

There is a tendency among economists to refer to “the market” as if it were like a clock. You can find all sorts of clocks, running on different kinds of power, analog or digital, mounted on walls or on your wrist, with different numbering styles and decorations, but in the end they are all just clocks. They tell you what time it is on a 12 or 24 hour cycle. If you need this basic information, any clock will do.

Not so with markets, however. Markets do bring buyers and sellers together, but beyond that they differ not only in appearance but also in purpose and performance. Markets comprise a class of social institutions, like families and governments, and their characteristics have a profound effect on economic outcomes. This has become particularly apparent in the course of the transformation of the formerly Communist countries of Asia and eastern Europe to more capitalist principles. All are becoming “market economies” in some sense, but their success depends crucially on the *kinds* of markets they are developing. Fortunately, economics has some useful things to say about the different ways markets solve the problems put to them. Based on this, we can not only choose intelligently between markets and other systems of allocation, but also among different market types.

7.1 Markets in History

As far as we are able to know, there have been markets for as long as there has been organized human society. Archeological evidence from Europe, Asia and the Americas indicates that prehistoric humans used objects that were gathered or produced hundreds or even thousands of miles away, and presumably some form of exchange was the basis for transferring these goods from one location to another. With the dawn of civilization in China, India, Egypt and Mesopotamia we find coins and designated trading sites.

Markets in all such societies were *places*. They were central locations where goods were brought together for sale and where negotiations over transfers of ownership were carried on. Nevertheless, markets fulfilled other, less directly

economic functions. The Greek agora, for instance, was a marketplace, but it was also where political decisions were made and social relationships developed.

Only in the last few centuries have people begun to think of markets as pure, disembodied realms of exchange, placeless institutions with no other function than to facilitate trading and establish prices. Indeed, some markets appear to have those characteristics today. Think of financial markets, which take their ethereal existence across global computer networks. Trading in standardized assets like stocks and government bonds is entirely anonymous; no one knows or cares anything about the person whose trade shows up on the computer screen. All that matters is the price and quantity.

Actually, markets have never entirely separated themselves from their social context. As we will see, they still have to insert themselves into networks of information and expectation, and these are the products of complex social organization. Even the financial markets are less disembodied than they appear. (This becomes apparent during episodes of fraud, when the hidden role of those who make markets function rises to the surface.) All markets have to solve, or try to solve, certain problems, and all such solutions, as we will see, require active human intervention to produce results. The rest of this chapter will identify the most common problems and strategies for coping with them.

7.2 The Enforcement Problem

Consider this story: Fred and Ginger run into each other at the market. Fred looks at Ginger and says, “I like your backpack”; Ginger looks at Fred and says, “I like your watch.” As it happens, Fred has a box full of watches back at home, and Ginger keeps a backpack collection in her closet, so Fred would rather have Ginger’s backpack than the watch he is wearing, and Ginger would rather have Fred’s watch than her backpack. A mutually profitable trade appears to be in the offing.

Now let’s make two further assumptions. First, as is common in economics, we will assume for convenience that both individuals are purely self-interested. Fred gets no pleasure from Ginger’s feelings, nor Ginger from Fred’s. Second, we assume that each person’s action is entirely independent of the other’s. That is, both Fred and Ginger decide separately what to do with their possessions, and whether or not a trade occurs will only be known after the fact. Based on this, and with two goods and two individuals, we can envision four possibilities, each of which can be ranked by Fred and Ginger according to their desirability. The most desirable outcome for Fred, which we will designate as a 4, is that he keeps his watch and also acquires Ginger’s backpack. The second most desirable is that he gives up his watch but gets her backpack. The third is that he keeps his watch and fails to get her backpack. The fourth, of course, is that he loses both. A similar ranking holds for Ginger, except that she prefers having a watch and no backpack to a backpack and no watch.

Each person has the option of taking one of two actions, supplying what they have to the other or keeping it for him or herself. Using the game theory notation of

| | | GINGER | |
|------|---|--------|--------|
| | | D | C |
| FRED | D | (2, 2) | (4, 1) |
| | C | (1, 4) | (3, 3) |

Fig. 7.1 Payoff matrix for a potential trade between Fred and Ginger. Fred and Ginger are the players; C (supplying the good already possessed) and D (not supplying it) are the choices. There are four possibilities depending on which choices are made, ranked from 1 (worst) to 4 (best). Within each set of parentheses, the payoff for Fred is given first, then the payoff for Ginger

Chap. 3, call the first option cooperation and the second defection. The result is the payoff matrix shown in Fig. 7.1.

As this diagram makes clear, Fred and Ginger face a Prisoner's Dilemma. Each has an incentive to try to acquire what the other has without giving up what they have. Whether or not Ginger gives Fred her backpack, Fred is better off by hanging onto his watch, and the reverse is true for Ginger. If each follows a purely individualist motivation, they will end up in the cell in the upper left-hand corner, keeping everything and exchanging nothing. Only through mutual cooperation can they make themselves both better off by each giving what they have to the other. **A market exchange is a solved Prisoner's Dilemma.**

This last statement may seem a little extreme. Surely, you would think, the billions of transactions that occur daily in a modern economy can't all be intricate negotiations of collective action problems. Above all, doesn't the assumption of independent choice (Fred's decision to part with his watch is completely independent of Ginger's decision to give Fred her Backpack) contradict how markets usually work? When you buy a loaf of bread in a store, you do not give up your money independently of getting the bread, nor does the store hand you the bread until you pay for it at the cash register. If decisions are not independent, then the upper-right and lower left-hand cells (where one person gets everything and the other nothing) are not possible, and the only choice is whether to exchange (lower right) or stand pat (upper left).

These are valid criticisms, but experience indicates that it would be a mistake to reject the independence assumption altogether. Many market transactions involve promises which come due at different times. It is certainly possible for people to break their promises, either on purpose or because they are not as careful as they might be to avoid circumstances that will lead to the promises being broken. Loans provide obvious examples, but much the same could be said of other longer-term contracts, as when one company promises to make future shipments of supplies to another. Moreover, it is often possible for sellers to cut corners on quality, and this is comparable to a sort of withholding. Suppose you agree to buy two machines from me, each of which costs me \$500 to make. If I quietly cut back on quality, lowering my cost to \$250 apiece, what is the difference (for me) between this form

of deceit and only shipping one rather than both at the original quality? As for you, it depends: maybe two lower-quality machines are better than one higher-quality one and maybe not.

