EQUILIBRIUM PAYOFFS IN STOCHASTIC GAMES WITH GRADUAL STATE CHANGES

MARCIN PĘSKI, THOMAS WISEMAN

Abstract. We characterize perfect public equilibrium payoffs in discounted stochastic games with a continuous state space, in the case where the length of the period shrinks. The probability (conditional on the action profile and the current state) that the state changes in any given direction is fixed with respect to the period length, but the *magnitude* of the change is proportional to the period length.

1. Introduction

We consider discounted stochastic games with imperfect public monitoring and a continuous state space, where the magnitude of state transitions are proportional to the period length. Players' actions and the current state influence state transitions, and the stage-game payoffs and monitoring technology vary with the state. We characterize perfect public equilibrium payoffs for short period lengths, holding the player's rate of time discounting. In the limit, the discounting between periods shrinks to zero, but the rate at which the state changes remains fixed. In particular, the set of equilibrium payoffs typically depends on the initial state, because the state changes only gradually.

The paper makes two main contributions. First, we present a meaningful definition of the feasible and individually rational payoff sets in this setting, for each initial state. To do this, we define the "pseudo-instantaneous payoff" of an action in a given state, given a function mapping tomorrow's state to a continuation value, as the sum of the current stage-game payoff and the expected rate of change in the continuation value

Date: October 8, 2014.

PRELIMINARY AND INCOMPLETE..

We thank Drew Fudenberg, Johannes Hörner, Nicolas Vieille, and Jim Wiseman for helpful conversations.

Pęski: Department of Economics, University of Toronto, mpeski@gmail.com. Wiseman: Department of Economics, University of Texas at Austin, wiseman@austin.utexas.edu.

that the action generates. (That is, an action that yields a high payoff today may have a lower pseudo-instantaneous payoff than an action that yields a low payoff today but is more likely to lead to a state with a high continuation value.) We say that a collection $F = \{F(s)\}_s$ of available continuation payoffs in each state s is *self-feasible* if each payoff in each s can be achieved as the psuedo-instantaneous payoff from some current action and a continuation payoff function that maps each new state $s^{'}$ to a value in $F(s^{'})$. The collection is *self-individually rational* if for each s and each player, every payoff in $F(s)$ exceeds the minmax pseudo-instantaneous payoff (again using continuation payoffs that take values in F). We define V^{δ} as the largest self-feasible and self-individually rational collection, given per-period discount factor δ, and show that for any initial state s, the set of perfect public equilbirium payoffs starting from s must lie in V^{δ}.

The second contribution of the paper is to strengthen the definition of a self-feasible and self-individually rational collection in a parametrized way, and to show that the maximal collection satisfying any such strengthened notion is contained in the equilibrium set for high enough δ (that is, for short enough period lengths.) That result requires that Fudenberg, Levine, and Maskin's (1994) "FLM" conditions on public monitoring hold uniformly across states. We also show that for any fixed δ, the maximal collection under a strengthened notion converges (as we "weaken" the strengthening) to the maximal collection under the original definition, for generic games. In that sense, the upper and lower bounds on the equilibrium set approach each other.

Stochastic games relax the strong restriction in repeated games that the conditions under which players interact in one period are independent of outcomes in previous periods. In many important economic examples of such settings, players can adjust their actions frequently, while the state changes only gradually. For instance, competing oligopolists may set prices daily, while market demand depends on slow shifts in the macroeconomy. Similarly, firms can adjust advertising expenditures quickly, but brand awareness changes slowly. Another example is natural resource depletion: the owners of fishing boats can decide every day whether or not to go out, and each such decision has only a very small effect on the current future stock of fish.

Relatively little previous work on stochastic games has focused on this case. Instead, much of the literature has examined an alternative limiting case, with a finite state space: fix the period length (and transition rates), and let players become very patient. In that case, the discounted time to transition between states shrinks to zero. Dutta (1995) derives a perfect-monitoring folk theorem for that environment; Fudenberg and Yamamoto (2011), and Hörner et al. (2011) extend that result to imperfect monitoring. All three results require, in contrast to the setting here, that the set of PPE payoffs be independent of the initial state as the discount factor δ approaches 1. To guarantee that independence, those authors use the condition that no single player can prevent the Markov process governing the state from being irreducible. We do not assume irreducibility, because, as we mention above, in our model the set of equilibrium payoffs typically depends on the initial state regardless. Without irreducibility, we can allow, for example, individual firms to go bankrupt and permanently exit the game.

The paper most closely related to this one is Pęski and Wiseman (forthcoming). There, the state space is finite (or countably infinite), and the probability of state transition is proportional to the period length. That is, there transitions of fixed magnitude occur with probabilities proportional to the period length, while here transitions whose magnitude is proportional to the period length occur with fixed probabilities. In both cases the rate of time preference is fixed. For the limiting case as period length shrinks, Pęski and Wiseman (forthcoming) obtain a somewhat sharper characterization of equilibrium payoffs than we provide here. We discuss the difference in Section 4.2.

The rest of the paper is organized as follows: we describe the model in Section 2. In Section 3 we define pseudo-instantaneous payoffs, self-feasibility and self-individual rationality, and the strengthened versions of those concepts, and we present a simple example of how to compute the maximal self-feasible and self-individually rational payoff collection. We present our two main theorems in Section 4, and Section 5 is the conclusion.

2. Model

2.1. **Stochastic games with a continuous state space.** There are N expected-utility maximizing players playing an infinite-horizon stochastic game. The time between periods is given by $\Delta > 0$, and all players discount the future at rate $r > 0$, so that the per-period discount rate is $e^{-r\triangle} \equiv \delta$. We will consider the case that the time between periods shrinks to zero, so that the length of a period is proportional to $1 - \delta$. The set of stage-game action profiles is $A = A_1 \times A_2 \times \cdots \times A_N$, where A_i is the set of actions for player i. After each period, players observe a public signal of actions, y, drawn from a finite set Y.

There is a K-dimensional set S of states of the world; S is a closed, bounded, convex subset of $\mathbb{R}^K$. (The value of K is described below.) At the end of each period, the state changes randomly. There is a finite set Z of K-dimensional vectors representing possible state changes (both direction and magnitude). The function $\Phi : S \to \mathbb{R}^K$ (also described below) maps each state s to a vector that specifies, for each action profile, the resulting probabilities of each public signal and state change, as well as the vector of payoffs for each action profile-public signal pair.

At the start of each period, the state s is publicly observed. Then each player i chooses an action $a_i \in A_i$, and then all players observe the public signal y. In state s, the public signal is distributed according to $\rho_Y(a, s)$, where a is the profile of actions of all players.

Player i's payoff in state s when action a is played and signal y occurs is $u_i\left(a, y, s\right)$; let $g_i(a, s)$ be player i's expected payoff when action profile a is played in state s:

$$g_i(a, s) = \sum_{y \in Y} \left[\rho_Y\left(a, s\right)\left(y\right)\right] u_i\left(a, y, s\right).$$

Denote by $g(a, s)$ the vector of expected payoffs for each player. We assume that the payoffs are uniformly bounded by $M \equiv \max_{a,s} \|g(a, s)\| < \infty$ (where $\|\cdot\|$ denotes the Euclidean norm). It will be convenient to work directly with the expected payoff function $g(a, s)$ rather than with $u_i\left(a, y, s\right)$.

When action profile a is played in state s, the probability that state change $z \in Z$ is drawn is $\rho_Z\left(a, s\right)[z]$, and next period's state is then $s + (1 - \delta)z$. That is, the magnitude of the state transition is proportional to the length of the period. Let $z_{max} = \max_{z \in Z} \|z\| < \infty$.

Let $\rho(a, s)$ denote the joint distribution over public signal-state change pairs. Let m_i denote the number of actions available to player i ($m_i \equiv \# A_i$), m_Y the number of public signals in Y, and m_Z the number of state changes in Z. Then $\Phi(s) = (\rho(\cdot, s), g(\cdot, s))$. That is, for each state, Φ specifies an element of $(\Delta(Y \times Z))^{\# A} \times \left(\mathbb{R}^N\right)^{\# A}$, so K is equal to $(m_Y \cdot m_Z - 1 + N) \cdot \# A$.

To summarize, a game consists of a set of N players; a collection (S, A, Y, Z) of states, action profiles, public signals, and state changes; the mapping Φ from states to the parameters (including transition probabilities) of the stage game; an initial state $s_0 \in S$; and δ.

All these definitions extend in a natural way to mixed actions. This structure is common knowledge. We assume that a public randomization device is available to the players.

The set of *public histories* in period t is equal to $H_t \equiv Y^{t-1} \times Z^{t-1} \times S^t$, with element $h_t = (s_1, y_1, z_1, \ldots, y_{t-1}, z_{t-1}, s_t)$, where s_t denotes the state at the beginning of period t, and y_t and z_t denote the public signal and state change, respectively, realized at the end of period t. Player *i's private history* in period t is $h_t^i = (s_1, y_1, z_1, a_{i,1}, \ldots, y_{t-1}, a_{i,t-1}, z_{t-1}, s_t)$, where $a_{i,t}$ is player i's action in period t; $H_t^i \equiv (Y \times Z \times A_i)^{t-1} \times S^t$ is the set of such private histories. Define $H \equiv \bigcup_{t \geq 1} H_t$ and $H^i \equiv \bigcup_{t \geq 1} H_t^i$. For any history h_t, let $s(h_t) \equiv s_t$ denote the current state.

A strategy for player i is a mapping $\alpha_i : H^i \to \triangle A_i\left(s\left(h_t^i\right)\right)$, and a public strategy for player i is a mapping $\sigma_i : H \to \triangle A_i\left(s\left(h_t\right)\right)$. Let Σ_i and Σ_i^P, respectively, denote the set of strategies and the set of public strategies for player i; let Σ and Σ^P denote the sets of strategy profiles and of public strategy profiles, respectively. Given per-period discount factor $\delta < 1$, a profile of strategies $\alpha \in \Sigma$, and an initial state $s \in S$, the vector of expected payoffs in the dynamic game is given by

$$v^\delta(\alpha, s) = (1 - \delta) E \sum_{t=1}^{\infty} \delta^{t-1} g(a_t, s_t),$$

where the expectation is taken with respect to the distribution over actions and states induced by the strategy α and initial state s. For each public strategy $\sigma \in \Sigma^P$ and public history $h \in H$, the continuation payoffs $v^{\delta}(\sigma, h)$ are calculated in the usual way.

Given $\delta < 1$, define the (convex hull of the) set of feasible payoffs in initial state s, $\hat{V}^{\delta}(s)$, as

$$\hat{V}^{\delta}(s) \equiv \text{co}\left\{v^{\delta}(\sigma, s) : \sigma \in \Sigma^{P}\right\}.$$

Note that the set of feasible payoffs $\hat{V}^{\delta}(s)$ varies with the state even in the limit as δ approaches 1 (that is, as the time between actions shrinks to 0): the payoffs to the stage game vary with the state, and both the discount rate and the rate at which the state changes are fixed per unit of time as δ grows. The set $\hat{V}^{\delta}(s)$ is *not*, in general, equal to the set of feasible payoffs for the stage game in state s.

A *public perfect equilibrium* in a game with per-period discount factor δ and initial state s is a strategy profile σ such that for each public history h with $s_1 = s$, each player i, and each strategy σ_i' of players i, $v_i^{\delta}(\sigma, h) \geq v_i^{\delta}(\sigma_i', \sigma_{-i}, h)$. Let $E^{\delta}(s)$ be the set of payoffs obtained in public perfect equilibria of the game, given δ and initial state s. In most of this paper, "equilibrium" means public perfect equilibrium.

