

Yale University

EliScholar – A Digital Platform for Scholarly Publishing at Yale

Cowles Foundation Discussion Papers

Cowles Foundation

1-1-2016

How Fast Do Equilibrium Payoff Sets Converge in Repeated Games?

Johannes Hörner

Satoru Takahashi

Follow this and additional works at: <https://elischolar.library.yale.edu/cowles-discussion-paper-series>



Part of the [Economics Commons](#)

Recommended Citation

Hörner, Johannes and Takahashi, Satoru, "How Fast Do Equilibrium Payoff Sets Converge in Repeated Games?" (2016). *Cowles Foundation Discussion Papers*. 2473.

<https://elischolar.library.yale.edu/cowles-discussion-paper-series/2473>

This Discussion Paper is brought to you for free and open access by the Cowles Foundation at EliScholar – A Digital Platform for Scholarly Publishing at Yale. It has been accepted for inclusion in Cowles Foundation Discussion Papers by an authorized administrator of EliScholar – A Digital Platform for Scholarly Publishing at Yale. For more information, please contact elischolar@yale.edu.

HOW FAST DO EQUILIBRIUM PAYOFF SETS CONVERGE
IN REPEATED GAMES?

By

Johannes Hörner and Satoru Takahashi

January 2016

COWLES FOUNDATION DISCUSSION PAPER NO. 2029



COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

<http://cowles.yale.edu/>

How Fast Do Equilibrium Payoff Sets Converge in Repeated Games?*

Johannes Hörner[†], Satoru Takahashi[‡]

January 26, 2016

Abstract

We provide tight bounds on the rate of convergence of the equilibrium payoff sets for repeated games under both perfect and imperfect public monitoring. The distance between the equilibrium payoff set and its limit vanishes at rate $(1 - \delta)^{1/2}$ under perfect monitoring, and at rate $(1 - \delta)^{1/4}$ under imperfect monitoring. For strictly individually rational payoff vectors, these rates improve to 0 (*i.e.*, all strictly individually rational payoff vectors are exactly achieved as equilibrium payoffs for δ high enough) and $(1 - \delta)^{1/2}$, respectively.

Keywords: Repeated games, rates of convergence.

JEL codes: C72, C73.

1 Introduction

Most of the results in repeated games apply to the case of low discounting, especially once imperfect monitoring is considered. The paper by Abreu, Pearce and Stacchetti (1990) is the exception that confirms the rule: in the quarter century since its publication, no paper has managed to improve on its fixed-point characterization of the set of perfect public equilibrium payoffs under imperfect monitoring. In contrast, asymptotic analysis has not only progressed

*This work is partially supported by Singapore MOE AcRF Tier 1 Grant R-122-000-212-112.

[†]Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA, johannes.horner@yale.edu.

[‡]National University of Singapore, 1 Arts Link, Singapore 117570, Singapore, ecsst@nus.edu.sg.

since then (Fudenberg, Levine and Maskin, 1994; Hörner, Sugaya, Takahashi and Vieille, 2011, etc.), but it has also been successfully extended to games with private monitoring.

One of the downsides of asymptotic analysis is its lack of nuance. Whether the folk theorem holds or not depends on linear algebraic properties of the monitoring structure (*e.g.*, the rank of the matrix that summarizes it). Fixing this rank, the amount of noise is irrelevant for the limit characterization. To be sure, Kandori (1992) shows that for any fixed discount factor, the equilibrium payoff set becomes weakly larger as signals become more informative about players' actions. As remarkable as this result may be, it does not help quantify how noise affects equilibrium outcomes. Yet, it is important to know whether the limit characterization is a reasonable approximation for the game with a fixed but high discount factor, in the range of values that are usually used in calibration exercises.

Similar frustration arises in other sciences, in which asymptotic methods allow clear-cut results that are seemingly out of reach for fixed discounting. Undiscounted models are routinely used in engineering and operations research, in particular. However, it is also standard practice in these disciplines to complement such models with an analysis of the rate of convergence (see, for instance, Whitt, 1974, for a survey in queueing; Altman and Zeitouni, 1994, more recently).¹ As imperfect a measure as such a rate might be, it is the next best thing to an exact characterization for fixed discounting.

The goal of this paper is to provide such an analysis for repeated games under perfect as well as imperfect public monitoring. Throughout, we maintain standard assumptions that ensure that the limit (as discounting vanishes) equilibrium payoff set is the entire feasible and individually rational payoff set.²

Our results are as follows. First, under perfect monitoring: the (subgame perfect) equilibrium payoff set converges to its limit at a rate at least as fast as $(1 - \delta)^{1/2}$ (Proposition 1). We show by an example that this rate is tight (Proposition 2). However, there are two caveats to this result. First, all strictly individually rational payoff vectors, often including a large part of the Pareto frontier, are exactly achieved as equilibrium payoffs for high enough discount factors. Second, the example showing that the rate $(1 - \delta)^{1/2}$ is tight is non-generic. Generically (in the sense of

¹This is not to say that rates of convergence have not been studied in economics (and not only with respect to the discount rate, in particular, with respect to growth); see Barro and Sala-i-Martin (1992) and Ortigueira and Santos (1997), among many other examples in macroeconomics. Rates of convergence also play a critical role in econometric theory.

²That is, the “folk theorem” holds as stated by Fudenberg and Maskin (1986) in the case of perfect monitoring and by Fudenberg, Levine and Maskin (1994) in the case of imperfect public monitoring.

stage-game payoffs), the rate can be improved to $1 - \delta$ (Proposition 3).

Second, under imperfect public monitoring: the (perfect public) equilibrium payoff set converges to its limit at a rate at least as fast as $(1 - \delta)^{1/4}$ (Proposition 5). Here as well, we show by an example that this rate is tight (Proposition 6). In this case also, the result comes with two caveats. First, strictly individually rational payoff vectors can be approached at rate $(1 - \delta)^{1/2}$ under mild conditions (Propositions 7 and 8 give precise statements and assumptions), and this rate cannot be improved in general (Proposition 4). Second, the rate $(1 - \delta)^{1/4}$ can be improved to $(1 - \delta)^{1/3}$ if the stage game is the prisoner's dilemma (Proposition 9).

We also show that faster rates can be obtained for an important class of monitoring structures, namely, all-or-nothing monitoring, in which the action profile is either perfectly observed (with some probability) or not at all. For such structures, the rate is $(1 - \delta)^{1/2}$, and generically $1 - \delta$ (consequences of Lemma 7).³

The spread between these rates is significant for applications. For instance, using a discount rate of $1 - \delta = 2\%$, as is frequently done in applications to macroeconomics, it holds that $(1 - \delta)^{1/2} \approx 14\%$, and $(1 - \delta)^{1/4} \approx 38\%$. In a standard prisoner's dilemma (Section 4.2), the efficient payoff vector is achieved under perfect monitoring (given that particular choice of δ), but the loss under imperfect monitoring is of the order $(1 - \delta)^{1/2} \approx 14\%$.⁴ From a theoretical point of view, our results establish that the type of monitoring structure makes a quantitative difference in terms of rates of convergence, even under those assumptions that guarantee that there is no asymptotic difference whatsoever.

Related Literature: Surprisingly, we seem to be the first to study the rate of convergence in non-zero-sum discounted repeated games. Obviously, a similar exercise has been carried out in related areas. We start with a brief overview of those.

There is a sizable literature on the rate of convergence of the value in zero-sum games (both repeated and stochastic games). In the case of one-sided incomplete information and perfect monitoring, Aumann and Maschler (1967) establish that the value converges at rate $(1 - \delta)^{1/2}$.⁵

³On the other hand, product monitoring, as is often assumed, does not improve the rate of convergence. To the contrary, to the extent that, under full support, Nash equilibria and perfect public equilibria are payoff-equivalent for such games, our results extend to Nash equilibrium payoffs.

⁴Clearly, the constant of proportionality matters. Our analysis leaves open the question on how this constant can be bounded in terms of stage-game payoffs and the monitoring structure.

⁵As far as we can tell, the coincidence of this rate with the one we derive for non-zero-sum games with perfect monitoring is fortuitous, although, as in our case, the rate can be substantially improved in specific classes of games, see Zamir (1971–72).

Mertens (1998) extends the analysis to the case of imperfect monitoring. The rate $(1 - \delta)^{1/3}$ then plays a prominent role.

The rate of convergence of the value of Markov decision processes (the special case of a stochastic game with only one player) has been thoroughly investigated. It has been shown that, under fairly weak assumptions (*e.g.*, Yushkevich, 1996), the rate of convergence is $1 - \delta$.⁶

In discounted zero-sum stochastic games under perfect monitoring, Bewley and Kohlberg (1976) show that the value admits a Puiseux series in the discount rate $1 - \delta$. More specifically, as a function of the dimension of the game (the cardinality of the state and action sets), there exists an integer N such that the value admits an expansion in powers of the form i/k , $i = 0, \dots, k$, for some integer $k \leq N$, in some neighborhood of $\delta = 1$. As is well known, there is surprising flexibility in reverse-engineered games with specific power expansions.

In repeated games, some questions related to the rate of convergence have been examined before. In particular, authors have asked whether (i) the Pareto frontier can be exactly achieved under imperfect monitoring for high discount factors, and (ii) payoff vectors that exactly give a player his minmax payoff (as opposed to strictly more) can be achieved for high discount factors under perfect monitoring.

Regarding (i), see Fudenberg, Levine and Takahashi (2007), who provide a sufficient condition (see their Section 4.4), and Azevedo and Moreira (2007), who (in addition to providing sufficient conditions as well) establish an “anti-folk theorem:” an exact folk theorem is valid only for a zero measure set of (two-player) games under imperfect monitoring (see their Theorem 4 for a formal statement). Relative to these results, we provide a measure of the rate at which the equilibrium payoff set approaches the Pareto frontier. Regarding (ii), see Thomas (1995) as well as Berg and Kärki (2014).

Finally, some papers consider particular games and derive exhaustive analyses of the equilibrium payoff set under perfect monitoring, as a function of the discount factor. For the prisoner’s dilemma, these include Stahl (1991) and Mailath, Obara and Sekiguchi (2002).

We discuss some of these papers further as we proceed.

⁶Obviously, in the case of a single state, as in the case of a repeated game, the rate is 0.

2 Notation

We fix a finite stage game. There are $n \geq 2$ players. Each player $i = 1, \dots, n$ has a finite set A_i of actions and a stage-game payoff function, or reward, $u_i: A \rightarrow \mathbf{R}$, where $A = \times_{i=1}^n A_i$. Throughout, we assume that $|A_i| \geq 2$ for all i . The domain of u_i is extended to mixed action profiles $\alpha \in \Delta(A)$ as usual. Let $\bar{u} := \max_{i,a} |u_i(a)|$. For each player i , we normalize his minmax payoff (based on mixed actions) to 0, that is,

$$0 = \min_{\alpha_{-i} \in \times_{j \neq i} \Delta(A_j)} \max_{a_i \in A_i} u_i(a_i, \alpha_{-i}).$$

Throughout this paper, we let $\underline{\alpha}^i = (\underline{a}_i^i, \underline{\alpha}_{-i}^i)$ refer to any of the action profiles that minmaxes player i (in particular, \underline{a}_i^i is a best reply to $\underline{\alpha}_{-i}^i$). We let

$$F_+ := \text{conv } u(A) \cap \mathbf{R}_+^n$$

denote the set of feasible and individually rational payoff vectors.⁷ Throughout, we maintain the assumption that F_+ has non-empty interior. Let

$$\underline{v}_i := \min \{v_i : v \in F_+\}$$

denote the lowest payoff that player i can receive in F_+ . Obviously, $\underline{v}_i \geq 0$ for all i , and the inequality can be strict for some i .⁸

This stage game is repeated infinitely many times. Monitoring is imperfect and public. We let Y denote the finite set of public signals, and $\pi: A \rightarrow \Delta(Y)$ describe the monitoring structure. Section 3 is concerned with the special case of perfect monitoring, which obtains when $Y = A$ and $\pi(\cdot | a)$ assigns probability 1 to $y = a$ for all a . We say that a mixed action profile α has *individual full rank for player i* if $\{\pi(\cdot | a_i, \alpha_{-i}) : a_i \in A_i\}$ is linearly independent; it has *pairwise full rank for distinct players i and j* if $\{\pi(\cdot | a_i, \alpha_{-i}) : a_i \in A_i\} \cup \{\pi(\cdot | a_j, \alpha_{-j}) : a_j \in A_j\}$ contains $|A_i| + |A_j| - 1$ linearly independent vectors.

