

Economics is not sociology, psychology or politics, but it relies on assumptions about society, mental and emotional processes, and the political and legal environment. Until recently, however, these assumptions didn't come from the other disciplines which take them as their fields of study; instead, they were largely inherited from the eighteenth century worldview out of which Adam Smith and his followers fashioned their early renditions of economic theory. That is to say, they reflected the prejudices of the Enlightenment in England around the time of the American revolution. They are rationalist, individualist and concerned with subduing nature for the greater benefit of civilization. In this chapter we will look carefully at several of the most important conceptual building blocks, explaining exactly how they appear in modern economics and subjecting them to critical scrutiny.

3.1 Choice and Exchange: Metaphors for Economic Life

Think about a day in the life of anyone taking part in an economic system—a day in *your* life, perhaps. A list of economic activities might include:

- which consumer goods, like toothpaste or breakfast food, you use, and how much
- which household goods you use, and which you leave for others in your household
- commuting to work and getting there on time (or not)
- working
- studying and attending classes
- paying bills
- shopping
- housecleaning
- raising children

- borrowing money or buying financial assets as investments
- quitting old or accepting new employment
- searching for goods, jobs, housing or other items, without necessarily buying, renting, enrolling, etc.

All of these activities are economic in the sense that they involve the production, distribution or use of goods and services. If we added up this entire list for everyone in our society, the sum would look very much like “the economy” as a whole.

This is how the real world is, but so many qualitative differences make it impossible to do the sort of systematic analysis economics aims at. Somehow, we must simplify. There are many ways to do it, but economics selects just one: it treats every form of economic activity as a *choice*, and every economic interaction between people as an *exchange*. These are such shocking simplifications, with such profound implications for all that follows, that we need to examine them very carefully.

Choice and exchange in their economic usages are metaphors. A metaphor draws our attention to some aspects of a complex phenomenon by referring to something else that shares them. For instance, one famous metaphor is: Time is a river that flows endlessly through space. Time is a difficult concept to wrap one’s mind around, but this metaphor does help somewhat. It points to the one-way movement of time as we experience it, and the water imagery reminds us that time can “carry” things along with it. On the other hand, time is not at all a river in other respects, nor is a river time. Time is not made up of a physical substance the way rivers are made up of water. Moreover, rivers can be dammed or even have their direction of flow reversed; try that with time! The point is that metaphors are helpful as long as we remember their limitations.

What about the metaphors used in economics? First consider choice. Much of what we do in the economy does involve choosing: we choose where to work, where to live, and paper or plastic in the check-out line. No doubt many of the choices we make are unconscious, but it might not be too far off the mark to think about them as if they were conscious and “rational” as we will describe in the following section. Nevertheless, the metaphor of choice can be misleading in some instances. There are two reasons for this.

First, many of the actions we undertake are governed by a process very different from conscious choice. In fact, quite often choices are made for us by others. Many goods are consumed institutionally, for example, such as lunch offerings in a work or school cafeteria. Often one member of the household purchases goods for other members. We are also subject to pressure, sometimes intense pressure, from people we are close to over questions like employment, major purchases, etc. In some cases the pressure comes from society in general, through judgments of what is fashionable or signals high or low prestige. In these cases, treating an economic decision as if it were a free individual choice may be a mistake. We will discuss this problem in greater detail in the chapter on consumption.

Second, many activities are not choices at all. You can choose whether to buy white or red potatoes, but cooking the potatoes is an act of (household) production, not a choice. Working, doing the actual tasks that make up a job, is not choosing; it is working. Spending days or weeks searching for a new house is not making a choice; it’s doing a search. Of course, subject to the qualification we made in the previous paragraph, all these activities lead up to or follow from a choice. In other

words, what the metaphor of choice is telling us is that what is deemed important about any economic activity is the element of choice connected to it. This is a simplification of great power, because it enables us to make general statements that apply to the many aspects of life through their common element of choice, but it downplays the economic importance of the non-choice element. For example, as we will see, economists until recently have reduced the experience of work to the moments at which a worker chooses to begin or quit a job. Other aspects of work entered in only to the extent that they provide information used in making this choice. Real jobs, however, are usually social situations in which individuals interact continuously; they have communication and power structures that play an important role in both the productivity of work and its impact on the worker. To use the metaphor of choice is to direct attention to the discrete moments in which workers decide whether to join or quit; it directs attention away from the ongoing interactions on the job.

Now consider the other metaphor, exchange. This one is, if anything, more elusive. In the view of economics, the market is the primary mechanism that governs economic life, and markets are the place where buyers and sellers exchange with one another. We think we know what an exchange is: I have something you want, you have something I want, and we exchange with each other. If it is a freely chosen exchange, each of us will be at least as well off as before, and at least one of us (the one who initiated the exchange) will be better off. Add that up over millions upon millions of exchanges each day in a large country like the United States, and you appear to have a recipe for continuous economic improvement. In addition, exchange has the added benefit of being fairly simple in structure and therefore not too difficult to analyze. A provides B with something; B provides A with something else. It happens in an instant, and then it's over. It can be fully described by listing the parties to the exchange, what they provided each other, and the exact place and time the exchange took place. As you might expect, it is easy to translate this into the language of mathematics.

The preceding discussion of choice should have already alerted you to a potential drawback with exchange as a metaphor for economic relationships: economic interactions can take place before or after moments of exchange, and something important is lost if we don't incorporate them. This is true of work relationships, household relationships, landlord-tenant relationships, and so on. A more complete economics (which doesn't really exist yet) would combine the aspect of exchange with the institutional, cultural and social structures that bring people together in economic life.

There is additional wrinkle, however. The simple model of exchange, instantaneous and unambiguous, is not characteristic even of most transactions economists call exchange. How can this be? I believe that here, as in so many other areas, the root of confusion can be found in language. We use the same word, exchange, to refer to many different types of transactions, similar in some respects but different in others. An exchange of the simplest type is a simultaneous trade of one good for another that completely exhausts the transaction. But most real-world transactions differ at least a little from this stylized type, either because they take place over a

substantial period of time, or because the terms of the exchange are not simple “things” that can pass from one hand to another, or both.

Consider, for instance, the employment relationship. Economics speaks of labor markets, in which workers sell their labor, firms purchase it, and the two parties agree on the price (wage). But what do workers actually sell? They can’t sell their labor in the sense of divesting themselves of it, since it is inseparable from them. In a sense, they are selling a promise to submit to the authority of the employer, but this is hardly an open-ended promise, and should they violate it there is usually no recourse for the employer other than severing the contract. But more fundamentally, how can any human being sell his or her own *future* submission? At each instant in time we remain not only capable of choice, but *forced* to choose what to do with ourselves. Whatever the formal trappings, short of selling oneself into slavery there is no way a worker can suppress his or her future freedom of choice in return for money. Of course, employment contracts are based on this fiction, and because they are people normally try to adhere to them as best they can. For our purposes, it is enough to point out that it is simply impossible to swap today’s money for tomorrow’s obedience in the same way we might swap used paperbacks in an instantaneous exchange at a flea market.

You might think that at least consumer markets display the simpler form of simultaneous exchange. In a sense they do, but in another sense they don’t. Certainly something like an instantaneous exchange occurs when a good, like a pair of blue jeans, exchanges for money at a precise moment in time at a store’s cash register, but there are aspects to this transaction that linger on into the future. On a mundane level, the consumer has the option of returning the blue jeans under certain conditions, voiding the sale, and the store owner faces the possibility that a check might not clear. Beyond this, however, more subtle maneuvering may be taking place. Most consumer goods like blue jeans are now branded; they are sold under a well-known brand name as part of a larger strategy to increase the sales of the brand in a variety of markets. Each purchase is a moment in a larger effort on the part of the brand to expand its share; so the price might be set, for example, to influence future behavior and not simply to make money at one point in time. To treat the transaction as if it were a single disconnected moment might be to miss the forces that really determine what goods are offered to the market at what prices.

