

A FOLK THEOREM FOR STOCHASTIC GAMES WITH INFREQUENT STATE CHANGES

MARCIN PESKI, THOMAS WISEMAN

ABSTRACT. Dutta (1995) studies dynamic stochastic games with finite states, and proves a folk theorem that holds as players become very patient (so that players discount vanishingly little both the time until the next period and the expected time until the next state transition). Fudenberg and Yamamoto (2011) and Hörner, Sugaya, Takahashi, and Vieille (2011) extend that analysis to the case of imperfect monitoring. Here, we consider the case where the length of a period shrinks, but players' rate of time discounting remains fixed. In this case, the discounting between periods shrinks to zero in the limit, but the discounting of the expected time until a state transition does not. Our main result is a version of a folk theorem that holds under Fudenberg, Levine, and Maskin's (1994) monitoring conditions. We do not require that the stochastic game be irreducible. We show, though, that even if the game is irreducible, the set of equilibrium payoffs may be smaller than the set of payoffs generated by strategies that give each player at least his minmax payoff after every history. There may exist equilibrium payoffs that cannot be achieved with Markov strategies that are individually rational in every state.

1. INTRODUCTION

Stochastic games are generalizations of repeated games in which the payoffs in a period depend not only on the current action profile but also on the value of a state variable, whose random evolution is itself influenced by players' actions. Stochastic games allow for dynamic interaction between players, but do not impose the strong restriction that the parameters of the interaction in one period are independent of outcomes in previous periods (or, indeed, that the parameters are identical across periods). Important economic examples are models with stock variables such as

Date: February 17, 2012.

Preliminary and Incomplete. Please do not cite.

capital, savings, technology, or brand awareness; models with persistent shocks to demand, productivity, or income; models of durable goods markets; and political economy models where government policy changes at discrete intervals.

Dutta (1995) derives a folk theorem for stochastic games with finite states and perfect monitoring. In particular, he shows that as players become very patient (so that the discounted time until the state changes shrinks to zero), then any feasible vector of payoffs that guarantees each player at least his minmax payoff can be attained (more precisely, it can be approximated arbitrarily closely) in equilibrium. Fudenberg and Yamamoto (2011) (from now on, “FY”) and Hörner, Sugaya, Takahashi, and Vieille (2011) (“HSTV”) provide conditions on imperfect public monitoring (in each period players observe the state and a noisy signal of the action profile just played) under which Dutta’s (1995) folk theorem extends. All three results require that the set of perfect public equilibrium payoffs be independent of the initial state as the discount factor δ approaches 1. That condition will be satisfied if the game is *irreducible* – that is, if no single player’s deviation can prevent the Markov process governing the state variable from being irreducible.

In this paper, we consider an alternative limiting case for stochastic games: the length of a period shrinks, but the players’ rate of time discounting remains fixed. In the limit, the discounting between periods shrinks, but the discounted time until a state transition does not (since we interpret the transition rate as a rate per unit of time, not per period). In many economic settings, where players get frequent opportunities to adjust their actions, but the state changes more rarely, our model may be more natural. For example, competing firms must set prices and choose how much to invest in research each period, but technological breakthroughs that reduce costs occur infrequently. Similarly, regulatory policies may change only when a new government is elected, but firms or individuals interact with each frequently in the meantime.

In our setting, the set of feasible and individually rational payoffs does depend on the initial state, even as the time between periods shrinks, because the discounted fraction of the game that players expected to spend in that state before the first transition is non-negligible. Similarly, minmax payoffs vary with the current state. A

strategy profile can only be part of an equilibrium, then, if it delivers continuation payoffs that are individually rational, given the current state, after each history – the payoffs thus achievable are *ex post individually rational*, in Dutta’s (1995) terminology. In fact, though, even a feasible, ex post individually rational payoff vector may not be achievable in equilibrium. The intuition is roughly as follows: in standard repeated games, the set of available continuation payoffs is independent of today’s action profile. In a stochastic game, on the other hand, today’s action affects tomorrow’s state, and thus the set of available continuation payoffs. We must include this effect on expected continuation payoffs when we calculate minmax payoffs. (For example, a player may prefer a profile that yields a lower payoff today if it leads with higher probability to a state where his minmax payoff is high.) Player i ’s minmax payoff in a given state, then, depends on the values of continuation payoffs available in other states. In a subgame perfect equilibrium, continuation play after a state change must give *all* players at least their minmax payoffs; that requirement may imply that player i ’s continuation payoff is strictly above his minmax payoff. In that case, player 1 may be able to assure himself a payoff above his minmax payoff in the current state.

We illustrate the possibility that a feasible, ex post individually rational payoff may be outside the equilibrium set with an example later in the paper. We conclude that the set of feasible, ex post individually rational payoffs is not the appropriate generalization of the feasible, individually rational set in standard repeated games. Instead, we define for each player and each state the minmax payoff *relative* to a collection F of available continuation payoffs in each state, and say that F is *relatively (ex post) individually rational* if for every state s , each payoff in $F(s)$ exceeds the state- s minmax payoff relative to F for each player. It is straightforward to show that the set of equilibrium payoffs is relatively individually rational. Our main result is a folk theorem. We show that under Fudenberg, Levine, and Maskin’s (1994) conditions on public monitoring, if F is feasible and relatively individually rational, then any payoff vector $v \in F(s)$ can be achieved in a subgame perfect equilibrium starting from initial state s , as long as players are patient enough.

Dutta (1995), FY, and HSTV use irreducibility to ensure that in the limit the set of feasible and individually rational payoffs is independent of the state. Since that property will in general not hold in our model, we do not require irreducibility. We can allow, for example, multiple absorbing states, each reachable from the initial state – the first firm to achieve a technological breakthrough might permanently capture the market, for instance. We do require that the maximal feasible and strictly relatively individual rational set have full dimension. We show that generically that set either has full dimension or is empty (because the vector of minmax payoffs is efficient). Interestingly, that set may include payoffs that cannot be achieved through relatively individually rational Markov strategies.

The organization of the rest of the paper is as follows: in Section 2 we describe the model, and in the following section we present the folk theorem. In Section 4 we investigate the properties of the ex post individually rational sets, and give some sufficient conditions for full dimensionality. Section 5 is the conclusion.

2. MODEL

2.1. Stochastic game. 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 $e^{-r\Delta} \equiv \delta$). There is a finite set S of K payoff-relevant states of the world. In each different state $s \in S$, there is a stage game $G(s)$ with set of action profiles $A(s) = A_1(s) \times A_2(s) \times \cdots \times A_N(s)$, where $A_i(s)$ is the (finite) set of actions for player i . Let $m_i(s)$ denote the number of actions available to player i : $m_i(s) \equiv |A_i(s)|$. At the start of each period t , the state s is publicly observed. Then each player i chooses an action $a_i \in A_i(s)$, and then all players observe a public signal y drawn from a finite set Y , which has m elements. The public signal is distributed according to $\rho(a, s)$, where a is the profile of actions of all players.

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

$$g_i(a, s) = \sum_{y \in Y} [\rho(a, s)(y)] u_i(a_i, y, s).$$

Denote by $g(a, s)$ the vector of expected payoffs for each player. Define M as the greatest length of the vector of expected stage-game payoffs from any action in any state: $M \equiv \max_{a,s} \|g(a, s)\|$.

At the end of a period in which action profile a is played in state s , the probability that the state changes to state $s' \neq s$ is equal to $\Delta \hat{\gamma}(s'; a, s)$: the constant rate $\hat{\gamma}(s'; a, s)$ multiplied by the period length Δ . That is, the transition probability per unit of *time* is fixed, so (for small Δ) the transition probability per *period* is proportional to the length of the period. With the remaining probability $1 - \Delta \hat{\gamma}(a, s)$, where

$$\hat{\gamma}(a, s) \equiv \sum_{s' \neq s} \hat{\gamma}(s'; a, s),$$

the state does not change. Note that when the period length Δ is close to zero, then the discount rate per period δ is approximately equal to $1 - r\Delta$, and so the probability of transition from state s to state s' given action profile a in each period, $\Delta \hat{\gamma}(s'; a, s)$, is approximately equal to $(1 - \delta) \frac{1}{r} \hat{\gamma}(s'; a, s)$. Since we will focus on the limiting case as $\Delta \rightarrow 0$, for notational simplicity we will re-specify the per-period transition probabilities as $(1 - \delta) \gamma(s'; a, s)$, where

$$\gamma(s'; a, s) \equiv \frac{1}{r} \hat{\gamma}(s'; a, s).$$

Similarly, we re-specify the probability that no transition occurs as $1 - (1 - \delta) \gamma(a, s)$, where

$$\gamma(a, s) \equiv \sum_{s' \neq s} \gamma(s'; a, s).$$

We will then consider the limit as $\delta \rightarrow 1$. In particular, we assume throughout that $(1 - \delta) \max_{a,s} \gamma(a, s) < 1$, so that the maximum probability of transition in any period is less than 1.

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 $H_t \equiv Y^{t-1} \times S^t$, with element $h_t = (s_1, y_1, \dots, y_{t-1}, s_t)$, where s_t denotes the state at the beginning of period t and y_t denotes the public signal realized at the end of period t . Player i 's *private history* in

period t is $h_t^i = (s_1, y_1, a_{i,1}, \dots, y_{t-1}, a_{i,t-1}, s_t)$, where $a_{i,t}$ is player i 's action in period t ; $H_t^i \equiv (Y \times A_i)^{t-1} \times S^t$ 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 $\alpha_i : H^i \rightarrow \Delta A_i(s(h_t^i))$, and a public strategy for player i is a mapping $\sigma_i : H \rightarrow \Delta A_i(s(h_t))$. 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 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.

2.2. Pseudo-instantaneous payoffs. Given a state $s \in S$, let $u = (u_{s'})_{s' \neq s} \in R^{N \times (K-1)}$ specify a vector of continuation payoffs as a function of next period's state. For each such u and action profile a , we define a vector of payoffs

$$\psi^\delta(a, s, u) \equiv \frac{1}{1 + \delta \gamma(a, s)} \left[g(a, s) + \delta \sum_{s' \neq s} \gamma(s'; a, s) u_{s'} \right]. \quad (2.1)$$

We refer to $\psi^\delta(a, s, u)$ as the *pseudo-instantaneous payoff* (or just the pseudo-payoff) from playing action profile a in state s , given continuation payoffs u . To motivate the name, we observe that in state s the expected payoff v from playing the same profile a in each period until the state changes is exactly $\psi^\delta(a, s, u)$:

$$\begin{aligned}
 v &= (1 - \delta) g(a, s) + \delta \left(\sum_{s' \neq s} (1 - \delta) \gamma(s'; a, s) u_{s'} \right) + \delta [1 - (1 - \delta)(a, s)] v \\
 &\Rightarrow \\
 v (1 - \delta [1 - (1 - \delta)(a, s)]) &= (1 - \delta) g(a, s) + \delta \left(\sum_{s' \neq s} (1 - \delta) \gamma(s'; a, s) u_{s'} \right) \\
 &\Rightarrow \\
 v (1 - \delta) (1 + \delta \gamma(a, s)) &= (1 - \delta) g(a, s) + \delta \left(\sum_{s' \neq s} (1 - \delta) \gamma(s'; a, s) u_{s'} \right) \\
 &\Rightarrow \\
 v &= \frac{1}{(1 + \delta \gamma(a, s))} \left[g(a, s) + \delta \left(\sum_{s' \neq s} \gamma(s'; a, s) u_{s'} \right) \right] = \psi^\delta(a, s, u).
 \end{aligned}$$

Notice that (2.1) is well-defined for all $\delta \leq 1$. More generally, we can represent the expected payoff from playing profile a for one period, given continuation payoffs u if the state changes and continuation payoff u_s if the state does not change, as a convex combination of the pseudo-instantaneous payoff and u_s :

$$\begin{aligned}
 &(1 - \delta) g(a, s) + \delta \left(\sum_{s' \neq s} (1 - \delta) \gamma(s'; a, s) u_{s'} \right) + \delta [1 - (1 - \delta)(a, s)] u_s \\
 &= (1 - \delta) (1 + \delta \gamma(a, s)) \psi^\delta(a, s, u) + [1 - (1 - \delta) (1 + \delta \gamma(a, s))] u_s \\
 &= \beta^\delta(a, s) \psi^\delta(a, s, u) + [1 - \beta^\delta(a, s)] u_s, \tag{2.2}
 \end{aligned}$$

where $\beta^\delta(a, s) \equiv (1 - \delta) (1 + \delta \gamma(a, s)) \in (0, 1)$.

