

Al-Najjar, Nabil I.; Smorodinsky, Rann

Working Paper

Large Non-Anonymous Repeated Games

Discussion Paper, No. 1250

Provided in Cooperation with:

Kellogg School of Management - Center for Mathematical Studies in Economics and Management Science, Northwestern University

Suggested Citation: Al-Najjar, Nabil I.; Smorodinsky, Rann (1998) : Large Non-Anonymous Repeated Games, Discussion Paper, No. 1250, Northwestern University, Kellogg School of Management, Center for Mathematical Studies in Economics and Management Science, Evanston, IL

This Version is available at:

<https://hdl.handle.net/10419/221606>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Discussion Paper No. 1250

**LARGE NON-ANONYMOUS
REPEATED GAMES**

By

Nabil I. AL-NAJJAR*

And

Rann Smorodinsky⁺

December 27, 1998

<http://www.kellogg.nwu.edu/research/math>

Abstract

Saborian [8], following Green [4], studies a class of repeated games where a player's payoff depends on his stage action and an anonymous aggregate outcome, and shows that long-run players behave myopically in any equilibrium of such games. In this paper we extend Sabourian's results to games where the aggregate outcome is not necessarily an anonymous function of players' actions, and where players strategies may depend non-anonymously on signals of other players' behavior. Our argument also provides a conceptually simpler proof of Green and Sabourian's analysis, showing how their basic result is driven by bounds on how many pivotal players there can be in a game.

* MEDS/KGSM, 2001 Sheridan Road, Evanston, IL 60208, U.S.A. <al-najjar@nwu.edu>.

⁺ Faculty of Industrial Engineering and Management, Technion—Israel Institute of Technology, Haifa 32000, Israel <rann@ie.technio.ac.il>.

1 Introduction

Long-term strategic interactions among a large number of players are characterized by two competing forces. Since players in a repeated setting consider the long-term consequences of their actions, their behavior may substantially differ from the one predicted by short-term, myopic considerations. On the other hand, a player facing a large number of opponents may judge his impact on future play to be too small to offset the loss from not taking a short-term optimal action.

Green [4] studied the effects of these two competing forces in an attempt to provide adequate foundation for competitive behavior in markets where firms interact strategically in a repeated setting. Sabourian [8] extended Green's model in several important directions. In particular, Sabourian drops Green's restriction to trigger strategies and relies solely on Green's continuity assumption. Both papers consider an anonymous repeated game with random outcomes, and where a player's stage game payoff may depend on his own action as well as an aggregate outcome that may be influenced by other players' actions. Sabourian proves that if the mapping from the action profile to outcomes is anonymous and continuous in an appropriate sense, then in any (subgame perfect) equilibrium of the game, players play their myopic ϵ -best reply at each node of the game.

In this paper we extend Green's and Sabourian's results to the case where outcomes are not necessarily anonymous functions of players' actions, and where players may condition on individualized signals of their opponents not reflected in the collective outcome. In addition, we provide a conceptually simpler proof which clarifies the intuition underlying Green-Sabourian's analysis and a potentially useful tool for carrying it out in new contexts.

Specifically, we consider a setting in which players observe noisy signals of their opponents' actions, as well as a collective outcome that is an arbitrary function of the signal realizations. Players' payoffs depend on their own actions and the collective outcome, but their strategies may also depend on their opponents' signals. We establish

an analogue of Green-Sabourian’s result that most players play non-strategically. We do this using a notion of *influence* that measures the impact of a unilateral change of one agent’s signal on the expected value of the collective outcome. For a given threshold $\alpha > 0$, a player is α -pivotal if his influence exceeds α . We use a non-pivotalness result established in earlier work (Al-Najjar and Smorodinsky [1]) to conclude that the number of players who are α -pivotal relative to the collective outcome in a given period is bounded by some integer K_α^* . This bound holds uniformly over all strategy profiles, and independently of the total number of players in the game.

Our version of Green-Sabourian’s result then follows from this non-pivotalness theorem: for any equilibrium and at any node of the repeated game, no more than K_α^* players are α -pivotal; i.e., believe their actions have important enough impact on the common outcome to offset the cost of deviating from taking an approximate myopic best response. Thus, in any equilibrium and at any node of the repeated game, the remaining $N - K_\alpha^*$ players play approximate one-shot best responses. This, in essence, is Green-Sabourian’s point, except that their conclusion that ‘all players’ play myopic ϵ -best reply is slightly weakened to ‘all but a vanishing fraction of players’.

