Introduction

What I Am Doing

If there were only one man in the world, he would have a lot of problems, but none of them would be legal ones. Add a second inhabitant, and we have the possibility of conflict. Both of us try to pick the same apple from the same branch. I track the deer I wounded only to find that you have killed it, butchered it, and are in the process of cooking and eating it.

The obvious solution is violence. It is not a very good solution; if we employ it, our little world may shrink back down to one person, or perhaps none. A better solution, one that all known human societies have found, is a system of legal rules explicit or implicit, some reasonably peaceful way of determining, when desires conflict, who gets to do what and what happens if he doesn't.

The legal rules that we are most familiar with are laws created by legislatures and enforced by courts and police. But even in our society much of the law is the creation not of legislatures but of judges, embedded in past precedents that determine how future cases will be decided; much enforcement of law is by private parties such as tort victims and their lawyers rather than by police; and substantial bodies of legal rules take the form, not of laws, but of private norms, privately enforced.

Going farther afield in time and space we encounter a much greater diversity, both in the sources of legal rules and in the ways in which they are enforced. If we are considering all systems of legal rules in all times and places, the ways in which legal rules are created and enforced in America in this century are simply data—one out of many possible solutions to the problem of human conflict, one out of many possible systems of legal rules. This book directs most of its attention to the past century or two of Anglo-American law not because it is more important than other legal systems but because the author, most readers, and most of the scholars whose ideas I will be talking about know more about that legal system than about the legal rules of Homeric Greece, Papua New Guinea, Saga period Iceland, or Shasta County, California. But the ideas I am discussing are as relevant to those systems as to ours—as we will see when we take a brief look at several of them in chapter 17.

There are many ways of looking at a legal system, among them the perspective of a legal historian, a legal philosopher, or a lawyer interested in creating arguments courts will accept or contracts they will enforce. This book is written by an economist. My approach is to try to under-

stand systems of legal rules by asking what consequences they will produce in a world in which rational individuals adjust their actions to the legal rules they face.

While this is not the only possible approach, it is one with very general application. Legal rules exist, at least in large part, in order to change how the people affected by them act. A speed limit exists because someone wants people to drive more slowly. The legal rule that holds that any ambiguity in a contract is to be interpreted against the party who drafted it exists because someone wants people to write contracts more carefully.

The economic approach works in two directions. Starting with an objective, it provides a way of evaluating legal rules, of deciding how well they achieve that objective. Starting with a legal rule, better, a system of legal rules, it provides a way of understanding it—by figuring out what objective it is intended to achieve.

The central assumption of economics is *rationality*—that behavior can best be understood in terms of the purposes it is intended to achieve. The secondary assumption running through this book is that systems of legal rules, or at least large parts of systems of legal rules, make sense—that they can be understood as tools with purposes. The rationality assumption will not be questioned here, although there is an extensive literature elsewhere on the subject, of which the most interesting part, in my judgment, is the recent work in evolutionary psychology. The secondary assumption will be questioned repeatedly. One of the questions running through this book is to what degree the legal rules we observe can be explained as tools—in particular, as tools designed to achieve the particular purpose, economic efficiency, that economic analysis of the law most commonly ascribes to them. In chapter 19 I sum up the evidence and deliver a mixed verdict.

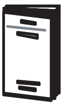

What Is Wrong with It

A system of legal rules is not entirely, perhaps not chiefly, the product of deliberate human design; to a considerable extent it represents the unplanned outcome of a large number of separate decisions, by legislators bargaining over particular provisions in the law or judges trying to find and justify verdicts for particular cases. It is therefore possible that such a system may have no objective for us to find. There is no guarantee that we will be able to make sense of any particular system of legal rules, since there is no guarantee that it makes sense. Human beings are born equipped with a superb pattern-recognition engine—so good that not only can we find patterns that even a well-designed computer would miss,

we can sometimes find patterns that aren't there. One of the questions you should be asking yourself, especially as you approach the end of the book, is to what degree economics discovers order in law and to what degree it imposes it.

