Polinsky, "The Coase Theorem"

A. Mitchell Polinsky, *An Introduction to Law and Economics* (2003), Chs. 3-4

One of the central ideas in the economic analysis of law was developed in an article by Ronald H. Coase in 1960. This idea, which has since been named the *Coase Theorem*, is most easily described by an example. Consider a factory whose smoke causes damage to the laundry hung outdoors by five nearby residents. In the absence of any corrective action each resident would suffer $75 in damages, a total of $375. The smoke damage can be eliminated in either of two ways: A smokescreen can be installed on the factory's chimney, at a cost of $150, or each resident can be provided an electric dryer, at a cost of $50 per resident. The efficient solution is clearly to install the smokescreen because it eliminates total damages of $375 for an outlay of only $150, and it is cheaper than purchasing five dryers for $250.

Zero Transaction Costs

The question asked by Coase was whether the efficient outcome would result if the right to clean air is assigned to the residents or if the right to pollute is given to the factory. If there is a right to clean air, then the factory has three choices: pollute and pay $375 in damages, install a smokescreen for $150, or purchase five dryers for the residents at a total cost of $250. Clearly, the factory would install the smokescreen, the efficient solution. If there is a right to pollute, then the residents face three choices: suffer their collective damages of $375, purchase five dryers for $250, or buy a smokescreen for the factory for $150. The residents also would purchase the smokescreen. In other words, the efficient outcome will be achieved regardless of the assignment of the legal right.

It was implicitly assumed in this example that the residents could costlessly get together and negotiate with the factory. In Coase's language, this is referred to as the assumption of *zero transaction costs*. In general, transaction costs include the costs of identifying the parties with whom one has to bargain, the costs of getting together with them, the costs of the bargaining process itself, and the costs of enforcing any bargain reached. With this general definition of transaction costs in mind, we can now state the simple version of the Coase Theorem: If there are zero transaction costs, the efficient outcome will occur regardless of the choice of legal rule.

Note that, although the choice of the legal rule does not affect the attainment of the efficient solution when there are zero transaction costs, it does affect the distribution of income. If the residents have the right to clean air, the factory pays $150 for the smokescreen, whereas if the factory has the right to pollute, the residents pay for the smokescreen. Thus, the choice of the legal rule redistributes income by the amount of the least-cost solution to the conflict. Because it is assumed for now that income can be costlessly redistributed, this distributional effect is of no consequence—if it is not desired, it can be easily corrected.

Positive Transaction Costs

The assumption of zero transaction costs obviously is unrealistic in many conflict situations. At the very least, the disputing parties usually would have to spend time and/or money to get together to discuss the dispute. To see the consequences of positive transaction costs,

suppose in the example that it costs each resident $60 to get together with the others (due, say, to transportation costs and the value attached to time). If the residents have a right to clean air, the factory again faces the choice of paying damages, buying a smokescreen, or buying five dryers. The factory again would purchase the smokescreen, the efficient solution. If the factory has a right to pollute, each resident now has to decide whether to bear the losses of $75, buy a dryer for $50, or get together with the other residents for $60 to collectively buy a smokescreen for $150. Clearly, each resident will choose to purchase a dryer, an inefficient outcome. Thus, given the transaction costs described, the right to clean air is efficient, but the right to pollute is not.

Note that in the example the preferred legal rule minimized the effects of transaction costs in the following sense. Under the right to clean air, the factory had to decide whether to pay damages, install a smokescreen, or buy five dryers. Because it was not necessary for the factory to get together with the residents to decide what to do, the transaction costs—the costs of the residents to get together—did not have any effect. Under the right to pollute, the residents had to decide what to do. Because the residents were induced to choose an inefficient solution in order to avoid the cost of getting together, the transaction costs did have an effect. Thus, even though no transaction costs were actually incurred under the right to pollute because the residents did not get together, the effects of transaction costs were greater under that rule.

We can now state the more complicated version of the Coase Theorem: If there are positive transaction costs, the efficient outcome may not occur under every legal rule. In these circumstances, the preferred legal rule is the rule that minimizes the effects of transaction costs. These effects include actually incurring transaction costs as well as the inefficient choices induced by a desire to avoid transaction costs.