For another example, consider an individual who takes out a loan from a bank in order to start a business. There is a sense in which the making of such a loan is an exchange according to the simple version used by economics: one party, the borrower, receives money in the present and offers in return a promise to pay back the original amount plus interest at some point in the future, perhaps putting up some of the business assets as collateral. (That promise, incidentally, is very much a “thing”: it is a piece of paper, and the bank can sell it to interested buyers just as it can sell anything else it owns.) But unlike a “real” exchange (such as the exchange of paperbacks), the process does not come to an end at the moment the loan is made. The bank may have legal rights to monitor the way the new business is run. It may be able to veto business decisions it doesn’t agree with, and it may have the right to shut down the business when it thinks it is in danger of not being repaid. This raises

a host of difficult issues: just what rights should the bank have? What control should they have over the collateral? Under what conditions should they be free to call in the loan before its term expires? These questions have to be resolved somehow, and they may even be negotiated by the two parties (although often they aren't), but it would be stretching things to describe the entire process as an "exchange". There is clearly an *aspect* of exchange in a credit relationship, but there are other aspects as well.

One way to highlight the role of exchange as a metaphor for market activity is to imagine other possible metaphors. Perhaps the best-known alternative is the notion that markets are the site of *combat* between competing interests. Joseph Schumpeter, the Austrian-born economic theorist, was particularly fond of this way of thinking about markets, and his work is suffused with images of war and contest. According to this view, the most important thing that markets have in common is not consensual exchange but eat-or-be-eaten competition. Businesses are like armies: within their own walls they mobilize for combat, in the outside world they throw all they have at their competitors. To the victors go the spoils: profits and market domination. The losers must tighten their belts; they may even be driven out of the market altogether. Schumpeter scoffed at the notion that markets are efficient or even orderly. The thrust and parry of the marketplace does not tend toward an equilibrium, he argued; markets are constantly assaulted by new business initiatives designed to change supply and demand, not merely adapt to them. While the competitive process, in Schumpeter's view, generates waste and dislocation, it more than makes up for it by spurring entrepreneurs to discover new, more productive ways of meeting society's economic needs. We could spend many more pages investigating the strengths and weaknesses of the Schumpeterian approach to economic analysis, but for our purposes there is a much simpler point. The market mechanism as a form of warfare is as plausible a metaphor as that of market "exchange"; in fact, business schools are more likely to adopt the combative than the consensual interpretation of markets.

Another famous metaphor can be found in the writings of Karl Marx. For Marx, at the heart of any economy is labor. It is the work that people put into goods and services that gives them value. Even the machines that are used to improve the productivity of labor impart value only because they themselves are the product of previous work. Thus the economic system can be seen as something like a hydraulic network of pipes and reservoirs. Labor flows from workers to commodities. Commodities are sold for money, some of which is returned to workers, who then buy back a portion of their aggregate labor in the form of consumer goods. The difference between these flows—value created by labor and value returned to it—accumulates in "holding tanks" called capital, where it can be used to construct new equipment, employ new workers or simply gratify the desires of those who own these pools of value. The entire analysis is built up from this metaphorical account of the flow and accumulation of value (labor time). Nothing much is changed by the social rituals that accompany these flows; the fact that employment agreements are contractual (take the legal form of exchanges) is nearly incidental. You could do the same type of analysis if labor relations were dictated by a court or government

agency. Here the metaphor built around the flow of labor value calls attention to aspects of the economy that are obscured by an emphasis on exchange, just as the metaphor of exchange highlights aspects (the role of worker choice) that Marxists tend to underplay.

It may be that, as linguistic creatures, we have no way to avoid the use of metaphor altogether. Nor can we switch every few moments from one metaphor to another without disrupting our train of thought. There is nothing fundamentally wrong with beginning with the metaphors of choice and exchange, so long as we keep in mind their limitations and remember to correct for them when they become glaring. Make a mental note that economics rests on metaphor; we will return to this fundamental insight many times.

3.2 Psychology: The Assumption of Rational Self-Interest

One of Smith's contemporaries was fellow-Scotsman David Hume. Hume elaborated a philosophy in which the human mind is divided neatly into two parts, the faculty of reason and the raw, uncontrollable force of emotions. Reason can tell us *how* to do something but never *why*, in the ultimate sense of "to what end?" Emotions alone provide the motive force of human behavior, the goals to strive after, but emotions are not governed by reason. Hence there is a sort of means-end dichotomy, where emotions, which are beyond the power of philosophy, provide the ends and reason the means.

This is also the framework for what is still the predominant economic view of psychology. People are said to have preferences which reflect their emotional makeup, but they have no control over these preferences; they simply are what they are. They are taken as given by the reasonable person, who then devotes his or her intellectual resources solely to the task of satisfying them. To be rational is to satisfy your preferences as fully and efficiently as possible.

To make matters easier, economists assume that all preferences adhere to self-interest. There is much confusion in the writings of economic methodologists on this issue. The question as it is most often put is this: by definition, do people ever act against their self-interest? After all, if I choose to do one thing rather than another, it must be because it is what I "want" to do; therefore it satisfies my own desire and responds to my self-interest. Even diving into an icy river to save a total stranger who is drowning is what I want to do if I do it and is therefore in my self-interest. Even doing this on an impulse when I don't know how to swim is acting in my self-interest. Putting aside all the hair-splitting, the matter comes down to a question of definition. If we define self-interest as whatever motivates me to do something, then the concept is purely tautological: it is true by definition. In that case, however, it doesn't tell us very much. The only interesting propositions are those that might be false; then there is a real question being answered. So in this case the relevant question is whether people are usually self-interested in another sense—whether they put their personal comfort and convenience ahead of the desires of others. This is an empirical question: we could examine the evidence,

and the answer we get might be different for different people, different cultures, different times.

In point of fact, economists nearly always assume that people are self-interested in this second sense as well. They are thought to respond to what they, individually, will get out of a course of action and to pay no attention to the consequences for others. Because of this, the basic unit of motivation in economics is the **incentive**, the extra reward or harm that will come to an individual for making a particular choice. Economics is typically the study of incentives: what they are in any given situation and how people can be expected to respond to them. Most economists do not think it is very effective to persuade people through appeals to religion or morality; what works is to give them the proper incentives that will lead them to do what you want them to do.

A good example is crime—for instance, car theft. Most of us think about crime in moralistic terms. Certain acts are wrong and people should not do them. The purpose of punishment is to achieve a measure of justice, in the sense that it would be unfair if people benefitted from acts like car-theft and were never at risk of any loss, even if caught. If it is unacceptable for people to benefit themselves by violating the rights of others, then either they should be prevented from doing so in advance, or they should be punished in some way after the fact. Different people have different conceptions of justice, but generally they involve some balancing of the “badness” of the violation with the degree of retribution. Nevertheless, the main way we would expect to reduce crime is through education and social influence. Children should be taught not to steal from one another, and adults who commit theft should be subject to public disapproval. A society in which many people *want* to steal is already in trouble, whatever measures are taken to control them.

