

5

Law and Economics

THE PREVIOUS FOUR CHAPTERS have explored several of the central issues in the philosophy of law. It should be clear by now both that there is a good deal about the law of interest to philosophers and that recent work in legal philosophy has greatly extended its frontiers. Nevertheless, there have been no attempts by philosophers to develop a unified, comprehensive moral or philosophic theory of law. No philosopher has claimed that the substance of such diverse areas of the law as crimes, torts, property, contracts, civil procedure, corporations, and family law could possibly derive from any one moral principle or set of consistent moral principles. There are several reasons for this. First, the claim that one could derive the substance of such complex and diverse bodies of law strikes the cautious intellect as a bit preposterous on its face. Second, most philosophers lack the requisite knowledge of law to attempt such a broad sweep. Third, it is highly unlikely that philosophers could ever agree upon one principle or set of principles as the one or ones the law seeks to enforce. To the extent there is widespread agreement among philosophers, we converge on method (on ways of framing issues and pursuing solutions to them) rather than on substance (on a particular principle that systematically ought to be employed to solve problems). Philosophy of law has no "research program" as such.

While there is nothing we can refer to as a research program in legal philosophy, such a program has emerged in the field of law and economics. Over the last several decades economists and lawyers trained in or enamored of economics have sought to explore the extent to which virtually all areas of the law could be understood as the institutional embodiment of the principle of economic efficiency. It is strange that economists do not find the claim that all of the law could derive from a single principle at all preposterous. The work in law and economics has had both analytic and normative dimensions. The analytic work has aimed at demonstrating that large areas of law could be explained by seeing them as concerned not so much with matters of justice but with the efficient allocation of resources.

The normative work in the field is concerned to give legislators and judges a framework for legislating and adjudicating cases so as to promote the goal of efficiency.

This chapter explains the economic analysis of law. It requires absolutely no prior knowledge of economics, nor does it presuppose any familiarity (by the student or the instructor) with the law and economics literature. Consequently, I have broken the chapter into three general sections. In the first section, I give definitions of the key concepts in the analysis and present the basic models that the analysis uses. In the second section, I show how the economic analysis can be used to elucidate various areas of the common law. In this section, I demonstrate how the economic analysis can be applied to problems of pollution, automobile accidents, and breach of contract. I also discuss briefly the results of the first important discussion of the economics of crime. In the third section, I consider several objections—some good, some popular but not so good—to various features of the economic analysis of law.

Basic Concepts and Models

Understanding Efficiency

One reason philosophers of law should take economic analysis seriously is because the most basic notion in the analysis—efficiency or Pareto optimality—was originally introduced to help solve a serious objection to the widely held moral theory, utilitarianism. Utilitarians hold that the principle of utility is the criterion of right conduct. Whether one advances total or average utilitarianism, applying the principle requires interpersonal comparisons of utility.¹ How can we evaluate policies in virtue of their effect on individual welfare or utility, if one person's utility cannot be compared with that of another? But the claim that we can compare utilities is quite controversial, and the Pareto criteria were first introduced to obviate the problem of interpersonal comparability. There are two Pareto criteria: *Pareto optimality* and *Pareto superiority*. Pareto superiority ranks or orders social states according to the following conditions:

Definition₁: One state of the world, S_1 , is Pareto superior to another, S , if and only if no one is worse off in S_1 than in S , and at least one person is better off in S_1 than in S .

Whether or not a person is better off in one state or another usually depends on his relative welfare, and each person is presumed to be the judge of his relative well-being. Pareto superiority is sometimes characterized in the following way as well:

Definition₂: S_1 is Pareto superior to S if and only if no one prefers S to S_1 and at least one person prefers S_1 to S .

The Pareto superiority criterion obviates the interpersonal comparability problem of classical utilitarianism. Because no one is made worse off, there are no losers in Pareto improvements whose losses are to be subtracted from, i.e., has a greater utility in S_1 than in S , and no one is worse off in S_1 than in S , then, going from S to S_1 , there is a net gain in total utility.

We can now introduce the derivative concept of Pareto optimality:

Definition₃: S_1 is Pareto optimal if and only if there exists no S_n such that S_n is Pareto superior to S_1 .

A Pareto optimal state has no states Pareto superior to it. When resources are distributed in a Pareto optimal fashion, there is no way of making anyone better off without making someone else worse off. Pareto optimal states are the eventual outcome of a sequence of Pareto superior moves, though one can reach a Pareto optimal state through a sequence of non-Pareto superior moves. Consider the following simple example. Suppose there is only one commodity, C ; and ten units of it; two persons, X and Y , both of whom prefer more C to less. Imagine the following two distributions of resources between X and Y .

S_2 : X has 10 units of C ; Y has 0 units of C .

S_1 : X has 0 units of C ; Y has 10 units of C .

Both S_2 and S_1 are Pareto optimal. Any change from S_2 , for example to S_1 , can enhance Y 's well-being only by diminishing X 's. Just the same, though in reverse, obtains in any move from S_1 to S_2 . In going from S_2 to S_1 or from S_1 to S_2 we secure an optimal outcome through a non-Pareto superior change. Suppose the initial distribution is represented as S_0 :

S_0 : X has 5 units of C ; Y has 5 units of C .

Here the initial distribution is Pareto optimal as well. It is easy to see that there are a large number of Pareto optimal states that can be attained from a given initial distribution of resources. Some of these involved a sequence of Pareto superior moves; some involve no Pareto superior moves; others involve a mixture of Pareto superior and non-Pareto superior moves.

