Chapter 1

SEQUENTIAL CONTRIBUTIONS TO PUBLIC GOODS

I examine games involving private contributions to a public good and show that less of the public good will be supplied if agents move sequentially than if they move simultaneously. If the agents bid for the right to move first, the agent who values the public good least will win. If each agent chooses the rate at which he will subsidize the other agent's contributions, the subsidies that support the Lindahl allocation are the unique equilibrium outcome. I also describe two related subsidy-setting games that yield Lindahl allocations in n-person games with general utility functions. This paper will appear in the Journal of Public Economics.

Several authors have examined the private provision of public goods in simultaneous-move games. The Nash equilibria in these games turn out to have several surprising and interesting properties. For details see Warr (1983), Cornes-Sandler (1986) and Bergstrom-Blume-Varian (1986). In this paper I investigate games where agents decide on their contributions to a public good sequentially. In this sort of "Stackelberg contribution game" the agent who moves first can credibly commit to his contribution in a way that is not possible in a simultaneous-move game.

Admati-Perry (1991) analyze a game in which agents alternate contributions to joint project. However, in their game the project is either completed or not, and no benefits are generated from a partially completed

This work was supported by the National Science Foundation Grant SES-8800114 and a Fulbright grant. I wish to thank Ted Bergstrom, Joel Guttman, Leif Danziger, Robert McClelland, Ig Horstman, the participants at the Michigan-Western Ontario economic theory workshop and an anonymous referee for their comments and suggestions. I also wish to thank the Santa Fe Institute for their hospitality during the period of this research.

project. In this paper, by contrast, the focus is on the amount of the public good that is provided in the contribution game.

It turns out that the ability to commit to a contribution exacerbates the free rider problem: I show that the total amount of the public good provided in a sequential game is never larger than the amount provided in a simultaneous-move game. Along the way, I establish several other interesting results concerning equilibria in sequential contribution games. In particular, I examine some mechanisms in which agents choose to subsidize the contributions made by other agents, and show that these mechanisms provide simple and natural solutions to the problem of implementing Lindahl equilibria.

1.1 An example with quasilinear utility

It is instructive to start with a simple example with two agents. Each agent i divides his wealth w_i between private consumption, $x_i \geq 0$, and a contribution to a public good, $g_i \geq 0$. The total amount of the public good is $G = g_1 + g_2$.

Each agent's utility function is linear in his private consumption and a concave increasing function of G, so that the utility of agent i is

$$u_i(G) + x_i = u_i(g_1 + g_2) + w_i - g_i.$$

Let $\overline{g}_i$ be the amount of the public good that maximizes agent i's utility when the other agent contributes zero. We call this amount agent i's *standalone contribution.* We say that agent i *likes the public good more than agent j* if $\overline{g}_i > \overline{g}_j$; i.e., if agent i would contribute more to the public good if he were the only contributor.[1] We assume that $w_i > \overline{g}_i$ so that consumption of the private good, x_i, is always strictly positive. This makes it possible to drop w_i, since it is an inessential constant in each agent's utility function.

The reaction function

The reaction functions and Nash equilibria for the simultaneous contribution game are reasonably well known. We sketch the derivation here for purposes of comparison to the sequential case.

First, we derive the reaction function for agent 2. The first-order condition if agent 2 contributes a positive amount to the public good is

$$u_2'(g_1 + g_2) = 1.$$

[1] It is easy to show that if one agent's marginal willingness to pay for the public good agent is uniformly larger than the other's, then his standalone contribution will be larger.

Letting $G_2(g_1)$ be agent 2's reaction function, we must have

$$u_2'(g_1 + G_2(g_1)) = 1.$$

Since $\overline{g}_2$ is the amount that agent 2 contributes when $g_1 = 0$, we have:

$$G_2(g_1) = \overline{g}_2 - g_1.$$

This derivation is valid only when agent 2 contributes a positive amount to the public good. Since $g_2 \geq 0$, we must have

$$G_2(g_1) = \max\{\overline{g}_2 - g_1, 0\}.$$

This "kink" in the reaction function is what makes the analysis interesting.

The Nash equilibrium

A Nash equilibrium is a pair of contributions (g_1, g_2) such that

$$\begin{aligned} g_1 &= G_1(g_2) \\ g_2 &= G_2(g_1). \end{aligned}$$

A Nash equilibrium is depicted in Figure 1. In the case depicted, agent 1 likes the public good more than agent 2. In this case agent 1 contributes the entire amount of the public good and agent 2 free rides. If both agents have the same tastes for the public good, the reaction functions overlap and there is a whole range of equilibrium contributions, although there is still a unique equilibrium amount of the public good.

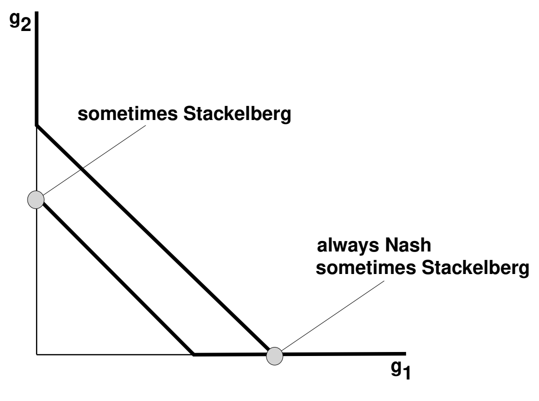

Figure 1. Nash and Stackelberg equilibria. In the Nash equilibrium the player who likes the public good the most contributes everything and the other player free rides. In the Stackelberg equilibrium the first contributor may free ride even if he likes the public good more than the other agent.

The Stackelberg equilibrium

We assume that agent 1 moves first. The utility of agent 1 as a function of his contribution is given by

$$\begin{aligned} V_1(g_1) &= u_1(g_1 + G_2(g_1)) - g_1 \\ &= u_1(g_1 + \max\{\overline{g}_2 - g_1, 0\}) - g_1. \end{aligned}$$

We can also write this as

$$V_1(g_1) = \begin{cases} u_1(\overline{g}_2) - g_1 & \text{for } g_1 \le \overline{g}_2 \\ u_1(g_1) - g_1 & \text{for } g_1 \ge \overline{g}_2. \end{cases}$$

It is clear from Figure 2 that there are two possible optima: either the first agent contributes zero and achieves utility $u_1(\overline{g}_2)$ or he contributes $\overline{g}_1$ and achieves utility $u_1(\overline{g}_1) - \overline{g}_1$.

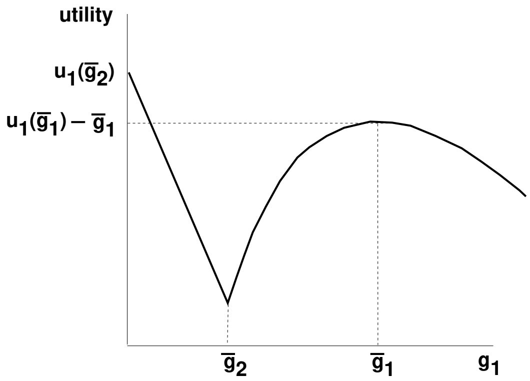

Figure 2. Utility of the first contributor as a function of his gift. In this example, agent 1's optimal choice is to contribute zero, but if the "hump" were higher, he would want to contribute $\overline{g}_1$.

Case 1. *The agent who likes the good least is the first contributor.* In this case the optimal choice by the first player is to contribute zero. This is true since

$$u_1(\overline{g}_2) > u_1(\overline{g}_1) > u_1(\overline{g}_1) - \overline{g}_1.$$

Case 2. *The agent who likes the public good the most is the first contributor.* In this case, either contributor may free ride. The easiest way to see this is by example. Suppose that agent i's utility for the public good is $u_i(G) = a_i \ln G$, so that agent i's standalone contribution is a_i. Then the first contributor will get utility $\ln a_1 - a_1$ or $\ln a_2$, and either of these may be larger. ∎

If the agents have tastes that are very similar, then the first contributor will free ride on the second's contribution. However, if the first agent likes the public good *much* more than the second, then the first agent may prefer to contribute the entire amount of the public good himself.

Referring to Figure 1 we see that there are two possible Stackelberg equilibria: one is the Nash equilibrium, in which the agent who likes the good most contributes everything. The other Stackelberg equilibrium is where the agent who likes the good *least* contributes everything. This equilibrium cannot arise as a Nash equilibrium since the threat to free ride by the agent who likes the public good most is not credible in the simultaneous-move game.

Note that it is always advantageous to move first since there are only two possible outcomes and the first mover gets to pick the one he prefers. Also note that the sum of the utilities is higher at the higher level of the public good. It follows that if you want to ensure that the higher level of the public good is provided, then you should make sure that the person who likes the good *least* moves first.

1.2 Examples

Here we briefly describe two examples of free riding in games involving sequential contributions to public goods. The first is a variation on the Samaritan's dilemma as described by Buchanan (1975), Varian (1982), and Lindbeck (1988). Consider a game between the young generation and the old generation. The old generation has two choices: they can save for their retirement, or they can squander their earnings on their consumption. The young generation likewise has two strategies: they can support the older generation when it retires, or they can let them starve.

A natural assignment of payoffs to these strategies implies that there are two Nash equilibria in this game. In one equilibrium, the older generation expects younger generation to support them and they therefore choose to squander. Given the choice between letting the older generation starve or providing support, the younger generation chooses to provide support. In the other Nash equilibrium, the older generation expects the younger generation not to provide support, and so they prudently choose to save. Since they reach retirement with adequate resources, the younger generation chooses not to support them.

However, only one of these equilibria survives when we take account of the fact that the older generation gets to move first. The unique subgame perfect equilibrium is for the older generation to squander, recognizing that the younger generation will be forced to provide for them in their old age.

Even though the older generation cares more about their consumption than the younger generation does, the older generation still finds it optimal to free ride on the younger generation's contribution.

The second example has to do with reproductive strategy in animals. Males and females normally contribute equal amounts of genetic material to their offspring. However, they often contribute unequally to the care of the offspring. The degree to which males and females provide care for their offspring depends, in part, on the timing of their choices. In mammals, the male is generally the first mover when he fertilizes the eggs of the female. Once the female egg is fertilized the male may well depart, leaving most childcare responsibilities to the female.

In some species of fish, the childcare responsibilities are completely reversed. Among fish, the female is the first mover when she lays the eggs. The male moves second by fertilizing them. At that point the female may well depart, leaving the male to care for the eggs.

Forsyth (1986) describes a particularly interesting case in which there is yet another level of free-riding. There are two types of bluegill sunfish: large "territorial" males and much smaller "satellite" males. The territorial males build nests in order to attract females, while the satellite males hover around the nest built by the territorial male. After a female selects a nest and releases her eggs, the satellite male rushes into the nest and releases its sperm at the same time as the territorial male attempts to fertilize the eggs. The territorial male then has no recourse but to care for the eggs, even though only a fraction of them carry his genetic material, since he is better able to fend off predators than the smaller satellite male. The female bluegill and the satellite male then blithely swim away, free riding on the childcare activities of the territorial male!

1.3 Bidding for the right to move first

Since the first mover always has an advantage in the sequential contribution game, we might consider auctioning off the right to move first. If the first mover likes the public good much more than the second, then he will provide the entire amount of the public good anyway, so it is no advantage to him to be first mover. The advantage to the first mover only arises when the players have similar tastes for the public good. In this case, each player would prefer to move first and free ride on the other's contribution.

Consider, then, the case where agent 1 likes the public good a bit more than agent 2, so that $\overline{g}_1 > \overline{g}_2$, but not so much more as to contribute everything himself. Agent 1 can get utility $u_1(\overline{g}_2)$ by moving first and free riding. If he moves second, he gets utility $u_1(\overline{g}_1) - \overline{g}_1$. The amount that he would be willing to bid to move first, b_1, is therefore

$$b_1 = u_1(\overline{g}_2) - u_1(\overline{g}_1) + \overline{g}_1.$$

We assume that the bids are either thrown away or given to the other agent. Since there are no income effects, it doesn't matter what is done with the bids.

The difference between the bids of the two agents is

$$b_1 - b_2 = [u_1(\overline{g}_2) - u_1(\overline{g}_1)] + [u_2(\overline{g}_2) - u_2(\overline{g}_1)] - [\overline{g}_2 - \overline{g}_1]. \tag{1.1}$$

It seems plausible to suppose that the agent who values the public good more would be willing to pay more to move first. However, this is exactly wrong! Under our assumptions, the agent who values the public good *least* is willing to pay more for the first-mover position.

To see this, note that concavity of the utility functions gives us the following inequalities:

$$\begin{aligned} u_1(\overline{g}_2) - u_1(\overline{g}_1) &\leq u_1'(\overline{g}_1)[\overline{g}_2 - \overline{g}_1] \\ u_2(\overline{g}_2) - u_2(\overline{g}_1) &\leq u_2'(\overline{g}_1)[\overline{g}_2 - \overline{g}_1]. \end{aligned}$$

Substituting these into (1.1) we have

$$b_1 - b_2 \leq [u_1'(\overline{g}_1) + u_2'(\overline{g}_1) - 1](\overline{g}_2 - \overline{g}_1).$$

Since $u_1'(\overline{g}_1) = 1$, this simplifies to

$$b_1 - b_2 \leq u_2'(\overline{g}_2)[\overline{g}_2 - \overline{g}_1] < 0,$$

where the last inequality follows since $\overline{g}_1 > \overline{g}_2$.

It follows that $b_1 < b_2$. That is, the agent who likes the public good the *least* will be willing to pay the *most* in order to move first. As we've seen, this will ensure that the largest amount of the public good will be provided. Essentially each agent is bidding for the right to free ride on the other agent, and it better to free ride on someone who will provide a lot of the public good than someone who will provide only a little.

1.4 Subsidizing the other agent

In the game considered above, one agent is able to commit to a contribution to the public good before the other agent makes his choice. The contribution by the first agent affects the *benefits* that the second agent receives from his contribution. In this section we examine public goods games in which the agents can influence the *cost* to other agents of their contributions. In particular, we examine what will happen if one agent has the opportunity to subsidize the other agent's contributions.

For simplicity we consider the case of identical utilities. As we have seen, the Nash equilibrium amount of the public good, G^n, is determined by the condition

$$u_i'(G^n) = 1,$$

and any set of contributions (g_1^n, g_2^n) such that $g_1^n + g_2^n = G^n$ is a Nash equilibrium.

Suppose that agent 1 offers to subsidize agent 2's contributions at rate s_2. The payoffs to the two agents become

$$u_1(g_1 + g_2) - g_1 - s_2 g_2$$
$$u_2(g_1 + g_2) - (1 - s_2)g_2.$$

Since agent 2 now faces a lower cost of contributing than agent 1, agent 2 will contribute the entire amount of the public good. The amount that agent 2 contributes, $G(s_2)$, is determined by the equation

$$u_2'(G(s_2)) = 1 - s_2.$$

Given our assumption that $u_2(G)$ is strictly concave, $G(s_2)$ will be a continuous increasing function. Agent 1's utility from offering the subsidy rate s_2 is

$$u_1(G(s_2)) - s_2 G(s_2).$$

Originally agent 1 had utility $u_1(G^n) - g_1^n = u_1(G(0)) - g_1^n$. The increase in agent 1's utility from offering the subsidy is

$$u_1(G(s_2)) - u_1(G(0)) + g_1^n - s_2 G(s_2). \tag{1.2}$$

As s_2 approaches zero, this expression converges to $g_1^n \geq 0$. This implies that if agent 1 was initially contributing a positive amount, he would be strictly better off by offering to subsidize agent 2's contribution by a sufficiently small amount. Intuitively, if agent 1 offers a very tiny subsidy to agent 2, agent 2 will end up contributing the entire amount of the public good in the second stage. But a tiny subsidy hardly costs agent 1 anything, which means that each agent will always want to subsidize the other agent's contribution.

Roberts (1987) and Bergstrom (1989) show that if a subsidy on contributions to a public good is financed by an equal lump-sum tax, then each agent prefers that the other agent be subsidized. We have shown that each agent prefers to subsidize the other agent even if he must pay for the subsidy himself. However, the Roberts-Bergstrom result holds for general preferences, while our result only holds for quasilinear preferences.

1.5 Equilibrium subsidies

We have seen that each agent will prefer to subsidize the other agent in our contribution game. Suppose that each agent simultaneously names a rate at which he is willing to subsidize the other agent. Then, given these subsidy rates, the agents play a simultaneous contribution game. What is the equilibrium of this game?

In order to answer this question, we need some facts about the Lindahl allocation in this model. In the case of quasilinear utility there will be a unique amount of the public good that the sum of the utilities. This amount, G^e, satisfies the first-order condition

$$u_1'(G^e) + u_2'(G^e) = 1. \tag{1.3}$$

This is just the familiar condition that maximizes the sum of the marginal willingnesses-to-pay must equal marginal cost.

Suppose that we choose rates s_i^e to support the efficient amount of the public good.

$$\begin{aligned} u_1'(G^e) &= 1 - s_1^e \\ u_2'(G^e) &= 1 - s_2^e, \end{aligned} \tag{1.4}$$

Note that (1.3) and (1.4) together imply

$$s_1^e + s_2^e = 1. \tag{1.5}$$

Equation (1.5) implies that the utility of agent 1 at the Lindahl allocation is

$$u_1(G^e) - (1 - s_1^e)g_1 - s_2^e g_2 = u_1(G^e) - (1 - s_1^e)G^e.$$

Hence $s_2^e = 1 - s_1^e$ is effectively a Lindahl price for agent 1; accordingly, we call (s_1^e, s_2^e) the *Lindahl subsidies.*

Theorem 1. *The unique subgame perfect equilibrium of the subsidy setting game yields the Lindahl subsidies, and the resulting allocation is the Lindahl allocation.*

Proof. If we have an interior equilibrium in the contribution stage, the amount of the public good must satisfy:

$$\begin{aligned} u_1'(g_1^* + g_2^*) &= 1 - s_1 \\ u_2'(g_1^* + g_2^*) &= 1 - s_2 \end{aligned} \tag{1.6}$$

Suppose that agent 2, say, is not contributing. There is no cost to agent 1 of increasing s_2 up to the point where agent 2 is just on the verge of contributing. We assume that this has been done, which means that (1.6) will apply even in the case of boundary solutions.

If agent 1 increases the subsidy agent 2 faces slightly then agent 2 will do all the contributing, and agent 1 will contribute zero. If we are in equilibrium, such a change cannot benefit agent 1. Using a limiting argument similar to that used to establish equation (1.2) this implies

$$u_1(g_1^* + g_2^*) - (1 - s_1)g_1^* - s_2 g_2^* \geq u_1(g_1^* + g_2^*) - s_2(g_1^* + g_2^*).$$

After simplification, we have

$$g_1^*(s_1 + s_2 - 1) \geq 0. \tag{1.7}$$

Similarly, if agent 1 reduces the subsidy rate that agent 2 faces, then agent 1 will do all the contributing. In equilibrium agent 1 cannot benefit from this change which implies

$$u_1(g_1^* + g_2^*) - (1 - s_1)g_1^* - s_2 g_2^* \geq u_1(g_1^* + g_2^*) - (1 - s_1)(g_1^* + g_2^*).$$

Simplification yields

$$g_2^*(s_1 + s_2 - 1) \leq 0. \tag{1.8}$$

The same arguments applied to agent 2 yield

$$g_2^*(s_1 + s_2 - 1) \geq 0 \tag{1.9}$$

$$g_1^*(s_1 + s_2 - 1) \leq 0. \tag{1.10}$$

At least one of (g_1^*, g_2^*) must be nonzero in equilibrium. This observation, together with (1.7)–(1.10), implies $s_1 + s_2 = 1$. It follows from (1.6) that the resulting allocation is the Lindahl allocation. ∎

The intuition behind this result is rather nice: if agent 1 contributes *directly* to the public good, it costs him a dollar for each dollar he contributes. But if he contributes *indirectly*, say by giving a ten percent subsidy, then it costs agent 1 only ten cents for each dollar that agent 2 contributes to the public good. Each agent will continue to subsidize the other until the marginal cost of contributing directly is equal to the marginal cost of contributing indirectly. In the case of identical consumers, this is where each agent gives the other a fifty percent subsidy—which is is the Lindahl price.

Related literature

There is a large literature on designing mechanisms to solve the public goods problem. Groves (1979), Groves-Ledyard (1987), Moore (1991), and Ledyard (1992) provide surveys of various aspects of this literature.

The classic solutions to the free-rider problem are due to Groves-Ledyard (1977), Hurwicz (1979), and Walker (1981). However, the most closely related result to ours is Guttman (1978, 1986). In the first stage of Guttman's game each agent announces a rate at which he will match the other agent's contribution. In the second stage each agent contributes a "flat" amount plus a matching amount tied to the other agent's direct contribution. Guttman argues that the equilibrium of this game is the Pareto efficient amount of the public good.[2] In Guttman's game, the matching contribution is goes directly to the public good; in my game, the subsidy is paid to the other contributor. However, the strategic nature of the two games are similar since each agent is effectively setting a price for the other agent's contribution. Guttman (1986) describes some experimental evidence in his matching game.

More recently, Jackson-Moulin (1992) describe a two-stage bidding game that yields a Lindahl-like allocation in the case of a *discrete* public good. In the first stage of the Jackson-Moulin game, each agent announces the total value of the public good. In the second stage, each agent announces his own valuation. The Jackson-Moulin game implements an efficient allocation in undominated Nash equilibrium, although they offer a variation that implements efficiency in subgame perfect equilibrium.

In Varian (1989b) I describe a subsidy-setting game for implementing Lindahl-like outcomes for completely general externalities problems. The mechanism in Varian (1989b) is more general than the subsidy setting game described here, but it is also more complicated. The nice thing about the subsidy game is that it is very natural.