⁷We let $\text{conv } S$ refer to the convex hull of $S \subseteq \mathbf{R}^n$.

⁸For example, consider the following two-player stage game:

0, -1	1, 1
-1, -2	2, 0

In this game, both players have minmax payoffs of 0, whereas $\underline{v}_1 = 1/3$.

The game unfolds as follows. In each of the rounds $t = 1, 2, \dots$, players simultaneously choose actions. Given the action profile $a \in A$, a public signal y is then drawn with probability $\pi(y | a)$. In addition, players privately know the actions that they have chosen in the past. We assume that no player observes his reward.⁹

Formally, a public history at the start of round t is a sequence $h^t = (y^1, \dots, y^{t-1})$. We set $H^1 := \{\emptyset\}$. The set of public histories at the beginning of round t is therefore $H^t := Y^{t-1}$. We let $H := \bigcup_{t \geq 1} H^t$ denote the set of all public histories. The private history for player i at the beginning of round t is a sequence $h_i^t = (a^1, y^1, \dots, a^{t-1}, y^{t-1})$. Similarly, we define $H_i^1 := \{\emptyset\}$, $H_i^t := (A_i \times Y)^{t-1}$ and $H_i := \bigcup_{t \geq 1} H_i^t$.

A (behavior) strategy for player i is a map $\sigma_i: H_i \rightarrow \Delta(A_i)$. A strategy profile σ generates a probability distribution over histories in the obvious way and thus also a distribution over sequences of players' rewards. Each player seeks to maximize his payoff, that is, the average discounted sum of his rewards, using a common discount factor $\delta < 1$. Given that players follow the strategy profile σ , the payoff of player i is defined as

$$\sum_{t=1}^{\infty} (1 - \delta) \delta^{t-1} \mathbf{E}_{\sigma} [u_i(a^t)].$$

A strategy σ_i is called public if it depends on the public history only, and not on player i 's private information. That is, a public strategy is a map $\sigma_i: H \rightarrow \Delta(A_i)$. A *perfect public equilibrium* is a profile of public strategies such that, given any period t and public history h^t , the strategy profile is a Nash equilibrium from that period on. Perfect public equilibrium is a restrictive equilibrium concept, but it is the only tractable one known to date. Most of the papers on public monitoring focus on such equilibria. So do we. Perfect public equilibrium reduces to subgame perfect equilibrium in the case of perfect monitoring.

We let $E(\delta)$ denote the (compact) set of perfect public equilibrium payoff vectors for any given discount factor $\delta < 1$. Because attention is restriction to perfect public equilibrium, $E(\delta) \subseteq F_+$. Furthermore, because F_+ has non-empty interior, $d(E(\delta), F_+) \rightarrow 0$ as $\delta \rightarrow 1$ under perfect monitoring (Fudenberg and Maskin, 1986) and under imperfect public monitoring and standard rank assumptions (Fudenberg, Levine and Maskin, 1994).¹⁰ Our goal is to understand the rate

⁹One way to interpret this assumption is to view u_i as an expectation: player i 's realized reward $g_i(a_i, y)$ only depends on his action and the public signal, and $u_i(a)$ is defined as $\mathbf{E}[g_i(a_i, y) | a]$. Given π , not all functions u_i can be factorized in this manner, however.

¹⁰We let $d(S, S')$ refer to the Hausdorff metric between two compact subsets $S, S' \subset \mathbf{R}^n$.

at which $d(E(\delta), F_+)$ vanishes as $\delta \rightarrow 1$.

3 Perfect Monitoring

This section focuses on perfect monitoring. In Section 3.1, we prove that the subgame perfect equilibrium payoff set converges to the set of feasible and individually rational payoffs at a rate at least as fast as $(1 - \delta)^{1/2}$ (Proposition 1). We also show by an example that the rate $(1 - \delta)^{1/2}$ is tight (Proposition 2).

Yet, for a given game, only some vertices of F_+ are problematic. In Section 3.2, we review sufficient conditions for particular vertices to be approached at faster rates. In particular, strictly individually rational payoff vectors are exactly achieved for low enough discounting.

Furthermore, not all stage games exhibit such slow convergence. Proposition 3 establishes that the rate can be improved to $1 - \delta$ for generic stage games.

3.1 A Tight Bound

To understand why the rate of convergence is precisely $(1 - \delta)^{1/2}$, consider the following heuristic argument. Suppose that there are two players, and that the feasible and individually rational payoff set is such that player 1's payoff, say, necessarily exceeds player 2's (*i.e.*, $\forall v \in F_+$, $v_1 \geq v_2$). Suppose also that minmaxing player 1 entails a per-period cost $c > 0$ to player 2. Suppose further that player 1 is receiving approximately his lowest payoff in $E(\delta)$, say v_1 , and must be incentivized not to deviate, which would yield an immediate gain of $g > 0$. If such a deviation is punished by T periods of minmaxing, the cost (to player 1) inflicted by this minmaxing phase must exceed the gain g , for the payoff after this phase cannot serve this purpose (player 1 is receiving his lowest payoff already anyhow). That is, we must have

$$(1 - \delta^T) v_1 \geq (1 - \delta) g, \tag{1}$$

where the left-hand side represents the foregone flow payoff during the minmaxing phase which the deviation triggers, and the right-hand side is the one-period benefit from the deviation.

On the other hand, subgame perfection requires player 2 to be willing to carry out the minmaxing. Hence, right after the deviation, his continuation payoff must be positive. Hence, it must be that $-(1 - \delta^T) c + \delta^T v_2 \geq 0$. But we have $v_1 \geq v_2$ for all $v \in F_+$, and so it is necessary

that

$$\delta^T v_1 \geq (1 - \delta^T) c. \quad (2)$$

By multiplying each side of (1) and (2), canceling out $1 - \delta^T$ and taking the square root, we obtain

$$v_1 \geq \left(\frac{(1 - \delta) gc}{\delta^T} \right)^{1/2}.$$

If $(0, 0) \in F_+$, then this payoff vector can only be approached at rate $(1 - \delta)^{1/2}$.

Proposition 1 *Fix a finite stage game under perfect monitoring. Then for any vertex v of F_+ , there exists $K_v > 0$ such that for any $\delta < 1$, there exists a subgame perfect equilibrium whose payoff is within the distance $K_v(1 - \delta)^{1/2}$ of v . In particular, there exists $K > 0$ such that $d(E(\delta), F_+) < K(1 - \delta)^{1/2}$ for any $\delta < 1$.*

Proof. For simplicity, we assume a public randomization device. Further, we assume that mixed actions are observable. We can dispense with both assumptions, following standard arguments (Sorin, 1986; Fudenberg and Maskin, 1991).¹¹

For the first statement, fix any vertex v of F_+ . Because the interior of F_+ is non-empty, we can find $k_1, \dots, k_n, l_1, \dots, l_n \in \mathbf{R}$, independently of δ , such that for every i , (i) $k_i > \bar{u}$ if $v_i = 0$ (hence $\underline{v}_i = 0$), (ii) $l_i > \bar{u}$, and (iii) $v^i(\delta) = (v_1^i(\delta), \dots, v_n^i(\delta)) \in F_+$ for δ close to 1, where

$$\begin{aligned} v_i^i(\delta) &= v_i + k_i(1 - \delta)^{1/2}, \\ v_j^i(\delta) &= v_j + (k_j + l_j)(1 - \delta)^{1/2}, \quad j \neq i. \end{aligned}$$

For example, we can let $(k_1, \dots, k_n) = K_0(v' - v)$ and $l_1 = \dots = l_n = 2\bar{u}$ with any interior point v' of F_+ and sufficiently large $K_0 > 0$. See Figure 1. Note that $v^1(\delta)$ converges to v at rate $(1 - \delta)^{1/2}$.

¹¹More precisely, replacing a public randomization device by a sequence of pure action profiles could affect continuation payoffs by $O(1 - \delta)$. To see this, note that the key inequality in the proof of Fudenberg and Maskin (1991) is that $(1 - \delta^d)2\bar{u} < 3\varepsilon'$ (see p. 433), where d is a fixed integer, to ensure that the continuation payoff remains at all later dates within $4\varepsilon'$ of the target payoff v . It suffices to take $\varepsilon' > 2d\bar{u}(1 - \delta)$ to ensure that this inequality holds for δ high enough, and so the payoff distance is of order $1 - \delta$. Also, without mixed action observability, we could control payoffs in each reward phase R^i by $O(1 - \delta^{T(\delta)}) = O((1 - \delta)^{1/2})$ so that, during the corresponding punishment phase P^i , each punisher $j \neq i$ would become indifferent among all pure actions in the support of the minmax action \underline{a}_j^i . [Note that Proposition 1 requires punishment phases of length at most $1 - \delta^{T(\delta)} \approx (1 - \delta)^{1/2}$, and hence this (times \bar{u}) is the maximum adjustment to make.] Neither argument would affect the target equilibrium payoff vector by more than $O((1 - \delta)^{1/2})$.




Figure 1: Payoff vectors $v^1(\delta), v^2(\delta) \in F_+$ within a distance $O((1-\delta)^{1/2})$ of v such that $v_1^1(\delta) < v_1^2(\delta)$ and $v_2^2(\delta) < v_2^1(\delta)$.

Following closely Fudenberg and Maskin (1986), we shall construct a strategy profile that achieves payoff $v^1(\delta)$ for δ close to 1. The strategy profile can be described in terms of reward phases and punishment phases, as in their construction. Unlike in Fudenberg and Maskin (1986), however, the length of each punishment phase depends on δ . More specifically, given δ , pick the smallest integer $T = T(\delta)$ satisfying

$$1 - \delta^T > (1 - \delta)^{1/2}.$$

Since $1 - \delta^{T(\delta)-1} \leq (1 - \delta)^{1/2}$, we have $1 - \delta^{T(\delta)} \leq 1 - \delta + (1 - \delta)^{1/2}$. Combined with $1 - \delta^{T(\delta)} > (1 - \delta)^{1/2}$, this implies $1 - \delta^{T(\delta)} \approx (1 - \delta)^{1/2}$ as $\delta \rightarrow 1$.

Play starts in phase R^1 in the initial period. In each phase, play proceeds as follows, according to phases $P^i, R^i, i = 1, \dots, n$.

Reward phase R^i : action profile $\alpha^i \in \Delta(A)$ is played (possibly taking advantage of the public randomization device), where $u(\alpha^i) = v^i(\delta)$.¹² Play stays in phase R^i unless a player j unilaterally deviates, in which case play moves to phase P^j .

Punishment phase P^i : The minmaxing action profile $\underline{\alpha}^i$ is played. Play stays in phase P^i for either $T(\delta)$ periods, or until there is a unilateral deviation by some j , whichever comes first. In the former case, play transits to phase R^i .¹³ In the latter, it transits to phase P^j .

To prove that this strategy profile is a subgame perfect equilibrium, we use the one-shot

¹²Thereafter, we will sometimes suppress the dependency of variables on δ if they stay bounded (or bounded away from “baseline” constants) as $\delta \rightarrow 1$ and do not affect our asymptotic analysis.

¹³In case $j = i$, we restart the count of periods.

deviation principle in different phases in turn.

Phase R^i : player i 's incentive constraint is satisfied if

$$v_i + k_i(1 - \delta)^{1/2} \geq (1 - \delta)\bar{u} + \delta^{T(\delta)+1} (v_i + k_i(1 - \delta)^{1/2}),$$

that is,

$$(1 - \delta^{T(\delta)+1}) (v_i + k_i(1 - \delta)^{1/2}) \geq (1 - \delta)\bar{u}.$$

If $v_i > 0$, then this condition is satisfied for δ close to 1, since the term $(1 - \delta^{T(\delta)+1})v_i \approx (1 - \delta)^{1/2}v_i$ dominates all other terms as $\delta \rightarrow 1$. If $v_i = 0$, then this condition reduces to $(1 - \delta^{T(\delta)+1})k_i(1 - \delta)^{1/2} \geq (1 - \delta)\bar{u}$, which is satisfied since $k_i > \bar{u}$ and $1 - \delta^{T(\delta)+1} > 1 - \delta^{T(\delta)} > (1 - \delta)^{1/2}$. Player $j \neq i$'s incentive constraint in phase R^i follows, since the left-hand side is even higher.

