Discussion Paper No. 1114*

WHEN ARE NON-ANONYMOUS PLAYERS NEGLIGIBLE?

Drew Fudenberg, David Levine and Wolfgang Pesendorfer**

January 5, 1995

We examine games played by a single large player and a large number of opponents who are small, but not anonymous. If the play of the small players is observed with noise, and if the number of actions the large player controls is bounded independently of the number of small players, then as the number of small players grows, the equilibrium set converges to that of the game where there is a continuum of small players. The paper extends previous work on the negligibility of small players by dropping the assumption that small players actions are "anonymous". That is, we allow each small player's actions to be observed separately, instead of supposing that the small players' actions are only observed through their effect on an aggregate statistic.

* The authors are grateful for financial support from National Science Foundation Grants SBR9320695, SBR9409180 and the UCLA Adademic Senate.

** Departments of Economics Harvard, UCLA and Northwestern.

1. Introduction

This paper examines a game played by a single large player and a number of small opponents. Our interest is in finding conditions under which the small players are negligible, in the sense that the actions of any individual small player have little or no effect on the subsequent play of others. In contrast to previous work (discussed below) on this "negligibility" question, we do not wish to impose the condition that the small but non-infinitesimal players are anonymous. Instead, we will suppose that each small player's action influences the distribution of a distinct signal, which may be thought of as a noisy signal of that player's action. Thus, even when there are many players, the large player may be able to make a fairly precise inference about the actions of any small opponent. Nevertheless, we will argue that the large player will not be able to exploit this information, and that the outcome of the game is as if the actions of individual small players could not be observed.

The key assumption we use is that the number of actions the large player has available is bounded independently of the number of small players. Under this assumption, the large player is not able to implement separate rewards and punishments for each individual small player. As we show, this implies that the small players (to a good approximation) ignore the effect their actions might have on the large player's play. Consequently, in the one-shot case the large player can do no better than to commit to an uncontingent choice of action (as in a Stackleberg equilibrium) even if commitments to contingent "threat" strategies are allowed. In the case of a repeated game, a similar argument shows that the small players to a good approximation behave myopically, that is, they ignore the effect of their current actions on the future play of their opponents. This provides a rationale for using reputation models with myopic opponents in situations where the opponents are actually long-lived but small.

To see why the negligibility of small players requires a proof, recall that in general dynamic games, equilibria can be radically different in a model with a finite number of agents than in the standard model used to describe the continuum-of-agents limit. In the standard model it is assumed that the play of any measure-0 set (and hence of any single agent) is ignored. The problem is that in any finite game each agent has positive measure, and so in principle his actions can be observed, but in the continuum limit this information "vanishes."[1] Consequently the games with finite agents may have equilibria where the small players are induced to play non-myopically by the consideration of future rewards or punishments that can be triggered by their current action, yet such equilibria are ruled out by the "negligibility" assumption of the continuum model.

Green (1980), Sabourian (1990), and Levine and Pesendorfer (1994) develop a justification of the negligibility assumption in the setting of "anonymous" games, where the observed signals depend only on an aggregate statistic (typically the sum) of the small players' actions.[2] Note that this anonymity condition on its own does not explain why the small agents are negligible, since, with a finite number of players, there is a change in the

[1] The assumption that measure-0 sets cannot be observed is usual in this context, but is not necessary. See Fudenberg and Levine (1988) for a model of non-anonymous players and observed actions in which information is preserved in passing to the continuum of players limit.

[2] Dubey and Kaneko (1985) assume that small deviations are unobservable, but large deviations are observed perfectly.

aggregate play whenever any player deviates. Consequently, a slight deviation of aggregate play from the equilibrium outcome indicates that *someone* must have deviated, and so the large player can design a strategy that deters deviations regardless of the size of the small players. Yet in the continuum of players limit, a deviation by a single small player deviation will not change the aggregate statistic, and so there cannot be equilibria that rely on punishments. Intuitively, we would expect if we drop the assumption that individual play is perfectly observable, this discontinuity would go away, so that the equilibria in the finite player case would be similar to that in the limit. We show that this is indeed the case.

Why does the lack of anonymity matter? If the large player knows the small players by name, he may pick on a particular player or players, and punish and reward them to manipulate their behavior. For example if you know that the government will only build a large public project if you volunteer to pay for it your behavior will be quite different than if you are one of many anonymous voters. If the number of large player's actions is proportional to the number of small players, so that the large player can reward or punish each one individually, then relaxing the anonymity condition overturns the negligibility result. The essential point of this paper is that non-anonymous small players do remain negligible if the large player is limited in his number of actions, so that most small players must necessarily be treated as if they are anonymous.

To get a rough intuition for this, consider a game where each small player chooses whether to work or shirk, and the large player only has two possible actions, "reward" and "punish;" the small players prefer working to shirking but they would rather work and be rewarded than shirk and be punished. If the large player observes the small players' actions without noise: he can induce them all to work by threatening to punish if even one player shirks. Next suppose that each small player's action is observed with noise, so that with some small probability ε the large player thinks a small player has shirked when that player really worked. The noise wouldn't matter if the large player had the ability to give separate rewards and punishments to each player, for then the strategy "punish small player i if it looks as though that player shirked" will still induce every player to work. However, when the large player only has two actions available, he is unable to implement this sort of personalized punishments.