1.6 General utility functions

I have described several results for public goods problems with two stages in which a price or a quantity is set in the first stage, and further quantity choices are made in the second stage. However, the quasilinear case is very special: do these results generalize to a model with income effects?

Let $u_i(G, x_i)$ be a utility function with G is the level of the public good and x_i is the private consumption of agent i. As before, I assume that utility is a differentiable, strictly concave function.

I first derive the form of the reaction function. Agent 2's maximization problem is

$$\begin{aligned} &\max_{x_2, g_2} u_2(g_1 + g_2, x_2) \\ \text{such that } &g_2 + x_2 = w_2 \\ &g_2 \geq 0. \end{aligned}$$

We can add g_1 to each side of the constraints and use the definition $G = g_1 + g_2$ to rewrite the problem as

$$\begin{aligned} &\max_{x_2, G} u_2(G, x_2) \\ \text{such that } &G + x_2 = w_2 + g_1 \\ &G \geq g_1. \end{aligned}$$

[2] Guttman (1978) did not contain a complete proof of this result due to space limitations.

In this problem, agent 2 is choosing the level of the public good, subject to the constraint that the level that he chooses is at least as large as the contribution of agent 1.

Following Bergstrom-Blume-Varian (1986) we note that this problem is a standard consumer demand problem except for the inequality constraint. Let $f_2(w)$ be agent 2's demand function for the public good. It follows from the above remarks that

$$G = g_1 + g_2 = \max\{f_2(w_2 + g_1), g_1\}.$$

Subtracting g_1 from each side of this equation, we have the reaction function

$$G_2(g_1) = \max\{f_2(w_2 + g_1) - g_1, 0\}.$$

According to this reaction function, agent 1 will either contribute zero or the amount of the public good that he would demand if his wealth were $w_1 + g_2$ *minus* the amount contributed by the other agent. The following assumption is quite natural:

Normal goods. Both the public and the private good are strictly normal goods at all levels of wealth.

Given this assumption it is easy to see the general shape of the reaction function. When $g_1 = 0$, agent 2 will contribute $f_2(w_2)$. As g_1 increases, the contribution of agent 2 will decrease, but less than one-for-one. For some g_1^c we may have $f_2(w_2 + g_1^c) = 0$; at this point agent 1 contributes so much that agent 2 chooses to free ride. We call this amount the *complete crowding out* contribution.

We summarize some properties of the reaction function in the following fact, the proof of which follows immediately from the assumption.

Fact 1. *The reaction function $G_2(g_1)$ is a nonincreasing function. It will be strictly decreasing when it is not equal to zero. The function $H(g_1) = g_1 + G_2(g_1)$ is a strictly increasing function.*

As before, we can use this reaction function to calculate the Nash equilibria and the Stackelberg equilibrium. A Nash equilibrium is a solution (g_1^n, g_2^n) to

$$\begin{aligned} g_1^n &= \max\{f_1(w_1 + g_2^n) - g_2^n, 0\} \\ g_2^n &= \max\{f_2(w_2 + g_1^n) - g_1^n, 0\}. \end{aligned}$$

A Stackelberg equilibrium is a pair $(g_1^s, G_2(g_1^s))$ for which g_1^s solves

$$\max_{g_1} \; u_1(g_1 + \max\{f_2(w_2 + g_1) - g_1, 0\}, w_1 - g_1).$$

We want to compare the solutions of these two sets of equations. This comparison is made simpler by noting that Bergstrom-Blume-Varian (1986) have proved that under the normality assumption we have made there is a *unique* Nash equilibrium. There will also be one Stackelberg equilibrium for each ordering of the agents.

1.7 Results for general utility functions

We have three sets of results. The first set of results concerns who contributes and who free rides. The second set of results concerns the effect of redistributions of wealth. The third set of results concerns how the amount of the public good provided in the Stackelberg equilibrium compares to the amount provided in the Nash equilibrium.

Free riding

Fact 2. *If the standalone contribution is less than the complete crowding out contribution ($\overline{g}_1 < g_1^c$) then both agents must contribute in the Stackelberg equilibrium.*

Proof. Evaluate the right derivative of agent 1's utility function at g_1^c. We have

$$\frac{\partial u_1(g_1^c, w_1 - g_1^c)}{\partial G} - \frac{\partial u_1(g_1^c, w_1 - g_1^c)}{\partial x_1} < 0.$$

The inequality follows since the derivative equals zero at $\overline{g}_1$, and $g_1^c > \overline{g}_1$. (Recall that $u_1(g_1, w_1 - g_1)$ is a concave function.) It follows that agent 1's utility will increase if he contributes less than g_1^c, even if he is the only one to contribute. The fact that the other agent will also contribute can only increase the first agent's utility. Hence the Stackelberg equilibrium must involve contributions by both agents. ∎

Fact 3. *If there is a Nash equilibrium with $g_1^n = 0$, then this is also a Stackelberg equilibrium.*

Proof. By definition of Nash equilibrium, agent 2 is on his reaction curve, so we only need to show that agent 1 is on his reaction curve. If agent 1 contributes 0, then agent 2 will contribute $\overline{g}_2$. Let $g_1 > 0$ be any other possible contribution by agent 1. Then we have :

$$u_1(\overline{g}_2, w_1) > u_1(g_1 + \overline{g}_2, w_1 - g_1) > u_1(g_1 + G_2(g_1), w_1 - g_1).$$

The first inequality follows from the Nash assumption. The second inequality follows since $G_2(0) = \overline{g}_2$ and $G_2(g_1)$ is a nonincreasing function. ∎

Wealth redistribution

Fact 4. *Suppose that we have a Stackelberg equilibrium (g_1^s, g_2^s). Let $(\Delta w_1, \Delta w_2)$ be a redistribution of wealth from the leader to the follower such that $g_i + \Delta w_i \geq 0$ for $i = 1, 2$ with $\Delta w_1 < 0$. Then the Stackelberg equilibrium after this redistribution is $(g_1^s + \Delta w_1, g_2^s + \Delta w_2)$ and the total amount of the public good remains unchanged.*

Proof. The first-order condition for the Stackelberg equilibrium is

$$\frac{\partial u_1(g_1 + g_2, w_1 - g_1)}{\partial G} f_2'(w_2 + g_1) - \frac{\partial u_1(g_1 + g_2, w_1 - g_1)}{\partial x_1} = 0.$$

Let (Δw_i) be a redistribution of wealth and let agent i change his contribution by $\Delta g_i = \Delta w_i$ for $i = 1, 2$. Note that since $\Delta w_1 + \Delta w_2 = 0$ we must have $\Delta g_1 + \Delta g_2 = 0$. Now simply observe that if each agent changes his contribution in this way the first-order condition is still satisfied. ∎

Warr (1983) and Bergstrom-Blume-Varian (1986) show that essentially the same result holds in an (interior) Nash equilibrium. Interestingly, a transfer from the follower to the leader may not be neutral since the leader's objective function is not concave and a transfer may induce a discontinuous response.[3]

Bergstrom-Blume-Varian (1986) also investigate the boundary cases in some detail. In the two-agent context we are investigating here the analysis of the boundary cases is quite straightforward so we simply state the result.

Fact 5. *Suppose that person 1 is contributing and person 2 is not. Then a redistribution from 2 to 1 will increase the amount of the public good, while a redistribution from 1 to 2 can decrease or increase the amount of the public good.*

Comparison to the Nash equilibrium

Theorem 2. *The amount of the public good contributed by agent 1 in the Stackelberg equilibrium is never larger than the amount provided by agent 1 in the Nash equilibrium. That is, $g_1^s \leq g_1^n$.*

[3] See Rudolf Kerschbamer and Clemens Puppe, "Sequential Contributions to Public Goods: A Note on the Structure of the Equilibrium Set," Working Paper, University of Vienna, July 1998.

Proof. There are two cases to consider, the case where $g_2^n = 0$, and the case where $g_2^n > 0$. In the first case, it follows from Fact 3 that the Nash equilibrium is also a Stackelberg equilibrium. Hence, $g_1^n = g_1^s$.

As for the second case, let $V_1(g_1) = u_1(g_1 + G_2(g_1), w_1 - g_1)$. The derivative of $V_1(g_1)$ evaluated at the Nash equilibrium is:

$$\begin{aligned} V_1'(g_1^n) &= \frac{\partial u_1(G^n, w_1 - g_1^n)}{\partial G}[1 + G_2'(g_1^n)] - \frac{\partial u_1(G^n, w_1 - g_1^n)}{\partial x_1} \\ &= \frac{\partial u_1(G^n, w_1 - g_1^n)}{\partial G} G_2'(g_1^n) < 0. \end{aligned}$$

This follows since $\partial u_1/\partial G - \partial u_1/\partial x_1 = 0$ at the Nash equilibrium and $G_2'(g_1^n)$ is strictly negative when $g_2^n > 0$. This shows that agent 1's utility will increase by decreasing his contribution from which it follows that $g_1^s < g_1^n$. ∎

Corollary. *The total amount of the public good in the Stackelberg equilibrium is less than or equal to the total amount provided in the Nash equilibrium.*

Proof. According to Fact 1, the function $H(g_1) = g_1 + G_2(g_1)$ is an increasing function. Therefore,

$$H(g_1^n) = g_1^n + G_2(g_1^n) = g_1^n + g_2^n \geq g_1^s + G_2(g_1^s) = H(g_1^s).$$

The corollary follows. ∎

1.8 Public goods and general utilities

Here we describe some generalizations of the two-person subsidy-setting game described earlier. Each mechanism is a variation on the compensation mechanism described in Varian (1989b).

Subsidy setting with more than two agents

The subsidy-setting game described earlier only works for two agents with quasilinear utility. If there are more than two agents, each agent still prefers that the other agents' contributions be subsidized—but each agent wants someone else to do the subsidizing. However, by adjusting who subsidizes whom, it is possible to overcome this sort of free-riding in subsidy setting.

The trick is to have agent 1 set the rate at which agent 2 will subsidize agent 3's contributions. Agent 2, in turn sets the rate at which agent 3 will subsidize 1's contributions, and so on. More specifically, let s^i_j be the subsidy facing agent j as set by agent i. As before, let x_1 be agent 1's private consumption, and g_1 be his contribution to the public good. If there are 3 agents doing the contributing, agent 1's budget constraint will be:

$$x_1 + (1 - s^2_1 - s^3_1)g_1 = w_1 - s^3_2 g_2 - s^2_3 g_3.$$

Note that, unlike the two-person case, the prices agent i faces are independent of the prices he sets.

In order to establish the result, we need a regularity assumption on demand functions: they must be be locally invertible functions of the subsidy rates. This means that if agent 1 wants agent 2 to contribute a little more or a little less to the public good, there is some subsidy rate that he can set that will induce agent 2 to do this. Local invertibility will be satisfied if the derivative of each agent's demand function with respect to price is not zero at the equilibrium.[4]

Theorem 3. *Let each agent have continuous convex preferences. Assume that the demand functions are locally invertible. Then the subgame perfect equilibria of the subsidy setting game are Lindahl allocations.*

Proof. We prove the theorem for 3 agents, but the idea extends to an arbitrary number of agents. Let (x^*_i, g^*_i) be a subgame perfect equilibrium of this game, and let (x'_i, g'_i) be an allocation that Pareto dominates it. Since preferences are convex and continuous, we can assume that (x'_i, g'_i) is arbitrarily close to (x^*_i, g^*_i).

Agent i chooses his own contribution directly and chooses the other agents' contributions through the subsidy rates that he sets for them. If each agent prefers (x'_i, g'_i) to (x^*_i, g^*_i), then this allocation must not be affordable for each agent at the equilibrium subsidy rates. This yields three inequalities:

$$\begin{aligned}
x'_1 + (1 - s^2_1 - s^3_1)g'_1 &> w_1 - s^3_2 g'_2 - s^2_3 g'_3 \\
x'_2 + (1 - s^1_2 - s^3_2)g'_2 &> w_2 - s^1_3 g'_3 - s^3_1 g'_1 \\
x'_3 + (1 - s^1_3 - s^2_3)g'_3 &> w_3 - s^1_2 g'_2 - s^2_1 g'_1.
\end{aligned}$$

Adding these inequalities gives

$$\sum_{i=1}^{3} x'_i + \sum_{i=1}^{3} g'_i > \sum_{i=1}^{3} w_i,$$

which shows that the Pareto dominating allocation is not feasible.

[4] This is automatically satisfied if utility is quasilinear and strictly concave.

Next we show that the equilibrium is Lindahl. Agent 1 can choose g_1 directly and choose g_2 or g_3 indirectly, through the subsidy he sets for the other agents. Since the contributions are perfect substitutes in consumption, they must have the same equilibrium price. Hence we can define $p_1 = 1 - s_1^2 - s_1^3 = s_3^2 = s_2^3$. From this it follows that agent 1's budget constraint can be written as

$$x_1 + p_1(g_1 + g_2 + g_3) = x_1 + p_1 G = w_1,$$

which shows that the allocation is Lindahl. ∎

It is worth observing that this mechanism is balanced both in and out of equilibrium. The proof that the equilibrium is Lindahl is essentially the proof given in Varian (1989b) for general externalities, specialized to the case of public goods.

Another mechanism for implementing Lindahl allocations

Here is another two-stage mechanism that implements Lindahl allocations. The message space in this mechanism is much smaller than in the previous mechanism since each agent has to name only a single number.

The price setting stage. Each agent i announces a number q_i. The price for agent i's contribution to the public good is the average of the numbers named by the other agents:

$$p_i = \frac{1}{n-1} \sum_{j \neq i} q_j.$$

The contribution stage. Each agent i chooses (x_i, g_i) to maximize his utility subject to the budget constraint

$$x_i + p_i g_i = w_i - p_i \sum_{j \neq i} g_j - Q(p),$$

where the penalty term is given by $Q(p) = \left(\sum_{j=1}^{n} p_j - 1\right)^2$.

In the second stage we have each agent choosing a contribution to the public good for which he pays a price $p_i g_i$. But agent i must also make a payment based on the amount of the good contributed by the *other* agents.

Since the total amount of the public good is $G = g_i + \sum_{j \neq i} g_j$, we could also write agent i's budget constraint as $x_i + p_i G = w_i - Q(p)$.

It is easy to show that the Lindahl allocation (p_i^e, G^e) is *an* equilibrium of this game. To do this we only need to show that if every agent makes announcements which lead to the Lindahl prices, agent i cannot increase his utility by announcing something that yields non-Lindahl prices for the other agents.

To prove this we first observe that since the Lindahl prices (p_i^e) result in an efficient amount of the public good, the penalty term, $Q(p)$, must be zero by the standard Samuelson efficiency condition.[5] Suppose that agent i, announces some numbers that change the price vector p^e to p' and that this change results in some possibly different amount of the public good, G'. Note that agent i can only affect the prices facing the other agents, not the price he faces.

We have

$$u_i(G^e, w_i - p_i^e G^e) \geq u_i(G', w_i - p_i^e G') \geq u_i(G', w_i - p_i^e G' - Q(p')).$$

The first inequality comes from the fact that we start with a Lindahl allocation; the second inequality comes from the fact that $Q(p') \geq 0$. This argument shows that agent i is at least as well off announcing prices that lead to the Lindahl prices (p_j^e) as any other prices; i.e., that Lindahl allocations are an equilibrium to the mechanism. In the appendix I show that there can be no other equilibria of the mechanism.

I've described this mechanism in terms of setting prices. However, it can also be described in terms of setting subsidies, which makes it look more like the mechanisms described earlier. In the subsidy-setting framework, each agent i names a number $1 - s_i$ which turns out in equilibrium to be the rate at which agent i's contributions are subsidized *and* the rate at which agent i subsidizes everyone else's contributions. We denote the contributions by everyone except agent i by $G_{-i} = \sum_{j \neq i} g_j$. The budget constraint facing agent i then can be written as

$$(1 - \sum_{j \neq i} s_j) g_i + x_i = w_i - (1 - \sum_{j \neq i} s_j) G_{-i} - (1 - \sum_{j=1}^{n} s_j)^2.$$

It can be shown that in equilibrium $s_i + \sum_{j \neq i} s_j = 1$, so that agent i ends up being subsidized at the rate $1 - s_i$ and subsidizing the other agents at the same rate.

Related literature

Danziger-Schnytzer (1991) have independently examined a subsidy-setting mechanism similar in spirit to this one. In the Danziger-Schnytzer mechanism, each agent names a rate at which he will subsidize the other agents'

[5] We assume that it is efficient to provide a positive amount of the public good.

contributions. In in my mechanism, the rate at which each agent subsidizes the other agents is set by the *other* agents. Furthermore, my mechanism requires a penalty function while a penalty is not necessary in the Danziger-Schnytzer mechanism. The Danziger-Schnytzer mechanism is therefore a bit simpler in the case of pure public goods. However, I show in Varian (1989b) that my mechanism also works for general externalities problems.

1.9 Incomplete information

Up until now we have examined games where each agent knows the preferences and wealth of the other agent. Here we consider a model where each contributor has incomplete information about the other contributor. In our game the second contributor reacts passively, making his optimal choice given the first agent's contribution. Hence it is irrelevant whether or not he knows anything about the first contributor. The only interesting uncertainty concerns the first contributor's knowledge of the second contributor's preferences.

Consider the quasilinear model examined earlier. In this case all that is relevant from the first contributor's point of view is the value of $\overline{g}_2$—how much the second person will contribute if the first person contributes zero. Suppose that the first contributor has a prior distribution on how much the other person will contribute and seeks to maximize expected utility.[6]

The expected utility of the first contributor is

$$\begin{aligned} V_1(g_1) &= \int_0^\infty [u(g_1 + \max\{\overline{g}_2 - g_1, 0\}) - g_1] f(\overline{g}_2)\, d\overline{g}_2 \\ &= \int_0^{g_1} [u(g_1) - g_1] f(\overline{g}_2)\, d\overline{g}_2 + \int_{g_1}^\infty [u(\overline{g}_2) - g_1] f(\overline{g}_2)\, d\overline{g}_2 \\ &= [u(g_1) - g_1] F(g_1) + \int_{g_1}^\infty [u(\overline{g}_2) - g_1] f(\overline{g}_2)\, d\overline{g}_2. \end{aligned}$$

Differentiating this expression with respect to g_1 and simplifying yields

$$V_1'(g_1) = u_1'(g_1) F(g_1) - 1.$$

Note that when $g_1 = 0$, the probability that $\overline{g}_2$ is less than g_1 is zero, so that $V_1'(0) = -1$. If g_1 is large enough so that agent 1 is certain that $\overline{g}_2$ is less than g_1, then $V_1'(g_1) = u_1'(g_1) - 1$. Hence, agent 1's utility as a function of his gift is similar to the shape depicted in Figure 2. Depending on the beliefs of agent 1 about agent 2's maximum contribution, agent 1

[6] We assume that the von Neumann-Morgenstern utility function takes the quasilinear form. This is restrictive, but seems necessary for a simple analysis.

will either choose to free ride, or to contribute an amount g_1^* that satisfies the condition

$$u_1'(g_1^*)F(g_1^*) - 1 = 0.$$

This marginal condition is quite intuitive. If agent 1 decides to contribute a bit more of the public good, he will get $u_1'(g_1^*)$, but only if agent 2 has $\overline{g}_2 < g_1^*$. Otherwise, agent 1 will get no incremental utility from his contribution—since his contribution would just crowd out some of the public good that agent 2 would have contributed anyway. Hence the *expected* marginal utility of agent 1's contribution is $u_1'(g_1^*)$ times the probability that $\overline{g}_2 < g_1^*$, which is just $u_1'(g_1^*)F(g_1^*)$. The optimal contribution is determined by the condition that this expected marginal utility must equal the (certain) marginal cost of the contribution.

How does this amount compare to $\overline{g}_1$, which is what agent 1 would contribute under certainty? Note that $V_1'(\overline{g}_1) = u_1'(\overline{g}_1)F(\overline{g}_1) - 1 = F(\overline{g}_1) - 1$. As long as there is some possibility that agent 2 will have $\overline{g}_2 > \overline{g}_1$, we will have $F(\overline{g}_1) < 1$ and $V_1'(\overline{g}_1)$ will be negative. If $V_1(g_1)$ is concave this implies that the equilibrium contribution in the presence of uncertainty is less than the contribution under certainty. Intuitively, the *possibility* that agent 2 may value the good more than agent 1 leads agent 1 to reduce his contribution to the public good, hoping to free ride on agent 2's contribution.

1.10 Summary

We have examined some sequential games involving contributions to a public good. If preferences are quasilinear, then:

1. The sequential equilibrium of the contribution game will provide the same or less of the public good than the simultaneous move game.

2. The player who likes the public good least will bid the most to move first.

3. Each player would like to subsidize the other player's contributions. If both players choose subsidy rates and then play the voluntary contribution game, a Lindahl equilibrium is the unique subgame perfect equilibrium of this two-stage game.

4. The equilibrium of the sequential move game is independent of small redistributions of wealth.

5. If the first agent is uncertain about the type of the second agent, he will tend to contribute less to the public good.

In the case of general utility functions, the amount of the public good supplied in the Stackelberg game will never be more than in the Nash game. Furthermore, there are several generalizations of the subsidy setting game that yield Lindahl allocations.

1.11 Appendix. Proof of uniqueness

Here we provide a proof of the claim made in the text that the only equilibria of our price-setting mechanism are Lindahl allocations. We assume that for some agent the marginal rate of substitution between the private and the public good is infinite when $G = 0$ so that someone will always want to make a positive contribution to the public good. We also assume local invertibility of the sort described earlier: there is a set of prices that will result any desired (small) change in behavior. See Varian (1989b) for more discussion of the role of this assumption.

