When Are Nonanonymous Players Negligible?

Citation

Fudenberg, Drew, David K. Levine, and Wolfgang Pesendorfer. 1998. When are nonanonymous players negligible?. *Journal of Economic Theory* 79, no. 1: 46-71.

Published Version

<http://dx.doi.org/10.1006/jeth.1997.2373>

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:3203775>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

When are Non-Anonymous Players Negligible?^{*}

Drew Fudenberg, David Levine and Wolfgang Pesendorfer^{**}

December 7, 1996

© This document is copyrighted by the authors. You may freely reproduce and distribute it electronically or in print, provided it is distributed in its entirety, including this copyright notice.

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 SBR-9424013, SBR-9320695, SBR-9409180, and the UCLA Academic Senate. We also thank the Editor and two anonymous referees for useful suggestions and comments.

^{**} 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 not too large relative to 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 Stackelberg 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 each agent's actions can be observed, but in the continuum limit this information "vanishes."¹ Consequently the games with finitely many 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 (1995) show that the negligibility assumption holds approximately in all equilibria of games with small but non- infinitesimal agents provided that the agents are anonymous and their aggregate play is observed with some noise. More formally, these papers assume that the distribution of observed signals depends on some aggregate statistic of the small players' actions, such as the sum, and that this distribution is non-degenerate.²

In this paper, we drop the anonymity assumption, and suppose that the large player receives a noisy signal of each small player's action. Why does the lack of anonymity matter? If the large player knows the small players by name, she may pick on a particular player or players, and punish and reward them to manipulate their behavior.

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 Note that the anonymity condition on its own is not sufficient for this conclusion, , since, with a finite number of players, there is a change in the 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 will not change the aggregate statistic, and so there cannot be equilibria that rely on punishments. Intuitively, we would expect this discontinuity to vanish if individual play is not perfectly observable, so that the equilibria in the finite player case would be similar to that in the continuum limit. Levine and Pesendorfer (1995) show that this is indeed the case. See also Dubey and Kaneko (1985) who assume that small deviations are unobservable, but large deviations are observed perfectly.

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 large player can reward or punish each small player independently, 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 her 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 everyone ” and “punish everyone.” 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, she can induce them all to work by threatening to punish everyone 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. Suppose first that the large player uses a strategy that is feasible when the small players are anonymous, so that the probability of punishment depends only on the number of observations of the outcome “shirk”, and not on the identities of the apparent shirkers. Since there will always be a fraction of size approximately ϵ that appears to have shirked even if they did not, it is difficult (and in the limit as $N \rightarrow \infty$ impossible) for any anonymous policy of the large player to induce a distribution of responses that is different if all small players work than if one small player shirks while the others work, so the small players cannot be induced to work.

If the small players are not anonymous, and the large player’s strategy space is large enough to permit separate rewards and punishments for each small player, the strategy “punish small player i if it looks as though that player shirked” *will* induce every

player to work, so that the introduction of noise does not imply negligibility if players are not anonymous *and* personalized punishments are feasible.

However, if the large player only has two actions available, she is unable to implement these sorts of personalized punishments, and we show that the result is as if the small players were anonymous. This is true even though the large player has the option of basing her action solely on the observed outcomes of a subset of the small 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 mathematical lemma that has an independently interesting interpretation. Consider a discrete random variable with fixed support, such as categorical³ data on per capital GNP or tax rates. The lemma says that the average effect that individuals expect to have on the random variable is small if the number of individuals is large and the signals of each individual's action are somewhat noisy. In a sense this lemma implies that there is a limit on the number of "important" individuals.

In section 2 we set up the basic model, and define the concept of a "naïve" equilibrium in which the small players act as if most of them ignore contingent responses by the large player. Section 3 discusses the mathematical lemma summarized in the previous paragraph, and establishes the results about naïve equilibrium that are used in our characterizations of equilibria.

³ That is, the data is reported in categories such as \$10,000 to \$11,999.

Instead of proceeding directly to that characterization, section 4 examines the extent to which the result of section 3 can be obtained in limits where the number of large player's actions grows along with the number of small players. Readers interested primarily in the application of our results to the equilibria of games may want to skip this section, as it is not used in the remainder of the paper.

Our first application of the theory is in section 5, which considers the case of one-shot games in which a single large player may choose a precommitment strategy against a number of small players.⁴ The essential point we make is similar to that in Levine and Pesendorfer (1995): without noise the large player can appropriate all the surplus, but with noisy observation of the small players, she can only get the Stackelberg payoff.

The second application applies the results to dynamic games. We show that when the play of non-myopic small players is observed with noise, to a good 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. Section 6 considers the finitely repeated case, and section 7 considers the case of infinite repetition with discounting.

Section 8 contains a concluding discussion of our results.

2. 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

⁴While 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 (1996) and others have extended the scope of this result in a variety of ways.

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 $b_0 \in B$. 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 $a_{-i} = (a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_N)$, and so forth. 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. The probability of individual outcome y_i is denoted $\rho_i(y_i | a_i)$. We also write $\rho(y | \alpha)$ for the joint probability of the vector of individual outcomes given the profile of 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 $\sigma_0: \times_{i=1}^N Y_i \rightarrow \Delta(B)$, where $\Delta(B)$ is the set of probability distribution over 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 $u_i(a_i, b)$. The utility of the large player $u_0(a_1, a_2, \dots, a_N, b)$, on the other hand, depends on the actions of all the small players as well as the large player's own action.

This assumption is more restrictive than necessary, as our main results also hold if each small player's utility 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.

Definition 1 (best response profiles): A best response profile to σ_0 is a mixed small player profile α such that for all small players i , and all a'_i ,

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

We wish to compare this to the situation where the small players “naively” ignore the fact that their actions influence the observed outcomes and hence the play of the large player.

Definition 2 (naive best response profile): A naive best response profile to σ_0 is a mixed small player profile α such that for all small players i and all a'_i ,

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

We will show that under certain circumstances a best response profile is like a naive best response profile. Because we are considering games with finitely many players, 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 the definitions here do not impose anonymity; in particular they allow the large player to choose her 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, it is only possible for the large player to target a small number of small players this way; more precisely, 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 best response profile): An ε -naive best response given σ_0 is a mixed profile α for the small players such that there are individual losses $\varepsilon_i \geq 0$ with $\sum_{i=1}^N \varepsilon_i / N \leq \varepsilon$, and such that for all small players i and all a'_i ,

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

We will also make use of the analogous notion of approximate best response.