With respect to one another, Pareto optimal states are Pareto noncomparable; that is, they cannot be compared by the Pareto superiority criterion. Look at S_0 , S_1 , and S_2 . Each is optimal, but none is Pareto superior to either of the others.

The set of Pareto optimal states attainable from an initial distribution of resources is represented as points on the so-called utility/possibility, or Pareto, frontier. Consider Figure 5.1: The origin represents the initial distribution of resources. Every move to the northeast of the origin represents a Pareto improvement. A move to the north represents a gain for B ; a move east represents a gain for A ; a move northeast is a gain for both, or at

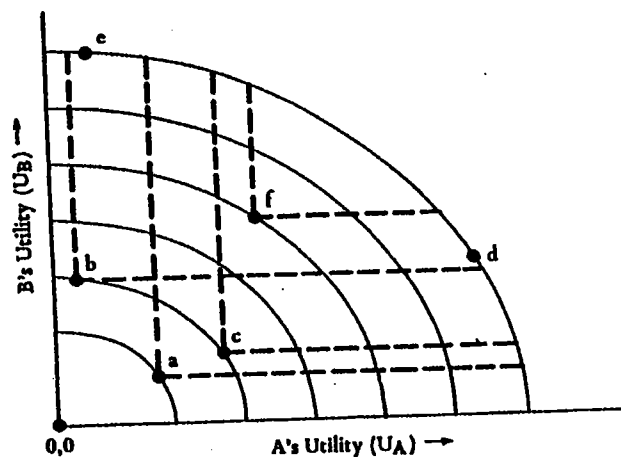


Figure 5.1 Utility/possibility frontier

least a gain for one at no one's expense. The points *a*, *b*, *c*, *d*, *e*, and *f* represent improvements from the origin for both A and B. Compare *a* and *b* and *a* and *c*. In going from *a* to *c* the lot of both A and B is improved. In going from *a* to *b*, however, B's utility is increased, but A's is not.² Now compare the moves from *a* to *d* and *a* to *e*. Both *d* and *e* are Pareto optimal, and Pareto noncomparable with respect to one another. In going from *a* to *d*, A and B attain an optimal state through a Pareto superior move.³ In going from *a* to *e*, A and B attain an optimal outcome through a non-Pareto superior move. So if A and B attain *d* first by moving to *a*, then *c*, they follow a Pareto superior path. If, however, they reach *e* first by going to *a*, then *b*, they secure a Pareto optimal outcome through a mixture of Pareto superior and non-Pareto superior steps.⁴

The graph of the utility/possibility frontier enables us to illustrate and explain a point economists make about the relationship between economics and political morality. The economist sees himself as trying to frame policies or rules that insure that a society reaches the Pareto frontier. Having done that, he is satisfied that the relevant rules, policies, or institutions are *efficient*—that is, no one can then secure a gain except at another's expense. But there are a large number of places on the frontier. Each is optimal, and each therefore satisfies the economist's efficiency requirement. They differ from one another *distributionally*, from the point of view of who has what. Recall our original example: *S*₀, *S*₁, and *S*₂ are all optimal; they differ in regard to the distribution of C between X and Y. For the economist, once we attain the frontier, it is the job of the social theorist/philosopher to determine which efficient outcome is "best." To use what turns out to be a somewhat inaccurate dichotomy, the economist sees the domain of efficiency as the size of the "pie," and the domain of justice or morality as the shape

and distribution of its pieces. In the economist's view, efficiency is prior to justice; first make the pie as large as possible, then slice it as "fairly" as possible.

This may be as good a time as any to discuss briefly one objection to this way of putting the relationship between efficiency and justice. Suppose that a point *within* the frontier is more "just" than any point *on* the frontier. Then in order to do justice one must forego efficiency. Justice is therefore prior to efficiency. The economist's response is that for every point within the frontier, there exists a point on the frontier that makes neither of the parties worse off and enhances the welfare of at least one of them. The parties could trade with one another from that or any other point within the frontier to the frontier.⁵ In the absence of adverse third-party effects,⁶ what grounds could we have for preferring the point within the frontier to the point on the frontier, i.e., an outcome that starts from a just distribution and involves only voluntary exchange. The same argument could be made with respect to all points within the frontier; i.e., for every point within the frontier that is desirable for its "justness" there exists at least one point on the frontier attainable by voluntary transfer. If a point within the frontier is just, we can describe a process that increases efficiency and preserves justice.

Pareto superior policy changes increase net utility, thus obviating the interpersonal comparability problem of classical utilitarianism. What is the relationship between Pareto optimality and utility? A move to a Pareto optimal state need not increase net utility. That is easy to see, since what makes a state Pareto optimal depends on whether there are any Pareto superior moves that can be made *from* it, not whether the move *to* it is itself Pareto superior. In our graph, the move from *f* to *e* is a move to a Pareto optimal state from a nonoptimal one, yet there is no reason to believe that in making such a move overall utility is increased: B is made worse off, A better off. Because the move is not to the northeast of the origin, it is not a Pareto superior one. We cannot know then whether there has been a gain in utility unless we can compare A's and B's utility.