Let (p_i^*) be the set of prices that result from the equilibrium announcements (q_i^*). We first show that in equilibrium we must have $\sum_{k=1}^n p_k^* = 1$. Suppose not. We have assumed that least one agent i will choose a positive amount of the public good. Suppose that this agent changes his announcement q_i. This will change agent i's utility by

$$\left[\frac{\partial u_i}{\partial G} - \frac{\partial u_i}{\partial x_i} p_i^*\right] \frac{\partial G}{\partial q_i} - 2\frac{\partial u_i}{\partial x_i}\left(\sum_{k=1}^n p_k^* - 1\right).$$

Since agent i is contributing a positive amount the bracketed expression vanishes by the envelope theorem, leaving us with a term that is, by assumption, nonzero. Hence there is some change in i's announcement that will increase his utility, contradicting the assumption of equilibrium. It follows that $\sum_{k=1}^n p_k^* = 1$ in equilibrium.

Now suppose that the equilibrium allocation is not Pareto efficient. Then there is some other feasible allocation (x_i', G') that all agents prefer. By continuity and convexity, we can take this allocation to be arbitrarily close to the equilibrium allocation. Local invertibility implies that there is some announcement of prices that each agent can make that will implement this dominating allocation. If they choose not to do so, it must be because the dominating allocation violates their budget constraint,

$$x_i' + p_i G' > w_i.$$

Summing over the agents and using the fact that $\sum_{k=1}^n p_k^* = 1$, we have

$$\sum_i x_i' + G' > \sum_i w_i,$$

which contradicts the assumption that the dominating allocation is feasible.

References

Admati, A. & Perry, M. (1991). Joint projects without commitment. *Review of Economic Studies*, *58*, 259–276.

Bergstrom, T. (1989). Love and spaghetti, the opportunity cost of virtue. *Journal of Economic Perspectives*, *3*(2), 165–173.

Bergstrom, T., Blume, L., & Varian, H. (1986). On the private provision of public goods. *Journal of Public Economics*, *29*, 25–49.

Buchanan, J. M. (1975). The Samaritan's dilemma. In Phelps, E. S. (Ed.), *Altruism, Morality, and Economic Theory.* The Sage Foundation.

Cornes, R. & Sandler, T. (1986). *The Theory of Externalities, Public Goods, and Club Goods*. Cambridge University Press, Cambridge, England.

Danziger, L. & Schnytzer, A. (1991). Implementing the Lindahl voluntary-exchange mechanism. *European Journal of Political Economy*, *7*, 55–64.

Forsyth, A. (1986). *A Natural History of Sex*. Charles Scribner's Sons, New York.

Groves, T. (1979). Efficient collective choice when compensation is possible. *Review of Economic Studies*, *46*, 227–241.

Groves, T. & Ledyard, J. (1977). Optimal allocations of public goods: A solution to the 'free rider problem'. *Econometrica*, *45*, 783–809.

Groves, T. & Ledyard, J. (1987). Incentive compatibility since 1972. In Groves, T., Radner, R., & Reiter, S. (Eds.), *Information, Incentives and Economic Mechanisms*. University of Minnesota Press, Minneapolis.

Guttman, J. (1978). Understanding collective action: Matching behavior. *American Economic Review*, *68*, 251–255.

Guttman, J. (1986). Matching behavior and collective action: Some experimental evidence. *Journal of Economic Behavior and Organization*, *7*, 171–198.

Guttman, J. (1987). A non-Cournot model of voluntary collective action. *Economica*, *54*, 1–19.

Hurwicz, L. (1979). Outcome functions yielding Walrasian and Lindahl allocations at Nash equilibrium points. *Review of Economic Studies*, *46*, 217–225.

Jackson, M., & Moulin, H. (1992). Implementing a public project and distributing its cost. *Journal of Economic Theory*, *57*, 125–140.

Ledyard, J. (1995). Public goods: A survey of experimental research. in *Handbook of Experimental Economics*, John Hagel and Alvin Roth (eds.), Princeton University Press.

Lindbeck, A. & Weibull, J. W. (1988). Altruism and time consistency: the economics of fait accompli. *Journal of Political Economy*, *96*(6), 1165–1182.

Moore, J. (1991). Implementation in environments with complete information. Tech. rep., Suntory-Toyota International Centre for Economics, London School of Economics.

Roberts, R. (1987). Financing public goods. *Journal of Political Economy*, *95*, 420–437.

Varian, H. (1982). Pensions and public policy. Tech. rep., University of Michigan.

Varian, H. (1989). A solution to the problem of externalities when agents are well-informed. Tech. rep., CREST Working Paper.

Walker, M. (1981). A simple incentive compatible scheme for attaining Lindahl allocations. *Econometrica*, *48*, 56–73.

Warr, P. (1983). The private provision of a public good is independent of the distribution of income. *Economic Letters*, *13*, 207–211.

Chapter 2

A SOLUTION TO THE PROBLEM OF EXTERNALITIES WHEN AGENTS ARE WELL-INFORMED

I describe a class of simple two-stage mechanisms that implement efficient allocations as subgame perfect equilibria for economic environments involving externalities. These mechanisms, known as compensation mechanisms, solve a wide variety of externalities problems including implementation of Lindahl allocations, regulation of monopoly, and efficient solutions to the Prisoners' Dilemma.

Consider an economic environment in which agents take actions that impose benefits or costs on other agents. The agents involved know the relevant technology and the tastes of all other agents. However the "regulator," who has the responsibility for determining the final allocation, does not have this information. How can the regulator design a mechanism that will implement an efficient allocation?

In this paper I describe a class of simple two-stage games whose subgame perfect equilibria implement efficient allocations in this sort of environment. In addition to implementing efficient outcomes, the mechanisms also achieve desirable distributional goals. In the case of public goods, the mechanisms implement Lindahl allocations; in the case of a negative externality the injured parties are compensated. Because payment of "compensation" is an important feature of the mechanisms I describe, I refer to the general class of mechanisms as *compensation mechanisms.* These mechanisms appear to work in a broad variety of economic environments and do

This work was supported by National Science Foundation grant SES-8800114, a Fulbright grant and the IRIS Research Program. I wish to thank Mark Bagnoli, Ted Bergstrom, Ken Binmore, Alan Kirman, Paul Milgrom, Diego Moreno, Arthur Robson, Steve Salant, Mark Walker, Michael Whinston and two anonymous referees for their comments and suggestions. I also wish to thank the European University Institute and the University of Siena for their hospitality during the period in which I conducted this research.

not involve substantial restrictions on tastes or technology. They are also quite simple to describe and analyze.

The fact that multi-stage games and subgame perfect equilibria may be useful in implementation problems was suggested by Crawford (1979) , Moulin (1979, 1981) and extensively analyzed by Moore and Repullo (1988). Moore and Repullo show that in economic environments, almost any choice rule can be implemented by subgame perfect equilibria. However, as Moore and Repullo point out, "…the mechanisms we construct …are far from simple; agents move simultaneously at each stage and their strategy sets are unconvincingly rich. We present such mechanisms to show what is possible, not what is realistic." (p. 1198) Moore and Repullo also show that in certain "economic environments" it is possible to use somewhat simpler mechanisms. However, the compensation mechanism appears to be much simpler than the examples Moore and Repullo examined. For a thorough review of the recent literature on implementation in complete information environments see Moore (1991).

It should be emphasized that the solution concept of subgame perfection requires the agents to be informed about the technology and tastes of the other agents. This is, of course, more restrictive than one would like. However, there is a broad set of cases for which this assumption may be plausible. For example, consider a group of agents who must design a constitution that describes a mechanism to make group decisions for problems that will arise in the future. At the time the mechanism is chosen, the agents may not know the relevant tastes and technologies, but they will know these things when the mechanism is actually used. In this circumstance, the compensation mechanism may be a useful mechanism. See Moore and Repullo (1988) and Maskin (1985) for further discussion of these issues.

I first describe a very simple example of the compensation mechanism in a two-agent externalities problem and discuss in an intuitive way why the method works. The following sections show how the method can be extended to work in more general environments.

2.1 A simple example of the compensation mechanism

Consider the following externality problem involving two agents. For simplicity we think of each agent as a profit-maximizing firm. Firm 1 produces output x so as to maximize profit

$$\pi_1 = rx - c(x),$$

where r is the competitive price of output and $c(x)$ is a differentiable, positive, increasing, and convex cost function.

Firm 1's choice of output imposes an externality on firm 2; in particular, firm 2's profits are

$$\pi_2 = -e(x),$$

where $e(x)$ is a differentiable, positive, increasing, and convex function of x. All of this information is known to both agents, but is not known by the regulator. In general, the level of output chosen by firm 1 will not be efficient, since firm 1 ignores the social cost its choice imposes on firm 2. There are three classic solutions to this problem of externalities.

One class of solutions, associated with Coase (1960), involves negotiation between the agents. Coase claims that if transactions costs are zero and property rights are well-defined, agents should be able to negotiate their way to an efficient outcome. But this is an incomplete solution to the problem of externalities since Coase does not describe a specific mechanism for negotiation. The compensation mechanism described below provides a structure for such negotiations, and therefore can be viewed as being complementary to the Coase approach.

A second class of solutions, associated with Arrow (1970), involves setting up a market for the externality. If a firm produces pollution that harms another firm then a competitive market for the right to pollute may allow for an efficient outcome. From the Coasian point of view, a competitive market is a particular institution that allows agents to "negotiate" their way to an efficient outcome. However, as Arrow points out, the market for allocating a particular externality may be very thin—in many cases of interest such markets involve only two participants.

However, a thin market does not necessarily mean a non-competitive market. There are both theoretical and empirical reasons to believe that certain kinds of market interaction can be competitive even though only a small number of agents are involved. For example, a Bertrand model of oligopoly yields a more-or-less competitive outcome with only two firms. The real-life implementation of Bertrand competition—competitive bidding—seems to work reasonably well even if there are only a small number of bidders. This suggests that markets for externalities with price-setting agents may be a useful model for negotiations among agents. This is a key insight behind the compensation mechanism.

A third class of solutions, associated with Pigou (1920), involves intervention by a regulator who imposes a Pigovian tax. The difficulty with this solution is that it requires the regulator to be able to compute the correct level of the Pigovian tax; in many cases the regulator may not have access to this information, so the Pigovian solution is also incomplete. The compensation mechanism solves this problem, since it gives the regulator a method to induce the participants to reveal the information necessary to construct the optimal Pigovian tax.

Returning to our example, we note that if the regulator had full information, internalizing the externality would be easy. One solution would be for the regulator to impose the costs of the externality on firm 1 by charging it a "tax" of $e(x)$ if it produces x units of output. Firm 1 would then solve the problem

$$\max_x \; rx - c(x) - e(x).$$

Let x^* be solution to this problem; then x^* satisfies the first-order condition

$$r - c'(x^*) - e'(x^*) = 0.$$

Because of our curvature assumptions on $e(x)$, the regulator could just as well set a "Pigovian tax," $p^* = e'(x^*)$ and let firm 1 solve the problem

$$\max_x \; rx - c(x) - p^* x.$$

However, we have assumed that the regulator doesn't know the externality cost function and therefore cannot determine the appropriate value of p^*. The regulator's problem is to design a mechanism that will induce the agents to reveal their information about the magnitude of the externality and achieve an efficient level of production. Here is a version of the compensation mechanism that solves the regulator's problem.

Announcement stage. Firm 1 and firm 2 simultaneously announce the magnitude of the appropriate Pigovian tax; denote the announcement of firm 1 by p_1 and the announcement of firm 2 by p_2.

Choice stage. The regulator makes sidepayments to the firms so that the two firms face profit maximization problems:

$$\Pi_1 = rx - c(x) - p_2 x - \alpha_1 (p_1 - p_2)^2$$
$$\Pi_2 = p_1 x - e(x).$$

The parameter $\alpha_1 > 0$ is of arbitrary magnitude.

In this mechanism firm 1 is forced to pay a tax based on the marginal social cost of the externality as reported by firm 2, and firm 2 receives compensation based on the marginal social cost of the externality as reported by firm 1. Firm 1 must also pay a penalty if it reports a different marginal social cost than firm 2 reports. Any penalty that is minimized when the reports are the same will work, but we have chosen a quadratic penalty for simplicity. Note in particular that the penalty can be arbitrarily small.

2.2 Analysis of the compensation mechanism

There are many Nash equilibria of this game; essentially any triple (p_1, p_2, x) such that $p_1 = p_2$ and x maximizes firm 1's objective function is a Nash equilibrium. However, if we use the stronger concept of subgame perfect equilibrium we get a much smaller set of equilibria. In fact, the *unique* subgame perfect equilibrium of this game has each agent reporting $p_1 = p_2 = p^*$ and firm 1 producing the efficient amount of output.

In order to verify this, we must work backwards through the game. We begin with the choice stage. Firm 1 maximizes its profits, given the Pigovian tax announced in stage 1, which implies that firm 1 will choose x to satisfy the first-order condition

$$r = c'(x) + p_2. \tag{2.1}$$

This determines the optimal choice, x, as a function of p_2, which we denote by $x(p_2)$. Note that $x'(p_2) < 0$—the higher the tax that firm 2 announces, the less firm 1 will want to produce.

We next examine the price-setting stage of the game. We first examine firm 1's choice problem. If firm 1 believes that firm 2 will announce p_2, then firm 1 will want to announce

$$p_1 = p_2 \tag{2.2}$$

This is clear since p_1 only influences firm 1's payoff through the penalty term, and the penalty is minimized when $p_1 = p_2$.

Consider now firm 2's pricing decision. Although firm 2's announcement has no *direct* effect on firm 2's profits, it does have an *indirect* effect through the influence of p_2 on firm 1's output choice in stage 2. Differentiating the profit function of firm 2 with respect to p_2, and setting it equal to zero we have

$$\Pi_2'(p_2) = [p_1 - e'(x)]x'(p_2) = 0. \tag{2.3}$$

Since $x'(p_2) < 0$ we must have $p_1 = e'(x)$.

Combining (2.1), (2.2) and (2.3) we have

$$r = c'(x) + e'(x),$$

which is the condition for social optimality. Hence, the unique subgame perfect equilibrium to this game involves firm 1 producing the socially optimal amount of output.[1]

[1] We have not considered the possibility of mixed strategies. However, since the stage 1 game is supermodular and has a unique pure strategy equilibrium, the results of Milgrom and Roberts (1991) can be applied to show that there are no mixed strategy equilibria.

2.3 Why the compensation mechanism works

The intuition behind the mechanism is not particularly difficult. Firm 2 effectively chooses x by setting the price firm 1 faces. If we are to have an equilibrium, we must have $p_1 = e'(x)$; otherwise firm 2 would want to change its announcement of p_2 in order to induce firm 1 to change x. Furthermore, firm 1 will always want to set $p_1 = p_2$ so as to minimize its penalty. The only configuration compatible with these conditions is the efficient outcome.

For example, suppose that firm 1 thinks that firm 2 will report a large price for the externality. Then, since firm 1 is penalized if it announces something different from firm 2, firm 1 will also want to announce a large price. If firm 1 announces a large price, firm 2 will be "overcompensated" for the externality—so it will want firm 1 to produce a large amount of output. But the only way firm 2 can give firm 1 an incentive to produce a large amount of output is by reporting a *small* price for the externality. This contradicts our original assumption that firm 1 thinks that firm 2 will report a large price for the externality. The only equilibrium for the mechanism is where firm 2 is just compensated (on the margin) for the cost that firm 1 imposes on it; at this point firm 2 doesn't want firm 1 to increase or decrease its level of production.

2.4 Extensions of the basic example

Balance

The compensation mechanism, in the form presented above, is balanced in equilibrium but not out of equilibrium. However, if there are at least three agents it is easy to choose transfers to balance the mechanism. As Moore and Repullo point out, we can simply distribute the surplus or deficit generated by each agent's choice among the other agents. Since this lump sum distribution is independent of agent i's choice, there are no resulting incentive effects.[2]

Let's see how this works in the simple example considered above. We now suppose that agent 1 imposes an externality on agents 2 and 3. Use the notation p_{ij}^k to represent the price announced by agent k that measures (in equilibrium) the marginal cost that agent j's choice imposes on agent i.

[2] This idea seems to have been first used by Groves-Ledyard (1987). Since then it has been used by a number of other authors.

The basic compensation mechanism for this problem has payments of the form

$$\begin{aligned}
\Pi_1 &= rx - c(x) - [p_{21}^2 + p_{31}^3]x - \|p_{21}^1 - p_{21}^2\| - \|p_{31}^1 - p_{31}^3\| \\
\Pi_2 &= p_{21}^1 x - e_2(x) \\
\Pi_3 &= p_{31}^1 x - e_3(x).
\end{aligned}$$

If we distribute payments so as to balance the budget out of equilibrium, the payoffs become

$$\begin{aligned}
\Pi_1 &= rx - c(x) - [p_{21}^2 + p_{31}^3]x - \|p_{21}^1 - p_{21}^2\| - \|p_{31}^1 - p_{31}^3\| \\
\Pi_2 &= p_{21}^1 x - e_2(x) + [p_{31}^3 - p_{31}^1]x + \|p_{31}^1 - p_{31}^3\| \\
\Pi_3 &= p_{31}^1 x - e_3(x) + [p_{21}^2 - p_{21}^1]x + \|p_{21}^1 - p_{21}^2\|.
\end{aligned} \tag{2.4}$$

Straightforward addition shows that this game is balanced. Using the same sort of arguments as before, it is possible to verify that the unique equilibrium of this mechanism is the efficient outcome. In fact, it is not necessary to have penalty terms when there are more than two agents. To see this, set the penalty terms in (2.4) equal to zero and differentiate the relevant objective functions with respect to the choice variables x, p_{21}^2, and p_{21}^3:

$$\begin{aligned}
r - c'(x) - [p_{21}^2 + p_{31}^3] &= 0 \\
[p_{21}^1 - e_2'(x) + p_{31}^3 - p_{31}^1]x'(p_{21}^2 + p_{21}^3) &= 0 \\
[p_{31}^1 - e_3'(x) + p_{21}^2 - p_{21}^1]x'(p_{21}^2 + p_{31}^3) &= 0.
\end{aligned}$$

Since the derivative of x must be non-zero due to strict convexity of the cost function, the terms in brackets must be zero. Adding the bracketed expressions in the last two equations together and substituting into the first equation shows that the equilibrium is efficient.

Yet a third way to balance the mechanism is to allow agent 2 to name the cost that 1 imposes on 3 and vice versa. This is a bit less natural in terms of the information requirements but it yields a very simple mechanism:

$$\begin{aligned}
\Pi_1 &= rx - c(x) - (p_{21}^3 + p_{31}^2)x \\
\Pi_2 &= p_{21}^3 x - e_2(x) \\
\Pi_3 &= p_{31}^2 x - e_3(x).
\end{aligned}$$

Differentiating with respect to each of the choice variables as above shows that the equilibrium of this mechanism is efficient, and it is obviously balanced. Each of these ways of balancing the compensation mechanism works in general as we shall see below.

Adjusting to equilibrium

There is a natural adjustment process for the compensation mechanism that will lead naive agents to the subgame perfect equilibrium. Suppose that two agents play the game repeatedly. In period $t+1$ agent 1 sets p_1 to be whatever price agent 2 announced last period, and agent 2 moves p_2 in a direction that increases its profits if agent 1 sets the same price as it did last period. In the choice stage, agent 1 chooses output to maximize profits, given the current prices. This gives us a simple discrete dynamical system:

$$\begin{aligned} p_1(t+1) &= p_2(t) \\ p_2(t+1) &= p_2(t) - \gamma[p_1(t) - e'(x(p_2(t))]. \end{aligned} \tag{2.5}$$

Here $\gamma > 0$ is a speed of adjustment parameter. The differential equation analog of this system is

$$\begin{aligned} \dot{p}_1 &= p_2 - p_1 \\ \dot{p}_2 &= -\gamma[p_1 - e'(x(p_2))]. \end{aligned} \tag{2.6}$$

It is easy to show that that (2.6) is locally stable; the difference equation version described in (2.5), will be locally stable if γ is small enough to avoid "overshooting." Note that if the agents use this adjustment procedure neither one needs to know anything about the other agent's technology. All that information is subsumed in the price messages that the agents send back and forth.[3]

Nonlinear taxes and compensation functions

The basic compensation mechanism described above uses linear pricing. Linear prices are fine in a convex environment, but if the environment is not convex, linear prices will not in general be able to support efficient allocations. However, this difficulty is no problem for a suitable generalization of the compensation mechanism.

Announcement stage. Firm 1 and firm 2 each announce the externality cost function for firm 2. Call these announcements $e_1(\cdot)$ and $e_2(\cdot)$.

[3] This is, of course, a very special adjustment process. However, Milgrom and Roberts (1991) show that for dominance solvable games every adjustment process consistent with adaptive and/or sophisticated learning converges to the dominance solvable equilibrium. Hence it may reasonably be expected that a wide class of adjustment mechanisms will work when the second-stage game is dominance solvable (as it is in this case.)

Choice stage. Firm 1 chooses x and each firm receives payoffs given by:

$$\Pi_1(x) = rx - c(x) - e_2(x) - \|e_1 - e_2\|$$
$$\Pi_2(x) = e_1(x) - e(x)$$

Here $\|e_1 - e_2\|$ signifies any norm in the appropriate function space. All that is required is that it is minimized when both agents report the same function.

To see that this works, simply note that in equilibrium, firm 1 will always want to report the same function as firm 2, so $e_1(x) \equiv e_2(x)$. Maximization of profit by firm 1 in the choice stage implies

$$rx^* - c(x^*) - e_2(x^*) \geq rx - c(x) - e_2(x) \quad \text{for all } x. \tag{2.7}$$

However, in the announcement stage, firm 2 can induce any level of x that it wants by appropriate choice of the function e_2. Hence, the equilibrium choice of x must also maximize firm 2's profits:

$$e_1(x^*) - e(x^*) \geq e_1(x) - e(x) \quad \text{for all } x. \tag{2.8}$$

Adding (2.7) and (2.8) together, and using the fact that $e_1(x) \equiv e_2(x)$ in equilibrium, we have

$$rx^* - c(x^*) - e(x^*) \geq rx - c(x) - e(x) \quad \text{for all } x, \tag{2.9}$$

which shows that x^* is the socially optimal amount.

