Decentralization, Duplication, and Delay

Patrick Bolton; Joseph Farrell

Stable URL:
http://links.jstor.org/sici?sici=0022-3808%28199008%2998%3A4%3C803%3ADDAD%3E2.0.CO%3B2-6

The Journal of Political Economy is currently published by The University of Chicago Press.

Decentralization, Duplication, and Delay

Patrick Bolton

Ecole Polytechnique

Joseph Farrell

University of California, Berkeley

We argue that although decentralization has advantages in finding low-cost solutions, these advantages are accompanied by coordination problems, which lead to delay or duplication of effort or both. Consequently, decentralization is desirable when there is little urgency or a great deal of private information, but it is strictly undesirable in urgent problems when private information is less important. We also examine the effect of large numbers and find that coordination problems disappear in the limit if distributions are common knowledge.

I. Introduction

Most economists believe that decentralized economic systems are more efficient than centralized ones. But the theoretical basis for this belief is not entirely clear. Certain theoretical results—notably the welfare theorem and the Coase theorem—state that, in idealized models, decentralized outcomes are Pareto efficient. But as Lange, Lerner, and others noted long ago, even if we accept the assumptions required for these results, this alone cannot justify a preference for decentralization since centralization might be equally efficient.

In response, Hayek (1945) argued that the missing element is an understanding of how decentralized and centralized systems deal with what we now call private information. He suggested that, in informational terms, an economic system must either put "at the disposal of a single central authority all the knowledge which ought to be used but is initially dispersed among many different individuals, or [convey] to the individuals such knowledge as they need in order to enable them to fit their plans in with those of others" (p. 521). Hayek argued that centralization does the former poorly and that a decentralized market system does the latter well, notably through prices.

As we know, however, the price system does not always work perfectly. For instance, prices may fail to reflect some "external" costs. As Pigou pointed out, this deficiency may in theory be repaired by central intervention in the form of corrective taxes. But if the central authority does not have the (considerable) information needed to calculate the right corrective tax, a clumsy intervention may do more harm than good, and it may be better on balance to put up with the externality. Whether this is so depends, of course, on just how ignorant the central authority is, on how harmful the externality is, and on other factors.

Another, much less studied, problem that a laissez-faire system may not solve perfectly is *coordination*. Quite different from ordinary externality or monopoly problems, coordination problems involve more than one relatively desirable (pure-strategy) equilibrium, and other outcomes are undesirable. Examples include Schelling's (1960) well-known where-to-meet puzzles, problems of the choice of compatibility standards, and the problem faced by two potential entrants into a natural monopoly. In these problems, prices alone typically do not guide actions well. For instance, in the entry problem, the desirable outcomes involve one firm entering and the other not. Because this is an asymmetric outcome of an initially symmetric problem, anonymous market prices cannot suffice.

As with externalities, simple complete-information models would suggest that central intervention could easily solve these problems. Indeed, in this case, no taxes or coercion would be needed: merely "indicative planning," in which the planner makes suggestions without coercive power.[1] But just as we should not take Pigou's analysis as proving that central intervention is necessarily desirable, so here, if a central planner lacks important private information (for instance, about the relative merits of the equilibria), intervention may be unde-

---

[1] For example, this is true in the coordination problems analyzed by Schelling (1960), Dixit and Shapiro (1986), Farrell (1987), and Farrell and Saloner (1988).

sirable. In this paper, we study an example of a coordination problem with private information and show how the desirability of uninformed central coordination varies according to the importance of private information and the importance of coordination.

In our example, two firms contemplate sinking costs in order to enter a natural-monopoly market. We suppose that these costs are private information, so that neither firm (nor anyone else) knows which has the lowest costs. The best outcome is that the lower-cost firm enters and the other stays out. "Decentralization" in our example is an incomplete-information game of timing between the firms: In any period, if neither firm has already sunk the costs of entry, each decides whether to sink costs ("enter") or to wait another period. Waiting is sensible if the firm fears that the other is likely to enter.

In a symmetric equilibrium of this game, each firm is uncertain about when the other will enter. Lower-cost firms are less worried about the possibility that their rival will enter and therefore choose to enter sooner than higher-cost firms do; consequently, if the firms have (sufficiently) different costs, then the lower-cost firm enters and preempts its higher-cost rival. Thus the laissez-faire system effectively "uses" private information, even in this model without prices. But this decentralized coordination is far from perfect. If the firms have equal costs, then they enter simultaneously: there is inefficient *duplication*. If both have high costs, both wait: there is inefficient *delay*.

In our example, our model of "central intervention" consists of nominating one of the two firms to enter; in the spirit of Hayek, we suppose that (for reasons we do not model) a central planner cannot or will not collect and use the private information in the way that mechanism design theory might suggest. In other words, we model a clumsy central solution to the coordination problem. We compare the expected social cost (relative to the first-best) of the fact that half the time, the uninformed central planner will nominate the higher-cost firm to enter, with the cost of duplication and delay incurred under decentralization.

Not surprisingly, we find that the less important the private information that the planner lacks and the more essential coordination is, the more attractive the central planning solution is. Perhaps more interesting is the effect of urgency. The decentralized solution, since it requires time to work, performs poorly if time is pressing. Centralization—in our model—does not involve delay and consequently is a good system for "emergencies," when delay is very costly. This prediction of our model is consistent with the fact that almost all organizations rely on central direction in emergencies. In particular, in warfare, laissez-faire decentralization is little favored: armies work by

command. Studying economic organization in World War II, Mil-
ward (1977, p. 99) found that

> the differences between the political societies involved in the
> Second World War were very wide. Once commitment to a
> war involving a high level of economic effort was accepted,
> however, the economic problems to be solved were often
> very similar. A comparative study of the machinery by which
> these economic decisions were taken and enforced indicates
> that common decisions impose to some extent common ad-
> ministrative solutions. . . . *Everywhere the price mechanism came
> to be regarded as a method of allocating resources which was too slow
> and too risky.* [Italics added]

Similarly, Scitovsky, Shaw, and Tarshis, in their study of economic
mobilization for war (1951, p. 139), found that the market system
"takes time—sometimes a considerable length of time—in effecting
adjustments to changes in supply and demand conditions."

In this paper we use an example to address two questions that we
think receive too little attention from economic theorists. The first
concerns adjustment costs: How do different mechanisms perform
while approaching equilibrium? Our example illustrates some of the
inefficiencies that may arise as a decentralized system gropes for an
efficient static "equilibrium" in which high-cost firms are "out" while
low-cost firms are "in." In particular, the time required for decen-
tralized screening may be costly. The second question concerns *coordi-
nation:* If agents simply do not know what others are expecting them
to do, then coordination problems arise, and these are not well cap-
tured by standard notions of externalities and incentives (which de-
scribe why agents may knowingly act in a way inconsistent with
efficiency). As our example suggests, central direction may resolve
such coordination problems, but probably imperfectly.