The distributional consequences of legal rules are somewhat more complicated when there are transaction costs. It is no longer true, as it was when there were zero transaction costs, that the choice of the rule redistributes income by the amount of the least-cost solution. In the example, if the residents have the right to clean air, the factory pays $150 for the smokescreen, whereas if the factory has the right to pollute, the residents pay $250 for five dryers.

Although the simple version of the Coase Theorem makes an unrealistic assumption about transaction costs, it provides a useful way to begin thinking about legal problems because it suggests the kinds of transactions that would have to occur under each legal rule in order for that rule to be efficient. Once these required transactions are identified, it may be apparent that, given more realistic assumptions about transaction costs, one rule clearly is preferable to another on efficiency grounds. The more complicated version of the Coase Theorem provides a guide to choosing legal rules in this situation. All of the applications investigated in this book—nuisance law, breach of contract, automobile accidents, law enforcement, pollution control, products liability, principal-agent liability, and litigation—can be approached in this way, although some fit more naturally into the Coasian framework than others.

First Application – Nuisance Law

One area of law that can be readily examined in terms of the Coase Theorem is nuisance law. Nuisance cases result from incompatible land uses and typically involve a small number of individuals bargaining with each other, as when emissions from a factory fall upon neighboring

property, bright lights or noise disturb someone's sleep, or an unsightly building spoils an attractive residential neighborhood.

Adopting a framework first suggested by Guido Calabresi and A. Douglas Melamed, the resolution of a nuisance dispute may be viewed as involving two steps. First, an *entitlement* must be chosen—that is, a determination must be made regarding who is entitled to prevail. The injurer can be granted the right to engage in the activity that causes harm, or the victim can be granted the right to be free from harm. Then, a decision must be made about how to protect the entitlement. One possibility is to grant the holder of the entitlement an *injunction*. If the victim holds the entitlement, protecting it by an injunction means that he can prohibit the injurer from causing harm.[1] Thus, the injurer can cause damage only if he "buys off" the victim. Similarly, if the injurer holds the entitlement, protecting it by an injunction means that the victim must buy off the injurer if he wants damages reduced.[2]

An alternative method of protecting entitlements is to give the holder of the entitlement an amount of money—*damages*—that some governmental body, such as a court, determines. If the victim has the entitlement, he has the right to be compensated, but he cannot prohibit the injurer from causing harm as he could under an injunctive remedy. Analogously, if the injurer holds the entitlement, protecting it by a damage remedy would mean that the victim could restrict the injurer's activity but would have to compensate the injurer for the injurer's "damages" (for example, forgone profits). . . .[3] Thus, there are four possible solutions, corresponding to who is given the entitlement and how it is protected.

In this chapter, we will examine whether the efficiency criterion can determine which entitlement to choose and which remedy to use to protect it. The analysis will be based on an example of a polluting factory located next to a single resident. The facts of the example are described in Table 1. The factory can produce zero, one, two, or three units of output. Increasing output results in additional profits for the factory and additional damages to the resident. If the factory produces one unit the factory obtains $10,000 in profits and the resident suffers $1,000 in damages. Total profits less total damages are $9,000. If the factory produces a second unit, the factory's additional profits are $4,000—so the factory's total profits are $14,000—and the resident's additional damages are $15,000—so the resident's total damages are $16,000. Then total

[1] [Ed: An injunction allows the victim to prohibit the injurer from causing harm, because injunctions are backed up by the court's contempt power. If the injurer violates the injunction, the judge may imprison the injurer or impose other sanctions (such as fines whose amounts increase daily) to coerce compliance.]

[2] [Protecting the injurer's entitlement by an injunction means that the victim has no right to force the injurer to stop. If the victim tries to do so (e.g. by breaking into the factory and shutting it down), the injurer can get an injunction to stop the break in (or more realistically, to stop future break-ins). Or, more simply, the injurer can call the police to protect his property by arresting the victim, or, in appropriate circumstances, can use self-help to defend his property.]

[3] [This combination of entitlements and remedies might be implemented as follows. The victim would have a right to get an injunction against the injurer, but would have to compensate the injurer for the costs (including lost profits) of compliance. This combination of entitlements and remedies is seldom, if ever, used.]

profits less total damages are $2,000. If the factory produces a third unit, the results are described similarly.