There is, however, a field called “the economics of crime”, which is the application of standard economic concepts to the issues raised by various forms of crime. In the case of car-theft, the economist reasons like this: All people are self-interested, including potential thieves. They might choose to steal cars because of the prospect of personal gain, particularly the money they can make by stripping (“chopping”) and reselling them. Thus, in the absence of any policy, they have an *incentive* to do this. An effective crime deterrent policy would be one that erases this incentive through some combination of greater resources spent in policing and prosecuting and greater penalties for those found guilty. The economic problem is one of devising the most efficient mix of these measures and choosing the most efficient tradeoff between the cost to society of having cars stolen and the cost of measures to change the incentives to steal them. Justice, in the conventional sense, has nothing to do with it, nor is any consideration given to the cultural and social factors that might make people want to steal cars in the first place. The entire problem is one of incentives: incentives for people to steal or not steal, incentives for society to spend money to prevent theft or absorb its costs.

In this example we can see that the assumption of self-interest is crucial. Since the economist assumes that all people are self-interested, the only policies available are those that alter the calculations people will make; criminals will have to consider the risk of doing time as well as the rewards of driving off with a new

Mercedes. Nevertheless, in most societies crime of this sort is not commonplace because most people are *not* self-interested in this narrow sense. Most people will not steal things just because they think they can get away with them; they will be inhibited by an inner voice that reminds them that stealing is morally wrong, or that gets them to see the situation through the eyes of the person whose car is stolen. In fact, some people may steal even though it is *against* their self-interest in the personal-comfort-and-convenience sense we are using. They may be hostile toward the individual or group they are stealing from, or may be acting in revenge of some earlier deed. Both possibilities point to limits to the economic approach to crime.

At the same time, we should not be so quick to reject the self-interest hypothesis altogether. Even if many people are not self-interested about crime, certainly some are, and altering their incentives may be a reasonable way to frame public policy. Moreover, few people are completely non-self-interested, and so an incentive-based approach may work up to a point even if it is not the complete answer. This is not a question that can be decided at the level of theoretical abstraction; the role of self-interest has to be looked at on a case-by-case basis, with an open mind.

3.3 Rationality and Uncertainty

It would seem that we have now tackled the hard part of the economic conception of rationality, and that the rest, the use of reason to pursue interests (self or otherwise) should be fairly straightforward. Not so. In fact, economists have a very precise notion of what it means to be a rational individual, and this has produced a complex and fascinating debate. In the coming pages we will investigate this concept of **rational choice**— what it requires, what it tries to explain, the limits it is unable to overcome.

Let's begin with a situation that calls for a decision. It's five o'clock: what should I make for dinner? Now, what's interesting here is not what I will actually make (if that's what you'd rather read about right now, I'd suggest putting this text down and picking up a cookbook), but how I will make my choice. First, we might take note of the fact that there is nothing preordained about eating a large meal in the early evening. Other cultures do this differently, and this suggests that **cultural norms**, which I will discuss in a later chapter, have a role to play. Second, I might not think about what to make at all, because I had already planned this meal some time ago. I went to the store, bought the ingredients and solidified my intention of cooking a particular dish. While I am perfectly free to reconsider this plan, of course, I might be on automatic pilot and follow the plan without thinking about it. Third, I might feel an urge to make something I have the ingredients for but haven't eaten in a long time, or that reminds me of friendly dinners I've had in the past. In other words, I might act on impulse. Or fourth, maybe I make the same thing every other Tuesday, and here I am: it's that other Tuesday and why change now? In this case I'd be a creature of habit. What all of these approaches have in common is what they are not: they are not rational in the sense that economists use the term. Some of them are unconscious—I am not aware of making any choice at all. Others

may be governed by parts of my mind that are not particularly rational, such as my love of familiarity and repetition, my urge to imitate those around me and follow their cultural norms, or my sudden craving for mashed potatoes (which may vanish after the first forkful).

To be rational, on the other hand, is to calculate the costs and benefits of each course of action in light of a clearly defined set of goals. The goals in this case might be some combination of keeping to a healthy diet, eating something that tastes good, and using up the vegetables at the bottom of my refrigerator that are about to nurture new life forms. Whatever they are, I should, if I am rational, think through all my options and try to anticipate how each will help meet these goals. I must be fully aware of my choice (not unconscious) and must not give in to passing emotions that interfere with this systematic calculation.

This would be relatively easy if I knew with perfect certainty exactly what outcome would arise from every option. If I knew in advance exactly how each dish would come out, how I would feel about having made it a week later, whether anyone else was likely to drop by unannounced, how long the leftovers would keep, etc. I would be capable of a high degree of rationality without much effort. Unfortunately, most choices in life must be made under conditions of **uncertainty**. We have some knowledge, but we don't know everything. We can say that, even if we can't predict the future, we have a reasonable idea of which outcomes are possible and which are not. If I boil a pound of potatoes for myself and eat them, I will not be hungry 15 minutes later. The point is that we can, if we think things through, reduce the infinity of potential outcomes for each course of action to a relatively small number of outcomes worth paying attention to. What we don't know is which outcome will occur for sure.

To make my example more specific, let's suppose I want to make a salad if the dinner is just for me, but I would rather have prepared a pasta dish if friends come over to visit. So we will focus on just two options, salad or pasta, and two possible outcomes, one if I eat alone, the other if I eat with friends. To make things even easier, let's summarize the outcomes by giving them numbers on a scale of zero to ten, where 0 is absolute misery and 10 is perfect bliss. My decision could be portrayed in the matrix (Fig. 3.1) that appears on the following page.

Clearly, if I am eating by myself, I'd rather have a salad, but the salad won't work very well for a crowd. I like pasta somewhat by myself, but it would be the best choice if there will be others joining me. If I had reason to think there was exactly a 50–50 chance that I would have friends over, I could do the calculation required by this version of rational choice:

$$\text{Salad} = .5 \times 6 + .5 \times 3 = 4.5 \quad (3.1a)$$

$$\text{Pasta} = .5 \times 4 + .5 \times 8 = 6 \quad (3.1b)$$

Aha! I should put the water on for pasta right now. I will be a little disappointed if it's just for me, but I will be very happy if friends drop by. This is a very simple example of an **expected utility** calculation. The numbers that represent the values

		OUTCOMES	
		Alone	Friends
OPTIONS	Salad	6	3
	Pasta	4	8

Fig. 3.1 A matrix for calculating the expected benefit of dinner. Each cell (intersection of row and column) represents the benefit from dinner under that condition (alone or with friends) and based on that food choice (salad or pasta)

of the outcomes are referred to as their utility; the calculation is expected because I don't know for sure what the future will bring, but I factor in the probability of each outcome. Equations 3.1a and 3.1b represent the closest I can come to a rational anticipation or expectation of the future.

Of course, the odds of friends arriving may be less than or greater than 50 %, and I might not know off the top of my head how desirable any of the outcomes are. We could rewrite Eqs. 3.1a and 3.1b to be more general, so that they can apply to any possible pair of probabilities and any valuation of the outcomes. This means moving into the realm of algebra, replacing specific numbers by unspecific letters. To do this, let's invent some terminology. Let's call the first possible outcome, the one that occurs when I'm alone, outcome 1 and its utility for me v_{S1} for "the value of outcome #1 when I make salad" and v_{P1} for "the value of outcome #1 when I make pasta". Then the utility for me when friends come over is v_{S2} or v_{P2} depending on what I make. Each outcome has a probability; call the probability of the first outcome (alone) p_1 and the second (friends) p_2 . The formulas for my two options are now:

$$\text{Salad} = p_1 v_{S1} + p_2 v_{S2} \quad (3.2a)$$

$$\text{Pasta} = p_1 v_{P1} + p_2 v_{P2} \quad (3.2b)$$

For each of these, we could write it using the summation sign:

$$\text{Salad} = \sum_i p_i v_{si} \quad i = 1, 2 \quad (3.3a)$$

$$\text{Pasta} = \sum_i p_i v_{pi} \quad i = 1, 2 \quad (3.3b)$$

Now we are ready to graduate to the most general version of the expected utility formula, one which could apply to any option with any number of possible outcomes. It looks like this:

$$EU(B) = \sum_i p_i v_{Bi} \quad i = 1, 2, \dots, n \quad \sum_i p_i = 1 \quad (3.4)$$

The left-hand side reads “the expected utility of (option) B” and it equals the sum of the product of probability and value for every possible outcome. The middle part of Eq. 3.4 says that there are n possible outcomes, and each one of them is calculated. The short equation on the right says that the sum of all probabilities is equal to 1; this means you are not overlooking any potential outcome. (Whatever eventually happens has to be one of the possibilities you calculated.)