One objection to the economic approach to understanding the logic of law is that law may have no logic to understand. Another and very different objection is that law has a logic but that it is, or at least ought to be, concerned not with economic efficiency but with justice. We punish criminals not, or at least not entirely, because doing so achieves good consequences but because criminals deserve to be punished. We require tortfeasors to make their victims whole not because doing so gives people an incentive not to be tortfeasors but because it is just that he who did the damage should pay for it. On precisely the same grounds, we insist that if our child has made a mess, he should clean it up.

To this very persuasive line of argument I have two answers. The first is that justice does not give an adequate account of law, both because it is irrelevant to a surprisingly large number of legal issues and because we have no adequate theory of what makes some rules just and some unjust. To a considerable degree, our intuitions of justice are consequence, not cause—we think rules are just because they are the rules we have been brought up with.

My second answer is that in many, although probably not all, cases it turns out that the rules we thought we supported because they were just are in fact efficient. To make that clearer I have chosen to ignore entirely issues of justice going into the analysis. In measuring the degree to which legal rules succeed in giving everyone what he wants, and judging them accordingly, I treat on an exactly equal plane my desire to keep my property and a thief's desire to take it. Despite that, as you will see, quite a lot of what looks like justice—for example, laws against theft and the requirement that people who make messes should clean them up—comes out the other end. That, I think, is interesting.

And for Whom I Am Doing It

This book is aimed at three different sorts of reader. The first is the proverbial intelligent layman—someone who thinks it would be interesting to know about law and economics and what they have to do with each other, himself, and the world in which he lives and so is reading this book for the same sort of reasons that make me read *The Selfish Gene* or *The Red Queen.* The second is the legal professional who would like to know more about the economic approach to his field. The third is the student,

most probably in an economics department or a law school, who is reading this book because his professor told him to—and will, I hope, find that that is not the only reason to do so.

One problem in writing for different sorts of readers is that they want different sorts of books. Students, especially law students, and, to some degree, legal professionals expect a scholarly apparatus of footnotes, case cites, extensive bibliographic references, and the like that the intelligent lay reader is likely to find clumsy and unnecessary. I have dealt with that problem by moving the scholarly apparatus to cyberspace. This book is written for the lay reader, with no footnotes and few case cites or references. To go with it, I have produced a web site containing, I hope, everything that the student or legal professional will find missing in the hard copy currently in his hands. To link the two, the margin of the book contains icons representing links on the book's Web page

www.best.com/~ddfr/laws_order/:

Cite to a book or article

Link to a webbed book or article

Case

Math

Additional comments

One reason I wrote the book this way is that I have a somewhat mixed view of scholarly apparatus, even in an academic context. Certainly it is useful to have pointers to cases, articles, and the like readily available. But this book is fundamentally about a structure of ideas, and it is easy to lose track of that structure in a maze of academic detail—easy not only for reader but also for writer. I still remember with distaste the first chapter

of an early book in this field that consisted entirely of hooks to hang footnote references on, with scarcely a sentence that would convey any real information to a reader who did not already know what the book was supposed to be teaching.

My hope is that by paring the book down to what it is really about and taking advantage of modern technology to put everything else somewhere out of the way but within easy reach, I can achieve the benefit of the apparatus without the costs. At the same time I also provide myself a place for continued revision and expansion—without the need for any expensive resetting of type. Readers who want to help with that process will find my e-mail address readily available on the book's Web page.

And, Finally, a Road Map

There are two ways to organize the economic analysis of the law—economic or legal, by economic ideas or by areas of law. In this book I do both. The first part sketches basic economic concepts—rationality, economic efficiency, externalities, value of life, economics of risk allocation, et multae caetera—that can be used to understand a wide range of legal issues. It is followed by a one-chapter intermezzo in which I sketch out how our particular legal system is put together, primarily for the benefit of those readers who are neither lawyers nor law students. The second part then applies the economics to the analysis of the core areas of law—roughly speaking, the courses a law student will take in his first year—and is organized accordingly.

The concluding part applies what we have at that point learned in a variety of different ways: a chapter on legal systems very different from ours (including one located a few hours from where I am sitting), a chapter on the question of why we have two legal systems—tort law and criminal law—to do roughly the same thing in different ways and whether we could dispense with one of them, and a chapter considering the evidence for and against the claim that law, at least judge-made law, is economically efficient. The book ends with a final chapter in which I attempt to sum up what we have learned about systems of legal rules.

