When Are Agents Negligible?

By David K. Levine and Wolfgang Pesendorfer*

We examine the following paradox: in a dynamic setting, equilibria can be radically different in a model with a finite number of agents than in a model with a continuum of agents. We present a simple strategic setting in which this paradox is a general phenomenon. However, the paradox disappears when there is noisy observation of the players' actions, and the aggregate level of noise does not disappear too rapidly as the number of players increases. We give several economic examples in which this paradox has recently received attention: durable-goods monopoly, corporate takeovers, and time consistency of optimal government policy. (*JEL* C72, C73)

This paper examines a seemingly narrow technical puzzle: in a dynamic setting, equilibria can be radically different in a model with a finite number of agents than in a model with a continuum of agents. While seemingly narrow, this issue has broad economic importance: the continuum-of-agents model is widely used either explicitly or implicitly in applied economic situations ranging from competitive markets to public finance and political economy. The rationale for using the continuum-of-agents model is that it is a useful idealization of a situation with a large finite number of agents, but if equilibria in the continuum model are radically different from equilibria in the model with a finite number of agents, then this idealization makes little sense.

A good example in which this issue has arisen is the study of the Coase Conjecture for a durable-goods monopoly. Work by Drew Fudenberg et al. (1985), and Faruk Gul et al. (1986) showed that when a continuum of buyers is known to value a good more than the seller does, and when the trading period is short, to a good approximation, the monopolist sells the good immediately at her reservation price. On the other hand, Mark Bagnoli et al. (1989) observe that, when faced with a finite number of buyers, the seller can perfectly price-discriminate regardless of the length of the trading period.

This issue also arises in the discussion of the free-rider problem in corporate takeovers (Sanford Grossman and Oliver D. Hart, 1980). A continuum of negligible shareholders implies that any potential raider will not be able to appropriate the efficiency gains from his takeover, since the shareholders will refuse to tender their shares below their post-takeover value. On the other hand, a large finite number of shareholders may allow the raider to appropriate some or even all the efficiency gains from his takeover (see Bengt R. Holmstrom and Barry Nalebuff [1992] and the example below).

The paradox is caused by the following "disappearance of information" in the continuum. In a model taking place over time, agents have the chance to punish and reward other agents for their past play. If only the aggregate play is observed, then the play of one single agent does not affect the observed outcome in the continuum case, and hence individual deviations cannot be met with rewards or punishments. With a finite number of players, there is a change in the aggregate play whenever a player deviates. This change may be very small, but it is perfectly observable. Even though individual actions are unobservable, a slight deviation of aggregate play from the equilibrium outcome indicates that *someone must have deviated.* Therefore deviations can

* Levine: Department of Economics, University of California, Los Angeles, CA 90024; Pesendorfer: Department of Economics, Northwestern University, Evanston, IL 60208. We thank Kyle Bagwell, Drew Fudenberg, Andy Newman, Paul Milgrom, and anonymous referees for helpful comments. This work was partially supported by NSF grant SES90-23697 and the UCLA Academic Senate.

be punished or rewarded regardless of how many agents there are.[1]

For example, in the durable-goods monopoly case, Bagnoli et al. (1989) propose that in the initial period the monopolist sets a price equal to the reservation price of the highest-valued buyers and refuses to lower the price until all of them have bought. This condition forces the highest-valued buyers to buy at the initial price. Such a scheme is impossible with a continuum of anonymous buyers because the seller could not recognize the change in the sales volume or in the sales revenue if one buyer does not buy. More generally, since we have more information about individual play with finitely many players, we can construct incentive systems that are infeasible with a continuum of players and that allow us to induce aggregate outcomes that are unattainable in the continuum case.

There is an intuitive response to this paradox: if the play of a large finite number of agents cannot be perfectly observed, then deviations by a small player cannot be easily detected and responded to.[2] This intuition suggests that a model with a large finite number of agents and a small amount of aggregate noise may have equilibria very similar to those in the continuum limit.[3]

The contribution of this paper is to illustrate the truth of this intuition. We show that a deterministic dynamic model with a continuum of agents is a good idealization of a noisy model with a large finite number of agents.[4] We also argue that the aggregate level of noise must disappear as the number of players increases, but not too rapidly.

Our results are closely connected to the work of Edward J. Green (1980), Pradeep Dubey and Mamoru Kaneko (1985), and Hamid Sabourian (1990).[5] These papers point out the possibility of a discontinuity in extensive-form games when moving from the finite player to continuum limit. Dubey and Kaneko propose to fix the discontinuity by postulating that individual agents *assume* that their deviations cannot be detected unless they exceed some small threshold. This assumption is very strong and implies for games with continuous strategies and concave payoff functions that players will behave as if their actions are entirely unobservable *irrespective of the number of players*.