This is how it looks for any single option B. What economists mean by “rational choice” is that individuals should select B so as to maximize Eq. 3.4. To do this, you would have to identify all the possible outcomes that might arise, assign a probability to each of them, determine the value of each outcome given a particular course of action, add up all probability/value products (p times v), and do this for every option that presents itself. It’s not easy being rational, and by the time you’re finished it’s likely that everything has changed, and you’ll have to do your calculations all over again.

Narrowly speaking, all of this is simply absurd, and spelling it out as I have just done seems to be enough to discredit it. Nevertheless, there is a fallback position that most economists would embrace. Yes, no one has the time or obsession for detail to make truly rational decisions as defined by Eq. 3.4, but perhaps most people, most of the time make decisions that are reasonably close to this standard, even if they don’t know it. One version of this story is that we have rational compartments of our brains that are constantly cranking away, making calculations of this sort without any direct supervision by our conscious minds. Thus, we make decisions that are more rational than we realize. Another story says that we approximate Eq. 3.4 through trial and error. We make similar types of decisions over and over, and we learn from our experience what sorts of choices give us the greatest utility.¹ Thus we end up choosing *as if* we were rational, even if, at any moment, we aren’t. As we will see in future chapters, “as if” plays a large role in economic theorizing. There is a third story: even if this isn’t such a good description of how people actually act, it’s the ideal model of we should think about the choices we face. By basing itself on this notion of rationality, economics, in this case, would be about how people should be even if they aren’t—it would be a normative, not positive, theory of decision-making.

As we will see further on, however, there has been much interest among economists in recent years of other models of decision-making, typically drawn from cognitive and social psychology. This field of *behavioral economics* is growing rapidly and has applications to nearly every topic of economic research. It would take us too far afield to introduce it now, but the more general point is that the rigidity implied by expected utility as a depiction of ordinary, day-to-day

¹ Note that this argument is essentially the same as Darwinian natural selection, transposed to the realm of behaviors.

decision-making is breaking down. For later investigations of alternative theories of individual and organizational choice, however, we will need to have a clear understanding of expected utility, since the alternatives are normally defined as carefully specified departures from it.

Incidentally, there is a major alternative to the calculation of expected utility that is frequently employed in business and policy circles, even if economists largely ignore it: scenario analysis. Since this is a textbook on economics, I won't spend much time with it here. In general terms, however, scenario analysis involves four steps:

- Identify the key factors that are likely to influence future developments. At most, pick just handful of such factors; in many instances analysts pick just two. Examples could include whether new laws will be passed that alter the marketplace, whether public opinion shifts in one direction or another, whether new sources of energy are found, and so on.
- For each factor, pick a very small number of alternative developments to look at: a specific law will or will not be passed, public opinion goes in one particular direction or another, a new energy source with certain characteristics is or is not found. No attempt is made to incorporate all future possibilities; just a few representative ones are considered. Business gurus, for instance, like models with two factors and two possibilities for each, creating a 2×2 matrix of scenarios. This looks very nice in a slide presentation.
- Each combination of factor developments, such as each cell in a 2×2 matrix, constitutes a scenario. Analyze this scenario: what actions should be taken and what outcomes should we expect?
- Sum up all the scenario analyses. Which are the most desirable and which the least? Which scenarios are most likely? What actions taken today, before we know which scenario applies, will work best across different possibilities?

Scenario analysis does not try to boil everything down to a single number the way expected utility analysis does. It is purposely fuzzier. On the other hand, it makes fewer assumptions, such as attaching exact probabilities to each potential outcome, and it provides more "stories" that explain *how* outcomes occur and which factors may play the largest role, rather than just crunching numbers. The reason for bringing it up at this point is to demonstrate that the version of rational decision-making employed by economists is not the only one that exists, or even the one that decision-makers are most likely to use.

3.4 Individual and Collective Rationality

There is one more wrinkle in the theory of rationality that will play a large role in the chapters to come, the distinction between **individual** and **collective rationality**. Its most vivid representation is in a sort of game that goes by the name of Prisoner's Dilemma. This was created in 1950 by Merrill Flood and Melvin Drescher, who were working at the time for the RAND project funded by the US Air Force. It takes its name from a story that its authors told in order to flesh it out, although, as we will

see, it is highly adaptable to a wide variety of stories. It goes like this: suppose a crime has been committed and the police take in two suspects, who are held in separate cells. The evidence against them is limited, and the prisoners know this. It is clear that the best chance for a conviction is for one or both to confess, and the prisoners know this too. What each doesn't know, however, is what the other will do. The police, being sensible, offer each a deal. If you confess, they say, we will take that into account in sentencing. In fact, if you confess and finger your partner, and if we use this testimony to convict him, we'll let you off completely free. On the other hand, if you hold out against us and your partner confesses, you'll get the maximum we can throw at you. If you both confess your time in jail will be less than this maximum. What the police don't say, but what the prisoners know, is that if neither confesses they will be able to plea bargain to get an even lighter sentence, since the evidence is so meager. The dilemma in the Prisoner's Dilemma is deciding whether to give the police what they want.

Let's see how this can be expressed algebraically. To do this, we need ways to represent the elements of the story. The first step is to identify the players. That's easy; we can call them A and B. Then we need to assign letters to represent the two possible choices of confessing or not confessing. The standard language used by specialists in the field is to refer to them as *defecting* and *cooperating* respectively. That is, by confessing a prisoner is turning against his partner; by sealing his lips he is continuing to act in a partnerlike manner. (Note: the game explores the strategic interaction between the prisoners, where the police are in the background. For this reason, the refusal to speak with the police is called cooperation, even though the police wouldn't see it that way.) We shorten these to D (defect) and C (cooperate). Next we need representations of the consequences for the prisoners of their choice of actions. A simple way to convey this is with numbers, as we might for scores in a game. In this way, the higher the score the better the result. The numbers themselves are arbitrary; to make things easy we will select the integers 1, 2, 3 and 4.

The characteristic features of a Prisoners Dilemma game are summarized in the **payoff matrix** of Fig. 3.2 on the next page. It shows the outcomes to both players resulting from each combination of choices. A chooses between row 1 (D) and row 2 (C); B chooses between column 1 (D) and column 2 (C). The pair of numbers tells us what happens to the two of them (A's payoff, followed by B's) when A and B have made their particular choices. For instance, the upper-left cell says that, if both A and B choose to defect, each player will receive a "2". In the upper-right cell, A gets a "4" and B a "1".

To keep things as simple as possible, we are assuming that the game is perfectly *symmetrical*, that the payoffs for each player are the mirror image of the payoffs for the other. Three more assumptions will make the point of the game inescapable. First, we will assume that the prisoners are selfish; they care only about what happens to themselves individually and are unconcerned about the fate of their partners. In doing this, we are simply applying the economist's typical first-approximation assumption of rational self-interest to the players in our hypothetical game. Second, we will take it as given that only direct benefits recorded in the payoff matrix influence their decisions; that is, they are completely outcome-driven

		B	
		D	C
A	D	(2, 2)	(4, 1)
	C	(1, 4)	(3, 3)

Fig. 3.2 Payoff matrix for a two-player prisoners dilemma. A and B are the players; C (cooperation) and D (defection) 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 A is given first, then the payoff for B

and take no account of ethical principles, customs or other such considerations. Finally, we add the assumption that this is purely a one-time event: the police make their offer once, and then there is no further interaction between any of them. (This last restriction eliminates any possible influence of future “games” on this one.) All in all, this is a highly artificial world we have created, but, for now, its purpose is clarity of insight and not realism.