II. A Model of Decentralized Coordination
in Entry

We consider the following simple coordination problem. Two firms,
A and B, are contemplating entering a natural-monopoly industry.
Entry involves a sunk cost, S, which is borne by the entrant. A single
entrant will capture a fraction λ of the gross social benefits of entry,
V, which we normalize to be one, so that a sole entrant's net payoff is
$\lambda - S$. We suppose that each firm's value of S is known only to the
firm itself but is independently drawn from a known distribution $F(\cdot)$.
If both firms enter, each bears its cost S, and each can appropriate a

fraction μ of the social gross benefit $V = 1$.[2] (For instance, in the case of entry into a natural-monopoly industry, with Bertrand competition should both firms enter, we have $0 < \lambda \leq 1$ and $\mu = 0$.) We suppose that $\mu < S < \lambda$ for all S in the support of $F(\cdot)$, so that this is a coordination problem.[3] Each firm wants to enter if and only if the other does not enter. There are thus two pure-strategy equilibria; other outcomes are inefficient. The pure-strategy equilibria are typically not equally desirable, however: it is better to have the lower-cost firm enter.

We take "decentralization" here to mean that, at any date, each firm chooses to enter or not, with no explicit consultation with the other. If neither enters, then once each sees that the other has not entered, the game begins again (although with different beliefs). It is natural to suppose that there is a lag between a firm's irreversible commitment to sink the entry costs S and the unambiguous observation of that commitment by the other firm. We model this by supposing that there are infinitely many "periods" 1, 2, . . . and that in each period each firm observes what happened in the previous period and then chooses either to enter or to wait another period. Once entry by at least one firm has occurred, production takes place and the game ends. Benefits are discounted at a constant rate δ per period, both privately and socially.[4]

In this game, there are asymmetric subgame-perfect equilibria that "solve" this coordination problem: for example, firm 1 always enters immediately and firm 2 never enters. Much work on decentralization has focused on the constraints imposed by incentive compatibility, and if we were to take that approach, we would find that decentralization was consistent with these asymmetric equilibria. But it is far from clear how the firms would "find" one of those equilibria: we cannot assume that they do so without some process of explicit coordination. Moreover, those equilibria do not capture the informational benefits of decentralization since the private information on costs does not affect which firm enters. In analyzing decentralization, therefore, we focus on symmetric (Bayesian) equilibrium.[5]

[2] If competition ex post makes for a more efficient exploitation of the social opportunity modeled, then the *gross* social gain from the entry of both firms may exceed $V = 1$. But unless it exceeds $1 + S$, duplication is still socially undesirable.

[3] This is assumed for ease of exposition only; in an earlier draft, we considered the more general case in which some values of S may lie above λ or below μ.

[4] It would be possible in principle to model these delays within a continuous-time model; however, we suspect that it would be technically more complex and would yield little additional insight.

[5] Dixit and Shapiro (1986), Farrell (1987), and Farrell and Saloner (1988) argued informally that asymmetric pure-strategy equilibria are unconvincing and inappropriate for the study of decentralized coordination mechanisms. See Crawford and Haller (1990) for a more formal argument.

TABLE 1

FIRM A	FIRM B	
	Enter	Wait
Enter	$\mu - S_A, \mu - S_B$	$\lambda - S_A, 0$
Wait	$0, \lambda - S_B$	$\delta v(S_A), \delta v(S_B)$

As long as no entry has occurred, the payoff matrix in any period can be written as shown in table 1, where $v(S)$ is the continuation value of the game to a player of type S. We now show (i) that, in equilibrium, firms with lower values of S enter sooner (in a sense that we make precise) and (ii) that the process of entry has a declining hazard rate. Proposition 1, which fulfills claim i, is a statement about admissible strategies for a single firm; our other results assume a symmetric Bayesian equilibrium.

PROPOSITION 1. Low-cost types enter sooner than high-cost types do, in the following sense. Fix (say) firm A's beliefs about the (normal-form) strategy of firm B, and consider two possible values of S_A, say S_A^1 and S_A^2. Suppose that $S_A^1 < S_A^2$, that S_A^1 puts positive probability on entering in period t_1, and that S_A^2 puts positive probability on entering in period t_2. Then $t_1 \leq t_2$.

Proof. A normal-form strategy for firm A is a mixture of the elementary strategies $s(t)$: enter in period t if B has not entered before that. If firm A puts positive probability on $s(t)$, then its expected payoff from $s(t)$ must be at least as great as that from $s(t')$ for any other t'. Now suppose that A believes that there is a probability $\alpha(t)$ that B will not have entered before time t and a "hazard-rate" probability $h(t)$ that B will enter precisely at time t (if not preempted). Write $a(t) \equiv \delta^t \alpha(t)$. Then A's expected payoff from the strategy $s(t)$ is just

$$a(t)[\lambda - S - h(t)(\lambda - \mu)].$$

If the type S_A^i puts positive weight on strategy $s(t_i)$, then we have

$$a(t_1)[\lambda - S_A^1 - h(t_1)(\lambda - \mu)] \geq a(t_2)[\lambda - S_A^1 - h(t_2)(\lambda - \mu)]$$

and

$$a(t_2)[\lambda - S_A^2 - h(t_2)(\lambda - \mu)] \geq a(t_1)[\lambda - S_A^2 - h(t_1)(\lambda - \mu)].$$

Adding, we get

$$[a(t_1) - a(t_2)](S_A^1 - S_A^2) \leq 0.$$

Now $a(t + 1)/a(t) = \delta[1 - h(t)]$, so that $a(t)$ is strictly decreasing in t. Since $S_A^1 < S_A^2$ by hypothesis, this gives the result. Q.E.D.

COROLLARY A. In a symmetric Bayesian equilibrium, if just one firm enters, it is the lower-cost firm.

COROLLARY B. In a symmetric Bayesian equilibrium, if the firms' costs in fact differ, the lower-cost firm surely enters; the higher-cost firm may or may not enter.

Proposition 1 and its corollaries show how decentralization *sorts* potential entrants: if duplication is avoided, low-cost rather than high-cost firms enter. But this sorting is achieved only with delay and only through the threat of duplication, which makes high-cost firms hang back and allow low-cost firms to go first. And, of course, this fear is sometimes realized; the likelihood of that is determined by the firms' strategies. We turn next to a description of those strategies. From now on, we assume a symmetric Bayesian equilibrium.

The Fundamental Difference Equation

By proposition 1, there exist cutoffs $S_1, S_2, \ldots$ such that, for values of S between S_{t-1} and S_t, a firm of type S enters in period t (provided that its rival has not previously entered). Since the firm of type S_t is indifferent between entering in period t and waiting for period $t + 1$, we have

$$\lambda - S_t - h(t)(\lambda - \mu) = \delta[1 - h(t)][\lambda - S_t - h(t + 1)(\lambda - \mu)]. \quad (1)$$

Equation (1) is a difference equation relating S_t, S_{t-1}, and S_{t+1}. It describes the dynamic behavior of the decentralized equilibrium. Our next proposition summarizes this behavior.