If the large player uses a strategy where the probability of punishment depends only on the number of observations of the outcome "shirk", and not on the identities of the apparent shirkers, then the situation is as in the anonymous case where the large player does not observe individual signals: since there will be fraction about ε that appear to have shirked even if they did not, it is difficult (and in the limit as $N \to \infty$ impossible) for any anonymous policy of the large player to induce a different distribution of responses if all small players work than if one small player shirks while the others work.

Of course the large player has the option of basing its action solely or largely on the outcomes of a small number of players, and doing so will induce those agents to work. However, the number of agents who can be induced to work in this way is bounded independently of the total number of small payers, so that the fraction of players who can be induced to work converges to 0.

The more general statement of this conclusion is that as the number of small players grows to infinity, all but a vanishing fraction of them must play as if they are negligible. We use this conclusion to extend some results that have already been obtained for the anonymous case.

Our results follow from a basic mathematical lemma that has an independently interesting interpretation. Roughly, what this lemma says is that the average effect that an individual can have on a future variable is bounded by a number that is small if there are many decision makers and the probability of individual decisions is relatively random. This is of greatest interest where the future variable is an aggregate that is of equal interest to all the decision makers: per capita GNP or something of that sort. If individual decision determining this variable are relatively uncertain then at most a small number of individuals will be able to predict the effect that their individual decisions will have on this variable. In a sense this lemma implies that there is a limit on the number of "important" decision makers.

We first consider a one-shot game in which a single large player may choose a precommitment strategy against a number of small players.[3] The essential point we make is similar to that in Levine and Pesendorfer (1994): without noise the large player can appropriate all the surplus, but with noisy observation of the small players, he can only get the Stackleberg payoff.

The second part of the paper extends the results on one-shot games to demonstrate that when the play of patient small players is observed with noise, to a good

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

approximation they play "myopically" in the sense of not trying to influence the future actions of other players. This result is of particular interest when coupled with the literature on maintaining a reputation against myopic opponents, as it provides an alternative interpretation for the assumption of myopic play.

2. An Example

Before developing the model, we will use an example to illustrate the content of our results. From a game-theoretic viewpoint, this example is simply a more complex interesting version of the "reward" and "punish" game described in the introduction; the additional complexity makes the game more economically interesting.

In this example, the large player is a government who seeks to maximize its tax revenue. It can set a tax rate $b_i \in \{0,1\}$. Since the government can randomize, let β denote the probability that $b_i = 1$. Small players must decide whether to be unproductive or productive, that is, they choose $a_i \in \{0,1\}, i = 1, \ldots, n$.

First we consider a game with perfectly observable actions. The game is played in the following way. First the government commits to a strategy , then small players decide whether to be productive or unproductive and finally the government observes the small players' actions and executes its policy. The payoffs of the small players are given by

$$(1 - \beta_i) \cdot (1 + 2a_i) - a_i$$

The payoff of the government is given by

$$\beta_i \frac{1}{N} \sum_i (1 + 2a_i)$$

In the game in which the government can observe precisely whether a small player is productive or unproductive the following policy is optimal:

- If all small players choose $a_i = 1$ set $\beta_i = \frac{1}{3}$.
- If at least 1 small player chooses $a_i = 0$ set $\beta_i = 1$.

A best response to this policy is for all small players to choose $a_i = 1$ and hence the per capita tax revenue of the government is 2.

Suppose now that as a result of small player i's action the individual output $y_i \in \{0,1\}$ results. The probability of each individual output is $\rho_i(y_i|a_i)$, where

$$\rho_i(y_i|a_i = y_i) = 1-\varepsilon, \rho_i(y_i|a_i \neq y_i) = \varepsilon.$$

The payoff of a small player is his expected output minus his tax payment and the cost of effort.

$$u(a_i,\beta_i) = \left(1+2(a_i - \varepsilon(2a_i - 1))\right)(1-\beta_i) - a_i.$$

The tax revenue of the government depends on the output and tax rate

$$u_0(y,\beta_i) = \frac{1}{N}\sum_{i=1}^{N}(1+2y_i)(1-\beta_i)$$

Clearly, for a fixed number of small players n and for sufficiently small ε the policy described above for the case of perfect observability yields a tax revenue of $2-2\varepsilon$ per small player and hence the tax revenue of the government is almost as in the perfect information case. However, as we show in Proposition 2 for a fixed ε in the limit as $n \to \infty$ the average tax revenue is at most $1.5-2\varepsilon$ for any policy of the government. Hence the government cannot improve its payoff over the case where it simply announces a tax rate of $\frac{1}{2}$ that is independent of the observed outcome of the small players' actions.