Definition 4 (ε -best response profile): An ε -best response profile given σ_0 is a mixed small player profile α for the small players such that there are individual losses $\varepsilon_i \geq 0$ with $\sum_{i=1}^N \varepsilon_i / N \leq \varepsilon$ and such that for all small players i and all a'_i ,

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

3. Best Response Profiles are Approximately Naive

Our goal in this section is to give a theorem showing that Nash equilibria are approximate naïve equilibria. This proposition, as well as the ones that follow, is based on the following mathematical fact that is an immediate corollary of Lemma A' in the Appendix.

Lemma A: Let y_i be N independent random variables each taking on the values $\{y_i^0, y_i^1\}$, and let p be another random variable with finite support, which may be correlated with the y_i . Then

$$\frac{\sum_{i=1}^N pr(p \in \Pi | y_i = y_i^1) - pr(p \in \Pi | y_i = y_i^0)}{N} \leq \frac{2}{\sqrt{N} \min_{i,k} \sqrt{pr(y_i = y_i^k)}}.$$

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 outcomes 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) “votes” that determine $p \in \{-1, 1\}$. Further, suppose that the minimum probability that any individual assigns to any vote is bounded away from zero, so that the bound in the lemma has some force.⁵

One simple case is majority rule: if more than half the votes are yes then $p = 1$. If there are many voters no individual is likely to be decisive, and so the lemma is not

⁵ In our model, the “votes” are the outcomes, which are noisy signals of the players’ true actions, but the lemma holds for more general sources of noise, including “trembles” by the players (so that their true actions are not their intended ones) and more simply the use of non-degenerate mixed strategies by each player. All that is required is that no small player be able to predict the realized value of the signals that determine the aggregate outcome.

surprising. On the other hand, take the case where the rule is that $p = 1$ if there is an even number of yes's and $p = -1$ if there is an odd number of yes votes. *Ex post*, that is, conditional on the realized values of all votes, every player is 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 ante* decisiveness, i.e., the predictability of outcomes. If decision makers do not know one another's decisions in advance, but have a common prior that every other individual has probability $\frac{1}{2}$ chance of voting yes, independent of the others, then no individual has any effect on the probability that $p = 1$ whenever there are two or more players, and so each individual has negligible influence even away from the many-player limit. However if the probability of each player voting yes is $\frac{3}{4}$ then when there are only two players the chance that $p = 1$ is greater if an individual player votes yes than if the same individual votes no, and for any finite number of players we would expect that each individual's decision has some influence. However, Lemma A shows that the size of this influence converges to zero as the number of players gets large.

To make use of Lemma A, we must calculate the loss from naïve behavior. This is given in the next lemma, proven in the Appendix. In the following we denote by $\mathbf{B}$ a partition of the large player's action set and by $B \in \mathbf{B}$ a subset of the large player's actions contained in the partition.

Lemma B: Let $(\mathbf{B}_1, \dots, \mathbf{B}_N)$ be a collection of partitions of the large player's action set. If α is an ε -best response profile, it is an $\varepsilon + \varepsilon'$ -naïve best response profile with

$$\begin{aligned} \varepsilon' \leq & 2 \max_{i, a_i, b} |u_i(a_i, b)| \frac{1}{N} \sum_{i=1}^N \sum_{B \in \mathbf{B}_i} \max_{y_i, y'_i} |pr(b \in B | y_i, \alpha) - pr(b \in B | y'_i, \alpha)| \\ & + 4 \frac{1}{N} \sum_{i=1}^N \max_{a_i, B \in \mathbf{B}_i} [\max_{b \in B} u_i(a_i, b) - \min_{b \in B} u_i(a_i, b)] \end{aligned}$$

This bounds the naïve loss by two terms. The first is an average of conditional probabilities which is bounded by Lemma A. The second measures the utility variation

of small players as the large player's action varies within a member of the partition.

Thus the bound has force only if the partitions are constructed in such a way that for most players the utility variation within each member of the partition is small.

The easiest application of Lemma's A and B is when all the elements of the partitions $\mathbf{B}_i$ in Lemma B are sets consisting of a single action. In this case the final term in Lemma B drops out, as there is no utility variation at all within partitions. The bound from Lemma B is then

$$\varepsilon' \leq 2 \max_{i,a_i,b} |u_i(a_i,b)| \frac{1}{N} \sum_i \sum_b \max_{y_i,y'_i} |pr(b|y_i,\alpha) - pr(b|y'_i,\alpha)|.$$

Recall that $\rho_i(y_i|\alpha_i)$ is the probability of outcome y_i under the mixed action α_i , so that

$$\rho_i(y_i|\alpha_i) = \sum_{a_i} \rho_i(y_i|a_i) \alpha_i(a_i).$$

We can apply Lemma A to conclude:

Proposition 1: If α is an ε -best response profile, it is an $\varepsilon + \varepsilon'$ -naive best response profile with $\varepsilon' = \frac{1}{\sqrt{N}} \left[\frac{4m_b \max_{i,a_i,b} |u_i(a_i,b)|}{\min_{i,y_i} \sqrt{\rho_i(y_i|\alpha_i)}} \right]$.

The key fact to note is that the ε' of this proposition shrinks to 0 as $N \rightarrow \infty$, provided that the ρ_i have uniformly full support, that payoffs are uniformly bounded, and that the number of large player's actions, m_b , remains fixed.

If m_b increases rapidly with the number of small players, then ε' in Proposition 1 need not go to zero as the number of players goes to infinity. However, some cases in which m_b is large either for all N or in the limit can be accommodated by choosing a coarser partition than the singleton sets used in Proposition 1.