Green (1980) and Sabourian (1990) propose, as do we, that the problem can be fixed by postulating a model of noisy observation. In this model, the fact that small deviations are hard to detect is a conclusion rather than an assumption. Moreover, for a fixed number of players and "small noise" the resulting equilibria will be similar to the equilibria in the case with perfect observability.[6] The work of Green and Sabourian differs from ours in several respects, however. They both study infinitely repeated games, whereas we deal with a special class of three-period games. Green considers the case when players are restricted to trigger-type strategies, and Sabourian

[1] Roy Radner (1980) considers ε-equilibria in a finitely repeated Cournot game and finds that for a fixed number of repetitions, the ε equilibria converge to the competitive equilibrium as the number of players gets large. The argument is based on the assumption that firms have unlimited capacity; with a fixed capacity the described paradox arises also in this case: the continuum limit has the competitive equilibrium as the unique outcome, whereas with any arbitrarily large but finite number of players collusion can be sustained.

[2] A related idea may be found in a paper by Nabil Al-Najjar (1992).

[3] When the play of individual agents is observed, an alternative solution to this paradox is to drop the anonymity assumption in the continuum limit. This option results in equilibria in which deviations by a single infinitesimal player lead to a large reaction by other players. Such a notion of equilibrium is explored by Fudenberg and Levine (1988).

[4] Another implication of noise in the observation of the first mover is considered by Kyle Bagwell (1992), who shows that, when there is no precommitment, the implicit precommitment value of being the first mover is diminished when there is even a small amount of noise.

[5] We are grateful to Paul Milgrom for pointing out the connection between our results and the work of Dubey and Kaneko and to several referees for pointing out the connection to Green and to Sabourian.

[6] In Dubey and Kaneko's (1985) approach, the continuity with respect to the number of players is achieved at the cost of introducing a discontinuity in the solution concept: equilibrium behavior for any finite number of players when the threshold is zero will be radically different from equilibrium behavior when the threshold is strictly positive, even if arbitrarily small. Neither Green (1980) nor Sabourian (1990) considers the effect of noise on the set of equilibria for a fixed number of players.

generalizes to all strategies. In contrast to Green and Sabourian, we use a model of additive noise and assume that the standard error of the aggregate noise goes to zero more slowly than the inverse of the number of players. Green and Sabourian each make a more technical assumption than we do.[7] Rather than prove a general theorem applicable to the widest possible variety of dynamic games, we choose to focus on the simplest class of games in which the issue of reaction arises: we study games in which a single large agent can precommit to a reaction to a group of anonymous agents (either finite or a continuum). We further simplify by making the strong anonymity assumption that the large agent can observe only the average play of the small agents.

While such a structure may seem very special, recent theoretical literature shows that such a two-stage precommitment equilibrium is a consequence of reputation-building in a repeated setting even when precommitment is impossible. This was implicit in the work of David Kreps and Robert Wilson (1982) and of Paul R. Milgrom and John Roberts (1982) on reputational equilibrium and was made explicit in the work of Fudenberg and Levine (1989, 1992). Marco Celentani (1991), Celentani and Pesendorfer (1992), Klaus Schmidt (1993), and others have extended the scope of this result in a variety of ways.

We also specialize by making a linearity assumption on the payoff function of the small players. This assumption will always be satisfied if the large agent can play mixed strategies and the small agents have von Neumann-Morgenstern preferences. If it fails, the large agent can use even a very small amount of noise to randomize his play. This approach allows him to effectively commit to a mixed strategy in the noisy case, which is impossible if there is no noise and mixed strategies are not allowed.

We consider three illustrative applications of our theoretical results, each a simplification of a model that has been studied extensively. In connection with the Coase conjecture, we consider a model in which a monopolist can commit to a supply curve. In the continuum case and in the noisy finite case, this model leads to the simple monopoly payoff. Without noise and with a finite number of buyers, the monopolist can achieve the payoff corresponding to perfect price discrimination.[8]

Second, we consider a game between a potential corporate raider and a large number of small stockholders similar to that discussed by Grossman and Hart (1980). With finitely many stockholders, the raider will succeed in appropriating all the efficiency gains from his takeover, whereas with a continuum of stockholders, the raider cannot appropriate any of these efficiency gains.

Finally, we consider a simplification of a model introduced by Stanley Fischer (1980) to study the time consistency of government policy. In this model, the government must choose between a capital tax that is a lump sum after household decisions are made and a distortionary labor tax. If households anticipate the capital tax, they will not invest. Previous analysts of this model, such as Varadarajan V. Chari and Patrick J. Kehoe (1990) have assumed a continuum of households and considered the optimal precommitment capital tax, or "Ramsey equilibrium," as the benchmark "best possible" equilibrium. However, without noise and with a finite number of households, the government can actually achieve the first best outcome: if there is sufficient investment, capital is taxed, whereas if investment is insufficient, a punitive tax on labor is imposed. Each household realizes that if it provides insufficient investment it (and everyone else) will face a punitive tax.

I. The Deterministic Case

We study a game that takes place over three periods ($t = 0, 1, 2$) between a single large player (type L) who can precommit in the ini-

[7] Green (1980) and Sabourian (1990) both assume that the map from the distribution of strategies (endowed with the weak topology) to probability distributions over outcomes (endowed with the total variation norm) is continuous. As we note in what follows, our much simpler assumption implies a condition similar to the one assumed by Green and Sabourian.