It also does not follow that a point on the frontier represents a utility maximum. This is a somewhat more difficult point. To say that a state of the world is Pareto optimal is just to say that there are no Pareto improvements in store. Pareto improvements are one way of *validating* whether a change in policy is utility maximizing, a way that does not require interpersonal comparability. It does not follow that a move from one social state to another cannot be utility maximizing even if it is not Pareto superior. It is just that in *determining* whether a non-Pareto superior move is utility maximizing we need to compare utilities. And if we cannot compare utilities we cannot *know* if such a move increases utility. (Of course, if we could compare utilities there would be considerably less need for the Pareto superiority criterion.) In short, from the fact that a social state is Pareto optimal nothing follows about whether the move to it is a utility-maximizing one, nor does anything follow about whether any further utility-maximizing moves can

be made. What does follow is that no utility-maximizing move that increases net utility in virtue of its satisfying the Pareto superior criterion is possible. So it is the Pareto superior criterion only that entails any judgments at all regarding total utility.

The Pareto superiority criterion is limited, however, in a very basic way. It enables us to order or rank social states and thus to evaluate policies only when there are no losers. Policies can prevent losers in one of two ways: straightaway or by compensation *ex post*. If you and I exchange goods in a way that makes us both better off, we satisfy the Pareto superior criterion in a straightforward fashion. Suppose, however, you engage in a risky activity—e.g., blasting—because even if you had to compensate me for damages should an occasional blast cause me property damage, you would still be better off than if you did not engage in the activity at all. Now you blast, gain a \$100 thereby, and cause me \$20 in property damage. The move from the state in which you forego blasting to the state in which you blast and cause me damage is *not* Pareto superior because I am worse off. If, however, you compensate me fully for my loss, the move to the state wherein you blast, damage, and compensate from the state in which you do not engage in blasting is a Pareto improvement. Compensation *ex post* is often a key step in making Pareto improvements.

The *possibility* of compensation is important in a different way. The Pareto superiority standard applies only where there are no losers. Most social policies and legal rules produce losers as well as winners. The Pareto test is therefore nearly useless in regard to the evaluation of most activity of concern to the social, political, or legal theorist. The Kaldor-Hicks test, based on the *possibility of compensation*, was introduced to obviate this problem and to extend the usefulness of the Pareto rankings:

Definition: S_1 is Kaldor-Hicks efficient to S if and only if in going from S to S_1 the winners *could compensate* the losers so that no one would be worse than he or she was in S and at least one person would be better off than he or she was in S .

Another way to put this is to say that S_1 is Kaldor-Hicks efficient to S provided that were compensation paid no one would prefer S to S_1 and at least one person would prefer S_1 to S . The Kaldor-Hicks criterion enables us to evaluate social policies that produce winners *and* losers. The difference between Pareto superiority and Kaldor-Hicks efficiency is just the difference between *actual* and *hypothetical* compensation. If compensation were paid to losers a Kaldor-Hicks efficient move would become a Pareto superior one. Therefore, the Kaldor-Hicks criterion is often called the Potential Pareto superiority test.

If compensation could be paid, why isn't it? That is not as hard a question as it would seem. The reasons are usually of two sorts: First, some losers deserve to lose—for example, when policies are implemented to break up inefficient monopolies. There is no reason to render monopolists no worse

off after breaking up their monopolies than they were while engaged in monopolistic behavior. Second, it is often very costly actually to compensate losers. Compensation is a transaction and has associated with it certain costs: transaction costs. Suppose Jones gains \$8 in the move from S to S_1 , and Smith loses \$6. The move from S to S_1 is therefore a Kaldor-Hicks improvement, because Jones could compensate Smith \$6 and still be ahead \$2. Now suppose that Jones and Smith are unknown to one another. The search costs alone are likely to exceed \$2. If they do, then actually requiring Jones to compensate Smith would make Jones worse off than he was at S . Actually compensating Smith would not be Pareto superior. This brings us to a general point. When the Kaldor-Hicks criterion is employed, the "hypothetical compensation" condition assumes that compensation is to be *costlessly* rendered. Actual compensation is not costless, however, and that is primarily why it is not paid.

Let us conclude this subsection by reviewing the definitions of the key efficiency criteria in terms of the analytic relationships among them:

1. S_1 is Pareto superior to S provided no one prefers S to S_1 and at least one person prefers S_1 to S .
2. S_1 is Pareto optimal provided there is no S_n Pareto superior to S_1 .
3. S_1 is Kaldor-Hicks (K-H) efficient to S provided that the winners at S_1 , *could* compensate the losers so that no one would then prefer S to S_1 and at least one person would prefer S_1 to S . S_1 is K-H efficient to S provided S_1 is Potential Pareto superior (P-P-S) to S .

When economists talk about efficiency they almost invariably mean Pareto optimality. When lawyers who advocate economic analysis talk about efficient legal decisions, rules, or policies it is considerably less clear whether they have Pareto optimality, superiority, or Kaldor-Hicks efficiency in mind. The differences, as we shall see below, are by no means trivial.

The Coase Theorem

Much of the economic analysis of law grows up around the line of argument presented in Ronald Coase's "The Problem of Social Cost." Suppose Jones the rancher lives adjacent to Smith the farmer; Jones raises cows, Smith raises corn. There is no fencing separating their property. Jones's cows wander and destroy Smith's corn crop. For every additional cow Jones raises there is an associated reduction in Smith's corn crop. Each cow Jones raises imposes a *private cost* (the cost to Jones of raising it) and a *social cost* (the cost to Smith in damages). *Social costs* are external effects: effects of one person's conduct or consumption on the welfare of others. Some external effects are *positive*, i.e., they increase another's well-being. Others are *negative*; i.e., they decrease another's welfare. *Externalities* are *inefficient* external effects: i.e., non-Pareto optimal, external effects of one person's activity on another's welfare or utility.