As an example, consider the case in which the large player's action space is a compact subset of a finite-dimensional Euclidean space, and the small players' utility functions are uniformly continuous in the action of the large player. In this case the set of utility profiles for the small players that can be achieved by varying the large player's

actions can be well approximated by a finite set. Specifically, for every ε'' we may find a partition $\mathbf{B}$ of the large players action set into $m_{\mathbf{B}}(\varepsilon'')$ components so that for every player i

$$\max_{b \in B} u_i(a'_i, b) - \min_{b \in B} u_i(a'_i, b) \leq \varepsilon''$$

Note that, as in Proposition 1, the partition is independent of the player. The following proposition follows now immediately from Lemmas A and B

Proposition 1': If α is an ε -best response profile, it is an $\varepsilon + \varepsilon'$ -naive best response profile with $\varepsilon' = \frac{1}{\sqrt{N}} \left[\frac{4m_B(\varepsilon'') \max_{i, a_i, b} |u_i(a_i, b)|}{\min_{i, y_i} \sqrt{\rho_i(y_i | \alpha_i)}} \right] + 4\varepsilon''$.

4. Games with Individualized Punishment

If the large player can punish each individual small player without affecting the payoff of other small players, it is clear that a best response need not be even approximately naive, as each small player can correctly anticipate a personal punishment for deviating from the equilibrium profile. This raises the question of precisely how well targeted punishments must be before best responses are no longer approximately naïve. Readers interested primarily in the application of our results to the equilibria of games may want to skip this section, as it is not used in the remainder of the paper.

As an illustration and partial answer to this question, in this section, we focus attention on games with the following simple structure. The large player's action is the choice of a subset of small players (including the null set). The interpretation is that the chosen subset is the set of small players that will be punished. Small player utility depends only on the action of the large player. Specifically, let Ψ denote a class of subsets of small players with generic element ψ . For every $\psi \in \Psi$ there is a corresponding large player action b_ψ ; each small player's utility can take on one of two values, 1 if the small player is not in ψ or 0 if she is in ψ .

Each class of subsets gives rise to a particular game. In some of these games best responses will be approximately naïve, and in some they will not. The question we wish to pose is for which classes of subsets Ψ (or more precisely the corresponding games) best responses will be approximately naïve. As we already noted, if Ψ consists of all 2^N subsets of small players, best responses will typically not be approximately naïve. However, best responses need not be approximately naïve even if the number of subsets of small players that can be punished is not nearly this large. Consider a partition of the set of small players so that each element of the partition except one has k players. The class Ψ consists of unions of the elements of this partition. If Ψ is equal to *all* the sets that may be obtained through unions of elements of the partition then there are roughly $(2^{1/k})^N$ sets in Ψ . If k is fixed then the best response profiles may fail to be approximately naïve for all N .⁶ This implies that if the number of actions increases with the number of small players at any exponential rate, then best response profiles need not be naïve in the limit.

To find a positive result we restrict the size of the sets in the class Ψ . Let $\kappa(\Psi)$ denote the largest number of players in any member of Ψ ; that is the greatest number of players that can be punished simultaneously.

Proposition 1'': If α is an ε -best response profile, it is an $\varepsilon + \varepsilon'$ -naïve best response profile with $\varepsilon' \leq 4 \max_{a_i, b} |u_i(a_i, b)| \left[\frac{2}{\sqrt{N} \min_{i, y_i} \sqrt{\rho_i(y_i | \alpha_i)}} + \frac{\kappa(\Psi)}{N \min_{i, y_i} \rho_i(y_i | \alpha_i)} \right]$.

Proof: see Appendix.

⁶ To see this note that if the large player punishes every small player who appears to deviate from his prescribed action then the small players action has a non-negligible influence on his probability of being punished for all N as long as the fraction of the population punished in equilibrium is not close to one. Thus if $k \cdot \rho_i(y_i | \alpha_i)$ is sufficiently small then the small players action will not be approximately naïve even in the limit.

Remark: An analogous result could be obtained in the case where the large player can either leave all small players unpunished or must punish all but a small fraction of small players. Thus if $N - \kappa(\Psi)$ denotes the *minimum* number of small players that must be punished then Proposition 1'' continues to hold.⁷

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 she is able to publicly commit himself to a strategy σ_0 before the small players act.⁸

If the large player is able to observe the actions of the small players, she can use this commitment power to induce the small players to choose any action profile that gives every small player a payoff that is strictly larger than his minmax payoff: If any small player deviates from the desired action the large player simply minmaxes that small player. Thus the large player can essentially get the largest payoff consistent with individual rationality of all small players.

However, we will show that 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 from an uncontingent commitment to 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.

Formally, we consider a sequence of games indexed by N , the number of small players. As we increase the number of small players, we hold fixed the players' action sets, the utility functions $u_i(a_i, b)$, and the distribution over individual outcomes $\rho_i(y_i, a_i)$. So that we may apply Proposition 1, we assume

⁷ The modification of the proof given in the appendix for this case is straightforward.

⁸ Remember that a strategy here is a contingent plan.

$$\sup_{i,a_i,b} |u_i(a_i,b)| < \infty, \quad \inf_{i,y_i} \sqrt{\rho_i(y_i|\alpha_i)} > 0.$$

Games are defined by a sequence of utility functions $u_0^N(a_1, a_2, \dots, a_N, b)$ for the large player. We assume that these utility functions are uniformly bounded.

The *Stackelberg payoff for the large player*, denoted u_0^{NS} , is the highest payoff the large player can get from publicly committing to a fixed, possibly mixed, action.

Formally, this payoff is given by

$$u_0^{NS} = \max_{\alpha, \beta} [u_0^N(\alpha, \beta) | u_i(\alpha_i, \beta) = \max_{\alpha_i} u_i(\alpha_i', \beta) \text{ for all } i].$$

In general, the utility of the large player might depend only on the play of one small player, in which case the fact that there are many small players is irrelevant. The interesting case is one in which no small player can have too much impact on the large player's utility. Formally, define the *influence* of a small player i on the large player's utility by

$$v_i^N \equiv \max_{a,b,a_i'} |u_0^N(a,b) - u_0^N(a_i', a_{-i}, b)|.$$

This represents the greatest impact of a change in player i 's action on the large player's utility for any combination of play by other players.

Definition 5: Small players are *asymptotically negligible* if the combined influence of all players remains bounded, i.e., $\sum_{i=1}^N v_i^N \leq K$ for some $K < \infty$, and if the individual influences go to zero uniformly, i.e., $\max_{i \leq N} v_i^N \rightarrow 0$.