Aside from its simplicity, this argument also shows that Green-Sabourian’s result holds when the payoff-relevant outcome depend non-anonymously on the players’ signals, or when players condition their behavior on signals of their opponents’ actions. The key feature needed to preserve their result is that a player’s signals affect other players’ payoffs only indirectly through the common outcome. The relevance of this extension may be seen in the competitive industry example that motivated Green’s original work. It is quite natural to allow for the possibility that firms have considerable, disaggregated information about the behavior of other firms in their industry. Our model allows firms to condition in a non-anonymous manner on such signals, provided they affect payoff’s only indirectly through common outcomes such as market price, state of demand and technology, and so on. While Green-Sabourian analysis rules out the possibility of firms conditioning their actions on payoff-irrelevant signals of their competitors, we show that it is an implication of equilibrium that this added flexibility has essentially no effect on the behavior of most players. Thus, Green-Sabourian’s

requirements that firms condition only on a common outcome, and that this outcome depends only on the players' frequencies of actions, are unnecessary.

The paper proceeds as follows. Section 2 introduces the repeated game model, which is essentially that in Sabourian. Section 3 contains the main result for the special case of a finite outcome set, as well as an extension to a model with a large outcome space. Specifically, we define a family of repeated games for which the number of players who play strategically is bounded by an integer that holds uniformly over all game in that family. Finally, Section 4 contains concluding remarks which relate our results to the literature and suggests directions for future work.

2 Model

The model we study is an adaptation of a model introduced by Green [4] and later studied by Sabourian [8]. In the original model the action spaces and the outcome spaces are continuous. We on the other hand assume that all that the action space is finite. As for the outcome space, we provide results for the finite case as well as for the continuous case. We begin by introducing the stage game and move later to the repeated game.

2.1 Stage Game

A stage game H is defined as the tuple $H = (\mathcal{K}, A, S, C, X, F, \pi)$.

- $\mathcal{K} = \{1, 2, \dots, K\}$ is the set of players;
- $A = \prod_{k=1}^K A^k$, where A^k is a *finite* set of actions available to player k ;
- $S = \prod_{k=1}^K S^k$, where S^k is a finite set of signals of player k 's actions;
- $C = \{C^k\}_{k=1}^K$, where $C^k : A^k \rightarrow \Delta(S^k)$ is a mapping from actions to signals;¹
- X is a *finite* set of outcomes;²

¹For any finite set Y we denote by $\Delta(Y)$ the set of all probability distributions on Y .

²Section 3.2 extends the model to large outcome spaces.

- $F : S \rightarrow \Delta(X)$ maps action profiles into outcome distributions.
- $\pi = \{\pi^k\}_{k=1}^K$, where $\pi^k : A^k \times X \rightarrow \mathbb{R}$ denotes player k 's payoff function.

Note that a player's payoff depends on the actions taken by his opponents only via the outcome x .

A player's (mixed strategy) is a probability distribution $\sigma^k \in \Delta(A^k)$. We shall often abuse notation and write a^k to denote the Dirac measure on a^k . Let $\sigma = (\sigma^1, \sigma^2, \dots, \sigma^k) \in \Delta(A)$ denote the strategy profile. As usual we denote by σ^{-k} the strategy profile of all players but player k . A strategy profile σ induces a measure over $A \times S \times X$, denoted λ_σ . Player k 's expected payoff from the strategy profile σ is $E_{\lambda_\sigma}[\pi(a^k, x)]$.

Definition 1 A strategy σ^k is a best response (BR) to σ^{-k} if

$$E_{\lambda_{(\sigma^k, \sigma^{-k})}}[\pi(a^k, x)] - E_{\lambda_\sigma}[\pi(a^k, x)] \leq 0 \quad \forall \bar{\sigma}^k \in \Delta(A^k). \quad (1)$$

It is an α -BR if the right-hand side of (1) is replaced with an $\alpha > 0$.

Definition 2 Let J be a nonnegative integer and $\alpha > 0$. A strategy profile $\sigma = (\sigma^1, \dots, \sigma^k)$ is called a (J, α) -equilibrium for H if σ^k is not an α -BR for at most J players.

In other words, in a (J, α) -equilibrium, all players but at most J , play their α -BR, to their opponents' strategy profile. A $(0, 0)$ -equilibrium is commonly known as a (stage game) *Nash equilibrium*.

2.2 Repeated Game

Given a stage game H , let $H^\infty = H^\infty(\delta^1, \dots, \delta^K)$ denote the infinitely-repeated game with discount factor δ^k for player k .

A pure strategy for player k is an assignment of an action a^k at time t , which may depend on players' signals in the previous $t - 1$ stages.³ A behavior strategy, Σ^k , for player k is an assignment of a stage game mixed strategy instead.⁴ Formally,

$$\Sigma^k : \bigcup_{t=0}^{\infty} S^t \rightarrow \Delta(A^k)^5 .$$

Note that players are assumed not to observe opponents' strategies or even stage game actions.