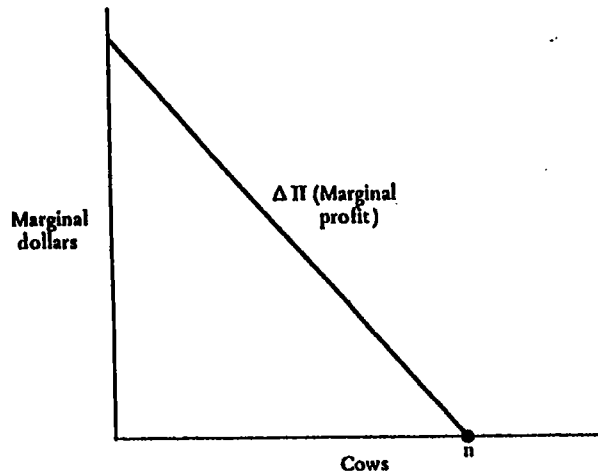


Figure 5.2 Marginal profit without internalized social costs

If Jones were *not* liable for the damage each of his cows imposes on Smith's corn crop, Jones would raise cattle until that point at which his marginal private costs equaled his marginal benefit (see Figure 5.2).

The economist distinguishes among total, average, and marginal benefit (or loss). If Jones profits \$50 from his first cow, \$40 from his second, \$30 from his third, \$20 from his fourth, \$10 from his fifth, and \$0 from his sixth, his *total* profit is $\$50 + \$40 + \$30 + \$20 + \$10 + \0 or \$150; his *average* profit is $150 \div 6$ or \$25; his *marginal* profit is \$50 for the first, \$40 for the second, and so on. Marginal benefit (loss) is that corresponding to each additional unit of production (consumption). Economic analysis of law is concerned primarily with the concept of marginal gain/profit/loss/welfare, and when we talk about economic gain (loss) we should be understood to mean marginal gain or loss, unless we explicitly state otherwise.

It is plausible to assume that the number of cows Jones would raise were he responsible only for his private costs exceeds the number he would raise were he liable for the damage his cows cause. That is because Jones ceases to raise cows when his marginal cost equals his marginal benefit, and imposing liability upon him for Smith's losses increases his marginal costs while holding his marginal benefit constant. On the other hand, in the absence of Jones's liability, Smith not only has to shoulder the private costs of farming corn, but the social costs of Jones's ranching as well. So were Jones liable in damages to him, Smith would grow more corn. The amount, and therefore the price of both beef and corn is directly affected by the question of who bears the social costs of Jones's ranching activity. Who

should bear the social costs of ranching on farming? First we must determine how much corn and how many cows ought to be ranched. Then we can determine if the external effects of Jones's cows cause an inefficient number of cows to be raised and corn to be farmed. But how can we figure out whether the amount of cows and corn is efficient without first knowing who is to bear the social costs of ranching and farming? We can, and here's how.

Instead of rancher Jones and farmer Smith on adjacent plots of land, imagine there is only Jones-Smith who is both a rancher and a farmer and owns both plots. The question for Jones-Smith is how much of each activity to engage in. Posing the problem this way makes the costs of ranching on farming part of Jones-Smith's *private* cost calculation. The very idea of a social cost is thereby eliminated. The process by which external or social costs are made part of one's private cost accounting is called "internalizing externalities."

At each stage in his decisionmaking Jones-Smith asks himself the following question: Will I secure more profit by raising the next cow (where my profit is equal to the benefit minus the cost of raising the cow *and* the cost of foregone corn crop profits) than I would secure from foregoing the cow for the corn? As long as he answers in the affirmative, he raises cows; when his answer is in the negative, he stops raising cows in favor of corn. He stops where his marginal benefit and cost curves intersect. Notice, he might have gone through the same process starting with corn. He would have asked himself: When does the marginal cost of farming corn exceed its marginal benefit? Here the costs are the sum of the standard farming costs *plus* the costs of foregone cattle. Again, Jones-Smith stops where his cost and benefit curves intersect (see Figure 5.3).

One gets the same result whether one begins with corn or cows. This is important as we shall see below. Moreover, the point at which the curves intersect marks the Pareto optimal allocation of Jones-Smith's resources since at that point any further cattle can make Jones-Smith the *rancher* better off only by making Jones-Smith the *rancher-farmer* on the whole worse off.⁸ The same holds in reverse. And there we have it. A fail-safe way of determining who should bear the social costs of ranching on farming.

Back to the example in which Jones and Smith are two distinct persons who own adjacent properties. How do we duplicate the result of this illustration? We could follow Pigou. Pigou argued that in order to internalize externalities a civil authority should impose a marginal tax on the offending party set equal to the marginal damage imposed by the offense. (Pigou also thought that subsidies should be awarded to encourage *positive* externalities.) To make our example concrete, assume the following schedule of profits and damages:

	Marginal Profit to Rancher	Marginal Damage to Farmer
Cow 1	\$50	\$10
Cow 2	\$40	\$20

Cow 3	\$30	\$29
Cow 4	\$20	\$40
Cow 5	\$10	\$49
Cow 6	\$ 0	\$57

If we impose a Pigouvian tax, Jones will raise the first cow and pay \$10 in taxes; the second and pay \$20; the third and pay \$29. He will not raise a fourth cow. This is the same result that would have been obtained given these figures were Jones and Smith one person. In other words, if Jones and Smith were in fact one person, then in order to raise a fourth cow he would have to incur \$40 in lost corn. Thus, because a fourth cow creates a net marginal loss of \$20 (\$20 gain, \$40 loss), he would not raise it. Thus, the Pigouvian tax secures the efficient outcome by getting Jones to internalize the relevant externalities.