Our main result in this section is that when small players are asymptotically negligible and the large player observes only the small players' outcomes, she can do no better in the limit than the Stackelberg payoff. This is analogous to a similar result in Levine and Pesendorfer (1995), who restrict the large player to play strategies that depend only on the average play of the small players.

Proposition 2: Let σ_0^N be a sequence of strategies for the large player, and suppose that α_i^N are a best response profile to σ_0^N . If small players are asymptotically negligible then

$$\limsup_{N \rightarrow \infty} \sum_{a,b,y} u_0^N(a,b) \sigma_0^N(b|y) \rho(y|a^N) \alpha^N(a) - u_0^{NS} = 0.$$

Before proving Proposition 2, we examine several examples that illustrate what it means for the small players to be asymptotically negligible. In all of these examples, the small players' actions are either 0 or 1, and the large player's utility does not depend on her own action. Note that this does not trivialize Proposition 2, since the large player action still indirectly has an impact on her own utility through the effect it has on small players' incentives.

In the first set of examples we study, the large player's utility is a weighted average of the small players' actions. That is,

$$u_0^N(a_1, a_2, \dots, a_N, b) = \sum_{i=1}^N w_i^N a_i,$$

where $\sum_{i=1}^N w_i^N = 1$. Notice that in this case the influences are the same as the weights, i.e., $v_i^N = w_i^N$. Consequently, the condition for asymptotic negligence reduces to the condition that the weights converge uniformly to zero. That is, $\max_{i \leq N} w_i^N \rightarrow 0$. The most obvious examples are $w_i^N = 1/N$, in which case asymptotic negligibility is obviously satisfied, and $w_1^N = 1, w_i^N = 0, i > 1$ in which case it is obviously not. There are also some less obvious examples, such as $w_i^N = 1/2\sqrt{iN}$ and $w_i^N = 1/i \log N$. It can easily be seen from integrating $\int_1^t (1/\sqrt{\tau}) d\tau$ and $\int_1^t (1/\tau) d\tau$ that $\sum_{i=1}^N w_i^N \approx 1$ in both cases. In both cases it is also obvious that $\max_{i \leq N} w_i^N = w_1^N \rightarrow 0$. This implies that asymptotic negligibility is satisfied. Notice, however that the ratio between the influence of the first and N th players goes to infinity, so that asymptotic negligibility does not imply anything about the relative significance of players.

Finally, it is useful to consider a case which can not be reduced to a simple average, and in which the location of the players plays a key role. Let us consider the case where the large player's utility is the (equally weighted) average of the number of small players i who chose 1 and whose predecessor $i-1$ did not choose 0. (In effect, choosing 0 not only gives nothing to the large player, but nullifies the choice of the next

player in line). In this case, if, for example, half the population is playing 1 and half 0, location is crucial. If players playing 1 are highly clustered, the large player gets around $\frac{1}{2}$; if the players playing 0 are every other player, or nearly so, then the large player gets almost nothing. Notice, however, that the asymptotic negligibility condition is satisfied in this example.

To prove Proposition 2, we need a technical Lemma.

Lemma C: If small players are asymptotically negligible then for all $\varepsilon > 0$ there exists $\bar{N}$ such that for all $N \geq \bar{N}$ and all profiles α^N there exists a pure profile a^N such that $|u_0^N(\alpha^N, b) - u_0^N(a^N, b)| < \varepsilon$.

Proof: Replace $(\alpha_1, \dots, \alpha_N)$ by $(a_1, \dots, a_N)$ by taking a random draw according to the probability distribution $(\alpha_1, \dots, \alpha_N)$. We show that the random variables

$$u_0^N(a_1, a_2, \dots, a_{N-1}, a_N, b) - u_0^N(\alpha_1, \alpha_2, \dots, \alpha_N, b)$$

converge in probability to zero, from which it follows that most choices of a (most meaning with respect to the probability distribution α) have the desired property. To show that the random variables converge to zero in probability, we compute

$$\begin{aligned} u_0^N(a_1, a_2, \dots, a_{N-1}, a_N, b) - u_0^N(\alpha_1, \alpha_2, \dots, \alpha_N, b) = \\ \sum_{i=1}^N [u_0^N(a_1, a_2, \dots, a_i, \alpha_{i+1}, \dots, \alpha_N, b) - u_0^N(a_1, a_2, \dots, a_{i-1}, \alpha_i, \dots, \alpha_N, b)] \end{aligned}$$

These are uncorrelated random variables, so

$$\begin{aligned} & \text{var}[u_0^N(a_1, a_2, \dots, a_{N-1}, a_N, b) - u_0^N(\alpha_1, \alpha_2, \dots, \alpha_N, b)] \\ &= \sum_{i=1}^N \text{var}[u_0^N(a_1, a_2, \dots, a_i, \alpha_{i+1}, \dots, \alpha_N, b) - u_0^N(a_1, a_2, \dots, a_{i-1}, \alpha_i, \dots, \alpha_N, b)] \\ &\leq \sum_{i=1}^N [v_i^N]^2 \leq \max_{i \leq N} v_i^N \sum_{i=1}^N v_i^N \end{aligned}$$

However, the assumption of asymptotic negligibility implies that $\max_{i \leq N} v_i^N \rightarrow 0$ and $\lim_{N \rightarrow \infty} \sum_{i=1}^N v_i^N < \infty$, so we have convergence in quadratic mean, implying convergence in probability. $\square$

Proof of Proposition 2: Let K^N be a sequence of integers that satisfies

(a) $K^N \rightarrow \infty$ and

(b) $\sum_{I^N} v_i^N \rightarrow 0$ for every subset I^N of agents that has the property $\#I^N \leq K^N$.

Clearly, by the assumption of asymptotic negligibility such a sequence of number exists.

