Forthcoming in: American Economic Review

When are Agents Negligible?

David K. Levine Wolfgang Pesendorfer*

November 1992

revised

August 1993

Abstract

We examine the following paradox: In a dynamic setting, an arbitrarily large finite number of agents and a continuum of agents can lead to radically different equilibrium outcomes. We show that in a simple strategic setting this paradox is a general phenomenon. We also show that the paradox disappears when there is noisy observation of the players' actions: The aggregate level of noise must disappear as the number of players increases, but not too rapidly. We give several economic examples in which this paradox has recently received attention: the durable goods monopoly, corporate takeovers, and time consistency of optimal government policy.

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. 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 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 (SLOW CONVERGENCE) implies that all the definitions are unchanged.

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

$$(9) \qquad g_S([\sigma_S], \sigma_L) \geq g_S([\sigma_S]\backslash_i x_S, \sigma_L)$$

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), $\pi_L([\sigma_S], \sigma_L) \geq \pi_L([\sigma'_S], \sigma'_L)$.

Noisy Paradoxical Theorem *If a game satisfies (LINEARITY), (CONCAVITY) and (NONDEGENERACY) then for any n and any $\epsilon > 0$ there is a $\gamma^n > 0$ such that the payoff to the large player in any noisy precommitment equilibrium is larger than $\pi_L^* - \epsilon$.*

The Noisy Paradoxical Theorem 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.

Not So Paradoxical Theorem *If a game satisfies (LINEARITY), (CONCAVITY) and (NONDEGENERACY) and if the noise satisfies (SLOW CONVERGENCE), 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 (SLOW CONVERGENCE). Clearly $(\alpha\gamma^n), \alpha > 0$, also satisfies (SLOW CONVERGENCE). We may view α as a measure of how much noise there is for small values of n while 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 (LINEARITY), (CONCAVITY), (NONDEGENERACY) and (SLOW CONVERGENCE). For any $\epsilon > 0$ and any $\underline{N} > 1$ there is an $\alpha > 0$ and a $\overline{N}, \underline{N} < \overline{N} < \infty$, such that if the noise is $(\alpha\gamma^n)$ and*

(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 $\pi_L^* - \epsilon$.

(ii) if $n \geq \overline{N}$ then the payoff to the large player in a precommitment equilibrium will be between $\bar{\pi}_L + \epsilon$ and $\bar{\pi}_L - \epsilon$.

The Corollary is an immediate consequence of the Noisy Paradoxical Theorem and the Not So Paradoxical Theorem. 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.

3 Economic Examples

We 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 (LINEARITY) is satisfied. The payoff of the monopolist is $\pi_L = x_L x_S$, so that (CONCAVITY) 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 our Not So Paradoxical Theorem 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 shareholders owns an equal amount 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 $x_L \cdot x_S \cdot (\eta - p)$. The payoff for the shareholder is $p \cdot x_L \cdot x_S + x_L \cdot \eta \cdot (1 - x_S)$. 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% 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 raiders policy, and for all $p < \eta$, every shareholder will set $x_S = 0$ in this case. The Not So Paradoxical Theorem 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

$$\pi_S(x_S, x_L) = (1-x_L)(1+rx_S-c) + x_L(1-x_S) \tag{11}$$

Since the government uses a mixed strategy, (LINEARITY) 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

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

which certainly satisfies (CONCAVITY).

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 1 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 (NONDEGENERACY) is satisfied.

In this context $\bar{\pi}_L$ is known as the second best, or Ramsey equilibrium. A calculation shows that this 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. This increases the utility when there is no capital tax, and reduces the penalty when there is a capital tax. The utility actually attained is $\bar{\pi}_L = r(1-c/(1+r)) < 0$.

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

Paradoxical Theorem *If a game satisfies (LINEARITY), (CONCAVITY) and (NONDEGENERACY) then for all 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 player 2 is $\bar{\pi}_L < \pi_L^*$.*

Proof: 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. (CONCAVITY) and (LINEARITY) ensure that there is a pair (x_L, x_S) such that $\pi_L(x_S, x_L) = \bar{\pi}_L$ and x_S is a best response to x_L. Moreover $\bar{\pi}_L < \pi_L^*$ by (NONDEGENERACY).□

Noisy Paradoxical Theorem *If a game satisfies (LINEARITY), (CONCAVITY) and (NONDEGENERACY) then for any n and any $\epsilon > 0$ there is a $\gamma^n > 0$ such that the payoff to the large player in any noisy precommitment equilibrium is larger than $\pi_L^* - \epsilon$.*

