30

TREATING EXTERNAL DISECONOMIES — MARKETS OR TAXES?

by DAVID STARRETT and RICHARD ZECKHAUSER

from John W. Pratt (ed.), Statistical and Mathematical Aspects of Pollution Problems. Dekker, N.Y., 1974

REPRINT NUMBER 59R

REPRINT SERIES

Public Policy Program — Kennedy School of Government — Harvard University

5

TREATING EXTERNAL DISECONOMIES - MARKETS OR TAXES?

David Starrett*
and
Richard Zeckhauser

Harvard University
Cambridge, Massachusetts

Sometimes the parties to an economic transaction do not reap all of the benefits and costs involved: there are externalities. Handling an externalities problem is like fighting an undiagnosed illness. At some times it will yield to any of a number of familiar treatments. At others, it will be resistant to any that are known. In this paper, we first consider those pathologies that can be adequately treated by standard remedies. Then we turn to examine some more resistant cases, cases that may require new diagnoses and unconventional methods of treatment.

1. TRACTABLE EXTERNALITY PROBLEMS

Thrust a price theorist into a world with externalities and he will pray for second best - many firms producing and many firms and/or consumers consuming each externality, with full convexity everywhere. No problem for the price theorist. He will just establish a set of artificial markets for externalities, commodities for which property rights were not previously defined.[†] Decision units, being small relative to the market, will take price as given. The resulting allocation will be a competitive

---

*Present address: Stanford University, Stanford, California.

[†]Due to their elusive physical natures, most externalities have not developed property rights. Markets for externality rights are artificial in the sense that natural forces have been insufficient to foster their existence.

outcome of the classical type. If artificial markets do not appeal, an equally efficient taxing procedure is available.

Where the externalities are beneficial, the price theorist's life is particularly easy; the more intractable aspects of externality problems do not appear. We concentrate here on external diseconomies, the more interesting case, and the case of greater empirical import.

For either artificial-market or taxing procedures it is essential to determine who has initial rights. Our common-law upbringings would lead us to expect that rights should go to the sufferers. It turns out, however, that it is just as efficient to give the rights instead to those who dump the sewage and soil the air.

Upton Paper Mills and Downley Baths are our representative producer and recipient of an externality. Upton dumps organic residue into the very river that provides the water for Downley's swimmers. Downley of course screens and purifies, but not without cost.

1.1. Artificial Markets*

Consider first the working of an artificial-markets scheme. If Downley has rights, he sells Upton the right to pollute. Upton, being but one of many polluters, can have no effect on the total pollution level, and hence no effect on its price. He takes price as given and continues polluting up to the point where his marginal gain from polluting (what is in effect his marginal cleanup cost) just equals the going price. Downley will sell rights to pollute up to the point where his marginal loss in profits from the pollutant equals the going price. The equilibrium level of production for the externality, the pollutant, is established where the marginal loss to the recipient equals the marginal cost to the producer of not providing the externality, these equal marginal amounts being the equilibrium price of the pollutant. This price times the amount of pollutant Upton dumps is the competitive charge Upton must pay Downley.

Give Upton the right to dump, and Downley must pay him a competitive charge equal to marginal cleanup cost times the amount of cleanup he provides. At the equilibrium, the equation of the marginals is the same as before. This second scheme presents a difficulty, however; it requires a benchmark for original filth in order to determine just how much cleanup has been provided. One possible benchmark is the amount Upton would

*Allocation schemes of this type were first proposed by Lindahl [6]. They have since been discussed by Coase [4] and Arrow [2], among others.

dump if there were no Downleys, if there were no affected parties. For now we merely assume that the benchmark, N, can be established.*

A key question is whether shifting property rights here can drive either Upton or Downley out of business. If not, the marginal efficiency conditions will define a unique equilibrium, and the level of pollution will be the same whether Upton or Downley has rights.†

But the optimal allocations are quite different if one or the other of the firms ceases to function. If we give Upton rights, might Downley not decide to go out of business to avoid paying the charge for Upton's cleanup? Alternatively, might not a potential polluter who previously could not turn a profit because of high effluent charges now find it advantageous to go into business? When you consider that the size of N is arbitrary, it is evident that the answer to these important questions is yes. What is surprising, however, is that neither answer can be affirmative so long as the traditional convexity assumptions are satisfied.** For those interested we provide a mathematical demonstration of this assertion; others may wish to skip the next four paragraphs.

The assessment or payment of a competitive charge for an externality (as determined on an artificial market) to a firm with a convex production

---

*There is a third possibility that also leads to the same result. The government can have rights. It will sell these rights, say N of them, on the open market to both producers and recipients. Any right purchased by a recipient is not available to a producer, and thus reduces externality production by one unit.

†There is a more general way to look at this process of transferring rights and shuttling rents. An equivalent procedure is one where a standard level of externality production is set. Everything above that level requires a payment from producer to recipient and conversely. The case we call producers'-rights sets that standard level at N. The recipients'-rights case sets it at 0.