In the Pigouvian tax approach to externalities, the state must identify one of the parties as the *cause* of the externality and impose and collect a marginal tax set equal to actual marginal damages. In "The Problem of the Social Cost," Ronald Coase presents an alternative to the Pigouvian approach that denies the necessity of both causation and taxation to efficiency. Indeed, Coase goes further than denying the relevance of causation to the goal of efficiency; he denies the existence of nonreciprocal causal relations. His argument, which economists by and large take seriously but which they ought not to, is as follows. The rancher's cows cause the farmer a loss. But if the state restricts the rancher from raising cows because cows trample




Figure 5.3 Marginal profit with internalized social costs

corn, it causes the rancher a loss. There is, then, a reciprocity in the causation of harm; either we cause the farmer harm by permitting the rancher's cows to graze or we cause the rancher harm by prohibiting his cows from trampling the corn. The central question is not which activity causes harm—they both do. The question is which harm (and how much of it) should we permit.⁹

Of course Coase is wrong. He treats the harm the rancher's cows cause the corn as if it were the same sort of harm the states does by reducing the level of the rancher's activity. Surely, whether the state interferes or not, cows destroy crops. That is the very plain sense in which causal relations are not reciprocal.

This blunder aside, Coase can be read sympathetically not as denying the existence of causal relations, but as denying their relevance to the pursuit of efficiency. The argument is this: Assume that the rancher and the farmer act cooperatively and that transactions between them are costless. The state might assign the rancher the right to raise as many cows as he would like or it might give the farmer the right absolutely to prohibit cows.

Consider first the case in which the rancher has the absolute right to raise cattle. The value to him of the fifth cow is \$10, but the cost to the farmer is \$49. The farmer will suffer \$49 in damage if the rancher raises the fifth cow. The farmer and the rancher have incentives to strike a deal. The rancher wants at least \$10 to forego the fifth cow, and the farmer is prepared to pay him up to \$48 to forego it. At a price somewhere between \$10 and \$48 a deal between them will be struck and there will be no fifth cow. Similar reasoning leads to the rancher foregoing a fourth cow. The fourth cow nets him \$20, but costs the farmer \$40. Again a deal will be struck. The farmer will pay the rancher something between \$20 and \$40 in exchange for which the rancher will forego a fourth cow. The third cow nets the rancher \$30 and causes the farmer \$29 in damage. No deal will be struck since to forego a third cow the rancher will not accept less than \$30 and the farmer will not offer more than \$29. The same holds for the first and second cows. The result is that the rancher raises three cows even if he is assigned the right to raise as many as he would like to.

Consider now the case in which the farmer is given the right to prohibit all cows. The first cow causes the farmer \$10 in damage but is worth \$50 to the rancher. The rancher stands prepared to offer the farmer more than \$10 and less than \$50 for the right to ranch that first cow. A deal will be struck that will enable the rancher to raise that cow but that provides at least full compensation to the farmer. The same argument holds for the second and third cows, but not for the fourth or fifth. In the latter two cases the damage to the farmer exceeds the benefit to the rancher. The farmer is unprepared to take less than full reimbursement for damages, but the rancher is unprepared to pay that much since it exceeds the value of each cow to him. Once again, the rancher will raise three cows even though we began by assuming that the farmer had the right to prohibit all cows.¹⁰

We secure the same result regardless of the assignment of rights with which we begin. The result duplicates the result of the Pigouvian tax, which

duplicates the result we obtain when we imagine only one person, rancher-farmer, who aims to maximize his profits by finding the optimal levels of ranching and farming. Coase has given us a way of internalizing externalities through *private exchange*. Indeed, there is a sense in which, when transactions are costless and people behave cooperatively, there are no externalities. In effect, the rancher and farmer are just like rancher-farmer. Just as he aims to maximize the joint profits of ranching and farming, the upshot of their negotiations is to do the same.

We can summarize Coase as follows. Where exchange is costless: (1) identifying one party as the cause of the externality is unnecessary to achieve efficiency; (2) there is no need for a state to impose a tax on one or the other party because the efficiency results from private exchange; (3) no matter who we assume has the relevant right, the efficient distribution of resources will result since the rights only define initial bargaining positions; (4) the assignment of legal rights therefore is otiose to efficiency; and (5) the assignment of rights can affect the relative wealth of the parties.

All of these points but (5) have already been illustrated; (5) is easy to show. If we give the farmer the right to prohibit cows, there will be three cows, but in order to obtain the right to raise them the rancher would have had to pay the farmer, thereby increasing the farmer's wealth. On the other hand, if the rancher is assigned the right to raise cattle he will end up raising three, but the farmer would have had to buy him down, thus increasing the rancher's wealth.

What is usually called the Coase Theorem can be put as follows:

Definition: When transactions are costless and individuals act cooperatively, any assignment of legal rights will be efficient.¹¹

Compare now the Pigouvian and Coasian approaches to externalities.¹² We have specified the conditions under which the Coasian approach is efficient: costless transactions and cooperative behavior. What are the conditions under which the Pigouvian approach is efficient?