Proof: Choose $\hat{x} = (\hat{x}_S, \hat{x}_L)$ so that $\pi_L(\hat{x}) \geq \pi_L^* - \epsilon$ and $\pi_S(\hat{x}) \geq \underline{\pi}_S + \epsilon, \epsilon > 0$. Clearly by continuity of the payoff functions and by (NONDEGENERACY) such a $\hat{x}$ exists. Choose δ such that $\pi_j(\hat{x}) \geq \pi_j(x_S, \hat{x}_L) - \epsilon^2$ for $j = L, S$ and for all $x_S \in [\hat{x}_S - n\cdot\delta, \hat{x}_S + n\cdot\delta]$. Further choose K, γ^n so that $\text{Prob}(|z - \hat{x}_S| > K) \leq 1 - \epsilon$ and $\text{Prob}(|z - x_S| > K) > 1 - \epsilon$ 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 (player i) small players choose $x_S \in N \equiv [\hat{x}_S - \delta, \hat{x}_S + \delta]$ it is optimal for player i to choose a x'_S such that $\sum x^i_S/n \in N$. To see this note that by choosing a $x_S \in [\hat{x}_S - n\cdot\delta, \hat{x}_S + n\cdot\delta]$ the small player can guarantee an aggregate action $\hat{x}_S$ which in turn implies that with probability $1-\epsilon$ the large player will choose $\hat{x}_L$. Thus the small player can guarantee himself a payoff of $\pi_S(\hat{x})(1-\epsilon) - \epsilon^2 + \epsilon\underline{\pi}_S$ 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 N. By a simple fixed point argument it follows that there exists a $\tilde{x}_S \in N$ such that if $n-1$ players choose $\tilde{x}_S$ then it is optimal for the n-th player to choose $\tilde{x}$. (For any $x_S \in N$ let $f(x_S) = \frac{n-1}{n}x_S + \frac{1}{n}x'_S$ where x'_S is the best response of a small player. f is a continuous function $f : N \to N$ 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 a x_S such that the aggregate action is $\hat{x}_S$. Thus the described strategy guarantees the large player a payoff of $\pi_L(\hat{x})(1-2\epsilon) + 2\epsilon\underline{\pi}_L \geq \pi^*_L(1-2\epsilon) - \epsilon + 2\epsilon\underline{\pi}_L$, were $\underline{\pi}_L$ is the payoff of the large player receives if he forces the small players payoff down to its minmax 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.□

Not So Paradoxical Theorem *If a game satisfies (LINEARITY), (CONCAVITY) and (NONDEGENERACY) and if the noise satisfies (SLOW CONVERGENCE), 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 $\tilde{\pi}_L$.*

Proof: Let σ^n_S be a sequence converging to σ_S, and $\sigma^n_L(z)$ be such that σ^n_S is a Stackelberg response to σ^n_L. The loss to the i-th small player from deviating from σ^n_S to x_S in the n-player game is

$$\begin{aligned}\int \pi_S(\sigma^n_S, \sigma^n_L(z))dF^n(z|[\sigma^n_S]) - \int \pi_S(x_S, \sigma^n_L(z))dF^n(z|[\sigma^n_S]\backslash x_S) &= \\ \int (\pi_S(\sigma^n_S, \sigma^n_L(z)) - \pi_S(x_S, \sigma^n_S(z)))dF^n(z|[\sigma^n_S]\backslash x_S) &+\end{aligned}$$

$$(13)\qquad \int \pi_S(\sigma_S^n, \sigma_L^n(z))(dF^n(z||\sigma_S^n|) - dF^n(z||\sigma_S^n|^n\backslash x_S)) \quad \geq \quad 0$$

Note that

$$(14)\int |dF^n(z||\sigma_S|) - dF^n(z||\sigma_S|\backslash x_S)| = \int \left| f^n(\frac{z-\sigma_S}{\gamma^n}) - f^n\left(\frac{z-(\sigma_S-(x_S-\sigma_S)/n)}{\gamma^n}\right)\right| dz$$

Using Taylor's Theorem, this equals

$$(15)\qquad \int \left| Df^n(h^n(z))\frac{x_S - \sigma_S}{n \cdot \gamma^n}\right|$$

Since we have assumed that the Df^n are uniformly bounded, (SLOW CONVERGENCE) implies that

$$(16)\qquad \lim_{n\to\infty} \int |dF^n(x||\sigma_S|) - dF^n z||\sigma_S|\backslash x_S)| = 0^{11}$$

Thus it follows that

$$(17)\qquad \int \pi_S((\sigma_S^n, \sigma_L^n(z)) - \pi_S(\sigma_S^n, \sigma_L^n(z)))dF^n(z||\sigma_S^n|\backslash x_S^n) \leq \epsilon^n$$

where $\lim_{n\to\infty} \epsilon^n = 0$. Therefore, there is a sequence of probability measures on X_S, $G^n(q)$ such that

$$(18)\qquad \int (\pi_S(x_S, q) - \pi_S(\sigma_S^n, q))dG^n(q) \leq \epsilon^n$$

from which it follows that

$$(19)\qquad \int (\pi_S(x_S, q) - \pi_S(\sigma_S, q))dG^n(q) \leq \epsilon^n + \sup_q |\pi_S(\sigma_S^n, q) - \pi(\sigma_S, q)|$$

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

$$(20)\qquad \int (\pi_S(x_S, q) - \pi_S(\sigma_S, q))dG(q) \leq 0$$

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

$$(21)\qquad \pi_S(x_S, x_L) - \pi_S(\sigma_S, x_L) \leq 0$$

On the other hand the large player gets

$$(22)\qquad \int \pi_S(\sigma_S^n, \sigma_L^n(z))dF^n(z||\sigma_S^n|) \to \int \pi_L(\sigma_S, q)dG(q) \leq \pi_L(\sigma_S, x_L),$$

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. □

Footnotes

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

[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 while with any arbitrarily large but finite number of players collusion can be sustained.