[8] Note that, in this example, buyer utility is linear in the seller action (price), and as a result mixed strategies are not called for.

tial period to a reaction in the final period and a large number of identical small players (a representative individual is denoted S) who must undertake an action after the precommitment but before the actual move of the large player.

Formally, we let $\mathcal{P}^n \equiv \{1, 2, \dots, n\}$ be a finite set of n small players who move in period 1, while $\mathcal{P}^\infty \equiv [0, 1]$ denotes a continuum of such players. Each small player undertakes an action chosen from $X_S = [0, 1]$. The single large player may precommit in period 0 to a contingent action in period 2 chosen from $X_L = [0, 1]^m$.

Payoffs for a typical small players are given by $\pi_S(x_S, x_L)$ where x_S is the action taken by the particular small player. Notice that we assume that each small player's payoff is independent of the actions of other small players. In addition, letting $x = (x_S, x_L)$, we assume the following.

ASSUMPTION 1 (linearity): $\pi_S(x_S, x_L) = a(x_S) + \sum_{j=1}^m b^j(x_S) \cdot x_L^j$, *where* $a, b^1, \dots, b^m$ *are continuous concave functions and* a *is strictly concave.*

Although the assumption of linearity in the large player's action seems strong, it is satisfied in the examples we study and, more importantly, will always be satisfied if the large player is allowed to play mixed strategies.[9]

The payoff of the large player is given by $\pi_L(x_S, x_L)$, where the first argument represents the average play of the small players.

ASSUMPTION 2 (concavity): $\pi_L(x_S, x_L)$ *is continuous in both arguments and concave in* x_L.

Define $\underline{\pi}_S \equiv \min_{x_L} \max_{x_S} \pi_S(x)$ to be the payoff that any small player can guarantee for himself. Define the best payoff for the large player to be

$$\pi_L^* = \max_x \pi_L(x) \tag{1}$$

subject to

$$\pi_S(x) \geq \underline{\pi}_S.$$

To avoid degeneracy, we make the following assumption.

ASSUMPTION 3 (nondegeneracy): *There is a unique point* x^* *such that* x^* *solves (1), and there is a point* $\hat{x}_S$ *so that* $\pi_S(x_S^*, x_L^*) < \pi_S(\hat{x}_S, x_L^*)$.

This assumption says that there is a unique pair of actions that give the large player his best payoff and at that pair of actions the small players are not playing a best response to the large player.

Also, define the simple Stackelberg payoff by

$$\overline{\pi}_L = \max_{x_S, x_L} \pi_L(x_S, x_L) \tag{2}$$

subject to

$$x_S = \arg\max_z \pi_S(z, x_L).$$

(Note that π_S is strictly concave in x_S, and hence the Stackelberg response of the small players is unique.)

A representative strategy for all small players is a $\sigma_S \in X_S$. If a single small player deviates to x_S this results in the average action of all small players changing from σ_S to

$$\sigma_S \backslash x_S \equiv \frac{n-1}{n} \sigma_S + \frac{1}{n} x_S. \tag{3}$$

A strategy for the large player is a measurable map $\sigma_L: X_S \to X_L$.

To define our notion of equilibrium, say the pair (σ_S, σ_L) is a *Stackelberg response* for the small player if

$$\pi_S(\sigma_S, \sigma_L(\sigma_S)) \geq \pi_S(x_S, \sigma_L(\sigma_S \backslash x_S)). \tag{4}$$

A pair (σ_S, σ_L) is a *precommitment equilibrium* if it is a Stackelberg response, and if for any Stackelberg response (σ_S', σ_L'), $\pi_L(\sigma_S, \sigma_L(\sigma_S)) \geq \pi_L(\sigma_S', \sigma_L'(\sigma_S'))$.

THEOREM 1 (Paradoxical Theorem): *If a game satisfies Assumptions 1–3, then for all*

[9] The linearity assumption is needed for our results in the noisy case: if the large player is constrained to play pure strategies in a setting in which the small players' payoffs are not linear in their action, Theorem 3 fails.

finite n a precommitment equilibrium exists, and the unique amount received by player L in any precommitment equilibrium is π_L^*; If $n = \infty$, a precommitment equilibrium exists, and the unique amount received by the large player is $\bar{\pi}_L < \pi_L^*$.

II. The Noisy Case

We now suppose that the large player can observe the play of the small players only with some degree of imperfection. Let $\mathbf{y} = (y_1, \ldots, y_n) \in X_S^n$ denote a vector of play by n small players. Suppose that the large player observes a random variable z, where

$$z = \sum_{i=1}^{n} \frac{y_i}{n} + \gamma^n \eta^n. \tag{5}$$

The γ^n are positive scalars and $\gamma^n \eta^n$ is an observational error.

We assume that η^n is a random variable with zero mean and unit variance and that it possesses a continuously differentiable density function f^n. Moreover, we assume that the derivatives of these density functions are bounded uniformly, including in n.

We impose the following condition on the noise process.

CONDITION 1 (slow convergence): $\gamma^n \to 0$ *and* $n\gamma^n \to \infty$.