The intuition for this is much the same as that in the earlier example: Suppose first that the government uses an "anonymous" strategy that depends only on the number of apparent deviations and not on the identities of the deviators. If the number of agents is very large, then there is always some fraction of agents that look as if they had deviated even if every agent actually choose the productive action. This makes it difficult (and in the limit as $N \to \infty$ impossible) for any anonymous policy of the government to induce a different distribution of responses in the situation in which all small players choose the productive action than in a situation in which one small player deviates and chooses the unproductive action. Hence such a policy will be ineffective in deterring small players from choosing $a_i = 0$.

The government can of course use a strategy which singles out certain players that will be punished for deviations. However, the fraction of players that the government is able to coerce into choosing $a_i = 1$ even if the tax rate is above $\frac{1}{2}$ goes to zero as the population size goes to infinity.

It is also of interest to consider the repeated version of this taxation game. Suppose that each period the tax rate is chosen at the same time as the effort decision (that is, a simultaneous move version of the sequential move a game above). In the case of perfect observability of the small players action the folk theorem applies, and so for sufficiently large discount factors the best equilibrium outcome for the government again gives the government a tax revenue of 2 per small player. However, our Proposition 4 shows that if we introduce imperfect observability of the small players actions then a typical small player will play a myopic best response to the tax rate chosen in the current period, and so the highest per capita per period tax revenue that the government can achieve is 1.5.

3. The "One-Shot" Model

This section develops a "one-shot" model in which each player acts only once. There are N small players $i = 1,\dots,N$, and one large player 0. In the first stage of the game, the small players simultaneously choose one of a finite number of actions $a_i \in A_i$; following this the large player chooses an action $b_i \in B$. There are m_a actions for each small player, and m_b for the large player. We let $a = (a_1, a_2, \dots, a_N)$ be a profile of actions for the small players only, and $a_{-i} = (a_1,\dots,a_{i-1},a_{i+1},\dots a_N)$. We denote mixed actions by α_i. As a result of small player i's action one of finitely many individual outcomes $y_i \in Y_i$ results. There are m_y outcomes for each player, with the probability of individual outcome y_i denoted $\rho_i(y_i|a_i)$. We also write $\rho(y|\alpha)$ for the joint probability of all individual outcomes given all the mixed actions. Somewhat loosely we use α for both the profile of mixed actions, and for the joint probability this profile induces over pure action profiles.

The large player moves after observing the outcomes of the small players, but not their actions, so that a strategy for the large player is a map $\sigma_0 : \times_{i=1}^{N} Y_i \to B$. We assume that the utility of each small players depends only on that player's actions and the action of the long-run player; this utility is denoted $u_i(a_i, b)$.

This assumption is more restrictive than necessary, as our main results also hold if each small player's payoffs depends on some aggregates of the play of all small players, so long as the support of these aggregates is bounded independently of the number of players. However, the results do not apply to games where small players interact directly with one another in subsets of small groups.

Our first set of definitions holds fixed the strategy of the large player. The standard game theoretic equilibrium concept is that of Nash equilibrium.

Definition 1 (Nash Equilibrium): A *Nash equilibrium* given σ_0 is a mixed small player profile α such that for all small players i, and all α'_i

$$\sum_{a,y,b} u_i(a_i,b)\sigma_0(b|y)\rho(y|a)\alpha(a) \geq$$
$$\sum_{a_{-i},y,b} u_i(a'_i,b)\sigma_0(b|y))\rho(y|a'_i,a_{-i})\alpha_i(a'_i)\alpha_{-i}(a_{-i})$$

We wish to compare this to the situation where the small players take the distribution of the large player's actions as given.

Definition 2 (Naive Equilibrium): A *naive equilibrium* given σ_0 is a mixed small player profile α such that for all small players i and all α'_i.

$$\sum_{a,y,b} u_i(a_i,b)\sigma_0(b|y)\rho(y|a)\alpha(a) \geq$$
$$\sum_{a,y,b} u_i(a'_i b)\sigma_0(b|y)\rho(y|a)\alpha(a)$$

In other words, the small players "naively" ignore the fact that their actions influence the observed outcomes and hence the play of the large player.

We will show that under certain circumstances Nash equilibrium is like a naive equilibrium. Because we are considering finite games, there will always be some loss to naive play; we will show that this loss must be small when there are many players. Notice that it is possible for the large player to choose his action based on the outcome of just one single small player. Such a small player will have a strong incentive not to play naively. However, we will show that it is only possible for the large player to target a small number of small players this way: specifically we will show that the loss from naive play averaged over the small players is small. This motivates the following definition.

Definition 3 (ε-Naive Equilibrium): An ε-*naive equilibrium* is a mixed profile α such that there are individual losses ε_i with $\sum_{i=1}^{N} \varepsilon_i / N \le \varepsilon$,

and such that for all small players i and all α'_i

$$\sum_{a,y,b} u_i(a_i,b)\sigma_i(b|y)\rho(y|a)\alpha(a) + \varepsilon_i \ge \sum_{a,y,b} u_i(a'_i,b)\sigma_i(b|y)\rho(y|a)\alpha(a)$$

We will also make use of the analogous notion of approximate Nash equilibrium.