Phase P^i : player $j \neq i$'s incentive constraint is satisfied if for all $t \leq T(\delta)$:

$$-(1 - \delta^t)\bar{u} + \delta^t (v_j + (k_j + l_j)(1 - \delta)^{1/2}) \geq (1 - \delta)\bar{u} + \delta^{T(\delta)+1} (v_j + k_j(1 - \delta)^{1/2}),$$

and it is sufficient to consider $t = T(\delta)$. It suffices that

$$-(1 - \delta^{T(\delta)})\bar{u} + \delta^{T(\delta)} (v_j + (k_j + l_j)(1 - \delta)^{1/2}) \geq (1 - \delta)\bar{u} + \delta^{T(\delta)+1} (v_j + k_j(1 - \delta)^{1/2}),$$

that is,

$$\delta^{T(\delta)}l_j(1 - \delta)^{1/2} \geq (2 - \delta - \delta^{T(\delta)})\bar{u} - \delta^{T(\delta)}(1 - \delta) (v_j + k_j(1 - \delta)^{1/2}).$$

This condition is satisfied for δ close to 1 since two terms $\delta^{T(\delta)}l_j(1 - \delta)^{1/2} \approx l_j(1 - \delta)^{1/2}$ and $(2 - \delta - \delta^{T(\delta)})\bar{u} \approx (1 - \delta)^{1/2}\bar{u}$ dominate all other terms as $\delta \rightarrow 1$ and $l_j > \bar{u}$. The argument for the minmaxed player ($j = i$) is standard, as he is taking a short-run best-reply and deviating would only postpone the date at which his flow payoff is positive again.

Note that F_+ is a convex polytope, and so has finitely many vertices, and note that $E(\delta)$ is convex because we assume the public randomization device. Hence, the second statement follows with $K = \max_v K_v$. ■

We now show by an example that the rate is tight. Consider the two-player stage game represented in Figure 2. Each player's minmax payoff is 0, and $F_+ = \text{conv}\{(0, 0), (\frac{1}{2}, 0), (1, 1)\}$.

This example has all the properties used in the heuristic of equations (1) and (2): player 2's payoff is always lower than player 1's by feasibility, and minmaxing player 1 drives player 2's payoff below his own minmax payoff, costing him one per period of punishment, while player 1

	L	R
U	$0, -1$	$1, 1$
D	$0, -1$	$0, 0$

Figure 2: A two-player game in which the rate of convergence is $(1 - \delta)^{1/2}$.




Figure 3: The shaded area depicts F_p , and the short solid line segment depicts line_p .

cannot get less than his minmax payoff even if he wanted to. Suppose for a simplified argument that (D, R) is the action profile that the equilibrium specifies. This exposes player 1 to a temptation equal to $1 - \delta$: minmaxing him for T periods scales down his payoff to $(1 - \delta^T)v_1$ if v_1 is already as low as it gets. Therefore, it holds that $(1 - \delta^T)v_1 \geq 1 - \delta$. Yet, minmaxing player 1 for T periods costs $1 - \delta^T$ to player 2, and so $1 - \delta^T$ must be less than $\delta^T v_2$, where v_2 is the payoff player 2 gets at the end of the punishment (and hence also less than v_1), else minmaxing is not individually rational. Hence, $\delta^T v_1^2 \geq (\delta^T v_2)v_1 \geq (1 - \delta^T)v_1 \geq 1 - \delta$, and so $v_1 \geq ((1 - \delta)/\delta^T)^{1/2}$.¹⁴

Proposition 2 *In the stage game of Figure 2 under perfect monitoring, there exists $\kappa > 0$ such that, for any $\delta < 1$, there is no subgame perfect equilibrium payoff within the distance $\kappa(1 - \delta)^{1/2}$ of $(0, 0)$. Therefore, it holds that $d(E(\delta), F_+) > \kappa(1 - \delta)^{1/2}$ for any $\delta < 1$.*

Proof. One can easily show that $(0, 0) \notin E(\delta)$ for any $\delta < 1$.

For each small $p > 0$, let $F_p := \text{conv}\{(p, 0), (\frac{1}{2}, 0), (1, 1), (\frac{p}{1-2p}, \frac{p}{1-2p})\}$, and line_p be the line segment that connects $(p, 0)$ and $(\frac{p}{1-2p}, \frac{p}{1-2p})$. See Figure 3. Note that the vector $(0, -\frac{1}{2})$ lies on the extension of line_p in the left-down direction. This geometric property will be used later. For each $\delta < 1$, since $(0, 0) \notin E(\delta)$, we can find the largest $p = p(\delta) > 0$ such that $F_p \supseteq E(\delta)$. Note that $p(\delta) \rightarrow 0$ as $\delta \rightarrow 1$. Since the distance between $F_{p(\delta)}$ and $(0, 0)$ is $p(\delta)$, it is enough to show that $p(\delta)$ is of order at least $(1 - \delta)^{1/2}$.

For any δ close to 1, it follows from the maximality of $p(\delta)$ and the compactness of $E(\delta)$ that $E(\delta) \cap \text{line}_{p(\delta)}$ is non-empty. Pick any $v(\delta) \in E(\delta) \cap \text{line}_{p(\delta)}$. Pick any subgame perfect equilibrium that achieves $v(\delta)$. We denote by α the (mixed) action profile played in the first period, and by $w(a, \delta) \in E(\delta)$ the continuation payoff vector after (realized) action profile a . Since $v(\delta)$ is a convex combination of $u(\alpha)$ and $\{w(a, \delta) : a \in A\}$, $v(\delta) \in \text{line}_{p(\delta)}$, and $w(a, \delta) \in F_{p(\delta)}$ for every a , we have $u(\alpha) \in \text{conv}\{(0, 0), (0, -\frac{1}{2}), (\frac{p(\delta)}{1-2p(\delta)}, \frac{p(\delta)}{1-2p(\delta)})\}$, and hence $\alpha_2(R) \geq \frac{1}{2}$.

Since player 1 has no incentive to deviate to U , we have

$$v_1(\delta) \geq (1 - \delta)u_1(U, \alpha_2) + \delta w_1((U, \alpha_2), \delta).$$

¹⁴This argument is not quite formal, as (D, R) need not be played when the lowest payoff vector obtains, nor need player 1 be minmaxed for T consecutive periods after his deviation; a correct bound is given in the proof of Proposition 2.

Since $v_1(\delta) \leq \frac{p(\delta)}{1-2p(\delta)}$, $u_1(U, \alpha_2) = \alpha_2(R) \geq \frac{1}{2}$, and $w_1((U, \alpha_2), \delta) \geq p(\delta)$, we have

$$\frac{p(\delta)}{1-2p(\delta)} \geq \frac{1-\delta}{2} + \delta p(\delta).$$

Solving this inequality, we have

$$p(\delta) \geq \frac{(1-\delta)^{1/2}}{2+2(1-\delta)^{1/2}},$$

and hence $p(\delta)$ is of order at least $(1-\delta)^{1/2}$. ■

3.2 Faster Rates for Particular Payoff Vertices or Generic Games

Arguably, the rate of convergence $(1-\delta)^{1/2}$ established in Section 3.1 is slow. This is a combination of two factors: considering the entire set of equilibrium payoff vectors, as opposed to specific payoff vectors, such as the Pareto frontier only; and considering all possible stage games, as opposed to “almost all” of them, in a sense defined below.

Thus, one might be interested in focusing on specific vertices of F_+ , or specific classes of games, for which convergence occurs at faster rates. But at what rate? Two obvious candidate rates are the rates $1-\delta$ and 0 (the latter means exactly achieving the desired payoff vector, or payoff set, for all high enough discount factors).

Rate 0: Fudenberg and Maskin (1986) show that the rate 0 is achieved if we focus on strictly individually rational payoff vectors.

Claim 1 (Fudenberg and Maskin, 1986) *Fix a finite stage game under perfect monitoring. Then for any $B \subseteq F_+$, closed, such that $v_i > 0$ for all $v \in B$ and i , there exists $\delta_B < 1$ such that $B \subseteq E(\delta)$ for any $\delta \in [\delta_B, 1)$.*

For which stage games is the rate 0 achieved for *all* vectors in F_+ ? The prisoner’s dilemma is one such game (Stahl, 1991; Mailath, Obara and Sekiguchi, 2002). Berg and Kärki (2014) analyze all symmetric 2×2 games and compute the lowest value of δ , if any, for which $E(\delta) = F_+$. For two-player games, Thomas (1995) provides a necessary and sufficient condition for the minmax payoff to be attained in a subgame perfect equilibrium, a requirement closely related to the rate 0.¹⁵

¹⁵It is an open problem for games with more than two players.

Rate $1 - \delta$: Fudenberg and Maskin (1990) examine under which conditions the set of Nash equilibrium payoffs coincides with the set of subgame perfect equilibrium payoffs for δ close to 1. Yet trivially, the set of Nash equilibrium payoff vectors converges to F_+ at rate $1 - \delta$ under perfect monitoring, as each player can be threatened by perpetual minmaxing, so that deviating yields at most a payoff of order $1 - \delta$.¹⁶

Claim 2 (Fudenberg and Maskin, 1990) *Fix a finite stage game under perfect monitoring in which either*

- $\underline{\alpha}_j^i$ is pure and $u_j(\underline{\alpha}^i) > 0$ for all i and j , $i \neq j$; or
- the minmax payoff vector $(0, \dots, 0)$ is in the interior of $\text{conv } u(A)$ and $\min_{a_i \in A_i} u_i(a_i, \underline{\alpha}_{-i}^i) < 0$ for all i .

Then there exists $K > 0$ such that $d(E(\delta), F_+) < K(1 - \delta)$ for any $\delta < 1$.

But weaker conditions suffice for the rate $1 - \delta$. The next lemma gives a set of such sufficient conditions.

Lemma 1 *Fix a finite stage game under perfect monitoring. Fix any vertex v of F_+ such that for each i , it holds that either*

- $v_i > 0$; or
- $\min_{a_i \in A_i} u_i(a_i, \underline{\alpha}_{-i}^i) < 0$; or
- for each $j \neq i$, either $v_j > \underline{v}_j$ or there exists $v' \in F_+$ such that $v'_i = v_i$ and $v'_j > v_j$.

Then there exists $K_v > 0$ such that for any $\delta < 1$, there exists a subgame perfect equilibrium whose payoff is within the distance $K_v(1 - \delta)$ of v .

In words, the first two conditions stated in the lemma are that any player who receives his minmax payoff at the vertex of interest must have an action that yields strictly less than his minmax payoff. Under these conditions, we can pick the duration T_i of player i 's punishment phase to be uniformly bounded in δ . If these conditions are violated, *i.e.*, if $v_i = \min_{a_i \in A_i} u_i(a_i, \underline{\alpha}_{-i}^i) = 0$, then we say that *long minmaxing of player i is required*. Even in this case, if the last condition

¹⁶Fudenberg and Maskin (1990) also provide several examples to illustrate the necessity of their assumptions for the two payoff sets (Nash and subgame perfect equilibrium payoff sets) to coincide.

is satisfied, then we can take $\delta^{T_i(\delta)}$ to be arbitrarily small, and the restriction on j 's payoff v_j ensures that i 's opponents can be incentivized to carry out the punishment. All three conditions play an important role in the proof of Proposition 3.

Proof. This is a modification of the construction of Proposition 1, and the modified construction is only sketched. We construct payoff vectors in a neighborhood of v that are associated with continuation payoff vectors in an equilibrium characterized by phases P^i and R^i , as in Proposition 1. We first observe that players j such that $v_j > \underline{v}_j$ can be incentivized to play any strategy whose continuation payoff vector remains in such a neighborhood, by Claim 1, for δ high enough. It suffices to pick a payoff vector in the interior of F_+ giving player j strictly less than the lowest payoff in the neighborhood (but strictly more than \underline{v}_j) in case he deviates, and use this payoff vector as a continuation equilibrium payoff vector. Hence, we may as well assume that there is no such j .