These definitions will be useful later.

2.3. Feasible and individually rational payoffs. Given discount factor $\delta < 1$, 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, in contrast to the setting in FY and HSTV, the set of feasible payoffs $\hat{V}^\delta(s)$ varies with the state even in the limit as δ approaches 1: the payoffs to the stage game vary with the state, and both the discount rate and the rate of transition between states are fixed per unit of time as δ grows. Further, we emphasize (at the risk of re-stating the obvious) that $\hat{V}^\delta(s)$ is *not* the set of feasible payoffs for the stage game in state s .

For each $\delta < 1$, let $E^\delta(s)$ be the set of payoffs obtained in subgame perfect equilibria of the game, given initial state s and discount factor δ .

The sets $E^\delta(\cdot)$ satisfy a pair of important properties. It is useful to state these properties for general collections of sets. Take any collection $F = \{F(s)\}_s$ of sets $F(s) \subseteq [-M, M]^I$ for each $s \in S$. (Recall that M is the maximal length of any stage-game payoff vector.) Say that collection F is δ -feasible if

$$F(s) \subseteq \text{co} \left\{ \psi^\delta(a, s, u) : a \in A(s), u \in \times_{s' \neq s} F(s') \right\}.$$

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

Second, for each player i , define the δ -minmax payoff relative to F for player i in state s as

$$e_i^\delta(s; F) = \inf_{\alpha_{-i} \in \times_{j \neq i} \Delta A_j(s), u \in \times_{s' \neq s} F(s')} \max_{a_i \in A_i(s)} \psi^\delta(a_i, \alpha_{-i}, s, u).$$

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

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

Proof. Check but should be obvious. □

Note that 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. For that reason, the collection $\left(\left\{ v \in \hat{V}^\delta(s) : v_i \geq e_i^\delta(s; \hat{V}^\delta) \forall i \in N \right\} \right)_s$ need not be δ -feasible: a strategy that, starting from the initial state s , yields a payoff above the state- s minmax vector may not give individually rational continuation payoffs.

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 $\epsilon \geq 0$. Say

that collection F is δ, ϵ -feasible if for each $v \in F(s)$,

$$B(v, \epsilon) \subseteq \text{co} \left\{ \psi^\delta(a, s, u) : a \in A(s), u \in \times_{s' \neq s} F(s') \right\}.$$

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

Lemma 1. *For each $\delta \leq 1$ and $\epsilon \geq 0$, there exists the largest collection V_ϵ^δ such that $V_\epsilon^\delta(s) \subseteq [-M, M]^N$ for each $s \in S$, and V_ϵ^δ is δ, ϵ -feasible and self- δ, ϵ -individually rational.*

Proof. Show that the union of two self-generated and individually rational collections is feasible and individually rational. \square

Lemma 2. *For each $\delta \leq 1$ and $\epsilon \geq 0$, each $V_\epsilon^\delta(s)$ is compact and convex.*

We will refer to elements of $V_\epsilon^\delta(s)$ as δ, ϵ -feasible and individually rational payoff in state s . We will refer to $V_0^1(s)$ as the set of feasible and individually rational payoffs in state s . Whenever we want to emphasize the dependence of set $V_\epsilon^\delta(s)$ on the game G , we will write $V_\epsilon^\delta(s; G)$.

Note that the collection of equilibrium payoffs E^δ is not necessarily δ, ϵ -feasible. The folk theorem in the next section will establish conditions under which V_ϵ^δ is contained in E^δ ; that is, when any δ, ϵ -feasible and individually rational payoff in state s can be achieved in a subgame perfect equilibrium from initial state s . For that result to be interesting, we need to establish that for large δ and small ϵ , V_ϵ^δ is approximately equal to E^δ . We turn to that issue in Section 4.

2.4. Identifiability. The definitions of individual full rank and pairwise full rank, conditions relating to the identifiability of players' actions by public signals, are the same as in FLM. Recall that $m_i(s)$ is the number of actions available to player i in state s , and that m is the number of public signals. For each state s , player i , and (mixed) action profile by the other players α_{-i} , let $\Pi_i(\alpha_{-i}, s)$ be the $m_i(s) \times m$ matrix whose rows correspond to the probability distribution over public signals induced by each of player i 's actions, given s and α_{-i} : $\Pi_i(\alpha_{-i}, s) \equiv \rho((\cdot, \alpha_{-i}), s)$. Similarly, for each state s and action profile α , let $\Pi_{ij}(\alpha)$ be the $(m_i(s) + m_j(s)) \times m$ matrix whose first $m_i(s)$ rows are $\Pi_i(\alpha_{-i}, s)$ and whose last $m_j(s)$ rows are $\Pi_j(\alpha_{-j}, s)$.

Action profile α has *individual full rank* in state s if $\Pi_i(\alpha_{-i}, s)$ has rank $m_1(s)$ for each player i . Action profile α has *pairwise full rank* for players i and j in state s if $\Pi_{ij}(\alpha)$ has rank $m_1(s) + m_j(s) - 1$. 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.

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

Definition 1. *Identifiability Condition:* For each state s , i) every pure action profile has individual full rank in state s , and ii) for all pairs of players i and j , there exists a profile $\alpha(s)$ that has pairwise full rank for i and j in state s .

That is, the identifiability condition requires that FLM's Condition 6.2 and 6.3 hold in each state.¹

3. FOLK THEOREM

FLM's folk theorem requires that the set of feasible and individually rational payoffs has nonempty interior. Here, we require not only that the set of feasible and individually rational payoffs has nonempty interior in each state, but also that there exists a strategy that yields continuation payoffs in that interior after every history. Formally, we require that $V_\epsilon^\delta(s)$ has full dimension for each state s . The role of this condition (which is equivalent to FLM's condition in the case of a single state) is the same as in FLM: 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. (We go into more detail below.)

Definition 2. *Uniform Interior Condition:* There exist scalars $\epsilon_0 > 0$ and $\delta_0 < 1$ such that for all $\epsilon \in [0, \epsilon_0)$ and $\delta \in (\delta_0, 1]$, each $V_\epsilon^\delta(s)$ is nonempty.

¹FLM's folk theorem (Theorem 6.2) requires either the pairwise full rank condition or that every pure-action, Pareto-efficient profile is *pairwise identifiable* for all pairs of players. In our setting, the appropriate analog would be Pareto efficiency in terms of the pseudo-instantaneous payoffs, which are endogenous (since they depend on continuation values after a state change). Thus, we do not focus on that condition.

With that definition, we can state our first result:

Theorem 1. *Suppose that the Uniform Interior Condition and the Identifiability Condition hold. Then for ϵ small enough, there exists $\delta^* < 1$ such that the following holds: for any initial state s , any $\delta \geq \delta^*$, and any $v \in V_\epsilon^\delta(s)$, there exists a perfect public equilibrium strategy profile σ such that $v(\sigma, s) = v$.*

In other words, given the assumptions, for sufficiently high δ , any δ, ϵ -feasible and individually rational payoff at state s can be attained in a subgame perfect equilibrium.

The proof is based on the techniques in the proof of FLM's folk theorem for games with imperfect public monitoring. That 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. The definition of V_ϵ^δ ensures that for each state s , there is a pseudo-payoff outside $V_\epsilon^\delta(s)$ in each direction.

Given a state s , let $V \subseteq \mathbf{R}^N$ be a set of payoffs, and let $W = \{W(s')\}_{s' \neq s}$, where each $W(s') \subseteq \mathbf{R}^N$, 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 , payoffs $u = (u_{s'})_{s' \neq s}$ such that $u_{s'} \in W(s')$ for each $s' \neq s$, and a function $w : Y \rightarrow V$ such that for each player i ,

$$\begin{aligned} v_i &= E_\alpha \left(\beta^\delta(a, s) \psi_i^\delta(a, s, u) + [1 - \beta^\delta(a, s)] \sum_{y \in Y} \rho(\alpha, s) [y] w_i(y) \right) \\ &\geq E_{a_i, \alpha_{-i}} \left(\beta^\delta(a, s) \psi_i^\delta(a, s, u) + [1 - \beta^\delta(a, s)] \sum_{y \in Y} \rho(\alpha, s) [y] w_i(y) \right). \end{aligned} \quad (3.1)$$

Expression 3.1 says that i) playing profile α in state s , followed by continuation payoffs u (if the state changes) or by public-signal-contingent continuation payoffs $w(y)$ (if the state does not change) yields expected payoff v , and that 2) given those continuation payoffs, playing α is optimal for all players. (Note that the continuation payoff after a state change does not depend on the public signal y ; we discuss that point later.)

A subset of \mathbf{R}^N is *smooth* if i) it is closed and convex; ii) it has nonempty interior; iii) at each boundary point v there is a unique tangent hyperplane P_v ; and iv) P_v varies continuously with v . We say that a smooth V is *decomposable on tangent hyperplanes with respect to δ and W in state s* if for each v on the boundary of V , there exist a mixed action profile α , payoffs $u = (u_{s'})_{s' \neq s}$ such that $u_{s'} \in W(s')$ for each $s' \neq s$, a translate P' of P_v , and a function $w : Y \rightarrow P' \cap V$ such that Expression 3.1 is satisfied and $\psi^\delta(a, s, u)$ is separated from V by P_v . In Lemma 4, we will demonstrate that decomposability on tangent hyperplanes implies decomposability.

To establish our folk theorem, we will first show (in Lemma 3) that for each s and each $v_0 \in V_\epsilon^\delta(s)$, the set $B(v_0, \epsilon/2)$ is decomposable on tangent hyperplanes with respect to δ and the collection $\{V_\epsilon^\delta(s')\}_{s' \neq s}$ for high enough δ . That fact, together with Lemma 4, will imply our result.

Lemma 3. *Suppose that the Uniform Interior Condition and the Identifiability Condition hold. Then for ϵ small enough, the following holds: for each s and each $v_0 \in V_\epsilon^\delta(s)$, there exists $\delta^*(v_0) < 1$ such that if $\delta \geq \delta^*(v_0)$, then the set $B(v_0, \epsilon/2)$ is decomposable on tangent hyperplanes with respect to δ and the collection $\{V_\epsilon^\delta(s')\}_{s' \neq s}$.*

Proof. Let payoff vector v on the boundary of $B(v_0, \epsilon/2)$ be given, let P be the hyperplane tangent to $B(v_0, \epsilon/2)$ at v , and let λ be the unit vector orthogonal to P . We will consider three cases separately:

If λ is regular: By the definition of V_ϵ^δ , we can choose payoffs $u = (u_{s'})_{s' \neq s}$ such that $u_{s'} \in V_\epsilon^\delta(s')$ for each $s' \neq s$ and an action profile α^* such that $\lambda \cdot \psi^\delta(\alpha^*, s, u) \geq \lambda \cdot v + \epsilon/2$ (so that $\psi^\delta(\alpha^*, s, u)$ is separated from $B(v_0, \epsilon/2)$ by P). FLM's Lemma 6.2 implies that, since the Identifiability Condition holds, we can additionally require that α^* has pairwise full rank for all pairs of players. We want to find a translate P' of P and a continuation payoff function (for when the state does not change)

$w : Y \rightarrow P' \cap B(v_0, \epsilon/2)$ such that Expression 3.1 holds. It is sufficient if for all i and all $\alpha \equiv (a_i, \alpha_{-i}^*)$,

$$v_i = \beta^\delta(\alpha, s) \psi_i^\delta(\alpha, s, u) + [1 - \beta^\delta(\alpha, s)] \sum_{y \in Y} \rho(\alpha, s) [y] w_i(y). \quad (3.2)$$

We will show that FLM's argument applies directly to a perturbed version of Expression 3.2, and then rely on a result about the continuity of solutions of linear equations (Mangasarian and Shiau, 1987) to argue that a nearby w works for Expression 3.2.