Condition 1, our key condition, guarantees that, while the uncertainty vanishes when the number of small players goes to infinity, it does so at a rate less than $1/n$. To make matters concrete, assume that

$$z = \frac{1}{n} \sum_{i=1}^{n} (y_i + \phi_i) \tag{6}$$

where the ϕ_i are independently and identically distributed and follow a normal distribution with zero mean. In other words, the large player observes all the small players but makes an observational error in observing each one play. In this case,

$$\gamma^n = \sqrt{\operatorname{Var}(\phi_i)} / \sqrt{n}$$

and $\eta^n = \sum_{i=1}^{n} \phi_i / (\gamma^n n)$ follows a standard normal distribution. Consequently, Condition 1 is satisfied. It is also possible to consider other models, in which, for example, the large player only observes a random sample of the small players.

The strategy of the large player is now a map from realizations of z to actions in X_L, that is, $\sigma_L : \mathbb{R} \to X_L$. Let $F^n(z|y)$ denote the distribution of z if $\mathbf{y} = (y_1, \ldots, y_n)$ is the action taken by the small players. The payoff of the ith small player is given by

$$g_S^n(y, \sigma_L) = \int \pi_S(y_i, \sigma_L(z))\, dF^n(z|y) \tag{7}$$

while the payoff of the large player is

$$g_L^n(y, \sigma_L) \equiv \int \pi_L\left(\sum y_i/n, \sigma_L(z)\right) dF^n(z|y). \tag{8}$$

Note that the large player's payoff depends on the actual play of the small players and not on the noisy signal. The idea is that the large player must respond to the small players at a time when he does not yet have complete information about their play, although later, after he has moved, he will find out what they did, and this will determine his payoff. Note also that in the infinite case, Condition 1 implies that all the definitions are unchanged from the previous section.

If all small players choose σ_S we denote this by $[\sigma_S] = (\sigma_S, \ldots, \sigma_S)$; if all small players except one choose σ_S and the deviant player chooses x_S we write $[\sigma_S]\backslash x_S = (\sigma_S, \ldots, x_S, \ldots, \sigma_S)$. A pair (σ_S, σ_L) is a *noisy Stackelberg response* for small players if, for all x_S,

$$g_S^n([\sigma_S], \sigma_L) \geq g_S^n([\sigma_S]\backslash x_S, \sigma_L). \tag{9}$$

A pair (σ_S, σ_L) is a *noisy precommitment equilibrium* if it is a Stackelberg response, and if for any Stackelberg response (σ_S', σ_L'), $g_L^n([\sigma_S], \sigma_L) \geq g_L^n([\sigma_S'], \sigma_L')$.

THEOREM 2 (Noisy Paradoxical Theorem): *If a game satisfies Assumptions 1–3, then for any n and any $\varepsilon > 0$ there is a $\gamma^n >$*

0 such that the payoff to the large player in any noisy precommitment equilibrium is larger than $\pi_L^* - \varepsilon$.

Theorem 2 shows that, for a given number of players and for small noise, the noisy precommitment equilibrium will be very close to the precommitment equilibrium in the case with perfect observability. This result should be seen as a "continuity check" for a fixed number of players. It shows that the equilibrium outcome of the game with perfect observability is preserved for a fixed number of players when the noise is small, and hence there is no discontinuity in the solution concept when moving from the precommitment equilibrium with perfect observability to the noisy precommitment equilibrium.

THEOREM 3 (Not-So-Paradoxical Theorem): *If a game satisfies Assumptions 1–3 and if the noise satisfies Condition 1, a noisy precommitment equilibrium exists. In this case, any limit of noisy precommitment payoffs to the large player as* $n \to \infty$ *is equal to* $\bar{\pi}_L$.

Fix a sequence (γ^n) satisfying Condition 1. Clearly $(\alpha\gamma^n)$, $\alpha > 0$, also satisfies that condition. We may view α as a measure of how much noise there is for small values of n, whereas the tail of (γ^n) determines the noise for large values of n. We summarize our results by saying that if α is sufficiently small, for small n, the noisy precommitment payoff will be close to the maximum possible π_L^*, while if n is large, it will be close to the simple Stackelberg payoff $\bar{\pi}_L$.

COROLLARY: *Suppose a game satisfies Assumptions 1–3 and Condition 1. For any* $\varepsilon > 0$ *and any* $\underline{N} > 1$ *there is an* $\alpha > 0$ *and an* $\bar{N}$, $\underline{N} < \bar{N} < \infty$, *such that, if the noise is* $(\alpha\gamma^n)$,

(i) *if* $n \leq \underline{N}$ *then the payoff to the large player in a precommitment equilibrium with n small players will be between* π_L^* *and* $\pi_L^* - \varepsilon$;
(ii) *if* $n \geq \bar{N}$ *then the payoff to the large player in a precommitment equilibrium will be between* $\bar{\pi}_L + \varepsilon$ *and* $\bar{\pi}_L - \varepsilon$.

The Corollary is an immediate consequence of Theorems 2 and 3. It shows that the equilibrium payoff to the large player is continuous in both directions: if the noise is small, then for a small number of players the equilibrium payoff will be similar to the perfect-observability case, whereas if the number of players is large, then the payoff to the large player will be similar to the continuum case.