2.2. **Identifiability.** The definitions of individual full rank and pairwise full rank, conditions relating to the identifiability of players' actions from public information, are the same as in FLM, except that the public information here includes both the public signal y and the state change z (as in Hörner et al. (2011)). Recall that m_i is the number of actions available to player i, m_Y is the number of public signals, and m_Z is the number of stage changes. For each state s, player i, and (mixed) action profile α, let $\Pi_i(\alpha, s)$ be the $m_i \times (m_Z \times m_Y)$ matrix whose rows correspond to the probability distribution over (public signal, state change) pairs induced by each of player i's actions, given s and α_{-i}: $\Pi_i(\alpha, s) \equiv \rho((\cdot, \alpha_{-i}), s)$. Similarly, for each state s and action profile α, let $\Pi_{ij}(\alpha, s)$ be the $(m_i + m_j) \times (m_Z \times m_Y)$ matrix whose first m_i rows are $\Pi_i(\alpha, s)$ and whose last m_j rows are $\Pi_j(\alpha, s)$.

Action profile α has *individual full rank* in state s if $\Pi_i(\alpha, s)$ has rank m_i for each player i. Action profile α has *pairwise full rank* for players i and j in state s if $\Pi_{ij}(\alpha, s)$ has rank $m_i + m_j - 1$. If an action profile α has individual full rank, then any player i's actions are distinguishable probabilistically (given that the other players are playing α_{-i}). If α has pairwise full rank for i and j, then deviations from α by player i are distinguishable from player j's deviations.

Hörner et al.'s (2011) identifiability requirements are based on those full-rank conditions. In order to ensure that the rank conditions hold uniformly over all states, we require some more notation (introduced in Pęski and Wiseman (forthcoming)).

Let $\mathcal{M}_{kl}$ be the set of $k \times l$ matrices, and let $\mathcal{M}_l$ be the set of square l-matrices. Given $j \leq k, l$ and matrices $A \in \mathcal{M}_{kl}$ and $B \in \mathcal{M}_j$, we write $B \subseteq A$ if matrix B can be obtained from A by crossing out $k - j$ rows and $l - j$ columns. Let

$$d_j (A) = \max_{\{B \in \mathcal{M}_j : B \subseteq A\}} |\det B| .$$

Thus, $d_j (A) > 0$ if and only if the rank of matrix A is not smaller than j. Individual full rank for action α in state s is equivalent to the condition $d_{m_i} (\Pi_i (\alpha, s)) > 0$ for each player i, and pairwise full rank for players i and j is equivalent to the condition $d_{m_i+m_j-1} (\Pi_{ij} (\alpha, s)) > 0$.

Given scalar $d > 0$, say that action α has *individual d-rank* in state s if $d_{m_i} (\Pi_i (\alpha, s)) \geq d$ for each player i. Similarly, say that α has *pairwise d-rank* for players i and j in state s if $d_{m_i+m_j-1} (\Pi_{ij} (\alpha, s)) \geq d$.

With those definitions, we can state the identifiability condition on the monitoring structure:

Definition 1. *Identifiability Condition:* There exists $d > 0$ such that for each state s,

(1) every pure action profile has individual d-rank in state s,
(2) for each pair of players i and j, there exists a profile $\alpha(s)$ that has pairwise d-rank for i and j in state s.

Our identifiability condition is an extension of Hörner et al.'s (2011) Assumptions F1 and F2 to an infinite state space.

2.3. **Examples.** In the first example, we consider a repeated Cournot oligopoly where a firm's costs depend on its capital stock.

Example 1. There are N firms, and each firm can produce either high output q^H or low output q^L. A firm also chooses its level of capital investment, x^H or x^L. The per-period profits of firm i are equal to

$$q_i P\left(q_1 + ... + q_N\right) - c_i\left(q_i, k_i\right) - x_i,$$

where $k_i \in [0, \bar{k}_i]$ is firm i's publicly observed capital stock. Its evolution depends on the investment x_i and on the current stock. In particular, capital may increase by a lot, increase a little, stay the same, decrease a little, or decrease a lot: $Z = \{-2\Delta, -\Delta, 0, \Delta, 2\Delta\}$. The probability that next period's capital stock is $k_i + (1-\delta)z$ is equal to $\rho_Z^i\left(x_i, k_i\right)[z] \geq 0$ for $z \in Z$, where $\rho_Z^i\left(x^H, k_i\right)$ first-order stochastically dominates $\rho_Z^i\left(x^L, k_i\right)$ for all i and all k_i. Firms' capital stocks evolve independently of each other: $\rho_Z\left(x, s\right)[z_1, \ldots, z_N] = \prod_i \rho_Z^i\left(x_i, k_i\right)[z_i]$.

The second example is a reciprocal effort game with two players. At each extreme state, one player is able to do a favor for the other, generating 7 units of utility at a cost of 1 unit of utility to himself. The payoffs at intermediate states are the weighted averages of those at the extremes. Total payoffs are maximized when both players do favors. Because of the simple structure of monitoring, transition rates, and stage-game payoffs in this example, we can reduce the state space to one dimension.

Example 2. There are two players, and each player has two actions available in the stage game, 1 (corresponding to exerting effort) and 0 (no effort). Monitoring is perfect. The state space is the unit interval, and the space of possible transitions is $Z = \{-1, 0, 1\}$. Transition probabilities are independent of actions: $\rho(a, s)[1] = \rho(a, s)[-1] = \frac{1}{3}$ for all a and all $s \in (0, 1)$. At the extreme states, $\rho(a, 1)[1] = \rho(a, 0)[-1] = 0$ and $\rho(a, 1)[-1] = \rho(a, 0)[-1] = \frac{1}{3}$. Stage-game payoffs are linear in the state s and are given by

$$\underset{\text{State } s=0}{\begin{array}{c|c|c|} & 0 & 1 \\ \hline 0 & 0,0 & 0,0 \\ \hline 1 & -1,7 & -1,7 \\ \hline \end{array}}, \quad \underset{\text{State } s=1}{\begin{array}{c|c|c|} & 0 & 1 \\ \hline 0 & 0,0 & 7,-1 \\ \hline 1 & 0,0 & 7,-1 \\ \hline \end{array}}, \text{ and, in general, } \underset{\text{State } s}{\begin{array}{c|c|c|} & 0 & 1 \\ \hline 0 & 0,0 & 7s,-s \\ \hline 1 & s-1, 7(1-s) & 8s-1, 7-8s \\ \hline \end{array}}.$$

We will return to Example 2 in Section 3.3.

Finally, in a standard repeated game, play never leaves the initial state.

Example 3. We say that a stochastic game is a *(standard) repeated game* if for each state $s \in S$ and each action profile $a \in A$, $\rho_Z(a, s)(0) = 1$.

3. Characterizing Payoffs

As a preliminary, we define *pseudo-instantaneous payoffs*, a piece of notation that is useful in capturing the effects that an action profile has on both today's payoff and expected continuation payoffs (through influencing transition rates).[1] In particular, using pseudo-instantaneous payoffs will be convenient because both the magnitude of a state transition and the weight on today's payoff in overall utility are proportional to $1 - \delta$.

3.1. **Pseudo-instantaneous payoffs.**

Let $u : S \to \mathbb{R}^N$ specify a vector of continuation payoffs, as a function of next period's state. For each such u, state s, and action profile a, we define a vector of payoffs

$$\psi^\delta (a, s, u) \equiv g(a, s) + \delta \sum_{z \in Z} \rho_Z (a, s) [z] \frac{u(s + (1 - \delta)z) - u(s)}{1 - \delta}. \tag{3.1}$$

We refer to $\psi^\delta(a, s, u)$ as the *pseudo-instantaneous payoff* from playing action profile a in state s, given continuation payoff function u. To motivate this construction, let state s be given, and choose any mapping u from $S \setminus \{s\}$ to $\mathbb{R}^N$. Let v denote the expected payoff in state s from playing the same profile a in each period until the state changes, followed by continuation payoffs specified by u. Complete the definition of u by setting $u(s) = v$. Then v is exactly equal to $\psi^\delta(a, s, u)$:

$$v = (1 - \delta) g(a, s) + \delta \left(\sum_{z \neq 0} \rho_Z (a, s) [z] u(s + (1 - \delta)z) \right) + \delta \rho_Z (a, s) [0] v$$

$$\Rightarrow$$

$$v = g(a, s) + \delta \sum_{z \in Z} \rho_Z (a, s) [z] \tfrac{u(s + (1 - \delta)z) - u(s)}{1 - \delta} = \psi^\delta (a, s, u).$$

[1] These pseudo-instantaneous payoffs are similar in spirit to those in Pęski and Wiseman (forthcoming), which have the same name but a somewhat different definition.

More generally, we can represent the expected payoff from playing profile a for one period, given continuation payoffs u, as a convex combination of the pseudo-instantaneous payoff and $u(s)$

$$(1-\delta)\, g\,(a,s) + \delta \left(\sum_{z \in Z} \rho_Z\,(a,s)\,[z] u(s + (1-\delta) z) \right)$$
$$= (1-\delta)\, \psi^{\delta}\,(a,s,u) + \delta u(s).$$

Observe that for δ close to 1, the pseudo-instantaneous payoff is roughly the stage-game payoff plus the expected derivative of continuation payoffs. Since the magnitude of the state change is proportional to $(1-\delta)$, the expected change in continuation payoffs is proportional to $(1-\delta)$times that derivative. The weight on the stage-game payoff also is $(1-\delta)$, so the pseudo-instantaneous payoff allows us to combine to payoff effects of the same order. We use pseudo-instantaneous payoffs in characterizing equilibrium payoffs.

3.2. **Feasible and individually rational payoffs.** Define a collection of payoff sets F as a correspondence mapping each state $s \in S$ into $[-M,M]^N$. (Recall that M is the maximal length of any stage-game payoff vector.) For each state s and payoff vector $v \in F(s)$, let $U_{F,s,v}$ denote the set of functions u from S to $[-M,M]^N$ such that i) $u(s) = v$, and ii) $u(s') \in F(s')$ for all s'. Say that collection F is *self-δ-feasible* for $\delta < 1$ if for each s,

$$F(s) \subseteq \operatorname{co} \left\{ v : v = \psi^{\delta}\,(a,s,u) \ \text{ for some } a \in A(s), u \in U_{F,s,v} \right\}.$$

Self-δ-feasibility means that each payoff in $F\,(s)$ can be generated as the expected payoff from some action profile in the state-s stage game followed by continuation payoffs that belong to collection F. The definition has a fixed point flavor.

For each player i and $\delta < 1$, define the *δ-minmax payoff relative to F* for player i in state s as

$$e_i^{\delta}\,(s;F) \equiv \inf_{v \in F(s)} \left\{ \inf_{\alpha_{-i} \in \times_{j \neq i} \Delta A_j, u \in U_{F,s,v}} \left\{ \max_{a_i \in A_i} \psi_i^{\delta}\,((a_i, \alpha_{-i}), s, u) \right\} \right\}.$$

Say that the collection F is *self-δ-individually rational* if for each state s, player i, and $v \in F\,(s)$, $v_i \geq e_i^{\delta}\,(s;F)$.

The (straightforward) proof of the following claim is in Appendix A:

Remark 1. The collection of feasible payoffs $\hat{V}^\delta \equiv \left\{\hat{V}^\delta(s)\right\}_s$ is self-δ-feasible, and the collection of equilibrium payoffs E^δ is both self-δ-feasible and self-δ-individually rational.

Note that self-δ-feasibility and self-δ-individually rationality together imply an ex post notion of individual rationality: each payoff above the minmax payoffs can be generated using continuation payoffs that are themselves above the minmax levels.

Next, we define stronger versions of these concepts to be used in constructing equilibria. Let $B(v, \epsilon)$ denote the closed ball centered at v with radius $\epsilon \geq 0$. Given a set V, let $\bar{B}(V, \epsilon) \equiv \bigcup_{v \in V} B(v, \epsilon)$ denote the closed ϵ-neighborhood of V. For any constant $C > 0$, let $U^C_{F,s}$ denote the set of functions u from S to $[-M, M]^N$ such that for all s' (including state s), i) $u(s') \in F(s')$, and ii) $\|u(s) - u(s')\| \leq \frac{1}{C}\|s - s'\|$. Extending the definition for $C = 0$, define $U^0_{F,s}$ as the set of functions u from S to $[-M, M]^N$ such that $u(s') \in F(s')$ for all s' (including state s).