PROPOSITION 2. The hazard rate $h(t)$ is nonincreasing. As $t \to \infty$, the cutoff type, S_t, converges to the upper limit $S^{\max}$ of the support of $F(\cdot)$. The fraction $F(S_t)$ of all types that have entered by date t converges to one but is bounded above by $1 - [1 - h(1)]^t$.

Proof. First, note that each side of (1) must be nonnegative since it represents the equilibrium expected payoff to a firm of type S_t, and the firm can always guarantee zero payoff by staying out. This nonnegativity implies (by [1]) that the hazard rate is nonincreasing. The upper bound on the fraction of firms that have entered by a given date t follows immediately.

If S_t's equilibrium payoff is zero, then (1) implies that the hazard rate $h(\tau)$ is constant from date t onward and that $S_\tau = S_t$ for all $\tau \geq t$. Since $S_t < \lambda$, (1) implies that the (constant) hazard rate $h(\tau)$ must be strictly positive. This implies that the fraction $F(S_\tau)$ of all types that have entered by date τ converges to one as $\tau \to \infty$, and $\lim_{\tau \to \infty} S_\tau = S^{\max}$.

If the equilibrium payoff expressed by (1) is strictly positive, we consider two subcases. First, if $h(t)$ converges to zero, then every type of firm can ensure a strictly positive expected payoff by waiting long

enough before entering; hence every type eventually enters (if not preempted) and S_t converges to $S^{\max}$. Second, if $h(t)$ converges to a strictly positive $\bar{h} > 0$, then (as with a positive constant hazard rate above) $\lim_{\tau \to \infty} S_\tau = S^{\max}$. Q.E.D.

Intuitively, if there is a firm willing to enter, then entry will eventually occur, but it can take an arbitrarily long time. So decentralization achieves some sorting, but it is typically imperfect and takes time. To say more, we must simplify the model still further; we shall consider the case in which the distribution $F(\cdot)$ has a two-point support.

III. Two Types

We consider the simple case in which there are just two cost types, "high" (S_H) and "low" (S_L), where $\mu < S_L < S_H < \lambda$. Let the probability of S_L be q. The expected social payoff in the first-best is then simply

$$W^* = 1 - (1 - q)^2 S_H - q(2 - q)S_L.$$

We now consider symmetric Bayesian equilibrium of the laissez-faire entry game.

By proposition 1, there is a date T such that low-cost types enter at or before T and high-cost types enter at or after T. For some parameter values there will be a mixture of the two types entering at T, and the sorting process will be incomplete; that is, sometimes a high-cost firm will enter (at T) although its rival is low-cost. For other parameter values, there will be complete separation; that is, all low-cost types enter at or before T and all high-cost types enter strictly after T. In such a fully separating equilibrium, decentralization always finds the low-cost entrant in the sense that, if the two firms differ in costs, then the low-cost firm is the unique entrant. Thus decentralization is shown most favorably in such an equilibrium. It is simplest to focus on the case in which $T = 1$, so that when there is a low-cost entrant there is no delay.

A Simple Equilibrium with Two Types

For a range of parameter values and for a suitable value of p, it is a perfect Bayesian equilibrium for both firms to use the following Bayesian strategy:

> If low-cost, then enter immediately. If high-cost, then do not enter in period 1; in each period $t \geq 2$, if there has been no entry before t, enter with probability p.

It is easy to show that this "simple equilibrium" exists if and only if

$$\mu \leq S_L + \delta(1 - q)(S_H - S_L) \leq \lambda - q(\lambda - \mu) \leq S_H \leq \lambda. \quad (2)$$

In such an equilibrium, the outcome of decentralization is as follows. If both firms are low-cost (which happens with probability q^2), then there is immediate duplication; welfare is $1 - 2S_L$. If one is low-cost and the other is high-cost (which happens with probability $2q[1 - q]$), then there is immediate low-cost entry and no duplication; welfare is $1 - S_L$. If both firms are high-cost (probability $[1 - q]^2$), then the first period sees no entry and the two firms continue under complete information, as we now analyze.

Entry Game with Complete Information

Since the symmetric equilibrium involves mixed strategies, each firm must be indifferent between entering and waiting:

$$v = p(\mu - S) + (1 - p)(\mu - S) = (1 - p)\delta v,$$

where p is the probability of entry (for a given firm) in any given period, and v is the continuation value of the game. Since $(1 - p)\delta < 1$, $v = 0$ and

$$p = \frac{\lambda - S}{\lambda - \mu}. \tag{3}$$

Notice that we require $\mu < S < \lambda$: mathematically so that $0 < p < 1$ and economically so that neither enter nor wait is a dominant strategy.

We now calculate the expected social surplus. The probability q_t that nothing happens before period t and that then just one firm enters is

$$q_t = (1 - p)^{2(t-1)} \times 2p(1 - p),$$

while the probability r_t that both do so is

$$r_t = (1 - p)^{2(t-1)} \times p^2.$$

Consequently, expected social surplus,

$$W^G \equiv \sum_{t=1}^{\infty} \delta^{t-1}[q_t(1 - S) + r_t(1 - 2S)],$$

is given by

$$W^G = \frac{p(2 - p)}{1 - \delta(1 - p)^2} - \frac{2p}{1 - \delta(1 - p)^2}S \tag{4}$$

$$= (1 - S) - \left[1 - \frac{p(2 - p)}{1 - \delta(1 - p)^2}\right](1 - S) - \frac{p^2}{1 - \delta(1 - p)^2}S. \tag{5}$$

In equation (5), we see the losses (relative to the first-best welfare level $1 - S$) due to delay (the second term) and due to duplication (the third term). As $p \to 1$, the losses due to delay converge to zero, as they do when $\delta \to 1$. But large values of p make the losses due to duplication worse, and large values of δ do not help there. We can interpret $\delta \to 1$ as representing either a less urgent problem or else shorter periods (less reaction lag).[6]

We can ask whether the equilibrium value of p, given by equation (3), is above or below the (second-best) optimal level. That is, would a decentralized policy (constrained to symmetric equilibrium) do better by inducing earlier or later entry on average? Such changes could be accomplished, for example, by changing λ or μ or S. Simple calculations show that in general the answer is ambiguous. Firms bear all the costs of duplication and only some of the costs of delay ($\lambda < 1$); this might make one think that they would be too cautious. But at the same time there is an inefficient preemption motive for each firm to move too fast: it does not take proper account of the reduction in its rival's expected payoff when it enters. Formally, we have

$$[1 - \delta(1 - p)^2]W^G = p(2 - p - 2S).$$

Hence, the sign of $\partial W^G / \partial p$ is that of

$$(1 - p - S)[1 - \delta(1 - p)^2] - \delta p(1 - p)(2 - p - 2S).$$

In particular, for small δ, $\partial W^G / \partial p$ has the sign of

$$1 - p - S = \frac{(1 - \lambda + \mu)S - \mu}{\lambda - \mu},$$

so when a problem is very urgent the decentralized solution may be too hasty or too cautious, according to whether $S \gtrless \mu/(1 - \lambda + \mu)$. For large δ, $\partial W^G / \partial p < 0$ always: that is, for nonurgent problems, the preemptive effect dominates and firms are always too hasty. Turning to the effects of different sizes of S, we find that for small S (close to μ) there is excessive haste, while for large S (close to λ) there is excessive caution.

