1

Equilibrium, Efficiency, and Asymmetric Information

1. Asymmetric Information **10**

2. Taxi! **16**

Problem set 18

3. Acid Rain **18**

Problem set 21

4. Efficiency **23**

Problem set 28

5. Equilibrium **30**

5.1 Dominant strategy equilibrium 31

5.2 Nash equilibrium 32

5.3 The invisible hand 33

5.4 The incentive to harbor terrorists 34

5.5 Dissolving a partnership 35

5.6 The centipede game 37

5.7 Subgame-perfect Nash equilibrium 40

Problem set 44

6. The Prisoner's Dilemma Game **45**

6.1 Economic sanctions 47

6.2 Public opinion 48

6.3 Pollution 48

6.4 Beggar-thy-neighbor policies 49

6.5 Disarmament 49

6.6 Cartels 49

6.7 Hostile takeovers 50

Problem set 51

7. Repetition and Equilibrium **53**

7.1 Repeated prisoner's dilemma with terminal date 54

7.2 Infinitely repeated prisoner's dilemma 55
7.3 Equilibrium theorem for infinitely repeated games 60
7.4 Terminal date and unknown type 64
Problem set 70

A successful institution, whether large or small, must coordinate the activities of its individual members. In this book, I examine the incentives at work in a wide range of institutions, to see how—and how well—coordination is achieved by informing and motivating individual decision makers. Incentives work well when they result in a high level of individual welfare generally. This is problematic because each individual acts to maximize his or her individual payoff, regardless of the implications for the welfare of others. In other words, we examine incentives to determine the extent to which they prevent the pursuit of self-interest from being self-defeating. We look at an entire economy, as well as a single firm in that economy. Even two-person institutions receive attention: a car owner and a mechanic hired to repair the car, for instance. In all cases, a satisfactory outcome requires coordination among the participants, and coordination requires *information transmission* and *motivation*, as shown in Table 1.1.

The individual members of the institution cannot do their part unless they receive information telling them what their roles are. In the case of a market economy, much of the vital information is transmitted by prices. In a wide range of situations, the consumer's budget constraint and the firm's profit motive give the respective decision maker the incentive to use the information embodied in prices in a way that enhances the welfare of all households. However, in many significant political and economic interactions, the relevant information has been received by individuals but they have no incentive to use that information in a way that enhances the welfare of others. If everyone chooses a strategy that benefits himself or herself a little and harms others a lot, the outcome will leave everyone with a lower payoff than the system is capable of delivering. For instance, each individual in a town knows that everyone can benefit from an Independence Day fireworks display. But there is no incentive for anyone to use this information about the spillover benefit in deciding whether to finance the display. In most towns, no individual would gain by watching fireworks if that person also had to pay the entire cost. If the decision were left to the market system there would be no fireworks. This is typically not a good outcome. If the display would cost $100,000 and there are 50,000 townspeople, then the fireworks spectacular could be produced by having each person contribute $2. In most towns, everyone would be better off if he or she gave up $2 to watch a fireworks display. Although everyone knows that there would be a high level of total benefit from the display, no individual has an incentive to act on that information. The economic theory of incentives is devoted in part to the design of mechanisms that give the decision maker an incentive to use information about spillover benefits.

In rare cases there is a natural alignment of the incentives of the decision maker and the rest of the community. For instance, the pilot of an aircraft is just as determined as the passengers to arrive safely at the destination.