The argument given here that a change in property rights will not alter resource allocations only holds if final demand prices do not change. They might change if individuals with divergent tastes own the firms involved in the transfers, and if these firms are large relative to the rest of the economy. See Dolbear [5] for a discussion of the distributional effects of property rights transfers.

**This issue has a long but inconclusive literature. See, for example, Buchanan [3].

set can never determine whether it operates at a profit or at a loss. To see this, consider a simple case with a recipient firm (Downley in our model) that produces one output, b, using a single input, a, and that suffers from a single pollutant, z. The firm's production function may be written

$$b = f(a, -z),$$

with $f_1 > 0$, and $f_2 > 0$, where the subscripts refer to partial derivatives with respect to the first and second arguments. This function is assumed to satisfy the traditional convexity assumptions, i.e., it is concave throughout. We can measure the level of pollutant from the benchmark N so that the arguments of the production function are nonnegative. The translated production function,

$$b = g(a, N - z),$$

is clearly concave if f is.

The price of the output is normalized to unity. If p_a and p_z are the prices to the firm of the input and pollutant, respectively, then (regardless of whether it is an Upton- or a Downley-rights scheme) the marginal conditions

$$p_a = f_1(a, -z) = g_1(a, N - z), \text{ and}$$

$$p_z = f_2(a, -z) = g_2(a, N-z)$$

must be satisfied. In what follows, we suppress the arguments of f_1, g_1, f_2, and g_2.

The properties of concave functions imply that

$$g(a, N - z) - g(0, 0) \geq g_1 a + g_2 (N - z) = p_a a + p_z (N - z).$$

If recipients have rights, Downley's profits, π, will be

$$\pi = g(a, N - z) - p_a a \geq p_z (N - z) + g(0, 0).$$

To achieve the outcome of the producers'-rights case, we need merely transfer $p_z (N - z)$ from recipient to producer. After the transfer the recipient firm, Downley, makes $g(a, N - z) - p_a a - p_z(N - z)$, a still greater profit than he would make if he went out of business and got $g(0, 0)$. Indeed, Downley has no incentive to change his behavior in any way.

The amount $p_z (N - z)$ of the transfer that shuttles back and forth depending on who has rights is a pure rent that returns to the fortunate party. As such it cannot reverse either a profitable or unprofitable status for the externality recipient. A completely parallel argument establishes this same result for a firm, such as Upton, which is an externality producer.

Note that as N becomes larger, this result becomes increasingly counterintuitive. Surely the amount of rent that can be passed back and forth without forcing a shutdown decision of either the externality recipient or producer is finite. It would hardly seem possible that it could depend on an arbitrarily chosen benchmark N. The apparent paradox is explained in the second half of the paper, where we show that for sufficiently large N, full convexity is logically impossible; without convexity the preceding analysis is not relevant.

1.2. Taxes*

It is sometimes proposed that a taxing scheme be employed to overcome externality difficulties. With taxing schemes, the party with rights quotes its marginal loss to the government, which then makes the appropriate collection, the level of the externality times the marginal loss inflicted.† Given the conditions of (1) many-participant markets (assuring "honest" preference revelation), and (2) full convexity, the only difference between the artificial-markets scheme just described and a taxing scheme is that the latter does not provide positive payment for the party that has rights; the government makes a profit. A simple transfer from the government to the party with rights is all that is necessary to turn the taxing-scheme outcome to the one that would have resulted had an artificial market been established. The allocation of resources, following the argument of the previous subsection, is clearly the same. Furthermore, shutdown or operate decisions are not affected. Whether a firm is or is not paid the appropriate tax (the value on a competitive market of the loss it suffers) cannot reverse the profitability or unprofitability of its operation.

The most familiar proposal in the pollution world for this sort of tax scheme involves the imposition of effluent charges, the charge being a specific example of a unit tax on externality producers. It is sometimes proposed that instead of setting a per unit charge on effluent, a certain number of pollution rights should be auctioned among the potential polluters. Except for the mechanical way the equilibrium is reached, the two schemes are identical. The results they produce will be indistinguishable if the

*The first general discussion of tax schemes in this context was given by Pigou [8]. Further discussion can be found in Meade [7] and Coase [4].

†For a discussion of the mechanics of such a scheme, see Aoki [1].

quantity of rights auctioned under the latter scheme is just equal to the amount of pollutant dumped given the particular effluent charge established for the former scheme. They are both, in effect, unit-tax schemes.* We conclude, then, that the four schemes, artificial market or tax, recipient or producer having rights, all produce identical allocations of resources, assuming many participants and full convexity.†

In the discussion that follows, we detail some frequently occurring situations in which the assumption system that supported the foregoing analysis is not applicable. The equivalences we just developed will no longer hold, and, for the most part, the schemes we have presented will not lead to efficient outcomes.