When the Pigouvian tax is imposed the parties might be foreclosed from further negotiations or continued negotiations might remain an option.¹³ Consider the case in which the parties remain free to negotiate. There are two options regarding the tax revenue. It can be given to the damaged party in the form of compensation for marginal damage or as a lump sum transfer, or it can be put in the general coffers. The Pigouvian tax is efficient only if compensation is paid, not otherwise. This may seem counterintuitive since the point of the tax is to force the injurer to take the social costs of his activity into account, which it succeeds in doing in either case. What accounts for the difference in efficiency? The tax transfer is necessary to induce the correct behavior in the victim. The following abbreviated benefit/damage schedule of the rancher-farmer negotiations illustrates this:

	Marginal Benefit (before tax)	Marginal Damage
Cow 1	50	10
Cow 2	40	20
Cow 3	30	29

We determined before that three cattle would be optimal. When the ranches pays the tax into the general coffers the schedule for the third cow looks like this:


	Marginal Benefit	Marginal Damage
Cow 3	\$1 (30-29 tax)	29

A third cow is worth only \$1 after taxes to the rancher, but that cow will still bring about \$29 damage to the farmer. The farmer, therefore, has an incentive to pay the rancher to forego a third cow, leading to an inefficient outcome. They both gain by acting inefficiently.

In contrast, if the tax revenue is awarded to the farmer on an actual damage basis he is *always* indifferent between a cow and there not being one. That is precisely what full compensation means. Because he is in general indifferent for all levels of ranching, he is indifferent regarding the third cow. The rancher, however, is not indifferent; he wants a third cow because there is a marginal gain of \$1 (\$30 minus \$29) in it for him. Thus, it is to his advantage to raise the third cow.¹⁴

Now consider the efficiency of the tax when further negotiations between injurer and the victim are impossible. There are two other relevant variables. The first is whether the tax revenue is transferred to the damaged party or placed in the general tax coffer. Again, there are two ways in which the tax revenues can be transferred. The first is on a marginal damage basis; the second is by a lump sum transfer, the basic condition of which is that the amount of transfer not be set to reflect damages, marginal or total. The second variable is whether the behavior of the injured party can affect net damages. That is, can the injured party take precautions that will reduce total costs?

The general problem can be set out as follows:



Let's begin with case 4. If no negotiations are possible and the tax revenue is not transferred to the damaged party and the damaged party's behavior cannot affect efficiency, the tax will be efficient. The reason is obvious. The only person's behavior that can affect damages is the offending party and his behavior is being taxed to reflect its damages. The tax leads him to engage in his activity at the efficient level and only his behavior counts. Provided negotiations are impossible, this line of argument implies that whenever the damaged party's behavior cannot affect net damages, the Pigouvian tax will be efficient whether or not the tax revenues are transferred to the injured party. The tax is therefore efficient in case 2 as well.

What happens when the injured party's behavior can affect net damages as in cases 1 and 3? Where there is no transfer the tax will be efficient. The tax leads the offending party to efficient behavior and the absence of compensation leads the injured party to adopt optimal precautions. By not being reimbursed for his losses, the injured party has an incentive to take precautions to reduce them. Holding the offending party's behavior constant, the victim will take just those precautions the cost of which are less than or equal to the gain to him in reduced damages—i.e., the optimal precautionary behavior. On the other hand, if the tax revenue is transferred to the injured party in the form of marginal damages equal to the damages he suffers, he is indifferent for all levels of damages. He has no incentive therefore to reduce them. If a change in his behavior is required to accomplish optimal cost reduction, a transfer set equal to marginal damage will lead to inefficiency since he has no incentive to reduce damages.

In summary, the optimality of the Pigouvian tax depends on three factors: (1) the possibility of further negotiations; (2) the scope and nature of any transfer of the tax revenue to the damaged party; and (3) the contribution of the relevant parties to the level of damages. Where further negotiations are possible the optimality of the tax depends on its revenue being given on marginal damage basis to the injured party, otherwise both parties will perceive a joint advantage in reducing the offending activity to inefficient levels. Where negotiations are not possible, the Pigouvian tax is optimal whenever *only* the behavior of the taxed party contributes to the level of damages. The tax is also optimal when the absence of a transfer acts as if it were a tax on the damaged party whose behavior affects net damages and who is thereby given sufficient incentive to take optimal precautions. Finally, where a transfer equal to marginal damages does occur, the damaged party whose behavior does affect net damages has no incentive to take the necessary precautions and the tax is inefficient.

Assigning Entitlements: Law as Market Mimicker

A central difference between the Coasian and Pigouvian approaches to externalities concerns the role of the state. While the Pigouvian approach does not require the state to assign rights to the damaged party that would preclude various levels of the harmful activity, it does require that the state

or its agents identify one of the parties as the *cause* of the damage, impose and collect a tax from that party, and decide how to distribute its revenues.

In contrast, when transaction costs are trivial or nonexistent, the state need only secure the integrity of the bargaining process by enforcing the resource allocations to which individuals negotiate. Because either assignment of property rights will prove efficient, there is no party on whom the state is required to confer the relevant property right. It is natural then to associate the Coasian property right approach with minimal government.

The Coasian approach entails minimal government only when transaction costs are trivial, since only then are negotiations efficient. When transaction costs are nontrivial, efficient negotiations may not occur. For example, if the value of a third cow to the rancher is \$30, and the value to the farmer of prohibiting it is \$29 (his damages), then there will be a third cow (which is efficient) only if the rancher is initially assigned a property right to ranch that cow or if transaction costs are less than \$1. If, however, the farmer is entitled to prohibit a third cow, the rancher will not purchase the right to raise that cow whenever the *costs* of the transaction exceed \$1. The rancher will not pay more than \$30 (which he would have to when transaction costs exceed \$1) to obtain an entitlement that is worth only \$30 to him. The farmer maintains the right to prohibit a third cow—a right he exercises—thus reducing the number of cows to two, which, *ex hypothesi*, is inefficient.