Table 1.1

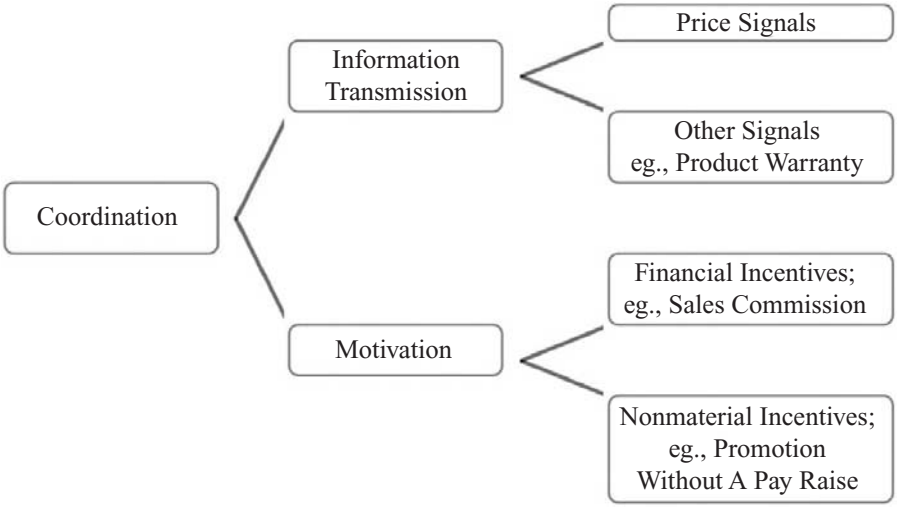

However, the welfare of an airport security guard or mechanic on the ground is not *directly* linked with that of the passengers. The passengers need to be reassured that the mechanic, say, has a strong incentive to act *as though* his or her chief concern is the passengers' well-being. With inappropriate incentives, a mechanic may succumb to the temptation to avoid hard work by doing a superficial job of inspection and repair.

In 1979 all DC-10 airplanes were temporarily grounded after one of them crashed upon takeoff. The crash was caused by a crack in one of the engine attachment assemblies. The crack resulted from the way that the engine was replaced after servicing. It had been reinstalled in a way that was not recommended or even anticipated by the plane's designer. Reattachment was henceforth done with special care (Petrosky, 1992, pp. 95–6).

Incentives are obviously of vital concern to air travelers and are worth studying for that reason alone. But they are also vital to society as a whole. Given the decisions made by others, a worker—whether a mechanic or professor or company president—may find it in his or her interest to expend little effort on the job while drawing a full salary. If a large fraction of the labor force can get away with shirking, the economy's output of goods and services will be greatly diminished and per capita consumption will be very low. In that case, each worker will wish that *everyone* had been prevented from shirking, to enable each to consume less leisure but more produced goods and services. *The pursuit of self-interest is self-defeating in this case.* A more appropriate system of incentives could have prevented this—making everyone better off, even though each individual is maximizing his or her own welfare *given the decisions of others* when everyone shirks as a result of poor incentives.

Appropriate incentives are crucial to the success of any institution, whether large or small. This book examines incentive environments and evaluates each in terms of its ability to promote individual welfare generally. In most cases, the pursuit of self-interest can lead to a high level of individual welfare generally only if the individual taking action incurs a cost equal to the cost that his or

her action imposes on the rest of society. We refer to this as *social cost pricing*. Here is an informal explanation of why social cost pricing works: Let U_i be the payoff (or utility) to individual i, who will act to maximize U_i. This will typically affect the payoffs of others, and we let C_i be the total decline in the payoffs of everyone but individual i, resulting from i's decision. Then C_i is the cost that i imposes on the rest of society. We modify the rules of the game so that the payoff to i is now $U_i - C_i$, which is what individual i will now maximize. But the change in $U_i - C_i$ is obviously equal to the change in the sum of the payoffs of everyone in society, including individual i. By imposing a cost on individual i equal to the cost that i's actions impose on the rest of society, we induce individual i to act to maximize the total social payoff, even though i is only directly interested in maximizing his or her own payoff.

Four hundred people died in January 1996 when the Indonesian ferry *Gurita* sank. The boat sailed even though the captain knew that the cement that had been used to patch holes in the hull had not dried. A government official had ordered the captain to sail or lose his job.

DEFINITION: ***Social cost pricing***

An institution uses social cost pricing if each decision imposes a cost on the decision maker that is equal to the total cost incurred by the rest of the group as a result of that decision. If there is in fact a net benefit realized by everyone else then the decision maker receives a reward equal to that net benefit.