2. RESISTANT EXTERNALITY PROBLEMS

Sometimes, alas, we find ourselves in third- and fourth-best worlds, where the externalities problems are not of the convenient type treated thus far, and new, more resistant externality problems arise. First we discuss briefly the problem of noncompetitive behavior, a problem which has been noted elsewhere by Arrow [2], due to thinness of markets and the public bad nature of the externality. Then we analyze difficulties involving nonconvexities. Finally we discuss parallel aspects of these difficulties when consumers are externality recipients.

2.1. Thinness of Markets

We have treated Upton and Downley, our representatives of the many producers and recipients, as buyers and sellers of the same externality. For this treatment to be meaningful on the selling side, it must be possible to have specific amounts of the externality conveyed privately to individual buyers. This would be the case, for example, with solid waste which could be disposed at any of a number of dumping grounds.

---

*There are (unfortunately) many frequently proposed schemes that do not achieve equivalent or efficient results even under our ideal conditions. For example, the oft-heard proposal that total costs of damage should be divided in proportion among the externality producers responsible for it, a scheme which is essentially average-damage taxation, is inefficient for the same reason that average-cost pricing is inefficient.

†We have demonstrated these results only for a particularly simple case. Starrett [11] has shown elsewhere that these results hold in a very general, abstract framework.

More commonly, unfortunately, any quantity of an externality received by one recipient is received by all. Thus, in the Upton-Downley example, Upton's effluent serves as a public bad for all downstream firms. The functioning of a market requires that the sellers compete with each other on the supply side. When the externality is a public bad, what is supplied by one (dumping rights in this instance) must be supplied by all. A traditional competitive market cannot exist. *

To achieve the equivalent of a competitive outcome, we would have to set up a separate competitive market for pollution rights between Upton and each one of the downstream firms. If each seller-recipient would take price as given in his particular market, and if Upton would act competitively and take as his given price the sum of these individual market prices, Upton would wind up paying a price per unit of externality equal to the sum of marginal downstream losses. To have such a price at the margin is the appropriate efficiency condition for the production of this public good (reduction in the level of externality). Given the public nature of the externality, the amount contracted on each market must be the same (the prices will differ in general). The simplest way to reach the efficient point may be to vary quantity on every market simultaneously and to ask each downstream firm to announce its going price, its marginal losses. Obviously, any scheme of this sort would require a substantial amount of centralized administration.

This scheme, as well as the tax proposals, will run into further difficulties given that there are not many producers and recipients dealing in the same commodity. Participants will find it in their interests to behave strategically. If recipients have rights, they will each be empowered to block any producer's production of the externality. A producer will be permitted to produce the minimum amount of the externality for which he buys rights from all sellers. Individual buyer-seller negotiations will probably result. Each seller will be responsible for but a small portion of the total price a producer (or all producers should they band together) must pay for increasing externality production. He will thus have little effect on the total amount the producer demands, and will have an incentive to sell at an extortionary price, a price well above his marginal losses. The likely outcome is that much less than an optimal

*There can still be competition on the demand side, all producers competing with each other to dump, the total amount dumped being the sum of what is dumped by each.

amount of externality rights will be sold, less than an optimal amount of externality production will take place. *

If producers have rights, the strategic shoe is on the other foot. A reduction in the amount of externality produced will benefit all recipients. The usual public goods, free-rider problem arises. Unless the recipients are capable of coordinating their activities, they will purchase less than an optimal amount of the public good (reduction of the externality).

Participants also have incentives to act noncompetitively in the tax schemes. If there are many downstream firms, they have an incentive to overstate their marginal losses to the government, since this may lead to a lower pollution level and there is nothing to lose. (Recall that recipients are not compensated in the tax scheme.)

2.2. Nonconvexities

Let us return to the river and assume that thinness-of-markets difficulties do not arise, or that the government can enforce the equilibrium through an administered planning process. Upton is still our representative producer. If property rights can be established for externalities, Downley is our representative recipient. Otherwise he is a single recipient whose losses from any level of externality production are the sum of the losses of all recipients. Both Upton and Downley act as price takers.

Our problems, alas, are not over. If Upton and his fellow polluters dump enough, it is absurd to expect that Downley will remain in the business of providing swimming facilities. Perhaps, if he had rights, he might stand around idle in order to collect his due from Upton. But from a resource allocation standpoint he would be out of business. This possibility was ruled out in Section 1 of this paper, where full convexity was assumed. Downley's quite appropriate shutdown decision must be associated with a nonconvexity in his production set.

To see that there is indeed such a nonconvexity, consider a single-product competitive firm that suffers from an external diseconomy. In the absence of externality markets or taxes, it could make a short-run,

---

* If the recipients were acting in concert, they would act as a monopolist and sell an amount which equates marginal revenue (derived in the usual manner) and marginal cost (the sum of their marginal cost curves). In the strategic situation described, unless the sellers can achieve some degree of coordination they will sell even less than they would if they were acting in this monopolistic fashion.

profit-maximizing decision for each potential level of externality. Plot its level of maximum profits as a function of the level of externality it receives. We know a lot about the shape of this function. It must be downward sloping, since z is a diseconomy. If, as assumed, the firm's production function is concave, this curve must be concave as well. This implies that it must reach any arbitrarily great negative value, since z is unbounded.* But the firm could always choose to produce nothing and thus incur at most fixed costs. This implies that the maximum-profits vs externality curve must be bounded below. The maximum loss floor that is established by the shutdown possibility rules out global concavity of the production function.