Define the functions

$$A(\alpha, s, u, \delta) \equiv g(\alpha, s) + \delta \sum_{s' \neq s} \gamma(s'; \alpha, s) u_{s'} \text{ and}$$

$$b(\alpha, s, \delta) \equiv 1 - (1 - \delta) \gamma(\alpha, s).$$

Note that $(1 - \delta) A(\alpha, s, u, \delta) = \beta^\delta(\alpha, s) \psi^\delta(\alpha, s, u)$ and $\delta b(\alpha, s, \delta) = 1 - \beta^\delta(\alpha, s)$, so Expression 3.2 can be rewritten as

$$\sum_{y \in Y} \rho(\alpha, s) [y] w_i(y) = \frac{1}{\delta b(\alpha, s, \delta)} [v_i - (1 - \delta) A_i(\alpha, s, u, \delta)] \quad (3.3)$$

for all i and all $\alpha \equiv (a_i, \alpha_{-i}^*)$.

Next, we consider the perturbed problem given by evaluating A and b at $\delta = 1$:

$$\sum_{y \in Y} \rho(\alpha, s) [y] w_i(y) = \frac{1}{\delta} [v_i - (1 - \delta) A_i(\alpha, s, u, 1)] \quad (3.4)$$

for all i and all $\alpha \equiv (a_i, \alpha_{-i}^*)$.

(Note that $b(\alpha, s, 1) = 1$.) Then FLM's Theorem 5.1 and their proof of Theorem 4.1 ensure the existence of a translate hyperplane P' and a function $w : Y \rightarrow P' \cap B(v_0, \epsilon/2)$ that satisfies Expression 3.4. (Given s and u , we can think of $A(\alpha, s, u, 1)$ as the payoff function of a stage game, and interpret Expression 3.4 as resulting from the corresponding repeated game.) We want to show that for δ close to 1, there are nearby P'_δ and $w_\delta : Y \rightarrow P'_\delta \cap B(v_0, \epsilon/2)$ that satisfies Expression 3.3. An argument analogous to FLM's proofs of Lemmas 5.4 and 5.3 establishes there exist a translate P'_δ and a function $w_\delta : Y \rightarrow P'_\delta$ that satisfies Expression 3.3, so we just need to show that the values of w_δ lie within $B(v_0, \epsilon/2)$.

FLM's proof of Theorem 4.1 shows that i) the distance from $E[w(y)|\alpha^*]$ (which lies in the interior of $B(v_0, \epsilon/2)$) to v is of order $(1 - \delta)/\delta$, ii) the distance from $E[w(y)|\alpha^*]$ to each realization $w(y)$ is also of order $(1 - \delta)/\delta$, and iii) the distance from $E[w(y)|\alpha^*]$ to the boundary of $B(v_0, \epsilon/2)$ along P' is of order $\sqrt{(1 - \delta)/\delta}$. (They can thus conclude that each $w(y)$ lies in $B(v_0, \epsilon/2)$ for large δ .) By analogous arguments, we have that i') the distance from $E[w_\delta(y)|\alpha^*]$ to v is of order $(1 - \delta)/\delta$, and iii') the distance from $E[w_\delta(y)|\alpha^*]$ to the boundary of $B(v_0, \epsilon/2)$ along P'_δ is of order $\sqrt{(1 - \delta)/\delta}$.

To show that the distance from $E[w_\delta(y)|\alpha^*]$ to each $w_\delta(y)$ is small enough, we argue that the distance from $w_\delta(y)$ to $w(y)$ is small. Mangasarian and Shiau's (1987) Theorem 2.2² ensures that the distance between w_δ and w (the solutions to Expressions 3.3 and 3.4) is proportional to the distance between the right-hand-sides of Expressions 3.3 and 3.4. (The Lipschitz coefficient depends only on ρ .) The partial derivative of the right-hand-side,

$$\begin{aligned} & \frac{\partial}{\partial \tilde{\delta}} \left(\frac{1}{\delta b(\alpha, s, \tilde{\delta})} \left[v_i - (1 - \delta) A_i(\alpha, s, u, \tilde{\delta}) \right] \right) \\ &= \frac{1 - b(\alpha, s, \tilde{\delta}) (1 - \delta) \sum_{s' \neq s} \gamma(s'; \alpha, s) u_{s'} - \left[v_i - (1 - \delta) A_i(\alpha, s, u, \tilde{\delta}) \right](\alpha, s)}{\delta \left[b(\alpha, s, \tilde{\delta}) \right]^2}, \end{aligned}$$

is continuous in δ and $\tilde{\delta}$, and is finite (equal to $-v_i(\alpha, s)$ at $\tilde{\delta} = \delta = 1$). Therefore, the distance between the right-hand-sides of Expressions 3.3 and 3.4 is bounded by a constant (the maximum value of that partial derivative between δ and 1) times the difference between δ and 1. Thus, the distance from $w_\delta(y)$ to $w(y)$ is of order $1 - \delta$. Since both $w_\delta(y)$ and $w(y)$ lie on translates of P' , the difference between the distance from $E[w_\delta(y)|\alpha^*]$ to $w_\delta(y)$ (along P'_δ) and the distance from $E[w(y)|\alpha^*]$ to $w(y)$ (along P') also is of order no greater than $1 - \delta$. Since (by condition ii above) the distance from $E[w(y)|\alpha^*]$ to $w(y)$ is of order $(1 - \delta)/\delta$, we conclude that the distance

²For the special case of a system of linear equalities (rather than a combination of equalities and inequalities), that theorem states that if the sets of solutions $S^i, i \in \{1, 2\}$ to the systems of equations $Cx = d^i, i \in \{1, 2\}$ are nonempty, then for any $x^1 \in S^1$, there exists $x^2 \in S^2$ such that $\|x^1 - x^2\| \leq K(C) \|d^1 - d^2\|$, where K is a particular, finite-valued function.

from $E[w_\delta(y)|\alpha^*]$ to $w_\delta(y)$ also is of order $(1 - \delta)/\delta$. Therefore, for high enough δ , each $w_\delta(y)$ lies within $B(v_0, \epsilon/2)$ (by condition iii'). PICTURES.

For v' on the boundary of $B(v_0, \epsilon/2)$ that are close to v , we again refer to Mangasarian and Shiau's (1987) Theorem 2.2 to conclude that there exists a continuation payoff function w'_δ close to w_δ that takes values (on a translate of P) that are still in the interior of $B(v_0, \epsilon/2)$, such that w'_δ solves Expression 3.3 for v' . That is, the same δ that is high enough for v is also high enough for nearby boundary points.

If $\lambda_i = 1$ for some i and $\lambda_j = 0$ for each $j \neq i$. (In this case, v maximizes player i 's payoff over $B(v_0, \epsilon/2)$.) Choose payoffs $u = (u_{s'})_{s' \neq s}$ such that $u_{s'} \in V_\epsilon^\delta(s')$ for each $s' \neq s$ and a pure action profile a^* such that

$$(a^*, u) \in \operatorname{argmax}_{a' \in A(s), u' \in \times_{s' \neq s} V_\epsilon^\delta(s')} \beta^\delta(a', s) \psi_i^\delta(a', s, u') + [1 - \beta^\delta(a', s)] v'_i \quad (3.5)$$

for all v'_i within some small neighborhood of v_i . (A solution exists because the optimal u is independent of v'_i , plus the objective function is continuous in v'_i and a is drawn from a finite set.) The definition of V_ϵ^δ ensures that the maximum value of the objective function is strictly greater than v_i (because $\psi_i^\delta(a', s, u') = v_i + \epsilon/2$ for some (a', u')), so it must be the case that $\psi_i^\delta(a^*, s, u) > v_i$. That is, $\psi^\delta(a^*, s, u)$ is separated from $B(v_0, \epsilon/2)$ by P .

The rest of the argument is similar to the previous case. Since a^* is a pure action profile, FLM's Lemma 5.1 and 5.2 and their proof of Theorem 4.1 ensure the existence of a translate P' of P and a continuation payoff function (for when the state does not change) $w : Y \rightarrow P' \cap B(v_0, \epsilon/2)$ such that the following condition holds:

$$\sum_{y \in Y} \rho(\alpha, s) [y] w_j(y) = \frac{1}{\delta} [v_j - (1 - \delta) A_j(\alpha, s, u, 1)] \quad (3.6)$$

for all $j \neq i$ and all $\alpha \equiv (a_j, a_{-j}^*)$. By the same continuity argument as in the previous case, for large enough δ there exist a translate P'_δ and a continuation payoff function $w_\delta : Y \rightarrow P'_\delta \cap B(v_0, \epsilon/2)$ that satisfy the following condition:

$$\sum_{y \in Y} \rho(\alpha, s) [y] w_j(y) = \frac{1}{\delta b(\alpha, s, \delta)} [v_j - (1 - \delta) A_j(\alpha, s, u, \delta)]$$

for all $j \neq i$ and all $\alpha \equiv (a_j, a_{-j}^*)$. That is, playing a^* is optimal for all players $j \neq i$. By Expression 3.5, action a_i^* is a best response for player i as well. As in the previous case, the same δ that is high enough for v is also high enough for nearby boundary points.

If $\lambda_i = -1$ for some i and $\lambda_j = 0$ for each $j \neq i$. (In this case, v minimizes player i 's payoff over $B(v_0, \epsilon/2)$.) Choose payoffs $u = (u_{s'})_{s' \neq s}$ such that $u_{s'} \in V_\epsilon^\delta(s')$ for each $s' \neq s$, an action profile for players $-i$, α_{-i}^* and a pure action for player i , a_i^* , such that

$$(\alpha_{-i}^*, u) \in \underset{\alpha_{-i} \in \times_{j \neq i} \Delta A_j(s), u' \in \times_{s' \neq s} V_\epsilon^\delta(s')}{\operatorname{argmin}} \left\{ \max_{a_i \in A_i(s)} \left\{ \beta^\delta((a_i, \alpha_{-i}), s) \psi_i^\delta((a_i, \alpha_{-i}), s, u') + [1 - \beta^\delta((a_i, \alpha_{-i}), s)] v_i' \right\} \right\}$$

and

$$a_i^* \in \operatorname{argmax}_{a_i \in A_i(s)} \left\{ \beta^\delta((a_i, \alpha_{-i}^*), s) \psi_i^\delta((a_i, \alpha_{-i}^*), s, u) + [1 - \beta^\delta((a_i, \alpha_{-i}^*), s)] v_i' \right\}$$

for all v_i' within some small neighborhood of v_i . (A solution exists for the same reason as in the previous case.) The definition of V_ϵ^δ ensures that this minmax value is strictly less than v_i (because $e_i^\delta(s, V_\epsilon^\delta) \leq v_{0i} - \epsilon < v_i$), so it must be the case that $\psi_i^\delta((a_i, \alpha_{-i}^*), s, u) < v_i$. That is, $\psi_i^\delta((a_i, \alpha_{-i}^*), s, u)$ is separated from $B(v_0, \epsilon/2)$ by P .