For $\delta < 1$, $C \geq 0$, $\epsilon \geq 0$, and $d \geq 0$, say that collection F is *self-δ, C, ϵ, d-feasible* if for each s,

$$\bar{B}\left(F\left(s\right), (1-\delta)\epsilon\right) \subseteq \mathrm{co}\left\{v : v = (1-\delta)\psi^\delta\left(\alpha, s, u\right) + \delta u(s) \text{ for some } a \in A^d(s), u \in U^C_{F,s}\right\},$$

where $A^d(s)$ denotes the set of action profiles that have pairwise d-rank for all players i and j in state s.

Similarly, define the *δ, C, d-minmax payoff relative to F* for player i in state s, $e_i^{\delta,C,d}(s; F)$, as follows. Let $A^i(s, d) \subseteq \times_{j \neq i} \Delta A_j$ denote the set of (independently mixed) strategy profiles α_{-i} for players $j \neq i$ such that for any $a_i \in A_i$, the profile (a_i, α_{-i}) has individual d-rank in state s for each player $j \neq i$. Then define

$$e_i^{\delta,C,d}(s; F) \equiv \inf_{\alpha_{-i} \in A^i(s,d), u \in U^C_{F,s}} \left\{\max_{a_i \in A_i} (1-\delta)\psi_i^\delta\left(\left(a_i, \alpha_{-i}\right), s, u\right) + \delta u_i(s)\right\}.$$

Say that collection F is *self-δ, C, ϵ, d-individually rational* if for each state s, player i, and $v \in F(s)$, $v_i \geq e_i^{\delta,C,d}(s; F) + (1-\delta)\epsilon$. For brevity, we will say that a collection is *self-δ, C, ϵ, d-FIR* if it is self-δ, C, ϵ, d-feasible and self-δ, C, ϵ, d-individually rational.

In words, a collection is self-δ, C, ϵ, d-FIR if for each state s, every payoff close (within $(1-\delta)\epsilon$) to $F(s)$ can be attained as the weighted average of a pseudo-instantaneous payoff (from an action α and a continuation payoff function u with values in F) and $u(s)$, where α must have pairwise d-rank for all pairs of players and u must not vary too much with the state (no more than $\frac{1}{C}$ times the distance between the states). In Section 4 and Appendix C.2, we highlight the roles that strictly positive values of C, ϵ, and d play in the proof of Theorem 1. In Section 4.2, we discuss how much we lose (in terms of describing equilibrium payoffs) when we require self-δ, C, ϵ, d-FIR payoffs for small C, ϵ, and d, relative to the set of self-$\delta, 0, 0, 0$-FIR payoffs.

The following lemma shows that for any non-negative C, ϵ, and d, there is a largest self-δ, C, ϵ, d-FIR collection. It also shows that the largest self-$\delta, 0, 0, 0$-FIR collection is exactly the largest self-δ-FIR collection. (Proofs are in the appendix.)

Lemma 1. *(1) For each $\delta < 1$, $\epsilon \geq 0$, $C \geq 0$, and $d \geq 0$, there exists the largest collection $V^{\delta}_{C,\epsilon,d}$ such that $V^{\delta}_{C,\epsilon,d}(s) \subseteq [-M, M]^N$ for each $s \in S$ and $V^{\delta}_{C,\epsilon,d}$ is self-δ, C, ϵ, d-FIR.*

(2) Each $V^{\delta}_{C,\epsilon,d}(s)$ is compact and convex.

(3) $V^{\delta}_{0,0,0}$ is equivalent to the largest collection V^{δ} such that $V^{\delta}(s) \subseteq [-M, M]^N$ for each $s \in S$ and V^{δ} is self-δ-FIR.

We will refer to elements of $V^{\delta}_{C,\epsilon,d}(s)$ and $V^{\delta}(s)$ as self-δ, C, ϵ, d-FIR payoffs and self-δ-FIR payoffs, respectively, in state s.

We know from Remark 1 that $E^{\delta}(s) \subseteq V^{\delta}_{0,0,0}(s)$ for each state s. That result implies that each $V^{\delta}_{0,0,0}(s)$ is nonempty, because the arguments of Mertens and Parthasarathy (1987, 1991) and Solan (1998) ensure that a PPE exists.

3.3. **Example.** In this section, we return to Example 2 and demonstrate how to calculate the collections of feasible payoffs $\hat{V}^{\frac{1}{2}}$ and self-δ-FIR payoffs V^{δ}. To simplify the calculations, consider the case where the period length is $\frac{1}{2}$ $(1-\delta = \delta = \frac{1}{2})$ and the initial state $s_0 \in \{0, \frac{1}{2}, 1\}$ - those three states are thus the only ones that are reachable.

First, we compute the set of feasible payoffs for each initial state, $\hat{V}^{\frac{1}{2}}(s)$. Those sets are spanned by the payoffs from Markov strategies (that is, those where a player's action depends only on the current state). There are 64 pure Markov strategies (2 actions per player in each of the three reachable states), but since player 2's action is irrelevant in state 0, and player 1's is irrelevant in state 1, there are effectively only 16 combinations. From those, it can be shown that the extreme payoffs are generated by the following four profiles: $a^{11} \equiv a_1(s) = a_2(s) = 1\ \forall s$, $a^{00} \equiv a_1(s) = a_2(s) = 0\ \forall s$, $a^{10} \equiv a_1(s) = 1, a_2(s) = 0\ \forall s$, and $a^{01} \equiv a_1(s) = 0, a_2(s) = 1\ \forall s$. (Recall that total payoffs are maximized when both players exert effort.) The sets of feasible payoffs from each intitial state are thus

$$\begin{aligned}
\hat{V}^{\frac{1}{2}}(0) &= \text{co}\left\{(0,6),(0,0),(\tfrac{-7}{8},6\tfrac{1}{8}),(\tfrac{7}{8},\tfrac{-1}{8})\right\},\\
\hat{V}^{\frac{1}{2}}(\tfrac{1}{2}) &= \text{co}\left\{(3,3),(0,0),(\tfrac{-1}{2},\tfrac{7}{2}),(\tfrac{7}{2},\tfrac{-1}{2})\right\},\\
\hat{V}^{\frac{1}{2}}(1) &= \text{co}\left\{(6,0),(0,0),(\tfrac{-1}{8},\tfrac{7}{8}),(6\tfrac{1}{8},\tfrac{-7}{8})\right\}.
\end{aligned}$$

For instance, the payoffs $v^{\frac{1}{2}}(a^{11},s)$ are found by solving the following three equations:

$$\begin{aligned}
v^{\frac{1}{2}}(a^{11},0) &= \tfrac{1}{2}g((1,1),0)+\tfrac{1}{2}\left[\tfrac{1}{3}v^{\frac{1}{2}}(a^{11},\tfrac{1}{2})+\tfrac{2}{3}v^{\frac{1}{2}}(a^{11},0)\right]\\
v^{\frac{1}{2}}(a^{11},\tfrac{1}{2}) &= \tfrac{1}{2}g((1,1),\tfrac{1}{2})+\tfrac{1}{2}\left[\tfrac{1}{3}v^{\frac{1}{2}}(a^{11},\tfrac{1}{2})+\tfrac{1}{3}v^{\frac{1}{2}}(a^{11},0)+\tfrac{1}{3}v^{\frac{1}{2}}(a^{11},1)\right]\\
v^{\frac{1}{2}}(a^{11},1) &= \tfrac{1}{2}g((1,1),1)+\tfrac{1}{2}\left[\tfrac{1}{3}v^{\frac{1}{2}}(a^{11},\tfrac{1}{2})+\tfrac{2}{3}v^{\frac{1}{2}}(a^{11},1)\right].
\end{aligned}$$

Next, we can derive the the self-δ-FIR payoffs, $V^{\frac{1}{2}}(s)$. Note that a player can guarantee himself a payoff of 0 by never exerting effort, so $V^{\frac{1}{2}}(s)$ cannot be greater than the positive quadrant of the feasible set $\hat{V}^{\frac{1}{2}}(s)$. We can show that $V^{\frac{1}{2}}(s)$ is in fact equal to that upper bound:

$$\begin{aligned}
V^{\frac{1}{2}}(s=0) &= \text{co}\left\{(0,6),(0,0),(\tfrac{6}{7},0)\right\},\\
V^{\frac{1}{2}}(s=\tfrac{1}{2}) &= \text{co}\left\{(3,3),(0,0),(\tfrac{-1}{2},\tfrac{7}{2}),(\tfrac{7}{2},\tfrac{-1}{2})\right\},\\
V^{\frac{1}{2}}(s=1) &= \text{co}\left\{(6,0),(0,0),(0,\tfrac{6}{7}),(6\tfrac{1}{8},\tfrac{-7}{8})\right\}.
\end{aligned}$$

Consider, for example, the payoffs $(0,6)$ in state 0. $(6,0)$ equals the pseudo-instantaneous payoff in state 0 from stage-game action $(1,1)$ and the continuation payoff function $u(0)=(0,6)$ and $u(\frac{1}{2})=(3,3)$. (Since state 1 is not reachable from state 0, the value $u(1)$ is irrelevant.)

$$\psi^{\frac{1}{2}}\left((1,1),0,u\right)=(-1,7)+\frac{1}{2}\left[\frac{1}{3}\frac{((3,3)-(0,6))}{\frac{1}{2}}\right]=(0,6).$$

Thus, $(6,0)$ is $\frac{1}{2}$-feasible with respect to $V^{\frac{1}{2}}$. To see that it is $\frac{1}{2}$-individually rational, note that $e_i^{\frac{1}{2}}\left(s;V^{\frac{1}{2}}\right)=0$ for each player in each state: a player can guarantee himself at least 0 in the stage game, and $(0,0)$ is available in $V^{\frac{1}{2}}(s')$ for each reachable state s'. Similar constructions yield the other payoffs in $V^{\frac{1}{2}}(0)$ and all the payoffs in $V^{\frac{1}{2}}(\frac{1}{2})$ and $V^{\frac{1}{2}}(1)$.

4. Equilibrium Payoffs for Large δ

In this section, we show that when the Identifiability Condition holds, then for any $C,\epsilon,d>0$, any self-δ,C,ϵ,d-FIR payoff at state s can be attained in a perfect public equilibrium from initial state s for sufficiently high δ: $V^{\delta}_{C,\epsilon,d}$ is contained in E^{δ}. The proof of that result is based on techniques in the proof of FLM's folk theorem for repeated games with imperfect public monitoring.

FLM's folk theorem r4equires that the set of feasible and individually rational payoffs have nonempty interior. The role of that condition is to guarantee that after any history, it is possible to provide incentives by constructing continuation payoffs that lie in any direction from the target payoffs. Here, that full dimensionality is implied by the definition of $V^{\delta}_{C,\epsilon,d}$: self-$C,\delta,\epsilon,d$-feasibility means that for any payoff $v\in V^{\delta}_{C,\epsilon,d}(s)$, payoffs near v can be generated by some action profile and some continuation payoffs in $V^{\delta}_{C,\epsilon,d}$.

Theorem 1. *Suppose that the Identifiability Condition holds. Then for each $C>0$, $\epsilon>0$, and $d>0$, there exists $\delta^*<1$ such that $V^{\delta}_{C,\epsilon,d}(s_0)\subseteq E^{\delta}(s_0)$ for any initial state s_0 and any $\delta\geq\delta^*$.*

4.1. **Proof of Theorem 1.** FLM's proof shows that any smooth set of payoffs W strictly in the interior of the feasible and individually rational set can be attained in equilibrium. A key step is to show that any payoff on the boundary of W can be achieved as the weighted average of a stage-game payoff in the current period that lies outside W (thus the requirement that W is strictly in the interior of the feasible set) and expected continuation payoffs that lie in W. Here, we want to do something similar, with pseudo-instantaneous payoffs taking the place of the stage-game payoffs. For $\epsilon > 0$, the self-C, δ, ϵ, d-feasibility of $V^{\delta}_{C,\epsilon,d}$ ensures that for each state s, there is a pseudo-instantaneous payoff outside $V^{\delta}_{C,\epsilon,d}(s)$ in each direction.