1

What Does Economics Have to Do with Law?

YOU LIVE IN A STATE where the most severe criminal punishment is life imprisonment. Someone proposes that since armed robbery is a very serious crime, armed robbers should get a life sentence. A constitutional lawyer asks whether that is consistent with the prohibition on cruel and unusual punishment. A legal philosopher asks whether it is just.

An economist points out that if the punishments for armed robbery and for armed robbery plus murder are the same, the additional punishment for the murder is zero—and asks whether you really want to make it in the interest of robbers to murder their victims.

That is what economics has to do with law. Economics, whose subject, at the most fundamental level, is not money or the economy but the implications of rational choice, is an essential tool for figuring out the effects of legal rules. Knowing what effects rules will have is central both to understanding the rules we have and to deciding what rules we should have.

The fundamental assumption of the economic approach, to law and everything else, is that people are rational. A mugger is a mugger for the same reason I am an economist: Given his tastes, opportunities, and abilities, it is the most attractive profession open to him. What laws are passed, how they are interpreted and enforced, ultimately depend on what behavior is in the rational interest of legislators, judges, and police.

Rationality does not mean that a burglar compiles an elaborate spreadsheet of costs and benefits before deciding whether to rob your house. An armed robber does not work out a precise analysis of how shooting his victim will affect the odds of being caught, whether it will reduce the chances by 10 percent or by 20. But if it is clear that it will reduce the risk of being caught without increasing the punishment, he is quite likely to pull the trigger.

Even in this weaker sense people are not always rational. I, for example, occasionally take a third helping of spaghetti when a careful calculation of my own long-run interests would lead me to abstain. I am well acquainted with my own irrationality and can take steps to deal with it. Having discovered that bowls of potato chips located within arm's reach empty themselves mysteriously, I at least sometimes take the precaution of putting the bowl somewhere else.

But I do not know other people—the vast masses of other people to whom economic analysis of law is intended to apply—well enough to incorporate their irrationalities into my analysis of the effect of legal rules on their behavior. What I do know about them is that they, like me, have purposes they wish to achieve and tend, albeit imperfectly, to correctly choose how to achieve them. That is the predictable element in human behavior, and it is on that element that economics is built.

Whether armed robbers should get ten years or life is not a burning issue for most of us. A question of considerably more importance is the standard of proof. In order for you to be convicted of a crime or to lose a civil case and have to pay damages, just how strong must the evidence against you be?

It is tempting to reply that nobody should be punished unless we are certain he is guilty. But by that standard nobody would ever be punished; the strongest evidence establishes only a probability. Even a confession is not absolute proof: While our legal system no longer permits torture, it does permit plea bargaining, and an innocent defendant may prefer a guilty plea on a minor charge to risking a long prison term on a major one. Scientific evidence is no more conclusive; even if we somehow had a perfect match between the DNA of the suspect and the criminal, there would still be the possibility that someone at the lab made a mistake or that somewhere, perhaps unknown to him, the suspect has an identical twin. If we are to convict anyone at all, we must do it on evidence short of absolute proof.

How far short? Raising the standard of proof reduces the chance of convicting an innocent defendant but increases the chance of acquitting a guilty one. Whether that is on net worth doing depends on the relative costs of the two kinds of mistakes. If, as Blackstone wrote more than two hundred years ago, it is better that ten guilty men go free than that one innocent be convicted, we should keep raising our standard of proof as long as doing so saves one more innocent defendant at the cost of freeing no more than ten guilty ones. We would end up with a high standard.

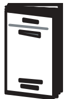

In fact, law in the United States and similar systems requires a high standard of proof ("beyond a reasonable doubt") in a criminal case but only a low standard ("preponderance of the evidence") in a civil case. Why? The answer cannot simply be that we are more careful with criminal convictions because the penalties are bigger. A damage judgment of a million dollars, after all, is a considerably more severe punishment for most of us than a week in jail.