It follows that when transaction costs are not trivial, it matters how property rights are assigned. Because the assignment of rights can make a difference in the efficiency of resource allocations, we need a principle to guide the assignment of property rights. It is at this juncture that the law and economics literature which derives from the Coasian property rights alternative to Pigouvian taxes begins to focus on the work of Richard Posner.

Posner's most basic contribution is the following. Where the conditions of the Coase Theorem—zero transaction costs and cooperative behavior—are satisfied, the law need not assign property rights in any particular way. Market exchange will always insure efficiency. When these conditions are not met, the law should promote efficiency by mimicking the market. By "mimicking the market," Posner means that the relevant legal authorities ought to assign property rights to those parties who would have secured them through market exchange.¹⁵ If the rancher would have secured the right to a third cow via negotiations regardless of the initial property right assignment, mimicking the market requires the courts to give the rancher that right straightaway.

Posner's principle for assigning legal rights—the principle that law should mimic the market—is a simple directive to courts to allocate resources as the market would have. The market would have allocated resources efficiently. When the market breaks down, the law should produce the result the market would have. But there are two ambiguities in the "mimicking the market" slogan. The first concerns whether the law should mimic or replicate the rational autonomous feature of market behavior, or the efficiency of market outcomes.¹⁶ This becomes important when we inquire into the *normative*

basis of economic analysis. The second concerns which exchange market the law ought to mimic. There are two choices. The first is the market in which the parties are in fact negotiating with its particular structures, costs of transactions, individual endowments, etc. The alternative is the Coasian market of costless exchange taking place between cooperative, fully informed, perfectly rational individuals.¹⁷

For now we will assume both that the feature of markets the law ought to replicate is the efficiency of exchange *outcomes* and that the market the law ought to mimic is the Coasian market of rational, fully informed individuals completely cooperating with one another in an effort to maximize joint welfare (or profits) through mutually beneficial exchange.

Protecting Entitlements

Once assigned, property rights need to be secured and enforced. In one of the most important papers in the literature, Guido Calabresi and Douglas Melamed distinguish among three ways of protecting entitlements: by (1) property rules, (2) liability rules, and (3) inalienability rules.¹⁸

Property rules protect entitlements by enabling the right bearer to enjoin others from reducing the level of protection the entitlement affords him except as he may be willing to forego it at a mutually acceptable price. If a right is protected by a *liability rule*, a nontitled party may reduce the value of the entitlement without regard to the right holder's desires provided he compensates *ex post* for the reduction in value. The value of the reduction, i.e., damages, is set by a collective body, usually a court, and need not coincide with what the entitled party would have been willing to accept for a reduction in the value of his entitlement. Liability rules give nontitled parties the license to purchase on a pay-as-you-damage basis at objectively set prices part or all of an entitlement held by another. Property rules prohibit such transfers of entitlements in the absence of an *ex ante* between the relevant parties.

An example might illustrate this difference. Suppose the right to my cabin in the mountains is protected by a property rule only. In that case, if you wanted my cabin or any part of it then you would have to seek me out and convince me that I should transfer a part of my right to you, otherwise you have no claim or liberty with respect to my house. In contrast, if the right to my cabin were protected by a liability rule only then you need not secure my consent in order to avail yourself of it. Instead, you would be subject to liability for your use of it and for whatever damages I might incur. On the other hand, I have no basis for prohibiting you from action contrary to my right. Instead, the structure of our relationship is as follows. The right is mine, but my having it does not entail that you must first seek to negotiate with me over its use. You may do as you please provided you pay a "user fee," the proceeds of which are transferred to me. The same right may be protected by either a property rule, a liability rule, or both. If my cabin is protected by a property rule only, I have no redress if you are not adequately induced to avoid reducing its value to me

without my consent.¹⁹ That is one good reason for protecting a right with both a property rule and a liability rule. Sometimes property rules may suffice if there are insufficient incentives to impose damage without consent, or if there are sufficient incentives not to, if, for example, the criminal sanction is employed successfully as a secondary means of enforcement.

In contrast, it is sometimes necessary to forego property rules as vehicles for protecting entitlements in favor of liability rules. This occurs most often when transaction costs are high. If transaction costs are high, a property rule is likely to prove inefficient since transfer to more valued use requires negotiations. If negotiations are costly, property rules may lead to entitlements being held by individuals who value them less. Liability rules may therefore be substituted for property rules where transaction costs are high. Under a liability rule regime, individuals who value entitlements more than the individuals upon whom the rights are initially conferred are induced to secure the entitlements without prior agreement and to pay damages. In such cases, the party who values the entitlement most secures it, which duplicates the outcome of the market exchange process. If damages under liability rule set by a court are equal to or greater than the decrease in the value of the entitlement to the injured party, the optimal outcome is secured through a Pareto superior *forced transfer*. If damages are set below the value of the entitlement to the injured party, the forced transfer is not Pareto superior. Property rules induce optimal transfers through Pareto improvements. Whether liability rules involve Pareto improvements depends on the level of compensation.