Let $I \subseteq \{1, \dots, n\}$ denote the set of players i such that $v_i = \min_{a_i \in A_i} u_i(a_i, \underline{a}_{-i}^i) = 0$ (note that I or its complement can be empty). For every $i \in I$, by the assumption, for each $j \neq i$, there exists $v' \in F_+$ such that $v'_i = v_i$ and $v'_j > v_j$. Thus we can pick $v^i(\delta) \in F_+$ such that

$$\begin{aligned} v_i^i(\delta) &= v_i + k_i(1 - \delta) (= k_i(1 - \delta)), \\ v_j^i(\delta) &\geq v_j + \eta, \quad j \neq i \end{aligned}$$

with some $k_i > 0$ (to be chosen later) and $\eta > 0$. For every $i \notin I$, we proceed as before, defining $v^i(\delta) \in F_+$ by

$$\begin{aligned} v_i^i(\delta) &= v_i + k_i(1 - \delta), \\ v_j^i(\delta) &= v_j + (k_j + l_j)(1 - \delta), \quad j \neq i \end{aligned}$$

with some $k_i, l_i > 0$ (to be chosen later).

For $i \in I$, the length $T_i(\delta)$ of punishment phase P^i (during which player i is minmaxed and best replies) is chosen such that $\lfloor \delta^{T_i(\delta)} \rfloor = 1 - \varepsilon$ for some $\varepsilon > 0$ such that $(1 - \varepsilon)(v_j + \eta) - \varepsilon \bar{u} > v_j$ for any $j \neq i$. This ensures that player $j \neq i$ has no profitable deviation during P^i for δ high enough. We pick $k_i > 0$ such that $\varepsilon k_i > 2\bar{u}$ and hence $\delta(1 - \delta^{T_i(\delta)})k_i > 2\bar{u}$ for δ high enough.

This ensures

$$- \underbrace{(1 - \delta)\bar{u}}_{\text{cost of a suboptimal action}} + \underbrace{\delta k_i(1 - \delta)}_{\text{continuation payoff}} \geq \underbrace{(1 - \delta)\bar{u}}_{\text{deviation gain}} + \underbrace{\delta^{T_i(\delta)+1} k_i(1 - \delta)}_{\text{postponed payoff given punishment}},$$

and so player i has no incentive to deviate from R^i either (the cost of a suboptimal action depends on the randomization device, and is at most $(1 - \delta)\bar{u}$). By construction, he has no incentive to deviate during P^i . Also, with fixed $\eta > 0$, player $j \neq i$ has no incentive to deviate from R^i for δ high enough.

For $i \notin I$, the construction is as in Proposition 1, with the adjustment that if $v_i = 0 > \min_{a_i \in A_i} u_i(a_i, \underline{\alpha}_{-i}^i)$, then we further specify that player i plays a pure action yielding a strictly negative payoff against $\underline{\alpha}_{-i}^i$ during P^i and that a deviation from the pure action restarts the punishment. Since $v_i > 0$ or the above case applies, flow payoffs during P^i are below $v_i - \eta$ with some fixed $\eta > 0$. Combined with the restart policy, this allows us to take the length T_i of punishment phase P^i to be independent of δ and yet a deviation from P^i unattractive to i . We can also pick $l_j > 0$, $j \neq i$, such that deviating from P^i is not profitable for players $j \neq i$. The incentive constraints for R^i are unchanged, which is guaranteed by the appropriate choice of k_i, l_i . ■

The assumptions in Lemma 1 can be further weakened, although a simple general criterion appears elusive. However, Lemma 1 suffices to establish the rate $1 - \delta$ in many familiar stage games. How many games fail all sufficient conditions? Not many: the next proposition states that the rate of convergence is generically at least as fast as $1 - \delta$.

For this proposition, we define genericity relative to $\mathcal{U}_0 \subseteq \mathbf{R}^{n|A|}$, the set of all payoff function profiles with full-dimensional F_+ , but without the normalization of minmax payoffs.¹⁷ The proof of the next proposition is in Appendix A.

Proposition 3 *Fix n and $|A_i|$ for each i . Then there exists an open and dense subset $\mathcal{U} \subseteq \mathcal{U}_0$ such that for every $u \in \mathcal{U}$, there exists $K > 0$ such that $d(E(\delta), F_+) < K(1 - \delta)$ for any $\delta < 1$.*

There are examples in which the rate of convergence is $1 - \delta$, and which are robust to payoff perturbations, and so the rate $1 - \delta$ cannot be improved generically. However, there are also robust examples in which the rate is 0. It is an interesting open question whether there exist

¹⁷Similar genericity results can be shown with appropriate normalization.

games for which convergence does not occur at any of these three rates, namely, $1 - \delta$, 0 or the non-generic rate $(1 - \delta)^{1/2}$.

4 Imperfect Public Monitoring

Our goal is to derive results analogous to those of Section 3 for the case in which monitoring is imperfect. To focus on the rate of convergence, as opposed to a possible failure of the folk theorem under imperfect monitoring, we impose standard rank assumptions that ensure that the folk theorem holds.

4.1 Quantifying Surplus Destruction

While under perfect monitoring, convergence can occur relatively slowly, this phenomenon is restricted to vertices of F_+ where at least one player receives his minmax payoff. (See Claims 1 and 2 in the previous subsection.) That is, under perfect monitoring, the cause for slow convergence is

- 1 Providing incentives to minmax other players.

Imperfect monitoring introduces two additional causes.

- 2 Surplus destruction induced by punishment being now carried out on the equilibrium path, approximately along the “hyperplane,” in the language of Fudenberg, Levine and Maskin (1994).
- 3 The randomization over actions that might help detecting deviations at the expense of efficiency.

As we will see, imperfect monitoring can make Cause 1 even worse (see Proposition 6 below). Proposition 4 illustrates Cause 2. Cause 3 is assumed away by the rank assumptions that we will impose.

Let us now illustrate Cause 2 and show that any punishment, even if it is approximately along the “hyperplane,” entails positive surplus destruction. We illustrate this surplus destruction by considering the prisoner’s dilemma described in Figure 4.¹⁸ Each player’s minmax payoff is 0,

¹⁸As the proof and the intuition that follows will make clear, the result is robust to payoff perturbations and holds for a variety of stage games.

	C	D
C	$2, 2$	$-1, 3$
D	$3, -1$	$0, 0$

Figure 4: A two-player game in which equilibrium payoffs are bounded away from efficiency by an order of at least $(1 - \delta)^{1/2}$.

and $F_+ = \text{conv} \{(0, 0), (\frac{8}{3}, 0), (2, 2), (0, \frac{8}{3})\}$. We give a lower bound on this surplus destruction, provided that $\pi(\cdot | a)$ has full support.

Proposition 4 *In the stage game of Figure 4 under imperfect public monitoring with full support, there exists $\kappa > 0$ such that for any $\delta < 1$, there is no perfect public equilibrium payoff vector within distance $\kappa(1 - \delta)^{1/2}$ of $(2, 2)$.*

Proof. Because of full support, we pick $\eta > 0$ such that $\pi(y | a) \geq \eta$ for any $a \in A$ and $y \in Y$. One can show that $(2, 2) \notin E(\delta)$ for any $\delta < 1$ (Fudenberg, Levine and Maskin, 1994, Theorem 6.5).

For each small $p > 0$, let Γ_p be the circle with center at $(2 - p, 2 - p)$ and radius $r = \sqrt{\frac{17}{10}}p$. Note that Γ_p crosses the Pareto frontier of F_+ at four points $(2 + \frac{3p}{10}, 2 - \frac{9p}{10})$, $(2 + \frac{p}{10}, 2 - \frac{3p}{10})$, $(2 - \frac{3p}{10}, 2 + \frac{p}{10})$ and $(2 - \frac{9p}{10}, 2 + \frac{3p}{10})$. Let arc_p denote the arc between $(2 + \frac{p}{10}, 2 - \frac{3p}{10})$ and $(2 - \frac{3p}{10}, 2 + \frac{p}{10})$.

Let F_p be the subset of F_+ that excludes the neighborhood of $(2, 2)$ separated by the arc arc_p . See Figure 5. Since $(2, 2) \notin E(\delta)$, for each $\delta < 1$ close to 1, we can find the largest $p = p(\delta) > 0$ such that $F_p \supseteq E_\delta$. Note that $p(\delta) \rightarrow 0$ as $\delta \rightarrow 1$. Let $r(\delta) = \sqrt{\frac{17}{10}}p(\delta)$. Since the distance between $F_{p(\delta)}$ and $(2, 2)$ is $\sqrt{2}p(\delta) - r(\delta) = (\sqrt{2} - \sqrt{\frac{17}{10}})p(\delta)$, it is enough to show that $p(\delta)$ is of order at least $(1 - \delta)^{1/2}$.

For any δ close to 1, it follows from the maximality of $p(\delta)$ and the compactness of $E(\delta)$ that we have $E(\delta) \cap \text{arc}_{p(\delta)} \neq \emptyset$. Pick any $v(\delta) \in E(\delta) \cap \text{arc}_{p(\delta)}$, and denote by λ the unit vector normal to the arc $\text{arc}_{p(\delta)}$ at point $v(\delta)$. Pick an arbitrary perfect public equilibrium that achieves $v(\delta)$. We denote by α the action profile played in the first period, and by $w(y, \delta) \in E(\delta)$ the continuation payoff vector after signal y . Note that for δ close to 1, both α_1 and α_2 put positive probabilities on C .

Since α_i puts a positive probability on C , for each $i = 1, 2$, we have

$$(1 - \delta)u_i(C, \alpha_j) + \delta \sum_y \pi(y | C, \alpha_j)w_i(y, \delta) \geq (1 - \delta)u_i(D, \alpha_j) + \delta \sum_y \pi(y | D, \alpha_j)w_i(y, \delta).$$




Figure 5: The shaded area depicts F_p , and the inside arc depicts arc_p .

Since $u_i(C, \alpha_j) = u_i(D, \alpha_j) - 1$, there exist $y, y' \in Y$ such that $w_i(y, \delta) - w_i(y', \delta) \geq \frac{1-\delta}{\delta} \geq 1 - \delta$. Thus $w_i(\cdot, \delta)$ must vary by at least $1 - \delta$.

On the other hand, for each $y \in Y$, we have

$$\begin{aligned} \lambda \cdot v(\delta) &= (1 - \delta)\lambda \cdot u(\alpha) + \delta \left(\sum_{y' \neq y} \pi(y' | \alpha) \lambda \cdot w(y', \delta) + \pi(y | \alpha) \lambda \cdot w(y, \delta) \right) \\ &\leq (1 - \delta)\lambda \cdot (2, 2) + \delta((1 - \eta)\lambda \cdot v(\delta) + \eta\lambda \cdot w(y, \delta)), \end{aligned}$$

since $\lambda \cdot u(\alpha) \leq \lambda \cdot (2, 2)$, $\lambda \cdot w(y', \delta) \leq \lambda \cdot v(\delta)$ for any $y' \in Y$, and $\pi(y | \alpha) \geq \eta$. Therefore, we have

$$\lambda \cdot w(y, \delta) \geq \lambda \cdot v(\delta) - \varepsilon(\delta),$$

where

$$\varepsilon(\delta) := \frac{1 - \delta}{\delta\eta} \lambda \cdot ((2, 2) - v(\delta)).$$

Note that, since $v(\delta) \in \text{arc}_{p(\delta)}$, $\lambda \cdot ((2, 2) - v(\delta))$ is of order $p(\delta)$, and hence $\varepsilon(\delta)$ is of order $p(\delta)(1 - \delta)$. Note also that, for δ close to 1, each $w(y, \delta) \in F_{p(\delta)}$ is inside the circle $\Gamma_{p(\delta)}$. Therefore, $w(\cdot, \delta)$ can vary in the direction tangent to λ by at most $2(r(\delta)^2 - (r(\delta) - \varepsilon(\delta))^2)^{1/2} = 2\varepsilon(\delta)^{1/2}(2r(\delta) - \varepsilon(\delta))^{1/2}$, which is of order $p(\delta)(1 - \delta)^{1/2}$.