Figure 1 shows two possible shapes for the profits function when there is a nonconvexity in the production set.

A nonconvexity will obviously have no effect on any outcomes so long as N, the level at which the externality would be produced in the absence of any compensation arrangements, is less than K, the point at which the nonconvexity begins.

In the analysis that follows, we are worried about cases where nonconvexities matter, where N > K. Such a location reflects a situation

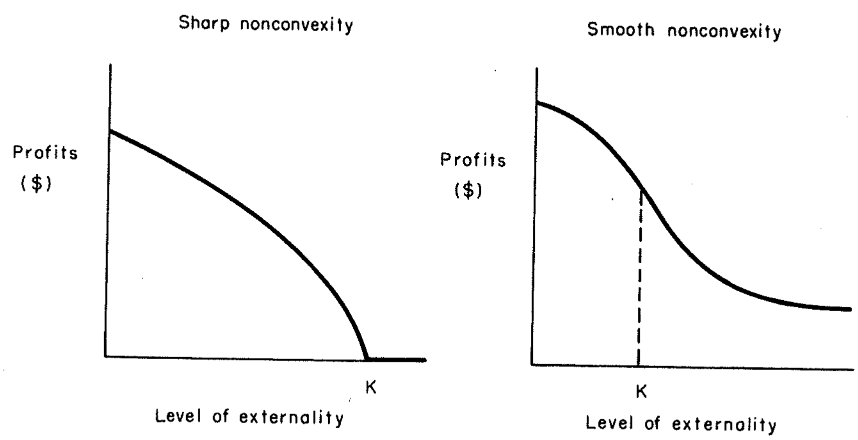

Fig. 1. Externality recipient's profits.

* A concave function always lies below its tangent line. If the tangent is ever downward sloping, it must cross the z axis, and so, therefore, would the profit function if it were everywhere concave.

where the potential harm is serious enough that it may induce a firm to shut down operations.

Drawing on historical precedent, we turn from swimming to agriculture and consider as our example the classic farmer-cowman illustration developed by Coase [4]. Coase observes that wandering cattle trample the neighbor farmer's crops. Consider two possible schedules of damage, representing, respectively, a sharp and a smooth nonconvexity (Table 1).

For full convexity the cows must do increasing marginal damage; something like Schedule A must apply. The difficulty is that the addition of a fifth cow to the herd induces the farmer to discontinue operations. Beyond four cows profits are identically zero,* and the profit function is not convex.

TABLE 1

Two Schedules of Damage to Farmer

	Schedule A: sharp nonconvexity		Schedule B: smooth nonconvexity	
Size of herd	Total profits	Marginal loss	Total profits	Marginal loss
0	8		8	
1	$7\frac{1}{2}$	$\frac{1}{2}$	6	2
2	6	$1\frac{1}{2}$	5	1
3	$3\frac{1}{2}$	$2\frac{1}{2}$	$4\frac{1}{2}$	$\frac{1}{2}$
4	0	$3\frac{1}{2}$	$4\frac{1}{4}$	$\frac{1}{4}$
5	$-4\frac{1}{2}$	$4\frac{1}{2}$	$4\frac{1}{8}$	$\frac{1}{8}$
	0[a]	0[a]		

[a] This figure applies if the farmer stops farming.

*Profit here should be thought of as return above and beyond variable cost, since in the short run the firm will quit only when it cannot cover variable cost. Returns to the fixed factor (land) are ignored in making short-run decisions.

Schedule B presents what we call a smooth nonconvexity. It illustrates a situation in which the marginal damage per cow is decreasing. Casual consideration of technological factors might lead us to suspect that this would be the case, that two cows would trample less apiece than one, and that a dozen would not multiply by six the damage done by but a pair.

Decreasing marginal damage over some range is not the only cause of the types of nonconvexities we have been considering. Nonconvexities will crop up if there are economies of scale in treating external diseconomies. For twice the price of a fence that keeps out one cow, the farmer could probably construct one that could keep out eight. Even if there were increasing marginal trampling damage over the one-to-eight range, once there are enough cows to make a fence worthwhile, a nonconvexity presents itself; extra cows will impose decreasing marginal costs (losses of profit) on the farmer. If, as seems not unlikely, the fence can keep out ten or one hundred, the nonconvexity will be of the sharp variety.

The results we are about to discuss hold equally whether nonconvexities are smooth or sharp. However, the exposition is somewhat easier when nonconvexities are of the sharp type; therefore that is the case we shall discuss.*

2.3. Artificial Markets

Consider Schedule A and the operation of an artificial market. Give the farmer, the recipient, rights. Say that the price, p^*, of the externality (cows that trample) is $2. The farmer, adopting traditional marginal optimization procedures, decides to accept two cows, the number for which his marginal loss does not exceed the price. Assume that the cowman's profit schedule leads him to demand rights to graze two cows at this price. Do we have an equilibrium? No! Surely, with the present price, this position is not optimal for the farmer. He would make greater total profits (infinite) if he accepted an infinite number of trampling cows at this price, or indeed at any positive price.