Definition 4 (ε-Nash Equilibrium): An ε-Nash equilibrium given σ_i is a mixed small player profile α such that there are individual losses ε_i with $\sum_{i=1}^{N} \varepsilon_i / N \le \varepsilon$

and such that for all small players i and all α'_i,

$$\sum_{a,y,b} u_i(a_i,b)\sigma_i(b|y)\rho(y|a)\alpha(a) + \varepsilon_i \ge \sum_{a_{-i},y,b} u_i(a'_i,b)\sigma_i(b|y))\rho(y|a'_i,a_{-i})\alpha_i(a'_i)\alpha_{-i}(a_{-i})$$

4. Nash Equilibria are Approximate Naive Equilibria

Our goal is to give a theorem showing that Nash equilibria are approximate naive. This proposition, as are all the following propositions, is based on the following mathematical fact that is proven in the appendix.

Lemma A: Let y_i be N independent random variables taking on values in a common finite set, and let p be another random variable with finite support, which may be correlated with the y_i. For any y_i^k, $k = 0,1$

$$\frac{\sum_{i=1}^{N} pr(p \in \Pi | y_i = y_i^0) - pr(p \in \Pi | y_i = y_i^1)}{N}$$
$$\leq \frac{4}{\sqrt{N} \min_{i,k} \sqrt{pr(y_i = y_i^k)}}$$

Roughly, what this lemma says is that the average effect that an individual can have on a future variable is bounded by a number that is small if there are many decision makers and the probability of individual decisions is relatively random. To understand why this lemma is true, it is useful to think of the y's as binary (yes, no) decisions that determine p, a binary outcome that may be positive or negative. One simple case is majority rule: if more than half the decisions are yes then the binary outcome is positive. Obviously in such a setting with many people no one is likely to be decisive, so the lemma is not surprising. On the other hand, take the case where the rule is that the outcome is positive if there are an even number of yes's and negative if there is an odd number. *Ex ante*, that is, conditional on the all the other decisions, everyone is completely decisive: each individual decision completely determines the outcome. How can our Lemma possibly hold in this setting? The answer is that the lemma is about *ex post* decisiveness; it is about the predictability of outcomes. If decision makers do not know in advance each other's decision, but have a common prior that every other individual has a 50-50 chance of saying yes, then no individual has any effect on the probability that the outcome is positive provided there are two or more player. Notice that in this 50-50 case we do not need large numbers of players. However if the odds of yes are not 50-50, say they are 75-25, then with an even number of players the chance of a positive outcome is greater if an individual player says yes (and conversely with an odd number of players), but according to the lemma, this chance must get smaller as the number of players gets large.

We can now state and prove our main proposition about naive equilibria.

Proposition 1: If α is an ε-Nash equilibrium, then it is an $\varepsilon+\varepsilon'$-naive equilibrium with

$$\varepsilon' = \frac{1}{\sqrt{N}}\left[\frac{4m_b m_y \max_{i,a_i,b} |u_i(a_i,b)|}{\min_{i,y_i}\sqrt{\rho_i(y_i|\alpha_i)}}\right];$$

Discussion: The key fact to note is that the ε' of this proposition shrinks to 0 as $N \to \infty$, provided that the ρ_i have uniformly full support, and that payoffs are uniformly bounded. We should emphasize the importance of holding the number m_b of large player's actions fixed for this conclusion. If instead m_b increases rapidly with the number of small players, then ε' need not go to zero. For example, if . the large player can punish each small player individually (so $m_b \geq 2^N$), it is clear that a Nash equilibrium need not be even approximately naive, as each small player can correctly anticipate a personal punishment for deviating from the equilibrium profile. One question of interest is exactly how rapidly we can let the number of actions grow with the number of small players. According to the bounds above $\frac{m_k}{\sqrt{N}}$ must go to zero. There is a substantial gap between $\frac{m_k}{\sqrt{N}} \to 0$ and $m_k \geq 2^N$, but the bound we develop here does not provide any insight into what might happen in this range, nor do we have good examples.

Proof:

Fix an ε-Nash equilibrium α. Let α'_i be a best response for player i to α under the "naive" assumption that deviations won't change the play of the large player. Then player i's loss at a under naive expectations is

$$\eta_i = \sum_{a,y,b} u_i(a'_i,b)\sigma_0(b|y)\rho(y|a)\alpha(a) - \sum_{a,y,b} u_i(a_i,b)\sigma_0(b|y)\rho(y|a)\alpha(a).$$

Since α is an ε-Nash equilibrium,

$$\begin{aligned}&\sum_{a,y,b} u_i(a_i,b)\sigma_0(b|y)\rho(y|a)\alpha(a) \geq \\ &\sum_{a_{-i},y,b} u_i(a'_i,b)\sigma_0(p|y)\rho(y|a'_i,a_{-i})\alpha_i(a'_i)\alpha_{-i}(a_{-i}) - \varepsilon_i\end{aligned}$$

from which we see that

$$\eta_i \leq \sum_{a,y,b} u_i(a'_i b)\sigma_0(b|y)\rho(y|a)\alpha(a) - \sum_{a_{-i},y,b} u_i(a'_i,b)\sigma_0(b|y)\rho(y|a'_i,a_{-i})\alpha_i(a'_i)\alpha_{-i}(a_{-i}) + \varepsilon_i.$$

In words, the loss under naive expectations is bounded by the sum of the loss under Nash expectations and the approximation error.