Now, FLM's Lemma 5.2 and 6.3 and their proof of Theorem 4.1 ensure the existence of an action profile α' close to (a_i, α_{-i}^*) (so that $\psi_i^\delta(\alpha', s, u)$ is separated from $B(v_0, \epsilon/2)$ by P), a translate P' of P , and a continuation payoff function $w : Y \rightarrow P' \cap B(v_0, \epsilon/2)$ such that i) Expression 3.6 holds, and ii) action a_i' is a best response for player i . The rest of the proof is the same as in the previous case.

Each payoff v on the boundary of $B(v_0, \epsilon/2)$ corresponds to one of those three cases. Thus, we have shown that for each such v , there is a $\delta_v < 1$ such that every point in a neighborhood of v can be decomposed on a tangent hyperplane if $\delta \geq \delta_v$. These open neighborhoods form a cover of the boundary. Since the boundary is compact, there is a finite subcover. Take the maximum $\underline{\delta}$ of the δ_v 's associated with that finite

collection, and we conclude that if $\delta \geq \underline{\delta}$, then the set $B(v_0, \epsilon/2)$ is decomposable on tangent hyperplanes with respect to δ and the collection $\{V_\epsilon^\delta(s')\}_{s' \neq s}$. \square

Next, we will show that decomposability on tangent hyperplanes implies decomposability. The formal statement is as follows:

Lemma 4. *Suppose that for each s and each $v_0 \in V_\epsilon^\delta(s)$, there exists $\delta^*(v_0) < 1$ such that if $\delta \geq \delta^*(v_0)$, then the set $B(v_0, \epsilon/2)$ is decomposable on tangent hyperplanes with respect to δ and the collection $\{V_\epsilon^\delta(s')\}_{s' \neq s}$. Then there exists $\delta^* < 1$ such that if $\delta \geq \delta^*$, then for each s and each $v_0 \in V_\epsilon^\delta(s)$, the set $B(v_0, \epsilon/2)$ is decomposable with respect to δ and the collection $\{V_\epsilon^\delta(s')\}_{s' \neq s}$.*

Proof. Pick any s and $v_0 \in V_\epsilon^\delta(s)$. Since every point on the boundary of $B(v_0, \epsilon/2)$ can be decomposed with respect to $\delta^*(v_0)$ and $\{V_\epsilon^\delta(s')\}_{s' \neq s}$, so can every point in the interior of $B(v_0, \epsilon/2)$: each such payoff is achievable (in expectation) through public randomization before the start of play. Since $V_\epsilon^\delta(s)$ is compact (by Lemma 2), it can be covered with a finite collection of such balls, each with an associated $\delta^*(v_0)$. Define $\underline{\delta}(s)$ as the maximum of those $\delta^*(v_0)$'s, and define $\underline{\delta}$ as the maximum $\underline{\delta}(s)$ over $s \in S$. That $\underline{\delta}$ is the required δ^* . \square

Now we can complete the proof of Theorem 1.

Proof of Theorem 1. For each state s , let $\bar{B}(V_\epsilon^\delta(s), \epsilon/2) \equiv \bigcup_{v \in V_\epsilon^\delta(s)} B(v, \epsilon/2)$ denote the generalized closed ball of radius $\epsilon/2$ around the set $V_\epsilon^\delta(s)$. Lemmas 3 and 4 show that the collection of payoff sets $\{\bar{B}(V_\epsilon^\delta(s), \epsilon/2)\}_{s \in S}$ is “self-decomposable” for high enough δ , in the sense that each $\bar{B}(V_\epsilon^\delta(s), \epsilon/2)$ is decomposable with respect to δ and the collection $\{\bar{B}(V_\epsilon^\delta(s), \epsilon/2)\}_{s' \in S}$. Lemma 2 shows that each $V_\epsilon^\delta(s)$ is compact and convex, so each $\bar{B}(V_\epsilon^\delta(s), \epsilon/2)$ is as well. An argument analogous to the second paragraph of FLM’s proof of Lemma 4.2 establishes the result. \square

4. GENERIC FOLK THEOREM

In this section, we show that for generic class of games, Theorem 1 describes essentially all equilibrium payoffs.

We start by describing our notion of genericity. Notice that, given the sets of players, states, and actions available in each state and the monitoring structure, a game G can be identified as a tuple (g, γ) (of stage-game payoffs and transition rates), hence as an element of $\times_{s \in S} (R^{N \times A(s)} \times R^{S \times A(s)})$. Let \mathcal{G}^M be the set of stochastic games with an absolute bound $0 < M < \infty$. Thus, \mathcal{G}^M is a subset of finitely dimensional space that is convex and that has non-empty interior. Let Λ be the Lebesgue measure on \mathcal{G}^M .

Let $\mathcal{G}_0 \subseteq \mathcal{G}^M$ be the class of all stochastic games such that $V_0^1(s)$ has non-empty interior for each state s . (Note that this class has positive measure: any stochastic game G such that the stage game in each state has a feasible and individually rational payoff set with full dimension, and such that all transition rates are 0, satisfies the condition, as does an open set around G .) 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) $\Lambda(\mathcal{G}_0 \setminus \mathcal{G}') = 0$, and ii) the claim holds for each game $G \in \mathcal{G}'$.

Theorem 2. (*Generic Folk Theorem*)

- (1) *For every game, for each $\eta > 0$, there exists $\delta^* < 1$ such that if $\delta \geq \delta^*$, then for each s and each $v \in E^\delta(s)$ there exists $v' \in V_0^1(s)$ such that $\|v - v'\| \leq \eta$.*
- (2) *Suppose that the Identifiability Condition holds. Then for generic games $G \in \mathcal{G}_0$, for each $\eta > 0$, there exists $\delta^* < 1$ such that for each s and each $v \in V_0^1(s)$ and each $\delta \geq \delta^*$, there exists $v' \in E^\delta(s)$ such that $\|v - v'\| \leq \eta$.*

The first part of Theorem 2 says that any for sufficiently patient players, all equilibrium payoffs are close to the feasible and individually rational set. (Recall that we refer to $V_0^1(s)$ as the set of feasible and individually rational payoffs in state s .) The second part of the theorem says that, for generic games, all elements of the feasible and individually rational set can be approximated by payoffs in subgame perfect equilibria, as long as that set has nonempty interior and the Identifiability Condition holds.

The first part of the Theorem follows from Remark 1 and the fact that for any convergent sequence $v_n \in V_0^{\delta_n}(s)$, where $\delta_n \rightarrow 1$, its limit belongs to the feasible and individually rational set: $\lim_{n \rightarrow \infty} v_n \in V_0^1(s)$. (To see the latter fact, notice that the set that consists of all limits of such sequences, $\limsup_{\delta \rightarrow 1} V_\epsilon^\delta(s)$, is $(1, 0)$ -feasible and self- $(1, 0)$ -individually rational, so it must be contained in $V_0^1(s)$.)

The proof of the second part of Theorem 2 consists of a few steps. First, we define the following property:

Definition 3. *Property A:* Game $G \in \mathcal{G}_0$ has *Property A* if for each $\eta > 0$, there exists $\epsilon > 0$ such that for each state s and each $v \in V_0^1(s)$, there exists $v' \in V_\epsilon^1(s)$ such that $\|v - v'\| \leq \eta$.

We show (in Lemma 5) that if the game satisfies property A, then it also satisfies the thesis of the second part of Theorem 2. (The proof relies on Theorem 1.) Thus, it is enough to show that Property A is generic. It is useful to notice that Property A is trivially satisfied in standard (not stochastic) games whose feasible and individually rational payoff sets have nonempty interior.

Remark 2. For each game $G = (g, \gamma) \in \mathcal{G}_0$, if $\gamma(a, s) = 0$ for each a and s , then G has Property A.

Proof. The result follows from the fact that for such games, any ϵ -interior of $V_0^1(s)$ is $(1, \epsilon)$ -feasible and individually rational. Moreover, because $G \in \mathcal{G}_0$, $V_0^1(s)$ is the limit of such ϵ -interiors. \square

Next, for each game G , we define a class of games G^η parametrized by a one-dimensional parameter η such that $G = G^1$. We show (in Lemma 7) that there are at most countably many $\eta \geq 1$ such that game G^η *does not* have Property A, so such games have zero-measure in the one-dimensional space $\{G^\eta\}$. We expand on this observation to show that the subset of games without Property A has zero measure.

Lemma 5. *Suppose that game $G \in \mathcal{G}_0$ satisfies Property A and the Identifiability Condition. Then for each $\eta > 0$, there exists $\delta^* < 1$ such that for each $v \in V_0^1$ and each $\delta \geq \delta^*$, there exists $v' \in E^\delta(s)$ such that $\|v - v'\| \leq \eta$.*

Proof. Notice that for each $\epsilon > 0$ and each $\delta > 1 - \frac{\epsilon}{2M}$, the set V_ϵ^1 is $(\delta, \epsilon/2)$ -feasible and self- $(\delta, \epsilon/2)$ -individually rational. It follows that $V_\epsilon^1(s) \subseteq V_{\epsilon/2}^\delta(s)$. Since $G \in \mathcal{G}_0$, it follows that if game G has Property A, then the Uniform Interior Condition holds. Thus, there exist $\epsilon > 0$ and $\delta^* = 1 - \frac{\epsilon}{2M}$ such that for each $v \in V_0^1(s)$, each $\delta \geq \delta^*$, there exists $v' \in V_\epsilon^\delta(s)$ such that $\|v - v'\| \leq \eta$. If δ is high enough, then v' is a payoff in subgame perfect equilibrium by Theorem 1. \square

Next, we show that Property A holds for generic games $G \in \mathcal{G}_0$. The proof relies on the following observation. For each $d \in R^{S \times N}$, let $\mathcal{G}_0(d) \subseteq \mathcal{G}_0$ be a class of games such that for each state s , $d(s) \in \text{int}V_0^1(s)$. For each game $G \in \mathcal{G}_0(d)$, we define a one-dimensional class of games $G^{\eta;d} = (g^{\eta;d}, \gamma^{\eta;d}, \pi)$, where $\eta \leq 1 + (\max_{a,s} \gamma(a,s))^{-1}$. Let

$$\begin{aligned} g^{\eta;d}(a,s) &= \frac{\eta}{1 - (\eta - 1)\gamma(a,s)}g(a,s) - \frac{(\eta - 1)(1 + \gamma(a,s))}{1 - (\eta - 1)\gamma(a,s)}d(s) \\ \gamma^{\eta;d}(a,s) &= \frac{\eta}{1 - (\eta - 1)\gamma(a,s)}\gamma(a,s). \end{aligned}$$

Lemma 6. *For each $\eta, \nu \geq 1$, each $G \in \mathcal{G}_0(d)$,*

- (1) $(G^{\eta;d})^{\nu;d} = G^{\eta\nu;d}$,
- (2) *for each action profile a , each state s , all continuation payoffs $u \in R^{S\{s\}}$, $\psi^1(a, s, u; G^{\eta;d}) - d = \eta(\psi^1(a, s, u; G) - d)$,*
- (3) *if $\eta > 1$, and $G \in \mathcal{G}_0(d)$, then there exists $\epsilon > 0$ such that for each state $s, V_0^1(s) \subseteq V_\epsilon^1(s; G^{\eta;d})$.*
- (4) *if $G \in \mathcal{G}_0(d)$, then $G^{\eta;d} \in \mathcal{G}_0(d)$.*

Proof. Part (1). Notice that

$$\begin{aligned} & \frac{\eta}{1 - (\eta - 1)\frac{\nu}{1 - (\nu - 1)\gamma(a,s)}\gamma(a,s)} \frac{\nu}{1 - (\nu - 1)\gamma(a,s)} \\ &= \frac{\eta\nu}{1 - (\nu - 1)\gamma(a,s) - (\eta - 1)\nu\gamma(a,s)} \\ &= \frac{\eta\nu}{1 - (\eta\nu - 1)\gamma(a,s)}. \end{aligned}$$

This implies that $(\gamma^{\eta;d})^{\nu;d} = \gamma^{\eta\nu;d}$. In a similar way, we show that $(g^{\eta;d})^{\nu;d} = g^{\eta\nu;d}$.