The probability of (eventual) coordination (i.e., just one firm enters) is

$$\frac{2p(1 - p)}{2p(1 - p) + p^2} = \frac{2 - 2p}{2 - p} = \frac{2(S - \mu)}{\lambda + S - 2\mu},$$

---

[6] Indeed, if we analyzed the problem in continuous time following Simon (1987), we would find instant entry but with a positive probability of duplication. But we believe that such a continuous-time model would be inappropriate since it would assume that players can react instantly to one another's choices. In reality, there are typically observation and decision lags; we believe that these are better analyzed in a model such as ours with discrete periods, although in principle one could build observation and decision lags into a continuous-time model.
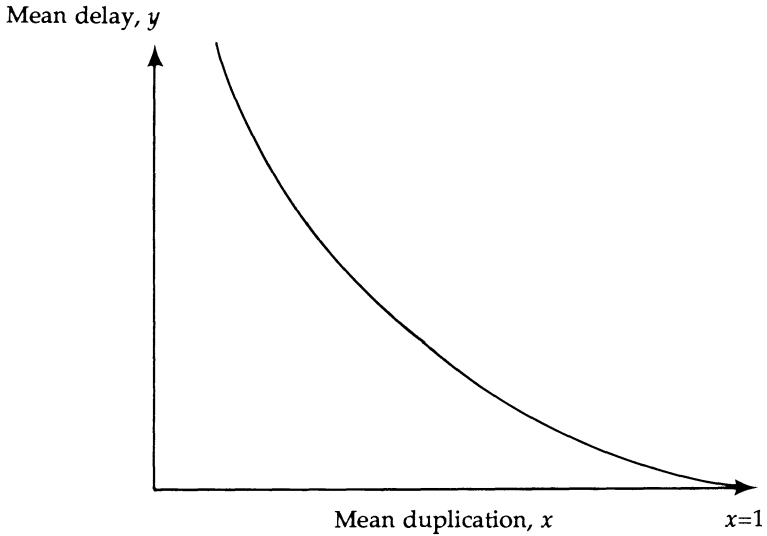
Mean delay, y

Mean duplication, x $x=1$

FIG. 1

which is increasing in S, decreasing in λ, and decreasing in μ, as one might expect: eventual coordination is more likely when firms are more cautious about sinking costs. Such parameters, however, simultaneously exacerbate the delay in reaching an outcome. *There is a trade-off between avoiding duplication and avoiding delay.* Indeed, we can see that trade-off quite clearly by expressing the probability of duplication and the mean delay in terms of the single variable p: the probability of duplication is just

$$x = \frac{p}{2 - p}, \tag{6}$$

while the mean delay is

$$y = \sum_{t=1}^{\infty} (q_t + r_t)t - 1 = \frac{1}{p(2 - p)} - 1. \tag{7}$$

Eliminating p from (6) and (7), we find that

$$y = \frac{(x - 1)^2}{4x}. \tag{8}$$

Policies such as subsidies for entry, which operate through changing p, move us along the curve (8), which is shown in figure 1.

This completes the analysis of the complete-information continuation game. We now return to the analysis of the simple Bayesian perfect equilibrium in the two-type case. Once the first period has

passed without entry, so that it is common knowledge that both firms have $S = S_H$, welfare is given by (4), or

$$W^G = \frac{p(2 - p)}{1 - \delta(1 - p)^2} - \frac{2p}{1 - \delta(1 - p)^2} S_H,$$

and p is given by (3) with $S = S_H$. Consequently, the ex ante expected social payoff from decentralization is

$$W^D = q^2(1 - 2S_L) + 2q(1 - q)(1 - S_L)$$

$$+ (1 - q)^2 \delta \left[\frac{p(2 - p)}{1 - \delta(1 - p)^2} - \frac{2p}{1 - \delta(1 - p)^2} S_H \right],$$

where $p = (\lambda - S_H)/(\lambda - \mu)$. Equivalently we can write the loss relative to the first-best, W^*, as

$$W^* - W^D = \left[q^2 S_L + (1 - q)^2 \frac{\delta p^2}{1 - \delta(1 - p)^2} S_H \right]$$

$$+ (1 - q)^2 \left[1 - \frac{\delta p(2 - p)}{1 - \delta(1 - p)^2} (1 - S_H) \right].$$

Here we see the losses due to duplication and to delay. We next compare these with the losses from a natural model of a highly imperfect central planning system.

Central Planning with Incomplete Information

While decentralized equilibrium in our model often selects efficient entrants, it involves duplication and delay. One of our central points in this paper is that, for some parameter values, *even a completely ignorant central planner outperforms decentralization* because speed and reliability of coordination outweigh the need to find the lower-cost entrant. In other cases, of course, the comparison is reversed, for under decentralization the lower-cost firm always enters (although the higher-cost firm may do so as well!). Evidently, the (gross) advantage of decentralization is related to the difference $E[S] - E[\min(S_A, S_B)]$, which is small when the distribution of S is concentrated or when the two firms' costs, S_A and S_B, are highly correlated.

Suppose then that a central planner does not use private information at all, but merely chooses an entrant at random. This will achieve immediate entry without duplication, but there will be no tendency to select a low-cost entrant: S will on average take its expected value.

In the two-type case, expected welfare under a random planner is simply

$$W^R = 1 - [q S_L + (1 - q) S_H].$$

Equivalently, we can write the loss relative to the first-best from random choice as

$$W^* - W^R = q(1 - q)(S_H - S_L). \tag{9}$$

Comparing Centralized and Decentralized Outcomes

Comparing these expressions for expected welfare is complex in general. Since we are considering only an example in any case, we shall make the comparison only for the polar cases $\delta = 0$ and $\delta = 1$.

For very urgent problems, when $\delta = 0$, decentralization yields an expected payoff of

$$2q(1 - q)(1 - S_L) + q^2(1 - 2S_L) \tag{10}$$

in the simple equilibrium, which exists when $\mu \leq S_L \leq q\mu + (1 - q)\lambda \leq S_H \leq \lambda$.[7] We can rephrase (10) to show how decentralization fails to achieve the first-best: the difference between the first-best expected payoff and (10) is

$$W^* - W^D = q^2 S_L + (1 - q)^2(1 - S_H), \tag{11}$$

where the first term is the loss due to duplication and the second term is the loss due to delay. Comparing (11) with (9), we find the following proposition.

PROPOSITION 3. In the two-type case, when there is great urgency ($\delta \approx 0$) and conditions (2) are satisfied, decentralization is superior to random choice if and only if

$$q(1 - q)(S_H - S_L) > q^2 S_L + (1 - q)^2(1 - S_H). \tag{12}$$

Evidently, (12) is less likely to hold when S_H and S_L are close, that is, when the private information that decentralization may be able to use is relatively unimportant. Likewise, it is more likely to hold when $1 - S_H$ is large (so that delay is very costly) and when S_L is large (duplication is costly).