From this we may now compute $\eta = \sum \eta_i$

$$\eta \le \frac{1}{N}\sum_{i}^{N} \sum_{y,b} u_i(a_i',b)\sigma(b|y)\rho(y_{-i}|\alpha_{-i})[\rho(y_i|\alpha_i)-\rho(y_i|a_i')] + \varepsilon$$

$$= \frac{1}{N}\sum_{i}^{N} \sum_{y,b} u_i(a_i',b)[\sigma(b|y)-\sigma(b|y_{-i},y_i')]\rho(y_{-i}|\alpha_{-i})[\rho(y_i|\alpha_i)-\rho(y_i|a_i')] + \varepsilon$$

$$< \frac{1}{N}\sum_{i}^{N} \max_{i,a,b} u_i(a_i,b)\sum_{y,b}|\sigma(b|y)-\sigma(b|y_{-i},y_i')|\rho(y_{-i}|\alpha_{-i}) + \varepsilon$$

$$= \frac{1}{N}\sum_{i}^{N} \max_{i,a,b} u_i(a_i,b)\sum_{b,y_i}|\sum_{y_{-i}} \sigma(b|y)\rho(y_{-i}|\alpha_{-i}) - \sum_{y_{-i}} \sigma(b|y_{-i},y_i')\rho(y_{-i}|\alpha_{-i})| + \varepsilon$$

$$= \max_{i,a,b} u_i(a_i,b)\sum_{b,y_i} \frac{1}{N}\sum_{i}^{N}|\sum_{y_{-i}} \sigma(b|y)\rho(y_{-i}|\alpha_{-i}) - \sum_{y_{-i}} \sigma(b|y_{-i},y_i')\rho(y_{-i}|\alpha_{-i})| + \varepsilon$$

$$\le m_b m_y \max_{i,a,b} |u_i(a_i,b)| \max\left[\frac{1}{N}\sum_{i}^{N}|\sum_{y_{-i}} \sigma(b|y)\rho(y_{-i}|\alpha_{-i}) - \sum_{y_{-i}} \sigma(b|y_{-i},y_i')\rho(y_{-i}|\alpha_{-i})|\right] + \varepsilon$$

The desired result now follows from Lemma A in the Appendix, which bounds the term in square brackets. Intuitively, this expression measures the average effect any one small player can have on the play of large player. Since the number of actions the large player has available is fixed, the extent of this average reaction shrinks as the number N of small players grows.

□

5. Contingent and Uncontingent Commitments

To illustrate the importance of the small players playing naively, we consider the best payoff the large player can obtain if he is able to publicly commit himself to a strategy σ_0 before the small players act. If the large player were able to observe the actions of the small players, she could use this commitment power by threatening that any individual who deviates from the desired action profile will be held to his minmax level. However, when actions are observed with noise and there are many small players, the best payoff the large player can obtain from this commitment to a contingent strategy is the payoff the player could obtain from an uncontingent commitment to always play a fixed action. Thus in the limit the large player cannot do better than her Stackelberg payoff of the stage game, even though the ability to make contingent commitments leads to higher payoffs in many other settings.

We now specialize to the case where the small players are all identical, so that $u_i(a_i, b) = u(a_i, b)$, and the distribution over individual outcomes has full support $\rho_i(y_i; a_i) = \rho(y_i; a_i) > 0$.

For a fixed distribution β over large player's actions, we say that α_i is a best response to β if it maximizes the small player's expected payoff $u_i(\alpha'_i, \beta)$ over all choices α'_i. We say that α_i is *individually rational with respect to* β if $u_i(\alpha_i, \beta)$ is at least the minmax payoff, which is

$$\min_{\beta'} \max_{\alpha_i'} u_i(\alpha_i', \beta').$$

The *Stackelberg payoff for the large player*, denoted u_0^S, is the highest payoff the large player can get from publicly committing to a fixed action (or equivalently from a perfect equilibrium of the Stackelberg game where the large player moves first and chooses an action which the small players observe before they act.) Formally, this payoff is given by

$$u_0^S = \max_\beta [u_0(\alpha_i, \beta) | \alpha_i \text{ is a best response to } \beta].$$

The *best feasible payoff*, denoted $\hat{u}_0$, equals

$$\max_{\beta}[u_1(\alpha_2,\beta)|\alpha_2 \text{ is individually rational with respect to } \beta].$$

This is clearly at least as large as the Stackleberg payoff, and in general is larger.

If the large-player can directly observe the mixed strategies of the small-players, then there is a precommitment strategy and best-response for the small players such that in the limit as the number of small players grows without bound, the large player indeed gets $\hat{u}_1$. To see this, choose a pair that gives the long-run player the desired payoff. The equilibrium consists of all small players following the given mixed strategy and, provided the small players do so, the large player doing likewise. If any small-player deviates, the large player minmaxes all small players. Obviously no small player has any incentive to deviate. Moreover, by the law of large numbers the average small player outcome converges to the expected outcome, so as the number of small players becomes infinite, the payoff to the large player approaches the desired quantity.