Another key step in FLM's proof is to construct continuation payoffs that make players indifferent among all their actions (and thus willing to play the equilibrium action). Here, when $C > 0$ and $d > 0$, we can ensure that those continuations payoffs are close enough to each other that they can be made to lie within $V^{\delta}_{C,\epsilon,d}\left(s'\right)$ for each possible "tomorrow's" state s'. The intuition, roughly, is that the more detectable deviations are (higher d), the less variation in continuation payoffs is needed to deter deviations. And the less continuation payoffs vary with the state (higher C), the less tempted a player is to deviate in a way that increases the probability of transitioning to a state where continuation payoffs are high. (See Corollary 1 and Lemma 4.)

Given a state s, let $V \subseteq \mathbb{R}^N$ be a set of payoffs, and let $W = \left\{W\left(s'\right)\right\}_{s'}$, where each $W\left(s'\right) \subseteq \mathbb{R}^N$, be a collection of payoff sets. Extending FLM and Abreu, Pearce, and Stacchetti (1986, 1990), we say that V is *decomposable with respect to* δ *and* W *in state* s if for each $v \in V$, there exist a mixed action profile α and a function $w : Y \times Z \to [-M, M]^N$ satisfying $w(y, z) \in W(s + (1-\delta)z)$ such that for each player i and each action $a_i \in A_i$,

$$\begin{aligned} v_i =&(1-\delta)g_i(\alpha, s) + \delta \sum_{(y,z)\in Y\times Z} \rho(a, s)[y, z]w_i(y, z) \\ \geq&(1-\delta)g_i((a_i, \alpha_{-i}), s) + \delta \sum_{(y,z)\in Y\times Z} \rho((a_i, \alpha_{-i}), s)[y, z]w_i(y, z)\,. \end{aligned} \tag{4.1}$$

Expression 4.1 says that i) playing profile α in state s, followed by continuation payoffs $w(y, z)$ that depend on the realized public signal and state change, yields expected payoff v, and that ii) given those continuation payoffs, playing α is optimal for all players.

The proof of Theorem 1 relies on the following lemma. Recall that $\bar{B}(V, \epsilon)$ is the closed ϵ-neighborhood of the set V.

Lemma 2. *Suppose that the Identifiability Condition holds. Then for each* $C > 0$*,* $\epsilon > 0$*, and* $d > 0$*, there exists* $\delta^* < 1$ *such that for each state* s*, each* $\delta \geq \delta^*$*, and each* $v^* \in$

$V^{\delta}_{C,2\epsilon,d}(s)$, the set $B(v^*,\epsilon)$ is decomposable with respect to $\left\{\bar{B}\left(V^{\delta}_{C,2\epsilon,d}\left(s'\right),\epsilon\right)\right\}_{s'\in S}$ and δ in state s.

A key feature of Lemma 2 is that for any ϵ, there is a single δ^* that works in every state. That uniformity allows us to cover the infinite state space. Using Lemma 2, we can complete the proof of Theorem 1.

Proof of Theorem 1. Lemma 2 shows that the collection of payoff sets $\left\{\bar{B}\left(V^{\delta}_{C,2\epsilon,d}(s),\epsilon\right)\right\}_{s\in S}$ is "self-decomposable" for high enough δ, in the sense that each $\bar{B}\left(V^{\delta}_{C,2\epsilon,d}(s),\epsilon\right)$ is decomposable with respect to δ and the collection $\left\{\bar{B}\left(V^{\delta}_{C,2\epsilon,d}\left(s'\right),\epsilon\right)\right\}_{s'\in S}$. Lemma 1 shows that each $V^{\delta}_{C,2\epsilon,d}(s)$ is compact and convex, so each $\bar{B}\left(V^{\delta}_{C,2\epsilon,d}(s),\epsilon\right)$ is as well. An argument analogous to the second paragraph of FLM's proof of Lemma 4.2 establishes the result. □

4.2. **Approximating $E^{\delta}(s)$.** Theorem 1 establishes that for any $C,\epsilon,d>0$, every payoff in $V^{\delta}_{C,\epsilon,d}(s_0)$ can be achieved in PPE if δ is large. Combining that result with Remark 1, we see that for large δ,

$$V^{\delta}_{C,\epsilon,d}(s_0)\subseteq E^{\delta}(s_0)\subseteq V^{\delta}_{0,0,0}(s_0).$$

(Recall that s_0 is the initial state.) That is, we have both an upper bound and a lower bound on the equilibrium set. Next, we show that as C, ϵ, and d shrink to zero, $V^{\delta}_{C,\epsilon,d}(s_0)$ converges to $V^{\delta}_{0,0,0}(s_0)$ for generic games. We define *generic* as follows:

Given the sets of players, actions, state changes, and public signals, the period length, and the initial state, a game G can be identified as $\Phi=(\rho,g)$, a pair of functions specifying, for each state and action profile, firstly the distribution over state changes and public signals, and secondly the stage-game payoffs. Let $\mathcal{G}$ be the space of games. We measure the distance between two games $G,G'\subseteq\mathcal{G}$ using the supremum norm:

$$D(G,G')=\sup_{a,s,y,z,i}\max\left(\left|\rho(a,s)[y,z]-\rho'(a,s)[y,z]\right|,\left|g_i(a,s)-g'_i(a,s)\right|\right).$$

Let $\mathcal{G}_0\subseteq\mathcal{G}$ be the class of games that satisfy the following uniform version of the interiority condition: there exists $\epsilon>0$ such that for each state s, $B(v,\epsilon)\subseteq V^{\delta}_{0,0,0}(s)$ for some $v\in\mathbb{R}^N$. We say that a claim holds *for generic games* $G\in\mathcal{G}_0$ if there exists a subset $\mathcal{G}' \subseteq \mathcal{G}_0$ such that i) the claim holds for each game $G \in \mathcal{G}'$, and ii) $\mathcal{G}_0 \setminus \mathcal{G}'$ is of the first category with respect to $\mathcal{G}_0$.[2] (Observe that failure of the uniform interiority condition is not generic in $\mathcal{G}$: any game G such that the transition rate is zero for all actions in all states, and such that the stage game in each state has a feasible and individually rational payoff set with full dimension, satisfies the condition, as does an open set around G.)

The following theorem describes the limiting behavior of $V^{\delta}_{C,\epsilon,d}(s_0)$.

Theorem 2. *If the Identifiability Condition holds, then for generic $G \in \mathcal{G}_0$,*

$$\liminf_{C,\epsilon,d \to 0} \mathrm{cl} V^{\delta}_{C,\epsilon,d}(s_0) = V^{\delta}_{0,0,0}(s_0).$$

Note that Theorems 1 and 2 together do not quite fully characterize equilibrium payoffs in the limit as the period length shrinks to zero. The gap arises because Theorem 1 shows that for any fixed $C, \epsilon, d > 0$, the equilibrium set contains $V^{\delta}_{C,\epsilon,d}(s_0)$ for δ close enough to 1, while Theorem 2 shows that for a *fixed* δ, $V^{\delta}_{C,\epsilon,d}(s_0)$ converges (generically) to $V^{\delta}_{0,0,0}(s_0)$ as C, ϵ, and d shrink to zero. This difference leaves open the possibility that for any fixed $C, \epsilon, d > 0$, $V^{\delta}_{C,\epsilon,d}(s_0)$ may be a strict subset of $V^{\delta}_{0,0,0}(s_0)$ for large δ.

Pęski and Wiseman (forthcoming) give a (generically) complete characterization of limit equilibrium payoffs for the case where transition probabilities (across a countable state space) are proportional to the period length. The technical difficulty of extending that result - roughly, to show that $V^{\delta}_{C,\epsilon,d}$ converges to a limit $V^{1}_{C,\epsilon,d}$ as δ grows, and that $V^{1}_{C,\epsilon,d}$ converges to $V^{1}_{0,0,0}$ as C, ϵ, and d shrink - is that here the set of states accessible from any given initial state varies with δ. Thus, even if the payoff sets in every other state converge (as δ grows or C, ϵ, and d shrink), there is no guarantee that the self-δ, C, ϵ, d-FIR payoff set at state s will converge. That problem persists even if $V^{\delta}_{C,\epsilon,d}(s)$ is upper hemicontinuous in the state s.

What sort of restriction on the parameters of the dynamic game will allow a complete characterization of E^{δ} for high δ is an open question.

[2]A subset X of a topological space Y is of the first category, or meager, if it is the union of countably many nowhere dense subsets of Y.

5. Summary and Discussion

This paper characterizes PPE payoffs of stochastic games with a continuous state space and imperfect public monitoring in the limit when there is little time between periods, when the magnitude of state transitions is proportional to the period length. We provide upper and lower bounds for the equilibrium payoff set, and describe how the gap between those bounds shrinks.

In deriving our results, we define pseudo-instantaneous payoffs. We hope that that tool, which captures the effect of current actions on both today's payoff and future continuation values, will be useful to applied modelers in industrial organization and macroeconomics.

The bounds on the equilibrium set that we obtain (V^{δ} and $V^{\delta}_{C,\epsilon,d}$) are defined implicitly as the maximal fixed points of particular correspondences, rather than explicitly constructed. For that reason, we do not have a clear understanding of the link between the parameters of the dynamic game (stage-game payoffs and transition rates) and the properties of the equilibrium set. Exploring that issue is a topic for future research.

Appendix A. Proof of Remark 1

Proof. Pick any nonzero vector $\lambda \in \mathbb{R}^N$. For each state s, let $v^{\lambda}(s) \in \operatorname{argmax}_{v \in \hat{V}^{\delta}(s)} \lambda \cdot v$, and let $\sigma^{\lambda}(s) \in \Sigma^P$ be a strategy that yields payoff $v^{\lambda}(s)$ from initial state s. Strategy $\sigma^{\lambda}(s)$ induces mappings $w : Y \times S \to [-M, M]^N$ that specify the continuation payoff $w(y, s') \in \hat{V}^{\delta}(s')$ if public signal y is observed in period 1 and the state in period 2 is s'. Strategy $\sigma^{\lambda}(s)$ also specifies the (mixed) action profile α^{λ} to be played in the first period. Then

$$\begin{aligned}
\lambda \cdot v^{\lambda}(s) &= \lambda \cdot (1-\delta)\, g\left(\alpha^{\lambda}, s\right) \\
&\quad + \lambda \cdot \delta \sum_{y \in Y, z \in Z} \rho\left(\alpha^{\lambda}, s\right)[y, z] w(y, s + (1-\delta) z) \\
&\leq \lambda \cdot (1-\delta)\, g\left(\alpha^{\lambda}, s\right) + \lambda \cdot \delta \sum_{z \in Z} \rho_Z\left(\alpha^{\lambda}, s\right)[z] v^{\lambda}(s + (1-\delta) z),
\end{aligned}$$

where the inequality is implied by the definition of $v^\lambda(\cdot)$. It follows that

$$\begin{aligned} \lambda \cdot v^\lambda(s) &\leq \lambda \cdot g\left(\alpha^\lambda, s\right) \\ &\quad +\lambda \cdot \tfrac{\delta}{1-\delta} \sum_{z\in Z} \rho_Z \left(\alpha^\lambda, s\right)[z]\left[v^\lambda(s+(1-\delta)z) - v^\lambda(s+(1-\delta)z)\right] \\ &= \lambda \cdot \psi^\delta\left(\alpha^\lambda, s, v^\lambda\right), \end{aligned}$$

Thus, $\lambda \cdot v^\lambda(s) \leq \max\limits_{a\in A(s)} \lambda \cdot \psi^\delta\left(a, s, v^\lambda\right)$ for all λ, and so we conclude that

$$\hat{V}^\delta(s) \subseteq \text{co}\left\{\psi^\delta(a, s, u) : a \in A(s) \text{ and } \exists v \in \hat{V}^\delta(s) \text{ s.t. } u \in U_{\hat{V}^\delta, s, v}\right\};$$

that is, $\hat{V}^\delta$ is self-δ-feasible.