The strategy profile $\Sigma = \{\Sigma^k\}_{k=1}^K$ induces a probability distribution over $(A \times S \times X)^\infty$, denoted λ_Σ . We use μ_Σ to denote the marginal over S^∞ . Let $\lambda_{\Sigma,t}$ and $\mu_{\Sigma,t}$ denote the corresponding marginals over the t^{th} coordinate. The expected stage-game payoff at time t to player k is therefore $E_{\lambda_{\Sigma,t}}[\pi(a^k, x)]$. The total expected payoff to player k from a strategy profile $\Sigma = (\Sigma^1, \dots, \Sigma^K)$ is $U^k(\Sigma) = \sum_{t=0}^{\infty} (\delta^k)^t E_{\lambda_{\Sigma,t}}[\pi(a^k, x)]$.

A Nash equilibrium of the repeated game, H^∞ , is a strategy profile $\Sigma = (\Sigma^1, \dots, \Sigma^K)$ satisfying $U^k(\Sigma) \geq U^k(\hat{\Sigma}^k, \Sigma^{-k})$ for any $k \in K$ and any behavior strategy $\hat{\Sigma}^k$ of player k .

2.3 A Uniform Family of Games

In this paper we prove results that hold uniformly for families of games parameterized by $Q < \infty$, $\delta < 1$, $M < \infty$, $\epsilon > 0$ as follows:

$$\mathcal{H}(Q, \delta, M, \epsilon) =$$

³In our model F is deterministic. Therefore we did not formally allow for strategies to depend on past outcomes, as this is redundant given the dependence on past signals. However one can easily extend our model to the case where F may be random and allow for strategies to depend on past outcomes as well.

⁴In this paper we restrict attention to behavior strategies. This restriction, from the seemingly more general mixed strategies, defined as probability distributions over the set of all pure strategies, is without loss of generality (see Kuhn [5] and Aumann [2]).

⁵With some expositional complications the model and results of this paper can be extended to the case where a player's strategy may depend on his own past actions.

$$= \{H^\infty : |\pi^k| < Q, \delta^k < \delta \forall k, |X| < M, C^k(a^k) \in \Delta^\epsilon(S^k) \forall k \text{ and all } a^k \in A^k\},$$

where $\Delta^\epsilon(S^k) \subset \Delta(S^k)$ denotes the set of all probability distributions on S^k which assign probability at least ϵ to any $s^k \in S^k$.

In words, a repeated game H in such family is one where stage payoffs cannot exceed $Q < \infty$, players' discount factors are bounded away by $\delta < 1$, the cardinality of the outcome set is bounded by $M < \infty$ and probabilities of signals are bounded away from zero in all games, where Q, δ, M, ϵ are uniform over the entire family.

Note that for any $H^\infty \in \mathcal{H}(Q, \delta, M, \epsilon)$, any strategy profile Σ assigns a positive probability to every finite history $\bar{s} \in \bigcup_{t=1}^\infty S^t$.⁶

3 Results

We begin this section with our main result which refers to the model introduced in the previous section. We then show how to extend this result to the case of a large outcome space, and finally we make a brief comparison with Sabourian's results.

For a behavior strategy Σ^k and a finite history $\bar{s} \in \bigcup_{t=1}^\infty S^t$ of length \bar{t} , we denote by $\Sigma_{\bar{s}}^k$ the continuation strategy after the history \bar{s} .⁷ A (J, α) -myopic equilibrium of H^∞ is a strategy profile Σ such that for any finite history $\bar{s} \in S^t$, which has positive probability according to λ_Σ , the stage game strategy profile $\Sigma_{\bar{s}}(\emptyset)$ is a (J, α) -equilibrium of H .

3.1 Main Result

The following theorem states that in a Nash equilibrium of any Game $H^\infty \in \mathcal{H}(Q, \delta, M, \epsilon)$ most players play myopically.

⁶Consequently, any Nash equilibrium is also a subgame perfect equilibrium. However, had we considered the slightly more general model where players' strategies may depend on own past signals, we would have needed to differentiate between the two solution concepts.

⁷Namely, $\Sigma_{\bar{s}}^k(\tilde{s}) = \Sigma^k(\bar{s}\tilde{s})$, for any finite history, $\tilde{s} \in \bigcup_{t=0}^\infty S^t$, of length \tilde{t} and where $(\bar{s}\tilde{s}) \in \bigcup_{t=0}^\infty S^t$ denotes a new history of length $\bar{t} + \tilde{t}$ generated by concatenation.

Theorem 1 For any Q, δ, M, ϵ and any $\alpha > 0$, there exists an integer $J = J(Q, \delta, M, \epsilon, \alpha)$ such that if $H^\infty \in \mathcal{H}(Q, \delta, M, \epsilon)$ then any Nash equilibrium of H^∞ is a (J, α) -myopic equilibrium of H^∞ .

The sense in which ‘most players play myopically’ is formalized in the theorem by asserting the existence of an upper bound, independent of K , for the number of players who do **not** play their α -BR of the stage game.