Part (2). For each $\eta > 0$,

$$\begin{aligned}
 & \psi^1(a, s, u; G^{\eta; d}) - d \\
 &= \frac{g^{\eta; d}(a, s) + \sum_{s' \neq s} \gamma^{\eta; d}(s'; a, s) u(s')}{1 + \gamma^{\eta; d}(a, s)} - d(s) \\
 &= \frac{\eta}{1 + \gamma(a, s)} g(a, s) - (\eta - 1) d(s) \\
 &\quad + \frac{\sum_{s' \neq s} \eta \gamma(s'; a, s) u(s')}{1 + \gamma(a, s)} - d(s) \\
 &= \eta \left(\frac{g(a, s) + \sum_{s' \neq s} \gamma(s'; a, s) u(s')}{1 + \gamma(a, s)} - d(s) \right) \\
 &= \eta (\psi^1(a, s, u) - d)
 \end{aligned}$$

Part (3). Let $r > 0$ be such that for each s , $B(d(s), r) \subseteq V_0^1(s)$. Let $\epsilon = (\eta - 1)r$. We show that collection $V_0^1(\cdot; G)$ is $(1, \epsilon)$ -individually rational in game $G^{\eta; d}$. Using part (2), we get

$$\begin{aligned}
 & e_i^1(s; V_0^1(\cdot; G), G^{\eta; d}) \\
 &= \sup_{\alpha_i \in \Delta A_i} \inf_{\alpha_{-i} \in \Delta A_{-i}} \inf_{u \in \times_{s' \neq s} V_0^1(s'; G)} \psi_i^1(\alpha_i, \alpha_{-i}, s, u; G^{\eta; d}) \\
 &= \sup_{\alpha_i \in \Delta A_i} \inf_{\alpha_{-i} \in \Delta A_{-i}} \inf_{u \in \times_{s' \neq s} V_0^1(s'; G)} d_i(s) + (\psi_i^1(\alpha_i, \alpha_{-i}, s, u; G^{\eta; d}) - d_i(s)) \\
 &= \sup_{\alpha_i \in \Delta A_i} \inf_{\alpha_{-i} \in \Delta A_{-i}} \inf_{u \in \times_{s' \neq s} V_0^1(s'; G)} d_i(s) + \eta (\psi_i^1(\alpha_i, \alpha_{-i}, s, u; G) - d_i(s)) \\
 &= -(\eta - 1) d_i(s) + \eta \sup_{\alpha_i \in \Delta A_i} \inf_{\alpha_{-i} \in \Delta A_{-i}} \inf_{u \in \times_{s' \neq s} V_0^1(s'; G)} \psi_i^1(\alpha_i, \alpha_{-i}, s, u; G) \\
 &= d_i(s) + \eta (e_i^1(s; V_0^1(\cdot; G), G) - d_i(s)).
 \end{aligned}$$

Because $e_i^1(s; V_0^1(\cdot; G), G) \leq d_i(s) - r$, it must be that $e_i^1(s; V_0^1(\cdot; G), G^{\eta; d}) \leq e_i^1(s; V_0^1(\cdot; G), G) - \eta r$. This implies that $V_0^1(\cdot; G)$ is $(1, \epsilon)$ -individually rational in game $G^{\eta; d}$.

Next, we show that collection $V_0^1(\cdot; G)$ is $(1, \epsilon)$ -feasible. Take any $v_0 \in V_0^1(s)$ and $v \in B(v_0, \epsilon)$. For each $\lambda \in R^I$ such that $\max_i |\lambda_i| = 1$,

$$\begin{aligned} \lambda \cdot v &\leq \epsilon + \lambda \cdot v_0 \leq (\eta - 1)r + \sup_{v' \in V_0^1(s)} \lambda \cdot v' \\ &= (\eta - 1)r + \psi^1(a, s, u; G) \end{aligned}$$

for some action profile a^λ , and continuation payoffs $u^\lambda \in \times_{s' \neq s} V_0^1(s'; G)$, where

$$(a^\lambda, u^\lambda) \in \arg \max_{a \in A, u \in \times_{s' \neq s} V_0^1(s')} \lambda \cdot \psi^1(a, s, u; G).$$

Because $V_0^1(\cdot; G)$ is $1, 0$ -feasible in game G , and $B(d(s), r) \subseteq V_0^1(s)$, it must be that $r \leq \lambda \cdot (\psi^1(a^\lambda, s, u^\lambda; G) - d(s))$ and

$$\begin{aligned} \lambda \cdot v &\leq (\eta - 1)r + \lambda \cdot \psi^1(a^\lambda, s, u^\lambda; G) \\ &\leq (\eta - 1)\lambda \cdot (\psi^1(a^\lambda, s, u^\lambda; G) - d(s)) + \lambda \cdot (\psi^1(a^\lambda, s, u^\lambda; G) - d(s)) + \lambda \cdot d(s) \\ &= \eta\lambda \cdot (\psi^1(a^\lambda, s, u^\lambda; G) - d(s)) = \lambda \cdot (\psi^1(a^\lambda, s, u^\lambda; G^{\eta;d}) - d(s)), \end{aligned}$$

where the last equality comes from part (2). Because $d(s) \in V_0^1(s; G)$, the above implies that

$$v \in \text{co} \left\{ \psi^1(a, s, u; G^{\eta;d}) : a \in A \text{ and } u \in \times_{s' \neq s} V_0^1(s') \right\},$$

and that $V_0^1(s)$ is $(1, \epsilon)$ -feasible in game $G^{\eta;d}$.

Part (4) follows from part (3) and the fact that $V_\epsilon^1(s; G^{\eta;d}) \subseteq V_0^1(s; G^{\eta;d})$. \square

Lemma 7. *For each game $G \in \mathcal{G}_0(d)$, there exists at most countably many $\eta \in [1, 1 + (\max_{a,s} \gamma(a, s))^{-1}]$ such that game $G^{\eta;d}$ does not have Property A.*

Proof. Let Λ_I be the Lebesgue measure on R^I . Define two functions of $\eta \in [1, 1 + (\max_{a,s} \gamma(a, s))^{-1}]$:

$$\begin{aligned} v_0(\eta) &= \sum_s \Lambda_I(V_0^1(s; G^{\eta;d})), \\ v(\eta) &= \sum_s \Lambda_I\left(\bigcup_{\epsilon > 0} V_\epsilon^1(s; G^{\eta;d})\right). \end{aligned}$$

Because $V_\epsilon^1(s) \subseteq V_0^1(s)$, we have $v_0(\eta) \geq v(\eta)$ for each η and game $G^{\eta;d}$ has Property A if and only if $v_0(\eta) = v(\eta)$. Let $S = \{\eta : v_0(\eta) > v(\eta)\}$. Part (3) of Lemma

6 implies that function $v_0(\eta)$ is increasing; identical argument shows that $v(\eta)$ is increasing. Moreover, part (3) implies that $\lim_{\eta' \searrow \eta} v(\eta') = v_0(\eta)$ for each η . It follows that

$$v(\eta) \geq v(1) + \sum_{s \in S, s < \eta} (v_0(s) - v(s)).$$

If there are uncountably many elements in S , then there exists η such that the right-hand side of the above inequality is unbounded. But this contradicts the fact that $v(\eta)$ is a well-defined function for each $\eta < 1 + (\max_{a,s} \gamma(a,s))^{-1}$. \square

We require one more lemma for the proof of Theorem 2. Let $\mathcal{A} \subseteq \mathcal{G}_0$ be the set of all games with Property A. Recall that Λ is a Lebesgue measure on the space of games.

Lemma 8. *For each $d \in R^{S \times I}$, $\Lambda(\mathcal{G}_0(d) \setminus \mathcal{A}) = 0$.*

Proof. Let $\mathcal{G}_{00} \subseteq \mathcal{G}_0$ be the subclass of all games such that $\max_{a,s} \gamma(a,s) > 0$ and let $\mathcal{G}_{00}^* \subseteq \mathcal{G}_{00}$ be the subclass of all games such that $\max_{a,s} \gamma(a,s) = 1$. Then, there exists mapping $h : \mathcal{G}_{00} \rightarrow \mathcal{G}_{00}^* \times (0, 2)$: for each $G = (g, \gamma) \in \mathcal{G}_{00}$, let $h(G) = (G^{\eta;d}, \eta^{-1})$, where

$$\eta = \frac{1 + \max_{a,s} \gamma(a,s)}{2 \max_{a,s} \gamma(a,s)}.$$

(Notice that the choice of η implies that $G^{\eta;d} \in \mathcal{G}_{00}^*$.) By Lemma 6, the mapping h is bijective, and for each $G \in \mathcal{G}_{00}^*$, $h^{-1}(G, \eta) = G^{\eta;d}$. It is easy to see that the mapping and its inverse are differentiable everywhere on its domains. Moreover, Lemma 6 implies that there exists function j :

Notice that $\mathcal{G}_{00}^* \times (0, 2)$ is a convex subset of finitely dimensional space, and hence it can be equipped with Lebesgue measure Λ^* . By Lemma 7, $\Lambda^*(h((\mathcal{G}_0(d) \cap \mathcal{G}_{00}) \setminus \mathcal{A})) = 0$. Because mapping h is differentiable with differentiable inverse, it follows that $\Lambda((\mathcal{G}_0(d) \cap \mathcal{G}_{00}) \setminus \mathcal{A}) = 0$.

Finally, Remark 2 shows that all games in $\mathcal{G}_0 \setminus \mathcal{G}_{00}$ have Property A. \square

Theorem 2 follows from the above lemma, and the fact that

$$\mathcal{G}_0 = \bigcup_{d \in Q^{N \times S}} \mathcal{G}_0(d)$$

where Q is the set of rational numbers.

5. EXAMPLES

In this section, we present several examples to illustrate properties of stochastic games with infrequent transitions.

Other applications - Acemoglu has this political economy papers.

5.1. Games without Property A. Lemma 8 established that Property A holds for generic games. Here, we present an example where Property A fails. Consider the following game: there are two players and two states, $S = \{s_1, s_2\}$. In each state s_i , player i has two actions. The payoffs are given in the tables below:

Example 1.

H	$-1, 2$	H	L
L	$0, 0$	$2, -1$	$0, 0$
	State s_1	State s_2	

The transition frequencies in each state do not depend on actions and are equal to 1: $\gamma(s_j; a, s_i)$ for all $i, j \in \{1, 2\}, i \neq j$ and all $a \in \{H, L\}$. It is easy to see that the vector of minmax payoffs in each state is $(0, 0)$, for any discount factor. At $\delta = 1$, the sets of feasible payoffs in each state are

$$\begin{aligned} \hat{V}^1(s_1) &= \text{co} \left\{ (0, 0), (0, 1), \left(\frac{2}{3}, \frac{-1}{3}\right), \left(\frac{-2}{3}, \frac{4}{3}\right) \right\}, \\ \hat{V}^1(s_2) &= \text{co} \left\{ (0, 0), (1, 0), \left(\frac{-1}{3}, \frac{2}{3}\right), \left(\frac{4}{3}, \frac{-2}{3}\right) \right\}. \end{aligned}$$

(For example, playing a_1 in state s_1 and a_2 in state s_2 yields $v^1((a_1, a_2), s_i)$ in state s_i , where $v^1((a_1, a_2), s_i) = \psi^1(a_i, s_i, v^1((a_1, a_2), s_j)) = \frac{1}{2}g(a_i, s_i) + \frac{1}{2}v^1((a_1, a_2), s_j)$. Solving yields $v^1((a_1, a_2), s_i) = \frac{2}{3}g(a_i, s_i) + \frac{1}{3}g(a_j, s_j)$.) The following proposition shows that although $V_0^1(s)$ has non-empty interior for each state s , this game violates Property A. That is, there is a discontinuity of the largest $1, \epsilon$ -feasible and self- $1, \epsilon$ -individually rational set with respect to ϵ at $\epsilon = 0$. In particular, at $\epsilon = 0$ the set has full dimension, but for any positive ϵ the set is empty. (The nature of the non-genericity that allows Property A to fail is that the pure-strategy profile of playing H in both states after every history gives player i a payoff, starting from state s_i , exactly equal to his minmax payoff.)