In many situations individuals must be sheltered from uncertainty if high levels of individual welfare are to be achieved. Full social cost pricing then would leave maximum exposure to risk or uncertainty. In other words, in the presence of uncertainty, incentives have to be less than fully efficient, to allow for insurance.

We look at incentive schemes currently in use, and we also consider the prospects for designing superior schemes in particular situations. The starting point is the realization that, although the decision maker's actions affect the welfare of a wider group, the decision maker has private information that is not available to members of that wider group—nor to a representative of the group, such as a government agency—*and* that the decision maker will act to maximize his or her payoff, without taking into consideration any resulting side effects on the other members of the group. For example, the manager of a factory has much better information about the production process and product quality than the firm's consumers or the residents of the neighborhood in which the factory is located. If the government attempts to regulate the firm—to affect product quality or the emission of toxic waste—it can do a much better job if it taps the manager's private information instead of issuing direct commands. If the government orders each factory to modify its production process in the same specific way, it may achieve the desired level of pollution abatement. However, it will usually be possible to achieve the same pollution reduction at a lower total cost in resources that have to be diverted from other uses by having the individual factories adjust in quite different ways, depending on their specific input

requirements and production technologies. Doing so requires the provision of incentives to harness the factory manager's self-interest and inside information. We refer to this as *incentive regulation*, and it is coming into increasing use, replacing the old *command and control* approach.

Transmission of information goes hand in hand with incentives. Market prices have their limitations as conduits of information, but they do a superb job in a wide range of situations. For example, wages are important determinants of individual career choices, and wages contain information about the value of various skills to all consumers. An occupation will command a high wage if it contributes significantly to the production of highly valued goods and services. That's because the high demand for a consumer good translates into high prices and profit for the producer. There will be great demand for workers who are crucial to the production process because they generate substantial revenues for their employers. The high demand for these workers leads to high wages. Competitive bidding in the labor market raises the wage of the most productive workers above that of other workers. A particular wage signals information to the economy as a whole concerning the value of the associated skill. We not only acquire the information that a particular occupation is valuable to consumers as a whole; at the same time, an individual has a strong incentive to take this information into consideration in choosing a career, because higher wages provide more income and thus more consumption opportunities.

In general, the way prices enter our budget constraints gives us the incentive to use the information embodied in those prices. All individuals maximize their own payoffs, but because the prices embody information about the welfare of others, the pursuit of self-interest induces individuals to take the welfare of others into consideration, without realizing that they are doing so.

Information transmission and motivation do not always go hand in hand. Commuters know that traffic is congested during rush hour. If individual driver A joined a car pool, other drivers would benefit from the reduction in the number of cars on the road. But the benefit to A is slight, and A's own welfare would decrease because of the inconvenience of not having his or her own car. Self-interest leads all motorists—well, almost all motorists—to drive their own cars to work. It's plausible that if everyone joined a car pool the improved traffic flow would leave everyone better off, *net* of the inconvenience of carpooling. As it is, everyone knows about the social value of carpooling but no one has an incentive to act on that information. However, information technology now allows municipalities to charge for the use of designated high-speed lanes. Such lanes remain uncongested because their user fee gives motorists for whom time is relatively less valuable the incentive to use the lanes that are free but more crowded.

Information transmission can be more or less costly. Low-cost information transmission is problematic. If the institution is the entire economy, the delivery of information throughout the economy can be exceedingly costly. For one thing, contracts must be enforced, and legal costs can be very high. Prices transmit information at low cost but, as Table 1.1 indicates, other devices such as warranties are important. An extensive warranty on a manufactured

"More than 2000 television sets a year exploded in Moscow alone" before the collapse of the Soviet Union (Milgrom and Roberts, 1992, p. 13).

good is a signal that the manufacturer believes that the likelihood of a defect is small. If an entrepreneur set out to deceive customers by manufacturing low-quality television sets and passing them off as high-quality sets, he could not offer a good warranty without losing the profits that his deception was designed to yield. He would know that a very high number of sets would be returned for replacement or repair under the warranty. Competition between manufacturers in a private ownership market economy induces each producer to make a high-quality appliance and offer an extensive warranty.