To illustrate these points, reconsider the rancher-farmer example. Suppose that the farmer is entitled to prohibit all cows and that this right is protected by a property rule. In that case, before the rancher can raise any cattle he would have to purchase the right to do so from the farmer. Now suppose that the costs of transacting between rancher and farmer are \$5. In that case, the rancher would be unable to buy the right to a third cow. (He would have to pay \$34 to purchase something worth \$30 to him.) Thus there will be two cows at most, which would be inefficient. If the farmer's right were protected by a liability rule instead of a property rule, the rancher would simply raise the third cow and compensate the farmer \$29 *ex post*. The liability rule allows for resources to move to their efficient uses in some circumstances when a property rule would not. Of course, this entire line of argument depends on the costs of compensation *ex post* being low. If those costs are high, then neither a property rule nor a liability rule will be efficient.

Next, liability rules can lead to efficient outcomes either by Pareto superior or non-Pareto superior paths. If the damages the rancher has to pay are set equal to the value of the corn to the farmer, the farmer is no worse off, the rancher is better off, and a third cow is raised. Thus the Pareto superior path to the Pareto optimal outcome is taken. Suppose, however, that the farmer has a special attachment to the corn he would lose by the rancher's having a third cow, and that he would not trade away *ex ante* the

right to prohibit that cow for less than \$29.50. The damage to the corn, we have agreed, is only \$29. If he is compensated only for his "pecuniary" loss (\$29) and not his "nonpecuniary" loss (\$.50), he is made worse off by the liability rule. The third (optimal) cow is raised, but the liability rule did not induce a Pareto superior path to the Pareto optimal outcome.

Note as well that the liability rule functions much like the Pigouvian tax, with the revenue being transferred to the victim. Both are contemplated when a market for internalizing externalities is too costly to be feasible. We might say, then, that both are potential solutions to the problem of market failure. Moreover, notice how the liability rule is efficient only when the costs of compensation are low. Thus, it is analogous to a property rule that is efficient only when *ex ante* transaction costs are low. Note also how high transaction and compensation costs make the taking of Pareto superior paths to Pareto optimal outcomes infeasible. Precisely this sort of consideration makes the Kaldor-Hicks criterion (which does not require either *ex ante* or *ex post* compensation) so attractive to legal economists.

Liability rules and property rules differ from inalienability rules in that when a right is protected by an inalienability rule transfers of any sort are prohibited. Rights protected by inalienability rules are not transferable. The right to one's freedom from servitude and the right to vote are examples of rights protected by inalienability rules. Protecting a right by an inalienability rule may amount to a decision to forego efficiency in favor of promoting or protecting some other social good. Some people might well be induced to exchange their rights. Doing so might be efficient; blocking such transfers might then be inefficient. On the other hand, if there is a reason to believe that a willingness to exchange a right like that to freedom from servitude for monetary gain indicates a lack either of full information or rationality, protecting rights by inalienability rules might be justified on the grounds that such transfers would not occur in a costless market populated by fully informed rational persons. So there is an argument from efficiency, albeit a somewhat attenuated one, against permitting certain exchanges by the use of an inalienability rule.

Considerations of efficiency can generate a principle for allocating initial entitlements—e.g., Posner's mimicking the market principle—as well as various vehicles for protecting entitlements once assigned—e.g., the Calabresi-Melamed property rule/liability rule/inalienability rule distinction. It is time to turn our attention to the application of this framework to central problems in law.

Applying Economics to Law

Nuisance

Imagine a typical nuisance, a polluting feedlot, whose level of pollution increases with its output.²⁰ The damages incurred by its neighbors also

depends on output. Let us make the example concrete by the use of the following schedule of benefits and damages.

Output in 25 Unit Intervals	Marginal Profit (marginal benefit— marginal private cost)	Marginal Damage
25	200,000	10,000
50	100,000	20,000
75	75,000	30,000
100	50,000	40,000
125	30,000	50,000

Suppose a case initiated by the neighbors comes before a court. They seek compensatory damages or an injunction against the pollutor. What would be the efficient thing for the court to do? How should it rule?

Ignoring the unlikely use of an inalienability rule strategy, there are at least four things (following the Calabresi-Melamed distinctions) a court might do: (1) it could assign to the neighbors the right to be free from pollution and secure the right by a property rule; (2) it could assign the neighbors the right to be free from pollution, but protect it by a liability rule; (3) it could assign to the manufacturer the right to pollute and protect it by property rule; or (4) it could assign to the manufacturer the right to pollute and protect it by a liability rule.

If the neighbors' right to be free from pollution is protected by a property rule, then the manufacturer has to purchase from them the right to impose any pollution-related damages. If that right is protected by a liability rule, the manufacturer is in effect free to pollute provided it pays damages to its neighbors. If the court assigns to the manufacturer the right to pollute and protects that right by a property rule, then the neighbors must negotiate with the manufacturer to induce it to reduce the level of pollution. If the court assigns to the manufacturer the right to pollute but protects it by a liability rule only, then the neighbors are free to force a reduction in the level of output provided they pay damages. Here damages would be measured by foregone profits.

Which, if any, of these combinations is efficient? If transaction costs are zero (or trivial) all approaches can be efficient depending on the level of compensatory damages set by the court. In other words, if the court has full information and transactions between the manufacturer and its neighbors are free, then all four combinations are efficient.

Consider first the cases in which rights are protected by property rules. In effect, this amounts in the one case to the neighbors securing an injunction against the polluting manufacturer; in the other it amounts to the manufacturer securing an injunction against any efforts its neighbors might make to reduce the level of its output without its consent. The table indicates that 100 units of output is efficient.²¹ Where transaction costs are zero, the property right assignment backed by a property rule simply follows the argument for the Coase Theorem. If the right is assigned to the neighbors, the manufacturer