TABLE 1

Nuisance Law Example

Output of Factory	*Additional Profits of Factory*	*Total Profits of Factory*	*Additional Damages of Resident*	*Total Damages of Resident*	*Total Profits Less Total Damages*
0	–	$0	–	$0	$0
1	$10,000	$10,000	$1,000	$1,000	$9,000
2	$4,000	$14,000	$15,000	$16,000	-$2,000
3	$2,000	$16,000	$20,000	$36,000	- $20,000

Maximizing the size of the pie in the nuisance law example is equivalent to maximizing the factory's profits net of the resident's damages. Given the data in Table I, this occurs when the factory produces one unit of output (see the last column). Thus, one unit of output is the efficient solution.

Zero Transaction Costs

We know from the discussion in the previous chapter that if there are zero transaction costs, then the factory will end up producing the efficient output regardless of the choice of remedy or entitlement. It will be useful to see how this comes about in the example before considering more realistic assumptions about transaction costs.

Under the injunctive remedy, suppose, for example, that an entitlement to clean air is given to the resident. This corresponds to giving the resident the right to force the factory to produce zero output. The factory, however, would gain $10,000 in profits from producing one unit, and the resident would suffer only $1,000 in damages. Thus, it is in each party's interest to reach an agreement in which the factory would pay the resident some amount between $1,000 and $10,000 for permission to produce one unit. Assuming zero transaction costs, which is interpreted to imply cooperative behavior, such an agreement will be reached. It will not be mutually beneficial for the factory to produce a second unit because the factory would gain only an additional $4,000, whereas the resident would suffer an additional $15,000 in damages. Similarly, it would not be mutually beneficial for the factory to produce three units. Thus, the parties would remain at one unit, the efficient solution.

Under the damage remedy, suppose that an entitlement to clean air is given to the resident, as before, and that the court makes the factory liable for the resident's actual damages. Because the factory would gain $10,000 from producing one unit and would be liable for only $1,000, the factory clearly will choose to produce at least one unit. It will not be in the factory's interest to produce the second unit because the increase in the factory's profits is $4,000 and its additional

liability is $15,000. Similarly, the factory would be worse off if it produced three units. Thus, the factory would choose to produce the efficient output.

Strategic Behavior

The assumption of zero transaction costs obviously is unrealistic in many respects. We will first consider the possibility that the parties will behave strategically; for instance, to establish reputations as tough bargainers, they may "hold out" for a disproportionate share of the gains from any agreement. If both parties are stubborn, they may not be able to reach an agreement even when both could be made better off. Such behavior is not uncommon; for example, parties often go to court rather than settle out of court more cheaply. We will now reconsider the nuisance remedies assuming that the parties behave strategically.

Under the injunctive remedy with an entitlement to clean air, we saw that it would be in each party's interest to reach an agreement in which the factory paid the resident some amount between $1,000 and $10,000. However, because of strategic behavior, the resident may, for example, hold out for $8,000, while the factory may refuse to pay anything over $5,000. As a result, the resident might enforce the injunction and shut down the factory (at least for some period), an inefficient outcome.

The problem of strategic behavior under the injunctive remedy can be overcome by the appropriate choice of the entitlement. Instead of an entitlement corresponding to zero output of the factory—an *absolute* entitlement to clean air—suppose that the court were to choose an entitlement corresponding to one unit of output—an *intermediate* entitlement. Under the injunctive remedy, this would mean that the factory could produce one unit, but no more, without having to obtain the permission of the resident. Starting at one unit of output, it would not be mutually beneficial to produce a second or third unit because the factory's gains are less than the resident's losses. Likewise, it would not be mutually beneficial to reduce output to zero because the resident's gain (in the form of reduced damages) is less than the factory's loss (in the form of reduced profits). Thus, starting at an intermediate entitlement of one unit, the parties will remain there. Strategic behavior cannot upset this outcome because no beneficial changes can be made that would require negotiation.

This discussion illustrates a general principle: Under the injunctive remedy, to overcome strategic behavior it is necessary to choose an entitlement corresponding to the efficient outcome.[4] This is because, starting from any other entitlement, the parties must reach an agreement to get to the efficient outcome; and strategic behavior may prevent this agreement from being reached.

Under the damage remedy with an entitlement to clean air and liability equal to actual damages, we saw that the factory would choose to produce the efficient output of one unit. The presence of strategic behavior does not affect this result because there are no bargains that the parties have to reach. Also, there are no threats the factory can make because, given that the factory is liable for actual damages, the resident is indifferent among all levels of the factory's output.