The general point is that, for markets to succeed and exchanges benefitting both parties to take place, there has to be a reliable expectation that buyers and sellers will keep their promises. The expense of bringing about this expectation, and having it borne out in most cases, is one type of **transaction cost**. A transaction cost is the cost of using market exchange to achieve an economic objective. This category covers everything from searching for suppliers and consumers to comparing various offers to writing up contracts to enforcing them. It is perhaps the single most important concept for understanding markets as institutions. If it is too costly for people to use a market, they won't and the potential advantages of this system of allocation will be lost.

To return to our problem of promise-keeping, we can identify several approaches that have worked in different times and places to ensure that markets do not break apart in an epidemic of mutual defection.

Social norms. As we saw earlier, markets did not arise in isolation from other human activities, and they are not usually isolated today. People who take part in markets have social as well as economic lives. They are raised to internalize the values of their culture, and they respond to signals from those they interact with that tell them whether they are behaving appropriately or not. Recall the argument from Chap. 4, for example, that if most members of a society are calculating the costs and benefits of engaging in crime, the battle has already been lost; the first line of defense is always the inculcation of values that favor honesty and respect. Fred is not likely to try to cheat Ginger, partly because it is likely that he has been raised to feel uncomfortable about cheating (even if he gets away with it), and also because he would feel bad if Ginger found out about it and let him know. Multiplied by millions and billions of people and transactions, this is an irreplaceable foundation for successful markets.

Social scientists use the term **social norm** to refer to a rule for behavior that has the properties of being nearly universal within a society and binding on its members. Once such a rule gets established it can have considerable inertia behind it. If everyone else expects you to observe a particular custom, you pay a price for violating it—but if you go along, you become part of the “everyone else” for other members of your society. Healthy societies have rules for promise-keeping, among other behaviors, that facilitate the operation of markets. Much depends on the maintenance of these norms.

Unfortunately, it is also possible for the expectation of promise-keeping to unravel. If a few individuals violate this norm visibly and are seen to prosper from it, they will attract imitators. After all, in situations with characteristics approximating a Prisoner's Dilemma, there is a powerful material incentive to take advantage of the cooperation of others. Before long, as defection becomes the normal expectation, it simply becomes too costly for most of those clinging to the cooperative norm to hold out. Worse, a situation of generalized defection (a dog-eat-dog world) can also be self-reinforcing and therefore stable. The

message here is that it is extremely important for societies to police the boundaries of social norms relating to honesty and mutual respect. Violators cannot act with impunity.

Reputation. Adam Smith was well aware of the problem posed by dishonesty in an incipient market economy like Britain's. He argued that most people would realize that it was in their interest to act fairly, since the short-term advantage of opportunistic behavior would be overshadowed by the long run cost to the scoundrel's reputation. In more modern terms, he recognized that the Prisoner's Dilemma game of the market, if that was indeed what it was, would be played repeatedly among a small enough group of players that each would be able to remember who had behaved honorably in the past. In this way the healthy virtues on which useful markets depend would be inculcated naturally. Indeed, Smith expected that individuals would come to see their own behavior through the eyes of those they dealt with, so that reputation would be truly internalized: people would experience the views of others as their own. In other words, if you have a positive reputation, this would be reflected in positive feelings you would have about yourself, and you would do the right thing in order to enjoy these good feelings.

Unfortunately, unless this internalization is permanent and very powerful, Smith's argument depends crucially on the assumption that the number of players is small enough that reputations will follow each of them around. This was plausible in the Britain of his day. Small-scale transactions, such as those between landlords and tenants or shop-keepers and customers, were largely confined to local markets based on face-to-face recognition. Larger financial deals were conducted among a small class of property-owners who knew of one another despite being relatively more dispersed. As markets grew larger and more anonymous, however, it became evident that neither the direct or indirect effects of social reputation would play a sufficient disciplinary role without some measure of conscious support.

Modern economies use a variety of means to reinforce the power of norms and reputation. Think of the role played by newspapers and television, which aggressively report on businesses accused of offending the public trust. Most of us are familiar with the mass media, but in some ways even more important work is done by the specialized business press, which digs up the dirt on culprits within its domain. Every major sector of a modern economy has such a publication, and readers pay close attention to reports of dishonesty or incompetence. Another example is the trade or professional association. In principle, doctors, attorneys, accountants and other professionals police their own ranks; if they find someone has stepped out of bounds, they revoke his or her license to practice. Trade associations also take complaints and may publicize or expel violators. Governments may also maintain lists of approved or licensed firms based at least in part on adherence to a minimal set of standards for lawfulness and honesty.

The importance of the roles played by such institutions is illuminated for us when they fail to do their job. During the US corporate scandals in the early part of the previous decade, for example, information emerged that major accounting firms were acquiescing in dishonest bookkeeping by the companies that hired them,

hiding their expenses and reporting fictitious income. This harmed the interests of their legitimate business partners, investors and employees. What failed in this case was not just the integrity of individual accountants and their companies, but the oversight of the profession as a whole. Among insiders in the field there were certainly suspicions, and perhaps even hard evidence, that standards were being violated, but the information was not made public. Reputations that should have been tarnished at the first whiff of impropriety were allowed to remain intact.

Contract design. For interactions that take place over time, it is sometimes possible to introduce elements into a contract that provide an incentive for one party to perform to the satisfaction of the other. There are many familiar examples: a contractor engaged to build a house will be paid in stages or even in a lump sum at the end of the project rather than upfront, a worker may be hired at a lower pay scale with raises tied to performance, and a bank may charge borrowers a penalty if they are late in their payments. What all of these have in common is that one party to the contract, typically the one in a better bargaining position, gets the other to agree to an arrangement in which the stronger will be in a position to evaluate and reward the weaker at regular intervals. From the point of view of the side that is able to extract such concessions, the contract goes a long way toward solving the enforcement problem. The weaker side may not be as pleased, however. For instance, most employment contracts are written in such a way that firms have more leeway to alter wages in response to their evaluation of the work an employee is doing than the employee has to alter, say, effort in response to the firm's follow-through on matters like working conditions and training.

In recent years there has been a flurry of interest among economists in the study of contract mechanisms that can embody their own internal incentive systems. A poorly designed contract can be self-defeating, encouraging one side or the other to behave in a way that makes the relationship unproductive. Consider, for example, a contract in which a producer of upholstery agrees to supply an automobile manufacturer with seat covers over the course of a year. If the agreement provides for bonuses for on-time delivery but no penalties for defects, you can imagine what the supplier will try to do if production starts to fall behind schedule. Similarly, is it a good or a bad thing to reward someone for success over which they had only partial control? A company that rewards its workers when profits go up would seem to be on the right track, except that profits depend on a great many things in addition to workers' effort and skill. What effect would paying this reward have if profits rose as a result of the mistakes of the competition, and not for anything the company, including its workforce, did on its own? As we will see in the next chapter, there can be large costs from failing to think through all the consequences of contract design.