III. Economic Examples

We will now consider three examples. All are simplified versions of games that have been studied in the literature. The first is connected with the Coase conjecture, the second with the free-rider problem in corporate takeovers, and the third with time-consistency of government policy.

Example 1: The large player is a monopolist, and the small players are potential buyers. The monopolist sets the price x_L as a function of the realized demand, and buyers choose a quantity to purchase, x_S. The buyers must choose how much to purchase before knowing the price but after the monopolist has decided on the price as a function of average demand. Let $p(x_S) \leq 1$ be the (downward-sloping) inverse demand curve of a typical buyer. We assume that $p(1) = 0$. There is no cost of production.

With perfect observation and a finite population, the monopolist can effectively extract all the consumer surplus by setting a "take it or leave it" price and by committing to not selling anything if demand is insufficient. With noise and a large population, such a commitment is not feasible, because the monopolist cannot tell if there is sufficient demand *of every consumer* and so must settle for the monopoly price.

To see why this is the case, we simply cast the model into our framework. The utility of a buyer is given by consumer surplus:

$$\pi_S(x) = \int_{z \leq x_S} p(z)dz - x_L x_S. \tag{10}$$

Notice that this function is linear affine in x_L, so that Assumption 1 is satisfied. The payoff of the monopolist is $\pi_L = x_L x_S$, so that Assumption 2 is satisfied.

In this game, π_L^* corresponds to the profit realized by the monopolist when he gets all the consumer surplus, while $\bar{\pi}_L$ corresponds to the simple monopoly profit. Without noise, the monopolist can commit to the policy of charging a price equal to the total consumer surplus if demand is 1 (i.e., if every consumer demands exactly one unit) and charging a choke price such as 1 otherwise. It is then an equilibrium for each individual buyer to purchase one unit, since each buyer realizes that by purchasing less, he will in fact face the choke price. When there is noise and the monopolist can observe his demand only imperfectly, such an extreme policy by the monopolist will not work, and Theorem 3 shows that in this case the monopolist (to a good approximation) can do no better than the simple monopoly profit.

Remark: In Bagnoli et al. (1989), the durable-goods monopolist is able to extract all the surplus by using the following "Pacman strategy": Every period the durable good is offered at a price equal to the highest reservation price of the remaining buyers. Every buyer with a reservation price equal to the current price realizes that if he decided to wait instead of purchasing today he would face the same price in any future period until he finally purchases. This strategy has the same spirit as the "choke-price strategy" described above. Similar to the monopoly example above, slightly imperfect observation of the realized demand will guarantee that the "Pacman strategy" cannot be successful.

Example 2: The large player is a potential raider of a corporation. The takeover will increase the value of the firm by $1 > \eta > 0$ due to better management. The initial value of the corporation is 0. Suppose every shareholder owns an equal number of shares. Each shareholder decides a fraction of his shares $x_S \in [0, 1]$ to be offered on the market (at a given price p, where $0 \leq p < \eta$). The raider chooses a takeover probability $x_L \in [0, 1]$ that is conditional on the total fraction of the company's shares offered. If he decides to take over the company, then he purchases the offered shares and implements the improvements in the corporation that lead to the increase in value of the company.[10] The payoff for the raider is $x_L x_S(\eta - p)$. The payoff for the shareholder is $p x_L x_S + x_L \eta(1 - x_S)$. Note that all our assumptions are satisfied in this case.

In the finite case without noise, the raider can commit to a policy of taking over the company only if 100 percent of the shares are offered. This will allow him to appropriate almost all the efficiency gains due to his takeover if the price p was set close to zero. Hence $\pi_L^* = \eta - p$ in this case.

With a continuum of agents, the precommitment equilibrium will not allow the raider to appropriate any of the efficiency gains due to his takeover. Irrespective of the raider's policy, and for all $p < \eta$, every shareholder will set $x_S = 0$ in this case. Theorem 3 shows that the observation of Grossman and Hart (1980) approximately carries over to the finite case, if the raider can only imperfectly observe the number of shares offered.

Example 3: The large player is a government that must choose whether to place a tax on capital or use a distortionary tax in an effort to raise adequate revenue. Small players are households endowed with a single unit of capital; if capital is not taxed, it may be invested to yield a return of $(1 + r)x_S$ where $x_S \leq 1$ is the investment and $r > 0$; if capital is taxed, investment yields no return. If each household invests to the maximum and capital is taxed, the government collects $1 + r$. To raise the same amount of revenue by an alternative tax (on labor, say) costs each household $c > 1 + r$, since the tax is distortionary. Let x_L be the probability that the government taxes capital. Household utility is

$$(11) \quad \pi_S(x_S, x_L) = (1 - x_L)(1 + r x_S - c) + x_L(1 - x_S).$$

[10] For simplicity we assume that the raider can implement the efficiency improvements even if he controls less than 50 percent of the corporation.