Theorem 1 follows from the next lemma regarding influence. Let Σ^k be a behavior strategy for k . Denote by $\Sigma_{(a)}^k$ the behavior strategy which takes the pure action $a \in A^k$ at stage 0, and otherwise follows Σ^k . We denote by $(\Sigma^{-k}, \Sigma_{(a)}^k)$ the profile of strategies where all players $j \neq k$ play the strategy Σ^j and player k plays $\Sigma_{(a)}^k$.

Lemma 1 For any Q, δ, M, ϵ and any $\hat{\alpha} > 0$, there exists an integer $\hat{J} = \hat{J}(Q, \delta, M, \epsilon, \hat{\alpha})$ such that if $H^\infty \in \mathcal{H}(Q, \delta, M, \epsilon)$ and Σ is an arbitrary profile of behavior strategies, then for any stage t there exists a subset \hat{K}_t of at most \hat{J} players such that for any $k \notin \hat{K}_t$ and pair of actions $a, b \in A^k$:

$$|E_{\mu_{(\Sigma^{-k}, \Sigma_{(a)}^k), t}} \pi^k(d, x) - E_{\mu_{(\Sigma^{-k}, \Sigma_{(b)}^k), t}} \pi^k(d, x)| < \hat{\alpha} \quad \forall d \in A^k.$$

In other words, Lemma 1 asserts that for all but \hat{J} players a change of actions at the beginning of the game can only have a limited effect on their expected payoff at stage t , as long as the action at stage t is left unchanged.

The proof of Lemma 1, which builds on a result from an earlier work (Al-Najjar and Smorodinsky [1]), is postponed to the appendix.

Proof of Theorem 1: If Σ is a Nash equilibrium for $H^\infty \in \mathcal{H}(Q, \delta, M, \epsilon)$ then $\Sigma_{\bar{s}}$ is also a Nash equilibrium for H^∞ , for any $\bar{s} \in S^t$ (remember that any history is assigned positive probability according to μ_Σ). Therefore, in order to prove Theorem 1 it is sufficient to show that for any Nash equilibrium the induced stage game strategies for the first stage of the game is a (J, α) -equilibrium. In other words it is sufficient to prove that $\Sigma(\emptyset) = \{\Sigma^k(\emptyset)\}_{k=1}^K$ is a (J, α) -equilibrium of the stage game H .

This follows from two observations. First, as the discount factors and stage payoffs are uniformly bounded by δ and Q respectively one can find a stage $T = T(Q, \delta, \alpha)$, independent of K , such that the discounted sum of stage payoffs after T lies in the interval $[-\frac{\alpha}{4}, \frac{\alpha}{4}]$ for any sequence of action tuples taken by the players. Thus, a change in action taken by player k at the initial stage of the game cannot change his accumulated discounted payoff from time T on by more than $\alpha/2$.

Second, by Lemma 1 there are only \hat{J} players for which a change of action at stage 0 can increase their payoff at any stage t by more than $\hat{\alpha} = \frac{\alpha}{2T}$. Therefore we conclude that at most $J = T\hat{J}$ players can increase their accumulated payoff between stages 1 and T , by more than $\alpha/2$, by a change of action at stage 0.

We conclude that for all, but J players, a change of action at stage 0 will not change the total discounted sum from stage 1 on by more than α . As Σ^k is a best response to Σ^{-k} we learn that the stage game strategy $\Sigma^k(\emptyset)$ must be an α -BR to $\Sigma^{-k}(\emptyset)$ for all but J players. This implies that $\{\Sigma^k(\emptyset)\}_{k=1}^K$ is a (J, α) -equilibrium for H . \square

3.2 Large Outcome Spaces

In this section we extend the results of the model to the case of a large outcome space, X . Specifically, we shall assume $X = [0, 1)$.⁸ Theorem 1 does not hold in this case without additional assumptions, as shown in the following example:

Example 1 *Let $A^k = S^k = \{0, 1\}$, let $C^k(a^k)$ assign probability $1 - \epsilon$ to $s^k = a^k$ and ϵ to $s^k = 1 - a^k$. Let $F : S \rightarrow X$ be any one-to-one function (e.g., $F(s^1, \dots, s^K) = \sum_k s^k \frac{1}{2}^k$). As F is one-to-one we can allow for the following payoff function, $\pi^k(x) = s^{k+1} + \frac{1}{5}(1 - s^k)$. In words, an agent's payoff depends on his predecessor's signal as well as his own signal. From a myopic perspective players should therefore play $a^k = 0$ at all stages. However, consider the following strategies for all players: at stage 1 let $a^k = 1$*

⁸The use of the half open interval is without loss of generality, and is done for notational convenience, as becomes clear in the proof of Theorem 2.

and at stage t let $a^k = 1$ if and only if s^{k-1} was equal 1, at all periods $1, \dots, t-1$. Note that, independently of the number of players, these strategies constitute a Nash equilibrium (if players are sufficiently patient), yet all players do not take myopic best responses at the first stage.