Now define for each player i

$$\varepsilon_i = \max_{a'} \sum_{a,y,b} u_i(a',b) \sigma_0^N(b|y) \rho^N(y|a) \alpha^N(a) - \sum_{a,y,b} u_i(a_i,b) \sigma_0^N(b|y) \rho^N(y|a) \alpha^N(a)$$

Let J^N be the $N - K^N$ players with the smallest values of ε_i . Then $\max_{i \in J^N} \varepsilon_i$ is no greater than $(1/K^N) \sum_{i \notin J^N} \varepsilon_i$. We may apply Proposition 1 to the game consisting just

of the players in the complement of J^N to conclude

$$\max_{i \in J^N} \varepsilon_i \leq (1/K^N) \sum_{i \notin J^N} \varepsilon_i \leq \frac{1}{\sqrt{K^N}} \left[\frac{4m_b \sup_{i,a_i,b} |u_i(a_i,b)|}{\inf_{i,y_i} \sqrt{\rho_i(y_i|\alpha_i)}} \right] \xrightarrow{N \rightarrow \infty} 0$$

This in turn implies that there is a best response $\hat{\alpha}_i^N$ to the mixed action

$\beta_0^N(b) \equiv \sum_{a,y} \sigma_0^N(b|y) \rho^N(y|a) \alpha^N(a)$ that the large player uses on the equilibrium path with the following property: for all $i \in J^N$ there is $\bar{\varepsilon}^N \rightarrow 0$ such that $\|\alpha_i^N - \hat{\alpha}_i^N\| \leq \bar{\varepsilon}^N$.

We may now calculate for the large player using the definition of β_0^N

$$\begin{aligned} u_0^{NS} &\geq \sum_{a',b} u_0^N(a',b) \beta_0^N(b) \hat{\alpha}^N(a') = \sum_b \left[\sum_{a'} u_0^N(a',b) \hat{\alpha}^N(a') \right] \beta_0^N(b) \\ &= \sum_{a,b} \left[\sum_{a'} u_0^N(a',b) \hat{\alpha}^N(a') \right] h(a,b) \alpha^N(a) \end{aligned}$$

where we have defined $h(a,b) = \sigma_0^N(b|y) \rho^N(y|a)$.

Let $\mu^N = \sum_{i \notin J^N} v_i^N$ and recall that $\mu^N \rightarrow 0$ by construction. Using uniform boundedness of

the payoff function of the large player conclude that for some sequence $\tilde{\varepsilon}^N \rightarrow 0$

$$\begin{aligned} u_0^{NS} + \tilde{\varepsilon}^N + \mu^N &\geq \sum_{a,b} \left[\sum_{a'} u_0^N(a',b) \hat{\alpha}^N(a') \right] h(a,b) \alpha^N(a) + \tilde{\varepsilon}^N \\ &\geq \sum_{a,b} \left[\sum_{a'} u_0^N(a',b) \alpha^N(a') \right] h(a,b) \alpha^N(a) \end{aligned}$$

Next by Lemma B it follows that for some pure profile a^N and some sequence $\tilde{\tilde{\varepsilon}}^N \rightarrow 0$

$$|\sum_a u_0^N(a^N,b) \alpha^N(a) - \sum_a u_0^N(a,b) \alpha^N(a)| \leq \tilde{\tilde{\varepsilon}}^N.$$

Therefore we can apply Lemma C twice to conclude

$$\begin{aligned}
u_0^{NS} + \tilde{\epsilon}^N + 2\tilde{\tilde{\epsilon}}^N + \mu^N &\geq \sum_{a,b} \left[\sum_{a'} u_0^N(a',b) \alpha^N(a') \right] \alpha^N(a) h(a,b) + 2\tilde{\tilde{\epsilon}}^N \\
&\geq \sum_{a,b} \left[\sum_{a'} u_0^N(a^N,b) \alpha^N(a') \right] \alpha^N(a) h(a,b) + \tilde{\tilde{\epsilon}}^N \\
&= \sum_b \left[\sum_a u_0^N(a^N,b) \alpha^N(a) \right] h(a,b) + \tilde{\tilde{\epsilon}}^N \\
&\geq \sum_b \left[\sum_a u_0^N(a,b) \alpha^N(a) \right] h(a,b) \\
&= \sum_{a,b,y} u_0^N(a,b) \sigma_0^N(b|y) \rho^N(y|a) \alpha^N(a)
\end{aligned}$$

allowing $N \rightarrow \infty$ now yields the desired result.

6. Finitely Repeated Games

Reputation effects are often analyzed in models with a single patient player who faces 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 undiscounted, finitely repeated case; in the next section we examine an infinite horizon model in which the small players discount future payoffs. We first describe the stage game played in every period. In the stage game, 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 $t = 1, \dots, T < \infty$. As before, as the result of a small player’s action one of finitely many individual outcomes $y_i \in Y$ occurs. (Note that at each date t the distribution of that period’s outcome y_i^t given player i ’s action a_i^t is distributed independently of all past outcomes and of the contemporaneous outcomes of other players.) The probability of each individual outcome is again $\rho_i(y_i|a_i)$. All players, when taking an action in period t , have observed the profile $y = (y_1, \dots, y_N)$ 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:

$$\sum_{t=0}^T u_i^t(a_i^t, x^t).$$

The payoffs of the large player are not relevant 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 $(\tau_1, \dots, \tau_N, \sigma_0)$ 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 6: The small players' profile $(\tau_1, \dots, \tau_N)$ is an ε myopic best response profile to σ_0 if for each player i there is an $\varepsilon_i \geq 0$ such that for every history h^t that has positive probability $\sum_{a_i, x} u_i^t(a_i, x) \tau_i^t(h^t, h_i^t)[a_i] \sigma_0(h^t)[x] + \varepsilon_i \geq \sum_x u_i^t(a_i', x) \sigma_0(h^t)[x]$, for all $a_i' \in A_i$ 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 profile to σ_0 then $(\tau_1, \dots, \tau_N)$ is an $\varepsilon + \varepsilon'$ myopic best response profile to σ_0 where

$$\varepsilon' = \frac{1}{\sqrt{N}} \left[\frac{4T^2 (m_X)^T \max_{i, a_i, x, t} u_i^t(a_i, x)}{\min_a \sqrt{\rho_i(y_i | a_i)}} \right].$$

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

Proof: We denote by $\varepsilon'(t)$ the maximum average utility loss of small players from playing a myopic best response in periods $t, t+1, \dots, T$.