Thus, unless there is some artificial restriction on the number of rights which can be sold, equilibrium cannot exist for this artificial-market scheme. The farmer who is offered a positive price for accepting cows which can cause him at most a finite amount of damage (his farming profit), will surely want to supply an unlimited number of rights; but if he is offered a zero price, he supplies no rights, so equilibrium is impossible.

*The two types of nonconvexities may be combined in a single instance if the total profits curve of a smooth nonconvexity cuts the horizontal axis, the axis on which the externality is measured. But either nonconvexity, by itself, is sufficient to hinder the effectiveness of the corrective measures outlined above.

Nonexistence is somewhat surprising, since the presence of arbitrary nonconvexities does not ordinarily rule out the possibility of equilibrium.

There may still exist equilibria if we modify the rules by introducing an upper limit on the number of rights which can be sold. Indeed, an equilibrium will surely exist if we set the upper limit at the level K in Fig. 1, since then the problem is convex and we can appeal to a standard existence theorem. However, such an equilibrium no longer has any efficiency properties; for example, if it happens that the best solution is to have the farmer out of business and pollution levels above K, any equilibrium which excludes this possibility is surely inefficient. To be sure that the upper limit does not exclude the best outcome, we would have to set that limit outside the "feasibility" region. In our example, we would set the limit at N, the maximum number of cows our rancher would put on the land if the farmer were absent.

Having taken this precaution, special circumstances are required in order that an equilibrium exist. Let us return to our example. When the externality market is introduced, the farmer receives an externality payment, p^*c, where c gives the number of cows grazed. His aggregate profits, π^a, are equal to his profits from farming, π, plus the externality payment. The line through the point (2, 6) in Fig. 2 is an iso-aggregate-profits line; it offers a value of 10. The point (2, 6) is a local optimum for the farmer, but he could increase his profits by stopping all production and selling N externality rights. In the case shown, where N = 8, the farmer could secure aggregate profits of 16 at the present externality price by ceasing farming operations. This means that (2, 6) cannot be a global equilibrium; there is no price that would lead the farmer to choose to operate at that point.

No other price can be an equilibrium either. At prices above $\hat{p}$, the farmer will always want to sell N externality rights, while the cowman will not want to have as many as N cows.* At prices below $\hat{p}$, the farmer will want fewer than two cows (the number he would choose when the price is p^*), while the cowman would want at least the two he had when the price was higher. In terms of supply-and-demand analysis, the farmer's supply curve of cattle externalities is discontinuous and the cowman's demand

---

*In the numerical example, the externality level can only take on integer values; $\hat{p}$ turns out to be 1 1/14, the price that makes the farmer's profits equal whether he accepts eight cows or one, the latter being his optimal inferior point.

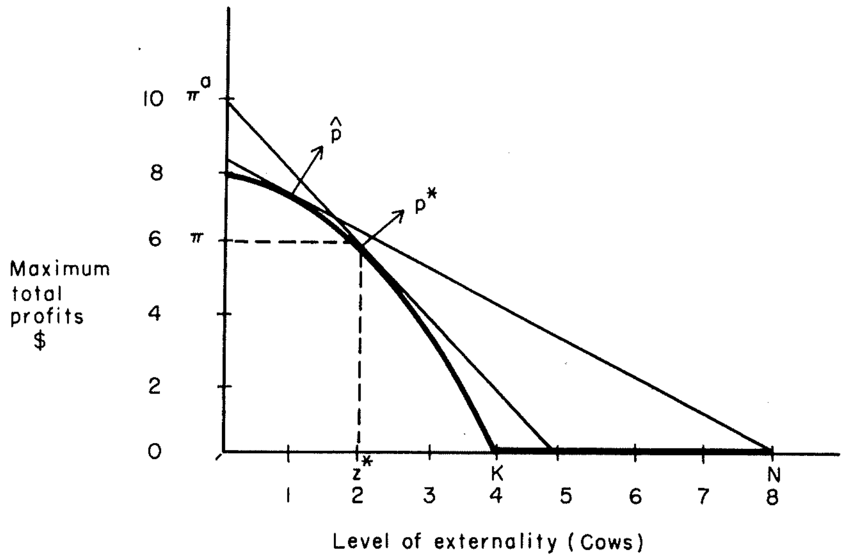

Fig. 2. Externality recipient's profits - sharp nonconvexity. The p's in the diagram represent the gradients of the total profits function for their respective price vectors.

curve passes through the gap of the discontinuity. Equilibrium will exist only if N is below 5 (so that at price p^* our farmer cannot offer enough rights to improve on the position (2, 6)).