Our key result is that because of the naive behavior of the small players, if the large-player observes only the small-player outcomes, he can do no better in the limit than the Stackleberg payoff. This is analogous to a similar result in Levine and Pesendorfer(1994), who restrict the large player to play strategies that depend only on the average play of the small players.

Proposition 2: Let σ_1^N be a sequence of strategies for the large player, and suppose that α_2^N are a Nash equilibrium with respect to σ_1^N. If the distribution of small player outcomes has full support, then

$$\limsup_{N\to\infty} \sum_{a,b,y} u_1(a,b)\sigma_1^N(b|y)\rho(y|\alpha^N)\alpha^N(a) \le u_1^S.$$

Proof: From proposition 1 it follows that α_i^N form an ε^N-naive equilibrium where $\varepsilon^N \to 0$. By definition this means that $\sum_i \frac{\varepsilon_i^N}{N} \to 0$, and implies that there is a sequence of sets of players I^N and numbers $\tilde{\varepsilon}^N$ with the property that $\frac{\# I^N}{N} \to 1$ and if $i \in I^N$

$$\sum_{a,y,b} u_i(a_i, b)\sigma_i^N(b|y)\rho(y|a)\alpha^N(a) - \max_{a_i'} \sum_{a,y,b} u_i(a_i', b)\sigma^N(b|y)\rho(y|a)\alpha^N(a) \le \tilde{\varepsilon}^N$$

This in turn implies that there is $\bar{\varepsilon}^N \to 0$ and $|\alpha_i^N - \hat{\alpha}_i^N| \le \bar{\varepsilon}^N$ where $\hat{\alpha}_i^N$ is a best response to the mixed action $\beta^N \equiv \sum_y \sigma_i^N(b|y)\rho(y|a)\alpha^N(a)$ that the large player uses on the equilibrium path. Since an increasing fraction of the population is playing a best response an increasing fraction of the time, applying the weak law of large numbers to the large player's payoff yields the desired result.

□

6. Finitely Repeated Games

Reputation effects are often analyzed in models with a single patient player who facing many myopic opponents. One justification for the myopia assumption is that the myopic players are really small players who play "as if" they were myopic because they are too small to have a significant effect on other players' future play. In the next two sections we give a precise interpretation of this.

We begin by considering the finitely repeated case; in the next section we examine an infinite horizon model in which the small players discount. Every period, each small player and the large player simultaneously take an action. (Note that this differentiates the model from that of previous sections, where the large player moves after the small ones.) Actions of the small players continue to be denoted by $a_i \in A_i$ actions of the large player are denoted by $x \in X$. Let m_X denote the number of actions of the large player. Periods are denoted by $t = 1, \ldots, T < \infty$. As before, as the result of a small player's action one of finitely many individual outcomes $y_i \in Y$ occurs. The probability of each individual outcome is $\rho(y_i; a_i)$. All players, when taking an action in period t observe these outcomes and the play of the large player in all previous periods. For a sequence of actions of a small player and the large player and outcomes (a_i^t, y_i^t, x^t) the payoff of the small player is given by:

$$\sum_{t=1}^{T} u_i^t(a_i^t, x^t)$$

Let $y = (y_1, \ldots, y_N)$ be a profile of outcomes for all small players. The payoffs of the large player are not relevant for the time being as the strategy of the large player is held fixed.

We denote by h_i^t the private history of small player i and by h^t the public history up to but not including period t. Pure strategies are maps from private and public histories to actions. Let $(\tau_1, \ldots, \tau_N, \sigma_l)$ denote a mixed strategy profile for the small players and the large player respectively. The assumption that the small players play myopically is captured in the next definition.

Definition 5: Small player profile $(\tau_1, \ldots, \tau_N)$ is an ε myopic best response to σ_0 if for each player i there is an $\varepsilon^i > 0$ such that for every history h^t that has positive probability

$$\sum_{a,x} u_i(a_i,x)\tau_i(h^i,h_i^i)[a_i]\sigma(h^i)[x]-\varepsilon_i \geq$$
$$\sum_{a,x} u_i(a_i',x)\sigma(h^i)[x]$$

for all $a_i' \in A$ and $\sum_i \varepsilon^i / N \leq \varepsilon$.

Our goal is to show that approximate best-responses are approximately myopic.

Proposition 3: If $(\tau_1,\dots,\tau_N)$ is an ε best response to σ, then $(\tau_1,\dots,\tau_N)$ is an $\varepsilon-\varepsilon'$ myopic best response to σ, where

$$\varepsilon' = \frac{1}{\sqrt{N}}\left[\frac{4T^3 m_X m_Y \max_{i,a_i,x,t} u_i^t(a_i,x)}{\min_{a} \sqrt{\rho_i(y_i|a_i)}}\right]$$

Remark: The crucial point is that $\varepsilon' \to 0$ as $N \to \infty$, provided that the ρ_i have uniformly full support, and that payoffs are uniformly bounded.