Contract law. The last line of defense against fraud and abuse in the market is the legal system. It can't be the first line, because it would be too expensive to litigate every possible violation, but no system of norms and reputation will deter all potential opportunists. The main branch of the law that adjudicates complaints arising from failure to keep promises is contracts. Attorneys for the plaintiff argue that the terms of a contract were violated, resulting in harm to their client. Defendants' attorneys dispute the interpretation of the contract or the facts of its

performance. Judges, usually with extensive training in contract law, conduct the trial and render an opinion. The entire system is a backstop for the effective functioning of the market system.

There are many ways for such a system to malfunction. The law itself may be undeveloped, as has been the case in some countries emerging from Communism. (Hungary, for example, delayed its privatization of state enterprises until its entire body of property and contract law could be rewritten.) Attorneys may not adhere to high professional standards; they may fail to make effective arguments or may try to use political or other pressures rather than the force of legal evidence. Finally, and perhaps most important, the judiciary itself may not be independent or competent. Judges may have a personal stake in the disputes before them (a conflict of interest) or may be under the control of outside political or economic factions. Perhaps they just don't know how to sort through the complexities typical of modern contract law. All of these shortcomings have appeared in countries that have tried to rapidly introduce market systems during the past two decades, but they also appear from time to time in countries with longstanding market economies.

7.3 The Complexity Problem

Legally, market transactions normally take the form of either bills of sale (immediate exchanges of goods for money) or contracts (exchanges of promises). The first of these is simple and straightforward; take a look at your next cash register receipt. Contracts, however, can be complicated to draw up, and if they are they will also be costly to execute and enforce. The smooth functioning of a market economy depends on its ability to keep most contracts to a manageable level of complexity.

An example will reveal what is at stake. The real estate market is the venue for the buying and selling of land and buildings erected on it. Such transactions tend to be lengthy, complicated and expensive. Much may depend, for instance, on exactly where the boundaries of a property are located, so this must be ascertained as part of the agreement. Houses must be inspected for hidden as well as visible features, and detailed instructions must be drawn up for items that may or may not be conveyed, such as appliances, fixtures and even decorative elements (like stained-glass windows). There are many ways third parties can have a claim on real estate (as collateral, through easements, etc.), so all such claims must be researched. It usually takes many weeks from the initial sale agreement to the final closing, and, between them, the parties should expect to pay a few percent of the sale price to the small army of lawyers, accountants, appraisers, insurers and others who process the complexity of the transaction. If every market in the economy were as complex and costly to operate as real estate, the system would soon grind to a halt.

At its heart, what makes real estate such a complicated business is that its product is nonstandard. You can't just say "acre" (or "hectare"), "house" or even "three-storey Victorian" and leave it at that; every parcel and structure has to be scrutinized characteristic-by-characteristic. Think how different this is from "car". Buying a car is also complicated, but much less so. If it is a new car, you write up an

agreement specifying make, model and any additional features desired. Also to be agreed on are the terms of payment—how much down, what interest rate and over what period of time. It is a bit inconvenient, but nowhere near as costly and cumbersome as real estate.

The difference can be summed up in a single word: standardization. A few specifications are sufficient to indicate everything that needs to be indicated about the car, because each is standardized. Every car with those specifications should be identical to every other. If cars were built by hand, one at a time, according to the passing whim of the builder, buying cars would be more like buying houses. (There is indeed a small market in custom-built cars, but even here there is more standardization than in real estate.) Standardization, which in some ways we lament as bleeding serendipity and richness from life, is essential to a modern economy.

How do economies achieve standardization of most goods and services traded in markets? Here are three broad approaches:

Private sector standardization. Standards are often set by the producers themselves, either separately or through associations. Since converging on a few standards can help expand the market for their products, firms are usually well aware of the issue. If a single firm (a monopoly) can command a decisive share of the market, it is in a position to make its own standards the ones all others must follow. (We will take a closer look at the advantages and disadvantages of this phenomenon in Chap. 13.) The most familiar contemporary example of monopoly standardization is the Microsoft Windows operating system (and related Microsoft software). For just over a decade, from around the time of World War I to the late 1920s, the Ford Motor Company had similar standardization influence in the US auto industry.

Trade associations have long sought to standardize the products of their members. This can be seen in the regulations promulgated by German producer guilds going back to the Middle Ages—consider the German Beer Purity Law, dating from 1487—but an interesting American example can be found in the grain trade. When white settlers moved to the middle west, such as Illinois and Iowa, they found some of the most fertile land for growing corn, wheat and similar crops in the entire world. They began to farm, and bags of grain were soon loaded onto ships and trains heading east to the large urban markets. Initially, each farmer's grain was bagged separately, with the grower's name attached. Since the commodity in question was "farmer X's grain", which would normally be different from farmer Y's (due to differences in seed, growing conditions and farming methods), each farmer had to negotiate separately for his price. The system was complicated and burdensome to both growers and buyers. A farmer had to consider whether the price he could get from a different buyer or in a different region might be better than the one offered him at the moment. There was no way to find out except to cart his bags somewhere else and negotiate all over again.

This logjam was broken by an association formed by the railroad companies. They had no interest in keeping track of which bag was grown by whom; they were building giant grain elevators and wanted to just dump the cargo en masse. Their solution was to devise a system of grading. Rather than being identified by its

grower, each shipment would be classified by general variety (such as summer or winter wheat) and by rough standards of quality (number 1, number 2 and so on). This immediately had several effects. It removed the uncertainty over price, since all grain of a standardized type and quality would be paid the same. It made shipping far more efficient, since grain could be moved in common containers and mechanically dumped into elevators. And of course it also gave the railroad companies more leverage over the growers, since they were the ones who decided what the quality categories would be and did the actual judging. Beyond this, it stimulated the growth of secondary markets, such as trading in grain futures. (A futures contract is an agreement to buy or sell a specified quantity of a specified quality at a specified price at a specified time in the future.) With the standardization of grain it became possible to speak of, say, number 2 winter wheat as the “same” commodity now or in 90 days, even though the actual wheat might be grown by two different farmers a thousand miles apart. In reflecting on this system, we can recognize both gains and losses. The farmers lost a bit of clout, and subtle differences in the product itself were erased through the imposition of broad standards. On the other hand, the far greater efficiency attained by standardization ultimately benefitted most farmers, if not to the same degree as others in the industry (shippers, middlemen, speculators).

Without looking too far, you should be able to find many other trade groups that have succeeded in constructing standards to facilitate trade.