This argument shows that all equilibria of the mechanism are efficient. However, in general there will be many equilibria of this game. To see this, observe that if e_1 and e_2 are equilibrium announcements, so are $e_1 + F$ and $e_2 + F$ for arbitrary values of F. In order to get uniqueness of equilibrium, it is necessary to restrict the class of allowable messages.[4]

One way to do this is to parameterize the cost function.[5] Suppose that the set of possible externality costs is $e(x, t)$, where t is a real-valued index of type. Suppose that the true type of firm 2 is t_0. In the announcement stage of the game, each firm simply announces the type of firm 2, and firm 1 pays a penalty if its announcement is different from that of firm 2. If t_1 is firm 1's announcement and t_2 is firm 2's announcement, the payoffs will be

$$\Pi_1(x) = rx - c(x) - e(x, t_2) - (t_1 - t_2)^2$$
$$\Pi_2(x) = e(x, t_1) - e(x, t_0)$$

[4] One could also refine the solution concept. In this example it may be reasonable for agent 1 to assume that agent 2 will announce the largest possible value of F consistent with agent 1's participation.

[5] In the convex case one can think of the efficiency prices as being a particularly convenient parameterization for the type space.

Differentiating with respect to x, t_1, and t_2, we have

$$r - c'(x) - \frac{\partial e(x, t_2)}{\partial x} = 0$$
$$t_1 - t_2 = 0$$
$$\left[\frac{\partial e(x, t_1)}{\partial x} - \frac{\partial e(x, t_0)}{\partial x}\right] x'(t_2) = 0.$$

Assuming that $x'(t_2) \neq 0$, it is easy to see that these equations imply

$$r = c'(x) + \frac{\partial e(x, t_0)}{\partial x},$$

which is the condition for social efficiency.

Of course, this argument requires sufficient regularity so that the various derivatives exist. If the environment is not suitably convex an argument can be constructed along the lines given above in inequalities (2.7)–(2.9). Note that in a nonconvex environment we need to assume that firm 2 can induce firm 1 to choose any desired level of x by choosing an appropriate value of t_2; this is simply a "global" version of the assumption that $x'(t_2) \neq 0$.

2.5 A general externalities problem

The externalities problem we have examined up until now is rather special. Only one agent makes a choice, and both agents have quasilinear objective functions so there are no income effects. In this section we consider a more general externality problem. For simplicity, we continue to examine a two-agent problem but the argument is easily generalized to n agents.

In the general model there are two choices, x_1 and x_2, and one transferable good, y. Agent i makes choice x_i, and has a quasiconcave utility function $u_i(x_1, x_2, y_i)$. Initially, agent i has w_i units of the transferable good, which can be thought of as money.

Efficient choices

In the absence of any transfers between the agents, agent i will choose x_i to maximize his own utility. The first-order condition characterizing these choices can be written as

$$\frac{\partial u_1/\partial x_1}{\partial u_1/\partial y_1} = 0$$
$$\frac{\partial u_2/\partial x_2}{\partial u_2/\partial y_2} = 0.$$

By contrast, an efficient allocation of choices must satisfy the first-order conditions

$$\begin{aligned}\frac{\partial u_1/\partial x_1}{\partial u_1/\partial y_1} + \frac{\partial u_2/\partial x_1}{\partial u_2/\partial y_2} &= 0 \\ \frac{\partial u_2/\partial x_2}{\partial u_2/\partial y_2} + \frac{\partial u_1/\partial x_2}{\partial u_1/\partial y_1} &= 0.\end{aligned} \tag{2.10}$$

These conditions simply require that the sum of the marginal-willingnesses-to-pay for activity i should be zero.

Define

$$p_{ij} = \frac{\partial u_i/\partial x_j}{\partial u_i/\partial y_i} \quad \text{for } i \neq j$$

Then we can write the efficiency conditions (2.10) as

$$\begin{aligned}\frac{\partial u_1/\partial x_1}{\partial u_1/\partial y_1} + p_{21} &= 0 \\ \frac{\partial u_2/\partial x_2}{\partial u_2/\partial y_2} + p_{12} &= 0\end{aligned}$$

This form suggests that the efficient allocation can be achieved if each agent faced the correct "price" for his choice. The problem is how to determine the correct price. Here is a description of the general compensation mechanism that solves this problem.

Announcement stage. Agent 1 announces p_{12}^1 and p_{21}^1, and agent 2 announces p_{12}^2 and p_{21}^2.

Choice stage. Each agent chooses x_i and y_i so as to maximize utility subject to a budget constraint:

$$\begin{aligned}&\max_{x_1, y_1} u_1(x_1, x_2, y_1) \\ &\text{such that } p_{21}^2 x_1 + y_1 = w_1 + p_{12}^2 x_2 - \|p_{21}^1 - p_{21}^2\|,\end{aligned}$$

and

$$\begin{aligned}&\max_{x_2, y_2} u_2(x_1, x_2, y_2) \\ &\text{such that } p_{12}^1 x_2 + y_2 = w_2 + p_{21}^1 x_1 - \|p_{12}^2 - p_{12}^1\|.\end{aligned}$$

We show below that the subgame perfect equilibria of this game are precisely the efficient allocations that satisfy the budget constraints. However, before providing that proof, let us make a few observations.

First, each agent i is facing a price, p_{ji}^j, for his own choice x_i. He is also receiving compensation $p_{ij}^j x_j$ for the choice that the other agent makes.

Both prices p_{ji}^j and p_{ij}^j are set by the *other* agent. Each agent i also pays a penalty based on how different his announced price, p_{ji}^i, is from the price that the other agent j announced for i's choice, p_{ji}^j.

As one might suspect, in equilibrium we must have $p_{ji}^i = p_{ji}^j$. This means that no penalties will be paid and that the payment made by agent i for his action will just be equal to the compensation paid to agent j. Hence, in equilibrium, the aggregate budget constraint will balance.

Let us now show that the equilibrium of this game must be efficient. We provide two proofs. The first proof simply involves writing down the first-order conditions for the utility maximization problems. There are three choice variables for each of the two agents, x_i, p_{ij}^i, and p_{ji}^i, so we have six first order conditions. Choosing the quadratic norm for computational simplicity, the first-order conditions are:

$$\frac{\partial \mathrm{u}_1}{\partial \mathrm{x}_1} - \frac{\partial \mathrm{u}_1}{\partial \mathrm{y}_1}\mathrm{p}_{21}^2 = 0 \tag{2.11}$$

$$\left(\frac{\partial u_1}{\partial x_1} - \frac{\partial u_1}{\partial y_1}p_{21}^2\right)\frac{\partial x_1}{\partial p_{12}^1} + \left(\frac{\partial u_1}{\partial x_2} + \frac{\partial u_1}{\partial y_1}p_{12}^2\right)\frac{\partial x_2}{\partial p_{12}^1} = 0 \tag{2.12}$$

$$\left(\frac{\partial u_1}{\partial x_1} - \frac{\partial u_1}{\partial y_1}p_{21}^2\right)\frac{\partial x_1}{\partial p_{21}^1} + \left(\frac{\partial u_1}{\partial x_2} + \frac{\partial u_1}{\partial y_1}p_{12}^2\right)\frac{\partial x_2}{\partial p_{21}^1}$$
$$-2\frac{\partial u_1}{\partial y}(p_{21}^1 - p_{21}^2) = 0 \tag{2.13}$$

$$\frac{\partial u_2}{\partial x_2} - \frac{\partial u_2}{\partial y_2}p_{12}^1 = 0 \tag{2.14}$$

$$\left(\frac{\partial u_2}{\partial x_2} - \frac{\partial u_2}{\partial y_2}p_{12}^1\right)\frac{\partial x_2}{\partial p_{21}^2} + \left(\frac{\partial u_2}{\partial x_1} + \frac{\partial u_2}{\partial y_2}p_{21}^1\right)\frac{\partial x_1}{\partial p_{21}^2} = 0 \tag{2.15}$$

$$\left(\frac{\partial u_2}{\partial x_2} - \frac{\partial u_2}{\partial y_2}p_{12}^1\right)\frac{\partial x_2}{\partial p_{12}^2} + \left(\frac{\partial u_2}{\partial x_1} + \frac{\partial u_2}{\partial y_2}p_{21}^1\right)\frac{\partial x_1}{\partial p_{12}^2}$$
$$-2\frac{\partial u_2}{\partial y_2}(p_{12}^2 - p_{12}^1) = 0 \tag{2.16}$$

Here we have assumed that the equilibrium choices in the second stage are differentiable functions of the price announcements made in the first stage. As we shall see in the next section, this is not necessary for the argument, but it does help to see why the method works. Note that when agent 1 chooses p_{12}^1, for example, he recognizes that both his own choice, x_1, and the other agent's choice, x_2, may respond to changes in p_{12}^1.

We must assume that $\partial x_2/\partial p_{12}^1$ and $\partial x_1/\partial p_{21}^2$ are not zero. Now simply observe that this assumption and (2.11)–(2.13) together imply that $p_{21}^1 = p_{21}^2$. Similarly, (2.14)—(2.16) imply that $p_{12}^1 = p_{12}^2$. Finally, combine (2.11) and (2.15) to get:

$$\frac{\partial u_1/\partial x_1}{\partial u_1/\partial y_1} + \frac{\partial u_2/\partial x_1}{\partial u_2/\partial y_2} = 0,$$

and combine (2.12) and (2.14) to get:

$$\frac{\partial u_2/\partial x_2}{\partial u_2/\partial x_2} + \frac{\partial u_1/\partial x_2}{\partial u_1/\partial y_1} = 0.$$

These are precisely the first-order conditions given in (2.10). Therefore, the subgame perfect equilibrium is efficient.

Note that the equilibrium is a *particular* efficient allocation, namely one that satisfies the "natural" budget constraint involving the efficiency prices. In general, such allocations will be a small subset of all efficient allocations. By analogy with the public goods literature, we call these allocations *generalized Lindahl allocations.* We show below that when the externalities problem is a public goods problem, the prices in the compensation mechanism are Lindahl prices.

A more general proof

The above proof shows clearly why an equilibrium of the compensation mechanism must be an efficient allocation. However, being a calculus proof, it doesn't deal very well with corner solutions, additional constraints, non-differentiabilities, etc. Here is another argument that handles these difficulties easily.

We need one assumption for our proof, an invertibility assumption that says that each agent can set a price for the other agent that will induce the other agent to make any desired choice. That is, if agent 1 would like agent 2 to make some choice there is some price that agent 1 can set that will induce agent 2 to make this choice. This is analogous to the assumption that $\partial x_2/\partial p_{12}^1 \neq 0$ in our previous proof. As in the differentiable case, the demand functions only need to be locally invertible if the environment is suitably convex.

Local Invertibility Assumption. Let $x = (x_1, x_2)$ be the outcome of some set of price announcements. Let $\hat{x}_i$ be a choice close to x_i that agent j prefers to x_i. Then there is some $\hat{p}_{ji}^j$ that agent j can announce that will make $\hat{x}_i$ an optimal choice for agent i.

Local invertibility says that agent i can manipulate agent j's choices through agent i's price announcements. This is a very weak assumption. If the agents' demands are differentiable functions of price with nonzero derivatives then the Inverse Function Theorem implies local invertibility.

Theorem. *Let preferences be convex and continuous. Then every subgame perfect equilibrium of the compensation mechanism is Pareto efficient.*

Proof. Let (x, y, p) be a subgame perfect equilibrium of the compensation mechanism. First we show that in equilibrium $p^1_{21} = p^2_{21}$. To see this, consider the agents' budget constraints:

$$p^2_{21}x_1 + y_1 = w_1 + p^2_{12}x_2 - \|p^1_{21} - p^2_{21}\|$$
$$p^1_{12}x_2 + y_2 = w_2 + p^1_{21}x_1 - \|p^2_{12} - p^1_{12}\|.$$

Note that agent 1 can influence agent 2's choice of x_2 through both the "income term," $p^1_{21}x_1$, and the "price term," p^1_{12}. However, by the Local Invertibility Assumption, any choice of x_2 that can be achieved through the income term can also be achieved by an appropriate choice of the price term, p^1_{12}.

Suppose that there were an equilibrium in which $p^1_{21} \neq p^2_{21}$. Let agent 1 set $p^1_{21} = p^2_{21}$ and adjust p^1_{12} so as to induce the original equilibrium value of x_2. This must reduce agent 1's penalty, and thereby increase agent 1's utility. This contradicts the assumption that we had an equilibrium. It follows that an equilibrium must exhibit zero penalty terms for all agents.

Suppose now that (x', y') is a feasible allocation that Pareto dominates the equilibrium allocation. We will show that the existence of such an allocation leads to a contradiction. By convexity and continuity of preferences, we can assume that (x', y') is arbitrarily close to the equilibrium allocation. According to the Local Invertibility Assumption agent 1 can induce agent 2 to choose x'_2 simply by choosing an appropriate level of p^1_{12}; furthermore, agent 1 can directly choose (x'_1, y'_1). If agent 1 decides *not* to choose this preferred allocation, it must be because it lies outside his budget set. The same argument applies to agent 2, and this gives us the inequalities

$$p^2_{21}x'_1 + y'_1 > w_1 + p^2_{12}x'_2$$
$$p^1_{12}x'_2 + y'_2 > w_2 + p^1_{21}x'_1.$$

Summing these inequalities and using the fact that $p^i_{ji} = p^j_{ji}$, we have

$$y'_1 + y'_2 > w_1 + w_2,$$

which shows that the Pareto dominating allocation must be infeasible. ∎

Note that the logic of this proof is quite general. In particular, the taxation and compensation functions do not need to be linear functions. All that is necessary is that each agent can manipulate the other agent's choice without incurring any cost himself. If the economic environment is convex, we only need local invertibility; if the economic environment is non-convex, we may need global invertibility.

It can also be shown that if the environment is convex, then any Pareto efficient allocation is an equilibrium of this game for a suitable choice of initial endowments. The proof is a simple variation on the second welfare theorem and is omitted for sake of brevity.

2.6 Balancing the mechanism

We have seen that in the simple example discussed earlier that the compensation mechanism can be balanced by distributing the budget surplus generated by each agent among the other agents. The same procedure works in general; here we examine the simple case of quasilinear utility.

The appropriate payoff to agent i is

$$u_i(x) + \sum_{j=1}^{n} B_{ij} - T_i,$$

where

$$B_{ij} = p_{ij}^j x_j - p_{ji}^j x_i - \|p_{ji}^i - p_{ji}^j\|$$

$$T_i = \frac{1}{n-2} \sum_{k \neq i} \sum_{j \neq i} B_{kj} = \frac{1}{n-2} \left[\sum_{k=1}^{n} \sum_{j=1}^{n} B_{kj} - \sum_{j=1}^{n} B_{ij} - \sum_{j=1}^{n} B_{ji} \right].$$

It is obvious that $B_{ii} = 0$ and it is not hard to show that

$$\sum_{i=1}^{n} \left[\sum_{j=1}^{n} B_{ij} - T_i \right] = 0. \tag{2.17}$$

Note that B_{ij} depends on the vector of prices and the vector of choices. It is important to observe that the only price term that agent i determines is the price in the penalty term, p_{ji}^i; all other prices are independent of i's choices. To emphasize the fact that the payment depends on x, we write $B_{ij}(x)$ in the following paragraphs.

By local invertibility, each agent can induce any desired allocation in the choice stage by choosing the appropriate prices in the announcement stage. Hence a subgame perfect equilibrium allocation must satisfy

$$u_i(x^*) + \sum_{j=1}^{n} B_{ij}(x^*) - T_i(x^*) \geq u_i(x) + \sum_{j=1}^{n} B_{ij}(x) - T_i(x),$$

for all x. Summing over the agents and using equation (2.17), we see that

$$\sum_{i=1}^{n} u_i(x^*) \geq \sum_{i=1}^{n} u_i(x),$$

for all x which shows that the subgame perfect equilibrium is Pareto efficient.

Note that this argument does not use the linear structure of the $B_{ij}(x)$ terms; indeed, the only feature used is that agent j can report a $B_{ij}(x)$ term that will induce agent i to make the choice x_i that agent j wants him to make. In a convex environment linear prices will generally have this property, but in other environments other sorts of pricing functions may be necessary.

Note further that this argument for efficiency does not use the penalty terms; if all the penalty terms are set equal to zero, the proof of efficiency still goes through. However, the penalty terms will in general be necessary if we want the equilibrium allocation to be a Lindahl allocation. Why? In order to be a generalized Lindahl allocation each agent must satisfy his budget constraint when each choice is priced at its supporting efficiency price. For this to be the case, the $T_i(x)$ term must be zero in equilibrium. If the penalty terms are present, each agent will have an incentive to set $p^i_{ji} = p^j_{ji}$ which will ensure that this will occur.

2.7 A different information structure

The compensation mechanism described above is appropriate for a "bilateral" information structure: if agent i imposes costs on agent j, both i and j know the magnitude of these costs. Another structure that one might imagine is that there is some third party, k, who knows the magnitude of these costs. In this case, we can use a slightly different type of compensation mechanism to achieve efficient outcomes.

Consider the following example with three agents. Agent i chooses x_i, holds "money" y_i, and has a quasilinear utility function $u_i(x_1, x_2, x_3) + y_i$. The prices that support an efficient allocation will have the form $p_{ij} = \partial u_i(x)/\partial x_j$. Let p^k_{ij} denote the report of person k about the appropriate magnitude of the price p_{ij}, and let $x = (x_1, x_2, x_3)$ be the vector of choices.

In this variant of the compensation mechanism the payoffs to the agents will be:

$$\begin{aligned} u_1(x) - (p^2_{31} + p^3_{21})x_1 + p^3_{12}x_2 + p^2_{13}x_3 \\ u_2(x) - (p^1_{32} + p^3_{12})x_2 + p^3_{21}x_1 + p^1_{23}x_3 \\ u_3(x) - (p^1_{23} + p^2_{13})x_3 + p^2_{31}x_1 + p^1_{32}x_2. \end{aligned} \tag{2.18}$$

Note the payoffs are balanced, even out of equilibrium. No sidepayments or penalties are necessary in this case.

One way to prove that the subgame perfect equilibrium is efficient is to differentiate the payoffs with respect to each of the choice variables. However, we can also apply the logic of the previous section. Simply replace the definitions used there with

$$\begin{aligned} B^k_{ij}(x) &= p^k_{ij}x_j - p^k_{ji}x_i \\ T_i(x) &= 0, \end{aligned}$$

where k takes on all possible values $1, \ldots, n$, but $k \neq i, j$. Note that when $n = 3$ we have the payoffs given in (2.18). These definitions imply that

$$\sum_{i=1}^{n} \sum_{j=1}^{n} B_{ij}^{k}(x) = 0,$$

and this is all that is required for the proof given in the previous section to work. The resulting allocation is automatically Lindahl.

2.8 Examples of the compensation mechanism

We have described the general form of the compensation mechanism; here we illustrate how it works in some specific cases.

Pure public goods

The special case of a pure public good is of some interest, since it is a well known and much studied example of a particular type of externality. Let x_1 and x_2 be two agents' monetary contributions to a public good. Let y_i be agent i's private consumption. In the absence of any transfer mechanism, agent 1's maximization problem in a public goods contribution game takes the form[6]

$$\begin{aligned} \max_{x_1, y_1} \; & u_1(x_1 + x_2, y_1) \\ \text{such that } & x_1 + y_1 = w_1 \\ & x_1 \geq 0. \end{aligned}$$

Since there is now a positive externality between the agents, it is natural to think of the agents as subsidizing each other rather than taxing each other. Applying the subsidy payments appropriate for the compensation mechanism, the budget constraint facing agent 1 becomes:

$$(1 - p_{21}^2)x_1 + y_1 = w_1 - p_{12}^2 x_2 - \|p_{21}^1 - p_{21}^2\|.$$

Here agent 1's contributions are subsidized at a rate p_{21}^2 which is chosen by agent 2; this subsidy is recovered by a tax on agent 2. Agent 1 also sets the rate at which 2's contributions should be subsidized, and in equilibrium he ends up paying $p_{12}^2 x_2 = p_{12}^1 x_2$ to cover this subsidy. In the compensation mechanism the taxes and subsidies that each agent faces are chosen by

[6] The nonnegativity constraint is natural in a model of voluntary contributions: one may choose to contribute a positive amount to a public good, but one is typically not able to make a *negative* contribution. The equilibrium of this contribution game has been studied extensively by Bergstrom, Blume and Varian (1986).

the other agent(s). See Danziger (1991) and Varian (1994) for similar mechanisms where the agents set some of the subsidy rates for themselves. Guttman (1978) describes a related mechanism where agents choose the rate at which they will match other agents' contributions to a public good. Guttman's mechanism is of some interest since matching contributions are a commonly used method to encourage contributions to a public good.

Since public goods are just a special kind of externality, the proof of efficiency given earlier still applies. Note that the non-calculus proof is the appropriate version here, due to the presence of the nonnegativity condition. However, given the special form of the public goods externality, we can say a bit more about the equilibrium prices. Suppose that we have an interior solution to the public goods game so that x_1 and x_2 are both positive. Since x_1 and x_2 are perfect substitutes in consumption, they must have the same price in equilibrium.