What then will our prisoners do?

Note that it is in the individual interest of each player to defect, no matter what the other one does. Suppose you are player A, for instance. Your opponent, B, will either cooperate or defect. If B cooperates you are better off by defecting, since $4 > 3$. If B defects, you will also be better off if you defect too, since $2 < 1$. So either way, you should choose D. (Formally, since D has a higher payoff than C for A whatever B does, there is no need to estimate the likelihood of B’s choice, as we would if we were to use the expected utility formula in Eq. 3.4; $EU(D) > EU(C)$ for all values of p .) Since the same logic applies to your opponent’s decision, however, the pair of you are likely to end up in the upper left-hand cell, each receiving 2 when, with some forethought and coordination, you might have been able to agree that both should cooperate, so that both would receive 3 instead.

Looking more closely, we see three distinguishing characteristics of the Prisoners Dilemma game. First, both players face a *cost to unilateral cooperation*. In other words, if you cooperate and the other player defects, you get the worst possible payoff. This is sometimes referred to as a “sucker’s reward”. (It is what baseball legend Leo Durocher no doubt had in mind when he said “Nice guys finish last.”) Second, both face a *benefit to unilateral defection*. If the other player is cooperative but you defect, you get the best possible payoff. There is a reward in this type of game to those who prey on the trust of others. Finally, *mutual cooperation is better for both players than mutual defection*. It doesn’t make sense for the players to end up defecting against one another, since they could both do better by cooperating instead. These three characteristics, in fact, define the Prisoners Dilemma as a particular type of game. (There are many other games that have been developed by economists and other social scientists to aid in the analysis of

complex strategic interaction. Some are close relatives of the Prisoners Dilemma, created by altering one or more of the three characteristics described in this paragraph.)

Can a Prisoners Dilemma have more than two players? Yes: for convenience let player A remain an individual, and have player B represent “everyone else”. From each player’s perspective, they are A and the others are B; so as long as the payoffs remain the same, the decision-making and the outcomes remain unchanged. Additional complications can be introduced by tying the payoffs to the proportion of B that cooperate or defect, but we will not take them up now. They don’t alter the essential characteristics of the game, although they do alter the prospects somewhat for arriving at mutual cooperation. (They tend to make explicit agreement more difficult, but also lower the cost to deploying cooperation as an individual strategy, as we will see in a later chapter.)

What makes the prisoners dilemma so interesting is its stark opposition of individual and social rationality. From an individual standpoint, the best choice is D, taking advantage of the benefit of unilateral defection or avoiding the cost of unilateral cooperation, depending on the choice taken by the other player. Since this is true for both, by following their individual self-interest, they end up in a state of mutual defection, each earning the inferior reward 2. It is in their collective self-interest, however to arrive at mutual cooperation, since in that case both would be better off with 3. In this respect, the Prisoners Dilemma represents a case in which the sort of self-centered calculation normally presupposed by economists is self-defeating.

It is important to underline this conclusion. Recall that, for economics, the governing metaphor for describing economic life is exchange. People are assumed to take part in exchanges based on their calculations of self-interest. If exchange is voluntary and sufficiently well-informed, at least one party will be better off for undertaking it, and no party will be worse off. Therefore all exchanges that meet these two conditions are “good”, and a major purpose of economic policy is to facilitate them. This is what “free market economics” is all about. We have questioned whether the metaphor of exchange leaves something out and whether rational self-interest is the right account to give of why people make the choices they do, but the Prisoner’s Dilemma poses a different challenge: *if a real-world situation has the characteristics summarized in Fig. 3.2, even perfectly well-informed, voluntary, self-interested exchange can lead to participants being worse off rather than better off.* To put it differently, the case for markets depends on there *not* being some aspect of the situation that rewards the sort of cooperation people engage in when they put the interests of others ahead of themselves. True, in many of the contexts in which markets function, they function well enough: there are minimal opportunities for cooperative gain. Nevertheless, it is important to be on the lookout for the exceptions, as we will see.

With the logic of the model under our belt, let’s switch to real-world examples. Consider first the problem of insuring that corporations honestly report their financial condition. Corporate record-keeping is extremely complex, and outside observers may not be able to tell whether financial reports accurately reflect current

business earnings and potential future liabilities. Businesses want to issue positive reports in order to make it easier to get new financing at low interest rates, and to keep their stock prices as high as possible. (We will look more closely at the incentive effects of financial markets in Chaps. 8 and 17.) But dishonest reporting complicates the problem of management and may squander whatever goodwill exists toward the company in the eyes of the public. To make things as simple as possible, assume there are only two companies and only two consequences of honest or dishonest reporting, stock prices and public image, with the first more important than the second. With just two companies and one stock market, relatively higher share prices for one company mean relatively lower prices for the other. We also (temporarily) assume that there are just two options, honesty and dishonesty, that the choice of which option to take is made independently by each company with no possibility of revision, and that there is no public regulation of corporate accounting.

For each company there are four possibilities:

- (a) It is dishonest but the other company is honest.
- (b) It is honest but the other company is dishonest.
- (c) Both are honest.
- (d) Both are dishonest.

How will they rank them? Given our assumptions, we expect (a) to be best, then (c), then (d), then (b). Here's why: (a) The best outcome arises when your own company is dishonest but the other is honest. This gives you a clear advantage in the stock market, which (by assumption) is more important than the risk of a negative public image if you are found out. (b) If both companies are equally honest or dishonest, their effects on the stock market cancel out, and the only difference has to do with public perception. In this case, mutual honesty is the better policy, since both companies enjoy more public approval. (c) Mutual dishonesty also preserves the status quo in the stock market, but comes at the cost of potential public disapproval. (d) The worst outcome occurs when you are honest and the other company isn't, since now your stock prices fall and theirs rise. Again, this is assumed to outweigh the benefits of potentially greater public approval.

This information can be displayed in a payoff matrix, as on the following page in Fig. 3.3, where dishonest reporting constitutes defection (D) and honest reporting cooperation (C).

As before, it doesn't matter what the other company chooses; dishonesty is preferred. If company B is dishonest, A can either be dishonest and receive (d) or be honest and receive (b)—but (d) is better than (b). Similarly, if B is honest, A can get (a) by being dishonest and (c) by being honest—and (a) is better than (c). So A does not have to read B's mind; it is better to fudge the accounting report in either case.

In the real world these assumptions are not entirely valid. Certainly there are more than two companies, and they have the opportunity to revise their accounting policies every time they prepare a new report. We will see in Chap. 10 that this moderates the collective action problem to some extent. If the reports are dishonest it is also possible that the financial markets will find out, and this could have devastating consequences for the company *and* its accountants. Nevertheless, despite all these qualifications, the historical record demonstrates that businesses,

		B	
		D	C
A	D	A gets (d) B gets (d)	A gets (a) B gets (b)
	C	A gets (b) B gets (a)	A gets (c) B gets (c)

Fig. 3.3 Payoff matrix for honesty in accounting as a two-player prisoners dilemma. A and B are two companies that can either report their financial condition honestly (C) or dishonestly (D). The ranking of outcomes is $(a) > (c) > (d) > (b)$, based on the assumption that the effects in the stock market of differences in honesty outweigh the costs of dishonesty for potential public disapproval—if the dishonesty is revealed. Each company therefore has an incentive to choose D, even though the result is that they are both worse off than they would be if both chose C

if not closely regulated, will often produce misleading accounting reports. (These acts of dishonesty often start out as small exaggerations and omissions and then gradually expand until they are out of control.) This is why industrialized countries *do* have regulations that impose penalties on firms and their accountants if dishonest reporting is brought to light. The goal is to alter the incentives so that, with possible penalties taken into consideration, (c) comes out on top of the ranking.