Turning to the other polar case, $\delta = 1$, we find that decentralization gives an expected social payoff of

$$W^D = 2q(1 - q)(1 - S_L) + q^2(1 - 2S_L) + (1 - q)^2 \frac{p(2 - p - 2pS)}{1 - (1 - p)^2}.$$

As before, it is useful to write this in terms of deviation from the first-best, that is,

$$W^* - W^D = q^2 S_L + (1 - q)^2 \frac{\lambda - S_H}{\lambda + S_H - 2\mu} (1 - S_H). \tag{13}$$

Comparing (13) with (9), we find the following proposition.

---

[7] Notice that, for any choice of S_L, S_H, and q, these conditions hold for suitable values of λ and μ. That is, we can choose the latter arbitrarily and still find genuine examples if we are prepared to choose λ and μ accordingly.

PROPOSITION 4. When there is little urgency ($\delta \approx 1$), decentralization is better than random choice if and only if

$$q(1 - q)(S_H - S_L) > q^2 S_L + (1 - q)^2 \, \frac{\lambda - S_H}{\lambda + S_H - 2\mu} \, (1 - S_H). \quad (14)$$

Since $S_H > \mu$, comparison of (12) and (14) shows that decentralization is more likely to be the better of these two systems when there is little urgency. We thus find some confirmation of the idea that decentralization is less desirable when a problem is urgent, as in emergencies.

Speed and Delay in Bureaucracy

It may seem surprising that our model depicts central bureaucracy as an insensitive but rapid decision maker. Most people would probably describe bureaucratic central decisions as slow, and often they are; but the fastest decisions are also made by central authorities.

While a central authority surely *can* make a rapid and random decision, it may not always do so. Instead, a planner may try to gather information. One natural way to do so is through an advocacy process of some kind, which we might model as a war of attrition between potential entrants. This conforms to some descriptions of the Japanese Ministry of International Trade and Industry, which tries to coordinate certain entry and major investment decisions, in an environment in which each of several firms would like to be selected. A similar kind of voluntary centralization in coordination was studied by Farrell and Saloner (1988), in the context of the choice of compatibility standards. They found that a centralized committee system is more reliable at coordinating than the "market" system, but is indeed slower. However, in evaluating centralization, one should remember that there is a choice of how much information to try to use; it would appear that rapid random choice is always an option.[8]

IV. Large Numbers

In this section we show that if we replicate our problem, so that there are many firms drawn independently from the same cost distribution (which is common knowledge), then coordination is no problem, al-

---

[8] A rapid and random choice might well have maximized social welfare in the choice of standards for AM stereo broadcasting (see Besen and Johnson 1986). Instead, the administrative process in the Federal Communications Commission became bogged down in an adversarial proceeding, and eventually the commission refused to choose standards. As this example suggests, although rapid choice is an option, it may not be taken even when it should be.

though of course incentive problems may remain. This suggests that decentralization might be most useful in large markets and large economies, *not* because central planners' information processing ability is limited or because large numbers make collusion less likely, but because coordination problems under decentralization are mitigated as the economy becomes large.

We can also interpret this result in terms of prices. Suppose that costs are independently drawn from a known distribution. Using only common knowledge, for any given fraction α, we can find an output price $p(\alpha)$ such that, when all agents who can make a profit at price $p(\alpha)$ enter and none who would make losses enter, then the supply is α *in expected value.* For small numbers, the realized supply will randomly differ from its expected value α, and, as Weitzman (1974) showed, it may therefore be better to use another means of coordination, such as centralized quantity setting. With large numbers, however, this problem disappears, and there is an almost deterministic "supply curve." This solves the coordination problem and also solves the "prices versus quantities" problem.

In a market setting with large numbers, no centralized quantity setting (as by a Soviet-style planner) or *centralized price setting* (as by a Walrasian auctioneer) is needed to solve the coordination problem: symmetric equilibrium determines a cutoff cost level such that precisely those with lower costs enter, and the equilibrium condition is that the (confidently) expected endogenous price is just the one to justify exactly those entering. Of course, that outcome need not be efficient in general, but any problems are standard incentive or externality problems, not coordination problems.

In a sense, this justifies the textbook treatment of entry into a competitive market: the right number of firms enter so that the profits for entrants are driven to zero, and no costly delay is imposed by the process. But we stress that, here, this is a *result,* requiring some restrictive assumptions.

Formally, consider n potential entrants, with sunk costs of entry S drawn independently from the same distribution $F(\cdot)$. Suppose that the *gross* benefits accruing to an entrant if a fraction f of all the firms enter are $b(f)$, where b is continuous and decreasing in f, with $b(0) \geq S^{\max}$ and $b(1) \leq S^{\min}$. (In the game above, we had $n = 2$, $b(1) = \mu$, and $b(\frac{1}{2}) = \lambda$.) As above, the potential entrants play a simple timing game: in each period, each chooses to enter or to wait. Because $b(F(\cdot))$ is decreasing, there exists a unique cutoff level $\overline{S}$ such that $b(F(\overline{S})) = \overline{S}$, and we have the following proposition.

PROPOSITION 5. In the limit, as $n \to \infty$, all firms with $S < \overline{S}$ enter in the first period, and there is no further entry.

Proof. In any symmetric perfect Bayesian equilibrium of the game

with n firms, there exists a cutoff level S_1^n such that all firms with $S < S_1^n$ enter in the first period, and others wait. Now for $i = 1, \ldots, n$, let x_i be a random variable defined by

$$x_i = \begin{cases} 1 & \text{if firm } i \text{ enters in period 1} \\ 0 & \text{if not,} \end{cases}$$

and let e^n be the random variable $\Sigma_{i=1}^n x_i$. The expected number of first-period entrants is simply $E(e^n) = nF(S_1^n)$. The x_i are independent and identically distributed. By the law of large numbers, therefore, for any $\epsilon > 0$,

$$\lim_{n \to \infty} \Pr\left[\left| \frac{e^n}{n} - F(S_1^n) \right| \geq \epsilon \right] = 0,$$

and hence by the continuity of $b(\cdot)$, for any $\epsilon' > 0$,

$$\lim_{n \to \infty} \Pr\left[\left| b\left(\frac{e^n}{n}\right) - b(F(S_1^n)) \right| \geq \epsilon' \right] = 0.$$

To see that this implies $S_1^n \to \bar{S}$, observe that otherwise, for some $S \neq \bar{S}$, we can take a subsequence of n's such that $S_1^n \to S$. But this implies that, with $\epsilon' = \frac{1}{2}|S - \bar{S}|$, say, and for large enough n in the subsequence, $b(e^n/n)$ is almost surely very close to $b(F(S))$. But this cannot be an equilibrium since, for $S \neq \bar{S}$, $S \neq b(F(S))$. Q.E.D.