Public standardization. Government agencies sometimes play a lead role in standardizing products. Many states, for instance, regulate apprenticeship programs in the building trades for such skills as electrician, carpentry, masonry, etc. There are multiple reasons for doing this, but an important one is to standardize what it is to be, say, a mason and what skills a client can expect to be purchasing.

A controversial recent example at the Federal level is the “organic” designation for food products. After months of hearings, the Agriculture Department issued a set of standards that all growers would have to meet if their produce was to carry an organic label. These included the number of years of continuous organic practice and the specific types of fertilizers and pest control substances that could be used. Certification has proved expensive, discouraging many smaller growers from taking advantage of the system. In addition, it has turned out that there isn’t a clear line separating organic from conventional; it is a matter of degree, and minimum standards fail to reward growers who go beyond the lowest common denominator. All the same, the standards appear to be facilitating the rapid expansion of this industry by removing confusion over which goods should be marketed under what set of guarantees.

Technological standardization. Quite often standardization is achieved not through any conscious intervention but through the technical dictates of the production process itself. Railroads have a common gage (distance between left and right wheels) because they have to in order to run on the same tracks. Air traffic control signals have to be standardized internationally because airplanes fly through multiple control regions. Above all, by its very nature, mass production standardizes products in order to achieve economies of scale. (This term will be

defined and discussed in Chap. 12.) That is the underlying reason why buying a car is not like buying a house. The car you buy will be nearly identical to thousands of others, not so the house. If we ever transform housing construction into a mass production industry with cookie-cutter products, real estate transactions will become simpler (and perhaps houses will become less personal).

7.4 Markets and Information

In many ways, markets are remarkable instruments for discovering and bringing together information about the economy. Each participant is in a position to make decisions about what and how to produce or consume based on the information that, perhaps, only they know. Drawing on what they have learned they can make and accept offers, and the prices that result in the marketplace will reflect, in some sense, the accumulated knowledge of everyone who participates. As this is being written, for instance, there is uncertainty in the global petroleum market due to the unsettled prospects for the world economy. If the major industrial countries are able to recover and resume economic growth, their demand for petroleum will go up, adding to already increased demand from countries like Brazil, India and China. If there is a second global slump or a slowdown in China in particular, however, demand will fall. There is also much debate about future supply, for reasons of both geology and feasibility (especially in deep-water deposits). Further down the road, there is uncertainty over actions governments might take to curtail energy use as a means to limit climate change. Meanwhile, around the world there is a constant buying and selling of oil contracts, and the price of the precious fluid is in flux. When traders anticipate greater demand, the price goes up; when they anticipate less, the price falls. In this way day-to-day price movements track the perceptions of those who are following the situation closely. If someone has an informed hunch that suggests a large new oil deposit is going to be announced, she might put in an order to sell, speculating that when this news gets around the price will fall (increase in supply) and she will be able to re-buy the oil she has just sold at a lower price, and therefore a profit. Sitting in front of my computer screen several time zones away, I have no knowledge of who this trader is or what she thinks she knows, but if she and several others like her have an effect on the market, I will see the price dip. This tells me that there has been a change in the perceptions of traders: they think supply will rise, demand will fall or perhaps both.

This episode tells us four things:

1. Markets reward people for having information that is better or sooner. Those who speculate in oil, or even just those who think carefully about major purchases or employment decisions, will fare better if they are better informed. Just as you would have an incentive to research the repair record of a used car you are considering buying, an oil trader has an incentive to learn all she can about the industry and its prospects. Information is valuable.
2. Market prices reflect the perceptions of all the people who participate through buying and selling. This is true for better or worse; if traders are misinformed or if they misunderstand the signals they are getting the price will reflect this too.

Market failures can also disrupt the process, as we will see in more detail in the chapters to come. Nevertheless, because of the first point, markets are somewhat more likely to be right than wrong in their perceptions. There will always be a range of views, of course, but markets summarize them through aggregation—yielding a price that reflects the relative pressures of overall supply and demand.

3. All the information that flows through a market is condensed into a single number, the price buyers are willing to pay and sellers willing to accept. What gets lost is all the richness: where the tip about the discovery of new oil fields came from, how reliable it is, and even that the information pushing down the price of oil is about potential future increases in supply and not, say, a drop in expected demand from a major consumer like China or Brazil. Nevertheless, with millions of prices being set every day, it is a great convenience to have information in this simplest of forms.
4. People in other sectors of the economy have to take prices into account, and when they do this they are, in effect, incorporating the information on which the prices were based. If the price of oil goes up, for instance, a community group may have to reconsider its decision to send several of its members to a conference in another city; perhaps they can no longer afford the cost of transportation. This group knows nothing about oil exploration or controls on offshore drilling or other arcana of the oil industry, but they know a higher price when they see one. This price induces them to make decisions that are consistent with the new information, however: if, in this example, the supply is likely to be reduced, consumers should look for alternatives. The point is that the price system not only summarizes information, it invests information with economic force.

That is the good news. The bad news is that this machine does not run all by itself. There is still a need for people to actually ferret out the information that markets trade on, and how the market is organized can have an effect on the costs entailed. Also, the process by which information is brought together in markets is fraught with pitfalls. We will now look at these two problems more closely.

7.5 Search Costs

Much of the information that people need is about what is happening in the marketplace. Who are the other buyers and sellers? What goods are they offering or trying to buy? At what prices? What are their reputations? In order to make a wise choice, first you have to know what the choices are. Market decisions are complicated when goods are less standardized and when producers are smaller and more numerous. Sometimes the crucial information is virtually invisible, such as the true creditworthiness of a particular borrower or the skill and character of a worker.

As information becomes harder to get, some people specialize in digging it up. This leads to a *market for information*. A simple example is the newspaper want ads or Craigslist; the service being provided is little more than a bulletin board. More sophisticated are trade journals (take a look at *Advertising Age* or *InfoWorld*)

and various firms that provide industry data for a price. These enterprises are business opportunities for those who own or work in them, but they are crucial elements of the market system as well. Market information services are to a market economy what political information—a free press—is to political democracy. If they are ineffective or simply underdeveloped, markets will do their job more poorly.

When information becomes costly, either in terms of your own time or someone else's (and therefore your own money), you will have to learn how to economize on it. Information that is essential will have to be acquired one way or another, but information of only moderate relevance or importance may not be worth the expense. If you consider the implications of this rather obvious point, you will see that it undermines the model of rational decision-making presented in Chap. 3. Recall that the decision-maker is viewed as maximizing her expected utility (EU), where

$$EU(B) = \sum_i p_i v_{Bi} \quad i = 1, 2, \dots, n \quad \sum_i p_i = 1$$

This formula says that the expected utility of an option B is the sum of all the values it might have, each multiplied by the probability that it will be the actual value. To perform this calculation, an individual would have to know and be able to evaluate all these values and also place a probability on each one. Then it would be necessary to do this for every available option, so that the one that maximizes expected utility is selected. If information is costly, however, there will be missing pieces in this formula. Some values and probabilities will simply not be identified or estimated. What then?