Next consider an important issue of public policy, state incentives for economic development. Suppose A and B are states—New York and California, for instance. A private company plans to make a major investment, but is unsure which state to locate in. Each state wants to be chosen, since the investment will create jobs and expand the tax base. If we want to cast this in the form of a Prisoners Dilemma, we can designate defection as negotiating individually with the company, offering subsidies like tax breaks and subsidized state-provided services. Cooperation means refusing to enter into negotiations and offering no special incentives to the company. It is plausible, then, that Fig. 3.4 on the following page might reflect the payoffs to each pair of choices, particular if they are evenly matched as potential investment sites. If both states offer equivalent subsidies they are on an equal competitive footing, but whoever wins the contest will have to provide costly benefits to the company. On the other hand, if neither makes offers the competitive situation is unchanged, and the winning state avoids having to pay for subsidies. Mutual cooperation is therefore better for both states than mutual defection. Unfortunately, both the incentive to unilateral defection and the penalty for unilateral cooperation exist in this example, provided that, for each state, the benefit of gaining the investment exceeds the costs of the subsidies. If New York cooperates by refusing to negotiate, California can (quietly) cut a deal and come out ahead. Moreover, if California defects, New York is worse off if it cooperates than if it defects too. For these reasons, it will be difficult to ensure that both states adopt a cooperative strategy. In the real world, of course, there are 50 states, and the prospects for achieving cooperation are even worse; this is why most states spend millions or even billions of dollars in a dubious competition over investment and

		B	
		D	C
A	D	<ul style="list-style-type: none"> • costly subsidies offered • equal competitive chances ◦ costly subsidies offered ◦ equal competitive chances 	<ul style="list-style-type: none"> • costly subsidies offered • high likelihood of winning the investment ◦ no subsidies offered ◦ little chance of winning the investment
	C	<ul style="list-style-type: none"> • no subsidies offered • little chance of winning the investment ◦ costly subsidies offered ◦ high likelihood of winning the investment 	<ul style="list-style-type: none"> • no subsidies • equal competitive chances ◦ no subsidies offered ◦ equal competitive chances

Fig. 3.4 Payoff matrix for investment competition as a two-player prisoners dilemma. A and B are two states competing to have a private investment sited. D (defecting) means offering subsidies to lure the investment; C (cooperating) means not offering them. The four cells describe payoffs resulting from the four combinations of choices. The first payoff listed is A’s; the second is B’s

jobs. (Note that public investment incentives serve no legitimate social function in this example, since, by hypothesis, the investment will take place somewhere with or without subsidies. There might be a case for subsidies if the investment itself, and not just the location, is at issue.)

Now let’s step back from these examples and take stock. For purposes of explanation (or perhaps shock value), I have emphasized the dilemma aspect of the Dilemma. Individual interest drives players into situations that are collectively irrational. Is there no escape then? In fact, the main direction of research into this game (theoretically, in controlled experiments and real-world case studies) is in solutions. There are several possible routes to a cooperative outcome, each with its costs and benefits:

- **Coercion.** An organization representing all the players can try to force them to cooperate by punishing defectors. This, of course, is one of the functions of organized crime in real life police scenarios. The prisoner who “sings” (confesses and implicates others) is in trouble.
- **Inducements.** Instead of punishments for defection, an organization can offer rewards (called “side payments” by game theorists) for cooperation. Some groups, for instance, enter the names of individuals who take on voluntary burdens into raffles, where they can win prizes.
- **Reputation.** If the game is played repeatedly by the same players (as it often is in real life), the players themselves can bring about cooperation by rewarding other cooperators and punishing defectors. “What goes around, comes around” is the

motto of such groups. Networks of cooperation and mutual support are common, in fact, in most stable communities that have been studied.

- Custom. Societies often promote social norms that require its members to cooperate. They are taught them (“socialized”) at an early age, and the result is that most people go through life abstaining from chances to cheat and take advantage of others without even consciously considering it. Of course, not everyone is equally well-trained, and social customs seem to allow loopholes for situations in which defecting is permitted. Thus, there is a tendency in many societies to tolerate self-interested calculation when the other parties are unseen and unknown, or when they are outsiders or disparaged castes or classes.
- Intrinsic satisfaction. There is some experimental evidence that human beings, being social creatures, are genetically equipped to experience satisfaction in cooperation, at least in some contexts. Virtue, it seems, *can* be its own reward.

We will return to these potential escape routes later on as we apply the Prisoners Dilemma model to various economic issues.

At this point you may be wondering how common such applications of the Prisoners Dilemma actually are. The answer is that it is one of the most widely used theories, not only in economics, but in several other fields as well. The reason this game has captured the imagination of researchers is that, despite its simplicity, it seems to get to the heart of many problems in modern life. Once the model has become familiar, you too may begin to see hints of it everywhere. Of course, complicated real-world situations are seldom fit exactly into the format of this game, and it is a matter of interpretation and judgment whether the similarities outweigh the differences. This is something to keep in mind when we return to the Prisoners Dilemma in future chapters.

3.5 Equilibrium: People as Particles

Ask most people what they think “equilibrium” means, and they will tell you something like “balanced”, “in the right place” or “satisfied”. We talk about people achieving an equilibrium in their lives, in which their different interests are each given a proper role, and the result is a general feeling of well-being. To be in disequilibrium is to be out of balance: missing something important, dissatisfied, unhappy.

As in so many other cases, the technical use of the word “equilibrium” departs significantly from its popular use. For economists, no concept is more important; nearly every theory they use depends on equilibrium as its organizing principle, and so the potential for misunderstanding is very large. In this section we will look carefully at the economic version of equilibrium, and identify the features that differ from the everyday variety.

Economists take their notion of equilibrium from classical mechanics (Newtonian physics). In any physical state there are a variety of forces at work on each “body” or physical unit. A billiard ball is struck, and the force of the cue initiates motion. Momentum propels it forward at a decreasing rate, as the effects of gravity

and friction take their toll. Eventually there is no force sufficient to cause the ball to change position, and it comes to a stop. Equilibrium is a state in which no existing force is sufficient to alter the system; it will remain in that state until a new force is applied.

This is also how economists think about equilibrium. Suppose I land a new job that pays twice as much as my old one. I now have lots more money, and I'm looking to spend some of it. My spending pattern—the amount I generally spend per month on food, entertainment, etc.—is now out of equilibrium, since a new force is bringing about a change. I will increase my spending until the level I reach is consistent (from my perspective) with my new income, at which point it will stabilize. Economists would say that I have reached a new equilibrium in my consumption, based on a change in income. This equilibrium is expected to continue until some new change—in income, prices, life circumstances, etc.—disrupts it.

To be very precise, equilibrium in its economics usage has two elements. First, it identifies a situation in which there is no “inner” tendency toward change. Once the people or institutions that make up the situation have achieved a common equilibrium (an equilibrium for each participant that takes into account everyone else's), there is no reason for any further change, unless some new event takes place. This is the “timeless” aspect of equilibrium: equilibrium as a cessation of time. (If time is motion, equilibrium, because it is motionless, has no time.) In the real world, of course, nothing ever stands still, and this means that equilibrium (in its economic sense) can apply only to models of things, not the things themselves. This is a limitation of models, but, as we have seen, limitation is the *point* of building models—they are conscious attempts to limit our vision in order to see particular things more clearly.