Thus large numbers may make decentralization work much better. Of course, large numbers may also affect the costs of centralization. In particular, Milgrom and Weber (1983)[9] consider a planner who does not know the distribution of costs across firms (formally, firms' costs need not be independently distributed) but can *sample* costs at different firms. They show that when there are many firms, the planner can almost completely learn about the distribution by sampling a very small proportion of the firms. This suggests that, in a large economy, a central planner may not find it very costly to gather information: he can learn a great deal at a cost that is very small in terms of the economy.[10]

V.   Related Literature

Many readers will find our ideas reminiscent of Weitzman's (1974) seminal paper. For instance, he wrote the following:

[9] Green and Laffont (1979) also consider the effects of statistical sampling on the costs of information acquisition in dominant-strategy mechanisms in large economies.

[10] Of course, this works not only with costly direct sampling but also with providing informational rents for revelation of private information: the individual rationality constraints can be ignored for a small sample since society can easily afford to subsidize a few randomly chosen "benchmark" firms to see what is feasible and to control the remainder.

> Suppose that fulfillment of an important emergency rescue operation demands a certain number of airplane flights. It would be inefficient to just order airline companies or branches of the military to supply a certain number of the needed aircraft because marginal (opportunity) costs would almost certainly vary all over the place. Nevertheless, such an approach would undoubtedly be preferable to the efficient procedure of naming a price for plane services. Under profit maximization overall output would be uncertain, with too few planes spelling disaster and too many being superfluous. [P. 478, n. 1]

Weitzman emphasized the static uncertainty about total output that stems from naming prices when supply curves are unknown. Our work differs in two ways. First, it is intrinsically dynamic: we are concerned not only with getting the wrong supply but with getting the supply too late.[11] Second, and more fundamental, Weitzman addressed a problem of *optimal centralization:* Should a central planner set prices or set quantities? He (consciously) did not ask whether centralization is desirable. In a *decentralized* market system, despite the misleading analogies of Walrasian auctioneers, prices are not centrally "named": they emerge over time from asymmetrically informed firms' initially uncoordinated entry and supply decisions. Thus when airplanes are needed, for instance, decentralization *does not* mean announcing a price that the planner hopes or guesses will elicit the correct supply any more than it means announcing who should supply which planes; it means waiting to see who will decide (not knowing the prices that will prevail) to bring planes to market, and the price and the supply are joint consequences of those decisions.

Recent work on decentralized coordination of the kind studied here includes Dixit and Shapiro (1986), Farrell (1987), and Farrell and Saloner (1988). These papers studied how decentralized coordination through unilateral decisions is possible but imperfect; Farrell and Saloner compared that with coordination through standardization committees, a form of voluntary centralization. All those papers, however, studied complete-information models and so could not analyze the "sorting" consequences of decentralization.

Crawford and Haller (1990) formally justify the use of symmetric equilibrium analysis in studying coordination. They focus on games in which players' interests completely coincide (so that cheap talk, if allowed, would resolve their problems), but their arguments could be extended to our case.

---

[11] The example above is not the only one in which it seems that Weitzman had in mind some notion of urgency, but his analysis did not reflect this.

Bliss and Nalebuff (1984) analyzed decentralized equilibrium with incomplete information (but in a war-of-attrition rather than a grab-the-dollar framework). They showed that equilibrium involves delay (though, in their continuous-time framework, no duplication).[12] They did not, however, draw out the second-best welfare implications of their equilibrium analysis.

Our work is inspired by the fundamental questions raised in the early "market socialism" debate. Lange, Lerner, Mises, and Hayek, among others, noted that welfare theorems under complete informa-tion are at most a first step toward welfare comparisons of centraliza-tion and decentralization. Their consensus was that the matter should turn on which system best dealt with private information and the problems that it causes; they debated whether decentralization and the price system or a central planning board (perhaps using pricelike algorithms) would do better. Our work shows, in a special model, how decentralization can use private information even without a price system; but it also warns that market screening involves efficiency costs.[13]

VI. Conclusion

We close by trying to put our very simple analysis into a broader perspective. It seems that laissez-faire systems often efficiently use dispersed private information, but that, in addition to the widely studied incentive problems (monopoly, externalities) that such sys-tems may encounter, they also may perform imperfectly on *coordina-tion* problems. On the other hand, it seems that centrally planned systems may be better at making rapid and arbitrary choices but do poorly at gathering and using dispersed information.

We have modeled this contrast very starkly in the present paper. In reality, both centralized and decentralized systems work in more sub-tle, and perhaps more efficient, ways than we have portrayed. We assumed that the decentralized solution involves no communication or other private coordination between the firms. In fact, private coor-dination may take place through communication (cheap talk),[14] through binding negotiations between potential entrants,[15] through joint ventures, or through competition for investment capital or other

[12] A continuous-time version of our game (analyzed according to the methods of Simon [1987]) would involve duplication but no delay. We do not stress this point, however, because we think that reactions are not instantaneous, and this is properly reflected in periods of positive length.

[13] For a recent related argument, see Stiglitz (1989).

[14] See Farrell (1987) for an analysis of this and of its limits.

[15] Such negotiations would have to be very circumspectly conducted if they are not to violate antitrust laws, however.

scarce factors of production. Similarly, we assumed that a centralized system simply ignores private information. An opposite view is reflected in the theory of mechanism design, which assumes that a central authority can gather and use all information, subject only to respecting incentive constraints. In particular, the problem that we analyzed may be very simply solved by such a planner, simply by auctioning a "franchise."

We think that, realistically, these and other mechanisms of laissez-faire coordination are likely to work only imperfectly in a world of incomplete information. For instance, private negotiations are subject to the usual inefficiencies of bargaining between privately informed agents. Likewise, if a central planner tries to collect private information, delays ensue. Collating information is time consuming, and the more information must be marshaled in one decision maker's mind, the slower the process. Because of information processing limits, information must be condensed before being compared.[16] Moreover, a central authority may be only imperfectly able to commit itself to a mechanism in advance, and private agents may not be fully committed to the mechanism. For instance, notably in defense procurement, bidders sometimes claim to have low costs, take on the contract, and later insist on renegotiating because costs have proved higher than "foreseen." Even the apparently simple first-best "franchise" solution to our entry problem may not work if these commitments are not available.[17] Because of these difficulties, we believe that the simple story we have described retains an element of truth even when these and similar refinements are allowed for.

Broad confirmation of these ideas seems to be provided by evidence on society's reactions to emergencies and to wars. Although many Western societies laud the laissez-faire system in peacetime, in an emergency they change their tune. Of course, there are many possible reasons for this, and our model touches on only one. But it is the one identified by Milward and by Scitovsky et al. as the prime defect of a market system's response to large new opportunities or problems. It is also consistent with some organizational choices made when speed is important but there is no "emergency."[18]

[16] Geanakoplos and Milgrom (1984) have modeled optimal organizational form taking into account this constraint.

