to encourage real competition, in a country where 90 percent of people live in urban areas. Some special treatments will wield

greater scarcity power—people fly from Australia and New Zealand to Hawaii for treatment with the Leksell Gamma Knife, a device for treating brain tumors. So there are some situations, but few, where scarcity power is a concern.

Externalities, too, are important only in select cases: for instance, for public health projects to restrict communicable diseases. (If everybody else had been using condoms to protect themselves from HIV/AIDS, I would not have needed to bother.) Yet neither externalities nor scarcity power are so severe or widespread that government provision becomes an attractive alternative. The keyhole solution would be some light-touch regulatory oversight to prevent the exploitation of scarcity power, coupled with focused subsidies to boost inoculation programs.

Fairness is not, strictly speaking, a market failure; it is something that even perfect markets do not necessarily provide. But we care deeply about fairness when it comes to medical care, both because we do not want the poor to be deprived, and because the cost of health care can vary dramatically, depending on the luck of each individual. In a civilized society we will want to make sure that everyone can afford some standard of medical care. The best way to do this is to tackle the general problem of poverty (think back to the "head start theorem" discussed in chapter 3) using redistributive taxes. After all, why spend so much to provide free medical care to poor people while ignoring the fact that they cannot afford healthy food or a safe house to live in?

That leaves inside information as the big obstacle to a well-functioning health-care system. The economic analysis we've done suggests that government provision is ineffective because decision making is out of patients' hands, and resources are rationed by political processes. Meanwhile, the overwhelming problem for the market provision of medical care is inside information, and more specifically its tendency to destroy insurance markets.

This diagnosis suggests a two-part keyhole treatment. The first part is to ensure the widespread availability of information: it should be easy to get a second opinion, easy to call a help-line, and easy to get information from libraries, clinics, the Internet,

even supermarkets. In the United Kingdom, people do not pay much attention to this information because doctors make the decisions. If we were asked to take responsibility for our own medical care we would pay much closer attention, and many more resources (public and private) would respond to our demands to know more.

The second part is to give patients an opportunity to use this information. In a privatized, insurance-based system the insurance company tends to make a lot of choices; in a government-provided system the government makes the choices. In a market-based system without insurance, the patient makes the choices. Much better. But the patient also has to pay for unpredictable and potentially catastrophic health-care costs.

How to give patients choice and responsibility without putting an unbearable burden on them? The best system would be one that compels patients to pay for many of the costs, thus providing an incentive to inform themselves and to make choices that are both in their interests and reasonably cost-effective but which leaves the most severe costs to the government or insurance. This might work, because most medical bills are not catastrophic and so do not need insurance.

How might such a system work, in more detail? The aim would be to give maximum responsibility and choice to patients, therefore requiring them to spend their own money rather than that of governments or insurers, but to make sure that nobody faced catastrophic medical bills and to make sure that even the poor had enough money to buy medical care.

These requirements suggest: people should pay for all medical care; but insurance should cover the largest bills; and that everyone should have a savings account dedicated to medical expenses, to which the government would contribute in the case of the poor or the chronically ill.

Catastrophe insurance, which pays out only when a particular course of treatment is very expensive, is fairly cheap. The savings are no problem either: simply reduce each person's tax bill by,

say $1,500 a year—this is very roughly the cost, in taxes, of both the UK and the US public health systems—and make them put the money in a savings account. For people who pay less than $1,500 in tax a year, the government would contribute money to make up the shortfall. Since the system is compulsory, no adverse selection takes place.

If you participated in such a program, how would it work for you? Your health-care savings would automatically go into a high-interest bank account. They would build up gradually throughout your life. For most people, medical bills are low in their younger years. So you could expect to have thirty thousand dollars in your account when you turn forty; more, if you've managed to keep your spending low and watched the money earn interest. Thirty thousand dollars buys a lot of medical care. Of course, it could all be consumed by a single expensive procedure, except that catastrophe insurance restricts your expenses.

If you reach retirement age with money still in your medical savings account beyond some minimum, you can put the excess toward your pension. When you die, you can pass the savings to other people's savings accounts (usually your spouse or children). So at every point in your life, you would have an incentive to spend money only on health care that you feel is absolutely necessary. If you felt that the right treatment for you was a bit of preventive maintenance—a course of shiatsu, say—then that would be your choice. You might well consider giving up smoking, given how much it would cost you in medical bills over the years. The catastrophe insurance would still pay for your lung transplant, of course, but no humane system can avoid moral hazard completely.

If one day your optician told you that you were suffering from age-related macular degeneration, but that a treatment with photodynamic therapy would increase your chances of maintaining your sight for a few more years . . . well, the choice would be yours. The photodynamic therapy drug, Visudyne, costs $1,500 a treatment. It would increase your chance of keeping good sight

from about 40 percent to 60 percent. No need to start talking about QUALYs: it's your money, and your choice.

The exception would be if you had a catastrophic expense, in which case the insurance company would prefer to pay for the cheapest treatment while you would want the best—a difficult problem, but no different from the conflict of interest faced for every single treatment in our medical systems today. The new system simply means that that inherent conflict of interest happens far more rarely.

It is common to provide goods and services on the private market, but one of the main alternatives—certainly, the main ideological alternative—has been to provide them using a political market instead. Medical goods and services are among the most difficult to distribute. We have tried using political markets, but they have let us down badly, and for obvious reasons.

At first sight, the private market failures, exemplified by the US system, are also obvious. But when we closely examine the market failures, it's the lack of information that is most serious, and the insurance market suffers the most serious consequences. American citizens receive much of their medical care through the intermediation of this badly malfunctioning market.

With some imagination, and some economics, we can step back from the troubles of our current systems and think about how to fix them. In Singapore, the system sketched in the last few pages has been successful for almost two decades. The typical Singaporean lives to the age of eighty, and the cost of the system (both public and private) is a thousand dollars per person—less than the cost of the bureaucracy alone in the United States. Each year, the typical Singaporean pays about seven hundred dollars privately (the average American pays twenty-five hundred dollars privately) and the government spends three hundred dollars per person (five times less than the British government and seven times less than the American government). Keyhole economics works.

The reason why Singapore's success is uncommon is probably that policy debates get stuck with one side claiming that we should rely on the market, and the other side asserting that the government would do a better job. So, government or market? We've learned that the question doesn't make any sense in isolation. To answer it we need to understand why markets might work, and how and why they fail.

We learned in chapter 3 exactly why markets work: because our choices as consumers between competing producers gives them both the right incentives and the right information to produce the right amount of exactly what we want. And we've also learned that scarcity power, externalities, and inside information can each ruin the way markets do this.

In the case of health care, the market works poorly because while we want the reassurance of knowing they can afford expensive medical bills, inside information eats away at the insurance by driving away low-risk customers and forcing premiums to rise. Private companies have developed ways to get around the problem, but they are expensive and bureaucratic. Singapore's government had the power to tackle the problem head on, by using forced saving and catastrophe insurance to make sure costs were manageable but keeping the power of patient choice at the heart of the system. Governments can replace markets, but they will often do better to try to fix them. They are unlikely to succeed unless they appreciate exactly what the problem is in the first place.