A plausible answer was given by Herbert Simon, an early winner of the Nobel Prize in economics. A decision-maker will first decide what outcome will be deemed good enough; this sets a benchmark. Then he will begin looking at options, not with the intention of calculating their exact expected value, but just to see whether it meets the good-enough standard or not. This usually requires less information. As soon as one option passes the test, the process stops. The decision-maker selects this one and declares himself satisfied—or, to use Simon's term, *satisficed*. **Satisficing** is a decision process whose goal is to identify at least one course of action that meets a sufficient level of adequacy. Imagine that you are buying a house in an urban area, for instance. There are a multitude of houses for sale, each unique in its location and features. You also have to consider future possibilities that might affect your choice, such as where new development will be located, transportation bottlenecks in the years ahead, and whether particular neighborhoods are likely to become more or less desirable. Since gathering the information to answer all these questions would either take you several years or require you to hire an army of real estate analysts, picking the single best house over the full range of scenarios is out of the question. Instead you should go into the process with a rough idea of what sort of house, at what price, you will accept and then start looking. When you find one that appears to meet your expectations, make an offer.

A good rule of thumb is this: the more complex the information for a given market and the less support one can get (or afford) from information services, the greater will be the difference between the satisficing standard and the maximum expected utility that could be had from a full-blown investigation. This is an important insight; it says that information burdens can reduce the value people can expect to get from participating in a market.

Thus far we have been treating information as a homogeneous good. You can have more or less of it, but there are no distinctions to be made between different sorts of information. Of course, in the real world it is exactly the opposite. It is always easier to get some types of information than others. Houses for sale in some neighborhoods are listed more prominently or are more convenient to visit. It takes less technical know-how to understand some differences between alternative computer configurations than others. It's easier to interview a job applicant who comes from a similar social background.

In this context it is important to bear in mind that markets do not exist in a vacuum. They are surrounded by a complex layering of social networks based on family relationships, ethnicity, friendship, shared activities, etc. Every social connection is a channel that transmits information, so the types of networks that market participants are involved in will determine to a considerable extent the types of information that will be abundant or scarce. Studies of geographically concentrated industries, such as Silicon Valley in California or high-end apparel in Italy, have shown that clubs, associations and other groupings play a key role in conveying information about job openings, potential markets, business opportunities, etc. Indeed, there is an interesting chicken-or-egg question here. Which comes first—do markets develop in a region and eventually lead to wider and deeper social ties, or do dense social networks provide the basis for markets to prosper? Many economists now think that the second possibility plays the greater role, and this is having an impact on policy research, particularly in developing countries.

7.6 Asymmetric Information

Every market transaction has two sides, and each requires information. Thus far we have looked at the information problem from just one person's point of view, but now we will consider the impact on markets if information is not just incomplete but *unequal*. There are many reasons why this might be the case, but one is especially important: people, typically sellers, often know something about themselves or their products that others, usually prospective buyers, don't. Job applicants, who want to sell their labor, normally know a lot more about their own abilities and intentions than someone who just reads their resumé would be able to figure out. Sellers of specialized equipment often know more about their specifications than buyers who have to make a wide range of purchases and don't have time to study each one in detail. Borrowers, who are effectively selling a loan to a bank or other lender, know more about their likelihood of repaying than outside credit analysts. This situation, in which one party to a transaction has private

information unavailable to the other, is called **asymmetric information**, and it is widespread throughout the economy.

One solution to this sort of problem would be a rule requiring everyone offering to sell a good or service to disclose all their private information, even (and especially) if it contradicts the impression they wish to create. Such regulations have been imposed, for instance in the used car market: sellers are required to inform prospective buyers of all defects they are aware of. Such rules cannot completely do away with private information, of course, but they can greatly reduce its scope.

The opposite tradition, however, has been dominant in most capitalist countries over the past several centuries. At its most extreme it takes the form of **caveat emptor**, “let the buyer beware”. It indicates that the seller is not obliged to disclose anything, and it is up to the buyer to find out the information on her own or to convince the seller to put guarantees of quality into the contract. To take one striking example, it is not against the law to lie on your resumé—for instance, to claim you have a college degree when you never attended college at all. It is up to employers, the potential buyers of your services, to find this out for themselves. If they do, they can refuse to hire (or can fire) you, but that is it. You are free to send your fraudulent resumé to the next job opening. (Do I have to say that I am not encouraging you to do this?) In the world of employment, it’s caveat emptor in nearly every case.

There is a longstanding debate among legal theorists over the merits of this doctrine. It has the virtue of encouraging people to be responsible for their own choices—to do research, to demand guarantees if appropriate, to determine what risks are worth taking. It reduces the burden on the judicial system, which might be overloaded if every disappointed buyer thought he deserved a day in court. On the other hand, there is always a case to be made for combating dishonesty and guile. Current debate focuses on the role of contract clauses like warranties that guarantee quality. If consumers are in a position to purchase them, they can decide whether such a guarantee is worth the price they have to pay to the seller. If not, they have no recourse, and strict caveat emptor may undermine the confidence on which markets rely. These comments are very general, however, and careful analysis of the issue usually turns on the particularities of different markets.

So let us assume that asymmetric information in much of the economy is unregulated; where can it lead? There are many possible outcomes, but we will focus on one that was first introduced by Joseph Stiglitz, a major figure in the economics of information (and other fields as well). To do this, we will construct a simple but telling model.

Here is the situation: an employer has a fixed sum of money M to spend on hiring workers to do a project. Each worker will be paid the same wage w , so M/w gives us the number of workers N the employer can afford:

$$N = M/w$$

The goal of the company is to produce a quantity of product designated as Q . The more workers, the more product, so it would appear that there is little to figure out:

the employer should pay the lowest possible wage that will attract enough workers. Let's say that there is a larger labor market that sets a standard wage w_m ; at this wage a large number of workers will apply for a job, but not below it since they would find other work that pays w_m . If this were the entire story, the firm will pay w_m and hire M/w_m workers to produce as much Q as possible.