By inspection of the budget constraint it follows that $(1 - p_{21}^2) = p_{12}^2$. Therefore the budget constraint facing agent 1 can, in equilibrium, be written as

$$p_{12}^2[x_1 + x_2] + y_1 = w_1.$$

It follows that an equilibrium value of p_{12}^2 is simply the Lindahl price of the public good for agent 1, and the equilibrium allocation is simply a Lindahl allocation. Hence, the compensation mechanism gives a way to decentralize Lindahl allocations by giving each agent the incentive to reveal the appropriate Lindahl prices.

Pure private goods

Agent 1 is a consumer who consumes an x-good and a y-good and has a quasilinear utility function $u(x) + y_1$. Agent 2 is a monopolist that can produce the x good at cost $c(x)$; it's objective function is $y_2 - c(x)$. How can we induce the monopolist to produce the socially optimal output?

If we are only interested in efficiency, this is not terribly difficult: simply have one of the agents dictate a production level and a transfer. In this full information environment we will get an efficient amount of x regardless of which agent chooses it; only the transfer will be different. However, if we want to get a *particular* efficient allocation—say the competitive outcome—it is not so obvious how to proceed. However, the compensation mechanism solves the problem quite readily.

Announcement stage. The consumer announces how much he values the good, p_1, and the producer announces how much the consumer values the good, p_2.

Choice stage. The producer chooses x and the payoffs are

$$\begin{aligned}\Pi_1 &= u(x) - p_2 x \\ \Pi_2 &= p_1 x - c(x) - \|p_2 - p_1\|\end{aligned}$$

Note that this problem is very similar to the simple externalities problem used to motivate the compensation mechanism, illustrating the Coasian point that externalities are just a special case of private goods.[7] Applying the standard argument shows that in equilibrium

$$p_1 = p_2 = u'(x) = c'(x),$$

which are the conditions that characterize the competitive allocation.

Regulation of duopoly

There are now three agents: the consumer (indexed by 0) and two firms. Firm 1 produces x_1 at cost $c_1(x_1)$, firm 2 produces x_2 at cost $c_2(x_2)$ and the consumer has utility function $u(x_1, x_2) + y_0$. The standard compensation mechanism involves payoffs of the form

$$\begin{aligned}\Pi_0 &= u(x_1, x_2) - p_{01}^1 x_1 - p_{02}^2 x_2 \\ \Pi_1 &= p_{01}^0 x_1 - c_1(x_1) - \|p_{01}^1 - p_{01}^0\| \\ \Pi_2 &= p_{02}^0 x_2 - c_2(x_2) - \|p_{02}^2 - p_{02}^0\|\end{aligned}$$

Here the consumer is setting the prices that the firms face, and the firms are setting the prices that the consumer faces.

However, in the case of duopoly it is natural to think that the firms may know more about each other's technology than the consumer knows. Hence it makes sense for each firm to report the price that the *other* firm should face. This yields payoffs of the form

$$\begin{aligned}\Pi_0 &= u(x_1, x_2) - p_{01}^2 x_1 - p_{02}^1 x_2 \\ \Pi_1 &= p_{01}^2 x_1 - c_1(x_1) \\ \Pi_2 &= p_{02}^1 x_2 - c_2(x_2).\end{aligned}$$

Note that the consumer chooses both x_1 and x_2 and that each firm sets the price for the *other* firms' product.

The arguments given earlier show that the competitive allocation is the unique equilibrium of this game. But it is useful to think about how it

[7] Or is it the other way around?

works here. Consider the classic Bertrand case where the two goods are perfect substitutes. Suppose that each firm has announced the competitive price. Why wouldn't firm 1 want to raise the price of firm 2's product, creating more demand for its own output? If firm 1 raised the price facing firm 2, then the consumer would demand more output from firm 1, which it would be *forced* to supply. But since the price that firm 1 faces equals its marginal cost this would reduce firm 1's profit.

Prisoners' dilemma

Consider the following asymmetric prisoners' dilemma:

		Column: Cooperate	Column: Defect
Row	Cooperate	$5, 5$	$2, 6$
	Defect	$7, 1$	$3, 3$

How can we induce the Pareto efficient outcome? Let $x_i = 1$ if agent i cooperates, and $x_i = 0$ if agent i defects and let $u_i(x_1, x_2)$ be the payoff to agent i taken from the above game matrix.

Announcement stage. Agent 1 names p_{12}^1, how much agent 1 should be paid if he cooperates, and p_{21}^1, how much agent 2 should be paid if he cooperates. Similarly agent 2 names p_{21}^2 and p_{12}^2.

Choice stage. Each agent chooses whether to cooperate or defect. The agents receive payoffs

$$\Pi_1 = u_1(x_1, x_2) + p_{21}^2 x_1 - p_{12}^2 x_2 - \|p_{21}^1 - p_{21}^2\|$$
$$\Pi_2 = u_2(x_1, x_2) + p_{12}^1 x_2 - p_{21}^1 x_1 - \|p_{12}^2 - p_{12}^1\|$$

Note the sign change: since there is now a positive externality between the two agents, it is natural to subsidize good behavior rather than penalize bad behavior.[8] Using the by now standard argument, it can be shown that the it is a subgame perfect equilibrium for both players to cooperate. The supporting prices satisfy the conditions

$$4 \geq p_{21}^1 = p_{21}^2 \geq 2$$
$$3 \geq p_{12}^1 = p_{12}^2 \geq 1.$$

(If the inequalities are strict, the cooperative equilibrium will be unique.) As usual the compensation mechanism produces an efficient outcome in this game—or in any game, for that matter.

[8] We could also formulate this problem so as to have each agent announce how the other agent should be fined if he defects.

2.9 Related literature

There is a vast literature on mechanism design that is concerned with how to implement various social choice functions. Much of this literature is concerned with whether a particular social choice function can be implemented by a decentralized game. Our concern is not so much with the *existence* of a mechanism, but rather finding a suitably simple mechanism. Most of the attempts to find "simple" solutions to externalities problems have been concerned with the case of public goods, so we provide a very brief review of that literature insofar as it relates to the work described here. Moore provides a thorough review of the recent literature.

The well-known demand revealing mechanism of Clarke (1971) and Groves (1976) implements the efficient amount of a public good via a dominant strategy equilibrium. However, this mechanism only works with quasilinear utility, and is not balanced, even in equilibrium. Furthermore, it does not in general yield a Pareto efficient outcome.

Groves and Ledyard (1977) describe a quadratic mechanism that yields efficient Nash equilibria for the public goods problem, but the equilibrium allocations are not Lindahl allocations. Hurwicz (1979) and Walker (1981) also describe mechanisms that implement Lindahl allocations. In the Hurwicz mechanism, each agent proposes an amount of the public good and a Lindahl price; agents pay a quadratic penalty if they announce different levels of the public good. Walker's (1981) mechanism avoids such penalty terms. Groves (1979) and Groves-Ledyard (1987) provide a nice survey of these results.

Turning to the more recent literature on simple mechanisms, Bagnoli-Lipman (1989) and Jackson-Moulin (1992) examine the special case of a discrete public good with quasilinear utility. The Bagnoli-Lipman mechanism is very simple: each agent offers a voluntary contribution. If the sum of the contributions covers the cost of the public good, it is produced, otherwise the contributions are returned. This mechanism implements the core of the public goods game in undominated perfect equilibria.

The Jackson-Moulin mechanism implements an efficient allocation using undominated Nash equilibria. They also describe a variation using subgame perfect equilibrium. Their mechanism is reasonably simple and works with a broad family of cost-sharing rules. However, it appears that both the Bagnoli-Lipman and Jackson-Moulin mechanisms work only in the special case of indivisible public goods and quasilinear utility.

Varian (1989a) describes some mechanisms for the public goods problem that are closely related to the compensation mechanism. In the case of two agents with quasilinear utility there is a very simple mechanism that achieves a Lindahl allocation: in the first stage each agent offers to subsidize the contributions of the other agent. In the second stage, each agent makes a voluntary contribution and collects the promised subsidies from the other agent.[9] Varian describes some other variations on the compensation mechanism for public goods problems involving many agents and general utility functions.

2.10 Summary

The compensation mechanism provides a simple mechanism for internalizing externalities in economic environments. Transfer payments can be chosen so that the compensation mechanism is balanced, and penalty payments, when they are used, can be chosen to be arbitrarily small. The main problem with the mechanism is that it requires complete information by the agents. In many cases a simple dynamic adjustment model will converge to the subgame perfect equilibrium.

References

Arrow, K. (1970). The organization of economic activity: Issues pertinent to the choice of market versus non-market allocation. In Haveman, R. H., & Margolis, J. (Eds.), *Public Expenditures and Policy Analysis*, pp. 59–73. Markham.

Bagnoli, M., & Lipman, B. (1989). Provision of public goods: fully implementing the core through private contributions. *Review of Economic Studies*, *56*, 583–602.

Bergstrom, T., Blume, L., & Varian, H. (1986). On the private provision of public goods. *Journal of Public Economics*, *29*, 25–49.

Clarke, E. (1971). Multipart pricing of public goods. *Public Choice*, *11*, 17–33.

Coase, R. H. (1960). The problem of social cost. *Journal of Law and Economics*, *3*, 1–44.

Crawford, V. (1979). A procedure for generating pareto efficient egalitarian-equivalent allocations. *Econometrica*, *47*, 49–60.

Danziger, L., & Schnytzer, A. (1991). Implementing the Lindahl voluntary-exchange mechanism. *European Journal of Political Economy*, *7*, 55–64.

[9] This mechanism is related to the mechanism of Guttman, which involves offering to match contributions. See also Danziger and Schnytzer for a related model.

Groves, T. (1979). Efficient collective choice when compensation is possible. *Review of Economic Studies*, *46*, 227–241.

Groves, T., & Ledyard, J. (1977). Optimal allocations of public goods: A solution to the 'free rider problem'. *Econometrica*, *45*, 783–809.

Groves, T., & Ledyard, J. (1987). Incentive compatibility since 1972. In Groves, T., Radner, R., & Reiter, S. (Eds.), *Information, Incentives and Economic Mechanisms.* University of Minnesota Press, Minneapolis.

Groves, T. (1976). Information, incentives, and the internalization of production externalities. In Lin, S. (Ed.), *Theory and Measurement of Economic Externalities.* Academic Press, New York.

Guttman, J. (1978). Understanding collective action: Matching behavior. *American Economic Review*, *68*, 251–255.

Hurwicz, L. (1979). Outcome functions yielding Walrasian and Lindahl allocations at Nash equilibrium points. *Review of Economic Studies*, *46*, 217–225.

Jackson, M., & Moulin, H. (1992). Implementing a public project and distributing its cost. *Journal of Economic Theory*, *57*, 125–140.

Maskin, E. (1985). The theory of implementation in Nash equilibrium. In Hurwicz, L., Schmeidler, D., & Sonnenschein, H. (Eds.), *Social Goals and Social Organization*, pp. 173–204. Cambridge Univeristy Press, Cambridge.

Milgrom, P., & Roberts, J. (1991). Adaptive and sophisticated learning in repeated normal form games. *Games and Economic Behavior*, 82–100.

Moore, J. (1992). Implementation, contracts, and renegotiation in environments with complete information. In Laffont, J. J. (Ed.), *Advances in Economic Theory: 6th World Congress*, pp. 182–281. Cambridge University Press, Cambridge, England.

Moore, J., & Repullo, R. (1988). Subgame perfect implementation. *Econometrica*, *56*, 1191–1220.

Moulin, H. (1979). Dominance solvable voting schemes. *Econometrica*, *47*, 1337–1351.

Moulin, H. (1981). Implementing just and efficient decision making. *Journal of Public Economics*, *16*, 193–213.

Pigou, A. C. (1920). *The Economics of Welfare.* Macmillan, London.

Varian, H. (1994). Sequential contributions to public goods. *Journal of Public Economics*, *53*(2), 165–186.

Walker, M. (1981). A simple incentive compatible scheme for attaining Lindahl allocations. *Econometrica*, *48*, 56–73.

Chapter 3

ECONOMIC INCENTIVES IN SOFTWARE DESIGN

I examine the incentives for software providers to design appropriate user interfaces. There are two sorts of costs involved when one uses software: the fixed cost of learning to use a piece of software and the variable cost of operating the software. I show that a monopoly provider of software generally invests the right amount of resources in making the software easy to learn, but too little in making it easy to operate. In some extreme cases a monopolist may even make the software too easy to learn.

The market for computer software is large and rapidly growing. Despite this, there has been little theoretical investigation of the unique economic features of the software market. In this paper I investigate an important aspect of software economics: the extent to which the providers of software have the right incentives to design an appropriate interface for their software.

3.1 User costs

An important feature of software is that there are large costs to the consumer of using it. First, one must learn how to use a particular software package. Even if one only wants to use the package occasionally, one has to read the documentation, practice a bit, and invest time and energy in learning the basics of how to use the package. This cost of learning the software is a *fixed cost* to the user: it is more-or-less independent of the amount of use that the software gets.

This should be contrasted with the *variable cost* of operating a software package. These are costs that are incurred every time one uses the software. The most obvious of these costs are time costs, such as a delay in loading or saving a file. If it takes 10 seconds to start the package every time you use it, this is 10 seconds of lost time each time the package is used. If one

This work was supported by the National Science Foundation. I also wish to thank the Dipartimento d'economia politica at the Universitá di Siena for their hospitality during the period of this research. I would like to thank Jim Adams, Paul Courant and Judy Olson for their comments on an earlier draft of this paper.

has to wade through an elaborate menu structure to perform a simple task, then this is a cost that must be borne every time the task is undertaken.

People who use a particular software package every day incur a large amount of these variable costs, while people who use this software rarely incur little variable costs. However, everyone incurs roughly the same fixed cost of learning the program.

Reviews of software often talk about "ease of use." The above distinction suggests that there are two dimensions to ease of use: ease of learning and ease of operation. Software that is easy to learn has a lot of menus and elaborate help screens. It provides user prompts and error messages. The documentation is easy to read. Software that is easy to operate generally has fewer menus, replacing them with command key combinations. This means that a given command can be executed quite quickly—once the user has made the investment in learning the appropriate keystrokes. In this sense, the command driven interface is easy to operate, although the menu interface is easy to learn.[1]

Another aspect of ease-of-operation is *performance:* how quickly and how well the software does the job it is supposed to do. When we turn to modeling consumer choice of software we will be interested in the *net performance* of the software: the difference between the benefits from the task the software does and the costs of making the software perform that task.

These two aspects of ease-of-use—being easy-to-operate versus being easy-to-learn—are not mutually exclusive. A well-designed software package can satisfy both goals. However the software designer still has to decide how much effort to put into improving each aspect of the user interface. If the software is supposed to be delivered in one week, is that week better spent improving the speed of some calculation or fine-tuning the menu structure? The answer presumably depends on how improvements in these two dimensions affect profits . . . which is where the economic analysis comes in.

In terms of our previous discussion, the cost of learning a piece of software is a fixed cost, while the cost of using a piece of software is a variable cost. The software provider would like to minimize both sorts of user costs, in order to make the software more attractive to consumers, but it is costly to do so. The question of interest to an economist is whether the market provides the right incentives to the software provider. Will the provider of software invest the socially correct amount of resources in minimizing each type of user cost?

Although we have discussed user costs in terms of software design, it is

[1] I first heard of this distinction in a discussion with Paul Scott. For some background concerning design of user-friendly software, see Nakamura (1990). See Mantei-Teorey (1988) for a discussion of benefit-cost analysis of software design from the viewpoint of software purchases by large corporations.

clear that it also applies to other types of goods. Consider for example, sporting equipment. Equipment designed for casual users may be very different from equipment designed for intensive users. An easy-to-learn tennis racket may be loosely strung, while an easy-to-use racket may be tightly strung. Or consider other sorts of hobbyist equipment such as cameras. Again, a camera for a casual user has a very different design than a camera for a professional user.

3.2 Market structure

However, there is an important distinction between the market for sporting equipment and the market for software. The market for tennis rackets and cameras seem to be reasonably competitive. There are a number of different types of products provided, and a given consumer can choose the type of product that is best for him. A casual user wants a product that is easy to learn; an intensive user wants one that is easy to operate, and the market provides both types of products.

The market for software is a bit different. It appears that for some products at least, the market is very highly concentrated. For example, in the database market, Ashton-Tate has over 50 percent of the market. In the wordprocessor market, WordPerfect appears to have over 60 percent of the market, and Microsoft controls another 25 percent. Until recently Lotus had a 75 percent share of the spreadsheet market.[2]

At a more aggregated level the market for personal computer software seems to be becoming more concentrated. According to the 1990 SoftLetter 100 list, Microsoft had 25% of industry revenue and the next 3 companies (Lotus, Novel, and WordPerfect) had another 25%. The top 12 companies had 77% of industry revenue, as compared to 66% in 1985.[3] This concentration has not gone unnoticed in Washington; the FTC is currently examining Microsoft looking for evidence of unfair trade practices.

In addition, there are compelling theoretical reasons to believe that the software market is unlikely to be a perfectly competitive market, due to the presence of increasing returns to scale. In the production of software, nearly all of the costs are fixed costs—the costs involved in designing, writing, debugging, documenting and marketing the software. Furthermore, most of these fixed costs are sunk costs—they are not recoverable if the firm exits the industry. The variable costs—the costs of duplicating, packaging, and distributing the software—are very small by comparison. The

[2] Data is for the MS-DOS market only. Database figures come from *Computer Reseller News,* July 15, 1991. Wordprocessing figures come from *Standard & Poor's Industry Surveys,* April 1991. Spreadsheet figures come from *Datamation,* December 15, 1990.

[3] Data from Bulkeley (1991); see also Fisher (1992).

fact that total costs are much larger than variable costs indicates that the likely equilibrium market structure will involve producers of mass-market software having a considerable amount of market power.

Another factor that suggests markets for software will be highly concentrated is the presence of network externalities among users. It is advantageous to me to have the same software as my colleagues since it makes it easier to share files, expertise, etc., and this tends to give the largest firm in the industry an advantage in selling more software.

Here we examine the admittedly extreme case of a monopolist. The case of monopolistic competition is certainly highly relevant, and I hope to examine it in future work, but it is beyond the scope of this paper. In any event, the phenomenon we examine here will apply in monopolistically competitive markets as well. Indeed, it will arise in any market where the producer has some degree of market power. The essential phenomenon we will discuss arises in any market where the producer faces a downward sloping demand curve for its product.

However, it is important for the following results that the manufacturer does not engage in product differentiation. That is, our model assumes that the producer sells only a single version of its product. This appears to be plausible in the case of software provision. There are some exceptions to this rule, such as WordPerfect and LetterPerfect, but these are rare. Generally, there is a single version of the product which is sold to a whole spectrum of users. This is quite different from the behavior of a typical manufacturer of tennis rackets, skis, or cameras.

3.3 The basic idea

The basic idea that I want to capture can be stated quite simply. We may think of ease-of-learning and ease-of-operation as two different dimensions of software "quality." It is well-known from the work of Spence (1975) that a monopolist does not in general have the right incentives to provide the appropriate amount of quality. Roughly speaking, the monopolist is interested in how a change in quality affects the willingness-to-pay of the marginal consumer, while the willingness-to-pay of the average consumer is the appropriate concern for social welfare.

Consider a monopolist contemplating investing an additional dollar in software design. Should the dollar go to making the software easier to operate or easier to learn? If the monopolist makes the software easier to learn then he will acquire additional customers—those consumers who previously weren't willing to invest in learning how to use the product, but now find the investment worthwhile. If the monopolist makes the software easier to operate, he will also acquire *some* new customers . . . but most of the benefits of the improvements in ease of operation accrue to the people who would have bought the software anyway. Since the monopolist cannot capture the full marginal benefits from making the product easier to operate, it will, in general, underinvest in this aspect of software design. However, since the monopolist can expand its market by making the software easy to learn, it will have the correct social incentives in this dimension.

3.4 Design of a word processor

Take, for example, the design of a word processor. Some users may have need for a word processor only once a week. Whether or not they buy a word processor depends on how difficult they think that it will be to learn to use it effectively. How quickly it reformats paragraphs or spell-checks is not of great significance to them. However, these features could be very important to a person who uses the word processor everyday. For an intensive user, the learning costs are small relative to the costs of operation; for a casual user, the learning costs are the dominant consideration.

Ideally, there would be "friendly" wordprocessors for casual users and "powerful" wordprocessors for intensive users. But if there is only one wordprocessor for both casual and intensive users there is an inevitable tradeoff in the design of such software. Should there be a special command to transpose two words? This could be useful to an intensive user, but probably not very useful to an occasional user. Providing and documenting such a command is costly to the software developer. In at least the first release of the software, the developer would probably concentrate more on the quality of the documentation and user interface rather than investing much time in adding rarely-used features. Similarly, the casual user would probably not be too concerned with how rapidly the document could be reformatted, or how quickly the spell-checker worked. But these factors could be very important to an intensive user.

People who use software on an occasional basis don't want a lot of choices—they are willing to give up some features in order to make the software easy to learn. People who use the software intensively are willing to invest in learning a variety of features since they will probably find occasion to use them.

Another important example of this distinction is in user support. "Hand-holding" support is critical to casual users, but not nearly so important to intensive users. In fact, intensive users would probably prefer to see a company devote more of its resources to improving software performance rather than providing increased handholding for new users. But it's the new users who bring in the new dollars—and that's why software companies invest in activities that can reduce the costs facing new users.

When we look at the evolution of personal computer software, we see improvements in both the ease-of-learning and the ease-of-use of software. The current behavior of software providers seems to be much more focussed on the ease-of-learning aspect of software design. In order to sell software to consumers who don't have it already, they have to make the software easier to learn. The attraction of user-friendly shells, such as MS-Windows, to software developers lies in the fact that once users have mastered the shell environment, the fixed costs of learning a new piece of software are much smaller for them. Hence one can expect that the demand for software products will increase.