Combining the above arguments, $p(\delta)(1 - \delta)^{1/2}$ is of order at least $1 - \delta$, hence $p(\delta)$ is of order at least $(1 - \delta)^{1/2}$. ■

The intuition is simple and can already be gleaned from a careful reading of the proof by Fudenberg, Levine and Maskin (1994). The key ingredient is the full support of the monitoring structure. It implies that some continuation payoffs (the vectors $w(y, \delta)$ in the proof) must lie at least $1 - \delta$ apart (for each player), given that this is the order of deviation gain in the one-shot game. At the same time, loosely speaking, if the overall payoff $v(\delta)$ is already within some distance $p(\delta)$ of the boundary (take the closest equilibrium payoff to the boundary in a given direction), then because continuation payoffs differ from the overall payoff by a flow payoff of order $1 - \delta$, we also have that the distance between $v(\delta)$ and $w(y, \delta)$ is of order $p(\delta)(1 - \delta)$ (along the relevant direction). To allow for the payoffs w to be both $1 - \delta$ “apart,” yet within $p(\delta)(1 - \delta)$ of the maximum, the best case scenario is that the boundary be nearly orthogonal to the direction under consideration. Yet, without loss, we can focus on a point at which the curvature of $E(\delta)$ is non-zero; indeed, at least of order $\frac{1}{p(\delta)}$. Hence, points that are $p(\delta)(1 - \delta)$ away from the the same boundary point can only be $p(\delta)(1 - \delta)^{1/2}$ apart from each other, and

hence $p(\delta)$ must be of order (at least) $(1 - \delta)^{1/2}$ as well.

Proposition 4 shows that the rate of convergence toward $(2, 2)$ is no faster than $(1 - \delta)^{1/2}$. In fact, Proposition 7 below implies that the rate is exactly $(1 - \delta)^{1/2}$ under standard rank assumptions. Yet, the rate of convergence towards $(0, \frac{8}{3})$ and $(\frac{8}{3}, 0)$ can be slower than $(1 - \delta)^{1/2}$. Proposition 9 below revisits the prisoner's dilemma and provides a definitive answer for the rate of convergence for the entire payoff set, namely $(1 - \delta)^{1/3}$.

4.2 A Tight Bound

We start with showing that the rate of convergence is at least as fast as $(1 - \delta)^{1/4}$. Recall that the folk theorem holds if, on top of our maintained assumption that F_+ is full-dimensional, first, for every player i , there exists a minmax action profile that has individual full rank for any $j \neq i$, and second, every pure action profile has pairwise full rank for any distinct pair of players.¹⁹ The first assumption (individual full rank) ensures that signals are sufficiently rich for deviations from a prescribed minmax action profile to be statistically detectable, and so to be punishable; the second (pairwise full rank) ensures that signals are so rich that deviations from pure action profiles are not only detected (statistically), but also ascribed to a specific player. This allows surplus to be redistributed rather than destroyed, in case a deviation gets detected, so that efficiency, for instance, does not need to be sacrificed despite the occurrence of punishments. The next result relies on these assumptions.

Proposition 5 *Fix a finite stage game under imperfect public monitoring in which for every player i , there exists a minmax action profile with individual full rank for any $j \neq i$, and that every pure action profile has pairwise full rank for any distinct pair of players. Then there exists $K > 0$ such that $d(E(\delta), F_+) < K(1 - \delta)^{1/4}$ for any $\delta < 1$.*

To prove this proposition, we first review Fudenberg-Levine's program. Let $S^{n-1} := \{\lambda \in \mathbf{R}^n : \|\lambda\| = 1\}$.²⁰ Recall that Fudenberg and Levine (1994) introduce the following program, for each $\lambda \in S^{n-1}$:

$$k(\lambda) = \sup_{\alpha, v, x} \lambda \cdot v,$$

¹⁹Weaker sufficient assumptions for the folk theorem to hold are known. In particular, it suffices that the relevant action profile can be approximated by a sequence of action profiles that have (individual, pairwise) full rank.

²⁰We let $\|v\| := \sqrt{\sum_{i=1}^n v_i^2}$ refer to the Euclidean norm of $v \in \mathbf{R}^n$.

where the supremum is taken over $\alpha \in \times_{i=1}^n \Delta(A_i)$, $v \in \mathbf{R}^n$, and $x: Y \rightarrow \mathbf{R}^n$ satisfying

$$\begin{aligned} v_i &= u_i(\alpha) + \sum_{y \in Y} \pi(y | \alpha) x_i(y) \\ &\geq u_i(a_i, \alpha_{-i}) + \sum_{y \in Y} \pi(y | a_i, \alpha_{-i}) x_i(y) && \forall i = 1, \dots, n, \forall a_i \in A_i, \\ \lambda \cdot x(y) &\leq 0 && \forall y \in Y. \end{aligned}$$

Intuitively, we may think of λ as the weights assigned to players' payoffs (although they can be negative), and v as the vector that maximizes weighted utilitarian welfare. The designer chooses transfers x to incentivize players to choose an action profile that delivers v (net of the transfers), subject to these transfers being budget-balanced, using the same weights λ in the constraint. The maximum value is referred to as the *score* (in direction λ).

For any $\bar{x} > 0$, we denote by $k(\lambda, \bar{x})$ the value of Fudenberg-Levine's program with the additional constraints $\|x(y)\| \leq \bar{x}$ for every $y \in Y$. Also, let

$$k_+(\lambda) := \max_{v \in F_+} \lambda \cdot v.$$

Under the assumptions in Proposition 5, for any λ , we have $k(\lambda, \bar{x}) \rightarrow k(\lambda) \geq k_+(\lambda)$ as $\bar{x} \rightarrow \infty$. We show that this convergence is uniform in λ and of rate $\frac{1}{\bar{x}}$.

Lemma 2 *There exists $K_0 > 0$ such that $k(\lambda, \bar{x}) > k_+(\lambda) - \frac{K_0}{\bar{x}}$ for any $\lambda \in S^{n-1}$ and any $\bar{x} > 0$.*

Proof. For a large L_0 (to be chosen, depending only on the stage-game structure, and not on \bar{x}), we consider the following two cases: (a) $|\lambda_i| > \frac{L_0}{\bar{x}}$ for at least two players, and (b) $|\lambda_i| \leq \frac{L_0}{\bar{x}}$ for all but one player.

In case (a), we assume $|\lambda_1|, |\lambda_2| > \frac{L_0}{\bar{x}}$ without loss of generality. We further assume $\frac{\lambda_1}{\lambda_2} > 0$. (The case with $\frac{\lambda_1}{\lambda_2} < 0$ is similar.) We will show that if \bar{x} is large enough, then every pure action profile a can be enforced with respect to the line $\lambda \cdot x(y) = 0$ such that $\|x(y)\| \leq \bar{x}$, and hence $k(\lambda, \bar{x}) = k(\lambda) \geq k_+(\lambda)$. To see this, first note that for each $i = 3, \dots, n$, since a has individual full rank for player i (implied by the pairwise full rank condition), there exists $x_i(\cdot)$ that satisfies player i 's incentive conditions:

$$\sum_{y \in Y} (\pi(y | a'_i, a_{-i}) - \pi(y | a)) x_i(y) \leq u_i(a) - u_i(a'_i, a_{-i}) \quad \forall a'_i \in A_i.$$

Note that such $x_i(\cdot)$ is independent of λ . Second, we substitute $x_2(y) = -\frac{1}{\lambda_2} \sum_{i \neq 2} \lambda_i x_i(y)$ and rewrite players 1 and 2's incentive constraints as

$$\begin{aligned} \sum_{y \in Y} (\pi(y | a'_1, a_{-1}) - \pi(y | a)) x_1(y) &\leq u_1(a) - u_1(a'_1, a_{-1}) && \forall a'_1 \in A_1, \\ \sum_{y \in Y} (\pi(y | a'_2, a_{-2}) - \pi(y | a)) x_1(y) &\geq -\frac{\lambda_2}{\lambda_1} (u_2(a) - u_2(a'_2, a_{-2})) \\ &\quad - \sum_{i=3}^n \frac{\lambda_i}{\lambda_1} \sum_{y \in Y} (\pi(y | a'_2, a_{-2}) - \pi(y | a)) x_i(y) && \forall a'_2 \in A_2. \end{aligned}$$

Since a has pairwise full rank for players 1 and 2, there exists $x_1(\cdot)$ that satisfies the above condition. Note that $|x_1(y)|$ is of order $\frac{1}{|\lambda_1|} + 1$ and $|x_2(y)|$ is of order $\frac{1}{|\lambda_2|} + 1$. Since $\frac{1}{|\lambda_1|}, \frac{1}{|\lambda_2|} < \frac{\bar{x}}{L_0}$, we can choose L_0 large enough to satisfy $\|x(y)\| \leq \bar{x}$ (for large enough \bar{x}).

In case (b), without loss of generality, we assume $|\lambda_2|, \dots, |\lambda_n| \leq \frac{L_0}{\bar{x}}$. We further assume $\lambda_1 < 0$. (The case with $\lambda_1 > 0$ is similar.) Note that $\lambda_1 < -1 + \frac{L_0}{\bar{x}}$ for large \bar{x} . Since there exists a minmax action profile $\underline{\alpha}^1$ with individual full rank for players $2, \dots, n$, there exists x such that $(\underline{\alpha}^1, (0, \dots, 0), x)$ is feasible in Fudenberg-Levine's program with $(-1, 0, \dots, 0)$. Note that $x_1(y) \geq 0$. Let $L_1 := \max_{i \neq 1, y} |x_i(y)|$. Let $v' = -\frac{(n-1)L_0 L_1}{\bar{x}} \lambda$ and $x'(y) = x(y) - \frac{(n-1)L_0 L_1}{\bar{x}} \lambda$. Since

$$\lambda \cdot x'(y) = \lambda_1 x_1(y) + \sum_{i=2}^n \lambda_i x_i(y) - \frac{(n-1)L_0 L_1}{\bar{x}} \leq 0,$$

$(\underline{\alpha}^1, v', x')$ is feasible in Fudenberg-Levine's program with λ . Therefore, for large enough \bar{x} , we have

$$k(\lambda, \bar{x}) \geq \lambda \cdot v' = -\frac{(n-1)L_0 L_1}{\bar{x}}.$$

Let $v_+ \in F_+$ be such that $\lambda \cdot v_+ = k_+(\lambda)$. Note that $v_{+,1} \geq 0$. Thus,

$$k_+(\lambda) = \lambda_1 v_{+,1} + \sum_{i=2}^n \lambda_i v_{+,i} \leq \frac{(n-1)L_0 \bar{u}}{\bar{x}}.$$

Therefore, we have $k(\lambda, \bar{x}) \geq k_+(\lambda) - \frac{(n-1)L_0(L_1 + \bar{u})}{\bar{x}}$. ■

Next, we record two "geometric" lemmas.

Lemma 3 For any $K_1 > 0$, there exists $K_2 > 0$ such that for any small $r > 0$, there exists a compact, convex, and smooth set W_r such that

1. the curvature of W_r is at most $1/r$; more precisely, for any boundary point v of W_r , there exists a closed ball with radius r that includes v (in its boundary) and is contained in W_r ;
2. $\max_{v \in W_r} \lambda \cdot v < k_+(\lambda) - K_1 r$ for any λ ;
3. $d(W_r, F_+) < K_2 r$.

Proof. With sufficiently large $L_2 > 0$, we shrink F_+ toward some interior point by a factor $1 - L_2 r$, and then smoothen the set by taking its r -neighborhood. ■

Lemma 4 For any $r > 0$ and $0 < \varepsilon < 2r$, if $w = (w_1, \dots, w_n)$ satisfies $w_1 \leq r - \varepsilon$ and $\|w - v\| \leq (2r\varepsilon)^{1/2}$ with $v = (r, 0, \dots, 0)$, then $\|w\| \leq r$.

Proof. We have $\|w\|^2 = \|v\|^2 + 2v \cdot (w - v) + \|w - v\|^2 \leq r^2 - 2r\varepsilon + 2r\varepsilon = r^2$, and hence $\|w\| \leq r$. ■

Proof of Proposition 5. Fix $K_0 > 0$ as in Lemma 2. Let $W(\delta) = W_{r(\delta)}$ as in Lemma 3, where we set $K_1 := K_0 + 3$ and $r(\delta) = (1 - \delta)^{1/4}$. It suffices to show that for any δ close to 1, $W(\delta)$ is self-generating, and hence $W(\delta) \subseteq E(\delta)$.