[2]A related idea may be found in 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 leads to equilibria in which deviations by a single infinitesimal player leads 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 Sabourian.

[6]In Dubey and Kaneko's 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 if the threshold is zero will be radically different from equilibrium behavior if the threshold is strictly positive, even if it is arbitrarily small.

Green and Sabourian do not consider the effect of noise on the set of equilibria for a fixed number of players.

[7]Green and Sabourian 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 below, our much simpler assumption implies a conditions 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.

[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, the Not So Paradoxical Theorem fails.

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

[11]The integral of the absolute value of the derivative to the measure is the total viariation 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 assumptions used by Green (1980) and Sabourian (1990).

[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. pg. 45.

References

[1] Al-Naijar, N. "The Coase Conjecture in Markets with a Finite Number of Consumers", Working paper, University of Quebec, Montreal, 1992.

[2] Ausubel, Lawrence M. and Deneckere, Raymond J."Reputation in Bargaining and Durable Goods Monopoly", Econometrica, 1989, 57, pp. 511-531.

[3] Bagnoli, Mark, Salant Stephen W. and Swierzbinski Joseph K. "Durable Goods Monopoly with Discrete Demand", Journal of Political Economy, 1989, 97, pp. 1459-78.

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

[5] Celentani, Marco. "Reputation in General Deterministic Stage Games", UCLA working paper No. 189, 1991.

[6] —– and Pesendorfer, Wolfgang. "Reputation in Dynamic Games", CMSEMS Discussion Paper No. 1009, Northwestern University 1992.

[7] Chari, Varadarajan V. and Kehoe, Patrick J. "Sustainable Plans and Debt". Journal of Political Economy, 1990, 98, pp. 783-802.

[8] Dubey, Pradeep and Kaneko, Mamoru. "Information Patterns and Nash Equilibria in Extensive Games II.", Mathematical Social Science, 1985, 10, pp. 247-262.

[9] Fischer, Stanley. " Dynamic Inconsistency, Cooperation and the Benevolent Dissembling Government", Journal of Economic Dynamics and Control, 1980, 1, pp. 93-107.

[10] Fudenberg, Drew and Levine, David K. "Open-Loop and Closed-Loop Equilibria in Dynamic Games with Many Players", Journal of Economic Theory, 1988, 44, pp. 1-18.

[11] —– ." Reputation and Equilibrium Selection in Games with a Single Patient Player", Econometrica, 1989, 57, pp. 759-78.

[12] —– 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. 143-149.

[13] Grossman, Sanford and Hart, Oliver D. "Takeover Bids, the Free Rider Problem, and the Theory of the Corporation", *Bell Journal of Economics*, 1980, 11, pp. 42-65.

[14] Green, Edward J. "Noncooperative Price Taking in Large Dynamic Markets", Journal of Economic Theory, 1980, 21, pp. 155-181.

[15] Gul Faruk, Sonnenschein, Hugo and Wilson, Robert. " Foundations of Dynamic Monopoly and the Coase Conjecture", Journal of Economic Theory, 1986, 39, pp. 155-190.

[16] Holmstrom, Bengt and Nalebuff, Barry. "To the Raider Goes the Surplus? A Reexamination of the Free-Rider Problem", Journal of Economics & Management Strategy, 1992, 1, pp. 37-63.

[17] Kreps, David and Wilson, Robert. " Reputation and Imperfect Information", Journal of Economic Theory, 1982, 27, pp. 253-79.

[18] Milgrom, Paul M. and Roberts, John. "Predation, Reputation, and Entry Deterrence", Journal of Economic Theory, 1982, 27, pp. 443-60.

[19] Parthasarathy, K. R., Probability Measures on Metric Spaces, New York: Academic Press, 1967, Theorem 6.4, p. 45.

[20] Radner, Roy. "Collusive Behavior in Noncooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives", Journal of Economic Theory, 1980, 21, pp. 136-154.

[21] Sabourian, Hamid. "Anonymous Repeated Games with a Large Number of Players and Random Outcomes", Journal of Economic Theory, 1990, 51, pp. 92-110.

[22] Schmidt, Klaus. "Reputation and Equilibrium Characterization in Repeated Games of Conflicting Interest", Econometrica, 1993, 61, pp. 325-352.