An analogous argument establishes that E^δ is self-δ-feasible. To see that E^δ is self-δ-individually rational, note that because a PPE strategy must specify, after any deviation, continuation payoffs that are themselves PPE payoffs, player i would have a profitable deviation from a strategy that did not give him a payoff of at least $e_i^\delta\left(s; E^\delta\right)$ starting from state s. □

Appendix B. Proof of Lemma 1

Proof. The first claim follows from the fact that if F and G are any two self-δ, C, ϵ, d-FIR collections, then their union $F \cup G$ is also self-δ, C, ϵ, d-FIR. Since, further, the convex hull of F, $\text{co}F \equiv \{\text{co}F(s)\}_s$, and the closure of F, $\text{cl}F \equiv \{\text{cl}F(s)\}_s$, are self-$\delta, C, \epsilon, d$-FIR collections, the second claim holds.

For the third claim, the existence of V^δ follows from the argument establishing the first claim. To see that $V^\delta \subseteq V^\delta_{0,0,0}$, note that for any s and any $v \in V^\delta(s)$, there exist an action profile α and a function $u \in U_{V^\delta,s,v}$ such that $v = \psi^\delta(\alpha, s, u)$. By the definition of $U_{V^\delta,s,v}$, $u(s) = v$. Thus, $v = (1-\delta)\psi^\delta(\alpha, s, u) + \delta u(s)$. Since, further, $U_{V^\delta,s,v} \subseteq U^0_{V^\delta,s}$, we conclude that v is $\delta, 0, 0, 0$-feasible with respect to V^δ. To see that v is $\delta, 0, 0, 0$-individually rational with respect to V^δ, note that (again because $U_{V^\delta,s,v} \subseteq U^0_{V^\delta,s}$) for each player i and state s, the $\delta, 0, 0$-minmax payoff relative to V^δ, $e_i^{\delta,0,0}\left(s; V^\delta\right)$, is at most $(1-\delta)e_i^\delta\left(s; V^\delta\right) + \delta \min\limits_{u\in V^\delta(s)} u_i(s)$. Because $u(s) = v \in V^\delta(s)$ and $\psi_i^\delta(\alpha, s, u) = v_i \geq e_i^\delta\left(s; V^\delta\right)$, we conclude that

$$v_i = (1-\delta)\psi_i^\delta(\alpha, s, u) + \delta u_i(s) \geq (1-\delta)e_i^\delta\left(s; V^\delta\right) + \delta \min_{u\in V^\delta(s)} u_i(s) \geq e_i^{\delta,0,0}\left(s; V^\delta\right).$$

Next we show that $V^{\delta}_{0,0,0} \subseteq V^{\delta}$. Choose any state s and any unit vector λ. Let $F^{\delta}(s)$ denote the set of payoffs that are δ-feasible with respect to $V^{\delta}_{0,0,0}$ in state s. We will demonstrate that $V^{\delta}_{0,0,0} \subseteq F^{\delta}$. Let $v^{\lambda} \in \operatorname{argmax}_{v \in F^{\delta}(s)} \lambda \cdot v$ and $v^{\lambda}_0 \in \operatorname{argmax}_{v \in V^{\delta}_{0,0,0}(s)} \lambda \cdot v$ denote the extreme points of F^{δ} and $V^{\delta}_{0,0,0}$, respectively, in direction λ. Note that because $v^{\lambda}_0 \in V^{\delta}_{0,0,0}(s)$, there exist an action profile α and a function $u \in U^0_{V^{\delta}_{0,0,0},s}$ such that $v^{\lambda}_0 = (1-\delta)\psi^{\delta}(\alpha, s, u) + \delta u(s)$. Define $\hat{u}(s')$ by $\hat{u}(s) \equiv \psi^{\delta}(\alpha, s, u)$ and $\hat{u}(s') \equiv u(s')$ for $s' \neq s$. Then $\hat{u} \in U_{V^{\delta}_{0,0,0},s,\psi^{\delta}(\alpha,s,u)}$, and so $\psi^{\delta}(\alpha, s, \hat{u}) \in F^{\delta}(s)$. A few steps of algebra yield $v^{\lambda}_0 = (1 - \delta\rho_Z(\alpha, s)[0])\,\psi^{\delta}(\alpha, s, \hat{u}) + \delta\rho_Z(\alpha, s)[0]u(s)$. Then

$$\begin{aligned}\lambda \cdot v^{\lambda}_0 &= (1 - \delta\rho_Z(\alpha, s)[0])\,\lambda \cdot \psi^{\delta}(\alpha, s, \hat{u}) + \delta\rho_Z(\alpha, s)[0]\lambda \cdot u(s) \\ &\leq (1 - \delta\rho_Z(\alpha, s)[0])\,\lambda \cdot v^{\lambda} + \delta\rho_Z(\alpha, s)[0]\lambda \cdot v^{\lambda}_0;\end{aligned}$$

the inequality holds because $\psi^{\delta}(\alpha, s, \hat{u}) \in F^{\delta}(s)$ and $u(s) \in V^{\delta}_{0,0,0}(s)$. It follows that $\lambda \cdot v^{\lambda}_0 \leq \lambda \cdot v^{\lambda}$, so $V^{\delta}_{0,0,0} \subseteq F^{\delta}$. Finally, a similar argument shows that $e^{\delta}_i\left(s; V^{\delta}\right) \leq e^{\delta,0,0}_i\left(s; V^{\delta}\right)$ for each player i and state s, so for any $v \in V^{\delta}_{0,0,0}(s)$, v is δ-individually rational with respect to $V^{\delta}_{0,0,0}(s)$ in state s. □

Appendix C. Proof of Lemma 2

C.1. **Identifiability and enforceability.** To prove Lemma 2, we need a few preliminary results. The first result (Lemma 3 in Pęski and Wiseman (forthcoming)) provides a bound on the size of solutions to a system of linear equations. For any vector $x \in \mathbb{R}^n$, let $\|x\|_{\infty} \equiv \max_i |x_i|$ denote the sup norm. (Recall that $\|x\|$ denotes the Euclidean norm, and notice that $\|x\|_{\infty} \leq \|x\| \leq n\,\|x\|_{\infty}$.) For each matrix A with generic element a_{ij}, let $\|A\|_{\infty} = \max_{ij} |a_{ij}|$.

Lemma 3 (Pęski and Wiseman's (forthcoming) Lemma 3)**.** *Let positive integers $j \leq n$, matrix $A \in \mathcal{M}_{jn}$, and vector $b \in \mathbb{R}^j$ be given. If either*

Case 1. $d_j(A) > 0$, *or*

Case 2. $d_{j-1}(A) > 0$ *and there exists a nonzero vector $a \in \mathbb{R}^j$ such that $a'b = 0$ and $a'A = \mathbf{0}$,*

then there exists $w \in \mathbb{R}^n$ such that $Aw = b$ and

$$\|w\|_{\infty} \leq \frac{1}{d_k(A)}\,\|A\|^n_{\infty}\,\|b\|_{\infty},$$

where $k = j$ in Case 1 and $k = j - 1$ in Case 2.

Let U be the set of unit vectors in $\mathbb{R}^N$: $U \equiv \left\{\lambda \in \mathbb{R}^N | \, \|\lambda\| = 1\right\}$. For each unit vector $\lambda \in U$, let $N(\lambda) = \{i : \lambda_i \neq 0\}$ and $b(\lambda) = \min_{i \in N(\lambda)} |\lambda_i|$. Say that vector $\lambda \in \mathbb{R}^N$ is *regular* if $\# N(\lambda) \geq 2$. Let $m \equiv m_Z \times m_Y$ denote the number of (state change, public signal) pairs. The following result follows directly from Case 1 of Lemma 3.

Corollary 1. *For each $d > 0$, player i, state s, profile α^* such that $d_{m_i}(\Pi_i(\alpha^*, s)) \geq d$, and vector $x \in \mathbb{R}^{m_i}$ such that $\|x\|_\infty \leq M$, there exists $w \in \mathbb{R}^m$ such that $\Pi_i(\alpha^*, s) w = x$ and $\|w\|_\infty \leq d^{-1} M$.*

The next result requires only slightly more work.

Lemma 4. *For each $d > 0$, state s, profile α^* such that $d_{m_i + m_j - 1}(\Pi_{ij}(\alpha^*, s)) \geq d$ for all players i and j, each regular unit vector $\lambda \in U$, and each collection of vectors $\{x_i\}_{i=1}^N \in \times_i \mathbb{R}^{m_i}$ such that $\|x_i\|_\infty \leq M$ and $\alpha_i^* \cdot x_i = 0$ for all i, there exists a mapping $w : Y \times Z \to \mathbb{R}^N$ such that for each player i,*

$$\Pi_i(\alpha^*, s) w_i = x_i,$$

and, for each (y, z), $\lambda \cdot w(y, z) = 0$, and $\|w(y, z)\|_\infty \leq d^{-1} N \frac{1}{b(\lambda)} M$.

Proof. Pick any $i, j \in N(\lambda)$ such that $i \neq j$. By definition of the matrices Π, $\alpha_i^* \cdot \Pi_i(\alpha^*, s) = \alpha_j^* \cdot \Pi_j(\alpha^*, s)$. Case 2 of Lemma 3 then implies that there exists $w^{i,j} \in \mathbb{R}^m$ such that

$$\Pi_{ij}(\alpha^*, s) w^{i,j} = \begin{bmatrix} \frac{1}{\# N(\lambda) - 1} x_i \\ -\frac{\lambda_j}{\lambda_i} x_j \end{bmatrix}$$

and $\|w^{i,j}\|_\infty \leq \frac{1}{b(\lambda)} d^{-1} M$. Using Case 1 of Lemma 3, for each player $i \notin N(\lambda)$, there exists $w^i \in \mathbb{R}^m$ such that for all actions a_i,

$$\Pi_i(\alpha^*, s) w^i = x_i$$

and $\|w^i\|_\infty \le d^{-1}M$. Fix $i \in N(\lambda)$ and define $w(y,z)$ as

$$\begin{aligned} w_i(y,z) &= \sum_{j\neq i} w^{ij}(y,z), \\ w_j(y,z) &= -\frac{\lambda_i}{\lambda_j} w^{ij}(y,z), \text{ for each } j \in N(\lambda)\setminus\{i\} \\ w_j(y,z) &= w^i(y,z) \text{ for each } j \notin N(\lambda). \end{aligned}$$

The result follows. □

C.2. **Proof of Lemma 2.** Lemma 2 follows easily from the following lemma.

Lemma 5. *Suppose that the Identifiability Condition holds, and let $C > 0$, $\epsilon > 0$, and $d > 0$ be given. Then for each each unit vector $\lambda^* \in U$, there exists $\delta_{\lambda^*} < 1$ and $\eta_{\lambda^*} > 0$ such that for each $\delta \ge \delta_{\lambda^*}$, each state s, each payoff vector $v^* \in V^\delta_{C,2\epsilon,d}(s)$, and each unit vector $\lambda \in U \cap B(\lambda^*, \eta_{\lambda^*})$, there exist payoff vector $v \in \mathbb{R}^N$, profile α, and a continuation payoff function $w : Y \times Z \to [-M,M]^N$ satisfying $w(y,z) \in \bar{B}\left(V^\delta_{C,2\epsilon,d}(s + (1-\delta)z), \epsilon\right)$ such that*

(1) (4.1) holds for each player i, and

(2) $\lambda \cdot v \ge \lambda \cdot v^* + \epsilon$.