Pick any boundary point $v(\delta)$ of $W(\delta)$, and denote by λ the unit vector normal to $W(\delta)$ at $v(\delta)$. Letting $\bar{x}(\delta) = (1 - \delta)^{-1/4}$ in Lemma 2, we can achieve the score $k(\delta, \bar{x}(\delta)) > k_+(\lambda) - \frac{K_0}{\bar{x}(\delta)}$ in Fudenberg-Levine's program by $(\alpha, v'(\delta), x(\cdot, \delta))$ with $\|x(y, \delta)\| \leq \bar{x}(\delta)$. Let

$$w(y, \delta) = \frac{1}{\delta}v(\delta) - \frac{1 - \delta}{\delta}v'(\delta) + \frac{1 - \delta}{\delta}x(y, \delta).$$

Then it is easy to show that α is enforceable with continuation payoff vector $w(\cdot, \delta)$, and the total payoff vector is $v(\delta)$.

All is left to show is that $w(y, \delta) \in W(\delta)$. By the construction of $W(\delta)$, it suffices to show that $w(y, \delta)$ belongs to the ball $B(\delta)$ with center $v(\delta) - r(\delta)\lambda$ and radius $r(\delta)$. Note first that

$$\begin{aligned} \lambda \cdot (w(y, \delta) - v(\delta)) &= \frac{1 - \delta}{\delta} \lambda \cdot (v(\delta) - v'(\delta)) + \frac{1 - \delta}{\delta} \lambda \cdot x(y, \delta) \\ &\leq \frac{1 - \delta}{\delta} \lambda \cdot (v(\delta) - v'(\delta)) \\ &< \frac{1 - \delta}{\delta} \left((k_+(\lambda) - K_1 r(\delta)) - \left(k_+(\lambda) - \frac{K_0}{\bar{x}(\delta)} \right) \right) =: -\varepsilon(\delta) \approx -3(1 - \delta)^{5/4}. \end{aligned}$$

On the other hand,

$$\begin{aligned}
\|w(y, \delta) - v(\delta)\| &= \left\| \frac{1-\delta}{\delta} (v(\delta) - v'(\delta) + x(y, \delta)) \right\| \\
&\leq \frac{1-\delta}{\delta} (\|v(\delta)\| + \|v'(\delta)\| + \|x(y, \delta)\|) \\
&\leq \frac{1-\delta}{\delta} (\bar{u} + (\bar{u} + \bar{x}(\delta)) + \bar{x}(\delta)) \approx 2(1-\delta)^{3/4}.
\end{aligned}$$

Therefore, for δ high enough, we have $\|w(y, \delta) - v(\delta)\| < (2r(\delta)\varepsilon(\delta))^{1/2}$. By Lemma 4, we have $w(y, \delta) \in B(\delta)$. ■

As a remark, notice that our individual and pairwise full rank conditions are used in the proof of Lemma 2 only. Therefore, even if the full rank condition holds only approximately, as long as Lemma 2 holds, the same rate $(1-\delta)^{1/4}$ applies. Moreover, if there exist $\beta > 0$ and $K_0 > 0$ such that $k(\lambda, \bar{x}) > k_+(\lambda) - \frac{K_0}{\bar{x}^\beta}$ for any $\lambda \in S^{n-1}$ and any $\bar{x} > 0$, then we can obtain the rate of convergence $(1-\delta)^{\beta/(2\beta+2)}$. (Proposition 5 can be regarded as a special case with $\beta = 1$.)

We now show that the rate $(1-\delta)^{1/4}$ is tight. Recall the two-player stage game of Figure 2. Assume that the monitoring structure has full support, and every pure action profile has pairwise full rank condition for players 1 and 2, and hence Proposition 5 applies. This example shows that the rate $(1-\delta)^{1/4}$ is tight. Roughly speaking, the slow rate arises because near coordinate directions, such as negative coordinate (minmax) directions, it is no longer possible to minimize surplus destruction by shifting continuation payoffs across players. To avoid continuation payoffs to interfere with the minmaxed player's incentives, we must pick these payoffs "well" inside the equilibrium payoff, and this minimum distance exacerbates the inefficiency. To put it differently, in non-coordinate directions, it is possible to "halve" the inefficiency of $(1-\delta)^{1/4}$ across players, but that is no longer possible in coordinate directions. The same issue arises in positive coordinate directions; see Propositions 7 and 8.

Proposition 6 *In the stage game of Figure 2 under imperfect public monitoring with full support, there exists $\kappa > 0$ such that for any $\delta < 1$, there is no perfect public equilibrium payoff within the distance $\kappa(1-\delta)^{1/4}$ of $(0, 0)$.*

Proof. We prove this proposition by contradiction. Let $d(\delta)$ denote the distance between $E(\delta)$ and $(0, 0)$. Suppose that there exists a sequence $\{\delta_m\}_{m \in \mathbf{N}}$ of discount factors such that $\delta_m \rightarrow 1$ and $d(\delta_m)(1-\delta_m)^{-1/4} \rightarrow 0$ as $m \rightarrow \infty$.




Figure 6: The shaded area depicts $F_{p,r}$, and the inside arc depicts $\text{arc}_{p,r}$.

Because of full support, we may pick $\eta > 0$ such that $\pi(y | a) \geq \eta$ for any $a \in A$ and $y \in Y$. One can show that $(0,0) \notin E(\delta)$ for any $\delta < 1$ (Fudenberg, Levine and Maskin, 1994, Theorem 6.5).

For each small $p, r > 0$ such that $p^{-1}r$ is large, let $\Gamma_{p,r}$ be the circle with center at $(p + \frac{r}{\sqrt{4p^2+1}}, -\frac{2pr}{\sqrt{4p^2+1}})$ and radius r . Let $\text{arc}_{p,r}$ denote the arc between $(p,0)$ and $\approx (p,p)$. Let $F_{p,r}$ be the subset of F_+ that excludes the neighborhood of $(0,0)$ separated by $\text{arc}_{p,r}$. See Figure 6.

Note that the tangent line of the arc has slope approximately $\frac{1}{2p}$, and that extending the tangent line in the left-down direction would cross the v_2 -axis at $(0, -\frac{1}{2})$ or above. For each $\delta < 1$ close to 1, since $(0,0) \notin E(\delta)$, we can find the largest $p = p(\delta) > 0$ such that $F_{p,r(\delta)} \supseteq E_\delta$ with $r(\delta) = (1 - \delta)^{1/4}$. Note that $d(\delta)$ and $p(\delta)$ vanish at the same rate as $\delta \rightarrow 1$, and $p(\delta_m)^{-1}r(\delta_m) \rightarrow \infty$ as $m \rightarrow \infty$.

For any δ close to 1, it follows from the maximality of $p(\delta)$ and the compactness of $E(\delta)$ that $E(\delta) \cap \text{arc}_{p(\delta), r(\delta)}$ is non-empty. Pick any $v(\delta) \in E_\delta \cap \text{arc}_{p(\delta), r(\delta)}$. Pick an arbitrary perfect public equilibrium that achieves $v(\delta)$, and denote by λ the unit vector normal to the arc $\text{arc}_{p(\delta), r(\delta)}$ at point $v(\delta)$. We denote by α the action profile played in the first period, and by $w(y, \delta) \in E(\delta)$ the continuation payoff vector after signal y . Since v is a convex combination of $u(\alpha)$ and $\{w(y, \delta) : y \in Y\}$, and $\lambda \cdot v \geq \lambda \cdot w(\delta, y)$ for any y , we have $u(\alpha) \in \text{conv} \{(0, 0), (0, -\frac{1}{2}), (\frac{1}{2}, \frac{1}{2})\}$, and hence $\alpha_1(D), \alpha_2(R) \geq \frac{1}{2}$.

Since player 1 has no incentive to deviate from α_1 to U , $\alpha_1(D) > 0$, and $u_1(D, \alpha_2) = u_1(U, \alpha_2) - \alpha_2(R) \leq u_1(U, \alpha_2) - \frac{1}{2}$, $w_1(\cdot, \delta)$ must vary by order at least $1 - \delta$.

On the other hand, for each $y \in Y$, we have

$$\begin{aligned} \lambda \cdot v(\delta) &= (1 - \delta)\lambda \cdot u(\alpha) + \delta \left(\sum_{y' \neq y} \pi(y' | \alpha) \lambda \cdot w(y', \delta) + \pi(y | \alpha) \lambda \cdot w(y, \delta) \right) \\ &\leq (1 - \delta)\lambda \cdot u(\alpha) + \delta((1 - \eta)\lambda \cdot v(\delta) + \eta\lambda \cdot w(y, \delta)), \end{aligned}$$

since $\lambda \cdot w(y', \delta) \leq \lambda \cdot v(\delta)$ for any $y' \in Y$ and $\pi(y | \alpha) \geq \eta$. Therefore, we have

$$\lambda \cdot w(y, \delta) \geq \lambda \cdot v(\delta) - \varepsilon(\delta),$$

where

$$\varepsilon(\delta) := \frac{1 - \delta}{\delta\eta} \lambda \cdot (u(\alpha) - v(\delta)).$$

Note that since $v(\delta) \in \text{arc}_{p(\delta), r(\delta)}$, $\lambda \cdot (u(\alpha) - v(\delta))$ is of order at most $d(\delta)$, and hence $\varepsilon(\delta)$ is of order at most $d(\delta)(1 - \delta)$. Note also that, for δ close to 1, each $w(y, \delta) \in F_{p(\delta), r(\delta)}$ is inside the circle $\Gamma_{p(\delta), r(\delta)}$. Therefore, $w(\cdot, \delta)$ can vary in the direction tangent to λ by at most $2(r(\delta)^2 - (r(\delta) - \varepsilon(\delta))^2)^{1/2} = 2\varepsilon(\delta)^{1/2}(2r(\delta) - \varepsilon(\delta))^{1/2}$, which is of order at most $d(\delta)^{1/2}(1 - \delta)^{5/8}$. Since the tangent line has a slope approximately equal to $\frac{1}{2p(\delta_m)}$ as $m \rightarrow \infty$, $w_1(\cdot, \delta_m)$ can vary by order at most $d(\delta_m)^{3/2}(1 - \delta_m)^{5/8}$.

Combining the above arguments, $d(\delta_m)^{3/2}(1 - \delta_m)^{5/8}$ is of order at least $1 - \delta_m$. Hence $d(\delta_m)$ is of order at least $(1 - \delta_m)^{1/4}$, which is a contradiction. ■

4.3 A Faster Rate toward Strictly Individually Rational Payoff Vectors

Although Proposition 6 shows that the rate of convergence can be as slow as $(1 - \delta)^{1/4}$, this result depends on both imperfect public monitoring and minmaxing. As we already discussed in Section 3, the rate of convergence is faster under perfect monitoring if we focus on strictly individually rational payoff vectors. Here, we show that under imperfect public monitoring, the rate of convergence is at least as fast as $(1 - \delta)^{1/2}$ toward strictly individually rational payoff vectors that do not maximize any player's payoff.

Proposition 7 *Fix a finite stage game under imperfect public monitoring with the same individual and pairwise full rank conditions as in Proposition 5. Then for any $B \subseteq F_+$, closed, such that $0 < v_i < \max_a u_i(a)$ for all $v \in B$ and i , there exists $K_B > 0$ such that B is contained in the $K_B(1 - \delta)^{1/2}$ -neighborhood of $E(\delta)$ for any $\delta < 1$.*

To prove this, pick small $\varepsilon > 0$ such that $\varepsilon \leq v_i \leq \max_a u_i(a) - \varepsilon$ for all $v \in B$ and i . For each $\lambda \in S^{n-1}$, we define $k_B(\lambda) := \max_{v \in B} \lambda \cdot v$.

Lemma 5 *There exists $\bar{x}_0 > 0$ such that $k(\lambda, \bar{x}) \geq k_B(\lambda)$ for any $\lambda \in S^{n-1}$ and any $\bar{x} > \bar{x}_0$.*

Proof. For a large L_0 (to be chosen, depending only on the stage-game structure, and not on \bar{x}), we consider the following two cases: (a) $|\lambda_i| > \frac{L_0}{\bar{x}}$ for at least two players, and (b) $|\lambda_i| \leq \frac{L_0}{\bar{x}}$ for all but one player.

In case (a), similarly to case (a) in the proof of Lemma 2, it follows from the pairwise full rank condition that there exists L_0 such that, if \bar{x} is large enough, then every pure action profile a can be enforced with respect to the line $\lambda \cdot x(y) = 0$ such that $\|x(y)\| \leq \bar{x}$, and hence $k(\lambda, \bar{x}) = k(\lambda) \geq k_B(\lambda)$.