To recover an analogue of Theorem 1 we add an assumption regarding the continuity of the payoff functions.

Definition 3 Fix a function $\rho : [0, 1] \rightarrow \mathbb{R}_+$. We say that the payoff function π^k is ρ -continuous if for any $a^k \in A^k$, $\epsilon > 0$ and $x, y \in X$ satisfying $|x - y| < \rho(\epsilon)$ then $|\pi(a^k, x) - \pi(a^k, y)| < \epsilon$.

In order to extend the results of Theorem 1 we replace the restriction on the cardinality of X with a *uniform continuity* assumption. Thus, we consider the following family of games. Fix of parameters $\delta < 1$, $\rho : [0, 1] \rightarrow \mathbb{R}_+$, $\epsilon > 0$, and define:

$$\begin{aligned} \mathcal{H}(\delta, \rho, \epsilon) = \\ = \{H^\infty : \pi^k \text{ is } \rho\text{-continuous, } \delta^k < \delta \forall k, C^k(a^k) \in \Delta^\epsilon(S^k) \forall k \text{ and all } a^k \in A^k\} . \end{aligned}$$

Note that our earlier uniform bound on payoffs is now redundant given uniform continuity. Note also that the force of uniform continuity is that it is required for all games in \mathcal{H} , independently of the number the players. We now show that players' equilibrium strategies in this new family of games has a myopic nature:

Theorem 2 Assume $X = [0, 1]$. For any δ, ρ, ϵ and any $\alpha > 0$, there exists an integer $J = J(\delta, \rho, \epsilon, \alpha)$ such that if $H^\infty \in \mathcal{H}(\delta, \rho, \epsilon)$ then any Nash equilibrium of H^∞ is a (J, α) -myopic equilibrium of H^∞ .

Similar to the proof of Theorem 1, the proof of Theorem 2 builds on a result on influence, which is equivalent of Lemma 1.

Lemma 2 *For any δ, ρ, ϵ and any $\hat{\alpha} > 0$, there exists an integer $\hat{J} = \hat{J}(\delta, \rho, \epsilon, \hat{\alpha})$ such that if $H^\infty \in \mathcal{H}(\delta, \rho, \epsilon)$ and Σ is an arbitrary profile of behavior strategies, then for any stage t there exists a subset \hat{K}_t of at most \hat{J} players such that for any $k \notin \hat{K}_t$ and pair of actions $a, b \in A^k$:*

$$|E_{\mu_{(\Sigma^{-k}, \Sigma_{(a)}^k), t}} \pi^k(d, x) - E_{\mu_{(\Sigma^{-k}, \Sigma_{(b)}^k), t}} \pi^k(d, x)| < \hat{\alpha}. \quad \forall d \in A^k$$

The proof of Lemma 2 is postponed to the appendix.

Given Lemma 2 the proof of Theorem 2 follows exactly in the steps of the proof of Theorem 1, and is therefore omitted.

3.3 A Comparison with Sabourian's Results

Our main result is different than Sabourian's [8] result in a few ways. It is weaker in some ways and stronger in another way. The relative weaknesses are: (a) We assume that agents actions are imperfectly observed by the aggregating mechanism, whereas Sabourian allows for players to recall their own actions. (b) In order to consider a large outcome space we need to impose a continuity assumption on the payoff function. (c) The action spaces are assumed finite in our model. We conjecture that this restriction is not necessary.

On the other hand our result holds without the constraint posed by Sabourian regarding the structure of the mechanism \hat{F} . In particular we do not require that \hat{F} be anonymous as Sabourian does. Furthermore we allow for all players to perfectly observe the history of signals of all players, whereas Sabourian assumes players can only recall their own actions.

4 Concluding Remarks

In this paper we provided an extension of Green-Sabourian’s analysis to games with non-anonymous outcomes and strategies. Using recent non-pivotalness results, we establish powerful bounds on how many players might act strategically. This provides a transparent and simple argument that sheds light on the forces driving Green-Sabourian’s analysis.