Economics suggests a simple explanation. The typical result of losing a lawsuit is a cash payment from the defendant to the plaintiff. The result of being convicted of a crime may well be imprisonment or execution. A

high error rate in civil cases means that sometimes I lose a case I should have won and pay you some money and sometimes you lose a case you should have won and pay me some money. On average, the punishment itself imposes no net cost; it is simply a transfer. A high error rate in criminal cases means that sometimes I get hanged for a murder I didn't commit and sometimes you get hanged for a murder you didn't commit. In the criminal case, unlike the civil case, one man's loss is not another man's gain. Punishment is mostly net cost rather than transfer, so it makes sense to be a good deal more careful about imposing it.

For an application of economics to a different part of the law, consider the nonwaivable warranty of habitability, a legal doctrine under which some courts hold that apartments must meet court-defined standards with regard to features such as heating, hot water, sometimes even air conditioning, whether or not such terms are provided in the lease—indeed, even if the lease specifically denies that it includes them. The immediate effect is that certain tenants get services that their landlords might not otherwise have provided. Some landlords are worse off as a result; some tenants are better off. It seems as though supporting or opposing the rule should depend mainly on whose side you are on.

In the longer run the effect is quite different. Every lease now automatically includes a quality guarantee. This makes rentals more attractive to tenants and more costly to landlords. The supply curve, the demand curve, and the price, the rent on an apartment, all shift up. The question, from the standpoint of a tenant, is not whether the features mandated by the court are worth anything but whether they are worth what they will cost.

The answer may well be no. If those features were worth more to the tenants than they cost landlords to provide, landlords should already be including them in their leases—and charging for them. If they cost the landlord more than they are worth to the tenant, then requiring them and letting rents adjust accordingly is likely to make both landlord and tenant worse off. It is particularly likely to make poorer tenants worse off, since they are the ones least likely to value the additional features at more than their cost. A cynical observer might conclude that the real function of the doctrine is to squeeze poor people out of jurisdictions that adopt it by making it illegal, in those jurisdictions, to provide housing of the quality they can afford to rent.

If my analysis of the effect of this legal doctrine seems implausible, consider the analogous case of a law requiring that all cars be equipped with sunroofs and CD changers. Some customers—those who would have purchased those features anyway—are unaffected. Others find that they are getting features worth less to them than they cost and paying for them in the increased price of the car.

This is a very brief sketch of a moderately complicated economic problem, and the result is not quite so clear as the sketch suggests. With a little effort one can construct possible situations in which a restriction on the terms of leases benefits some tenants and landlords at the expense of others, or most tenants, or most landlords. With more effort one could construct a situation in which the restriction benefits both landlords and tenants. The important point is not that restrictions on the terms of contracts are a good or a bad thing but that one cannot evaluate their effects by looking only at the terms that are restricted. You also have to look at the effect of the restriction on the other terms of the contract, in my example the rent.

In any particular law case it looks as though what is at stake is how the legal system will deal with this particular set of events, all of which have already happened. From that backward-looking point of view it is often hard to make sense out of existing law. The reason is not that law does not make sense but that we are facing in the wrong direction.

Suppose, for example, that I take advantage of a particularly good opportunity to push my rich uncle off a cliff. By extraordinary bad fortune a birdwatcher happens to have his camera pointed in my direction at just the wrong time, with the result that I am caught, tried, and convicted. During the sentencing phase of the trial my attorney points out that my crime was due to the conjunction of extraordinary temptation (he was very rich, I was very poor) and an improbably good opportunity—and I had only one rich uncle. Besides, once I have been convicted of this crime, potential future victims are unlikely to go rock climbing with me. Hence, he argues, the court should convict me and then let me go. Whatever they do, I will never kill again, and hanging or imprisoning me will not, he points out, bring my uncle back to life.

The conclusion is bizarre, but the argument seems logical. The reply many legal scholars would probably offer is that the law is concerned not only with consequences but also with justice. Letting me go may do no damage, but it is still wrong.