[4] [Of course, even if the entitlement corresponding to the efficient outcome is not chosen, the parties might still negotiate to achieve the efficient outcome. The possibility of strategic behavior means they might sometimes fail to bargain to the efficient outcome, not that they will always fail to do so.]

It is essential to the conclusion of the previous paragraph that liability equals actual damages. To see why, suppose that the factory's liability is $7,000 for the first unit—exceeding the resident's damages of $1,000—and, as before, is equal to the resident's damages for the second and third units. If the factory produces one unit of output it will gain $3,000 ($10,000 - $7,000). The resident also will gain $6,000, the amount by which the liability payment exceeds actual damages ($7,000 - $1,000). But the factory can deny this gain to the resident by not producing the first unit of output. Therefore, if the factory believes that it can bargain more effectively than the resident, it may threaten to not produce unless the resident pays some specified amount up to his full gain of $6,000. However, if the resident believes that he is the better bargainer, he may not give in to the factory's demand. As a result, the factory may carry out its threat, if only to make future threats credible, and produce at an inefficient output.

The kind of "extortion" threat just described cannot occur if liability is equal to actual damages. Because the resident then is not overcompensated, and so gains nothing from an increase in the factory's output, the factory's threat to not produce the first unit of output has no effect. With full compensation, the resident is indifferent to whether the first unit is produced and is likewise indifferent with respect to the second and third units. Thus, the factory will maximize its after-liability profits by increasing production to the efficient output of one unit.

This discussion illustrates another general principle: Under the damage remedy, to overcome strategic behavior it is necessary to set liability equal to actual damages. If liability exceeds actual damages, then the party who is liable has an incentive to threaten to deny the other party's overcompensation by choosing an inefficient outcome.

Imperfect Information

The analysis of strategic behavior under the injunctive and damage remedies suggests another way in which the assumption of zero transaction costs is likely to be unrealistic. For both remedies, it was seen that the court must have certain information about the nuisance dispute in order to achieve the efficient solution. Under the injunctive remedy, the court needs to know the efficient outcome to choose an entitlement corresponding to it. And under the damage remedy, the court needs to know the resident's damages to set liability equal to actual damages. It was implicitly assumed that the court had whatever information was required. We will now reconsider the remedies when this information is incomplete. The assumption of strategic behavior will be maintained in this discussion.

Suppose that the court has limited information of the following sort: It knows the resident's schedule of damages but does not know the factory's schedule of profits. For example, the court might easily be able to obtain information about the damage to property from pollution but not about the cost to the polluter of changing production methods to abate pollution.

Under the injunctive remedy, the court no longer can achieve the efficient outcome. To reach that outcome, strategic behavior must be avoided, which requires under the injunctive remedy that the entitlement coincide with the efficient output. But to determine the efficient output, the court must know when the factory's profits net of the resident's damages are maximized. Knowing the damage schedule alone obviously is insufficient to determine this level of output.

Although the court could attempt to estimate the efficient output, if it makes a mistake, as it generally will, strategic behavior may prevent the parties from bargaining to the efficient output.

The damage remedy can reach the efficient outcome despite the court's imperfect information. This result can be guaranteed, however, only if the court gives an absolute entitlement to clean air to the resident and sets liability equal to actual damages. Any other entitlement might lead to the efficient outcome, but it need not. For example, suppose that the court assigns an intermediate entitlement to pollute to the factory corresponding to two units of output and makes the factory liable thereafter for the resident's damages. Initially, the factory would choose to produce two units because there is no liability up to and including the second unit, and liability for the third unit—equal to the resident's damages of $20,000—exceeds the factory's additional profits—equal to $2,000. The resident then would have an incentive to "bribe" the factory to reduce output from two units to the efficient output of one unit, but because of strategic behavior the parties might not reach that output.

On the other hand, if the court chooses an entitlement corresponding to zero or one unit of output and sets liability thereafter equal to actual damages, the damage remedy will lead the factory to produce at the efficient output. In other words, with liability equal to actual damages, the damage remedy leads to the efficient outcome if, and only if, the entitlement is at or below the efficient output. However, because the court cannot determine the efficient output with its limited information, the only way it can guarantee the efficient result is to choose the entitlement corresponding to the lowest possible output—an absolute entitlement to the resident.