Two related strands of literature should be mentioned. First, non-pivotalness theorems have been used in the literature to address a variety of other questions. These include, most notably, Mailath and Postlewaite [6] result on the provision of public goods and Fudenberg, Levine and Pesendorfer [3] in their analysis of agency, reputation, and repeated games with a large player. While similar in spirit, the techniques developed in these papers do not allow for the derivation of a bound on the number of α -pivotal players independent of N as we do here (K_α^* defined above). This bound allows us to develop a clean statement of the result on approximate reversion to myopic play and provides sharp bounds on the rates of convergence.⁹

The second related strand of literature is that concerned with refining equilibria in dynamic games by appealing to Markovian restrictions on strategies. See Maskin and Tirole [7] for discussion and bibliography. The analytical and practical appeal of Markovian restrictions is obvious. However, justifying Markovian behavior hinges on the plausibility of the assumption that players’ believe their opponents limit their strategy choices to ones that condition only on payoff-relevant aspects of the environment. In our context, with no state variables, Markovian behavior reduces to playing an equilibrium of the one-shot game. In this case, our analysis may be restated as saying that most players optimally ignore payoff-irrelevant aspects of the environment in choosing their actions.¹⁰ Our paper makes explicit what sort of conditions are needed

⁹Fudenberg, Levine and Pesendorfer [3] informally suggest the possibility (p. 65) of a result in the spirit of our main results above. Their conjecture seems to claim that on average, over all players, stage game payoffs are α sub-optimal. This is strictly weaker than providing a bound on the number of individuals who do not play a myopic α -best reply.

¹⁰Note that Green-Sabourian’s analysis cannot address this issue as it *assumes* that players can

to justify players making approximate Markovian assumptions about the behavior of their opponents (e.g., noisy information about actions, large number of players, and interdependence of payoffs through a common collective outcome, and so on). The advantage of this analysis is that it gives a set of criteria for judging when the (often controversial) Markovian assumption is justified. We hope to pursue this in greater generality in the future.

condition only on payoff-relevant outcomes.

APPENDIX

The object of this appendix is to prove Lemmas 1,2. The proofs provided are based on an earlier result from Al-Najjar and Smorodinsky [1].¹¹ We begin by recalling this result.

Let $F : S \rightarrow [0, 1]$ be an arbitrary assignment from the signal space into the unit interval. Let μ^k be an arbitrary probability measure on S^k satisfying $\mu^k(\bar{s}) > \epsilon$ for any player k and any signal $\bar{s} \in S^k$. We denote by μ the product probability distribution of the μ^k s over S . Player k 's influence on F is

$$V^k = \max_{\bar{s} \in S^k} E_{\mu}(F|s^k = \bar{s}) - \min_{\bar{s} \in S^k} E_{\mu}(F|s^k = \bar{s}).$$

Note that V^k bounds the change in $E(F)$ induced by a unilateral change in the marginal μ^k . The result of Theorem 2 in Al-Najjar and Smorodinsky [1] focuses on the number of players for which this bound exceeds a given threshold $1 > \alpha > 0$:

Proposition 1 *There exists an integer $J_{\alpha}^*(\epsilon)$ such that for any $K < \infty$, any signal space S and any vector of probability measures, $\{\mu^k\}_{k=1}^K$,*

$$\#\{k : V^k > \alpha\} < J_{\alpha}^*(\epsilon).$$

Throughout the rest of the appendix we hold various parameters of the model fixed:

- The parameters $Q < \infty$, $\delta < 1$, $M < \infty$, $\rho, \epsilon > 0$ defining the families of games $\mathcal{H}(Q, \delta, M, \epsilon)$ and $\mathcal{H}(\delta, \rho, \epsilon)$.
- For any game $H^{\infty} \in \mathcal{H}(Q, \delta, M, \epsilon)$ or in $\mathcal{H}(\delta, \rho, \epsilon)$ and for any player in H^{∞} we fix an arbitrary action at stage zero, an arbitrary strategy Σ^k and an arbitrary stage t .

¹¹This manuscript is available at: <http://www.kellogg.nwu.edu/research/math/papers/1174R.pdf>

Note that we do not bound the number of players in the game.

These parameters, in turn, fix, for each game, the probability distributions $\lambda_{\Sigma,t}$ over S . As it is assumed that for each $H^\infty \in \mathcal{H}(Q, \delta, M, \epsilon)$ or $H^\infty \in \mathcal{H}(\delta, \rho, \epsilon)$ the distribution $C^k(a^k)$ is in $\Delta^\epsilon(S^k)$, we conclude that $\lambda_{\Sigma,t} \in \Delta^\epsilon(S^k)$ as well.

Let μ be an arbitrary distribution over S . We shall denote by $F(\mu)$ the corresponding distribution over X .

PROOF OF LEMMA 1:

Following Al-Najjar and Smorodinsky [1] we define the influence of player k , $V^k(H^\infty)$, by

$$V^k(H^\infty) = \max_{a,b \in A^k} \|F(\mu_{(\Sigma^{-k}, \Sigma_{(a)}^k), t}) - F(\mu_{(\Sigma^{-k}, \Sigma_{(b)}^k), t})\|.^{12}$$

Where $\|\cdot\|$ denotes the sup-norm metric on $\Delta(X)$.¹³

The number of $\bar{\alpha}$ -pivotal players in H^∞ , denoted $J_{\bar{\alpha}}(H^\infty)$, is

$$J_{\bar{\alpha}}(H^\infty) = \#\{k : V^k(H^\infty) > \bar{\alpha}\}.$$

Al-Najjar and Smorodinsky ([1], Theorem 2) provide a bound, independent of the number of players K , on $J_{\bar{\alpha}}(H^\infty)$:

Proposition 2 *For any $\bar{\alpha} > 0$ there exists an integer $J_{\bar{\alpha}}^*$ such that if $H^\infty \in \mathcal{H}(Q, \delta, M, \epsilon)$ then $J_{\bar{\alpha}}(H^\infty) < J_{\bar{\alpha}}^*$.*

In words, what Proposition 2 says is that all players but at most $J_{\bar{\alpha}}^*$ of them cannot change the distribution over X by more than $\bar{\alpha}$.¹⁴

¹²A minor difference between the original notion of influence and the one used here, is that here we take the maximum over pairs of distributions induced by actions in A^k , whereas in the original definition the maximum was over all pairs of Dirac measures. One should note that the influence defined here is, generally, smaller than the original notion.

¹³The sup-norm metric is defined as follows. For $F_1, F_2 \in \Delta(X)$, let $\|F_1 - F_2\| = \max_{B \subset X} |F_1(B) - F_2(B)|$. Note that $2\|F_1 - F_2\| = \sum_{x \in X} |F_1(x) - F_2(x)|$.

¹⁴Two comments are in place: First, because of the argument made in footnote 7 the proposition applies to the current notion of influence. Second, the statement of the result in [1] applies to the case $Q = 1$ and $M = 2$. However, as indicated by footnotes 4 and 8 in [1] the generalization to arbitrary Q and M is straightforward.

We now turn to show that small changes in the distribution over X cause small changes in the expected payoffs:

Proposition 3 *For any $\hat{\alpha} > 0$ and any action $d \in A^k$, if $F_1, F_2 \in \Delta(X)$ satisfy $\|F_1 - F_2\| < \hat{\alpha}/2Q$ then $|E_{F_1}\pi^k(d, x) - E_{F_2}\pi^k(d, x)| < \hat{\alpha}$.*

Proof:

$$\begin{aligned} |E_{F_1}\pi^k(d, x) - E_{F_2}\pi^k(d, x)| &\leq \sum_{x \in X} \pi(d, x) |F_1(x) - F_2(x)| \leq \\ &\leq Q \sum_{x \in X} |F_1(x) - F_2(x)| \leq 2Q \|F_1 - F_2\| \leq \hat{\alpha}. \end{aligned}$$

□

We can finally prove Lemma 1. Let $\bar{\alpha} = \hat{\alpha}/2Q$ and let $\hat{J} = J_{\bar{\alpha}}$, as in Proposition 2. This ensures that there exists a subset of players, denoted K_t , of size less or equal \hat{J} , such that if $k \notin K_t$ then for any pair of actions $a, b \in A^k$, $\|F(\mu_{(\Sigma^{-k}, \Sigma_{(a)}^k), t}) - F(\mu_{(\Sigma^{-k}, \Sigma_{(b)}^k), t})\| \leq \bar{\alpha}$. Proposition 3, in turn, ensures that for $k \notin K_t$

$$|E_{F(\mu_{(\Sigma^{-k}, \Sigma_{(a)}^k), t})}\pi^k(d, x) - E_{F(\mu_{(\Sigma^{-k}, \Sigma_{(b)}^k), t})}\pi^k(d, x)| < \hat{\alpha}. \quad \forall a, b, d \in A^k,$$

as required. □

PROOF OF LEMMA 2:

The analog of Proposition 2 for the large outcome space is:

Proposition 4 *For any δ, ρ, ϵ , any $\bar{\alpha} > 0$ and any measurable subset $\hat{X} \subset X$, there exists an integer $\bar{J} = \bar{J}(\delta, \rho, \epsilon, \bar{\alpha})$ such that if $H^\infty \in \mathcal{H}(\delta, \rho, \epsilon)$ and Σ is an arbitrary profile of behavior strategies, then for any stage t exists a subset of players, \hat{K}_t , of at most \bar{J} players such that for any $k \notin \hat{K}_t$ and pair of actions $a, b \in A^k$:*

$$|F(\mu_{(\Sigma^{-k}, \Sigma_{(a)}^k), t})(x \in \hat{X}) - F(\mu_{(\Sigma^{-k}, \Sigma_{(b)}^k), t})(x \in \hat{X})| < \bar{\alpha}.$$

In words, most players' actions at the first stage of the game has no effect on the probability that the outcome at time will be in an arbitrary subset of the outcome space. The claim of Proposition 4 is actually a rephrasing of Theorem 2 of Al-Najjar and Smorodinsky [1]. Consequently the proof is omitted.