Proposition 1. *For the game in Example 1,*

$$\begin{aligned} V_0^1(s_1) &= \text{co} \left\{ (0, 1), (0, 0), \left(\frac{1}{2}, 0\right) \right\}, \\ V_0^1(s_2) &= \text{co} \left\{ (1, 0), (0, 0), \left(0, \frac{1}{2}\right) \right\}, \\ V_\epsilon^1(s_1) &= V_\epsilon^1(s_2) = \emptyset \forall \epsilon > 0. \end{aligned}$$

Proof. It is straightforward to verify that the sets $V_0^1(s_1)$ and $V_0^1(s_2)$ as defined above are 1, 0-feasible and self-1, 0-individually rational. Further, it is easy to show that any 1, 0-feasible and self-1, 0-individually rational set must lie in

$$\left(\hat{V}^1(s) \cap \left\{ (v_1, v_2) \in \mathbb{R}^2 : v_1, v_2 \geq 0 \right\} \right)_s$$

(since minmax payoffs relative to $V_0^1(s)$ are weakly greater than minmax payoffs relative to the feasible set $\hat{V}^1(s)$). Since $V_0^1(s) = \hat{V}^1(s) \cap \{(v_1, v_2) \in \mathbb{R}^2 : v_1, v_2 \geq 0\}$ for both states, it follows that $V_0^1(s)_s$ is the largest 1, 0-feasible and self-1, 0-individually rational set.

Next, consider $\epsilon > 0$, and suppose that $V_\epsilon^1(s_1)$ is nonempty. It follows by the symmetry of the game that $V_\epsilon^1(s_2)$ is nonempty as well. Let

$$l = \max_{(v_1, v_2) \in V_\epsilon^1(s_1)} (2v_1 + v_2),$$

and note that (by symmetry again) l also satisfies

$$l = \max_{(v_1, v_2) \in V_\epsilon^1(s_2)} (v_1 + 2v_2).$$

Thus, any $(v_1, v_2) \in V_\epsilon^1(s_2)$ satisfies $v_2 \leq \frac{1}{2}l - \frac{1}{2}v_1$. Further, since $B((v_1, v_2), \epsilon)$ must lie in $\hat{V}^1(s) \cap \{(v_1, v_2) \in \mathbb{R}^2 : v_1, v_2 \geq 0\}$, it must be that $v_2 \geq \epsilon$. Therefore, $v_1 \leq l - 2\epsilon$, since $v_1 + 2v_2 \leq l$. Finally, define $V_\epsilon^1(s_2|1) \equiv \{v_1 \in \mathbb{R} : \exists v_2 \text{ s.t. } (v_1, v_2) \in V_\epsilon^1(s_2)\}$, and note that

$$\begin{aligned}
l &\leq \max_{a \in \{H,L\}, (v_1, v_2) \in V_\epsilon^1(s_2)} \{2\psi_1^1(a, s_1, (v_1, v_2)) + \psi_2^1(a, s_1, (v_1, v_2))\} \\
&= \frac{1}{2} \max_{a \in \{H,L\}} \{2g_1(a, s_1) + g_2(a, s_1)\} + \frac{1}{2} \max_{(v_1, v_2) \in V_\epsilon^1(s_2)} \{2v_1 + v_2\} \\
&\leq 0 + \frac{1}{2} \max_{v_1 \in V_\epsilon^1(s_2|1)} \left\{ 2v_1 + \frac{1}{2}l - \frac{1}{2}v_1 \right\} \\
&\leq \max_{v_1 \in V_\epsilon^1(s_2|1)} \left\{ v_1 + \frac{1}{4}l - \frac{1}{4}v_1 \right\} \\
&\leq \max_{v_1 \in V_\epsilon^1(s_2|1)} v_1 \\
&\leq l - 2\epsilon.
\end{aligned}$$

But this is a contradiction, and so both $V_\epsilon^1(s_1)$ and $V_\epsilon^1(s_2)$ must be empty. \square

For the game in Example 1, the second part of Theorem 2 fails: many payoffs in V_0^1 are not achievable in equilibrium. In fact, for any $\delta < 1$, the only subgame perfect equilibrium is to play L after every history, yielding the minmax payoffs $(0, 0)$ in each state. The proof is very similar to the proof of Proposition 1.

Proposition 2. *For the game in Example 1, $E^\delta(s) = \{(0, 0)\}$ for any $\delta < 1$ and each state s , regardless of the monitoring structure.*

Proof. First, note that playing L after every history is a subgame perfect equilibrium, so $\{(0, 0)\} \in E^\delta(s)$. Next, let

$$l = \max_{(v_1, v_2) \in E^\delta(s_1)} (2v_1 + v_2),$$

and note that (by symmetry) l also satisfies

$$l = \max_{(v_1, v_2) \in E^\delta(s_2)} (v_1 + 2v_2).$$

Thus, any $(v_1, v_2) \in E^\delta(s_2)$ satisfies $v_2 \leq \frac{1}{2}l - \frac{1}{2}v_1$. Since individual rationality requires that $v_1, v_2 \geq 0$, we must have $l \geq v_1 \geq 0$. Further, the condition that $v_1, v_2 \geq 0$ implies that if $E^\delta(s_2)$ contains any point other than $(0, 0)$, then $l > 0$. Define $E^\delta(s_2|1) \equiv \{v_1 \in \mathbb{R} : \exists v_2 \text{ s.t. } (v_1, v_2) \in E^\delta(s_2)\}$, and note that

$$\begin{aligned}
 l &\leq \max_{a \in \{H, L\}, (v_1, v_2) \in E^\delta(s_2)} \left\{ 2\psi_1^\delta(a, s_1, (v_1, v_2)) + \psi_2^\delta(a, s_1, (v_1, v_2)) \right\} \\
 &= \frac{1}{1+\delta} \max_{a \in \{H, L\}} \{2g_1(a, s_1) + g_2(a, s_1)\} + \frac{\delta}{1+\delta} \max_{(v_1, v_2) \in E^\delta(s_2)} \{2v_1 + v_2\} \\
 &\leq 0 + \frac{\delta}{1+\delta} \max_{v_1 \in E^\delta(s_2|1)} \left\{ 2v_1 + \frac{1}{2}l - \frac{1}{2}v_1 \right\} \\
 &\leq \frac{\delta}{1+\delta} 2l \\
 &\leq l.
 \end{aligned}$$

The final inequality is strict if $l > 0$, so we conclude that $l = 0$, and that $E^\delta(s_2)$ and (by symmetry) $E^\delta(s_1)$ are equal to $\{(0, 0)\}$. \square

5.2. Markov strategies are not enough. Another interesting property of stochastic games is that the set of payoffs generated by Markov strategies that are individually rational in each state may be strictly smaller than the set of equilibrium payoffs. Given the widespread use of Markov strategies in the applied literature on stochastic games, it is important to note that if we restrict attention to such strategies, we may not be able to describe all possible equilibrium payoffs (leaving aside the question of whether or not those payoffs can be achieved in a Markov perfect equilibrium). To illustrate this possibility, we modify Example 1 slightly, so that Property A is satisfied.

Example 2.

H	$-1, 3$	H	L
L	$0, 0$	$3, -1$	$0, 0$
	State s_1	State s_2	

As before, the transition frequencies in each state do not depend on actions and are equal to 1: $\gamma(s_j; a, s_i)$ for all $i, j \in \{1, 2\}, i \neq j$ and all $a \in \{H, L\}$. The vector of minmax payoffs in each state is $(0, 0)$, for any discount factor. At $\delta = 1$, the sets of feasible payoffs in each state are

$$\begin{aligned}
 \hat{V}^1(s_1) &= \text{co} \left\{ (0, 0), \left(\frac{1}{3}, \frac{5}{3}\right), \left(1, \frac{-1}{3}\right), \left(\frac{-2}{3}, 2\right) \right\}, \\
 \hat{V}^1(s_2) &= \text{co} \left\{ (0, 0), \left(\frac{5}{3}, \frac{1}{3}\right), \left(\frac{-1}{3}, 1\right), \left(2, \frac{-2}{3}\right) \right\}.
 \end{aligned}$$

Let $\tilde{V}^1(s)$ denote the set of feasible payoffs, starting from state s , that give both players at least their minmax payoffs: $\tilde{V}^1(s) \equiv \hat{V}^1(s) \cap \{(v_1, v_2) \in \mathbb{R}^2 : v_1, v_2 \geq 0\}$. Those sets are given by

$$\begin{aligned}\tilde{V}^1(s_1) &= \text{co} \left\{ (0, 0), \left(\frac{1}{3}, \frac{5}{3}\right), \left(\frac{8}{9}, 0\right), \left(0, \frac{16}{9}\right) \right\}, \\ \tilde{V}^1(s_2) &= \text{co} \left\{ (0, 0), \left(\frac{5}{3}, \frac{1}{3}\right), \left(0, \frac{8}{9}\right), \left(\frac{16}{9}, 0\right) \right\}.\end{aligned}$$

We can use Theorem 2 to show that all of these payoffs can be achieved (approximately) in a SPE when players patient.

Proposition 3. *Suppose that the Identifiability Condition holds for the game in Example 2. Then for each $\eta > 0$, there exists $\delta^* < 1$ such that for each $v \in \tilde{V}^1(s)$ and each $\delta \geq \delta^*$, there exists $v' \in E^\delta(s)$ such that $\|v - v'\| \leq \eta$.*

Proof. Without loss of generality, consider initial state s_1 . First, we show that $V_0^1(s_1) = \tilde{V}^1(s_1)$: any 1, 0-feasible and self-1, 0-individually rational set must lie in $(\tilde{V}^1(s))_s$, so it is sufficient to show that $(\tilde{V}^1(s))_s$ is 1, 0-feasible and self-1, 0-individually rational. It is clearly self-1, 0-individually rational, since all payoffs are weakly positive. It is also 1, 0-feasible: $(0, 0) = \psi^1(L, s_1, (0, 0))$, $(\frac{1}{3}, \frac{5}{3}) = \psi^1(L, s_1, (\frac{16}{9}, 0))$, $(\frac{8}{9}, 0) = \psi^1(H, s_1, (\frac{5}{3}, \frac{1}{3}))$, and $(0, \frac{16}{9}) = \psi^1(H, s_1, (1, \frac{5}{9}))$. (Note that $(1, \frac{5}{9}) = \frac{2}{5}(0, \frac{8}{9}) + \frac{3}{5}(\frac{5}{3}, \frac{1}{3})$, so $(1, \frac{5}{9}) \in \tilde{V}^1(s_2)$.) It is straightforward to verify that the collection $(V_\epsilon^1(s))_s$ is continuous in ϵ at $\epsilon = 0$. Thus, the conditions of the second part of Theorem 2 are satisfied, and the result follows. \square

Next, consider a Markov strategy $\sigma^M = (\sigma_1^M, \sigma_2^M)$, where $\sigma_i^M \in [0, 1]$ is the probability that player i plays action H in state i . Denote by $M^\delta(s)$ the set of payoffs in initial state s that are generated by some Markov strategy that yields both player at least their minmax payoffs in both states:

$$M^\delta(s_i) = \left\{ v^\delta \left((\sigma_1^M, \sigma_2^M), s_i \right) : (\sigma_1^M, \sigma_2^M) \in [0, 1]^2, v^\delta \left((\sigma_1^M, \sigma_2^M), s_j \right) \geq 0 \text{ for } j \in \{1, 2\} \right\}.$$