Even when the information transmission problem is solved, the motivation problem remains. As with the highway congestion example, there must also be an incentive for the individual to use the information in a way that promotes the goals of the institution—a high level of welfare by commuters generally, in the case of the traffic example. Incentives are essential because individuals' paramount concern is their own welfare, not the welfare of others. This book is devoted to the study of material incentives—incentives that have their impact on the decision maker's welfare through their impact on his or her consumption opportunities. How can they be designed to harness self-interest and prevent the pursuit of self-interest from being self-defeating?

An automobile repair shop illustrates nicely how incentives will come into play in this book. A car owner who brings his car to the shop for repair wants a reliable job done at low cost. He has neither the expertise nor the time required to monitor the mechanic. If the car owner suspects that the mechanic has cut corners he is likely to broadcast his suspicions to acquaintances. This implicit threat, along with the existence of other repair shops competing for business, gives the owner of a garage some incentive to ensure that the repairs are well done and that customers are not overcharged. But how does the garage owner motivate the mechanic that she employs? Competition and reputation effects may give the right incentives to the owners of firms, but they are just part of the solution. The owner—in general, the *principal*—now has the problem of providing appropriate incentives to the *agents* (mechanics) that she hires. We attempt to solve this problem—with considerable success. The private ownership market economy is very sophisticated when it comes to generating devices for solving these principal-agent problems. But there are serious limits to the ability of *any* institution to overcome incentive difficulties in many situations. The difficulties are compounded by the presence of random effects. If the car breaks down a week after it was repaired, should that be attributed to shirking on the part of the mechanic or to bad luck?

In World War II the United States won the race with Germany to develop the atomic bomb. Computers were not available, of course, and the United States depended on a team of high school graduates to do a staggering amount of calculating. The productivity of the calculators increased almost tenfold when they were told what they were working on (Gribbin and Gribbin, 1997, p. 97).

Although this book is almost exclusively concerned with material incentives, we acknowledge that nonmaterial incentives play

a role in any institution. In one of the first influential articles on the modern economics of information, Kenneth J. Arrow (1963b) noted that the information advantage possessed by physicians in treating their patients has led to the emergence of institutions based on trust and delegation to supplement market incentives. Hence, the code of medical ethics.

Each of us does things that benefit others, at some personal sacrifice. Nevertheless, we employ a model that assumes that each individual always pursues narrow self-interest. One reason for doing so is that we are alerted to potential difficulties if our model results in low levels of individual welfare generally. Moreover, we are much less likely to recommend policies that are naively utopian when we work within this framework.

As early as 1931 the Soviet ruler Joseph Stalin deviated from the egalitarian wage ethic, realizing that a high level of economic performance could not be achieved without material incentives. The opportunity to work for the common good was not sufficient motivation (Laffont and Martimort, 2002, p. 23).

The importance of incentives has been documented in many ways and in many contexts, although the specific contractual form derived from economic theory is not always reflected in contracts as actually written (Chiappori and Salanié, 2003). For the specific case of the relationship between a tenant farmer and the landowner, Allen and Lueck (2002) show convincingly that incentives are central to understanding the nature of the contracts that are employed.

In 1896 South Carolina enacted a law levying a fine on any county in which a lynching took place. No county that had been fined for this abuse ever had a second lynching (Dray, 2002). Lynching of African Americans by white mobs was common from the late nineteenth century until the middle of the twentieth and was one of the many devices by which African Americans were terrorized.