Since the government uses a mixed strategy, Assumption 1 is satisfied. Government utility is equal to the households' utility, except that if capital is taxed and if households invest less than the maximum, there will be a revenue shortfall resulting in a loss of $p(1 - x_S)$, where $p > 1$. Government utility is therefore

$$(12) \quad \pi_S(x_S, x_L) = (1 - x_L)(1 + rx_S - c) + x_L(1 - x_S)(1 - p)$$

which certainly satisfies Assumption 2.

In this example, $\pi_L^* = 0$ corresponds to the first best. Inspection shows that $\pi_L \leq 0$, and $\pi_L = 0$ only if households invest one unit and the probability of a tax on capital is 1. Note that $\underline{x}_L < 0$ and hence the first best allocation satisfies individual rationality. At the government maximum, households receive 0 and can improve their utility by reducing investment, so Assumption 3 is satisfied.

In this context, $\overline{\pi}_L$ is known as the second-best, or Ramsey, equilibrium. A calculation shows that this result is obtained when $x_L = 1/(1 + r)$, in which case households are indifferent to the level of investment. Moreover, it is best for the government if households invest to the maximum in this case. Maximum investment increases the utility when there is no capital tax and reduces the penalty when there is a capital tax. The utility actually attained is $\overline{\pi}_L = r[1 - c/(1 + r)] < 0$.

Previous analysts of the problem (see Fischer, 1980; Chari and Kehoe, 1989) have always dealt with the continuum case and concluded that the Ramsey equilibrium is the best possible. Questions have focused on whether the government can actually precommit and so achieve the payoff corresponding to the Ramsey equilibrium, or whether there is a time-inconsistency problem. The analysis here shows that, with perfect observation and finitely many households, the government can do significantly better than the Ramsey equilibrium: the government taxes only capital provided there is enough investment. If there is not enough investment, the government follows a punitive strategy of taxing only labor. Note that the payoff of the households when only labor is taxed is $(1 + r - c) < 0$.

APPENDIX

PROOF OF THEOREM 1 (Paradoxical Theorem):

In the finite case the large player can use the following policy. If x_S^* is observed then the large player chooses x_L^*. If any $\sigma_S \neq x_S^*$ is observed then the large player chooses $\underline{x}_L$, where $\underline{x}_L$ satisfies $\max_{x_S} \pi_S(x_S, \underline{x}_L) = \underline{\pi}_S$. By the construction of x^* no small player has an incentive to deviate from x_S^* since $\pi_S^* \geq \underline{\pi}_S$. Since π_L^* is the highest payoff the large player can get in any precommitment equilibrium, the above policy is optimal.

In the continuum case, since any single player deviation does not change the aggregate σ_S, the optimal policy for the large player can be taken to be a constant action. Assumptions 1 and 2 ensure that there is a pair (x_L, x_S) such that $\pi_L(x_S, x_L) = \overline{\pi}_L$ and x_S is a best response to x_L. Moreover $\overline{\pi}_L < \pi_L^*$ by Assumption 3.

PROOF OF THEOREM 2 (Noisy Paradoxical Theorem):

Choose $\hat{x} = (\hat{x}_S, \hat{x}_L)$ so that $0 < \hat{x}_S < 1$ and so that $\pi_L(\hat{x}) \geq \pi_L^* - \varepsilon$ and $\pi_S(\hat{x}) \geq \underline{\pi}_S + \varepsilon$, $\varepsilon > 0$. Clearly, by continuity of the payoff functions and by Assumption 3, such an $\hat{x}$ exists. Choose δ such that $0 < n\delta < \min[\hat{x}_L, 1 - \hat{x}_L]$ and so that $\pi_j(\hat{x}) \geq \pi_j(x_S, \hat{x}_L) - \varepsilon^2$ for $j =$ L, S and for all $x_S \in [\hat{x}_S - n\delta, \hat{x}_S + n\delta]$. Further choose K and γ'' so that $\Pr(|z - \hat{x}_S| > K) \leq 1 - \varepsilon$ and $\Pr(|z - x_S| > K) > 1 - \varepsilon$ for $|x - \hat{x}_S| > \delta$.

Let the large player use the following strategy: if $|z - \hat{x}_S| \leq K$ then he chooses $\hat{x}_L$ and if $|z - \hat{x}_S| > K$ he chooses $\underline{x}_L$. Given this strategy, note that whenever all but one small player choose $x_S \in \mathcal{N} \equiv [\hat{x}_S - \delta, \hat{x}_S + \delta]$, it is optimal for the deviant small player to choose an x_S' such that

$$\frac{n-1}{n} x_S + \frac{1}{n} x_S' \in \mathcal{N}.$$

$$\text{(A1)}\quad \int \pi_S(\sigma_S^n, \sigma_L^n(z))dF^n(z\,|\,[\sigma_S^n]) - \int \pi_S(x_S, \sigma_L^n(z))dF^n(z\,|\,[\sigma_S^n]\backslash x_S)$$

$$= \int (\pi_S(\sigma_S^n, \sigma_L^n(z)) - \pi_S(x_S, \sigma_L^n(z)))dF^n(z\,|\,[\sigma_S^n]\backslash x_S)$$

$$+ \int \pi_S(\sigma_S^n, \sigma_L^n(z))(dF^n(z\,|\,[\sigma_S^n]) - dF^n(z\,|\,[\sigma_S^n]\backslash x_S)) \geq 0$$