Contrast the reaction of software vendors to MS-Windows with the reaction, say, to an increase in CPU speed. This might make getting things done a lot easier for intensive users of some software packages, but it probably wouldn't sell much new software. Of course, developers might redesign their software to add help features that were not feasible before—but that simply shows what is important to the software producers.

3.5 Monopoly provision of quality

The producer of a product chooses both the price and the characteristics of the product he produces. Most work in economics is concerned with the pricing decision. However, considerations of product design are also of great importance. Spence (1975) and Sheshinski (1976) consider the incentives facing a monopolist in choosing the "quality" of its product. Here quality should be thought of as a variable that shifts the demand curve for the product; in our application, "quality" is the ease-of-learning and the ease-of-operation described above.[4]

Spence (1975) computes the derivative of consumers' surplus minus costs evaluated at the monopoly position and derives two conditions sufficient to sign this derivative. The first involves comparing the impact of a quality change on the marginal versus the average consumer: if the average consumer values the change in quality more than the marginal consumer then the monopolist underprovides quality. The second involves examining the sign of $\partial^2 p(x,q)/\partial x \partial q$. These conditions provide an answer to the question, but it is hard to interpret precisely what they mean. As Schmalensee (1979) puts it: "It is very hard to form any general intuition about the sign (let alone the magnitude) of the crucial cross-derivative P_{QX}."

In order to determine the sign it is helpful to develop a microeconomic model of consumer choice, a task I pursue below. However, before doing that it may be useful to derive the Spence-Sheshinski result. The derivation below is different from the method used by Spence and Sheshinski and has the advantage that it focuses attention on the crucial aspect of the problem relevant to the case at hand.

[4] The most common example of quality in the literature is "durability." See the seminal work of Swan (1970) and the survey by Schmalensee (1979). See also the textbook treatment by Tirole (1988).

Let x denote the quantity and q the quality of some product. Let $u(x, q)$ be the utility of the product and $c(x, q)$ be the cost of providing it. Let $p(x, q) = \partial u(x, q)/\partial x$ be the inverse demand curve for the product. The social objective function is defined to be

$$W(x, q) = u(x, q) - c(x, q)$$

which is simply benefits minus costs. The monopolist's objective function is given by profit:

$$P(x, q) = p(x, q)x - c(x, q).$$

Let (x^m, q^m) denote the monopolist's profit-maximizing choice of output and quality. We are interested in the derivative of welfare evaluated at the monopolist's choice.

Write the welfare function as

$$W(x, q) = [u(x, q) - p(x, q)x] + [p(x, q)x - c(x, q)] = CS + PS.$$

This is simply the sum of consumer surplus plus producer surplus. If we differentiate with respect to q and evaluate the derivative at the monopolist's optimum, we see that the derivative of producer surplus must be zero—since the monopolist is already maximizing profits. Hence the derivative of welfare with respect to quality is simply the derivative of consumers' surplus with respect to quality. This is a significant simplification since it means that we don't have to model the cost side of things at all.[5]

How does consumer surplus change as quality changes? This derivative is given by

$$\frac{\partial W(x_m, q_m)}{\partial q} = \frac{\partial u(x_m, q_m)}{\partial q} - \frac{\partial p(x_m, q_m)}{\partial q} x_m.$$

We can write this expression as

$$\frac{\partial W(x_m, q_m)}{\partial q} = x_m \frac{\partial}{\partial q} \left[\frac{u(x_m, q_m)}{x_m} - p(x_m, q_m) \right].$$

Hence the sign of the derivative of the welfare change is just the sign of the term in brackets. The first term in the brackets is the total willingness-to-pay divided by the number of consumers who purchase the good; this is the average willingness-to-pay. The second term in the expression is the price—the marginal willingness-to-pay. The welfare effect of the quality

[5] This observation also has global implications. Start at the profit-maximizing position and consider a "large" change in quality. Profits go down since we've moved away from the profit-maximizing choice; if consumers' surplus also goes down, welfare unambiguously decreases.

change depends on how a quality change affects the *difference* between these two terms. This proves the first of Spence's observations.

Note that no calculations are necessary; all that is required is the observation that the derivative of profit is zero at the monopoly solution and the observation that consumers' surplus is proportional to the difference between an average and a marginal quantity. We now go on to ask what it is about demand that determines the sign of this quantity. In other words, how does consumers' surplus change as the demand curve moves? As shown in Figure 1, we can decompose a movement of the demand curve into a parallel "shift" and a "tilt." The shift doesn't change consumers' surplus at all; only the tilt matters. It is easy to see that if the demand curve gets flatter consumers' surplus decreases and if it gets steeper, consumers' surplus increases.[6]

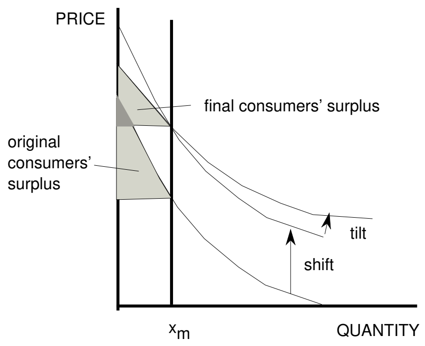

Figure 1. Decomposing the change in demand into a shift and a tilt.

Thus what matters is how a change in quality affects the *slope* of the demand curve; this is given by

$$\frac{\partial}{\partial q}\frac{\partial p}{\partial x} = \frac{\partial^2 p(x,q)}{\partial q \partial x}$$

This is, of course, simply the Spence-Sheshinki condition. However, the interpretation in terms of shifts and tilts turns out to be quite useful below.

We can also relate this back to the earlier discussion of marginal versus average valuations. Changes in quality that shift the demand curve have no effect on welfare since they don't affect the difference between the average

[6] The "tilt" terminology is slightly misleading since the normal usage of tilt implies a *constant* change in slope. Of course this is not necessary for the result; all that is required is that the change in q either increases or decreases the slope of the demand curve at every point—not that it changes the slope of the demand curve by the same amount at every point.

and marginal valuation. To affect the average and marginal consumer differently, the change in quality must affect the *slope* of the demand curve.

It follows that to answer the question of how a change in quality affects welfare, we need to construct a micro-model of consumer behavior and see how the quality variable enters the demand curve. Quality variables that shift the demand curve have no effect on welfare; variables that tilt the demand curve increase or decrease welfare depending on which way they tilt the demand curve.

3.6 The model

I now present a formal model of ease of learning versus ease of operation. I model the user costs in the following way. I suppose that there are a number of different users, each of whom uses the software more or less intensively. Let n be the number of times that a consumer uses a piece of software in some given time period, and let $g(n)$ be the number of consumers who use the software this often. For simplicity we take the frequency-of-use of the software to be independent of the ease-of-operation, although this can be relaxed.

Each time the software is run, the user bears a cost v. This is a variable cost of operation: it could refer to the time it takes to run the program, the complexity of the keystrokes necessary to run it, etc. High-intensity users—those who use the software a lot—pay a high variable cost.

Let F be the fixed cost of running the program. This is the cost that the user must pay regardless of his intensity of use. If she runs the program once or a hundred times, she must pay the same cost F. This should be thought of as the cost of learning to use the program. A program that is easy to learn has a low value of F; a program that is easy to operate has a low value of v.

Let $c(x, v, F)$ be the cost to the manufacturer of selling x copies of a program that has user costs of (v, F). For simplicity, we will suppose that the cost function has the separable form $c(x, v, F) = c_x(x) + c_v(v) + c_F(F)$. The term $c_x(x)$ measures the cost of producing x units of the software. The term $c_v(v)$ is the cost of designing software with variable costs v. The term $c_F(F)$ is the cost of designing software with fixed user costs F. This separable structure is not necessary for most of the results, but it makes the analysis simpler.

For simplicity we assume that the marginal costs of production are constant, and set $c_x(x) = c_x x + K$. Here c_x is the marginal cost of producing an extra copy of the software, once it has been created, and K is the fixed cost of producing the software. We should think of the fixed costs as being large relative to the variable costs of production. Note that the cost functions c_v and c_F should be *decreasing* functions of their argument since it should cost more to make a package with smaller user costs. It is natural to assume that both of these functions are convex, since the marginal cost of improving a package should increase the better the package is to start with.

Let b be the gross benefit to the user each time he or she uses the program. If a user runs the software n times, the net benefit accruing to the user is then $(b - v)n - F$. This is the gross benefit per use minus the user costs. If the package sells for a price of p, then a person who uses the software n times has a consumer surplus of $(b - v)n - F - p$.

The benefit, b, measures the *performance* of the software. In our formulation, all that matters to the consumer is the difference between the performance, b, and the ease-of-operation, v. We might think of this as the *net performance* of the software: the net benefit of the software per use. In general b is a choice variable—the producer can invest more or less effort in order to increase b. But since all that matters to the consumer is $b - v$, an increase b is equivalent to a decrease in v. Hence there is no need to carry out a separate analysis of the choice of b.

We suppose that a person who has positive consumer surplus will purchase the product, and a person who has negative consumer surplus will not. The marginal user will be the person who has a net surplus of zero. If n^* is the intensity of use by this consumer, then it must satisfy the equation

$$(b - v)n^* - F - p = 0,$$

which implies

$$n^* = \frac{F + p}{b - v}. \tag{3.1}$$

This gives us a relationship between the characteristics of the software, (v, F, p), and the number of *uses.* We want to convert this into a relationship between the price and the number of *users* in order to determine how many consumers will buy the product.

Let $G(n)$ be defined by

$$G(n) = \int_n^\infty g(t)\, dt.$$

This measures the number of people who use the software at least n times per period. If the number of uses by the marginal user is n^*, as defined in equation (3.1), and there are x users in total, x must satisfy the equation

$$G(n^*) = x. \tag{3.2}$$

Let $H(x)$ be the inverse function of $G(n)$. The function $H(x)$ measures the number of uses by the marginal person if x units are sold. Applying H to both sides of (3.2) and using (3.1), we can write

$$n^* = H(x) = \frac{F + p}{b - v}.$$

Solving for p as a function of x we have the inverse demand function

$$p(x) = (b - v)H(x) - F.$$

Since $G(n)$ is a monotonic decreasing function, so is its inverse, $H(x)$. Hence the inverse demand function is a decreasing function of price. Note that the variable cost affects the *slope* of the demand curve, while the fixed user cost merely *shifts* the demand curve.

This is quite reasonable. If the a software package becomes easier to operate, then all users are willing to pay more for it. But the high intensity users' willingness-to-pay goes up by more than the other users, since they use it more often. On the other hand, if the software becomes easier to learn, then everyone will be willing to pay more for it, regardless of their intensity of use.

We are now in a position to apply the preceding analysis concerning the welfare effect of changing v and F. However, it is useful to spell out the welfare analysis in slightly more detail.

In order to do this we first derive an expression for consumers' surplus. If x users buy the software the gross surplus (the area under the demand curve) is:

$$u(x) = \int_0^x p(t)\,dt = \int_0^x [(b - v)H(t) - F]\;dt.$$

If each package is sold at a price of $p(x)$, the net consumers' surplus is

$$\begin{aligned} u(x) - p(x)x &= \int_0^x [(b - v)H(t) - F]\;dt - [(b - v)H(x) - F]x \\ &= (b - v)\left[\int_0^x H(t)\,dt - H(x)\right]. \end{aligned} \tag{3.3}$$

Note that F drops out of this expression; it follows immediately that the derivative of consumers' surplus with respect to F is zero. Furthermore, since $H(x)$ is a decreasing function, it is easy to see that the expression in brackets is positive. A reduction in v helps the average consumer more than the marginal consumer since the average consumer uses the software more intensively than the marginal consumer.

The monopolist has the correct incentives with respect to ease-of-learning, but the wrong incentives with respect to ease-of-operation. Why? Essentially, the reason is the standard monopoly distortion pointed out by Spence (1975): the monopolist cares about the marginal consumer, not the average consumer. In our framework, the marginal consumer values ease-of-learning in exactly the same way as the average consumer; hence, there is no distortion in this aspect of the product design. But the marginal consumer in our model uses the product less intensively than the average consumer; hence the monopolist has too little incentive to invest in reducing this sort of user costs. From the monopolist's point of view, the high-intensity user will buy the product anyway, and the monopolist has no incentive to make the product easier to operate for them. But the monopolist has just the right incentive to make the product easy to learn, since this increases the size of its market, and makes *all* users willing to pay more for the product.

The problem discussed above arises due to the fact that in our model the monopolist doesn't have a way to extract any payment from the infra-marginal users, even though they would be willing to pay for improvements in ease-of-operation. In real life, the monopolist does have such an option: it can offer software upgrades. Intensive users will be willing to pay for those upgrades if they offer improved capabilities.[7]

However, typically a new release of the software is sold to both new and existing customers. When trading off investment in ease-of-use and ease-of-learning the software producer will still face the incentives described above: it will be willing to invest less in features valued by consumers who are sure to buy the product anyway. In any event, the fact that software can be upgraded is unique among products and is worth examining in its own right.[8]

The distortion in this model depends on the fact that the consumer cares about the number of *uses* while the monopolist cares about the number of *users.* If the monopolist could charge a price per use, there would be no distortion. To see this, imagine that the software is run on a mainframe computer so that the software provider can monitor the number of uses. The monopolist sets a schedule $\pi(n)$ that indicates the charge per use. The price schedule is given by

$$\pi(n) = \begin{cases} \infty & \text{if } n \leq F/(b-v) \\ (b-v)n + F & \text{if } n > F/(b-v). \end{cases}$$

It is easy to check that this price schedule extracts all the consumers' surplus from the users of the software. Hence the monopolist will choose the socially optimal levels of F and v.

3.7 Software that is too easy to learn

In the above analysis we've seen an example where the product of the monopolist has too little quality (too hard to operate) and just the right amount of quality (appropriate ease of learning). It would be nice to complement this with an example where a monopolist provides too *much* quality.

In order to do this, let us change the model slightly. Suppose now that there is no difference in intensity of use among consumers. For simplicity

[7] For some background on upgrades, see Bulkeley (1990).

[8] I intend to investigate software upgrades in future work.

suppose that all consumers use the program only once, and each gets the *same* net benefit $b - v$.[9] However, users differ in how difficult it is to learn to use a new program. To be specific, the net surplus from use of the computer program is

$$b - v - \gamma F - p.$$

Here F is a measure of the how easy the software is to learn—the fixed costs—and γ measures the capability of a given individual to learn the software. People with high values of γ find it more costly to learn a new piece of software than individuals with low values of γ. We suppose that γ is distributed in the population according to some cumulative distribution function $J(\gamma) = \int_0^\gamma j(t)\,dt$.

The marginal purchaser of the program satisfies the condition that benefits are just equal to the price of the software,

$$b - v - \gamma F - p = 0,$$

so

$$\gamma^* = \frac{b - v - p}{F}.$$

Everyone with a smaller γ buys the software, so the total sales are

$$x = J(\gamma) = J((b - v - p)/F).$$

Letting K be the inverse of J, we have

$$K(x) = \frac{b - v - p}{F},$$

which implies that the inverse demand function is

$$p = b - v - FK(x).$$

Note that $K'(x) > 0$ since it is the inverse of a cumulative distribution function.

For this form of demand, changes in v shift the demand function and changes in F tilt the demand function. According to our previous analysis, the monopolist produces the right ease-of-operation, but the wrong ease-of-learning. In fact $\partial^2 p/\partial x \partial F < 0$. From our previous analysis, this implies that welfare increases if F increases—that is, welfare goes up if the software is made harder to learn! In this model the monopolist *overinvests* in making the software easy to learn.

Why is this? In this model the marginal consumer is one who finds the software harder to learn than the average consumer. Hence making the software a little easier to learn benefits the marginal consumer more than the average consumer. Hence the monopolist tends to invest too many resources in attracting marginal consumers rather than, say, improving the functioning of the program for the inframarginal consumers.

[9] An example that fits this model might be tax preparation software. You use this software only once a year, so that there is little difference in intensity of use across the population. Whether or not you choose to use the software depends primarily on how easy it is to learn.

3.8 Policy implications

What are the implications of the easy-of-operation/ease-of-learning distortion from a policy perspective? Obviously it is premature to draw definitive conclusions from such a simple model, but it is worthwhile raising the question to see where a more in-depth analysis may lead.

Since we are examining a monopolist, it is always in the social interest to increase output. One possibility would be to pursue antitrust actions to eliminate the monopoly power. However, it is far from clear that this would be appropriate since it would affect incentives to innovate and perhaps lead to excessive product differentiation.

Accordingly, we adopt the viewpoint that the monopoly output distortion should be tolerated. However, the above arguments suggest that even if the monopoly output remains constant there still will be social benefits to encouraging changes in the monopolist's provision of "quality."

In different models of quality, different tools may be appropriate. For example, in some models, setting minimal (or maximal!) quality standards may be appropriate. However, this instrument seems implausible in our model of software.

One interesting policy choice is to subsidize the provision of software "quality." In practice this would be done by publicly sponsored research grants. We suppose that social policies such as this will reduce the cost of providing software that is easy-to-learn and software that is easy-to-use. What will be the impact of such subsidies on social welfare?

We first examine the original model where consumers differ in the intensity of use. Suppose that we subsidize the cost of developing easy-to-learn and easy-to-operate software at rates s_v and s_F respectively. Welfare can be written as

$$\begin{aligned} W(x,v,F) = [u(x,v,F) - p(x,v,F)x] + [&p(x,v,F)x - (1-s_v)c_v(v) \\ &-(1-s_F)c_F(F)] \quad - [s_v c_v(v) + s_F c_F(F)]. \end{aligned}$$

The three bracketed terms in this expression are consumer surplus, producer surplus, and government expenditure respectively. Differentiating this expression with respect to s_v and s_F and evaluating the derivative at the monopoly equilibrium with $s_v = s_F = 0$ we have

$$\begin{aligned} \frac{dW}{ds_v} &= -\frac{\partial p}{\partial x}\frac{\partial x}{\partial s_v} + \left[\frac{\partial u}{\partial v} - \frac{\partial p}{\partial v}x\right]\frac{\partial v}{\partial s_v} + \left[\frac{\partial u}{\partial F} - \frac{\partial p}{\partial F}x\right]\frac{\partial F}{\partial s_v} \\ \frac{dW}{ds_F} &= -\frac{\partial p}{\partial x}\frac{\partial x}{\partial s_F} + \left[\frac{\partial u}{\partial F} - \frac{\partial p}{\partial F}x\right]\frac{\partial F}{\partial s_F} + \left[\frac{\partial u}{\partial v} - \frac{\partial p}{\partial v}x\right]\frac{\partial v}{\partial s_F}. \end{aligned}$$

In calculating these derivatives several terms drop out due to utility maximization and profit maximization. The remaining terms are composed of two effects: the direct effect on consumers' surplus of changing v and F and the indirect effect of the induced output change and the cross effect of the subsidy.

We have already calculated the direct effect; it is zero for changes in F and positive for reductions in v. I show in the appendix that both the output effect and the cross effect have a positive effect on utility. Hence there is a case to be made for imposing (small) subsidies on both cost functions. However, in terms of the impact on output, a reduction in F is exactly equivalent to imposing an output subsidy on the monopolist. Hence the social benefits of making the software easier to learn are just the same as the benefits from an output subsidy. This is to be compared to the effects of subsidizing v. In this case the subsidy benefits consumers both through the increase in output *and* through the improvement in "quality."

Another way to observe this is to consider the case where $F = 0$. In this case the profit-maximization problem for the monopolist becomes

$$(b - v)H(x) - (1 - s_v)c_v(v).$$

The first-order condition for output is

$$(b - v)H'(x) = 0,$$

which is independent of s_v. Hence output doesn't change when ease-of-operation is subsidized. Nevertheless, welfare increases due to the impact of the quality change on the inframarginal consumers.

We turn now to the second model, where consumers differ in the cost of learning. We show in the appendix that $dx/ds_F > 0$ in this case as well. Hence the impact on social welfare is composed of two effects: the benefit from having more output and the cost from the monopolist investing "too much" in making the software easy-to-learn. The combination is ambiguous, but at least it forces attention on the proper tradeoff: one would have to expect a big output effect from a subsidy in order for it to be worthwhile from a social point of view.

I interpret this as saying that it is reasonable to use public funds to subsidize research on how to make software easy to operate. However, there is no particular argument, in the context of this model at least, to subsidize research on how to make software easier to learn—the market gives the monopolist the right incentives with respect to this choice already, at least conditional on the output chosen by the monopolist. Of course, there may be other reasons to subsidize research on this aspect of software design. For example, there may be economies of scale in research, or there may be problems with appropriability of intellectual property that could cause problems for developing easy-to-operate software in the private sector. Or it may be that there are lower costs to developing easy-to-learn software in an educational environment. It would clearly be premature to make policy pronouncements without careful consideration of these possibilities.

3.9 Appendix. Comparative statics of the profit maximization problem

Here we study the impact of subsidizing research on ease-of-learning and ease-of-operation. We assume that this research lowers the cost of providing easier-to-operate and easier-to-learn software. We model this cost reduction as being equivalent to subsidies of s_v and s_F on the cost functions. Letting $R(x) = H(x)x$, we can write the monopolist's profit maximization problem as

$$\max_{x,v,F} (b - v)R(x) - Fx - (1 - s_v)c_v(v) - (1 - s_F)c_F(F).$$

Note that we have set the cost of production equal to zero. Alternatively we could incorporate a constant marginal cost of production into F.