The second aspect of equilibrium is more subtle. Recall the example of my change in consumption after receiving an increase in income. Equilibrium occurs when the spending change is completed, but the story begins in a state a disequilibrium, before I change my spending. When the curtain rises, I have the extra money, but I am not spending it yet. In order for the concept of equilibrium to come into play, I have to make the transition from lower to higher spending. In the real world, this might take a bit of doing; spending requires effort. I might need to locate new stores or open up a new bank account. For the purposes of our highly simplified story, the exact process is not important. What is important is the *assumption* that I will manage it somehow, and go from a routine of spending less to one of spending more. In other words, for the concept of equilibrium to have any meaning in a world of disequilibrium, there must be a process—an impetus—that moves participants in the economy toward equilibrium when they are not in it.

Thus, every instance of equilibrium in economics has two components. First, it identifies a state at which the economy (in a model) comes to rest. Once this state is achieved it will persist as long as the conditions of the model persist. Second, it incorporates a story that describes why and how people, beginning in a state of disequilibrium, will move to equilibrium. In my spending example, the story might

be one of unsatisfied desires when I had less money, a backlog of items on my wish list, etc.

This is exactly what equilibrium means in economics and nothing more. *It does not imply happiness, balance or perfection.* Indeed, we will see that it is possible for an economy to be in equilibrium (according to commonly used models) and for massive problems to exist: high levels of unemployment, great disparities in income and wealth, life-threatening pollution and the failure to provide essential goods and services to those who need and want them.

In principle, economics is concerned with improving social well-being. The strategy used by most economists is outlined in the box below:

The Economist's Strategy for Addressing a Social Problem

1. Create a model representing the essentials of the current state of the economy or some important part of it.
2. Determine the equilibrium of this model.
3. Demonstrate the reasons why this equilibrium falls short of solving the larger problem under investigation.
4. Propose changes to economic policies or institutions that mitigate the shortcomings of the equilibrium in Step 2. These can either be changes in the way the economy operates, so that a different equilibrium will emerge with fewer social deficiencies, or an adjustment to be administered after the original equilibrium is reached.

Suppose, for example, that problem is global warming, the buildup of greenhouse gases that threatens catastrophic changes in the world's climate patterns. This is due to many factors, one of which is that people who own cars are driving them too much and burning too much gasoline. The economist would begin by demonstrating that this excessive driving is not a temporary aberration, but an equilibrium that can continue indefinitely unless something is done. This part of the analysis would be couched in some form of supply-and-demand theory: estimating the demand for gas as affected by prices, the supply provided by oil companies, etc. The current situation is an economic equilibrium if consumers are buying the quantity of gas they want to buy at the going price, and oil producers are producing the amount of oil they want to produce at that price. Both are doing what they would like, given what the others are also doing. Hence there is no reason for anyone to choose differently.

The third step is for the economist to show that the equilibrium in the market for gasoline is not optimal for "society"—in this case, the world as a whole. This can be done by assessing the costs and benefits of reducing global warming or, more simply, demonstrating that some of the cost of burning gasoline is not being taken into account by consumers and producers in the model; their behavior is "defective" from a social standpoint. This leads to the final step, which might result in a proposal for a tax on gasoline at the pump, for example, in order to reflect the

true cost of greenhouse gas emission, so that the new equilibrium (with the tax) would be better for society than the old one.

In this example, and in all the work that economists do, equilibrium is a concept whose only purpose is to facilitate modeling and other forms of analysis. It does not imply any normative judgment about whether the economy is doing what we want it to. It can be used normatively, of course, when economists employ models to help them devise solutions to economic situations they see as problems for society.

3.6 Systems of Allocation

As the previous sections of this chapter should make clear, economists regard the fundamental economic question to be that of choosing among alternative options available to us. This is typically done in the context of limited resources. Consumers do this when they decide what to buy (the resource being money), firms when they decide what to produce and how, and so on. In economics lingo, this is the problem of **allocation**. Any given economy offers us a system for making allocative decisions, but this system can vary dramatically depending on where and when we look.

Indeed, everywhere there are people there is an economy. Human economies have existed in all historical eras and on every continent. If we permit them to pass before us—the hunting and gathering societies, the irrigation-based civilizations of China and the Middle East, the native people of the Pacific Northwest and their potlatches, the early renaissance merchant economies of northern Italy and the Hanseatic League, modern industrial capitalism—it seems as though there is no common thread. (This is disconcerting: will future generations find anything in our own economic order that reminds them of their ways?) Yet, at a very high level of abstraction, we can say that only five types of economic decision-making have ever appeared:

1. Custom: we do it this way because this is how it is done. The basis for custom may be a religion, a body of traditional teachings, or simply the unconscious imitation of older generations.
2. Gift exchange: individuals or groups provide goods or services to others with no specified payment in return, but with the expectation that others will make similar gifts to them in the future. Some anthropologists might argue that gift exchange is really a form of custom, since the expectations that keep the system going are usually embedded in larger customary relationships. For instance a successful hunting party will share its bounty with others who were less successful, knowing that the roles may well be reversed after the next hunt—but this in turn depends on social norms regarding “proper” behavior that make it reasonable to assume that such reciprocity will actually take place.
3. Administration: one entity, either individual or collective, instructs another entity to do, or not do, something. Governments regulate modern economies according to this principle; this is also the basis for college admissions, where admissions officials decide which students can be enrolled.

4. Collective organization: the members of an entity decide on the actions of their *own* group. What distinguishes collective decision-making from administration is that a collective process is mutual—the decision-makers are the same ones who will have to abide by the decision—whereas administrators make decisions for others. (Administrators in one context, of course, can be collective decision-makers in another; think of a board of trustees. One moment they are collectively planning their own calendar, and then next they are making policy decisions for the university or other institution they are in charge of.)
5. Markets: individuals or groups enter into voluntary contracts with one another. Decisions in markets are made by decentralized agreement; each participant decides whether to accept the offers of others, and the choices are added up in the marketplace to produce a social outcome.

These are ideal types; in reality, most institutions are mixed. In fact it is not uncommon to find all five of these commingled in the same social nexus. Perhaps a very simple example will make this clear. Suppose I stop at a farm stand to buy some vegetables for dinner; what allocative mechanisms are at work? Certainly I am involved in a market, since I am offering to buy goods from a producer, but that is not all. The administrative mechanism of government regulation is there behind the scenes, determining, for instance, what chemicals the farmer was permitted to use in growing the food, and how much residue can be left on it. Even if I am willing to buy food with more chemical contamination for a lower price, I do not have that legal option. (Of course, the farmer might violate this law.) Custom also enters in. Am I permitted to “test” the produce before I buy it? Can I peel back the corn husks to look for pest damage? Can I sample the strawberries? This is entirely a matter of local custom. In some situations this behavior is expected; in others it would be taken as a violation of accepted practice. You probably won’t see any signs posted telling you what’s allowed; you just have to know. (Anyone who has shopped at farm markets in different countries will recognize how much customs can vary, and how this affects the quality and price of what is bought.) Finally, when I take the bag of fruits and vegetables home, the ultimate allocation is a combination of gift exchange and collective decision-making. A family, of course, is fundamentally a system of gift exchange; family members provide goods (like produce bought at a farm stand) and services (like cooking) without demanding immediate compensation, expecting that those who benefit today will make their own contributions tomorrow. There is also an element of conscious, collective decision-making. For example, as a family, we may discuss what to do with all those strawberries: should we just pile them into bowls and eat them as they are or use them to make a pie? So all five mechanisms are likely to come into play at some point, and all of them are “economic” in the sense that they help determine the production and distribution of food—a fundamental economic good.