In case (b), without loss of generality, we assume $|\lambda_2|, \dots, |\lambda_n| \leq \frac{L_0}{\bar{x}}$. We further assume $\lambda_1 < 0$. (The case with $\lambda_1 > 0$ is similar.) Note that $\lambda_1 < -1 + \frac{L_0}{\bar{x}}$ for large \bar{x} . Since there exists a minmax action profile $\underline{\alpha}^1$ with individual full rank for players $2, \dots, n$, there exists x such that $(\underline{\alpha}^1, (0, \dots, 0), x)$ is feasible in Fudenberg-Levine's program with $(-1, 0, \dots, 0)$. Note that $x_1(y) \geq 0$. Let $L_1 := \max_{i \neq 1, y} |x_i(y)|$. Similarly to case (b) in the proof of Lemma 2, for large enough \bar{x} , we have

$$k(\lambda, \bar{x}) \geq -\frac{(n-1)L_0L_1}{\bar{x}}.$$

Let $v_B \in B$ be such that $\lambda \cdot v_B = k_B(\lambda)$. Note that $v_{B,1} \geq \varepsilon$. Thus,

$$k_B(\lambda) = \lambda_1 v_{B,1} + \sum_{i=2}^n \lambda_i v_{B,i} \leq \left(-1 + \frac{L_0}{\bar{x}}\right) \varepsilon + \frac{(n-1)L_0 \bar{u}}{\bar{x}}.$$

Therefore, we have $k(\lambda, \bar{x}) \geq k_B(\lambda)$ for large enough \bar{x} . ■

The rest is the same as the proof of Proposition 5, except that we replace $(1 - \delta)^{1/4}$ by $(1 - \delta)^{1/2}$. That is, we find a compact, convex, and smooth self-generating set that approximates B at the rate $(1 - \delta)^{1/2}$. This corresponds to the remark after the proof of Proposition 5 with $\beta = \infty$.

Proposition 7 excludes max points. The exclusion of max points is not an issue in the prisoner's dilemma, because max points are not individually rational, but this need not be so in other stage games. However, under a mild genericity assumption on stage-game payoffs, we can extend Proposition 7 and obtain the convergence rate $(1 - \delta)^{1/2}$ even toward max points.

Proposition 8 *Fix a finite stage game under imperfect public monitoring with the same individual and pairwise full rank conditions as in Proposition 5. Assume that for every pure action profile $a \in A$, if $u_i(a) = \max_{a'} u_i(a')$, then $u_i(a) > u_i(a''_i, a_{-i})$ for every $a''_i \neq a_i$. Then for any $B \subseteq F_+$, closed, such that $v_i > 0$ for all $v \in B$ and i , there exists $K_B > 0$ such that B is contained in the $K_B(1 - \delta)^{1/2}$ -neighborhood of $E(\delta)$ for any $\delta < 1$.*

Proof. We need to modify case (b) in the proof of Lemma 5 regarding directions λ near $(1, 0, \dots, 0)$. Fix any λ such that $\lambda_1 > 0$ and $|\lambda_2|, \dots, |\lambda_n| \leq \frac{L_0}{\bar{x}}$. Pick any $a \in A$ that maximizes $\lambda \cdot u(a)$. Since λ is near $(1, 0, \dots, 0)$, a also maximizes u_i . Since a has individual full rank for players $2, \dots, n$, there exists x such that $(a, u(a), x)$ is feasible in the program of Fudenberg-Levine with $(1, 0, \dots, 0)$. Let $x'_1(y) = -\frac{1}{\lambda_1} \sum_{i=2}^n \lambda_i x_i(y)$ and $x'_i(y) = x_i(y)$ for every $i = 2, \dots, n$. Since player 1 has strict incentives to play a_i , and $|\lambda_2|, \dots, |\lambda_n|$ are sufficiently small (for \bar{x} large enough), we can enforce a , and achieve the score $\lambda \cdot u(a)$. ■

4.4 The Prisoner's Dilemma

Proposition 6 depends also on stage-game payoffs. Here we show that in the prisoner's dilemma of Figure 4 under imperfect monitoring, the exact rate of convergence is given by $(1 - \delta)^{1/3}$. Moreover, our proof suggests that this result is robust to payoff perturbations, implying

that the rate of convergence is $(1 - \delta)^{1/3}$ for a nonempty and open set of stage-game payoffs.²¹

Proposition 9 *In the stage game of Figure 4 under imperfect public monitoring in which (C, C) , (C, D) , and (D, C) have pairwise full rank for players 1 and 2, there exists $K > 0$ such that $d(E(\delta), F_+) < K(1 - \delta)^{1/3}$ for any $\delta < 1$.*

Once again, the proof is similar to that of Proposition 5. At this time, we use the following property of $k(\lambda, \bar{x})$.

Lemma 6 *There exists $L_0 > 0$ such that $k(\lambda, \bar{x}) > k_+(\lambda) + \frac{1}{2}$ if $\lambda_i < -\frac{1}{2}$ and $\lambda_j \geq \frac{L_0}{\bar{x}}$ for some $i \neq j$.*

Proof. This follows from the pairwise full rank condition of (C, D) and (D, C) . See case (a) in the proofs of Lemmas 2 and 5. ■

Lemma 6 gives a strictly positive lower bound of $k(\lambda, \bar{x}) - k_+(\lambda)$, which is used to show that continuation payoffs are “inward” from the total payoff by $O(1 - \delta)$ along direction λ . Unlike Lemmas 2 and 5, however, Lemma 6 does not give any estimate for λ close enough to $(-1, 0)$ or $(0, -1)$. For such directions, we will enforce static Nash equilibrium (D, D) by constant continuation payoffs.

Proof of Proposition 9. Let $r(\delta) = 5(1 - \delta)^{1/3}$ and $\bar{x}(\delta) = (1 - \delta)^{-1/3}$. For δ high enough, we modify Lemma 3 and construct a compact, convex, and “almost smooth” set $W(\delta)$ with curvature at most $1/r(\delta)$ everywhere except at $(0, 0)$, and the boundary at $(0, 0)$ consists of two half-lines normal to $\left(-\sqrt{1 - \left(\frac{L_0}{\bar{x}(\delta)}\right)^2}, \frac{L_0}{\bar{x}(\delta)}\right)$ and $\left(\frac{L_0}{\bar{x}(\delta)}, -\sqrt{1 - \left(\frac{L_0}{\bar{x}(\delta)}\right)^2}\right)$, respectively.²² (The choice of $K_1 > 0$ in Lemma 3 is arbitrary.) See Figure 7.

Pick any boundary point $v(\delta)$ of $W(\delta)$, and denote by λ the unit vector normal to $W(\delta)$ at $v(\delta)$. If $v(\delta)$ is close to $(0, 0)$, then we can enforce the static Nash equilibrium (D, D) with constant continuation payoffs $w(y, \delta) = \frac{1}{\delta}v(\delta)$. If $v(\delta) \neq (0, 0)$ and $\lambda_1, \lambda_2 \geq -\frac{1}{2}$, then we can adapt the proofs of Propositions 5 and 7 to generate $v(\delta)$ with respect to $W(\delta)$. Thus, we assume that $v(\delta)$ is bounded away from $(0, 0)$ and $\lambda_i < -\frac{1}{2}$ and $\lambda_j \geq \frac{L_0}{\bar{x}}$ for some $i \neq j$. Without loss

²¹We do not know whether this rate applies generically. The rate of convergence is likely to depend on whether there are vertices of $\text{conv } u(A)$ below the minmax payoffs, whether these vertices are strictly negative, whether the minmax profile is a static Nash equilibrium, whether the minmax profile gives minmaxing players positive, zero or negative payoffs, etc.

²²The “kink” at $(0, 0)$ is not essential. What is essential is that the tangent line at any boundary point of $W(\delta)$ near $(0, 0)$ weakly separates $(0, 0)$ and $W(\delta)$.




Figure 7: The shaded area depicts the set $W(\delta)$ of the proof of Proposition 8.

of generality, we assume that $\lambda_1 < -\frac{1}{2}$ and $\lambda_2 \geq \frac{L_0}{\bar{x}(\delta)}$. By Lemma 6, we can achieve the score $k(\delta, \bar{x}(\delta)) > k_+(\lambda) + \frac{1}{2}$ in Fudenberg-Levine's program by $(\alpha, v'(\delta), x(\cdot, \delta))$ with $\|x(y, \delta)\| \leq \bar{x}(\delta)$. Let

$$w(y, \delta) = \frac{1}{\delta}v(\delta) - \frac{1-\delta}{\delta}v'(\delta) + \frac{1-\delta}{\delta}x(y, \delta).$$

Then it is easy to show that α is enforceable with continuation payoff vector $w(\cdot, \delta)$, and the total payoff vector is $v(\delta)$.

All is left to show is that $w(y, \delta) \in W(\delta)$. By construction of $W(\delta)$, it suffices to show that $w(y, \delta)$ belongs to the ball $B(\delta)$ with center $v(\delta) - r(\delta)\lambda$ and radius $r(\delta)$. Note first that

$$\begin{aligned} \lambda \cdot (w(y, \delta) - v(\delta)) &= \frac{1-\delta}{\delta}\lambda \cdot (v(\delta) - v'(\delta)) + \frac{1-\delta}{\delta}\lambda \cdot x(y, \delta) \\ &\leq \frac{1-\delta}{\delta}\lambda \cdot (v(\delta) - v'(\delta)) \\ &< \frac{1-\delta}{\delta} \left((k_+(\lambda) - K_1 r(\delta)) - \left(k_+(\lambda) + \frac{1}{2} \right) \right) =: -\varepsilon(\delta) \approx -\frac{1}{2}(1-\delta). \end{aligned}$$

On the other hand, $\|w(y, \delta) - v(\delta)\|$ is approximately bounded above by $2(1-\delta)^{2/3}$. Therefore, for δ high enough, we have $\|w(y, \delta) - v(\delta)\| < (2r(\delta)\varepsilon(\delta))^{1/2}$. By Lemma 4, we have $w(y, \delta) \in B(\delta)$. ■

It is not difficult to modify the proof of Proposition 6 and show that the rate $(1-\delta)^{1/3}$ is tight for the prisoner's dilemma under imperfect public monitoring with full support.

4.5 All-or-Nothing Monitoring

Neither Proposition 4 nor 6 makes any special assumption on the imperfect monitoring structure. Assuming, for instance, that monitoring has a product structure does not “help” in terms of convergence rates; recall footnote 3.

However, these propositions assume that the monitoring has full support. Yet some natural examples of monitoring structures fail this assumption. All-or-nothing monitoring structures offer such an example. With probability q , players observe the chosen action profile perfectly, while with probability $1-q$, they see “nothing.” Formally, the set of public signals is given by $Y = A \cup \{\emptyset\}$, and $\pi(\cdot | a)$ assigns probability q to $y = a$ and probability $1-q$ to $y = \emptyset$.²³ Let $E(\delta, q)$ be the equilibrium payoff set under discount factor δ and all-or-nothing monitoring with

²³We thank the Associate Editor for suggesting this example.

q . Then we can show the following.²⁴

Lemma 7 *It holds that $E(\delta, q) \supseteq E(\delta q/(1 - \delta + \delta q), 1)$.*

Proof. It is enough to show that $E(\delta q/(1 - \delta + \delta q), 1)$ is self-generating under discount factor δ and all-or-nothing monitoring with q . Pick any $v \in E(\delta q/(1 - \delta + \delta q), 1)$. Then there exist $\alpha \in \times_{i=1}^n \Delta(A_i)$ and $w: A \rightarrow E(\delta q/(1 - \delta + \delta q), 1)$ such that

$$\begin{aligned} v_i &= \frac{1 - \delta}{1 - \delta + \delta q} u_i(\alpha) + \frac{\delta q}{1 - \delta + \delta q} w_i(\alpha) \\ &\geq \frac{1 - \delta}{1 - \delta + \delta q} u_i(a'_i, \alpha_{-i}) + \frac{\delta q}{1 - \delta + \delta q} w_i(a'_i, \alpha_{-i}) \end{aligned}$$

for all $i = 1, \dots, n$ and $a_i \in A_i$. Then we have

$$\begin{aligned} v_i &= (1 - \delta)u_i(\alpha) + \delta(qw_i(\alpha) + (1 - q)v_i) \\ &\geq (1 - \delta)u_i(a_i, \alpha_{-i}) + \delta(qw_i(a_i, \alpha_{-i}) + (1 - q)v_i) \end{aligned}$$

for all $i = 1, \dots, n$ and $a_i \in A_i$. Thus, v is enforceable under discount factor δ and all-or-nothing monitoring with q . ■

As a corollary, we can show that the rate of convergence under all-or-nothing monitoring is the same as that under perfect monitoring. More precisely, combined with Propositions 1 and 3, Lemma 7 implies that for any finite stage game under all-or-nothing monitoring with $q > 0$, there exists $K_q > 0$ such that $d(E(\delta, q), F_+) < K_q(1 - \delta)^{1/2}$, and generically the right-hand side improves to $K_q(1 - \delta)$. Imperfect monitoring in the form of all-or-nothing monitoring may affect the coefficient K_q , but does not affect the rate of convergence.²⁵

5 Discussion

Our paper has shown that even when imperfect monitoring has no cost in terms of the limit equilibrium payoff set, it has a cost in terms of the rate at which this set is approached.