The instability in any proposed equilibrium which derives from the farmer's threat to quit may seem somewhat silly, but, as we shall see below, it is a signal that an efficient allocation of resources may require that the farmer get out of farming.* We conclude that in cases where externalities are potentially strong enough to drive out firms, situations where N is large, the artificial-markets scheme will not work at all.†

---

*Note that instability is even more likely in the long run, since then the farmer will quit whenever he cannot cover his fixed as well as his variable cost.

†Our conclusion that equilibria will generally fail to exist seems to be related to some results of Shapley and Shubik [9]. They showed that in games with external diseconomies, the core frequently fails to exist. However, Starrett [10] has shown elsewhere that the two results are unrelated. The Shapley-Shubik result depends on a particular definition of property right, whereas ours does not.

2.4. Taxes

Can we expect a tax scheme to lead to the optimum equilibrium? Returning to our diagram, it appears that not one, but two distinct points are candidates for such an equilibrium. The point (2, 6) is a possibility. If the government tells the farmer that there are going to be two cows, the farmer's quoted marginal losses will be \$2, and at a tax (price) of \$2, the cowman will want to have exactly two cows. But the point (N, 0) is also a contender. At that point, with the farmer out of business, marginal losses (and hence taxes) are zero and the cowman will want to have N cows, his profit-maximizing number of cows assuming zero externality payments.

The presence of multiple equilibria should not surprise us. Pigou [8] points out that nonconvexities will lead to multiple tax equilibria. The point that we want to make here is that since external diseconomies logically imply nonconvexities, multiple equilibria are going to be the rule rather than the exception.

Should the farmer operate or not? To find out, we need to know whether society is a net gainer or loser from having the farmer operate at this location. Assuming that prices of private goods correctly reflect social opportunity costs, the best point is the point of maximum combined profit.* In Fig. 3, we plot the marginal profit loss of the farmer, ss, and the marginal profit gain, dd, for the cowman (derived from figures not given here), as functions of the number of cows.

Note that there are two points, A and N, where the marginal gain is equal to the marginal loss. These naturally correspond to the two potential tax equilibria. Whether we would prefer A or N depends on whether the total area under the ss curve between A and N (the farmer's profit loss moving from A to N) is greater or less than the total area under the dd curve over that range (the cowman's gain for the corresponding move).

We can conceive of situations in which the balance might be tipped either way. For example, if the soil is particularly rich, so that the farmer could earn considerable differential rent on it, and if it is the case that if the cowman intensified operations he would not secure much of a profits gain, then the farmer should continue farming. If the land is poor for farming, then the situation would probably be reversed, and the farmer should cease operations.

---

*We are assuming that our firms are small enough that their behavior does not affect final demand prices. Furthermore, the statement that everyone is better off at the point of maximum combined profit may require transfers among owners of the farm and the ranch.

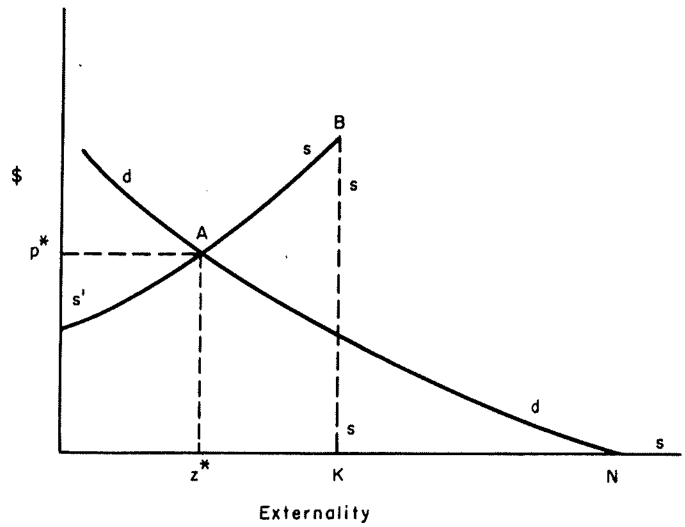

Fig. 3. Marginal profit and loss curves.

Unfortunately, if we set up the adjustment mechanism for taxes described earlier, there can be no assurance that the preferred equilibrium will be found. Assuming that we follow the conventional marginal prescriptions, we will end up at A or N depending on whether we start with externality levels less than or greater than K. Where we start, not the merits of the situation, will determine the outcome. The only sure way to find the optimum when there is the possibility that it will be appropriate for the externality recipient to cease operation, is to engage in the type of macro cost-benefit analysis described above. And in a general equilibrium framework with a highly interdependent externality structure, the required analysis would be incredibly complicated. Searching for shortcuts in such a procedure will provide an important, and hopefully productive, area for future research.

The above analysis suggests at least one important implication for current policy far beyond the narrow example on which it was based. Until recently, externality producers have not been forced to bear many of the costs they impose on others. Thus, current pollution levels are probably beyond optimal levels. What taxes should be charged? The above analysis suggests that using marginal losses may understate the relevant costs and lead to less than desirable adjustments. Firms, or households (see below), which have quit or moved will not be counted at all, even though they would benefit from substantially lower pollution levels. Furthermore, agents who have already absorbed the bulk of their potential damage from pollution (whose loss-of-profit functions show a smooth nonconvexity) may have small marginal losses now. Under these circumstances a nonoptimal equilibrium may be established with relatively low taxes and relatively high pollution levels.