To consider another example, we have already examined the ways that work in typical businesses has aspects of voluntary exchange (markets) and other aspects that do not fit so neatly into that framework. Certainly market allocation is involved when workers consider what job offer to accept, for example, or when wages are negotiated. But there is also a large role for administrative allocation, as in the

day-to-day authority invoked by supervisors. Work teams sometimes have scope for collective decision-making, and custom (for instance, “corporate culture”) is pervasive. It is also recognized that some businesses run on a type of gift exchange: employers make a “gift” of better wages and working conditions than the market requires, and workers respond with a “gift” of greater work effort and higher commitment. We could be justified in saying, then, that productive work in our society is not organized solely through markets, but also through the other four allocative mechanisms. Nevertheless, markets play a larger role in these aspects of the US economy than in most of Western Europe, and, as we will see in a later chapter, this role has continued to expand in recent years. Allocation is rarely left to just one process, but some processes may be more important than others.

The Main Points

1. Choice and exchange can be considered metaphors for economic life, in the sense that they foreground certain aspects of the economy—moments of personal choice, especially related to buying and selling—at the cost of giving less attention to other aspects. This can be made clearer by comparing the conventional metaphor of economics (choice) to others that have sometimes been advanced, like Marx’s view of the economy as the product of human labor or Schumpeter’s emphasis on combat between businesses competing to conquer each other’s markets.
2. Economists usually assume that individuals are rationally self-interested. This embodies two assumptions, that we tend to act on the basis of what serves our own personal interest (rather than taking into account the interests of others apart from how they affect us), and that our reasoning adheres to the model of expected utility maximization. To be rational in this sense is to consider all possible outcomes of every course of action, placing a probability and a value on each. The formula that expresses this is

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

where $EU(B)$ is the expected utility of option B , p_i is the probability of outcome I arising, and v_{Bi} is the value of this outcome should it occur. The rational individual calculates this for all possible options and then chooses the one whose expected utility is greatest. An alternative approach to the problem of uncertainty is scenario analysis.

3. There is an important distinction to be made between individual and collective rationality, where the first refers to choices that maximize expected utility for individuals one at a time, while the second represents choices that people might make in a coordinated way that could yield even more individual utility. One powerful demonstration of this distinction is the prisoner’s dilemma model. It is characterized by three features: the benefit of unilateral defection, the

disadvantage of unilateral cooperation, and the superiority of mutual cooperation to mutual defection, where cooperation refers to choices made by individuals that are in the interest of other players and defection to choices that are against the interest of other players. In a prisoner's dilemma it does not matter what other players do; each individual is better off defecting whether or not the other cooperates or defects. Nevertheless, by both acting in an individually rational manner, the players end up with the less desirable outcome of mutual defection. In practice, societies have evolved various mechanisms that sometimes steer individuals in the direction of cooperation in prisoner's dilemma-like situations.

4. As used by economists, "equilibrium" does *not* mean "desirable". Rather, it refers to a state of affairs in which there is no inbuilt tendency toward change, and to which individuals are likely to return if they deviate from it. The purpose of having such a concept is to facilitate prediction: by identifying a particular outcome as an equilibrium, economists are asserting that it is likely to arise, and the specific reasons why the equilibrium properties are believed to be met provide the basis for explaining this prediction.
5. Economists view economic systems as solving problems of allocation, determining how limited resources are devoted to competing uses. In general terms, there are five such allocative mechanisms—custom, gift exchange, administration, collective organization and markets. In practice, they tend to overlap.

► Terms to Define

Administration
Allocation
Collective organization
Cooperation vs defection (in a Prisoners Dilemma)
Custom
Equilibrium
Expected utility
Gift exchange
Incentive
Individual vs collective rationality
Markets
Payoff matrix
Prisoner's Dilemma
Rational choice

Questions to Consider

1. Make a list of all your economic activities (involvement with the production, distribution or use of goods and services) yesterday. Which of your activities could best be described as choices? Which were exchanges (money for goods and services)? To what extent were the non-choice and non-exchange activities foreseen and incorporated when you chose and exchanged? How close do the choice and exchange metaphors come to encompassing the factors that people other than yourself should take into account when trying to understand your economic life?
2. Increasingly, colleges and universities are being asked to regard themselves as businesses providing educational services to their student customers. In other words, students are viewed as exchanging money for a package of services including classes, student support, campus life activities, etc. This approach is then used to identify marketing, quality assessment and other initiatives for higher education modeled on management in the for-profit sector. Clearly, this vision depends on the appropriateness of the underlying metaphor of education as an exchange. What, in your view, are the advantages and disadvantages of thinking about education in this way? What aspects of the situation are captured in the exchange metaphor? What aspects are excluded or misrepresented?
3. How many decisions have you taken in the past week that were *not* based on self-interest? In other words, how often did you put the interests of others ahead of your own?
4. Look again at Fig. 3.1, with its utility matrix for the salad-or-pasta dinner choice. What is the break-even probability of friends coming over? That is, at what probability of friends dropping by is the expected utility of making a salad equal to the expected utility of making pasta? Can you show how this would be calculated using the expected utility formula represented by Eqs. 3.2a and 3.2b?
5. To what extent was your decision to take this economics course “rational” in the precise terms of expected utility theory? Did you consider all the alternative courses of action? Did you forecast the likelihood and desirability of possible outcomes resulting from this decision?
6. Do you think that rationality in the form of expected utility maximization (using the formula in Eq. 3.4) represents an ideal that we should aspire to, even if it is sometimes beyond our abilities? Should people be encouraged (or even taught in schools) to think this way?
7. One of the most-publicized problems in professional sports is the use of performance-enhancing drugs. Of course, not all such drugs constitute a problem, just those that are physically harmful, like steroids. (Steroids increase the risk of cancer later in life.) Try to fit the Prisoners Dilemma model to this problem. Assume a two-person contest (all other competitors can be folded into player B) and symmetrical payoffs. Crucially, assume that each athlete values winning so much that having a competitive advantage outweighs the future health costs of taking steroids. Construct a payoff matrix as in Figs. 3.2, 3.3,

and 3.4, and show that the three central characteristics of the Prisoners Dilemma apply. Show that the assumption that winning is valued above health is necessary for two of these characteristics. Do you think this assumption is warranted in real life? Which of the routes to cooperation sketched above is employed by sports organizations like Major League Baseball, FIFA (soccer) and the Olympic Games?

8. Most people in large cities live far from where they work or go to school. If a large proportion of them rely on cars for transportation, the result is rush hour traffic jams. Each driver presumably calculates the advantages and disadvantages of the various options available: drive or take mass transit, drive alone or carpool, leave and return during rush hour or travel on an off-peak schedule, etc. Driving every day, they are well-informed about the consequences. Under these circumstances, can daily traffic jams be an equilibrium outcome as we have defined equilibrium in this chapter? How could you find out? Suggest an approach using individual surveys, as well as one that looks at overall behavior (traffic flows) rather than individual statements. In each case, what would count as evidence of “equilibrium gridlock”? And why would it matter whether or not traffic tie-ups are equilibrium events? In other words, who might be able to use this information, and for what purposes?
9. In the household(s) you grew up in, what was the balance between the five systems of allocation? Which systems were employed for which goods or services? Do you wish the balance had been different? How?
10. Public libraries purchase books which are then made available to the community free of charge. There is little or no cost to acquiring a library card, and cardholders may borrow any books they choose. Of what system or systems of allocation is the library an example?