One measure of the importance for public policy of a formal study of incentives is provided by McAfee and McMillan (1988). They estimate that switching to appropriate contract design could reduce government costs by at least 8%, and sometimes by as much as 30% (p. 149). The switch to the *responsibility system* in Chinese agriculture in the 1980s resulted in a remarkable increase in productivity over a short period of time (McMillan, 1992, pp. 96–8). The responsibility system requires each farm to deliver a fixed amount of output to the state, but the farm keeps the proceeds of all output above this quota. This is an example of social cost pricing: The social cost of the farmer's leisure consumption is the output that society loses when the farmer consumes an hour of leisure. But that is also equal to the cost imposed on the farmer under the new system because the farmer would have been allowed to keep the harvest from that hour of labor. Under the old system, the cost to the farmer of an additional hour of leisure consumption was zero because all of the output from an additional hour of labor goes to the state. It was the farmer whose return was fixed.

Another reason why we assume selfish behavior at every turn is that, although it abstracts from important features of the real world, it gives us a simple model with a lot of explanatory power. We have come to accept abstract models in everyday life and should not be reluctant to employ them in economics. A road map, for instance, is a representation of a particular region. It abstracts from

almost everything that is important—scenery, the location of shops, and so on. Because it *is* so abstract it is very easy to use to work out a route from one location to another; it can even be used to compute a short route. Similarly, an economic model can be exceedingly abstract and still allow us to determine the effect of an excise tax on a commodity's price or the nature of a salary contract that will be offered when the employer can observe the quality of the employee's work but cannot validate that observation with evidence that would be credible to a third party, such as a judge. Conclusions are drawn from abstract, formal economic models via theorems.

Many people are impatient with economists for abstracting—and worse, employing assumptions that are at odds with reality. It may comfort you to know that this is standard practice in physics. It can even be useful for a physicist to assume that a cow is a sphere! (See Krauss, 1993, pp. 1–7.) "The set of tools physicists have to describe nature is limited. Most of the modern theories you read about began life as simple models by physicists who didn't know how else to start to solve a problem. . . . *Before doing anything else, abstract out all irrelevant details!* . . . Overcoming the natural desire *not* to throw out unnecessary information is probably the hardest and most important part of learning physics" (Krauss, 1993, p. 4). The classical model of the motion of the planets around the sun assumes that the mass of each planet is concentrated at a single point of zero breadth. That's absurd. Nevertheless, the model is extremely useful. It was used to predict the existence of the planet Pluto, for example, which was discovered in 1930.

We begin then by assuming that all individuals evaluate outcomes exclusively in terms their effect on their own well-being. This allows us to work out an individual's response to a change in the incentive environment. The assumption of selfish utility maximization implies that there *will* be a response. Not everyone is able to grasp this point. For example, a lot of people argue against long prison sentences for drunk drivers who kill or maim others: "It could happen to anyone." Well, wouldn't you make sure that it couldn't happen to you if a long prison sentence were the penalty for drunk driving? To adapt a phrase of Dr. Johnson's, the prospect of a long jail sentence focuses the mind wonderfully.

Public drunkenness is not uncommon in Japan, but drunk driving is very rare because of the severe penalties. A professional person can even be disqualified from practicing if convicted of driving while intoxicated.

We examine incentives at work to see whether we can expect outcomes that maximize individual welfare generally when individuals are motivated by selfish considerations. In each case study we assume that an individual takes whatever available course of action leads to the highest possible personal benefit for himself or herself. Of course, in real life there are situations in which some or all individuals behave altruistically, at least up to a point. But self-seeking behavior is pervasive enough to warrant independent study, particularly when the economy as a whole is our concern. Therefore, our goal is to work out the implications of self-motivated behavior, by means of examples and theorems, and we try to learn from them without being distracted by the many real-world features that

are left out of the models. We discover that the need to provide individuals with socially beneficial incentives imposes constraints on the economic system as a whole, forcing us to make trade-offs. For instance, giving individuals an incentive to truthfully reveal their preferences for public goods leads to government budget imbalance. By identifying such trade-offs we can design better public policies. In particular, we won't waste resources trying to accomplish goals that are mutually exclusive.