Proof. We consider three cases separately:

Case 1: Regular $\lambda^* \in U$. We first observe that the definition of collection $V^\delta_{C,2\epsilon,d}(s)$ implies that for each state s and $v^* \in V^\delta_{C,2\epsilon,d}(s)$, we can find continuation payoffs $u^* \in \times_{s'} V^\delta_{C,2\epsilon,d}(s')$ satisfying $\|u^*(s) - u^*(s')\| \le \frac{1}{C}\|s - s'\|$, and profile α^* satisfying $d_{m_i+m_j-1}(\Pi_{ij}(\alpha^*, s)) \ge d$ for all i, j, such that

$$(1-\delta)g(\alpha^*, s) + \delta \sum_z \rho_Z(\alpha^*, s)[z] u^*(s + (1-\delta)z) = v^* + 2(1-\delta)\epsilon\lambda^*.$$

Let $v \equiv v^* + \epsilon\lambda^*$. Now let

$$x_i(a_i) \equiv$$
$$\frac{1}{\delta}\left\{ \begin{array}{rl} & (1-\delta)g_i(\alpha^*, s) + \delta \sum_z \rho_Z(\alpha^*, s)[z] u_i^*(s + (1-\delta)z) \\ - & \left[(1-\delta)g_i\left((a_i, \alpha^*_{-i}), s\right) + \delta \sum_z \rho_Z\left((a_i, \alpha^*_{-i}), s\right)[z] u_i^*(s + (1-\delta)z)\right] \end{array} \right\}.$$

Note that $\|x_i\| \leq \frac{1-\delta}{\delta} 2(M + \frac{1}{C} z_{max})$ (because $\|u^*(s) - u^*(s')\| \leq 2\frac{1}{C} \|s - s'\|$), and $\alpha_i^* \cdot x_i = 0$.[3] Let $\Delta \equiv 2(M + \frac{1}{C} z_{max}) \frac{N}{b(\lambda^*)}$. By Lemma 4, then, there exists a mapping $\hat{w} : Y \times Z \to \mathbb{R}^N$ such that $\lambda^* \cdot \hat{w}(y,z) = 0$ for each (y,z), $\|\hat{w}(y,z)\|_\infty \leq \frac{1-\delta}{\delta} \frac{\Delta}{d}$, and for each player i,

$$\Pi_i (\alpha^*, s) \hat{w}_i = x_i.$$

[4]

Define $\hat{\epsilon} \equiv \frac{1}{\delta} \epsilon \{1 - 2(1-\delta)\}$, and let

$$w(y,z) \quad = \quad u^*(s + (1-\delta)z) + \hat{\epsilon}\lambda^* + \hat{w}(y,z).$$

Then simple computations show that for any player i and any action $a_i \in A_i$,

$$\begin{aligned}
& (1-\delta) g_i\left((a_i, \alpha_{-i}^*), s\right) + \delta \textstyle\sum_{(y,z) \in Y \times Z} \rho((a_i, \alpha_{-i}^*), s)[y,z] w_i(y,z) \\
= \; & (1-\delta) g_i\left((a_i, \alpha_{-i}^*), s\right) + \delta \textstyle\sum_z \rho_Z\left((a_i, \alpha_{-i}^*), s\right)[z] u_i^*(s + (1-\delta)z)] + \delta\hat{\epsilon}\lambda_i^* \\
& + \left\{ \begin{array}{rl} & (1-\delta) g_i(\alpha^*, s) + \delta \sum_z \rho_Z(\alpha^*, s)[z] u_i^*(s + (1-\delta)z) \\ - & [(1-\delta) g_i\left((a_i, \alpha_{-i}^*), s\right) + \delta \sum_z \rho_Z\left((a_i, \alpha_{-i}^*), s\right)[z] u_i^*(s + (1-\delta)z)] \end{array} \right\} \\
= \; & (1-\delta) g_i(\alpha^*, s) + \delta \textstyle\sum_z \rho_Z(\alpha^*, s)[z] u_i^*(s + (1-\delta)z) + \delta\hat{\epsilon}\lambda_i^* \\
= \; & v_i^* + 2(1-\delta)\epsilon\lambda_i^* + \epsilon\{1 - 2(1-\delta)\}\lambda_i^* \\
= \; & v_i^* + \epsilon\lambda_i^* \\
= \; & v_i.
\end{aligned}$$

That is, (4.1) holds with equality for all players and all actions. It remains to show that $w(y,z) \in \bar{B}\left(V_{C,2\epsilon,d}^\delta(s + (1-\delta)z), \epsilon\right)$ for each (y,z). Observe that

$$\epsilon - \hat{\epsilon} = \epsilon \left[1 - \frac{1}{\delta}\{1 - 2(1-\delta)\}\right] = \epsilon \frac{1-\delta}{\delta},$$

so the distance from $u^*(s + (1-\delta)z) + \hat{\epsilon}\lambda^*$ to $u^*(s + (1-\delta)z) + \epsilon\lambda^*$ (on the boundary of $B(u^*(s + (1-\delta)z), \epsilon))$ is $\frac{1-\delta}{\delta}\epsilon$. As a consequence, the distance from $u^*(s + (1-\delta)z) + \hat{\epsilon}\lambda^*$ to the boundary of $B(u^*(s + (1-\delta)z), \epsilon)$ along the hyperplane $\lambda^* \cdot u = \lambda^* \cdot u^*(s + (1-\delta)z)$ is $\sqrt{\frac{1-\delta}{\delta}\epsilon}$. Since $\lambda^* \cdot \hat{w}(y,z) = 0$ and $\|\hat{w}(y)\|_\infty \leq \frac{1-\delta}{\delta}\frac{\Delta}{d}$, if δ is close to 1, then $\frac{1-\delta}{\delta}\frac{\Delta}{d} < \sqrt{\frac{1-\delta}{\delta}\epsilon}$, and $w(y,z)$ ($\equiv u^*(s + (1-\delta)z) + \hat{\epsilon}\lambda^* + \hat{w}(y,z)$) lies in $B(u^*(s + (1-\delta)z), \epsilon) \subseteq \bar{B}\left(V_{C,2\epsilon,d}^\delta(s + (1-\delta)z), \epsilon\right)$. Because $b(\lambda)$ is continuous,

[3]This is the step where we require that $C > 0$.

[4]This is the step where we require that $d > 0$.

there is a $\delta_{\lambda^*} < 1$ that works for all unit vectors λ close enough (within some $\eta_{\lambda^*} > 0$) to λ^*.

Case 2: $\lambda_i^* = -1$ for some i, and $\lambda_j^* = 0$ for all $j \neq i$. The definition of collection $V_{C,2\epsilon,d}^{\delta}(s)$ implies that for each state s and $v^* \in V_{C,2\epsilon,d}^{\delta}(s)$, we can find continuation payoffs $\underline{u} \in \times_{s'} V_{C,2\epsilon,d}^{\delta}(s')$ satisfying $\|\underline{u}(s) - \underline{u}(s')\| \leq \frac{1}{C}\|s - s'\|$, and profile $\underline{\alpha}_{-i}$, such that for each a_i, $d_{m_i}(\Pi_i((a_i, \underline{\alpha}_{-i}), s)) > d$ for all $j \neq i$ and

$$(1-\delta)g_i((a_i, \underline{\alpha}_{-i}), s) + \delta \sum_z \rho_Z((a_i, \underline{\alpha}_{-i}), s)[z]\underline{u}_i(s + (1-\delta)z) \leq v_i^* - 2(1-\delta)\epsilon.$$

Let $\underline{a}_i$ be an action that maximizes

$$(1-\delta)g_i((a_i, \underline{\alpha}_{-i}), s) + \delta \sum_z \rho_Z((a_i, \underline{\alpha}_{-i}), s)[z]\underline{u}_i(s + (1-\delta)z),$$

let $\underline{\alpha} \equiv (\underline{a}_i, \underline{\alpha}_{-i})$, and let

$$\underline{v} \equiv (1-\delta)g(\underline{\alpha}, s) + \delta \sum_z \rho_Z(\underline{\alpha}, s)[z]\underline{u}(s + (1-\delta)z).$$

Define the vector v as $v_i \equiv v_i^* - \epsilon$, and $v_j \equiv \underline{v}_j$ for all $j \neq i$. Now, for each $j \neq i$, let

$$x_j(a_j) \equiv$$

$$\frac{1}{\delta}\left\{\begin{array}{ll} & \underline{v}_j \\ - & \left[(1-\delta)g_j\left((a_j, \underline{\alpha}_{-j}), s\right) + \delta \sum_z \rho_Z\left((a_j, \underline{\alpha}_{-j}), s\right)[z]\underline{u}_j(s + (1-\delta)z)\right] \end{array}\right\}.$$

Note that $\|x_j\| \leq \frac{1-\delta}{\delta} 2(M + \frac{1}{C}z_{max})$ (because $\|\underline{u}(s) - \underline{u}(s')\| \leq \frac{1}{C}\|s - s'\|$), and $\alpha_j^* \cdot x_j = 0$. Recall that $\Delta \equiv 2(M + \frac{1}{C}z_{max})N\frac{1}{b(\lambda^*)}$. By Corollary 1, then, there exists a mapping $\hat{w} : Y \times Z \to \mathbb{R}^N$ such that for each (y, z), $\|\hat{w}(y, dz)\|_\infty \leq \frac{1-\delta}{\delta}\frac{\Delta}{d}$, $\hat{w}_i(y, z) = 0$, and for each player $j \neq i$,

$$\Pi_j(\underline{\alpha}, s)\hat{w}_j = x_j.$$

Then let

$$w_j(y, z) \quad = \quad \underline{u}_j(s + (1-\delta)z) + \hat{w}(y, z)$$

for $j \neq i$. It is straightforward to see that for any player $j \neq i$ and any action $a_j \in A_j$,

$$(1-\delta)g_j\left((a_j, \underline{\alpha}_{-j}), s\right) + \delta \sum_{(y,z) \in Y \times Z} \rho((a_j, \underline{\alpha}_{-j}), s)[y, z]w_j(y, z) = \underline{v}_j = v_j.$$

That is, (4.1) holds with equality for all actions for all players $j \neq i$. For player i, define $\hat{\epsilon} \equiv \frac{1}{\delta}\{\epsilon - (v_i^* - \underline{v}_i)\}$, and let

$$w_i(y,z) = \underline{u}_i(s + (1-\delta)z) - \hat{\epsilon}.$$

Then $\underline{a}_i$ is a best response for player i, and

$$\begin{aligned}
& (1-\delta)g_i(\underline{\alpha}, s) + \delta \textstyle\sum_{(y,z)\in Y\times Z} \rho(\underline{\alpha}, s)[y,z] w_j(y,z) \\
= \; & (1-\delta)g_i(\underline{\alpha}, s) + \delta \textstyle\sum_z \rho_Z(\underline{\alpha}, s)[z] \underline{u}_i(s + (1-\delta)z) - \delta\hat{\epsilon} \\
= \; & \underline{v}_i - \{\epsilon - (v_i^* - \underline{v}_i)\} \\
= \; & v_i^* - \epsilon \\
= \; & v_i.
\end{aligned}$$

Thus, (4.1) holds for all players. It remains to show that $w(y,z) \in \bar{B}\left(V^\delta_{C,2\epsilon,d}(s + (1-\delta)z), \epsilon\right)$ for each (y,z). Observe that

$$\begin{aligned}
\epsilon - \hat{\epsilon} &= \epsilon - \tfrac{1}{\delta}\{\epsilon - (v_i^* - \underline{v}_i)\} \\
&\geq \epsilon\tfrac{1}{\delta}[\delta - 1] + \tfrac{1}{\delta}2(1-\delta)\epsilon \\
&= \epsilon\tfrac{1-\delta}{\delta}
\end{aligned}$$

(recall that $\underline{v}_i \leq v_i^* - 2(1-\delta)\epsilon$), so the distance from $\underline{u}_i(s + (1-\delta)z) - \hat{\epsilon}$ to $\underline{u}_i(s + (1-\delta)z) - \epsilon$ (on the boundary of $B(\underline{u}_i(s + (1-\delta)z), \epsilon)$) is at least $\frac{1-\delta}{\delta}\epsilon$. The rest of the proof is the same as in the previous case.