We compute $\varepsilon'(t)$ recursively from the end of the game and then set $\varepsilon' = \max_t \varepsilon'(t)$. 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 she 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 $[V_i^{T-t-1}(b^t), V_i^{T-t-1}(b^t) - \varepsilon_i - \varepsilon'_i(t+1)]$. Define

$$(*) \quad U_i(a_i, b^t) = \sum_x u_i^t(a_i, x) \sigma_0(h^t)[x] + V_i^{T-t-1}(b^t)$$

Since $(\tau_1, \dots, \tau_N)$ is an ε -best response profile in the repeated game, it follows that $(\tau_1^t(h^t, h_i^t), \dots, \tau_N^t(h^t, h_i^t))$ is an $\varepsilon + \varepsilon'(t+1)$ best response profile to U_i . Proposition 1 then implies that $(\tau_1^t(h^t, h_i^t), \dots, \tau_N^t(h^t, h_i^t))$ is an $\varepsilon + \varepsilon'(t)$ myopic best response where

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

Setting $\varepsilon' = \varepsilon(1) = T \frac{1}{\sqrt{N}} \left[\frac{4(m_X)^T \max_{i, a_i, x, t} (Tu_i^t(a_i, x))}{\min_a \sqrt{\rho_i(y_i | a_i)}} \right]$ then gives the result.

□

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 that $T = \infty$ and that

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

Proposition 4: Consider a sequence of games with m_X fixed and $\max_{i \leq N, a_i, x, t} u_i^t(a_i, x)$ bounded. If $(\tau_1, \dots, \tau_N)$ is a best response to σ_0 then $(\tau_1, \dots, \tau_N)$ is an ε^N -myopic best response to σ_0 where $\varepsilon^N \rightarrow 0$ as $\sqrt{N} \min_{a, i \leq N} \sqrt{\rho_i(y_i | a_i)} \rightarrow \infty$.

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

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

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

$$\varepsilon = \left(\delta^T + \frac{1}{\sqrt{N}} \left[\frac{4T^2(m_X)^T}{\min_a \sqrt{\rho_i(y_i|a_i)}} \right] \right) \max_{i, a_i, x, t} u_i^t(a_i, x).$$

As $\sqrt{N} \min_{a, i \leq N} \sqrt{\rho_i(y_i|a_i)} \rightarrow \infty$, we may let $T \rightarrow \infty$ sufficiently slowly that

$$\frac{1}{\sqrt{N}} \left[\frac{4T^2(m_X)^T}{\min_a \sqrt{\rho_i(y_i|a_i)}} \right] \rightarrow 0,$$

giving the desired result. □

8. Conclusions

There are several extensions and implications of our results that should be remarked on. First, 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.

Second, our basic assumption is that the information of the large player about the small players consists of a separate signal for each small player, but that the number of available actions is limited. It is also of interest to consider the opposite case in which the large player has many actions, but he has less information about the small players; an extreme case of this is that the large player's information is simply an aggregate of what the small players do. Levine and Pesendorfer (1995) considered this model with the additional assumption that the large player cannot target individual small players for punishments, and showed that the small players behave myopically. Since the availability of actions that punish particular players is irrelevant if only an aggregate is observed, the myopic behavior of the small players extends to this case. Thus there is a sense in which this paper and Levine and Pesendorfer (1995) have analyzed opposing endpoints along a spectrum of sufficient conditions for the negligibility of small players. The more refined question of tracing the frontier of such sufficient conditions is interesting but beyond the scope of this paper.⁹

Appendix A

Lemma A': Let y_i be N independent random variables each taking on the values $\{y_i^0, y_i^1\}$, and let p be another random variable with finite support, which may be correlated with the y_i . Suppose there is an $\varepsilon > 0$ such that for each vector $y = (y_1, \dots, y_N)$ there exists $0 \leq q(y) \leq 1$ with $\frac{1}{N} \sum_{i=1}^N |pr(p \in \Pi_i | y) - q(y)| \leq \varepsilon$. Then

⁹ We are grateful to Hyun Song Shin for raising this point.

$$\begin{aligned} & \frac{\sum_{i=1}^N pr(p \in \Pi_i | y_i = y_i^1) - pr(p \in \Pi_i | y_i = y_i^0)}{N} \\ & \leq \frac{2}{\sqrt{N} \min_{i,j} \sqrt{pr(y_i = y_i^j)}} + \frac{\varepsilon}{\min_{i,j} pr(y_i = y_i^j)}. \end{aligned}$$

Proof:

$$\begin{aligned} D &= \frac{\sum_{i=1}^N pr(p \in \Pi_i | y_i = y_i^1) - pr(p \in \Pi_i | y_i = y_i^0)}{N} \\ &= \frac{\sum_{i=1}^N \frac{pr(p \in \Pi_i, y_i = y_i^1)}{pr(y_i = y_i^1)} - \frac{pr(p \in \Pi_i, y_i = y_i^0)}{pr(y_i = y_i^0)}}{N} \\ &= \frac{\sum_{i=1}^N \frac{\sum_{\{y|y_i=y_i^1\}} pr(p \in \Pi_i, y)}{pr(y_i = y_i^1)} - \frac{\sum_{\{y|y_i=y_i^0\}} pr(p \in \Pi_i, y)}{pr(y_i = y_i^0)}}{N} \\ &= \frac{1}{N} \sum_y pr(y) \sum_{n=1}^N \begin{cases} \frac{pr(p \in \Pi_i | y)}{pr(y_i = y_i^1)} & y_i = y_i^1 \\ \frac{pr(p \in \Pi_i | y)}{pr(y_i = y_i^0)} & y_i = y_i^0 \\ 0 & y_i \neq y_i^1, y_i^0 \end{cases} \end{aligned}$$

Define

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

and note that $Ez_i^j = 1$.