To see this, note that by choosing an $x_S' \in [\hat{x}_S - n\delta, \hat{x}_S + n\delta]$ the small player can guarantee that

$$\frac{n-1}{n} x_S + \frac{1}{n} x_S' = \hat{x}_S \in \mathcal{N}$$

which in turn implies that with probability $1 - \varepsilon$ the large player will choose $\hat{x}_L$. Thus, by putting $[(n-1)/n]x_S + (1/n)x_S' \in \mathcal{N}$ the small player can guarantee himself a payoff of $\pi_S(\hat{x})(1-\varepsilon) - \varepsilon^2 + \varepsilon\underline{\pi}_S$, which for small ε is clearly better than what the small player could get by choosing an x_S for which the aggregate action is outside $\mathcal{N}$. By a simple fixed-point argument, it follows that there exists an $\tilde{x}_S \in \mathcal{N}$ such that if $n - 1$ players choose $\tilde{x}_S$ then it is optimal for the nth player to choose $\tilde{x}$. [For any $x_S \in \mathcal{N}$ let $f(x_S) = [(n-1)/n]x_S + (1/n)x_S'$, where x_S' is the best response of a small player; f is a continuous function $f : \mathcal{N} \to \mathcal{N}$ and hence there is a fixed point of f.] It remains to be shown that at $\tilde{x}$ the probability of punishment is small. A simple calculation shows that at $\tilde{x}$ the probability of punishment cannot exceed 2ε, since each small player always has the option of choosing an x_S' such that the aggregate action is $\hat{x}_S$. Thus the described strategy guarantees the large player a payoff of $\pi_L(\hat{x})(1-2\varepsilon) + 2\varepsilon\underline{\pi}_L \geq \pi_L^*(1-2\varepsilon) - \varepsilon + 2\varepsilon\underline{\pi}_L$, where $\underline{\pi}_L$ is the payoff the large player receives if he forces the small player's payoff down to its min-max value. Note that in any precommitment equilibrium, the large player has to get a payoff at least as large as the payoff of the indicated strategy. Therefore, since ε is arbitrary the proposition follows.

PROOF OF THEOREM 3 (Not-So-Paradoxical Theorem):

Let σ_S^n be a sequence converging to σ_S, and let $\sigma_L^n(z)$ be such that σ_S^n is a Stackelberg response to σ_L^n. The loss to the ith small player from deviating from σ_S^n to x_S in the n-player game is given by equation (A1) above. Note that

$$\text{(A2)}\quad \int |dF^n(z|[\sigma_S^n]) - dF^n(z|[\sigma_S^n]\backslash x_S)|$$

$$= \int \left| f^n\left(\frac{z - \sigma_S^n}{\gamma^n}\right) - f^n\left(\frac{z - (\sigma_S^n - (x_S - \sigma_S^n)/n)}{\gamma^n}\right)\right| dz.$$

Using Taylor's theorem, this equals

$$\text{(A3)}\quad \int \left| Df^n(h^n(z)) \frac{x_S - \sigma_S^n}{n\gamma^n} \right|.$$

Since we have assumed that the Df^n are uniformly bounded, Condition 1 implies that:[11]

$$\text{(A4)}\quad \lim_{n\to\infty} \int |dF^n(z|[\sigma_S^n]) - dF^n(z|[\sigma_S^n]\backslash x_S)| = 0.$$

[11] The integral of the absolute value of the derivative to the measure is the total variation of the measure. This condition essentially means that small changes in the distribution of players' actions have a small effect on the distribution of outcomes in the total variation norm. This is the assumption used by Green (1980) and Sabourian (1990).

$$(A5)\quad \int \pi_S((x_S, \sigma_L^n(z)) - \pi_S(\sigma_S^n, \sigma_L^n(z)))dF^n(z|[\sigma_S^n]\backslash x_S^n) \leq \varepsilon^n$$

Thus it follows that inequality (A5), above, must hold, where $\lim_{n\to\infty}\varepsilon^n = 0$. Therefore, there is a sequence of probability measures on X_S, $G^n(q)$ such that

$$(A6)\quad \int (\pi_S(x_S, q) - \pi_S(\sigma_S^n, q))dG^n(q) \leq \varepsilon^n$$

from which it follows that

$$(A7)\quad \int (\pi_S(x_S, q) - \pi_S(\sigma_S, q))dG^n(q) \leq \varepsilon^n + \sup_q |\pi_S(\sigma_S^n, q) - \pi(\sigma_S, q)|.$$

Let $G(\cdot)$ be a weak limit point of the sequence $G^n(\cdot)$.[12] Since by assumption $\sigma_S^n \to \sigma_S$, we conclude that

$$(A8)\quad \int (\pi_S(x_S, q) - \pi_S(\sigma_S, q))dG(q) \leq 0.$$

Set $x_L = \int q dG(q)$ to the expected value of the play of the large player according to G. Because π_S is linear affine in x_L, it follows that

$$(A9)\quad \pi_S(x_S, x_L) - \pi_S(\sigma_S, x_L) \leq 0.$$

On the other hand, the large player gets

$$(A10)\quad \int \pi_L(\sigma_S^n, \sigma_L^n(z))dF^n(z|[\sigma_S^n]) \to \int \pi_L(\sigma_S, q)dG(q) \leq \pi_L(\sigma_S, x_L)$$