The first-order conditions for this problem are

$$\begin{aligned} (b - v)R'(x) - F &= 0 \\ -R(x) - (1 - s_v)c'_v(v) &= 0 \\ -x - (1 - s_F)c'_F(F) &= 0. \end{aligned}$$

Totally differentiating this system and evaluating the derivatives at $s_v = s_F = 0$ we have

$$\begin{pmatrix} (b-v)R'' & -R' & -1 \\ -R' & -c''_v & 0 \\ -1 & 0 & -c''_F \end{pmatrix} \begin{pmatrix} dx \\ dv \\ dF \end{pmatrix} = \begin{pmatrix} 0 \\ -c'_v ds_v \\ -c'_F ds_F \end{pmatrix}.$$

In the case of a regular maximum the second-order conditions imply that the determinant of the Hessian matrix on the left-hand side of this expression will be negative and all principal minors of order 2 will be positive. This latter condition implies

$$\begin{aligned} (b - v)R''(x)c''_v(v) + R'(x)^2 &< 0 \\ (b - v)R''(x)c''_F(F) + 1 &< 0 \end{aligned} \tag{3.4}$$

These conditions are useful in signing the comparative statics effects.

Let $H < 0$ denote the value of the determinant of the Hessian and solve for the various differentials.

$$\begin{aligned} dF &= \frac{-ds_v R'(x)c'_v(v) + ds_F c'_F(F)\left[(b - v)R''(x)c''_v(v) + R'(x)^2\right]}{H} \\ dv &= \frac{-ds_F R'(x)c'_F(F) + ds_v c'_v(v)\left[(b - v)R''(x)c''_F(F) + 1\right]}{H} \\ dx &= \frac{ds_v R'(x)c'_v(v)c''_F(F) + ds_F c'_F(F)c''_v(v)}{H} \end{aligned}$$

It is straightforward to verify

- dv/ds_v and dF/ds_F are negative. Hence subsidizing ease-of-learning or ease-of-operation will tend to lead to improvements in those variables.
- dx/ds_v and dx/ds_F are positive. Hence the subsidies tend to increase output.
- dv/ds_F and dF/ds_v are negative. Hence subsidizing ease-of-learning will lead to an improvement in ease-of-operation and vice versa.

The other model discussed in the text was based on differences among the consumers in ease of learning. The profit maximization problem in that case is:

$$\max_{x,F} (b-v)x - FK(x)x - (1-s_F)c_F(F).$$

The first-order conditions are

$$\begin{aligned} (b-v) - FR'(x) &= 0 \\ -R(x) - (1-s_F)c_F'(F) &= 0, \end{aligned}$$

where $R(x) = K(x)x$. Totally differentiating this system and evaluating the result at $s_F = 0$ we have

$$\begin{pmatrix} -FR'' & -R' \\ -R' & -c_F'' \end{pmatrix} \begin{pmatrix} dx \\ dF \end{pmatrix} = \begin{pmatrix} 0 \\ -c_F' ds_F \end{pmatrix}.$$

From this we find

$$\frac{dx}{ds_F} = \frac{\begin{vmatrix} 0 & -R' \\ -c_F' & -c_F'' \end{vmatrix}}{\begin{vmatrix} -FR'' & -R' \\ -R' & -c_F'' \end{vmatrix}} > 0.$$

References

Bulkeley, W. (1990). Software users are beginning to rebel against the steady stream of upgrades. *Wall Street Journal*, B4.

Bulkeley, W. (1991). Software industry loses start-up zest as big firms increase their domination. *Wall Street Journal*, *August 27*, B1.

Fisher, L. (1992). Business turns tough in software. *New York Times*, *December 14*.

Mantei, M., & Teorey, T. (1988). Cost/benefit analysis for incorporating human factors in the software lifecycle. *Communications of the ACM*, *4*, 428–439.

Nakamura, R. (1990). The x factor. *Infoworld*, *November 19*, 51–55.

Schmalensee, R. (1979). Market structure, durability, and quality: a selective survey. *Economic Inquiry*, *42*, 177–196.

Sheshinski, E. (1976). Price, quality and quantity regulation in monopoly situations. *Economica*, *43*, 127–137.

Spence, M. (1975). Monopoly, quality and regulation. *Bell Journal of Economics*, *6*(2), 417–429.

Swan, P. (1970). Durability of consumer goods. *American Economic Review*, *60*, 884–894.

Tirole, J. (1988). *The Theory of Industrial Organization.* MIT Press, Cambridge, MA.

Chapter 4

PRICING CONGESTIBLE NETWORK RESOURCES

We describe the basic economic theory of pricing a congestible resource such as an ftp server, a router, a Web site, etc. In particular, we examine the implications of "congestion pricing" as a way to encourage efficient use of network resources. We explore the implications of flat pricing and congestion pricing for capacity expansion in centrally planned, competitive, and monopolistic environments.

The Internet is now involved in a major transformation from a government sponsored project to a private enterprise. Privatization and commercialization of the Internet means that providers of network connectivity and services will have to confront issues of pricing and cost recovery. When connectivity was provided to users via government subsidies, little attention was paid to these issues. Suddenly, they have become quite significant. At the same time, new problems in resource allocation are emerging as other telecommunication network technologies begin to converge.

We think that economic modeling can play a significant role in thinking about the consequences of various issues facing decisionmakers. Given the current paucity of economic data about the Internet, economic analysis is unlikely to give precise numerical answers to many questions of interest. Still explicit economic models can serve as a useful guide to "how to think" about some of these issues.

For example, consider the problem of providing bandwidth which will be shared by many users. As network technology and availability advances, there will likely be places and periods when bandwidth is scarce and periods when it is abundant. When the supply of bandwidth far exceeds the demand, there is little role for economics. But when the demand for bandwidth exceeds the supply, the fundamental issues of resource allocation become important.

There are many network resources whose performance suffers when there is "overuse": the switching capacity of the routers, the bandwidth of the transport medium, the disk and CPU capacity of popular servers, etc. When users access such resources they presumably take into account their

Co-authored with Jeff MacKie-Mason. We are grateful to Marvin Sirbu and Scott Shenker for their comments. This work was supported by the National Science Foundation grant SES-93-20481.

own costs and benefits from usage, but ignore the congestion, delay, or exclusion costs that they impose on other users. Economists refer to this phenomenon as a "congestion externality"; in ecology, it is known as the "problem of the commons" (Hardin (1968)).

There are many ways to deal with congestion externalities. One way is to establish social norms that penalize inappropriate behavior. Such norms can work well in small groups where there is repeated interaction, but they often do not scale well to a system with millions of users.

Another way to deal with congestion is to establish rationing or quota systems. (Bohn (1993)). One appeal of rationing is that is relatively easy to implement. Indeed, it is common today to see file servers, Web servers, and other network services that reject additional users when the load is too high.

Despite the simplicity of rationing and quotas, economists tend to favor pricing mechanisms as a way of alleviating congestion. One important feature of congestion prices is that they not only discourage usage when congestion is present, but they also generate revenue for capacity expansion. Indeed, it has long been recognized that under certain conditions the optimal congestion prices for a fixed amount of capacity will automatically generate the appropriate amount of revenue to finance capacity expansion.

In previous work we have proposed some simple pricing schemes to deal with congestion (MacKie-Mason (1993a, 1993b)). Here we examine the issue of how the pricing scheme chosen affects industry structure and performance. Our framework is that of "club theory," a term used by Buchanan (1965) to deal with the provision of shared goods. A textbook treatment of club theory can be found in Cornes (1986). The papers in the literature that are closest to the treatment here are Scotchmer (1985a, 1985b); we will describe the relationship of our work to this literature in more detail below.

4.1 Notation

Let x_i denote person i's use of the network resource and $X = \sum_{j=1}^{n} x_j$ the total use of the resource. The user cares about her own use, x_i, and the delay that she encounters. Delay should be interpreted as a general congestion cost: it can include the cost of exclusion, congestion, and so on. Delay depends on the *utilization* of the resource, which we define to be total use divided by capacity: $Y = X/K$. We summarize the preferences of the user by a utility function $u_i(x_i, Y) + m_i$, where m_i is money that the user has to spend on other things. We assume that $u_i(x_i, Y)$ is a differentiable, concave function of x_i and a decreasing concave function of Y.[1]

[1] Later on we consider a special form of this function, $u_i(x_i, Y) = v_i(x_i) - D(Y)$, where $D(Y)$ is interpreted as a delay cost. However, we will not introduce this specification until it is necessary.

The critical feature of this specification is the relationship between usage and capacity: if total usage (X) is doubled and capacity (K) is also doubled, then utilization $Y = X/K$ and hence delay remain constant.[2] We let $c(K)$ measure the cost of providing capacity. For simplicity we take this to be the only cost of providing the service.[3]

This specification is general enough to capture the essence of many network resources. Consider the specific example of an ftp server. In this context x_i could be the number of bytes transferred to user i, K would be the capacity of the server in terms of how many total bytes it can transfer in a given time period, and X would be total bytes transferred to all users. It is natural to suppose that user i cares about the total amount of material she retrieves and the delay involved in retrieving it. A router is another example. In this case x_i would be the bytes sent to (and/or received from) the router by user i, X would be the total use of the router, and K would be the maximum throughput of the router.

4.2 Efficient use and capacity

We first examine the efficient pattern of usage given some given capacity K. By definition, the efficient pattern maximizes the sum of benefits minus costs. Denoting aggregate net benefits by $W(K)$ we have:[4]

$$W(K) = \max_{x_j} \sum_{j=1}^{n} u_j(x_j, Y) - c(K). \tag{4.1}$$

The optimal solution must satisfy the first-order condition

$$\frac{\partial u_i(x_i, Y)}{\partial x_i} = -\frac{1}{K} \sum_{j=1}^{n} \frac{\partial u_j(x_j, Y)}{\partial Y} \tag{4.2}$$

This says that user i should use the system until the marginal benefit from her usage equals the marginal cost that she imposes on the other users.

[2] Delay is fully determined by average utilization only under certain traffic conditions. More generally delay may depend on peak utilization or the variance of utilization. Generalizing the model to account for such effects is clearly of interest, but is beyond the scope of this paper.

[3] In principle, costs could also depend on the amount of usage (X) and on the number of users (n), but we omit these in order to keep the model simple. Capacity costs are normally the dominant costs for most services of interest to us.

[4] We maximize total benefits minus total costs, without making any particular distributional judgments. We could, of course, allow for lump-sum transfer payments to the agents that reflected such concerns. However, such transfer payments would not modify the form of the solution to the benefit-cost problem considered here.

We can decentralize this solution by defining a "shadow price"

$$p_e = -\frac{1}{K} \sum_{j=1}^{n} \frac{\partial u_j(x_j, Y)}{\partial Y}, \tag{4.3}$$

which measures the total marginal congestion cost that an increase in x_i imposes on the users; note that this is independent of i. Suppose that consumer i is charged a price p_e for usage. Then she would want to solve the following problem

$$\max_{x_i} \ u_i(x_i, Y) - p_e x_i.$$

The solution to this problem is characterized by

$$\frac{\partial u_i(x_i, Y)}{\partial x_i} + \frac{1}{K} \frac{\partial u_i(x_i, Y)}{\partial Y} = p_e. \tag{4.4}$$

Referring to the definition of p_e in equation (4.3), we see that for large n the second term on the left-hand side will be negligible relative to the right-hand side of the equation. For large n this expression is essentially the same as the first-order condition for the social optimum given in (4.2), and thus the decentralized solution corresponds with the social optimum.

To see this more explicitly, consider the special case where $u_i(x_i, Y) = v_i(x_i) - D(Y)$. Then the social optimum in equation (4.2) is described by

$$u_i'(x_i) = \frac{n}{K} D'(Y),$$

and the individual optimization in equation (4.4) is

$$u_i'(x_i) = \frac{n+1}{K} D'(Y).$$

For large n these are virtually the same.

Economists say that the price p_e "internalizes" the externality by making the user face the costs that she imposes on the other users. The point of introducing the shadow price is to emphasize the fact that each user should face (essentially) the *same* price for usage—the sum of the marginal congestion costs that each user imposes on the other users.

Capacity expansion

In the maximization problem (4.1) we used $W(K)$ to denote the maximum welfare given an arbitrary capacity K. What happens to welfare as we expand capacity? Differentiating (4.1) with respect to K, we have[5]

$$W'(K) = -\sum_{j=1}^{n} \frac{\partial u_j(x_j, Y)}{\partial Y} \frac{X}{K^2} - c'(K).$$

Using the shadow price defined above, we can write this as

$$W'(K) = p_e \frac{X}{K} - c'(K). \tag{4.5}$$

From this it follows that $W'(K) > 0$ if and only if $p_e X - c'(K)K > 0$. This means that expanding capacity will increase welfare if and only if the revenue from the congestion fees ($p_e X$) exceeds the value of capacity ($c'(K)K$), where capacity is valued using the marginal cost of capacity.

Hence the shadow price p_e plays a dual role: it provides a measure of the social cost of increased usage for an given capacity, but it *also determines the value of a change in capacity*. The fact that congestion fees send the right economic signals to expand capacity under certain conditions was noted by Mohring (1962) and Strotz (1965); it takes various forms in the literature and is considered a classic principle of congestion pricing.

4.3 Pricing in a competitive market

The above discussion describes optimal pricing in a utopian world of welfare maximization. In the brave new world of deregulated, privately-provided information network services we would expect to see provision of network resources by profit-seeking firms. What kind of prices would emerge in such a market environment?

The answer depends on the details of market structure: clearly a monopoly or oligopoly structure will result in different (presumably higher) prices than a competitive market. We begin with the admittedly special case of a competitive market with many independent producers; later we examine monopoly provision.

We suppose that each producer uses a "two-part tariff" for pricing: a "subscription/attachment" fee of q per user per month, say, plus a usage fee of px_i. A representative producer's profits can then be written as

$$\pi = pX + nq - c(K)$$

Here pX is the revenue collected by usage-sensitive fees, nq is the revenue collected from connection fees, and $c(K)$ is the cost of providing capacity K. This appears to be a natural form for pricing network access and usage. Of course, pure connection pricing, in which $p = 0$, and pure usage pricing, in which $q = 0$, are special cases of this pricing form.

[5] Note that terms involving $\partial x_j / \partial K$ drop out due to the first-order conditions given in (4.2). This is an instance of what economists call the envelope theorem. (See Varian (1992)).

Consumer optimization

The utility maximization problem for consumer i is to choose which network resource to use and how much to use it. We suppose that there are (potentially) many suppliers with possibly different utilization levels. Suppliers with lower levels of utilization can charge more due to the better service they provide. We write the price offerings of a representative supplier with utilization Y as $(p(Y), q(Y))$, where $p(Y)$ is the usage fee and $q(Y)$ is the subscription fee.[6]

The utility maximization problem for a representative consumer now becomes

$$\max_{x_i, Y} \; u_i(x_i, Y) - p(Y)x_i - q(Y).$$

That is, the consumer chooses which provider to use (represented by Y) and how much to use (represented by x_i). For convenience, we assume that the menu of offered prices can be treated as a continuous and differentiable function of Y. The consumer's optimization problem has first-order conditions

$$\begin{aligned} \frac{\partial u_i(x_i, Y)}{\partial x_i} - p(Y) &= 0 \\ \frac{\partial u_i(x_i, Y)}{\partial Y} - p'(Y)x_i - q'(Y) &= 0. \end{aligned} \tag{4.6}$$

The first equation shows that each user will use the resource until the value of additional usage equals its price. The second equation shows that the consumer's choice of delay satisfies the condition that the marginal utility cost of increased delay must be compensated by a reduced expenditure, $p'(Y)x_i + q'(Y)$. Adding this last equation up across consumers gives us an expression that we will use below,

$$p'(Y)X + nq'(Y) = \sum_{j=1}^{n} \frac{\partial u_j(x_j, Y)}{\partial Y}. \tag{4.7}$$

Producer optimization

A representative producer chooses its capacity K and how much bandwidth to supply to users. We assume that there are many competing producers, each of whom takes the price-quality schedules $(p(Y), q(Y))$ as being outside of its control; i.e., determined by the competitive market.

[6] For simplicity, we assume that each firm offers only one class of service; this can easily be generalized.

The profit maximization problem facing a representative producer is to choose X and K to maximize profits given the price-quality schedules available in the market

$$\max_{X,K} \; p(Y)X + nq(Y) - c(K),$$

The first-order conditions are

$$p(Y) + p'(Y)\frac{X}{K} + n\frac{q'(Y)}{K} = 0$$
$$-p'(Y)\left(\frac{X}{K}\right)^2 - nq'(Y)\frac{X}{K^2} = c'(K).$$

Collecting terms we can write:

$$p(Y) + [p'(Y)X + nq'(Y)]\,\frac{1}{K} = 0 \tag{4.8}$$
$$-\,[p'(Y)X + nq'(Y)]\,\frac{X}{K^2} = c'(K). \tag{4.9}$$

Using equations (4.6) and (4.7), we can further simplify these equations to

$$p(Y) = \frac{\partial u_i(x_i, Y)}{\partial x_i} = -\frac{1}{K}\sum_{j=1}^{n}\frac{\partial u_j(x_j, Y)}{\partial Y} \tag{4.10}$$

$$-Y\sum_{j=1}^{n}\frac{\partial u_j(x_j, Y)}{\partial Y} = c'(K)K \tag{4.11}$$

Comparing (4.10) to (4.2) we see that the competitive price will result in the optimal degree of congestion. By combining (4.10) and (4.11) we can write

$$p(Y)X = c'(K)K, \tag{4.12}$$

which leads to the same rule for optimal capacity that we obtained in equation (4.5).

In this model a competitive supplier is forced to charge the socially optimal price for the quality of service that he offers. Why is the competitive market price equal to the sum of congestion costs? The term $-(1/K)\sum_j \partial u_j/\partial Y$ is how much the other users of the resource would be willing to pay the provider to *refrain* from selling additional usage. If this is less than the price a user is willing to pay for additional usage, the competitive supplier would want to allow more usage. The producer would stop providing additional usage when the price that a user is willing to pay for additional use is balanced with the amount that the other users are willing to pay for a reduction in total usage.

Free entry

If there are no restrictions on entry, firms will enter the industry until profits are driven to zero:

$$p(Y)X + nq(Y) - c(K) = 0.$$

Substituting the expression for $p(Y)$ derived above, we can write the zero-profit condition as

$$nq(Y) = c(K) - c'(K)K.$$

Dividing through by $c(K)$ we have

$$\frac{nq(Y)}{c(K)} = 1 - \frac{c'(K)}{c(K)/K} = 1 - \frac{1}{e},$$

where e is the *elasticity of scale* (average cost over marginal cost). If the marginal cost of capacity is small relative to the average cost, subscription fees will cover most of the cost of providing the service. If the marginal cost of capacity is large, then usage fees will contribute more to recovering total costs.

Scotchmer (1985a) examines a model of two-part pricing of a congestible resource that has some features in common with the one described above. In her model, congestion depends on the number of users, not the total usage, and the capacity of the club is fixed. (This is natural for the kinds of clubs that motivated her study: golf courses, ski clubs, swimming lanes, etc.; it is less natural in our context.) She considers an oligopolistic model with a finite number of firms and examines the limiting behavior as the number of firms increases. She finds that the connection fee goes to zero as the number of firms is increased. Although this result is derived under the assumption that the technology of each firm has fixed capacity, it appears that something similar to this will occur for the more general technology we consider here.

We should emphasize that we have examined the functioning of an idealized competitive industry—a market structure with many firms each having a small market share. There is a very real question of whether the provision of network services will actually have this industrial structure. A full examination of this issue is beyond the scope of this paper, but it is worth observing that a critical consideration is the nature of the cost function for providing network services. If larger firms can produce at lower unit cost, then a large firm will be able to undercut the price of a small firm and there will a tendency to end up with a single firm dominating the market.[7] We

[7] This is known as the case of a natural monopoly; see Sharkey (1982) for a detailed treatment.

should emphasize that this is only a tendency: antitrust policy, regulation, and other tools can be used to influence the ultimate outcome.

We conjecture that a primary factor in determining the industry structure of digital communication networks will be the ease of interconnection. If it is costly or difficult for networks to interconnect, large providers will have an automatic advantage in competing for customers. If this view is correct, there may well be a case for government policy to help coordinate interconnection agreements so as to ensure a healthy competitive environment.

Customer sorting and multiple qualities of service

Nothing in this model implies that there will be a single "optimal" quality of service offered. If all users were identical then the joint solution of equations (4.6) would yield a single quality Y^*, and associated prices $(p(Y^*), q(Y^*))$. However, user preferences for most services are often heterogeneous: some users may be very intolerant of delay while others may prefer to wait but pay low prices.

When customers have heterogeneous preferences for quality, social welfare is generally *not* maximized by having a single, "high quality" service or product available. Typically, there will be users who would prefer to accept lower quality in exchange for a price reduction—they value the quality difference less than they value the other goods and services they can buy with the savings. Competition with free entry will then force each quality level to be priced efficiently. Some suppliers will have low prices and high congestion, while others offer high prices and low congestion.

How does a competitive market arrive at the socially optimal variety of price-quality choices? Suppose that there are two types of user: delay-tolerant and delay-intolerant, but only one "average" quality of service is initially offered by all the firms. When would it pay for a firm to offer a different quality of service than its competitors?

By offering a quality of service optimized for one of the groups, a deviating firm could attract all the customers from that group. If the revenue from this deviation exceeds the cost of providing the new quality, this would increase the deviant firm's profits. If there are no fixed costs to creating different qualities we would expect to see as many different qualities as there are types of consumer preferences.

But what if there are large fixed costs to adding new service qualities? In this case it may well not be profitable for a deviant firm to provide a different quality since the entrant may have trouble extracting sufficient profits to cover its costs. Hence the equilibrium number of firms and variety of qualities of service offered will depend on the fixed costs of creating new qualities of service.[8]

When individual users have heterogeneous preferences

Thus far we have considered what happens if different users have different preferences for the resource. What if a single user has different willingness to pay for a resource when using it for different purposes? For example, a user may place a high value on the e-mail access from her network service provider, but a lower value on the ability to engage in real-time video conferencing. If there were small costs of connecting to more than one service provider, then we might see a "restaurant" equilibrium: various providers offering different service qualities at different prices, with a single consumer using more than one provider for different purposes.