$$\begin{aligned} D &= \frac{1}{N} \sum_y pr(y) \sum_{i=1}^N pr(p \in \Pi_i | y) (z_i^1(y) - z_i^0(y)) \\ &= \frac{1}{N} \sum_y pr(y) \left[\sum_{i=1}^N q(y) (z_i^1(y) - z_i^0(y)) + [pr(p \in \Pi_i | y) - q(y)] (z_i^1(y) - z_i^0(y)) \right] \\ &\leq \frac{1}{N} \sum_y pr(y) \left| \sum_{i=1}^N z_i^1(y) - \sum_{i=1}^N z_i^0(y) \right| \\ &\quad + \frac{1}{N} \sum_y pr(y) \sum_{i=1}^N [pr(p \in \Pi_i | y) - q(y)] (z_i^1(y) - z_i^0(y)) \end{aligned}$$

Note that

$$\begin{aligned} |\sum_{i=1}^N z_i^1(y) - \sum_{i=1}^N z_i^0(y)| &= |(\sum_{i=1}^N z_i^1(y) - 1) - (\sum_{i=1}^N z_i^0(y) - 1)| \\ &\leq |\sum_{i=1}^N z_i^1(y) - 1| + |\sum_{i=1}^N z_i^0(y) - 1| \end{aligned}$$

Therefore,

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

Using Jensen's inequality we get

$$\begin{aligned}
&E \left(\left| \frac{1}{N} \left(\sum_{i=1}^N z_i^1 - Ez_i^1 \right) \right| + \left| \frac{1}{N} \left(\sum_{i=1}^N z_i^0 - Ez_i^0 \right) \right| \right) \\
&\leq \left\{ E \left| \frac{1}{N} \left(\sum_{i=1}^N z_i^1 - Ez_i^1 \right) \right|^2 \right\}^{1/2} + \left\{ E \left| \frac{1}{N} \left(\sum_{i=1}^N z_i^0 - Ez_i^0 \right) \right|^2 \right\}^{1/2} \\
&= \left\{ \frac{1}{N^2} \left(\sum_{i=1}^N \text{var}(z_i^1) \right) \right\}^{1/2} + \left\{ \frac{1}{N^2} \left(\sum_{i=1}^N \text{var}(z_i^0) \right) \right\}^{1/2} \\
&= \left\{ \sum_{i=1}^N \frac{1 - pr(y_i = y_i^1)}{N^2 pr(y_i = y_i^1)} \right\}^{1/2} + \left\{ \sum_{i=1}^N \frac{1 - pr(y_i = y_i^0)}{N^2 pr(y_i = y_i^0)} \right\}^{1/2} \\
&\leq \frac{2}{\sqrt{N} \min_{i,j} \sqrt{pr(y_i = y_i^j)}}
\end{aligned}$$

which proves the claim. □

Lemma B: Let $(\mathbf{B}_1, \dots, \mathbf{B}_N)$ be N partitions of the large player's action set. If α is an ε -best response profile, it is an $\varepsilon + \varepsilon'$ -naive best response profile with

$$\begin{aligned}
\varepsilon' &\leq 2 \max_{i,a_i,b} |u_i(a_i,b)| \frac{1}{N} \sum_{i,B \in \mathbf{B}_i} \max_{y_i,y_i'} |pr(b \in B | y_i, \alpha) - pr(b \in B | y_i', \alpha)| \\
&\quad + 4 \frac{1}{N} \sum_{i=1}^N \max_{a_i, B \in \mathbf{B}_i} [\max_{b \in B} u_i(a_i, b) - \min_{b \in B} u_i(a_i, b)]
\end{aligned}$$

Proof: Fix an ε -best response profile α . Let a_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 α under naive expectations is

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

Since α is an ε -best response,

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

from which we see that

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

In words, the loss under naive expectations is bounded by the sum of the loss under

correct expectations and the approximation error. Note also that

$$\begin{aligned} & \sum_{y,b} u_i(a'_i, b) \sigma_0(b|y_{-i}, y'_i) \rho_{-i}(y_{-i}|\alpha_{-i}) [\rho_i(y_i|\alpha_i) - \rho_i(y_i|a'_i)] \\ &= \sum_{y_{-i}, b} u_i(a'_i, b) \sigma_0(b|y_{-i}, y'_i) \rho_{-i}(y_{-i}|\alpha_{-i}) \sum_{y_i} [\rho_i(y_i|\alpha_i) - \rho_i(y_i|a'_i)] \\ &= 0. \end{aligned}$$