But now let's introduce a wrinkle: workers differ by quality, and quality affects output. Specifically, if we designate the average quality of the workforce as v , output is a function of this variable and the number of workers:

$$Q = f(N, v)$$

To make matters a little more complicated, let's assume that, while each worker knows his own quality, this information is not available to the employer. Asking will not help, because workers will claim to be high quality whether or not they actually are. Fortunately for the employer, there is one clue available: no worker will work for less than what he knows he is worth. If he knows he is 10 % more productive than the average, he will hold out for a 10 % higher wage, thinking that his value will eventually be recognized somewhere. The employer's problem is to determine a wage w^* that results in the highest possible level of output P^* .

Variables in the Invisible Worker Quality Problem.

M = the fixed sum of money available to the employer for paying wages

w = the wage paid to all workers

w_m = the minimum wage that will draw applicants, set by the market

w^* = the wage that results in the highest production of output

N = the number of workers employed

Q = the quantity of output produced

Q^* = the maximum possible quantity of output

v = average quality of the workforce

For simplicity, we can divide this problem into two components. The first is establishing the size of the workforce. From this perspective, as we have seen, there is little to consider; lower wages mean more workers and more output. Figure 7.2 on the next page displays this relationship, assuming, for the moment, that v remains unchanged.

Now let's look at the impact of raising or lowering the wage on average worker quality, leaving aside its effect on the number of workers that can be hired. At any wage equal to or above w_m all workers of the lowest quality will apply, but better workers will apply only if the wage rises to meet their higher expectations. Thus a worker who is 10 % more productive than the average w_m worker will demand a wage that is 10 % higher. In this way, the further the wage rises above w_m the more and better workers will apply. This means that the average quality of the applicant pool also rises, so that, even if the employer doesn't know which individual workers are the best, over a large number of randomly selected hires she will find the average quality going up. This is reflected in Fig. 7.3.

Fig. 7.2 Wages and output in the invisible quality problem, holding worker quality constant. w_m is the lowest wage the employer can pay and attract a labor force. Given a fixed amount of money to spend, this wage leads to the highest number of workers and, holding worker quality constant, the highest level of output Q^*

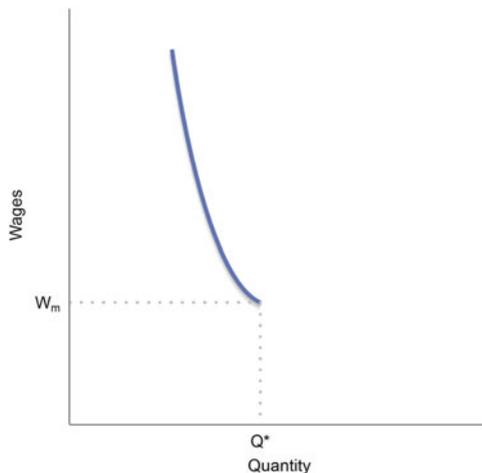
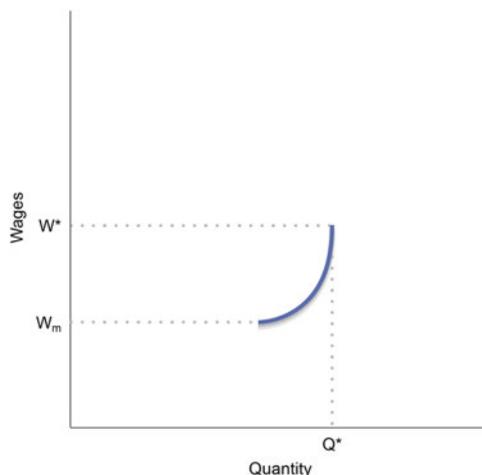


Fig. 7.3 Wages and output in the invisible quality problem, holding the number of workers constant. Average worker quality rises as the wage rises and therefore so does output. This relationship tails off as fewer workers remain to be discovered with very high quality levels. At w^* the very highest quality worker joins the applicant pool, and there are no further advantages to raising wages

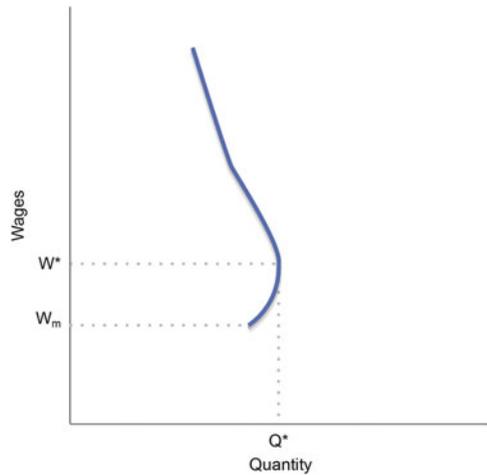


Now let's put both effects together in one diagram. Naturally, how they combine depends on their relative strength, and we have not provided any basis for determining this. One possibility is that, over an initial level of w above w_m the quality effect is stronger, but after that the numbers effect takes over. This would give us a combined wage-output relationship as in Fig. 7.4 on the following page:

As the wage rises above w_m output rises due to the predominance of the quality effect. Maximum output occurs when $w = w^*$. Above that level output declines. The employer in our story will therefore agree to pay a higher-than-market wage out of self interest.

Our story is completely hypothetical, of course, and it was tailor-made to yield its interesting outcome. It is fair to ask whether it has any real world counterparts. Many economists are inclined to say that it does. It is common for firms to pay

Fig. 7.4 Wages and output in the invisible quality problem, all effects included. Because of the quality effect, the wage that produces the most output, w^* , is higher than the lowest wage that attracts enough applicants, w_m



wages above the lowest level they would need to attract a workforce and, while other motives might account for this, there is some evidence that asymmetric information is a factor. Similarly, many firms are reluctant to lower their pay scales even during times of general wage slippage. Again, there are other explanations, but students of personnel practices think that maintaining applicant quality is a consideration. And the phenomenon extends beyond employment. A similar story can be told for lending: a bank that charges the highest possible interest rate is likely to attract borrowers who are, on average, less creditworthy than the lender might like. After all, if you think you have a high probability of defaulting on a loan, you may not care at all what the interest rate is; whereas if you expect to pay faithfully the interest rate is much more important. Since creditworthiness is, to a great extent, private information—something borrowers know a lot better than lenders—interest rates may influence the quality of borrowers in a manner similar to the effect of wages on the quality of workers. The lesson economists draw is that asymmetric information can fundamentally alter the way markets work, diminishing the role prices play in equilibrating supply and demand.

7.7 Market Efficiency

In recent years there has been increased attention given to exactly how markets work on a transaction-by-transaction level. This is because some markets hum along like well-oiled machines, while others operate slowly and clumsily, but it is also because economic policies often call upon us to *create* markets from scratch, where they have not existed before. Proposals to auction the public airwaves or allow trading in pollution permits, to mention two examples, depend for their success on whether the markets they require can be crafted to work efficiently and well.