Case 3: $\lambda_i = 1$ for some i, and $\lambda_j = 0$ for all $j \neq i$. The proof of this case is analogous to that of Case 2. □

Now we can complete the proof of Lemma 2.

Proof of Lemma 2. Lemma 5 associates a $\delta_{\lambda^*} < 1$ and $\eta_{\lambda^*} > 0$ with each unit vector $\lambda^* \in U$. Because the set of unit vectors U is compact, there is a finite collection $\hat{U}$ of λ^*'s such that $U = \cup_{\hat{U}} B(\lambda^*, \eta_{\lambda^*})$. Setting $\delta^* \equiv \max_{\hat{U}} \delta_{\lambda^*}$, we conclude that for each $\delta \geq \delta^*$, each state s, each payoff vector $v^* \in V^\delta_{C,2\epsilon,d}(s)$, and each unit vector $\lambda^* \in U$, there exist payoff vector $v \in \mathbb{R}^N$, profile α, and a continuation payoff function $w : Y \times Z \to [-M, M]^N$ satisfying $w(y,z) \in \bar{B}\left(V^\delta_{C,2\epsilon,d}(s + (1-\delta)z), \epsilon\right)$ such that (4.1) holds for each player i, and $\lambda \cdot v \geq \lambda \cdot v^* + \epsilon$.

Thus, when δ is high enough, then for each state s and each payoff vector $v^* \in V^{\delta}_{C,2\epsilon,d}(s)$, $B(v^*,\epsilon)$ is contained in the convex hull of the set of payoffs decomposable with respect to $\left\{\bar{B}\left(V^{\delta}_{C,2\epsilon,d}\left(s'\right),\epsilon\right)\right\}_{s'\in S}$ and δ in state s. The availability of a public randomization device implies that that set is convex, and so equal to its convex hull. □

Appendix D. Proof of Theorem 2

First, it is useful to define

$$V^{\delta}_{C,0+,d}(s) \equiv \bigcup_{\epsilon>0} V^{\delta}_{C,\epsilon,d}(s)$$

for each state s, $C \geq 0$, $d \geq 0$, and $\delta < 1$. Then Theorem 2 follows from Lemma 6:

Lemma 6. For each $\epsilon_0 \geq 0$, $C_0 \geq 0$, and $d_0 \geq 0$,

1. *for each game G, $\lim_{C\to 0} V^{\delta}_{C,\epsilon_0,d_0}(s_0) = V^{\delta}_{0,\epsilon_0,d_0}(s_0)$;*
2. *for each game G, if the Identifiability Condition holds, then $\liminf_{d\to 0} V^{\delta}_{C_0,0+,d_0}(s_0) = V^{\delta}_{C_0,0+,0}(s_0)$; and*
3. *for generic G, $\mathrm{cl}V^{\delta}_{C_0,0+,d_0}(s_0) = V^{\delta}_{C_0,0,d_0}(s_0)$.*

Proof of part 1. Let $z_{min} = \min_{z\in Z}\{\|z\| : z \in Z, z \neq 0\}$ denote the magnitude of the smallest possible nonzero state change. Given δ, then, if the state changes from one period to the next, the new state must be at least $(1-\delta)z_{min}$ from the old one. Also, recall that the distance between any two feasible payoffs is no greater than $2M$. Therefore, for $C \leq \frac{(1-\delta)z_{min}}{2M}$, the restriction that $\|u(s)-u(s')\| \leq \frac{1}{C}\|s-s'\|$ is always satisfied, and so $U^C_{F,s} = U^0_{F,s}$. Thus, $V^{\delta}_{C,\epsilon_0,d_0}(s_0) = V^{\delta}_{0,\epsilon_0,d_0}(s_0)$ for all $C \leq \frac{(1-\delta)z_{min}}{2M}$. □

To prove part 2, here is a preliminary result (Lemma 4 in Pęski and Wiseman (forthcoming)) that provides a lower bound on the local "variability" of non-zero polynomials. Let $m \equiv \#A$ denote the number of action profiles. For each positive integer n, let $\mathcal{F}_{n,m}$ be the space of polynomial functions $f : \mathbb{R}^m \to \mathbb{R}$ with m variables and of order not higher than n. We consider restrictions of such polynomials to the simplex Δ_m of probability distributions α over action profiles. For each $c \in (0,1)$, let $\mathcal{F}^*_{n,m}(c) \subseteq \mathcal{F}_{n,m}$ be the subspace of polynomials f such that $\sup_{\alpha\in\Delta_m}|f(\alpha)| \in [c,1]$.

Lemma 7 (Pęski and Wiseman's (forthcoming) Lemma 4)**.** *For each n, $c > 0$, and $\eta > 0$, there exists a constant $\bar{c} > 0$ such that for each polynomial $f \in \mathcal{F}^*_{n,m}(c)$ and each profile $\alpha \in \Delta_m$, there exists a profile $\alpha' \in \Delta_m$ such that $\|\alpha - \alpha'\| \leq \eta$ and $|f(\alpha')| \geq \bar{c}$.*

Part 2 of Lemma 6 follows from the following lemma:

Lemma 8. *If the Identifiability Condition holds, then for each $\eta > 0$, there exists $d_\eta > 0$ such that for each state s, each action profile α, and each player i,*

(1) there exists an action profile α'_{-i} for players other than i such that $\left\|\alpha_{-i} - \alpha'_{-i}\right\| \leq \frac{\eta}{M}$ and for each action $a_i \in A_i(s)$ and each player $j \neq i$, $d_{m_j}\left(\Pi_j\left(a_i, \alpha'_{-i}\right)\right) \geq d_\eta$.

(2) there exists an action profile α' such that $\|\alpha - \alpha'\| \leq \frac{\eta}{M}$ and $d_{m_i+m_j-1}\left(\Pi_{ij}\left(\alpha'\right)\right) \geq d_\eta$.

Proof. The existence of a $d'_\eta > 0$ satisfying the two conditions follows from the fact that the determinant of any of the relevant matrices is a non-zero polynomial in the mixed strategies of the players, Lemma 7, and the proofs of Lemmas 6.2 and 6.3 from FLM. □

Part 3 of Lemma 6 follows from the following lemma:

Lemma 9. *For any integer $k > 0$, the set*

$$\underline{\mathcal{G}}_0^k \equiv \left\{G \in \mathcal{G}_0 : \left\|V^\delta_{C_0,0,d_0}(s_0; G) \setminus V^\delta_{C_0,0+,d_0}(s_0; G)\right\| \geq \frac{1}{k}\right\}$$

is nowhere dense in $\mathcal{G}_0$.

The proof of Lemma 9, in turn, relies on Lemma 10, which shows that self-δ, C, ϵ, d-FIR payoffs for a game $G \in \mathcal{G}_0$ are self-$\delta, C, 0, d$-FIR for some nearby game.

Lemma 10. *For any game $G \in \mathcal{G}_0$ and any open neighborhood $U \subseteq \mathcal{G}_0$ of G, there exists $G' \in U$ and $\epsilon > 0$ such that for each state s, $V^\delta_{C_0,0,d_0}(s; G) \subseteq V^\delta_{C_0,\epsilon,d_0}\left(s; G'\right)$.*

Proof. By the definition of $\mathcal{G}_0$, we can find a function $v : S \to \mathbb{R}^N$ mapping states into payoff vectors and a scalar $\epsilon' > 0$ such that for each state s, $B(v(s), \epsilon') \subseteq$

$\text{int}V^{\delta}_{C_0,0,d_0}\left(s;G\right)$. For each$\eta \geq 1$, define the game $G^{\eta;v}$=$(\rho^{\eta;v}_Z, g^{\eta;v})$, where

$$\begin{aligned} g^{\eta;v}\left(a,s\right) &= \eta g\left(a,s\right) - \left(\eta-1\right)v\left(s\right) \\ \rho^{\eta;v}_Z\left(a,s\right)[z] &= \eta\rho_Z\left(a,s\right)[z]. \end{aligned}$$

Notice that $G^{1;v} = G$. We choose the parametrization so that as η increases, pseudo-instantaneous payoffs $\psi^{\delta}\left(a,s,u;G^{\eta;v}\right)$ expand radially from $v\left(s\right)$ relative to payoffs $\psi^{\delta}\left(a,s,u;G\right)$: for each action profile a, each state s, and all continuation payoffs $u: S \to \mathbb{R}^N$,

$$\psi^{\delta}\left(a,s,u;G^{\eta;v}\right) - v(s) = \eta\left(\psi^{\delta}\left(a,s,u;G\right) - v(s)\right). \tag{D.1}$$

To see that Expression D.1 holds, note that

$$\begin{aligned} & \psi^{\delta}\left(a,s,u;G^{\eta;v}\right) - v(s) \\ = \; & g^{\eta;d}\left(a,s\right) + \delta\textstyle\sum_{z\in Z}\rho^{\eta;v}_Z\left(a,s\right)[z]\frac{u(s+(1-\delta)z)-u(s)}{1-\delta} - v\left(s\right) \\ = \; & \eta g\left(a,s\right) - \left(\eta-1\right)v\left(s\right) + \delta\textstyle\sum_{z\in Z}\eta\rho_Z\left(a,s\right)[z]\frac{u(s+(1-\delta)z)-u(s)}{1-\delta} - v\left(s\right) \\ = \; & \eta\left(g\left(a,s\right) + \delta\textstyle\sum_{z\in Z}\rho_Z\left(a,s\right)[z]\frac{u(s+(1-\delta)z)-u(s)}{1-\delta} - v\left(s\right)\right) \\ = \; & \eta\psi^{\delta}\left(a,s,u;G\right) - v(s). \end{aligned}$$

To complete the proof, we claim that for any $\eta \geq 1$, there exists $\epsilon > 0$ such that for each state s, $V^{\delta}_{C_0,0,d_0}\left(s;G\right) \subseteq V^{\delta}_{C_0,\epsilon,d_0}\left(s;G^{\eta,\nu}\right)$. Let $\epsilon = \left(\eta-1\right)\epsilon'$. We show that collection $V^{\delta}_{C_0,0,d_0}\left(.;G\right)$ is self-$\delta, C_0, \epsilon, d_0$-individually rational in game $G^{\eta;\nu}$. Using

Expression D.1, we get

$$
\begin{aligned}
& e_i^{\delta,C_0,d_0}\left(s; V^{\delta}_{C_0,0,d_0}; G^{\eta;v}\right) \\
= & \inf_{\alpha_{-i}\in A^i(s,d_0),u\in U^{C_0}_{V^{\delta}_{C_0,0,d_0},s}} \left\{\max_{a_i\in A_i}(1-\delta)\psi_i^{\delta}\left(\left(a_i,\alpha_{-i}\right),s,u;G^{\eta;v}\right)+\delta u_i(s)\right\} \\
= & (1-\delta)v_i(s) \\
& + \inf_{\alpha_{-i}\in A^i(s,d_0),u\in U^{C_0}_{V^{\delta}_{C_0,0,d_0},s}} \left\{\max_{a_i\in A_i}(1-\delta)\left[\psi_i^{\delta}\left(\left(a_i,\alpha_{-i}\right),s,u;G^{\eta;v}\right)-v_i(s)\right]+\delta u_i(s)\right\} \\
= & (1-\delta)v_i(s) \\
& + \inf_{\alpha_{-i}\in A^i(s,d_0),u\in U^{C_0}_{V^{\delta}_{C_0,0,d_0},s}} \left\{\max_{a_i\in A_i}(1-\delta)\eta\left[\psi_i^{\delta}\left(\left(a_i,\alpha_{-i}\right),s,u;G\right)-v_i(s)\right]+\delta u_i(s)\right\} \\
= & \inf_{\alpha_{-i}\in A^i(s,d_0),u\in U^{C_0}_{V^{\delta}_{C_0,0,d_0},s}} \left\{\begin{array}{c} \eta\left[\max_{a_i\in A_i}(1-\delta)\psi_i^{\delta}\left(\left(a_i,\alpha_{-i}\right),s,u;G\right)+\delta u_i(s)\right] \\ +(1-\eta)\left[(1-\delta)v_i(s)+\delta u_i(s)\right] \end{array}\right\}
\end{aligned}
\tag{D.2}
$$