We conclude that

$$\begin{aligned} \eta_i &\leq \sum_{y,b} u_i(a'_i, b) [\sigma_0(b|y) - \sigma_0(b|y_{-i}, y'_i)] \rho_{-i}(y_{-i}|\alpha_{-i}) [\rho_i(y_i|\alpha_i) - \rho_i(y_i|a'_i)] + \varepsilon_i \\ &= \sum_{y_i, b} u_i(a'_i, b) \times \\ &\quad \sum_{y_{-i}} [\sigma_0(b|y_{-i}, y_i) - \sigma_0(b|y_{-i}, y'_i)] \rho_{-i}(y_{-i}|\alpha_{-i}) [\rho_i(y_i|\alpha_i) - \rho_i(y_i|a'_i)] + \varepsilon_i \\ &= \sum_{y_i, b} u_i(a'_i, b) [pr(b|y_i, \alpha) - pr(b|y'_i, \alpha)] [\rho_i(y_i|\alpha_i) - \rho_i(y_i|a'_i)] + \varepsilon_i \\ &= \sum_{y_i} \sum_{B \in \mathbf{B}_i} \sum_{b \in B} u_i(a'_i, b) [pr(b|y_i, \alpha) - pr(b|y'_i, \alpha)] [\rho_i(y_i|\alpha_i) - \rho_i(y_i|a'_i)] + \varepsilon_i \\ &= \sum_{y_i} \sum_{B \in \mathbf{B}_i} \max_{b \in B} u_i(a'_i, b) \sum_{b \in B} [pr(b|y_i, \alpha) - pr(b|y'_i, \alpha)] [\rho_i(y_i|\alpha_i) - \rho_i(y_i|a'_i)] \\ &\quad + \sum_{y_i} \sum_{B \in \mathbf{B}_i} \sum_{b \in B} [u_i(a'_i, b) - \max_{b \in B} u_i(a'_i, b)] \times \\ &\quad [pr(b|y_i, \alpha) - pr(b|y'_i, \alpha)] [\rho_i(y_i|\alpha_i) - \rho_i(y_i|a'_i)] + \varepsilon_i \\ &\leq \max_{a_i, b} |u_i(a_i, b)| \sum_{y_i, B \in \mathbf{B}_i} |pr(b \in B|y_i, \alpha) - pr(b \in B|y'_i, \alpha)| |\rho_i(y_i|\alpha_i) - \rho_i(y_i|a'_i)| \\ &\quad + \sum_{y_i} \sum_{B \in \mathbf{B}_i} |\max_{b \in B} u_i(a'_i, b) - \min_{b \in B} u_i(a'_i, b)| \times \\ &\quad \sum_{b \in B} [pr(b|y_i, \alpha) - pr(b|y'_i, \alpha)] [\rho_i(y_i|\alpha_i) - \rho_i(y_i|a'_i)] + \varepsilon_i \\ &\leq \max_{a_i, b} |u_i(a_i, b)| \sum_{B \in \mathbf{B}_i} \max_{y_i} |pr(b \in B|y_i, \alpha) - pr(b \in B|y'_i, \alpha)| \times \\ &\quad \sum_{y_i} |\rho_i(y_i|\alpha_i) - \rho_i(y_i|a'_i)| \\ &\quad + \max_{y_i, B \in \mathbf{B}_i} [|\max_{b \in B} u_i(a'_i, b) - \min_{b \in B} u_i(a'_i, b)|] \times \\ &\quad \sum_{b \in B} |pr(b|y_i, \alpha) - pr(b|y'_i, \alpha)| \sum_{y_i} |\rho_i(y_i|\alpha_i) - \rho_i(y_i|a'_i)| + \varepsilon_i \\ &\leq 2 \max_{a_i, b} |u_i(a_i, b)| \sum_{B \in \mathbf{B}_i} \max_{y_i} |pr(b \in B|y_i, \alpha) - pr(b \in B|y'_i, \alpha)| \\ &\quad + 4 \max_{B \in \mathbf{B}_i} [\max_{b \in B} u_i(a'_i, b) - \min_{b \in B} u_i(a'_i, b)] + \varepsilon_i \end{aligned}$$

Finally, we average over i to find the desired conclusion.

Proposition 1'': If α is an ε -best response profile, it is an $\varepsilon + \varepsilon'$ -naive best response profile with $\varepsilon' \leq 4 \max_{a_i, b} |u_i(a_i, b)| \left[\frac{2}{\sqrt{N} \min_{i, y_i} \sqrt{\rho_i(y_i | \alpha_i)}} + \frac{\kappa(\Psi)}{N \min_{i, y_i} \rho_i(y_i | \alpha_i)} \right]$.

Proof: We construct for each player i a partition consisting of 2 sets: $\mathbf{B}_i = \{\tilde{B}_i, \hat{B}_i\}$ where $b_\psi \in \hat{B}_i$ if player i gets punished, i.e., if $i \in \psi$, and $b_\psi \in \tilde{B}_i$ otherwise.

First we demonstrate that if α is an ε -best response profile, it is an $\varepsilon + \varepsilon'$ -naive best response profile with

$$\varepsilon' \leq 4 \max_{a_i, b} |u_i(a_i, b)| \left[\frac{1}{N} \sum_i \max_{y_i, y'_i} |pr(b \in \hat{B}_i | y_i, \alpha) - pr(b \in \hat{B}_i | y'_i, \alpha)| \right]$$

To demonstrate this, note that the utility variation within the elements of $\tilde{B}_i$ and $\hat{B}_i$ is zero for every small player. Therefore, we may apply Lemma B to conclude

$$\begin{aligned} \varepsilon' &\leq \frac{2 \max_{i, a_i, b} |u_i(a_i, b)|}{N} \times \\ &\sum_i \max_{y_i, y'_i} |pr(b \in \tilde{B}_i | y_i, \alpha) - pr(b \in \tilde{B}_i | y'_i, \alpha)| + \\ &\sum_i \max_{y_i, y'_i} |pr(b \in \hat{B}_i | y_i, \alpha) - pr(b \in \hat{B}_i | y'_i, \alpha)| \end{aligned}$$

Also note that $1 - pr(b \in \tilde{B}_i | y) = pr(b \in \hat{B}_i | y)$ and therefore

$$\begin{aligned} \varepsilon' &\leq 2 \max_{a_i, b} |u_i(a_i, b)| \left(\frac{1}{N} \sum_i \max_{y_i, y'_i} |1 - pr(b \in \hat{B}_i | y_i, a) - 1 + pr(b \in \hat{B}_i | y'_i, a)| \right) \\ &\quad + \frac{1}{N} \sum_i \max_{y_i, y'_i} |pr(b \in \hat{B}_i | y_i, a) - pr(b \in \hat{B}_i | y'_i, a)| \Big) \\ &\leq 4 \max_{a_i, b} |u_i(a_i, b)| \left[\frac{1}{N} \sum_i \max_{y_i, y'_i} |pr(b \in \hat{B}_i | y_i, \alpha) - pr(b \in \hat{B}_i | y'_i, \alpha)| \right] \end{aligned}$$

This establishes the claim above.

To see the proposition follows from the claim, observe that since the expected number of players punished must be bounded by the maximum number of players who can be punished it follows that for all y

$$\frac{1}{N} \sum_i pr(b \in \hat{B}_i | y) \leq \frac{\kappa(\Psi)}{N}$$

Thus setting $q(y) = 0$ we can apply Lemma A' to our previous claim to conclude that

$$\mathcal{E}' \leq 4 \max_{a_i, b} |u_i(a_i, b)| \left[\frac{2}{\sqrt{N} \min_{i, y_i} \sqrt{\rho_i(y_i | \alpha_i)}} + \frac{\kappa(\Psi)}{N \min_{i, y_i} \rho_i(y_i | \alpha_i)} \right]$$

□

Bibliography

- Celentani, M. (1991): "Reputation in General Deterministic Stage Games," forthcoming, *Economic Theory*.
- Celentani, M. and W. Pesendorfer (1996): "Reputation in Dynamic Games," forthcoming *Journal of Economic Theory*.
- 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 (1995): "When are Agents Negligible," *American Economic Review* 85, 1160-1170 .
- 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