The discussion thus far has shown that if the court knows the victim's schedule of damages but not the injurer's schedule of profits, it generally cannot implement an efficient injunctive remedy but it can implement an efficient damage remedy. However, in many nuisance situations the court might not be able to easily determine the victim's damages. For example, although a court might be able to accurately estimate the market price of someone's home, this price generally is less than the damages that would be suffered by the resident if he were forced to move, because it does not reflect the special attachment he might have for that location and house. Damages often include a subjective or idiosyncratic element of this sort that is difficult or impossible to measure. We will therefore briefly reconsider the remedies when the court is assumed to underestimate the resident's damages (and, as before, to not know anything about the factory's profits). For concreteness, suppose in the example that the court's estimate of damages is $500 per unit of output.

Because the court obviously still cannot implement an efficient injunctive remedy, the discussion will focus on the damage remedy. Suppose an absolute entitlement is awarded to the resident. If damages were measured accurately, then, as we saw above, the damage remedy would lead to the efficient outcome. Now, however, with damages underestimated, the factory generally will overshoot the efficient output. In the example, with liability equal to $500 per unit of output, the factory will choose to produce three units because its additional profit from producing each unit exceeds $500 (see Table 1).

Starting at an output of three units, the resident would be better off by $19,500 if output were reduced by one unit—he would lose a $500 liability payment, but his damages would decline

by $20,000 (see Table 1). The factory would lose only $1,500 by this change—its profits would fall by $2,000, but it would avoid a $500 liability payment. Thus, if the parties could reach an agreement in which the resident paid the factory some amount between $1,500 and $19,500 to reduce output by one unit, both parties would be better off. However, because of strategic behavior, such a deal will not necessarily occur. And even if an agreement were reached with respect to this unit, the parties might fail to reach an agreement when they negotiate over reducing output from two units to the efficient output of one unit.

The general point of this discussion can be simply stated. If courts underestimate the victim's damages, then the damage remedy initially will lead to an excessive output and this inefficiency may not be corrected because of strategic behavior. There is then no general reason to believe that a damage remedy would be preferable to an injunctive remedy. For example, suppose that, starting with an absolute entitlement to the resident, the damage remedy would lead to an output of three units for the reasons described in the previous paragraph, and the injunctive remedy would lead to an output of zero units because of strategic behavior. The injunctive outcome then is more efficient than the damage outcome because total profits less total damages are $0 rather than - $20,000 (see Table 1). In general, however, either remedy could be the more efficient one.

We can now summarize the results in this chapter regarding the efficiency analysis of nuisance remedies. If the parties can be expected to bargain cooperatively (and there are no other transaction costs), then every choice of entitlement and remedy will be efficient. If the parties are likely to act strategically, then the efficient outcome still can be achieved under both remedies if the court has adequate information. Strategic behavior can be overcome under the injunctive remedy by choosing the entitlement that corresponds to the efficient outcome, which can be determined only if the court knows the injurer's benefits from engaging in the harmful activity and the victim's damages. And strategic behavior can be overcome under the damage remedy by giving an absolute entitlement to the victim and setting liability equal to actual damages, which obviously requires knowledge of the victim's damages. If the court only knows the victim's damages, the injunctive remedy generally will fail because the court cannot accurately set the entitlement equal to the efficient outcome, but the damage remedy still can guarantee the efficient outcome. However, if the court underestimates the victim's damages, then the damage remedy generally will lead to excessive output and may be less desirable than the injunctive remedy.

Although in theory either remedy could therefore be more efficient, it may be apparent in particular circumstances that one remedy is likely to be better than the other. For example, suppose that a court is confident that its estimate of the victim's damages is close to the truth, but believes that its estimate of the injurer's benefits is inaccurate. Then an entitlement to the victim protected by a damage remedy generally would be preferred because this would be likely to lead to an outcome close to the efficient solution. Alternatively, suppose that a court has poor information both about the victim's damages and the injurer's benefits, but is confident that the efficiency loss from too little activity by the injurer is small relative to the efficiency loss from excessive activity. Then an entitlement to the victim protected by an injunctive remedy would be desirable because this would guarantee that the final outcome will not be too bad. Thus, the efficiency analysis of

nuisance law may be helpful even when there is some uncertainty about which entitlement and remedy to choose.