What does “efficient” mean in this context? How would we determine whether a market was efficient or not? First, we should be clear that an efficient market is not the same as an efficient economy. Economic efficiency is about the relationship between outputs and inputs, whereas market efficiency is about one piece of this, the extent to which markets expedite transactions, process information and achieve equilibrium. It is a more narrowly technical conception than economic efficiency, but not without significance.

Unfortunately, although the term “efficient markets” is used in many branches of economics, there is no consensus on the exact criteria that should be used to determine whether any given market should be deemed efficient or not. This means I am free to offer my own:

1. Cost. An efficient market has low overhead costs. In particular, search and enforcement costs should be minimal relative to the economic advantages available to those who participate in the market. The advent of computers, for instance, has greatly reduced the cost of operating financial markets, thereby increasing their efficiency. (What their effect on *stability* has been, however, is another story, as we will see in the book on macroeconomics.)
2. Speed of equilibration. An efficient market reaches equilibrium (as defined in Chap. 5) quickly, especially in comparison to the speed of events that dislodge it from equilibrium. In this respect as well, financial markets score highly. Even though the stream of news that upsets supply and demand is almost continuous, the adjustment is virtually instantaneous, again thanks to the computer. Labor markets are the opposite, slow and usually caught in disequilibrium. Excess supply or demand for labor can last for years, and the wage for identical jobs can vary from one employer to the next.

When markets are in disequilibrium, or divided into regions or other fragments with little cross-communication, it is common for multiple prices to appear for the same good. This creates an opportunity for **arbitrage**, which is buying a good in a market in which its price is low in order to resell it in a higher-priced market. Arbitrage is a consequence of disequilibrium, but it is also a force that tends to restore equilibrium, since any additional buying in the low-price market drives up the price to some extent, and selling in the high-price market drives it down. Antique dealers who go to estate sales in order to buy items that they resell in their shops to tourists at a large markup are engaging in arbitrage.

3. False trades. This concept is a bit more complicated. The basic idea is this: some buyers want an item (are willing to pay) more than others, and some sellers are willing to sell for less than others, but these inclinations are not visible. When trading begins, and before an equilibrium has had a chance to establish itself, transactions are more or less random: some buyers find sellers and vice versa, but others don't. There is only a weak relationship between having a greater desire to trade (high willingness to pay, willingness to accept a low price) and success at finding a suitable trading partner. Meanwhile, each transaction has its own price. Eventually the market comes to an equilibrium, with a price that registers the willingness to pay of the marginal buyer and willingness to accept of the

marginal seller. We can use this price as the benchmark, in the spirit of the Market Welfare Model.

Now ask the question, during the chaotic period during which the market equilibrium was being established, were any trades made that involved buyers whose willingness to pay was below the (future) equilibrium price or sellers whose minimum acceptable price was above it? This could happen. Suppose the eventual equilibrium price for a new strain of rice proves to be \$.20 a pound. Before this become clear, there may be trades above and below that price. In itself this is not a problem; If my willingness to pay is \$.70 and your marginal cost (and minimum acceptable price) is \$.15, even though we may exchange at \$.40 a pound, the trade itself is beneficial. Your marginal cost is less than the equilibrium price and my benefit is greater. But if we trade at \$.40 and your marginal cost is \$.25, then we have exactly the sort of exchange that an efficient market ought to prevent: once the price settles at \$.20, it will be below your marginal cost.

The underlying point, then, is straightforward: markets should encourage sales by those whose costs are less than the equilibrium price and discourage sales by those whose costs are greater; and they should encourage purchases by buyers whose willingness to pay is above the equilibrium price and discourage purchases by those whose willingness to pay is less. Since it takes time for the equilibrium to assert itself, perfect success is not at all guaranteed. Transactions that violate this stricture are referred to as **false trades**. To minimize them, markets should not only reach equilibrium sooner, they should also not depart too far from equilibrium during transition periods. To a considerable extent this is a function of the effectiveness of information flows between market participants. An example of a market that generates plenty of false trades is real estate: when buyers and sellers have spotty information about the other sales taking place within a large urban market, for instance, they are prone to make transactions they may later come to regret.

4. Utilization of information. Since markets are, among other things, information processing mechanisms, their efficiency depends on the extent to which they actually do incorporate all available information. This really has two levels, individual and collective. At the individual level, participants in an efficient market will acquire all the information that bears upon the goods being traded. As we have seen, this is called into question by the likely presence of search costs. In addition, at the collective level there should be incentives that lead better informed choices to replace poorer ones. For instance, suppose the market under examination is for next year's rice. At the individual level, if the market is to be efficient individuals must be motivated to dig up information pertinent to next year's supply and demand. But this is not all: some people will do this better than others. The ones who get the most, best and soonest information should be rewarded by the market; their expectations, rather than those of more poorly informed traders, should be the ones that set prices in the future. Suppose I am the first to observe a clue that indicates that next year's rice harvest will be below average. If the current price is \$.20 a pound, I will be happy to make a

commitment to buy at that rate next year, thinking that scarcity will make rice more expensive. If a further clue develops, more traders should come to the same conclusion. This will increase the future price to, say, \$.25 and reward my earlier instinct, for now I can already sell next year's rice that I bought for \$.20 and make an immediate profit of \$.05 a pound, even without waiting for next year to come around. That's how an efficient market *should* work: prescient expectations should drive out mistaken ones.

How can we tell, in fact, whether markets utilize information in this way? Consider a market whose participants are in the process of gathering information. At any moment their information is incomplete, but new information appears regularly. Sometimes current prices will prove to be too low, since the next batch of information shows that demand is (or will be) greater than thought, or supply less. Sometimes it's the opposite and prices are too high. Economic theory tells us that, in an efficient market, there will be no systematic bias; no one will be able to predict whether today's price is too high or low based on the previous record. If they could that would be information that someone should have utilized, and it should already be factored into the price. To put it differently, new information, if it is truly new, comes as a surprise, and we should not be able to predict it before it occurs. Since we can't predict what the information will be, we can't predict how it will affect prices.

This may seem like an arcane point, but we will see that it has important implications for the analysis of financial markets in particular. If the price of a bond or a share of stock incorporates all available information, then no one should be in a position to regularly outguess the market. Future price movements are simply unpredictable, and there is no point to paying someone for her investment advice. Is this actually true? We will look at this question in more detail in Chap. 18.