2.5. Consumers as Externality Recipients

What about consumers? They too suffer from external diseconomies. In fact, many would argue that the most important externalities are in consumption. There would seem to be two significant ways in which consumers as externality recipients are unlike firms. First, consumers do not present a convenient cardinal measure of their welfare which would correspond to profits for a firm. In a variety of capacities, the consumer's willingness to pay to reduce the level of an externality may play a parallel role to the marginal loss in profits concept when the externality recipient is a firm. But if we are interested, as many economists may be, in the question of whether an artificial-markets scheme can be counted on to lead to an efficient equilibrium, we must look deeper and see whether the consumer's indifference surfaces display nonconvexities.

Second, consumers, unlike producers, cannot shut themselves down. For them, going out of business is not an approved enterprise. But the consumer does have an equivalent to the shutdown option; he can cease consuming a particular good. Such an action would merely indicate a situation where an indifference curve met an axis. If the increased presence of an externality induces a consumer to stop consuming an affected good, we might expect to find a nonconvexity in the consumer's indifference map. We will find one if the consumer does not care about the level of externality once he ceases consumption, if in this sense his after-shutdown feelings about the externality parallel those of a firm.

That a nonconvexity will be present in such an instance can best be illustrated with reference to our polluted-river example. Consider, for instance, a consumer who swims not in the Downley Baths, but directly in the river. Beyond a certain level of filth, the consumer will swim elsewhere or take up another form of recreation. He will cease consuming a particular commodity.

A full geometric representation of the nonconvexity involved would require a complex three-dimensional diagram displaying the three arguments: pollution, swimming, and other goods (money). Fortunately, there is a simpler, though indirect, way in which we can demonstrate the presence of the nonconvexity. Look at a typical consumer's indifference curve on one swimming-pollution plane (Fig. 4). Every point on the curve represents some positive level of swimming and offers satisfaction level I. We assume that by giving the consumer a sufficiently large increment in his consumption of other goods we could just compensate him for a loss of all his swimming; that is, we can find a point (in general, there will be many) in commodity space that has zero swimming and offers satisfaction level I.

We will now argue that, given these assumptions, there is no point (such as Q) on the indifference curve displayed that can be a competitive equilibrium in the sense that there are prices (on real and artificial markets) for which this point is the least-cost way of achieving satisfaction level I. Since the consumer prefers less pollution, its price must be negative.

Now find the point in commodity space mentioned above, one that has zero swimming and offers satisfaction level I. Call it P. To purchase the goods bundle represented by P would cost a finite amount. But with zero swimming, the consumer no longer cares about the pollution level. With a negative price on pollution, the consumer would demand an arbitrarily large amount of it and would attain point P and satisfaction level I at as low a total cost as he wished. Thus, no point such as Q can represent a competitive equilibrium. This means that the consumer's indifference

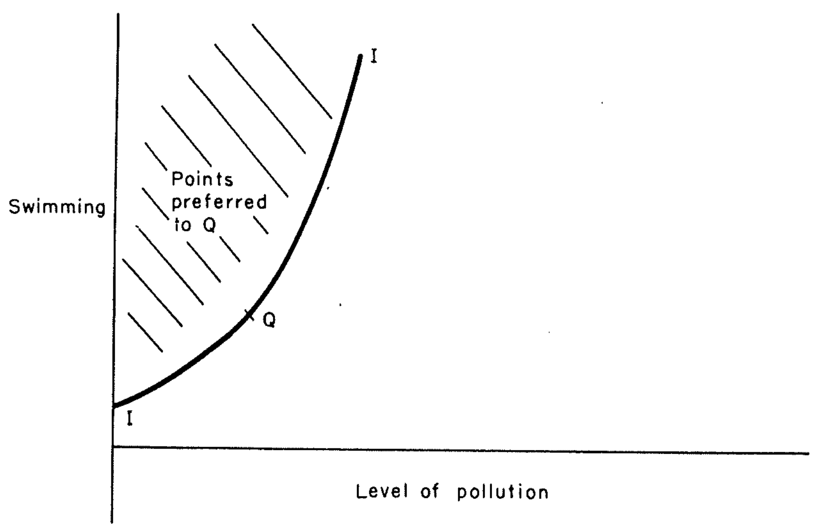

Fig. 4. Swimming-pollution indifference curve.

curves must present a nonconvexity; with full convexity of preferences, every commodity bundle is the least-cost way of achieving its own satisfaction level for an appropriate choice of prices.

With the consumer an externality recipient, there is no more assurance that an artificial-markets scheme will lead to an equilibrium than there is when the recipient is a firm. In effect, the two types of recipients display the same pathologies in these circumstances. It can be shown equivalently that if a tax scheme were employed in a consumer-recipient situation, the same form of macro cost-benefit analysis required with firms as recipients would be needed to find the efficient equilibrium among the multiple contenders.