However, there may be significant costs of accessing additional providers. For example, it might require having multiple lines running into the home or office, as we now have with telephone, cable and electric lines. If the costs of having multiple providers for multiple services get high enough, then we might expect to see single providers who offer multiple qualities of service. There has been considerable recent interest in the development of integrated services networks; see Braden-Clark-Shenker (1994) for a proposed multiple quality-of-service architecture for the Internet. Pricing is likely to be an effective mechanism for allocating different service qualities to appropriate uses, although the type of pricing that emerges will depend crucially on the evolution of the technological infrastructure (MacKie-Mason-Varian (1994a)).

Adding capacity

We saw earlier that the efficient congestion prices send the right signals for capacity expansion. Let us see how this works in a competitive market.

Suppose that a competitive firm must decide whether to add additional capacity ΔK. We consider two scenarios. In the first scenario, the firm contemplates keeping X fixed and simply charging more for improved quality of service due to the reduced delay. The extra amount it can charge user j is:

$$[q'(Y) + p'(Y)x_j]\,\frac{dY}{dK}\Delta K.$$

[8] Another factor that influences the number of firms is the presence of "network externalities." These occur when one consumer's utility of connecting to a network depends positively on the number of other users who are connected to the network. See Katz and Shapiro (1985) and Economides (1994) for an analytical treatment of this effect.

Using equation (4.6) this becomes

$$-\frac{1}{K}\frac{\partial u_j}{\partial Y}\frac{X}{K}\Delta K.$$

Summing this over all consumers and using equation (4.10) we have

$$p\frac{X}{K}\Delta K.$$

This will increase profits if the increase in revenue is greater than the cost of capacity expansion:

$$p\frac{X}{K}\Delta K - c'(K)\Delta K = \left[p\frac{X}{K} - c'(K)\right]\Delta K > 0.$$

Comparing this to equation (4.5) we see that profits will increase if and only if net social benefits increase.

In the second scenario, the firm expands its capacity and keeps its price fixed. In a competitive market it will then attract new customers due to the reduction in delay. In equilibrium this firm must have the same delay as other firms charging the same price. Suppose that in the initial equilibrium $X/K = Y$. Then the additional usage must satisfy $\Delta X = Y\Delta K$. It follows that the increase in in profit for this firm is given by

$$pY\Delta K - c'(K)\Delta K = \left[p\frac{X}{K} - c'(K)\right]\Delta K.$$

Again we see that capacity expansion is optimal if and only if it increases profits.

4.4 Equilibrium without usage fees

In this model usage fees play two critical roles—they determine both the efficient level of usage and the efficient level of capacity. However, usage-based pricing itself is expensive—it requires an infrastructure to track usage, prepare bills, and collect revenues. These transactions costs may be substantial, and a general examination of usage-based pricing must compare the benefits from improved resource allocation with the costs of accounting and billing. We do not attempt that exercise here. However, it is of considerable interest to examine how a model might function that has no usage fees, but only attachment/subscription fees.

It is convenient to specialize the model described above to a specific form for utility:[9]

$$u_i(x_i, Y) = v_i(x_i) - D(Y).$$

[9] We make this choice primarily to simplify the exposition; most of the results can be obtained without it, but with somewhat more effort.

Here $D(Y)$ is directly identified as the "delay costs" from congestion. We assume that $D(Y)$ is an increasing, differentiable, convex function. This says that the delay costs increase with utilization, and that they increase at an increasing rate. Note that this additive form implies that additional delay does not affect the marginal benefits from usage—an admittedly extreme assumption.

For this form of utility, the equilibrium values of (K^e, Y^e) in the world with usage based pricing can be written as

$$\begin{aligned} v_i'(x_i^e) &= \frac{n}{K^e} D'(Y^e) \\ nD'(Y^e)Y^e &= c'(K^e)K^e. \end{aligned} \tag{4.13}$$

The conditions are found simply by writing the conditions (4.10–4.11) for the special form of the utility function that we have adopted.

Let us now consider what would happen if only attachment pricing were available. Since access is priced, but there is no price for usage we assume that agent i satiates at some point x_i^a. This determines $X^a = \sum_{j=1}^n x_j^a$.

User i's utility maximization problem for Y is

$$\max_Y \; v_i(x_i^a) - D(Y) - q(Y),$$

which leads to the first-order condition

$$-D'(Y) = q'(Y).$$

Adding up across the consumers gives us

$$nq'(Y) = -nD'(Y). \tag{4.14}$$

The supplier's profit-maximization problem is

$$\max_K \; nq(Y) - c(K),$$

which has first-order conditions

$$-nq'(Y)\frac{X}{K^2} = c'(K).$$

Combining this with equation (4.14) we see that the equilibrium solution with no attachment pricing only must satisfy the equilibrium condition

$$nY^a D'(Y^a) = c'(K^a)K^a. \tag{4.15}$$

Comparing this to equation (4.13) we see that the form of the equation that determines equilibrium capacity is the same with and without usage-based pricing: in either case the amount of capacity will be determined by the willingness to pay for reduction in delay.

However there is one subtlety: even though the *form* of the equation is the same in both cases, it may be that the equilibrium magnitudes of the relevant variables are different. In particular, it can easily happen that the number of users is different with and without usage-based pricing. We must therefore compare the equilibria under the two different scenarios: when the number of users is the same and when the number of users is different.

The number of users is the same

For fixed X, the equilibrium capacity is determined by

$$nD'(X/K)\frac{X}{K} = c'(K)K. \tag{4.16}$$

The convexity of $D(Y)$ implies that the left-hand side of this equation is a decreasing function of K. The right-hand side will be an increasing function of K, as long as $c''(K)$ is not too negative. Putting these facts together, we have Figure 1A.[10] Certainly equilibrium usage with a zero usage price, X^a, is larger than the equilibrium usage with a positive usage price, X^e. Decreasing X shifts the $nD'(X/K)X/K$ curve down, so equilibrium capacity with usage-based pricing will be less than the equilibrium capacity without usage-based pricing.

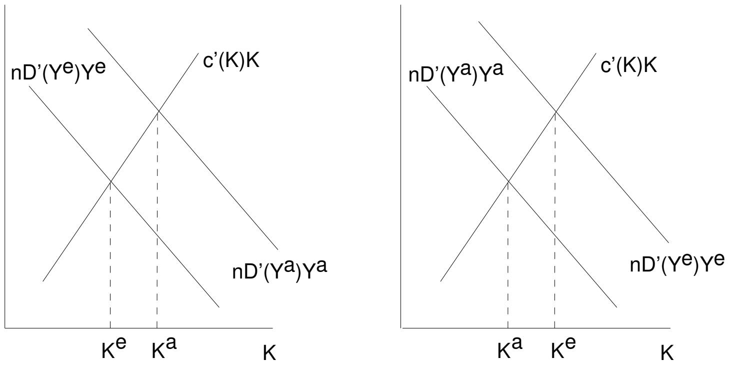

Figure 1. Determination of equilibrium capacity.

Will equilibrium congestion be higher or lower? With zero usage prices each user uses the resource more. But we have just shown that capacity will be higher, too, so it is not obvious what happens to utilization. Consider equation (4.16) again. Since $D(Y)$ is convex, $nD'(Y)Y$ is increasing in Y. If we write $K = X/Y$, it is easy to see that $c'(X/Y)X/Y$ is decreasing in Y as long as $c(K)$ is convex. Thus we can determine equilibrium congestion as in Figure 2. The increase from X^e to X^a causes $c'(X/Y)X/Y$ to move up, so with no usage pricing there is higher equilibrium congestion.

[10] The curves could be nonlinear; the straight lines are to simplify the presentation.

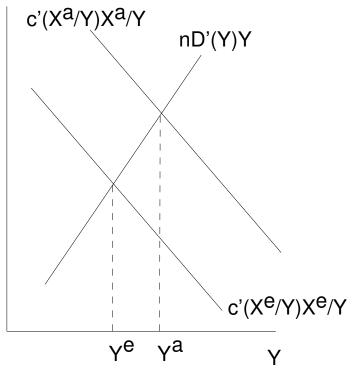

Figure 2. Determination of equilibrium congestion.

The number of users is different

Now we consider the case where the number of users changes. The equilibrium utility of a user without usage-based pricing is

$$v_i(x_i^a) - D(X^a/K^a) - q(Y).$$

This utility could be greater or less than the corresponding utility with usage-based pricing since there is more usage without prices, but there is also more congestion.

Suppose that there is some alternative service that provides the user with utility level u_i^*. Then voluntary participation requires that

$$v_i(x_i^a) - D(X^a/K^a) - q(Y^a) \geq u_i^*,$$

or,

$$v_i(x_i^a) - u_i^* \leq D(X^a/K^a) + q(Y^a).$$

That is, a user will stop using the network under access-only pricing if her net benefit from high usage is less than her congestion cost (including the access fee).[11]

Reducing the number of users will reduce $nD'(Y)Y$. This shifts *down* the corresponding curve in Figure 1B, and could result in an equilibrium amount of capacity that is *less* than one would have under usage-based

[11] We should note that there may also be users who do not consume a usage-priced resource, but *do* consume if there are only access prices. These would be users who want to generate a high volume of low-value traffic.

pricing.[12] One might call this a Yogi Berra equilibrium—after his famous remark that “it’s so crowded that no one goes there anymore.” In this case, however, the remark is apt: in this equilibrium there are a small number of intensive users with high tolerance for congestion, and therefore low willingness to pay for capacity expansion. The high-value users prefer to exit to alternative services.

4.5 Market power

What does utilization and capacity look like if there is market power? Suppose, for example, that a resource provider has a monopoly on the resource it provides: e.g., it is the only source for a certain kind of information. In this case it will typically have an incentive to restrict output in order to raise price. How does this affect its choice of optimal capacity?

If the provider prices only on the basis of usage, the answer is pretty straightforward. Generally output will be lower and price higher under a monopoly than under competition. Lower output means that the curve labelled $nD'(X/K)X/K$ will shift down in Figure 1, which implies less capacity.

However, this analysis is based on the assumption of usage pricing only. We have suggested that a combination of attachment and usage pricing would be a fairly common configuration for information and network service providers. The implications of such a two-part tariff are significant.

Identical tastes

For example, suppose that all users have the same tastes. In this case, the maximum connect fee that the monopolist can charge is the fee that makes the user indifferent between using the service and not using it. For simplicity, we normalize the utility of no use to zero, so the participation condition becomes

$$u(x, Y) - px - q = 0.$$

The profit maximization problem of the monopolist is

$$\max_{K,x} \; n[q + p(x, Y)x] - c(K).$$

Substituting from the participation condition we have

$$\max_{K,x} \; nu(x, Y) - c(K),$$

[12] Reducing n also reduces X^a (which is equal to the sum of satiation usage by all participating consumers), but the convexity of $D(Y)$ ensures that this indirect effect also works to shift $nD'(Y)Y$ downward.

which is just the problem of maximizing social welfare. It follows that the optimal policy of the monopolist is to set the use-price equal to the optimal congestion fee, charge the user $q = u(x^e, Y^e) - p^e x^e$ for usage, and make the socially optimal investment in capacity. This observation is the classic two-part tariff result of Oi (1971) See Schmalensee (1981) for a detailed exposition, and Varian (1989) for a survey of this and related results.

Different tastes

However, the assumption that all users—which really means all *potential* users—have identical tastes is rather unrealistic. Let us investigate the more realistic case of heterogeneous users. This case is well-treated in the literature on two-part tariffs cited above, but we need to see how it works for the congestion pricing problem we are examining here.

Let t be a parameter indexing tastes and write the utility function as $u(x, Y, t) = v(x, t) - D(Y)$. Let $f(t)$ be the density of type t and let $F(t)$ be the CDF. Choose the parameterization so that $u(x, Y, t)$ is decreasing in t.

The marginal consumer—the consumer who is just indifferent between using the service or not, denoted by T—is characterized by the condition

$$v(x, T) - D(Y) - q - px(p, T) = 0. \tag{4.17}$$

For any given p, the monopolist's choice of q is, effectively, a choice of the marginal consumer. Let $X(p, T)$ be the total demand of the consumers who use the service:

$$X(p, T) = \int_0^T x(p, t) f(t)\, dt.$$

The profit maximization problem of the monopolist is

$$\max_{T,p,K} \; qF(T) + pX(p, T) - c(K), \tag{4.18}$$

where q is defined in (4.17). Substituting, we have

$$\max_{T,p,K} [v(x, T) - D(Y)]F(T) + p[X(p, T) - x(p, T)F(T)] - c(K).$$

It is worth observing that if the demand of marginal consumer equals the demand of the average consumer, the bracketed term in the middle cancels out and we are back in the previous case.

The first-order conditions for p and K are

$$\begin{aligned} \left[\frac{\partial v}{\partial x}\frac{\partial x}{\partial p} - D'(Y)\frac{\partial X/\partial p}{K}\right] F(T) + p\left[\frac{\partial X}{\partial p} - \frac{\partial x}{\partial p}F(T)\right] + \\ [X(p, T) - x(p, T)F(T)] = 0 \\ D'(Y)YF(T) = c'(K)K. \end{aligned} \tag{4.19}$$

Define the elasticity of demand of the serviced customers as

$$\epsilon = -\frac{p}{X(p,T)}\frac{\partial X(p,T)}{\partial p}.$$

After some manipulations we can write the first-order condition as

$$\frac{p - c'(K)/Y}{p}\epsilon = 1 - \frac{x(p,T)}{X(p,T)/F(T)}.$$

The last term on the right-hand side is the ratio of the demand of the marginal consumer to the demand of the average consumer. If all consumers have the same tastes, then this fraction is 1, and we find that pricing at marginal congestion cost is optimal, as we have already observed. The interesting cases are when the marginal and the average consumer have different tastes.

Recall that by construction the marginal consumer has a lower total value for a given level of usage than the average consumer. Normally, one would think that a consumer with lower total value would want to consume less than a consumer with higher total value. In this case, the monopolist who uses a two-part tariff would set price higher than marginal congestion cost. However, if the marginal consumer wants to consume *more* than the average consumer, it is quite possible that the monopolist would want to set the price *lower* than marginal congestion cost. This is the "auto salesman equilibrium"—the monopolist prices the service so low that he loses money on every sale but makes up for it in volume! Unlike an automobile dealer, our monopolist really may be willing to do this since it collects a uniform subscription fee from each user which can make up for the lost usage fees.

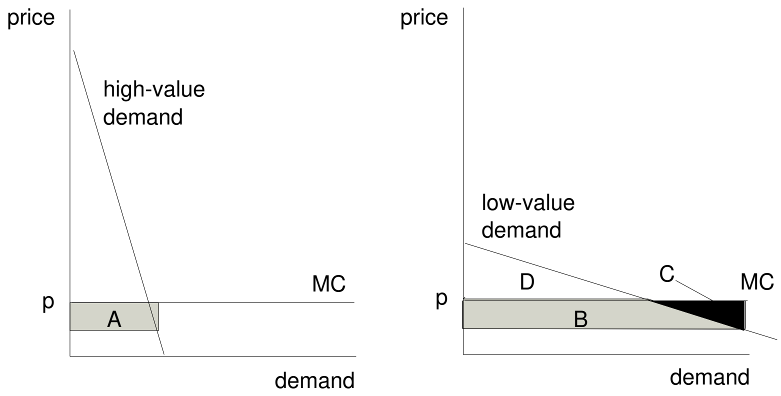

Figure 3. Pricing less than marginal cost may be optimal.

To see how this can happen consider Figure 3, which is based on Oi (1971). There are two users. One has a very high value for the service, but only wants to use a little of it. (Think of ASCII email.) The other user has a low willingness-to-pay for the service but wants to consume a very large amount of it. (Think of a teenager downloading MTV videos.) The teenager is the marginal user, and the connection fee—which is paid by both users—reflects his (relatively low) valuation.

For simplicity, we take the marginal cost of congestion to be constant. Suppose initially that the monopolist prices at the marginal congestion cost; we will show that under some circumstances monopoly profits will increase if the monopolist reduces its price.

If the monopolist sets price equal to marginal cost, the low-value user will achieve net consumer surplus of area D, while the high-value user achieves consumer surplus that is larger than D. The monopolist can therefore charge *each* of them a connection fee of D yielding profits of $2D$.

Now suppose that monopolist reduces its price to some amount *below* marginal cost. The monopolist can now increase the connection fee to $2(D + B)$. However, costs increase as well due to the increased use by both parties. The high-value user imposes additional costs of A and the low-value user imposes additional costs of $(B + C)$. The net increase in profits is $2B - A - (B + C) = B - A - C$. This area may easily be positive, as it is in the case illustrated.

The teenager's utility is larger since he can now download more videos, so he is willing to pay more for the connection; the monopolist extracts this additional surplus through the increased connection charge B. Although the teenager's utility increase (B) is less than the reduction in usage revenues $(B+C)$, the email user *also* has to pay the subscription increase (B). In addition the email user will impose costs on the monopolist of an amount A due to his additional usage. Hence if the subscription increase from the high-value user is greater than the usage revenue losses $(B - A - C > 0)$, profits will increase when price is set below the marginal congestion cost.

This is the same effect observed by Oi (1971) in his classic article. In the literature it is commonly regarded as a perverse effect that is unlikely to occur in reality. But in our context this effect appears to be quite plausible: it can easily happen that relatively low-valued services can require a huge amounts of bandwidth. In order to capture revenues from such uses, the monopolist may find it profitable to *underprice* the congestion they create, thereby imposing potentially significant congestion costs on high-value, low-bandwidth users.

4.6 Summary

We have argued that many network resources are congestible: that is, they can be used by more than one person but increasing usage degrades their quality. One person's use creates an externality: it lowers the value of usage for everyone else. Economists long have proposed pricing to internalize this externality: such a price should reflect both the direct and external costs of usage, so that consumers will use the resource efficiently.

In this paper we have developed this theory for a model of the type of congestible resources typically found in an information network. We found that if the resource is provided in a competitive market with connect fees and usage prices, the equilibrium price and capacity will maximize net social benefits. If there is a monopoly provider, however, the profit-maximizing usage price could be either higher or lower than the socially optimal price (with offsetting adjustments in the connection fee), depending on the value that different users put on the resource.

The extent to which the market is competitive ultimately depends on the cost structure of providing the network resource. Whether a given provider will offer a single or multiple qualities of service will depend both on the cost structure and the extent to which an individual user has preferences for multiple qualities of service.

Currently, the most common form of Internet pricing is pricing by access, with no usage-sensitive prices. With a fixed set of users, we expect to see greater capacity when usage is not priced, but also greater congestion. However, with greater congestion, congestion-sensitive users might not use the resource; the resulting "Yogi Berra" equilibrium might actually have lower usage (but higher congestion) than when usage is priced.

References

Bohn, R., Braun, H.-W., Claffy, K., & Wolff, S. (1993). Mitigating the coming Internet crunch: Multiple service levels via precedence. Tech. rep., UCSD, San Diego Supercomputer Center, and NSF.

Braden, R., Clark, D., & Shenker, S. (1994). Integrated services in the Internet architecture: an overview. Tech. rep., IETF. RFC 1633.

Buchanan, J. (1965). An economic theory of clubs. *Economica*, *32*, 1–14.

Cornes, R., & Sandler, T. (1986). *The Theory of Externalities, Public Goods, and Club Goods.* Cambridge University Press, Cambridge.

Economides, N. (1994). Critical mass and network size. Tech. rep., New York University Stern School of Business, New York.

Hardin, G. (1968). The tragedy of the commons. *Science*, *xx*, 1243–47.

Katz, M., & Shapiro, C. (1985). Network externalities, competition and compatibility. *American Economic Review*, *75*, 424–440.

MacKie-Mason, J. K., & Varian, H. (1993). Some economics of the Internet. Tech. rep., University of Michigan.

MacKie-Mason, J. K., & Varian, H. (1994a). Pricing the Internet. In Kahin, B., & Keller, J. (Eds.), *Public Access to the Internet*. Prentice-Hall, Englewood Cliffs, New Jersey.

MacKie-Mason, J. K., & Varian, H. R. (1994b). Economic FAQs about the Internet. *Journal of Economic Perspectives*, *8*(3).

Mohring, H., & Hartwiz, M. (1962). *Highway Benefits: An Analytical Approach*. Northwestern University Press, Evanston.

Oi, W. (1971). A Disneyland dilemma: two-part tariffs for a Mickey Mouse monopoly. *Quarterly Journal of Economics*, *85*, 77–96.

Schmalensee, R. (1981). Monopolistic two-part pricing arrangements. *Bell Journal of Economics*, *12*, 445–466.

Scotchmer, S. (1985a). Profit-maximizing clubs. *Journal of Public Economics*, *27*, 25–45.

Scotchmer, S. (1985b). Two-tier pricing of shared facilities in a free-entry equilibrium. *Rand Journal of Economics*, *16*(4), 456–472.

Sharkey, W. W. (1982). *The Theory of Natural Monopoly*. Cambridge University Press, New York.

Strotz, R. H. (1965). Urban transportation parables. In Margolis, J. (Ed.), *The Public Economy of Urban Communities*, pp. 127–169. Resources for the Future, Washington, D.C.

Varian, H. R. (1989). Price discrimination. In Schmalensee, R., & Willig, R. D. (Eds.), *Handbook of Industrial Organization*, Vol. I of *Handbooks in Economics*, chap. 10, pp. 597–654. Elsevier Science Publishing, New York.

Varian, H. R. (1992). *Microeconomic Analysis*. W. W. Norton & Co., New York.