The highest payoff for player i in $M^\delta(s_i)$ is strictly lower those in $\tilde{V}^1(s_i)$. In particular:

Proposition 4. *For the game in Example 2, $\max \{v_i : (v_1, v_2) \in M^\delta(s_i)\} \leq \frac{5}{9}$ for $i \in \{1, 2\}$ and all $\delta \leq 1$.*

Proof. Without loss of generality, consider $i = 1$. First, note that

$$\begin{aligned} v^\delta \left((\sigma_1^M, \sigma_2^M), s_1 \right) &= \psi^\delta \left(\sigma_1^M, s_1, v^\delta \left((\sigma_1^M, \sigma_2^M), s_2 \right) \right) \\ &= \frac{1}{1+\delta} (-\sigma_1^M, 3\sigma_1^M) + \frac{\delta}{1+\delta} v^\delta \left((\sigma_1^M, \sigma_2^M), s_2 \right). \end{aligned}$$

Symmetrically,

$$v^\delta \left((\sigma_1^M, \sigma_2^M), s_2 \right) = \frac{1}{1+\delta} (3\sigma_2^M, -\sigma_2^M) + \frac{\delta}{1+\delta} v^\delta \left((\sigma_1^M, \sigma_2^M), s_1 \right).$$

Solving yields

$$\begin{aligned} v^\delta \left((\sigma_1^M, \sigma_2^M), s_1 \right) &= \frac{1+\delta}{1+2\delta} (-\sigma_1^M, 3\sigma_1^M) + \frac{\delta}{1+2\delta} (3\sigma_2^M, -\sigma_2^M), \\ v^\delta \left((\sigma_1^M, \sigma_2^M), s_2 \right) &= \frac{1+\delta}{1+2\delta} (3\sigma_2^M, -\sigma_2^M) + \frac{\delta}{1+2\delta} (-\sigma_1^M, 3\sigma_1^M). \end{aligned}$$

Individuality rationality requires that $v^\delta \left((\sigma_1^M, \sigma_2^M), s \right) \geq 0$ for each state s . The necessary and sufficient condition is that $\sigma_2^M \in \left[\frac{1+\delta}{1+2\delta} \sigma_1^M, \frac{1+2\delta}{1+\delta} \sigma_1^M \right]$. Thus,

$$\begin{aligned} v_1^\delta \left((\sigma_1^M, \sigma_2^M), s_1 \right) &= \frac{1+\delta}{1+2\delta} (-\sigma_1^M) + \frac{\delta}{1+2\delta} (3\sigma_2^M) \\ &\leq \max_{\sigma \in [0,1]} \left\{ \frac{1+\delta}{1+2\delta} (-\sigma) + \frac{\delta}{1+2\delta} 3\min\left\{1, \frac{1+2\delta}{1+\delta} \sigma\right\} \right\} \\ &\leq \frac{5\delta^2 + \delta - 1}{4\delta^2 + 4\delta + 1} \\ &\leq \frac{5}{9}. \end{aligned}$$

□

Note that for any payoff $v \in \tilde{V}^1(s)$, there exists a Markov strategy that delivers payoff v from initial state s . For example, the strategy $(\frac{1}{6}, 1)$ yields $\frac{2}{3} \left(-\frac{1}{6}, \frac{1}{2} \right) + \frac{1}{3} (3, -1) = \left(\frac{8}{9}, 0 \right)$ at initial state s_1 and $\delta = 1$. Consistent with Dutta's (1995) Lemma 1, any feasible payoff can be achieved by a Markov strategy. The point of the preceding example is that those Markov strategies may not be individually rational after a state transition.

5.3. Ex post individual rationality is not enough. Here we show that even a strategy that delivers each player at least his minmax payoff after every history may yield a payoff that cannot be achieved in equilibrium. In Example 3, there are two states. State 1 is the initial state, and State 2 is absorbing. The discount factor is fixed at δ .

Example 3.

	A2	B2		C2	D2
A1	1, 1	1, 0	C1	4, 0	-1, -2
B1	0, 1	2, 2	D1	9, 2	0, -2
	State s_1			State s_2	

In state s_1 , the transition rate is $\frac{1}{8}$ for every action profile. In state s_2 , the transition rate for every profile is 0. That is, $\gamma(s_2; a, s_1) = \frac{1}{8}$ for all a , and $\gamma(s_1; a, s_2) = 0$ for all a .

In state s_2 , the stochastic game is reduced to a standard repeated game, and it is easy to see that player 1's minmax payoff is 0. We can then use that value to calculate his minmax payoff in state s_1 . Since

$$\begin{aligned}\psi_1^\delta((A1, A2), s_1, 0) &= \psi_1^\delta((A1, B2), s_1, 0) = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 0 = \frac{1}{2}, \\ \psi_1^\delta((B1, A2), s_1, 0) &= \frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 0 = 0, \\ \psi_1^\delta((B1, B2), s_1, 0) &= \frac{1}{2} \cdot 2 + \frac{1}{2} \cdot 0 = 1,\end{aligned}$$

that minmax value is $\frac{1}{2}$. Similarly, player 2's minmax payoff in state s_2 is also 0, and her minmax value in state s_1 is also $\frac{1}{2}$.

Thus, the Markov strategy profile where player 1 plays $B1$ in state s_1 and randomizes with probability $\frac{1}{10}$ on $D1$ in state s_2 , and player 2 plays $A2$ in state s_1 and $C2$ in state s_2 , yielding payoffs $(\frac{9}{4}, \frac{3}{5})$ in state s_1 and $(\frac{9}{2}, \frac{1}{5})$ in state s_2 , gives both players more than their minmax values after every history. There is no subgame perfect equilibrium, however, that gives a payoff close to $(\frac{9}{4}, \frac{3}{5})$ in state s_1 . In fact, any equilibrium must give player 1 a payoff of at least 2.5 in state s_1 .

Proposition 5. *Any subgame perfect equilibrium of the game in Example 3 gives player 1 an expected payoff of at least 2.5.*

Proof. In state s_2 , any feasible payoff that gives player 2 at least her minmax value of 0 gives player 1 a payoff of at least 4. Thus, any continuation equilibrium once state s_2 is reached must give player 1 a payoff of at least 4. By playing action $A1$ in state s_1 until a transition occurs, player 1 can assure himself an expected payoff of at least $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 4 = \frac{5}{2}$ against any equilibrium strategy of player 2. \square

In a standard repeated game, the fact that player 1 must get strictly more than his minmax payoff in any equilibrium (because, as in the game in state s_2 , individual rationality for player 2 requires a high payoff for player 1) does not affect the ability of player 2 to minmax player 1 for a *single period*. In a stochastic game, however, the relevant single-period minmax is the one that corresponds to pseudo-instantaneous payoffs (an action that gives a low payoff today is not an effective threat if it is likely to lead to another state with a high continuation payoff), and pseudo-instantaneous payoffs depend on the available continuation values in other states. Dutta's (1995) Example 2 provides a similar intuition for the case where transition probabilities are independent of δ .³

Note that the result does not depend on the fact that state s_2 is absorbing. Adding a small positive transition rate out of state s_2 would not qualitatively change the result.

5.4. Post-transition incentives are not enough. In the strategies constructed in the proof of Theorem 1, continuation payoffs after a state transition are independent of the public signal about actions. Players' incentives to choose the specified actions take the form of signal-dependent continuations contingent on the state not changing. In principle, we could, like HSTV, strengthen incentives by allowing post-transition payoffs to depend on the public signal as well. HSTV, in fact, take advantage of the

³The details of that example are not quite correct, however. In fact, in state σ , any feasible payoff vector x that gives both players more than 0 can be achieved in a subgame perfect equilibrium, if players are patient. First, note that in the absorbing state, s , patient players can get any payoff in $\text{co}\{(0, 2), (3, 3)\}$ in equilibrium. The payoff x can be written as $x = \alpha w + (1 - \alpha)v$, for some $\alpha \in [0, 1]$, where $v \in \text{co}\{(0, 2), (3, 3)\}$ and $w \in \text{co}\{(0, -1), (0, 3)\}$. (That is, w is the resulting of player 1 playing a_1 and player 2 randomizing between her actions.) Define T as $\delta^T \equiv \alpha$. Then here is a subgame perfect strategy profile that achieves (approximately) x : play the profile that yields w for the first T periods, then switch to state s (by playing (a_2, b_1)) and play the SPE that yields v . After any unilateral deviation by player 2 during the first T periods, restart. After any unilateral deviation by player 1 during the first T periods, switch to state s (if the deviation didn't already result in a switch) and play a SPE that gives player 1 a payoff below x_1 . (To achieve x exactly, the post-transition continuation payoff v would need to be adjusted slightly to compensate for the one period of payoff $(1, 0)$ when the players switch to state s , as well as for the fact that T may not be an integer.) Dutta's (1995) Example 1 has a similar problem.

property that transition rates vary with the action profile, and use the informativeness of state transitions about actions to weaken the sufficient conditions on the monitoring structure. In our setting, however, transitions do not occur frequently enough to allow such a weakening, in general. The following example illustrates the point.

Example 4.

	C	D
C	$1, 1$	$-L, 1 + G$
D	$1 + G, -L$	$0, 0$

State CC, CD, DC, DD

where $L > G > 0$. In each state, the transition rates are the same: for $s, s', a \in \{CC, CD, DC, DD\}$, $\gamma(s'; a, s) = \gamma > 0$ if $s' = a$ and $\gamma(s'; a, s) = 0$ if $s' \neq a$. There is no public signal of actions.

The two states are identical: the same Prisoners' Dilemma stage game is played in each, with the same transition rates. The only factor distinguishing this example from a standard repeated game is that actions are observed only in periods when a state transition occurs: the identity of the new state reveals the action profile that was just played. Clearly, the Identifiability Condition fails, so anything other than infinite repetition of the stage-game equilibrium (D, D) can be achieved only if the (perfect) information about actions revealed when state transitions occurred can be used to provide incentives. Those incentives are not generally strong enough, however: if the transition rate γ is low enough relative to the parameters L and G , then only (D, D) can be played in equilibrium, regardless of δ :

Proposition 6. *In the game in Example 4, if $\gamma < \frac{G(1+L)}{1+L+G}$, then no player plays C with positive probability in equilibrium.*

Proof. Without loss of generality, consider player 1. Since $\bar{v} \equiv \frac{1+L+G}{1+L}$ is the highest feasible payoff for player 1 consistent with player 2 getting at least her minmax value of 0, player 1's payoff in any equilibrium can be no higher than \bar{v} (and no lower than 0). A necessary condition for player 1 to be willing to play C in a period, then, is that

$$(1 - \delta) \cdot 1 + \delta(1 - \delta)\gamma \cdot \bar{v} \geq (1 - \delta)(1 + G) + \delta(1 - \delta)\gamma \cdot 0.$$

That condition holds only if $\delta\gamma \geq \frac{G}{\bar{v}} \geq \frac{G(1+L)}{1+L+G} > 0$. \square

The intuition, roughly, is that because the per-period probability of a transition is proportional to $1 - \delta$, the expected value of changing the infinite stream of future payoffs only if a transition occurs in the current period is on the same scale as the instantaneous payoff. Post-transition incentives, therefore, are only as effective as punishing a deviation for a single period would be in a standard repeated game.

5.5. Oligopoly in evolving world. We give an example of a stochastic game in which the description of the payoffs can be obtained as a solution to a simple dynamic problem.