[12] That is, the sequence has a subsequence that converges to this point in the weak topology. This topology is characterized by the convergence of the expectation of continuous functions, and the space of probability measures on a compact set is known to be compact in this topology (see e.g., K. R. Parthasarathy, 1967 theorem 6.4, p. 45).

where the final inequality follows from the assumption that the large player's payoff is concave in his own action. We conclude that the limit of the precommitment payoff to the large player in the finite games is not greater than in the limit game. Finally, we observe that since an optimal precommitment in the limit game is to precommit to a constant function, this is feasible and yields approximately the same payoff in the finite game for large n, so that the limit of precommitment payoffs is not smaller than the precommitment payoff in the limit game.

REFERENCES

Al-Najjar, Nabil. "The Coase Conjecture in Markets with a Finite Number of Consumers." Working paper, University of Quebec, 1992.

Bagnoli, Mark; Salant, Stephen W. and Swierzbinski, Joseph E. "Durable Goods Monopoly with Discrete Demand." *Journal of Political Economy,* December 1989, *97*(6), pp. 1459–78.

Bagwell, Kyle. "Commitment and Observability in Games." CMSEMS Discussion Paper No. 1014, Northwestern University, 1992.

Celentani, Marco. "Reputation in General Deterministic Stage Games." UCLA Working Paper No. 189, University of California-Los Angeles, 1991; *Economic Theory* (forthcoming).

Celentani, Marco and Pesendorfer, Wolfgang. "Reputation in Dynamic Games." CMSEMS Discussion Paper No. 1009, Northwestern University, 1992; *Journal of Economic Theory* (forthcoming).

Chari, Varadarajan V. and Kehoe, Patrick J. "Sustainable Plans." *Journal of Political Economy,* August 1990, *98*(4), pp. 783–802.

Dubey, Pradeep and Kaneko, Mamoru. "Information Patterns and Nash Equilibria in Extensive Games II." *Mathematical Social Science,* 1985, *10*, pp. 247–62.

Fischer, Stanley. "Dynamic Inconsistency, Cooperation, and the Benevolent Dissembling Government." *Journal of Economic Dynamics and Control,* February 1980, *2*(1), pp. 93–107.

Fudenberg, Drew and Levine, David K. "Open-Loop and Closed-Loop Equilibria in Dynamic Games with Many Players." *Journal of Economic Theory,* February 1988, *44*(1), pp. 1–18.

______. "Reputation and Equilibrium Selection in Games with a Single Patient Player." *Econometrica,* July 1989, *57*(4), pp. 759–78.

______. "Maintaining a Reputation When Strategies Are Imperfectly Observed." *Review of Economic Studies,* July 1992, *59*(3), pp. 561–80.

Fudenberg, Drew; Levine, David K. and Tirole, Jean. "Infinite Horizon Models of Bargaining with One-Sided Incomplete Information," in A. Roth, ed., *Bargaining with incomplete information.* London: Cambridge University Press, 1985, pp. 73–98.

Green, Edward J. "Noncooperative Price Taking in Large Dynamic Markets." *Journal of Economic Theory,* April 1980, *22*(2), pp. 155–81.

Grossman, Sanford and Hart, Oliver D. "Takeover Bids, the Free Rider Problem, and the Theory of the Corporation." *Bell Journal of Economics,* Spring 1980, *11*(1), pp. 42–65.

Gul, Faruk; Sonnenschein, Hugo and Wilson, Robert. "Foundations of Dynamic Monopoly and the Coase Conjecture." *Journal of Economic Theory,* June 1986, *39*(1), pp. 155–90.

Holmstrom, Bengt and Nalebuff, Barry. "To the Raider Goes the Surplus? A Reexamination of the Free-Rider Problem." *Journal of Economics and Management Strategy,* Spring 1992, *1*(1), pp. 37–63.

Kreps, David and Wilson, Robert. "Reputation and Imperfect Information." *Journal of Economic Theory,* August 1982, *27*(2), pp. 253–79.

Milgrom, Paul M. and Roberts, John. "Predation, Reputation, and Entry Deterrence." *Journal of Economic Theory,* August 1982, *27*(2), pp. 443–60.

Parthasarathy, K. R. *Probability measures on metric spaces.* New York: Academic Press, 1967.

Radner, Roy. "Collusive Behavior in Noncooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives." *Journal of Economic Theory,* April 1980, *22*(2), pp. 136–54.

Sabourian, Hamid. "Anonymous Repeated Games with a Large Number of Players and Random Outcomes." *Journal of Economic Theory,* June 1990, *51*(1), pp. 92–110.

Schmidt, Klaus. "Reputation and Equilibrium Characterization in Repeated Games of Conflicting Interest." *Econometrica,* March 1993, *61*(2), pp. 325–52.