[17] In the Appendix, we analyze how the central authority might perform in our problem if unable to make certain commitments.

[18] See, e.g., Chandler's (1977) discussion of how firms became more centralized as speed of reaction became more important. Also, H. Karatsu, of the Nippon Telegraph and Telephone Company, recently explained that the company "used a free bidding system until the U.S. occupation ordered it to do otherwise. That was because Japan had to develop a good telecommunications network system as quickly as possible" (*New York Times,* July 21, 1987).

We mention just one further important limitation of our analysis. We assumed that there is a clear social "opportunity" and that the only question is which of two contenders should exploit it. An alternative possibility is that, rather than (accurate) private *information* (here, on their own costs), private agents may have differing *views* on whether this new market is valuable at all. In that case, it may not be clear how the information should be used (aggregated) even if it were all public. Under decentralization, entry will occur if *some* agent is sufficiently optimistic (allowing, presumably, for a winner's-curse effect): an extreme-value statistic of the sample of views. When is this a good rule, and compared to what alternatives? Sah and Stiglitz (1985) have begun to address this question, but much remains to be done.

In this paper, we have tried to formalize some ideas of the "Austrian school" on how decentralization brings private information to influence market outcomes. Although such discussion is usually couched in terms of market prices, we believe that prices are not the essential element. As in our model, private information often affects decentralized outcomes even without prices. But the link is typically imperfect, and sometimes (especially, we argued, in urgent problems) the imperfections outweigh the advantages of decentralization.

Appendix

Central Direction versus Decentralization and Mechanism Design

As Lange and Lerner pointed out long ago, from the point of view of mechanism design theory, decentralization should never do better than a sophisticated planner: for the latter could always replicate the market institution and might do strictly better. In fact, in our model she can implement the first-best by auctioning a franchise and thus strictly outperform decentralization.

To represent the planner as a sophisticated and benevolent mechanism designer is clearly extreme. In practice, a planned economy must deal not only with the problem of individual incentives but also with those resulting from communication costs, commitment constraints, and the incentives of a self-interested planner. The latter incentive problem is all the more important since centralized organization of economic activity implies greater concentration of power. Discretionary central direction not only is abused but may also induce agents to engage in wasteful "influence activities" (Milgrom and Roberts 1988).

Although these problems are important, we shall concentrate on only one of the additional constraints besides private information faced by authoritarian systems: the planner's limited ability to commit. When production takes a long time, it is often impossible for the planner to be committed to just when she is to intervene and when not. Once she acquires the power to direct agents, she may intervene even when she ought not to, especially in an ex ante sense. Consequently, individual agents' incentives may be adversely affected: they may not undertake socially desirable investments for fear of being

expropriated ex post, or they may not disclose valuable private information for fear that the planner might use this information against them. This Appendix illustrates how lack of commitment in our model may prevent the planner from eliciting private information about firms' costs, and how this leads to inefficiency.

The simplest way of introducing noncommitment into our model is to suppose that some time elapses between the moment when a firm is selected for production and the time when it has completed production. Then, following the literature on dynamic incentive contracts (see, e.g., Laffont and Tirole 1988), we shall assume that the planner cannot commit to a long-term incentive scheme, meaning one that covers the entire relationship with the firm. Specifically, consider the following sequence of moves.

1. The planner sets up an initial incentive scheme in which she requires each firm to report its cost of production; the mechanism specifies how firms are to be selected on the basis of the vector of cost announcements. Payments may also be made at this point, on the basis of announcements.

2. The selected firms then prepare for production, which takes place in two stages. First, the firms build a plant that costs S. In a second stage, the firms use the plant to produce the good at cost c. Thus the total costs of a high-cost firm are given by $S_H = S + C_H$; for a low-cost firm, $S_L = S + C_L$.

3. The initial incentive scheme lasts only until the completion of the first stage of production. It can specify payments only contingent on the plant's being set up (and, of course, contingent on the cost announcements $\hat{c} = (\hat{c}_A; \hat{c}_B)$ of firms A and B, respectively).

We represent the initial incentive scheme by $\{b_i(\hat{c}); P_i^1(\hat{c})\}$, where $i = A, B$, $b_i(\hat{c})$ is the probability of selecting firm i given $\hat{c}$, and $P_i^1(\hat{c})$ is the transfer to firm i given $\hat{c}$, if it is selected and sets up a plant. Without loss of generality, we assume that a firm gets zero if not selected.

Between the first and second stages of production the planner must offer a new contract to the firms. If two firms have been selected by the initial incentive scheme, then the second incentive scheme is like the first one, but now each firm has sunk S and the planner knows each firm's initial cost announcement. If only one firm has been selected, the second incentive scheme simply offers a payment $P_2(\hat{c})$ for production in the second stage. The firm can, of course, reject the offer, in which case one period goes by before the planner can make a new offer. In other words, if only one firm has completed the first stage of production, that firm and the planner play a bargaining game with one-sided offers (by the planner) under possibly incomplete information. (If two firms have completed the first stage of production and both firms reject the planner's new offer, then again the planner must wait for one period before making another offer, etc.)

Finally, to complete the description of our game, we must specify what happens when the initial incentive scheme is rejected by both firms and what happens when neither firm is selected initially. In the first case, the planner must wait one period before offering a new mechanism; similarly in the second case, with the difference that when the planner makes a new offer she knows the initial cost announcements of each firm.

We shall assume that the planner's goal is to maximize consumer's surplus. This involves, among other things, minimizing payments to firms. (Alternatively, one may assume that the planner wishes to minimize payments to the firm because such payments involve a [resource] "cost of public funds." See Caillaud et al. [1988].) Thus if the planner learns from the initial cost announcements that the selected firm has low costs, she will make a low-price

offer in the future. Anticipating this, a low-cost firm may be tempted initially to conceal its low costs, even if it thereby reduces its chances of being chosen.

A complete analysis of this problem is complex. We shall derive only two propositions that show that *the planner cannot avoid selecting an inefficient firm (with positive probability)*. Thus the comparison between clumsy centralization and decentralization emphasized in the text is not too misleading if in fact the proper comparison is between decentralization and sophisticated planning without commitment. Our clumsy planner can be interpreted as a caricature of a sophisticated planner with limited commitment ability.

PROPOSITION A1. Any initial mechanism that never selects an inefficient firm is not incentive compatible.

Proof. First, any initial mechanism that never selects an inefficient firm must be a mechanism that induces full separation between the high- and low-cost types.

Starting from this observation, let $b_1(C_H; C_L)$ be the probability of selecting firm 1 when the profile of cost announcements is $(C_H; C_L)$. We claim that no fully separating mechanism with $b_1(C_H; C_L) = 0$ is incentive compatible. This follows from the incentive-compatibility constraints.