Proof: We compute $\varepsilon'(t)$ recursively from the end of the game. First, in period T it is obvious that $\varepsilon'(T)=0$. Now we suppose we are given $\varepsilon'(t+1)$ in a subgame that is reached with positive probability starting in period $t+1$.

In period t every small player i chooses action a_i. Conditional on a_i being realized equilibrium strategies induce a probability distribution over actions of the large player in each period and over public outcomes. Let $b^t=(x^{t+1},\dots,x^T)$ be the actions for the large player from period $t+1$ on and let $V_i^{T-t+1}(b^t)$ be the payoff of the small player i if he chooses a myopic best response in every period and if the large player plays b^t. By inductive hypothesis, the actual payoff of the small player in the continuation game starting in period $t+1$ is in the interval $[V_i^{T,t+1}(h^t), V_i^{T,t+1}(h^t)-\varepsilon_i-\varepsilon_i'(t-1)]$. Define

$$(*) \qquad U_i^t(a_i, h^t) = \sum_x u_i^t(a_i, x^t)\sigma(h^t)[x] + V_i^{T,t+1}(h^t)$$

Since $(\tau_1,\ldots,\tau_N)$ is an ε-best response it follows that $(\tau_1^t(h^t,h_1^t),\ldots,\tau_N^t(h^t,h_N^t))$ is an $\varepsilon+\varepsilon'(t-1)$ Nash equilibrium in the game (with only small players) defined by the payoffs in (*). Proposition 1 then implies that $(\tau_1^t(h^t,h_1^t),\ldots,\tau_N^t(h^t,h_N^t))$ is an $\varepsilon+\varepsilon'(t)$ price taking equilibrium where

$$\varepsilon'(t) = \varepsilon'(t-1) + \frac{1}{\sqrt{N}}\left[\frac{4(Tm_X)m_Y \max_{i,a_i,x,t}(Tu_i^t(a_i,x))}{\min_{y_i}\sqrt{\rho_i(y_i|a_i)}}\right].$$

❐

7. Infinitely Repeated Games with Discounting

We now specialize the model of the previous section to discounted payoffs, but consider the case in which the horizon is infinite. We assume

$$u_i^t(a_i, x) = (1-\delta)\delta^t u_i(a_i, x),$$

and $T=\infty$.

Proposition 4: If $(\tau_1,\ldots,\tau_N)$ is a best response to σ_0

then $(\tau_1,\ldots,\tau_N)$ is an ε-myopic best response to σ_0 where

$$\varepsilon = f_\delta\left(\frac{1}{\sqrt{N}}\left[\frac{4m_X m_Y}{\min_a \sqrt{\rho_i(y_i|a_i)}}\right]\right)\max_{i,a_i,x,t} u_i^t(a_i,x)$$

$$f_\delta(z) = \delta^{z^{-1/6}} + z^{1/2}$$

Remark: Again this implies that $\varepsilon \to 0$ as $N \to \infty$, provided that the ρ_i have uniformly full support, and that payoffs are uniformly bounded.

Proof: Since $(\tau_1,\ldots,\tau_N)$ is a best response to σ_0 for every history $(h^t, h_1^t,\ldots,h_N^t)$ that is reached with strictly positive probability $(\tau_1,\ldots,\tau_N)$ is an ε'-best response to σ_0 in the T-period game beginning with that history, where

$$\varepsilon' = \delta^T \max_{a_i,x,t} u_i(a_i,x).$$

We now apply Proposition 3 to find that $(\tau_1,\ldots,\tau_N)$ is an ε myopic best response where

$$\varepsilon = \left(\delta^T + \frac{1}{\sqrt{N}}\left[\frac{4T^3 m_X m_Y}{\min_a \sqrt{\rho_i(y_i|a_i)}}\right]\right)\max_{i,a_i,x,t} u_i^t(a_i,x)$$

Choosing

$$T = \left(\frac{1}{\sqrt{N}}\left[\frac{4m_X m_Y}{\min_a \sqrt{\rho_i(y_i|a_i)}}\right]\right)^{-1/6}$$

gives the desired result.❐

8. Implication for Reputation Effects Models

Propositions 3 and 4 give conditions under which myopic play is approximately optimal for fixed beliefs about the large player's strategy. All that matters for these results is the expected distribution of the large player's actions contingent on various observations, so the propositions apply to models where the large players has several possible "types", as in the reputation effects literature. If we consider such a reputation effects model, and let the number of small players grow as in propositions 3 and 4, then a limiting argument analogous to that of proposition 2 would show that the limiting value of the worst Nash equilibrium payoff for the large player equals the lower bound on this payoff in the "limit model" where the small players are exactly myopic. Consequently, the payoff bounds obtained in Fudenberg and Levine (1989,1992) for exactly myopic opponents are also the limiting value of the payoffs when facing many small opponents whose actions are observed with noise. Thus, as in the one-shot case, our model provides a justification for results derived under the assumption that small players act as if they are negligible.