To summarize, a consumer externality recipient may cease consuming a commodity and thus like a firm may present a nonconvexity that renders irrelevant conventional analyses about the existence and efficiency of competitive equilibria.

Despite these similarities, the consumer nonconvexity is less general than that for the firm. A consumer cannot always escape the effects of an externality. Suppose in our example that the consumer felt badly about river pollution even if he were not swimming in it. Then, although it is still likely that there will be a nonconvexity (since his dislike for pollution probably diminishes when he quits swimming), we can no longer demonstrate it logically. Furthermore, if an externality affects an individual through his consumption of an essential commodity, there will be no shutdown possibility and hence no essential nonconvexity. Air over New York City would not be such an essential commodity; one could always move away. But one cannot escape from radioactive contamination of our atmosphere. For this type of consumer externality there may not be a nonconvexity.

Even if full convexity is assured, the results for consumers differ from those for producers. With producers, the shuttling of rents that result from transfers of rights or switches from tax to market schemes do not affect resource allocations. But with consumers, these rents will exert income effects; these in turn will alter the efficient allocation of resources. Although tax and market schemes will both lead to Pareto-efficient outcomes when consumers are involved with externalities, they will lead to different efficient levels for the production of the externality.*

*The income effect on the purchase of externality production may well be negative. Rich people can swim in their own pools and will not have river pollution affect their natatory pleasures.

3. CONCLUSION

Ever since we left our opening section with its very restrictive assumptions, the problem of external diseconomies has looked more complex. In the real world we should expect to find externalities functioning as public bads for a variety of firms and consumers all of whom might exercise a shutdown option. The entire array of problems considered in the second half of this paper will be encountered.

Our conclusions, therefore, must be rather negative. We have shown that none of the decentralized schemes, even if they can be administered properly, can be expected to work in any straightforward manner when externalities are serious. An equilibrium may not even exist with artificial-markets setups. Even with tax schemes, our procedures of greatest promise, there may be multiple equilibria, only one of which will be efficient. To find and enforce the optimal solution may be a costly and complex procedure.

Despite these pessimistic findings, it would be wrong to conclude that a laissez-faire approach should be chosen by default. That mode of non-treatment has led to very serious present symptoms. Even rudimentary remedies, such as effluent charges or quota schemes, might be better than leaving matters untended. Surely we should conduct extensive research, both experimental and theoretical, to find positive treatments that might ameliorate existing conditions.

ACKNOWLEDGMENTS

An earlier version of this paper was presented at the Second World Econometrics Congress, September, 1970. This research was supported by NSF Grant GS-2797 (Starrett) and NSF Grant GR-88 to the Analytic Methods in Public Policy Faculty Seminar (Zeckhauser). That paper was distributed as Starrett and Zeckhauser [12].

REFERENCES

1. M. Aoki, Two planning processes for an economy with production externalities, Harvard Institute of Economic Research Discussion Paper No. 157, Harvard University, Cambridge, Mass., 1970.

2. K. J. Arrow, The organization of economic activity: issues pertinent to the choice of market versus non-market allocation, in The Analysis

and Evaluation of Public Expenditures: The PPB System, U. S. Congress, Joint Economic Committee, U. S. Government Printing Office, Washington, 1969, pp. 47-64.

3. J. Buchanan, External diseconomies, corrective taxes and market structure, Amer. Econ. Rev., 59, 174-177 (1969).

4. R. H. Coase, The problem of social cost, J. Law Econ., 3, 1-44 (1960).

5. F. Dolbear, On the theory of optimum externality, Amer. Econ. Rev., 57, 90-103 (1967).

6. E. Lindahl, Just taxation - a possible solution, in Classics in the Theory of Public Finance (R. A. Musgrave and A. T. Peacock, eds.), Macmillan, London, 1958, pp. 168-176.

7. J. Meade, External economies and diseconomies in a competitive situation, Econ. J., 62, 54-67 (1952); reprinted in Readings in Welfare Economics (K. J. Arrow and T. Scitovsky, eds.), Irwin, Homewood, Illinois, 1969, pp. 185-198.

8. A. C. Pigou, The Economics of Welfare, Macmillan, New York, 1920 (1st ed.) and 1932 (4th ed.).

9. L. Shapley and M. Shubik, On the core of an economic system with externalities, Amer. Econ. Rev., 59, 678-684 (1969).

10. D. Starrett, A note on externalities and the core, Harvard Institute of Economic Research Discussion Paper No. 198, Harvard University, Cambridge, Mass., 1971.

11. D. Starrett, Fundamental non-convexities in the theory of externalities, J. Econ. Theory, 4, 180-196 (1972).

12. D. Starrett and R. Zeckhauser, Treating external diseconomies - markets or taxes? Public Policy Program Discussion Paper No. 3, Harvard University, Cambridge, Mass., 1971.