We turn to the analog of Proposition 3 for the large outcome space. Fix an arbitrary $\hat{\alpha} > 0$. Let M be an integer such that for any $H^\infty \in \mathcal{H}$ and any player, k , in H^∞ , if $x, y, \in X$ satisfy $|x - y| < \frac{1}{M}$ then $|\pi^k(a^k, x) - \pi^k(a^k, y)| < \sqrt{\hat{\alpha}}$. Note that M is a function of the parameters of the family of games and is independent of the number of players in the game. Denote by $\hat{X}(m) = [\frac{m-1}{M}, \frac{m}{M})$, $m = 1, \dots, M$.

Proposition 5 *For any $\hat{\alpha} > 0$ and any action $d \in A^k$, if $F_1, F_2 \in \Delta(X)$ satisfy $|F_1(\hat{X}(m)) - F_2(\hat{X}(m))| < \sqrt{\hat{\alpha}}/M$, for all $m = 1, \dots, M$, then*

$$|E_{F_1} \pi^k(d, x) - E_{F_2} \pi^k(d, x)| < \hat{\alpha}.$$

Proof:

$$\begin{aligned} & |E_{F_1} \pi^k(d, x) - E_{F_2} \pi^k(d, x)| \leq \\ & \leq \sum_{m=1}^M |F_1(\hat{X}(m)) - F_2(\hat{X}(m))| (\max_{x \in \hat{X}(m)} \pi^k(d, x) - \min_{x \in \hat{X}(m)} \pi^k(d, x)) \leq \\ & \leq \sum_{m=1}^M |F_1(\hat{X}(m)) - F_2(\hat{X}(m))| \sqrt{\hat{\alpha}} < \sum_{m=1}^M \frac{\sqrt{\hat{\alpha}}}{M} \sqrt{\hat{\alpha}} = \hat{\alpha}. \end{aligned}$$

□

We can finally prove Lemma 2. Let M and $X(m)$, $m = 1, \dots, M$, be as in the the claim of Proposition 5. Let $\bar{\alpha} = \frac{\sqrt{\hat{\alpha}}}{M}$. By Proposition 4 there exist positive finite integers $\bar{J}(m) = J_{\bar{\alpha}}(m)$, $m = 1, \dots, M$ such that for each m there exists a subset of players, denoted $K_t(m)$, of size less or equal $\bar{J}(m)$, such that if $k \notin K_t(m)$ then for any pair of actions $a, b \in A^k$, $|F(\mu_{(\Sigma^{-k}, \Sigma_{(a)}^k), t})(x \in \hat{X}(m)) - F(\mu_{(\Sigma^{-k}, \Sigma_{(b)}^k), t})(x \in \hat{X}(m))| < \bar{\alpha}$. Let $K_t = \cup_{m=1}^M K_t(m)$. Note that K_t consists of at most $\hat{J} = \sum_{m=1}^M \bar{J}(m)$ players, and obviously \hat{J} does not depend on the number of players in the game. Proposition 5, in turn, ensures that for $k \notin K_t$

$$|E_{F(\mu_{(\Sigma^{-k}, \Sigma_{(a)}^k), t})} \pi^k(d, x) - E_{F(\mu_{(\Sigma^{-k}, \Sigma_{(b)}^k), t})} \pi^k(d, x)| < \hat{\alpha}. \quad \forall a, b, d \in A^k,$$

as required. □

References

- [1] Al-Najjar, N.I. and R. Smorodinsky (1998), "Pivotal Players and the Characterization of Influence", CMSEMS Discussion Paper no. 1174R, Northwestern Univ.
- [2] Aumann, R. J. (1964), "Mixed and Behavior Strategies in Infinite Extensive Games," in *Advances in Game Theory*, Annals of Mathematical Studies, 52 (eds. M. Dresher, L. S. Shapley and A. W. Tucker). Princeton, Princeton university press, 627-650.
- [3] Fudenberg, D., D. Levine, and W. Pesendorfer (1998), "When are Non-Anonymous Players Negligible?", *Journal of Economic Theory*, **79**, 46-71.
- [4] Green, E.J. (1980), "Non-Cooperative Price Taking in Large Dynamic Markets". *Journal of Economic Theory*, **22**, 155-182.
- [5] Kuhn, H. W. (1953), "Extensive Games and the Problem of Information," in *Contributions to the Theory of Games*, Vol. 2, Annals of Mathematical Studies, 28 (eds. R. H. Kuhn and A. W. Tucker). Princeton, Princeton university press, 137-156.
- [6] Mailath, G.J. and A. Postlewaite, (1990) "Asymmetric Information Bargaining Problems with Many Agents", *Review of Economic Studies*, **57**, 351-367.
- [7] Maskin, E. and J. Tirole (1994). "Markov Repeated Equilibrium," mimeo, Harvard university.
- [8] Sabourian, H. (1990), "Anonymous Repeated Games with a Large Number of Players and Random Outcomes". *Journal of Economic Theory*, **51**, 92-110.