With these four efficiency criteria under our belt, we can turn to features of markets that promote or interfere with their efficiency. This is an enormous topic, and to give you a feel for it I will look at just one issue, how prices are determined for individual transactions. What follows is a small taste of the fast-growing field of **market microstructure**, the study of the rules that different markets follow in their detailed operations.

The most natural pricing rule, the one that is honored in markets and bazaars around the world, goes like this: the buyer makes an initial price offer, the seller counteroffers, and they go back and forth until they agree somewhere in the middle—or not. The technical name for this procedure is a **double auction market**; parties on both sides of the market are free to bid their offers up or down. Research has shown something useful: double auction markets are highly efficient in their speed of equilibration. If there are many buyers and sellers all bidding with each other on the same item, the market as a whole will tend to arrive at an equilibrium price quickly with few false trades. If you want a picture of this to place in your mind, think of a farmers market. Dozens of farmers arrive early in the morning, each with the morning's harvest of lettuce. For simplicity, assume that every head of lettuce is identical. No price is posted at any stall, because the farmers expect to negotiate. Next a swarm of consumers descends on them, and individual consumers

begin to haggle with individual farmers. Exchanges begin to be made, and both farmers and consumers begin to notice the emerging price patterns. Before long they are using this information to guide their own bargaining, and before the morning has barely begun all are agreeing on a common price. This is what we would expect, given the results of experiments conducted by economists.

Unfortunately, there is a downside to double auction markets: what works for individual products and face-to-face relationships can be inapplicable to transactions at a distance or over a large, diversified range of products. In fact, markets of this sort play a limited role in allocating most consumer goods in industrialized economies. Far more common is another approach, in which the seller puts a price tag on the item and leaves it to the consumer to decide whether to buy or not. This is referred to as a **posted offer market**. Such markets require much less engagement on the part of those that participate in them. The negotiating over a bag of rice may be fun (or at least bearable) in a farmers market, but it is hard to imagine how one could bargain over every item in a large store. The downside of posted offer markets, on the other hand, is that they equilibrate very poorly. Excess supply and demand are fixtures in our economy largely because we rely, as we probably must, on market mechanisms that stumble slowly towards equilibrium.

The Main Points

1. Markets are less place-based and more abstract in modern times than they were long ago. Even so, there are important social and institutional differences between different kinds of markets.
2. When the goods that trade in markets have an element of promise—commitments to do something in the future—enforcement becomes important. An exchange under these circumstances is a solved prisoner's dilemma: individuals have succeeded in cooperating, making a mutually beneficial exchange, instead of taking advantage of the other party. Bringing about this sort of cooperation can be expensive, involving writing complex contracts, monitoring and enforcing them—what economists refer to as transaction costs. Fortunately, social norms and reputation effects often facilitate exchanges of promises so that cumbersome enforcement measures aren't necessary.
3. If goods are not standardized, market transactions become costly in other ways. Effort has to go into determining and communicating the unique characteristics of each item, special forms of protection against risk may be needed, and mass production technologies are difficult to implement. An important example in modern economies is the real estate sector. For these reasons, standardization plays an important role in the development of markets. Goods can be standardized by private agreement among producers or consumers, by monopoly control, government regulation or technological change.
4. Market prices reflect the information that participants have, or think they have, about the goods being traded. We often use expressions like “the markets think....” to express this. Summarizing all the information individuals possess

in a single number, the price, had advantages and disadvantages. We lose the richness of this information—what it consists of specifically, how credible it is, etc.—but it is easier to act on. In fact, since the price has real effects on producers and consumers, it invests the information flowing through markets with economic force.

5. Information is often costly to acquire, which makes it expensive for market participants to try to maximize their gains according to the expected utility formula. Instead, they are likely to satisfice, setting a benchmark level of value or benefit and accepting the first option they come across that meets this standard.
6. A widespread, complicated problem is asymmetric information, which occurs when one party has access to information that the other doesn't. The most common case is when the seller has private information about him or herself or about complex products that could vary a lot in quality. In situations like this, buyers will interpret the seller's offer price as a signal of hidden quality. The result is that prices will not be set simply by supply and demand, but may result in excess demand or (more likely) supply.
7. Markets differ tremendously in how efficient they are, where efficiency refers to (a) low transaction costs, (b) a high speed of reaching equilibrium, (c) little initial deviation from the ultimate equilibrium (reducing false trades), and (d) maximum use of all available information (and therefore unpredictability). Efficiency is strongly influenced by the rules that govern how markets work at the level of individual transactions—market microstructure. As an example, a double auction market tends to be extremely efficient, but most consumer markets follow the posted offer model.

► Terms to Define

Arbitrage
Asymmetric information
Caveat emptor
Efficient markets
False trades
Market microstructure
Satisficing
Social norm
Transaction cost

Questions to Consider

1. One reason we don't perceive most market exchanges as solutions to Prisoner's Dilemmas is that we take the legal guarantees for contract enforcement for granted. It is instructive to consider the pressures placed on markets that are not backed by law. One example is the trade in illicit drugs such as cocaine. A buyer who received inferior or insufficient product, or a seller who is not paid promptly and in full, cannot file a claim in court. Can you create a payoff matrix for a drug deal that conforms to the structure of a Prisoner's Dilemma? How well do you think this model applies? In the absence of legal recourse, what other factors make it possible for drug dealers to conduct their business?
2. When a company adopts a policy of branding, it is making its ownership of its various products more visible to consumers. Such a company wants you to know that everything it sells is a part of its enterprise, and it wants you to be more likely to buy because you know this. In a sense, the advertising campaigns to build up brands are creating a general corporate reputation that can be applied to all the products sharing the brand logo. Does this mean that branded products are more likely to fulfill the promises made for them because of the need to preserve this reputation, or is this function undermined by the advertising strategy on which branding is based? In answering this question, it may help to think of particular branded products, comparing them to similar products sold by companies, perhaps smaller, that do not engage in branding.
3. Can you think of another good or service in our economy, besides real estate, that is highly nonstandardized? Describe the process by which purchases are made. Is it as costly and time consuming as real estate?
4. Most countries evaluate their hotels, assigning them between one and four or five stars based on facilities, service etc. A tourist can then book a hotel without knowing anything else about it, just on the basis of how many stars it has been given. What are the advantages and disadvantages of such a system?
5. Does the asymmetric information problem apply to auto repair services? If so, what is the private knowledge? Do people who need their cars repaired sometimes pay extra for reasons analogous to the "invisible worker quality" story? Explain.
6. What experiences have you had with double auction markets? Do you think the transactions you were able to make were better (fairer, more likely to produce benefits for both you and your trading partner) as a result of the opportunity for both sides to negotiate? (Try to imagine how the same transaction would transpire under a posted offer system.)