Consider first the high-cost type. If it announces its true costs, then with probability q its rival has low costs and the high-cost firm will not be selected, since $b_1(C_H; C_L) = 0$. With probability $1 - q$, the rival firm also announces high costs, in which case our firm is selected with probability $b_1(C_H; C_H)$. If it is selected and sets up a plant, it gets a net payoff $P_1(C_H; C_H) - S$ from the first incentive scheme, and its continuation payoff is zero because the planner never offers more than C_H in the continuation game. Suppose now that the high-cost type lies and announces low costs. Then its expected payoff from the initial mechanism is $qb_1(C_L; C_L)[P_1(C_L; C_L) - S] + (1 - q)b_1(C_L; C_H)[P_1(C_L; C_H) - S]$. *Again, its continuation payoff is zero.* It is not the case (as one might think) that the high-cost firm that lies in the first period will suffer for it in the second period. For it can guarantee itself a continuation payoff of zero by refusing to complete production (that is what Laffont and Tirole [1988] call "grab the money and run"), and it will never get a positive payoff since transfers never exceed C_H. Thus the high-cost type's incentive constraint is given by

$$(1 - q)b_1(C_H; C_H)[P_1(C_H; C_H) - S]$$
$$\geq qb_1(C_L; C_L)[P_1(C_L; C_L) - S] + (1 - q)b_1(C_L; C_H)[P_1(C_L; C_H) - S]. \quad \text{(A1)}$$

The low-cost type's incentive constraint is given by

$$qb_1(C_L; C_L)[P_1(C_L; C_L) - S] + (1 - q)b_1(C_L; C_H)[P_1(C_L; C_H) - S]$$
$$\geq (1 - q)\left\{ b_1(C_H; C_H)[P_1(C_H; C_H) - S] + (C_H - C_L) \right. \quad \text{(A2)}$$
$$\left. + [1 - 2b_1(C_H; C_H)]\frac{\delta(C_H - C_L)}{2} \right\}.$$

The left-hand side follows from the fact that the low-cost type's continuation payoffs are zero under full separation, since the planner will never offer a transfer above C_L to what she believes is a low-cost firm. The right-hand side is explained by the fact that if a low-cost firm cheats, it can get an *informational rent* of $C_H - C_L$ if it is selected immediately and of $\delta(C_H - C_L)$ if it is selected one period later. (The firm may be selected one period later if the initial mechanism selects neither firm; this event occurs with probability

$1 - 2b_1(C_H; C_H) \geq 0$. It is then optimal for the planner to pick one of the [apparently] high-cost firms at random in the next period; thus each firm will be selected with probability $\frac{1}{2}[1 - b_1(C_H, C_H) - b_2(C_H, C_H)]$.)

Comparing (15) and (16), we can easily see that they both cannot hold. In other words, there does not exist a fully separating initial mechanism with $b_1(C_H; C_L) = 0$. Q.E.D.

Remark.—Among other things, proposition A1 says that a mechanism that initially does not select any firm at all is not incentive compatible. Such a mechanism would give zero continuation payoffs to both cost types, and a low-cost type can improve its payoff by lying and getting an information rent with positive probability.

COROLLARY. The first-best solution cannot be implemented by a planner facing the commitment constraints described above.

Proof. The first-best requires $b_1(C_H; C_L) = 0$. Q.E.D.

The commitment constraint not only prevents the planner from implementing the first-best but also prevents her from replicating the market solution. For one thing, she cannot commit not to intervene in the future. But, worse, she cannot even replicate the market outcome in the first period since this again involves $b_1(C_H; C_L) = 0$. By proposition A1, this is not incentive compatible. *Because of the commitment constraints, the set of feasible planning procedures does not include the market mechanism.* Emulating decentralization may not be an option for a central planner.

References

Besen, Stanley M., and Johnson, Leland L. *Compatibility Standards, Competition and Innovation in the Broadcasting Industry.* Santa Monica, Calif.: Rand Corp., 1986.

Bliss, Christopher, and Nalebuff, Barry. "Dragon-Slaying and Ballroom Dancing: The Private Supply of a Public Good." *J. Public Econ.* 25 (November 1984): 1–12.

Caillaud, Bernard; Guesnerie, Roger; Rey, Patrick; and Tirole, Jean. "Government Intervention in Production and Incentives Theory: A Review of Recent Contributions." *Rand J. Econ.* 19 (Spring 1988): 1–26.

Chandler, Alfred D., Jr. *The Visible Hand: The Managerial Revolution in American Business.* Cambridge, Mass.: Harvard Univ. Press, Belknap Press, 1977.

Crawford, Vincent, and Haller, Hans. "Learning How to Cooperate." *Econometrica* 58 (1990).

Dixit, Avinash, and Shapiro, Carl. "Entry Dynamics with Mixed Strategies." In *The Economics of Strategic Planning: Essays in Honor of Joel Dean,* edited by Lacy Glenn Thomas III. Lexington, Mass.: Lexington Books, 1986.

Farrell, Joseph. "Cheap Talk, Coordination, and Entry." *Rand J. Econ.* 18 (Spring 1987): 34–39.

Farrell, Joseph, and Saloner, Garth. "Coordination through Committees and Markets." *Rand J. Econ.* 19 (Summer 1988): 235–52.

Geanakoplos, John, and Milgrom, Paul. "A Theory of Hierarchies Based on Limited Managerial Attention." Manuscript. New Haven, Conn.: Yale Univ., 1984.

Green, Jerry R., and Laffont, Jean-Jacques. *Incentives in Public Decision-Making.* New York: North-Holland, 1979.

Hayek, Friedrich A. von. "The Use of Knowledge in Society." *A.E.R.* 35 (September 1945): 519–30.

Laffont, Jean-Jacques, and Tirole, Jean. "The Dynamics of Incentive Con-
    tracts." *Econometrica* 56 (September 1988): 1153–75.
Milgrom, Paul, and Roberts, John. "An Economic Approach to Influence
    Activities in Organizations." *American J. Sociology* 94 (suppl., 1988): S154–
    S179.
Milgrom, Paul, and Weber, Robert. "Organizing Production in a Large Econ-
    omy with Costly Communication." Cowles Foundation Discussion Paper
    no. 672. New Haven, Conn.: Yale Univ., July 1983.
Milward, Alan S. *War, Economy, and Society, 1939–1945.* Berkeley: Univ.
    California Press, 1977.
Sah, Raaj Kumar, and Stiglitz, Joseph E. "Human Fallibility and Economic
    Organization." *A.E.R. Papers and Proc.* 75 (May 1985): 292–97.
Schelling, Thomas C. *The Strategy of Conflict.* Cambridge, Mass.: Harvard
    Univ. Press, 1960.
Scitovsky, Tibor; Shaw, Edward; and Tarshis, Lorie. *Mobilizing Resources for
    War: The Economic Alternatives.* New York: McGraw-Hill, 1951.
Simon, Leo. "Basic Timing Games." Working Paper no. 8745. Berkeley:
    Univ. California, July 1987.
Stiglitz, Joseph E. "Incentives, Information, and Organizational Design." *Em-
    pirica* 16, no. 1 (1989): 3–29.
Weitzman, Martin L. "Prices vs. Quantities." *Rev. Econ. Studies* 41 (October
    1974): 477–91.