The economist offers a different response. The mistake is not in looking at consequences but in looking at the wrong consequences, backward at a murder that has already happened instead of forward at murders that may happen in the future. By letting me off unpunished, the court is announcing a legal rule that lowers the risk of punishment confronting other nephews faced, in the future, by similar temptations. Executing this murderer will not bring his victim back to life, but the legal rule it establishes may deter future murderers and so save those who would have been their victims. *Legal rules are to be judged by the structure of incentives they establish and the consequences of people altering their behavior in response to those incentives.*

Crime and contract are not the only parts of law in which the economic approach proves useful. Speeding fines are intended, not as an odd sort of tax, but as a way of making it in the interest of drivers to drive more slowly. Tort law determines what happens to people who get in auto accidents and thus affects the incentive to do things that might lead to being in an auto accident, such as not having your brakes checked, driving drunk, driving at all. The rules of civil procedure determine what sorts of information litigants are entitled to demand from each other and thus affect the incentive of firms to keep (or not keep) records, to investigate (or not investigate) problems with their products that might become the subject of litigation, to sue or not to sue. Divorce law determines under what circumstances you can get out of a marriage, which is one of the things relevant to deciding whether to get into it. The subject of economic analysis of law is law. All of it.

The Proper Application of High Explosives to Legal Theory

A physics student who has learned classical mechanics and the theory of electricity and magnetism has the basic equipment to deal with practically any pre-twentieth-century physics problem. Just add facts and mathematics and turn the crank. Throw in relativity and quantum mechanics and you can drop the "pre-twentieth-century" restriction. An economics student who has thoroughly mastered price theory is equipped to deal with very nearly every problem to which economic theory gives a clear answer, with the result that many of the courses offered by an economics department are simply applications of price theory to such particular areas as transportation, agriculture, trade, or law. A law student who has learned to understand tort law has the basic equipment to understand tort law. If he wants to understand criminal law, he must start over again.

Economics changes that. In the next few chapters you will be acquiring a set of intellectual tools. The rest of the book consists of the application of those tools to different areas of law. As you will see, once you understand property, or contract, or tort from the point of view of economics, you have done most of the work toward understanding any of the others. While each raises a few special issues, the fundamental analysis is common to all.

This is one explanation for the controversial nature of economics within the legal academy. On the one hand, it offers the possibility of making sense out of what legal academics do. On the other hand, it asserts that in order for legal academics to fully understand what they are doing, they must first learn economics. In the world of ideas, as in the world of geopolitics, imperialism is often unpopular with its targets.

A second reason economic analysis is controversial is that it sometimes produces conclusions with which many legal academics disagree—for example, that laws "protecting" tenants are quite likely to make tenants worse off. Scholars who apply economic analysis to law are routinely charged with conservatism, not in the literal sense of wanting to keep things unchanged (in that sense the traditional scholars in any field are the conservatives and the challengers the radicals) but in the current political sense.

There is some truth to this claim—more if "conservative" is changed to "libertarian." Part of the reason is the economist's underlying assumption that individuals are rational. While that assumption does not, as we will see, eliminate all reasons for wanting to interfere with market outcomes, it does eliminate many. And while rationality is an optimistic assumption when applied to individuals who are supposed to be acting for their own interest—buying and selling, signing contracts, getting married or divorced—it can be a pessimistic assumption when applied to people who are supposed to be acting in someone else's interests, such as judges or legislators. Their rationality may consist of rationally sacrificing the interests they are supposed to be serving, such as justice and the public good, to their own private interests.

But while economists are more likely to get some answers and less likely to get others than traditional legal scholars, the principal effect of economic analysis is to change not the conclusions but the arguments—for both sides of any controversial issue. It provides a powerful argument for the death penalty as deterrence but also, as we will see in chapter 15, a new argument against the death penalty. Applied to landlord tenant law, the most striking implication is that what legal rules you favor should depend very little on whether you care more for the interests of landlord or tenant. In most cases a bad law will hurt both groups and a good law help both, at least in the long run. In almost every application, economic analysis radically reshapes the arguments out of which legal conclusions come. One implication is that it is a tool or, if you prefer, a weapon, useful to people with a wide range of political agendas.

What the Law Has to Teach Economists

So far I have discussed economic analysis of law from the perspective of an economist, eager to show my legal colleagues why they must study economics if they hope to understand the law. The transaction is not, however, entirely one way. Economists have something to learn as well.