Appendix A

Lemma A: Let y_i be N independent random variables taking on values in a common finite set, and let p be another random variable with finite support, which may be correlated with the y_i. For any y_i^k, $k = 0,1$

$$\frac{\sum_{i=1}^{N} pr(p \in \Pi | y_i = y_i^1) - pr(p \in \Pi | y_i = y_i^0)}{N} \le \frac{4}{\sqrt{N}\min_{i,j}\sqrt{pr(y_i = y_i^j)}}$$

Proof:

$$D \equiv \frac{\sum_{i=1}^{N} pr(p \in \Pi | y_i = y_i^1) - pr(p \in \Pi | y_i = y_i^0)}{N}$$

$$= \frac{\sum_{i=1}^{N} \frac{pr(p \in \Pi, y_i = y_i^1)}{pr(y_i = y_i^1)} - \frac{pr(p \in \Pi, y_i = y_i^0)}{pr(y_i = y_i^0)}}{N}$$

$$= \frac{\sum_{i=1}^{N} \frac{\sum_{\{y: y_i = y_i^1\}} pr(p \in \Pi, y)}{pr(y_i = y_i^1)} - \frac{\sum_{\{y: y_i = y_i^0\}} pr(p \in \Pi, y)}{pr(y_i = y_i^0)}}{N}$$

$$= \sum_y pr(p \in \Pi, y) \sum_{n=1}^{N} \begin{cases} \frac{1}{pr(y_i = y_i^1)} & y_i = y_i^1 \\ -\frac{1}{pr(y_i = y_i^0)} & y_i = y_i^0 \\ 0 & y_i \ne y_i^1, y_i^0 \end{cases}$$

Consequently, setting

$$z_i^j \equiv \begin{cases} \dfrac{1}{pr(y_i = y_i^j)} & y_i = y_i^j \\ 0 & y_i \neq y_i^j \end{cases}$$

$$\begin{aligned}
&D \\
&= \frac{1}{N}\sum_y pr(p \in \Pi)\sum_{i=1}^N \left(z_i^1(y) - z_i^0(y)\right) \\
&\leq \frac{1}{N}\sum_y pr(p \in \Pi, y)\left|\sum_{i=1}^N z_i^1(y) - \sum_{i=1}^N z_i^0(y)\right| \\
&\leq \frac{1}{N}\sum_y pr(y)\left|\sum_{i=1}^N z_i^1(y) - \sum_{i=1}^N z_i^0(y)\right| \\
&\leq \frac{1}{N}\sum_y pr(y)\left(\left|\sum_{i=1}^N z_i^1(y) - 1\right| + \left|\sum_{i=1}^N z_i^0(y) - 1\right|\right) \\
&= E\left(\left|\frac{1}{N}\sum_{i=1}^N z_i^1 - Ez_i^1\right| + \left|\frac{1}{N}\sum_{i=1}^N z_i^0 - Ez_i^0\right|\right) \\
&\leq 2\left\{E\left|\frac{1}{N}\sum_{i=1}^N z_i^1 - Ez_i^1\right|^2\right\}^{1/2} + 2\left\{E\left|\frac{1}{N}\sum_{i=1}^N z_i^0 - Ez_i^0\right|^2\right\}^{1/2} \\
&= 2\left\{\sum_{i=1}^N \frac{1 - pr(y_i = y_i^1)}{N^2 pr(y_i = y_i^1)}\right\}^{1/2} + 2\left\{\sum_{i=1}^N \frac{1 - pr(y_i = y_i^0)}{N^2 pr(y_i = y_i^0)}\right\}^{1/2} \\
&\leq \frac{4}{\sqrt{N}\,\min_{i,j}\sqrt{pr(y_i = y_i^j)}}
\end{aligned}$$

which proves the claim.

□

Bibliography

Celentani, M. (1991): "Reputation in General Deterministic Stage Games," forthcoming, *Economic Theory*.

Celentani, M. and W. Pesendorfer (1992): "Reputation in Dynamic Games," CMSEMS Discussion Paper No. 1009, Northwestern University.

Dubey, P. and M. Kaneko (1985): "Information Patterns and Nash Equilibria in Extensive Games II," *Mathematical Social Science* 10, 247-262.

Fudenberg, D. and D. Levine (1988): "Open-Loop and Closed-Loop Equilibria in Dynamic Games with Many Players," *Journal of Economic Theory*, 44, 1-18.

Fudenberg, D. and D. Levine (1989): "Reputation and Equilibrium Selection in Games with a Single Patient Player," *Econometrica* 57, 759-78.

Fudenberg, D. and D. Levine (1992): "Maintaining a Reputation When Strategies Are Unobserved," *Review of Economic Studies*, 59, 561-580.

Green, Edward J. (1980): "Noncooperative Price Taking in Large Dynamic Markets," *Journal of Economic Theory*, 22, 155-181.

Kreps D. and R. Wilson, (1982): "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253-79.

Levine, D. and W. Pesendorfer (1994): "When are Agents Negligible," forthcoming *American Economic Review*.

Milgrom P. and J. Roberts, (1982): "Predation, Reputation, and Entry Deterrence," *Econometrica* 50, 443-60.

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

Schmidt, K.(1993), "Reputation and Equilibrium Characterization in Repeated Games of Conflicting Interest," *Econometrica* 61, 325-351