Let

$$
(\alpha^*_{-i},u^*)\in \operatorname*{arg\,inf}_{\alpha_{-i}\in A^i(s,d_0),u\in U^{C_0}_{V^{\delta}_{C_0,0,d_0},s}} \left\{\max_{a_i\in A_i}(1-\delta)\psi_i^{\delta}\left(\left(a_i,\alpha_{-i}\right),s,u;G\right)+\delta u_i(s)\right\}.
$$

Then the last line of D.2 is no greater than

$$
\begin{aligned}
& \eta\left[\max_{a_i\in A_i}(1-\delta)\psi_i^{\delta}\left(\left(a_i,\alpha^*_{-i}\right),s,u^*;G\right)+\delta u^*_i(s)\right]+(1-\eta)\left[(1-\delta)v_i(s)+\delta u^*_i(s)\right] \\
= & \eta e_i^{\delta,C_0,d_0}\left(s;V^{\delta}_{C_0,0,d_0};G\right)+(1-\eta)\left[(1-\delta)v_i(s)+\delta u^*_i(s)\right] \\
= & e_i^{\delta,C_0,d_0}\left(s;V^{\delta}_{C_0,0,d_0};G\right)-(\eta-1)\left[(1-\delta)v_i(s)+\delta u^*_i(s)-e_i^{\delta,C_0,d_0}\left(s;V^{\delta}_{C_0,0,d_0};G\right)\right].
\end{aligned}
$$

Because $e_i^{\delta,C_0,d_0}\left(s;V^{\delta}_{C_0,0,d_0};G\right)\le v_i(s)-\epsilon'$ and $e_i^{\delta,C_0,d_0}\left(s;V^{\delta}_{C_0,0,d_0};G\right)\le u^*_i(s)$, we get that

$$
\begin{aligned}
e_i^{\delta,C_0,d_0}\left(s;V^{\delta}_{C_0,0,d_0};G^{\eta;v}\right) &\le e_i^{\delta,C_0,d_0}\left(s;V^{\delta}_{C_0,0,d_0};G\right)-(\eta-1)(1-\delta)\epsilon' \\
&= e_i^{\delta,C_0,d_0}\left(s;V^{\delta}_{C_0,0,d_0};G\right)-(1-\delta)\epsilon.
\end{aligned}
$$

Thus, $V^{\delta}_{C_0,0,d_0}$ is self-δ,C_0,ϵ,d_0-individually rational in game $G^{\eta;\nu}$.

Next, we show that $V^{\delta}_{C_0,0,d_0}\left(.;G\right)$ is self-δ,C_0,ϵ,d_0-feasible in game $G^{\eta;\nu}$. Take any $v_0\in V^{\delta}_{C_0,0,d_0}(s)$ and $v'\in B(v_0,(1-\delta)\epsilon)$. For each unit vector λ,

$$
\lambda\cdot v'\le(1-\delta)\epsilon+\lambda\cdot v_0\le(\eta-1)\left(1-\delta\right)\epsilon'+\lambda\cdot\psi^{\delta}\left(a^{\lambda},s,u^{\lambda};G\right),
$$

where

$$\left(a^{\lambda},u^{\lambda}\right)\in\arg\max_{a\in A(s),u\in U^{C_0}_{V^{\delta}_{C_0,0,d_0},s}}\lambda\cdot\psi^{\delta}\left(a,s,u;G\right).$$

Because $V^{\delta}_{C_0,0,d_0}\left(.;G\right)$ is self-$\delta,C_0,0,d_0$-feasible in game G, and $B\left(v\left(s\right),\epsilon^{'}\right)\subseteq V^{\delta}_{C_0,0,d_0}\left(s\right)$, it must be that $\epsilon^{'}\leq\lambda\cdot\left(\psi^{\delta}\left(a^{\lambda},s,u^{\lambda};G\right)-v\left(s\right)\right)$ and

$$\begin{aligned}\lambda\cdot v'\leq&\left(\eta-1\right)\left(1-\delta\right)\epsilon^{'}+\lambda\cdot\left[(1-\delta)\psi^{\delta}\left(a^{\lambda},s,u^{\lambda};G\right)+\delta u^{\lambda}(s)\right]\\ \leq&\left(\eta-1\right)\left(1-\delta\right)\lambda\cdot\left[\psi^{\delta}\left(a^{\lambda},s,u^{\lambda};G\right)-v\left(s\right)\right]\\ &+(1-\delta)\lambda\cdot\left[\psi^{\delta}\left(a^{\lambda},s,u^{\lambda};G\right)-v(s)\right]+\lambda\cdot\left[(1-\delta)v(s)+\delta u^{\lambda}(s)\right]\\ =&\eta(1-\delta)\lambda\cdot\left[\psi^{\delta}\left(a^{\lambda},s,u^{\lambda};G\right)-v\left(s\right)\right]+\lambda\cdot\left[(1-\delta)v(s)+\delta u^{\lambda}(s)\right]\\ =&(1-\delta)\lambda\cdot\left[\psi^{\delta}\left(a^{\lambda},s,u^{\lambda};G^{\eta;\nu}\right)-v\left(s\right)\right]+\lambda\cdot\left[(1-\delta)v(s)+\delta u^{\lambda}(s)\right]\\ =&\lambda\cdot\left[(1-\delta)\psi^{\delta}\left(a^{\lambda},s,u^{\lambda};G^{\eta;\nu}\right)+\delta u^{\lambda}(s)\right],\end{aligned}$$

where the second equality comes from Expression D.1. The above implies that

$$v'\in\text{co}\left\{(1-\delta)\psi^{\delta}\left(a,s,u;G^{\eta;\nu}\right)+\delta u(s):a\in A(s)\text{ and }u\in U^{C_0}_{V^{\delta}_{C_0,0,d_0},s}\right\},$$

so $V^{\delta}_{C_0,0,d_0}\left(.;G\right)$ is self-δ,C_0,ϵ,d_0-feasible in game $G^{\eta;\nu}$. □

Now we can prove Lemma 9:

Proof of Lemma 9. Pick any k and any open neighborhood $U\subseteq\mathcal{G}_0$. We want to show that $U\cap\underline{\mathcal{G}}^k_0$ is not dense in U. Suppose, to the contrary, that $U\cap\underline{\mathcal{G}}^k_0$ is dense in U, and choose $G_1\subseteq U\cap\underline{\mathcal{G}}^k_0$ and $\eta>0$ such that $B(G_1,\eta)\subseteq U$. By the definition of $\underline{\mathcal{G}}^k_0$, $\left\|V^{\delta}_{C_0,0,d_0}(s_0;G_1)\right\|\geq\frac{1}{k}$.

Next, by Lemma 10, there is a $G_1^{'}\in B(G_1,\frac{\eta}{2})$ such that $V^{\delta}_{C_0,0,d_0}(s_0;G_1)\subseteq V^{\delta}_{C_0,\epsilon_1,d_0}(s_0;G_1^{'})$ for some $\epsilon_1>0$. Because $\psi^{\delta}\left(\cdot,\cdot,\cdot;G\right)$ is continuous in G, there exist $\eta_1\in\left(0,\frac{\eta}{4}\right)$ and $\epsilon_1^{'}\in(0,\epsilon_1)$ such that $V^{\delta}_{C_0,\epsilon_1,d_0}(s_0;G_1^{'})\subseteq V^{\delta}_{C_0,\epsilon_1^{'},d_0}(s_0;G)$ for all $G\in B(G_1^{'},\eta_1)$. Because $U\cap\underline{\mathcal{G}}^k_0$ is dense in U, there is a $G_1^{''}\in\underline{\mathcal{G}}^k_0$ within η_1 of $G_1^{'}$ (and therefore within $\frac{\eta}{2}+\frac{\eta}{4}<\eta$ of G_1). Since $\left\|V^{\delta}_{C_0,\epsilon_1^{'},d_0}(s_0;G_1^{''})\right\|\geq\left\|V^{\delta}_{C_0,0,d_0}(s_0;G_1)\right\|\geq\frac{1}{k}$, the definition of $\underline{\mathcal{G}}^k_0$ implies that $\left\|V^{\delta}_{C_0,0,d_0}(s_0;G_1^{''})\right\|\geq\frac{2}{k}$.

By a similar argument, there exists a $G_2^{''}\in\underline{\mathcal{G}}^k_0$ within $\frac{\eta}{8}+\frac{\eta}{16}$ of $G_1^{''}$ (and therefore within $\frac{\eta}{2}+\frac{\eta}{4}+\frac{\eta}{8}+\frac{\eta}{16}<\eta$ of G_1) such that $\left\|V^{\delta}_{C_0,0,d_0}(s_0;G_2^{''})\right\|\geq\left\|V^{\delta}_{C_0,0,d_0}(s_0;G_1^{''})\right\|+\frac{1}{k}\geq\frac{3}{k}$.

By repeating the argument, for any integer $n > 0$ we can find a G''_n with the property that $\left\| V^{\delta}_{C_0,0,d_0}(s_0; G''_n) \right\| \geq \frac{1+n}{k}$. But this is a contradiction: because stage-game payoffs are bounded in magnitude by M, $\left\| V^{\delta}_{C_0,0,d_0}(s_0; G) \right\|$ cannot exceed $(2M)^N$. □

References

[1] Abreu, D., Pearce, D., and Stacchetti, E. (1986). “Optimal Cartel Monitoring with Imperfect Information,” *Journal of Economic Theory*, **39**, pp. 251-269.

[2] Abreu, D., Pearce, D., and Stacchetti, E. (1990). “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, **58**, pp. 1041-1063.

[3] Dutta, P.K. (1995). “A Folk Theorem for Stochastic Games,” *Journal of Economic Theory*, **66**, pp. 1-32.

[4] Fudenberg, D., Levine, D., and Maskin, E. (1994). “The Folk Theorem with Imperfect Public Information,” *Econometrica*, **62**, pp. 997-1039.

[5] Fudenberg, D. and Yamamoto, Y. (2011). “The Folk Theorem for Irreducible Stochastic Games with Imperfect Public Monitoring,” *Journal of Economic Theory*, **146**, pp. 1664-1683.

[6] Hörner, J., Sugaya, T., Takahashi, S., and Vieille, N. (2011). “Recursive Methods in Discounted Stochastic Games: An Algorithm for $\delta \to 1$ and a Folk Theorem,” *Econometrica*, **79**, pp. 1277-1318.

[7] Mertens, J.-F. and Parthasarathy, T. (1987). “Equilibria for Discounted Stochastic Games,” C.O.R.E. Discussion Paper 8750.

[8] Mertens, J.-F. and Parthasarathy, T. (1991). “Nonzero-sum Stochastic Games,” in T. Raghavan, T. Ferguson, T. Parthasarathy, and O. Vrieze, eds., *Stochastic Games and Related Topics*, Boston: Kluwer Academic Publishers.

[9] Pęski, M. and Wiseman, T. (forthcoming). “A Folk Theorem for Stochastic Games with Infrequent State Changes,” *Theoretical Economics*.

[10] Solan, E. (1998). “Discounted Stochastic Games,” *Mathematics of Operations Research*, **23**, pp. 1010-1021.