Economics applies its general theory largely to abstract concepts—property, exchange, firms, capital, labor. Quite a lot of what lawyers and

law professors do involves dealing with the same concepts in their real-world incarnation.

An economist can talk about someone owning a piece of land and assume that that is the end of it. A lawyer dealing with property is brought face to face with the fact that ownership of land is not a simple concept. How does my ownership of a piece of land apply to someone else who wants to fly over my land, dig a hole next to it that my house might slide into, permit his cattle to wander onto my land and eat my vegetables, erect a structure on his land that shades my swimming pool? And if someone does do something that violates my property rights, just what am I allowed to do about it—ask him to leave, blow him up with a land mine, or sue him for damages?

These issues show up in real cases that real judges and lawyers have to deal with. The more you think about them, the clearer it becomes that what you own is not a piece of land but a bundle of rights related to a piece of land. For example . . .

Someone builds a new hotel in Florida that shades the swimming pool of the next hotel down the beach. The owners of the old hotel sue for damages. Conventional economic analysis holds that they should win. The new hotel imposes a cost on the old; making its builder liable forces him to include that cost—what economists call an external cost or *externality*—in deciding whether or not the new hotel is worth building.

But, as Ronald Coase pointed out in an article that laid an important part of the foundation for the economic analysis of law, that answer is too simple. It may be true that if the builder of the new hotel is not liable he need not consider the cost he imposes by locating his building where it will shade his neighbor's pool. But if he *is* liable, the neighbor at an earlier stage need not consider the cost he imposes by locating his swimming pool where a building on the adjacent lot will shade it—and thus forcing the owner of that lot to either leave it empty or build and pay damages. What we have are not costs imposed by one person on another but costs jointly produced by decisions made by both parties.

Part of Coase's solution to this problem is to restate it in terms not of external costs but of property rights. One of the rights of value to the owners of both hotels is ownership of the stream of sunlight currently falling on the pool. If that right belongs to the owner of the land on which the sunlight now falls—the owner of the old hotel—then the builder of the new hotel can be held legally liable for interfering with it. If it does not, he cannot. The right is of value to the owners of both adjacent pieces of property: One needs it to protect his swimming pool, the other to permit him to build a building that will shade it.

The solution suggested by Coase was not liability but trade. Define the relevant legal rules so that one of the parties has a clear right to the stream

of sunlight. If it is worth more to the other—if the gain from building the hotel is more than the cost of moving the swimming pool—the other can buy it. Thus Coase, by looking at real cases in which courts had to decide among competing uses of land, radically revised the economic analysis of externalities—a topic we will return to in chapter 4.

How She Growed: The Three Enterprises of Law and Economics

Economic analysis of law comprises three closely related enterprises: predicting what effect particular legal rules will have, explaining why particular legal rules exist, deciding what legal rules should exist.

The first is the least controversial. While many people believe that the consequences of a law are not the only thing determining whether it is good or not, very few believe that consequences are irrelevant. To the extent that economic analysis helps us perceive consequences of laws and legal decisions, especially consequences that are not obvious, it is useful to anyone trying to make or understand law. If imposing a life sentence for armed robbery results in more murders, that is an argument, although not necessarily a decisive argument, against doing it. If restrictions on the terms of leases make both landlords and tenants worse off, that is an argument, probably a decisive argument, for letting them set the terms themselves.

The second enterprise is using economics to explain the existence of the legal rules that we observe. This is a hard problem. Legal rules are created by legislatures and courts—and we have no very good theory, economic or otherwise, to explain the behavior of either. From a theoretical standpoint, the project is part of the field of economics known as public choice theory, an area still very much on the intellectual frontier. It contains some interesting first steps, such as Niskanen's model of the budget-maximizing bureau and Becker's analysis of the political market on which interest groups bid for legislation, with more concentrated and better organized groups typically using government to benefit at the expense of less concentrated and worse organized groups, but it has not yet provided a fully worked out and generally accepted economic theory of government.