Links

McMillan (2002) is a superb but non-technical account of how, and to what extent, markets can provide the incentives that lead to a high standard of living. The role of the CIA in supplying the Bush administration with evidence of Iraq's weapons of mass destruction, prior to the invasion of March 2003, is a reminder that the performance of an organization is a function of worker and management incentives. See *The Economist*, July 15, 2004 ("The weapons that weren't"). Baumol (1993) contains many examples of entrepreneurial responses to incentives, some of which reach back to ancient Greece and Rome. See Sappington (1993), Laffont (1994), and Sappington and Weisman (1996) for further discussion of incentive regulation. Stiglitz (1993) has a good discussion of the limits of prices in transmitting information. See Chapter 4 of Baumol (1993) for examples of the costs of contract enforcement.

Problem set

1. The services of garbage collectors have far more total value to the community than the services of heart surgeons: Compare a world without garbage collection—plagues, low life expectancy, only 50% of children surviving to the age of five—to a world without heart surgeons—no appreciable difference in life expectancy. But heart surgeons are paid far more per hour than garbage collectors. What information is being signaled by this wage rate differential?

2. I drive a car made in 1990, before air bags became mandatory in all cars sold in the United States. I could buy a safer car—a new Mercedes Benz, for example—but I prefer a basket of goods and services that includes my present car and an annual vacation on the ocean to a basket with a safer car but an annual vacation consisting of croquet in the backyard. Is it in society's interest for firms to devote enough resources to the production of consumer goods to ensure that there is absolutely no chance of a defective product injuring someone?

3. A barber will not stay in business long if he gives bad haircuts. Competition among barbers ensures that each attempts to build a reputation for high-quality service. What about an industry in which problems do not show up until long after the commodity has been purchased—housing construction, for instance? Is there a role for some form of government regulation in these cases?

1 ASYMMETRIC INFORMATION

When you hire a taxi you are employing an *agent* to carry out an assignment. You, the *principal*, want to get to your destination quickly and at low cost, but the taxi driver wants to maximize his revenue. The driver appears to have an incentive to overcharge, and your ability to monitor this is very limited because you know very little about traffic patterns and expedient routes, especially if you are a visitor to the city. This is an instance of a *hidden action* problem. The passenger cannot directly determine if the driver has acted in a way that minimizes the travel time.

DEFINITION: ***Hidden action problem***
A principal hires an agent to carry out a task, but it is impossible or extremely costly for the principal to monitor the agent.

In Section 2, we demonstrate that the conventional taxi fare schedule induces the driver (the agent) to choose the route that the principal (the passenger) would select if the principal had as much information about routes and traffic patterns as the agent—even though the principal in fact has very little information, and the agent knows it. In general, we investigate the possibility of providing appropriate incentives to agents to induce them to behave in the way the principals would instruct them to act if the principals themselves possessed the relevant information—even though the principal is in fact unable to monitor the agent, and the agent knows this. There are three reasons why the principal may want to employ an agent: The agent may possess a skill that is particularly appropriate to the task at hand. (I hired a specialist to remove a tree that had fallen over my driveway during a storm.) The principal may not have the time to carry out the task herself. (I sometimes eat in restaurants where the chef is not as good a cook as I am.) Finally, even if the principal and the agent are "twins," economies of scale can justify the delegation of some tasks by one to the other.

Providing the agent with an incentive that is optimal from the standpoint of the principal requires us to choose the incentive scheme that maximizes the principal's payoff subject to constraints. These constraints embody the notion that agents will act to maximize their payoffs subject to the incentive scheme governing their behavior, and the notion that agents have alternative job opportunities and hence must do at least as well working for principals as they would in the next best alternative. These constraints on principals' choices of incentive schemes result in principals achieving lower payoffs than if principals had all of the information available to agents and could simply instruct agents to carry out the actions that maximized the principals' payoffs. In some cases, the principal's loss is the agent's gain, but in other cases the constraints result in a net loss to the principal-agent duo. If there were no asymmetry of information, the