Consider a generalized oligopoly model. There are N firms. In each period, each firm chooses an action $a_i \in A_i$, and receives an instantaneous payoff $g_i(a, e, s_i)$ that depends on actions of all firms a , current state of the economy e , and firm-specific state s_i . We assume that

- each player i has an *inactive* action 0_i such that $g_i(0_i, a_{-i}, e, s_i) = 0$ for all a_{-i}, e , and s_i ,
- the payoffs of all other players (weakly) increases if player i changes to the inactive action, $g_j(a_i, a_{-i}) \leq g_j(0_i, a_{-i})$ for each
- the minimax payoff of each player in each stage game is equal to his inactive payoff,

$$\min_{\alpha_{-i}} \max_{a_i} g_i(a_i, \alpha_{-i}, e, s_i) = 0.$$

The state of the game is a vector $s = (e, s_1, \dots, s_N)$. We assume that the economy evolves independently from the behavior of the firms with the transition rate from e to e' equal to $\gamma^E(e', e)$. The evolution of firm's private state does not depend on the states or actions of other firms and the transition rate is equal to $\gamma_i(s'_i, a_i, e, s_i)$ from state s_i to s'_i . Because the probability of a simultaneous transition of the economy state and one (or more) of the private states is negligible in the limit $\delta \rightarrow 1$, the joint transition rate from state $s = (e, s_1, \dots, s_N)$ to state $s' = (e', s'_1, \dots, s'_N)$ given action

profile a , is equal to

$$\gamma(s', a, s) = \begin{cases} \gamma_E(e', e) & \text{if } e' \neq e \text{ and } s'_i = s_i \text{ for all } i, \\ \gamma_i(s'_i, a_i, e, s_i) & \text{if } s'_i \neq s_i, e' = e, \text{ and } s'_j = s_j \text{ for all } i \neq j, \\ 0 & \text{otherwise.} \end{cases}$$

The above description fits various oligopoly models. For example, one can interpret actions as tuples of quantities and R&D investment, the state of the economy as market conditions, and the private state of the firm as its patent pool. Because the firms always have an option not to produce and not to invest, they can ensure the payoff at least 0. Because any other firm can flood the market with its own goods, the profits can be forced to remain below any positive number. Finally

Let Λ_+ be the set of vectors λ such that $|\lambda| = 1$, and $\lambda_i \geq 0$ for each i . For each $\lambda \in \Lambda_+$, define

$$c^\lambda(s) = \max_a \frac{\lambda \cdot g(a, s) + \sum_{s' \neq s} \gamma(s', a, s) c^\lambda(s')}{1 + \gamma(a, s)}. \quad (5.1)$$

Proposition 7. *The oligopoly model has Property A, and*

$$V_0^1(s) = \left\{ v : v_i \geq 0 \text{ and } \lambda \cdot v \leq c^\lambda(s) \text{ for each } \lambda \in \Lambda_+ \right\}. \quad (5.2)$$

5.5.1. *Proof of Proposition 7.* For each $\lambda \in \Lambda_+$, let $p(\lambda) = \{i : \lambda_i > 0\}$. For each function $b : S \rightarrow B$, say that b does not depend on the private state of firm i if for each $s'_i \neq s_i$, $b(e, s_i, s_{-i}) = b(e, s'_i, s_{-i})$.

Lemma 9. *For each $\lambda \in \Lambda_+$, there exists a unique solution $c^\lambda(\cdot)$ to the system of equations (5.1) and $c^\lambda(s)$ does not depend on the private state of firms $i \notin p(\lambda)$. Moreover, one can choose action profile $a^\lambda(s)$ that maximizes (5.1) so that it does not depend on the private states of firms $i \notin p(\lambda)$ and $a_i^\lambda(s) = 0_i$ for any such firm.*

Proof. Define mapping $X : R^S \rightarrow R^S$ so that

$$X(c) = \max_a \frac{\lambda \cdot g(a, s) + \sum_{s' \neq s} \gamma(s', a, s) c(s')}{1 + \gamma(a, s)}.$$

Because the transition rates γ s are bounded, mapping X is a contraction. The uniqueness of the fixed point follows from the contraction mapping theorem. In order to

show that $c^\lambda(s)$ does not depend on the private states of firms $i \notin p(\lambda)$, define

$$\gamma^p(s', a, s) = \begin{cases} \gamma_E(e', e) & \text{if } e' \neq e \text{ and } s'_i = s_i \text{ for all } i, \\ \gamma_i(s'_i, a_i, e, s_i) & \text{if } s'_i \neq s_i \text{ for } i \in p(\lambda), \text{ and } e' = e, \text{ and } s'_j = s_j \text{ for all } i \neq j, \\ 0 & \text{otherwise.} \end{cases}$$

Then, the contraction mapping theorem shows that there exists a unique solution to

$$c^{p,\lambda}(s) = \max_a \frac{\lambda \cdot g(a, s) + \sum_{s' \neq s} \gamma^p(s', a, s) c^{p,\lambda}(s')}{1 + \gamma^p(a, s)}. \quad (5.3)$$

and it does not depend on the private states of firms $i \notin p(\lambda)$. We verify that $c^{p,\lambda}$ is a fixed point of mapping X :

$$\begin{aligned} & \max_a \frac{\lambda \cdot g(a, s) + \sum_{s' \neq s} \gamma(s', a, s) c^{\lambda,p}(s')}{1 + \gamma(a, s)} - c^{\lambda,p}(s) \\ &= \max_a \left[\frac{\lambda \cdot g(a, s) + \sum_{s' \neq s} \gamma^p(s', a, s) c^{\lambda,p}(s')}{1 + \gamma(a, s)} + \left(1 - \frac{\gamma(a, s) - \gamma^p(a, s)}{1 + \gamma(a, s)} \right) c^{\lambda,p}(s) \right] \\ &= \max_a \frac{1 + \gamma^p(a, s)}{1 + \gamma(a, s)} \left[\frac{\lambda \cdot g(a, s) + \sum_{s' \neq s} \gamma^p(s', a, s) c^{\lambda,p}(s')}{1 + \gamma^p(a, s)} - c^{\lambda,p}(s) \right] \end{aligned}$$

Let $W(a, s)$ be the value of the expression in the square bracket. Then, $\max_a W(a, s) = W(a^\lambda(s), s) = 0$, where $a^{p,\lambda}(s)$ is an action profile that maximizes (5.3). Because $\frac{1+\gamma^p(a,s)}{1+\gamma(a,s)} > 0$, the above expression is maximized by $a^\lambda(s) = a^{p,\lambda}(s)$ and it is equal to 0.

By construction, $a^{p,\lambda}(s)$ does not depend on the private states of firms $i \notin p(\lambda)$. Because γ^p does not depend on the private states of such firms, and because the payoff of player $j \neq i$ is maximized if player i is inactive, we can assume that $a^{p,\lambda}(s) = 0_i$ for each $i \notin p(\lambda)$. \square

For each λ , assume that action profile $a^\lambda(s)$ that maximizes (5.1) is chosen as in Lemma 9 so that it does not depend on the private states of firms $i \notin p(\lambda)$ and $a_i^\lambda(s) = 0_i$ for any such firm. For each state, find payoffs $v^\lambda(s)$, as a solution to the system of equations:

$$v^\lambda(s) = \frac{g(a^\lambda(s), s) + \sum_{s' \neq s} \gamma(s', a, s) v^\lambda(s')}{1 + \gamma(a^\lambda(s), s)}.$$

Lemma 10. *For each $\lambda \in \Lambda_+$, for each player i , $v_i^\lambda(s) \geq 0$ with equality for $i \notin p(\lambda)$.*

Proof. Because of the choice of $a^\lambda(s)$, $g_i(a, s) = 0$, and $v_i^\lambda(s) = 0$ for each player $i \notin p(\lambda)$.

We prove the Lemma by the induction on the cardinality of $p(\lambda)$. Suppose that the Lemma holds for all λ' such that $|p(\lambda')| < k$ and that there exists λ , $|p(\lambda)| = k$, and player $i \in p(\lambda)$ (i.e., $\lambda_i > 0$) such that $v_i^\lambda(s) < 0$. Define $\lambda^* \in \Lambda_+$. For each player $j \in p(\lambda) \setminus \{i\}$, let

$$\lambda_j^* = \frac{1}{\sqrt{\sum_{j' \in p(\lambda) \setminus \{i\}} \lambda_{j'}^2}} \lambda_j.$$

For any $j \notin p(\lambda)$ or $j = i$, let $\lambda_j = 0$. Then,

$$\begin{aligned} \lambda \cdot v^{\lambda^*}(s) &= \left(\sqrt{\sum_{j' \in p(\lambda) \setminus \{i\}} \lambda_{j'}^2} \right) \lambda^* \cdot v^{\lambda^*}(s) \\ &\geq \left(\sqrt{\sum_{j' \in p(\lambda) \setminus \{i\}} \lambda_{j'}^2} \right) \lambda^* \cdot v^\lambda(s) \\ &= \lambda \cdot v^\lambda(s) - \lambda_i v_i^\lambda(s) > \lambda \cdot v^\lambda(s) \end{aligned}$$

where the last inequality follows from the fact that $\lambda_i v_i^\lambda(s) < 0$. But this leads to a contradiction with the choice of $v^\lambda(s)$. \square

We can finish the proof of Proposition 7. First, notice that because collection V_0^1 is $(1, 0)$ -feasible, it must be that $\max_{v \in V_0^1(s)} \lambda \cdot v \leq c^\lambda(s)$. Thus, this shows that $V_0^1(s)$ is included in the right-hand side of (5.2). Second, by Lemma 10, collection

$$F(s) = \text{con} \left(\{(0, 0, \dots, 0)\} \cup \{v^\lambda(s) : \lambda \in \Lambda_+\} \right),$$

is $(1, 0)$ -individually rational. Because it is clearly $(1, 0)$ -feasible, it must be that $F(s) \subseteq V_0^1(s)$. Finally, because $v_i^\lambda(s) = 0$ for each λ and each firm $i \notin p(\lambda)$, $F(s)$ is equal to the right-hand side of (5.2). This ends the proof of the Proposition.

5.6. General sequences of stochastic games. The results of this paper can be easily generalized to sequences of stochastic games in which the payoffs g_δ and the transition function γ_δ is parametrized with discount factor $\delta < 1$. We assume that $g_\delta \rightarrow g_1 = g$ and $\gamma_\delta \rightarrow \gamma_1 = \gamma$. We can redefine pseudo-instantaneous payoffs ψ^δ , and collections V_ϵ^δ of (δ, ϵ) -feasible and individually rational actions using functions g_δ and

γ_δ instead of g and γ . The statements of Theorems 1 and 2 and their proofs apply with obvious changes. In particular, the definition of property A is not changed.

6. SUMMARY AND DISCUSSION

This paper

7. GENERIC PROPERTIES OF STOCHASTIC GAMES

REFERENCES

References

- Abreu, D., Pearce, D., and Stacchetti, E. (1986). "Optimal Cartel Monitoring with Imperfect Information," *Journal of Economic Theory*, **39**, pp. 251-269.
- Abreu, D., Pearce, D., and Stacchetti, E. (1990). "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, **58**, 1041-1063.
- Dutta, P.K. (1995). "A Folk Theorem for Stochastic Games," *Journal of Economic Theory*, **66**, pp. 1-32.
- Fudenberg, D., Levine, D., and Maskin, E. (1994). "The Folk Theorem with Imperfect Public Information," *Econometrica*, **62**, pp. 997-1039.
- 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.
- Hörner, J., Sugaya, T., Takahashi, S., and Vieille, N. (2011). "Recursive Methods in Discounted Stochastic Games: An Algorithm for $\delta \rightarrow 1$ and a Folk Theorem," *Econometrica*, **79**, pp. 1277-1318.
- Mangasarian, O.L., and Shiau, T.-H. (1987). "Lipschitz Continuity of Solutions of Linear Inequalities, Programs and Complementarity Problems," *SIAM Journal of Control and Optimization*, **25**, pp. 583-595.