There is, however, one conjecture about law that has played a central role in the development of law and economics. This is the thesis, due to Judge Richard Posner, that the common law, that part of the law that comes not from legislatures but from the precedents created by judges in deciding cases, tends to be economically efficient. I was implicitly relying on that conjecture when I explained the difference between the standard

of proof required for criminal conviction and that required for civil conviction; my argument took it for granted that legal rules were somehow shaped in a way that properly traded off the costs of false convictions against those of false acquittals.

Why might one expect legal rules to be like that? One answer offered by Posner is that the two central issues with which we might expect judges to be concerned are efficiency (the effect of legal rules on the size of the pie) and distribution (their effect on who gets how much of it). Common law consists, in large part, of the legal framework for voluntary transactions. The result, as suggested by the earlier example of rental contracts, is that most distributional effects of changes in the law are illusory; when we compel a change in one term of a contract in favor of one party, other terms, such as the price, shift in the opposite direction, wiping out the distributional effect. If using the law to redistribute is difficult, it seems plausible that judges might leave redistribution to legislatures and concern themselves with efficiency instead.

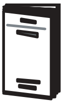

A very different argument offered by others for the same conclusion is that inefficient rules generate litigation, and litigation, eventually, generates changes in the rules. If some rule of the common law prevents people from doing things that are in their mutual interest, those affected will try either to change the law or work around it. Eventually they succeed. We are left with a common law shaped, "as if by an invisible hand," to maximize economic efficiency.

In addition to these theoretical arguments for why we might expect common law to be efficient, there is also the empirical argument, the claim that the common law legal rules we observe are, in most although not all cases, the rules we would get if we were trying to design an economically efficient legal system. Posner's immensely productive career as a legal theorist has largely consisted of piling up evidence for that argument. One of the things we will be doing in future chapters is examining that evidence, comparing the implications of economic theory with the laws we observe. In chapter 19 we will return to the Posner thesis in order to sum up the theoretical and empirical arguments for and against.

The Posner thesis that the common law is efficient leads naturally to the third and most controversial part of law and economics: using economic analysis to decide what the law should be. If we conclude that some particular common law rule—say, the nonwaivable warranty of habitability discussed a few pages back—is economically inefficient, that it makes us on net poorer, one conclusion is that Posner is wrong. Another might be that we should change it.

As a matter of simple logic, the claim that legal rules are efficient is entirely separate from the claim that they ought to be efficient. One might believe that laws should be efficient but are not, or that they are but

should not be, that other values should have greater weight than economic efficiency in determining the law. In practice, however, the two claims are easily confused and often combined. Posner makes both, although in each case with substantial qualifications.

I have repeatedly referred to "economic efficiency" without ever explaining precisely what the words mean. In this case as in many others it is dangerous to assume that a word used as a technical term has the same meaning as in other uses. "Strike" means very different things in baseball, bowling, and labor relations. "Efficient" means very different things applied to engines, employees, and economies.

Economic efficiency can most usefully be thought of as the economist's attempt to put some clear meaning into the metaphor "size of the pie." What makes doing so difficult is that the relevant pie is not a single object that we can weigh or measure but a bundle of many different sorts of goods and services, costs and benefits, divided among hundreds of millions of people. It is not obvious how those can be all put in common units and summed to tell us whether some particular change in legal rules (or anything else) increases or decreases the total. Solving that problem will be the subject of the next chapter.

One question that should have occurred to you by now is whether any of this has anything to do with the real world. One way to answer that is to go back to the two examples I started the chapter with, encouraging robbers to kill their victims and making apartments more expensive. The question you should be asking is not whether you are convinced that my analysis of those examples is correct—what I was offering, after all, was only a sketch of an argument. The question is whether you understand more about those issues than you did before you read this chapter. If the answer is yes, then the economic analysis of law has something to do with the real world.

A second way to answer the question is to consider whether you believe that people are, on the whole, rational. If we know that doing something will make someone better off, is that a good—not certain, but good—reason to expect him to do it? If the answer is "yes," are you willing to generalize, to apply it to police, judges, legislators, burglars, muggers, and potential victims? If the answer is still "yes," then you are in agreement with the fundamental assumption on which the theory is built.