²⁴This result extends to more general monitoring structures: the probability that each action profile be perfectly observed can depend on the action profile as long as it is positive, and signals in the complementary event need not be completely uninformative.

²⁵Note that imperfect monitoring does not improve the rate of convergence beyond perfect monitoring. This follows from Kandori (1992) if the public randomization device is available. Even without public randomization, one can easily extend Proposition 2 to *all* monitoring structures.

Specifically, we show that the set of equilibrium payoffs converges to its limit at rate $(1 - \delta)^{1/2}$ under perfect monitoring, and at rate $(1 - \delta)^{1/4}$ under imperfect monitoring, under the standard rank assumptions ensuring the folk theorem.

Throughout the paper, we have made a number of simplifying assumptions. Among them is the full dimensionality of F_+ . Even if the full dimensionality condition is violated, the set of equilibrium payoff vectors converges to some compact and convex set Q in the Hausdorff metric. Moreover, Q is characterized by finitely many steps of a version of Fudenberg-Levine's program, where continuation payoff vectors are restricted to lower-dimensional affine spaces (Fudenberg, Levine and Takahashi, 2007). In this case, if the last step of the program satisfies a counterpart of Lemma 2, then we obtain the same rate $(1 - \delta)^{1/4}$. See also the remark after the proof of Proposition 5.

In finite-horizon repeated games, Benoît and Krishna (1985) establish the pure-strategy folk theorem under perfect monitoring, the full dimensionality condition, and an additional condition that requires every player to have at least two stage-game Nash equilibrium payoffs. That is, any feasible and individually rational payoff vector can be approximately sustained in equilibrium as the horizon T becomes longer. We conjecture that the rate of convergence is $T^{-1/2}$. However, to show the mixed-strategy folk theorem (without assuming mixed action observability), Gossner (1995) uses a statistical test whose length becomes larger as the test becomes more precise. In his case, the rate of convergence may be $T^{-1/2}$ or strictly slower.²⁶ We have already referred in introduction to the literature on zero-sum games with incomplete information, in which the rates of convergence for the discounted case are obtained as corollaries of those for the finite-horizon case. In the context of Markov Decision Processes, White (1963) is the first to derive a rate of convergence as $T \rightarrow \infty$. A systematic analysis of the convergence rate in the finite-horizon case requires an investigation of the rate of convergence of the Bellman operator (see Hörner and Renault, 2014, for a derivation of the folk theorem based on the operator approach).

One limitation in our analysis of imperfect public monitoring, as already pointed out, is the restriction to public strategies. Kandori and Obara (2006) show that the equilibrium payoff set may expand strictly if we change the equilibrium concept from perfect public equilibrium to sequential equilibrium, thereby allowing for private strategies (although the difference in equilibrium concepts is irrelevant for product monitoring, as discussed in footnote 3 and Section 4.5). The ideas behind the proofs of Propositions 4 and 6 rely on perfect public equilibrium. It

²⁶Presumably, this question can be answered, given that the rate of convergence in approachability theory (on which he relies) has been thoroughly investigated, see Chen and White (1996).

is certainly possible that regarding imperfect monitoring, allowing equilibria in private strategies could accelerate the rate of convergence beyond the results that we have derived. Delayed punishments in the spirit of Compte (1998) might allow to economize on the cost of deterring deviations. This is left for future research.

Furthermore, Awaya and Krishna (2014) show that under private monitoring, communication helps expand the equilibrium payoff set. One may wonder whether communication could have an impact on the rate of convergence as well, even when the folk theorem holds with imperfect monitoring. However, to the extent that we use Tomala’s (2009) notion of *perfect communication equilibria*, which yields a recursive structure, all of our results survive with F_+ replaced by F_C , the set of feasible payoff vectors above the *correlated* minmax payoff vector. That is, the set of perfect communication equilibrium payoff vectors converges to F_C at rate $(1 - \delta)^{1/2}$ under perfect monitoring, and at rate $(1 - \delta)^{1/4}$ under imperfect (possibly private) monitoring. Moreover, these bounds are tight. On the other hand, Sugaya and Wolitzky (2015) allow for general sequential equilibria, and show that under mediated perfect monitoring, the set of sequential equilibrium payoff vectors converges to F_C at rate $1 - \delta$ (even without the standard full dimensionality condition).

Finally, we note that we have taken the classical “discrete-time” limit. That is, we have analyzed the limit of the equilibrium payoff set as $\delta \rightarrow 1$, keeping the monitoring structure fixed. Instead, one can take the “continuous-time” limit, where signals become less and less informative as $\delta \rightarrow 1$ (Abreu, Milgrom and Pearce, 1991, Sannikov, 2007, Sannikov and Skrzypacz, 2010). We are not aware of any result on the rate of convergence in the continuous-time limit, but some of the tools and proof techniques in this paper may be applicable.

References

- [1] Abreu, D., P. Milgrom and D. Pearce (1991). “Information and Timing in Repeated Partnerships,” *Econometrica*, **59**, 1713–1733.
- [2] Abreu, D., D. Pearce, and E. Stacchetti (1990). “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, **58**, 1041–1063.
- [3] Altman, E., and O. Zeitouni (1994). “Rate of Convergence of Empirical Measures and Costs in Controlled Markov Chains and Transient Optimality,” *Mathematics of Operations Research*, **19**, 955–974.

- [4] Aumann, R.J. and M. Maschler (1967). “Repeated games with incomplete information: A survey of recent results,” Reports to the U.S. Arms Control and Disarmament Agency, ST-116, Chapter III, 287–403.
- [5] Awaya, Y. and V. Krishna (2014). “On Tacit versus Explicit Collusion,” *mimeo*, Penn State University.
- [6] Azevedo, E. and H. Moreira (2007). “On the Impossibility of an Exact Imperfect Monitoring Folk Theorem,” *mimeo*, Fundação Getulio Vargas.
- [7] Barro, R.J. and X. Sala-i-Martin (1992). “Convergence,” *Journal of Political Economy*, **100**, 223–251.
- [8] Benoît, J.-P. and V. Krishna (1985). “Finitely repeated games,” *Econometrica*, **53**, 905–922.
- [9] Berg, K. and M. Kärki (2014). “How Patient the Players Need to Be to Obtain All the Relevant Payoffs in Discounted Supergames?,” *mimeo*.
- [10] Bewley, T. and E. Kohlberg (1976). “The Asymptotic Theory of Stochastic Games,” *Mathematics of Operations Research*, **1**, 197–208.
- [11] Chen, X. and H. White (1996). “Laws of Large Numbers for Hilbert Space-Valued Mixingales with Applications,” *Econometric Theory*, **12**, 284–304.
- [12] Compte, O. (1998). “Communication in Repeated Games with Imperfect Private Monitoring,” *Econometrica*, **66**, 597–626.
- [13] Fudenberg, D. and D. Levine (1994). “Efficiency and Observability with Long-Run and Short-Run Players,” *Journal of Economic Theory*, **62**, 103–135.
- [14] Fudenberg, D., D. Levine, and E. Maskin (1994). “The Folk Theorem with Imperfect Public Information,” *Econometrica*, **62**, 997–1040.
- [15] Fudenberg, D., D. Levine, and S. Takahashi (2007). “Perfect Public Equilibrium when Players are Patient,” *Games and Economic Behavior*, **61**, 27–49.
- [16] Fudenberg, D. and E. Maskin (1986). “The Folk Theorem for Repeated Games with Discounting or with Incomplete Information,” *Econometrica*, **54**, 533–554.

- [17] Fudenberg, D. and E. Maskin (1990). “Nash and Perfect Equilibria of Discounted Repeated Games,” *Journal of Economic Theory*, **51**, 194–206.
- [18] Fudenberg, D. and E. Maskin (1991). “On the dispensability of public randomization in discounted repeated games,” *Journal of Economic Theory*, **53**, 428–438.
- [19] Gossner, O. (1995). “The folk theorem for finitely repeated games with mixed strategies,” *International Journal of Game Theory*, **24**, 95–107.
- [20] 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**, 1277–1318.
- [21] Hörner, J. and J. Renault (2014). “A folk theorem for finitely repeated games with public monitoring,” mimeo, Toulouse School of Economics.
- [22] Kandori, M. (1992). “The use of information in repeated games with imperfect monitoring,” *Review of Economic Studies*, **59**, 581–594.
- [23] Kandori, M. and I. Obara (2006). “Efficiency in Repeated Games Revisited: The Role of Private Strategies,” *Econometrica*, **74**, 499–519.
- [24] Mailath, G.J., Obara, I. and T. Sekiguchi (2002). “The Maximum Efficient Equilibrium Payoff in the Repeated Prisoners’ Dilemma,” *Games and Economic Behavior*, **40**, 99–122.
- [25] Mertens, J.-F. (1998). “The speed of convergence in repeated games with incomplete information on one side,” *International Journal of Game Theory*, **27**, 343–357.
- [26] Ortigueira, S. and M.S. Santos (1997). “On the Speed of Convergence in Endogenous Growth Models,” *American Economic Review*, **87**, 383–399.
- [27] Sannikov, Y. (2007). “Games with imperfectly observable actions in continuous time,” *Econometrica*, **75**, 1285–1329.
- [28] Sannikov, Y. and A. Skrzypacz (2010). “The role of information in repeated games with frequent actions,” *Econometrica*, **78**, 847–882.
- [29] Sorin, S. (1986). “On Repeated Games with Complete Information,” *Mathematics of Operations Research*, **11**, 147–160.

- [30] Stahl, D.O. II (1991). “The graph of Prisoners’ Dilemma supergame payoffs as a function of the discount factor,” *Games and Economic Behavior*, **3**, 368–384.
- [31] Sugaya, T. and A. Wolitzky (2015). “On the Equilibrium Payoff Set in Repeated Games with Imperfect Private Monitoring,” mimeo, Stanford GSB.
- [32] Thomas, J.P. (1995). “Subgame perfect attainment of minimax punishments in discounted two-person games,” *Economics Letters*, **47**, 1–4.
- [33] Tomala, T. (2009). “Perfect Communication Equilibria in Repeated Games with Imperfect Monitoring,” *Games and Economic Behavior*, **67**, 682–694.
- [34] White, D.J. (1963). “Dynamic Programming Markov Chains and the Method of Successive Approximations,” *Journal of Mathematical Analysis and Applications*, **6**, 373–376.
- [35] Whitt, W. (1974). “Heavy Traffic Limit Theorems for Queues: A Survey,” Lecture Notes in Economics and Mathematical Systems, Vol. 98, Springer.
- [36] Yushkevich, A.A. (1996). “A Note on Asymptotics of Discounted Value Function and Strong 0-Discount Optimality,” *Mathematical Methods of Operations Research*, **44**, 223–231.
- [37] Zamir, S. (1971–72). “On the relation between finitely and infinitely repeated games with incomplete information,” *International Journal of Game Theory*, **1**, 179–198.

A Proof of Proposition 3

We denote player i ’s minmax payoff by \underline{v}_i . Various objects, such as \underline{v}_i , F_+ , and $E(\delta)$, depend on payoff function profile $u = (u_i)_{i=1}^n$, but for notational simplicity, we keep such dependence implicit.

For each $i \neq j$, let \mathcal{U}_{ij} be the set of payoff function profiles that satisfy the sufficient conditions stated in Lemma 1 for the rate of convergence to be at most $1 - \delta$, *i.e.*, if long minmaxing of player i is required, then for each $j \neq i$, either $v_j > \underline{v}_j$ or there exists $v' \in F_+$ such that $v'_i = v_i$ and $v'_j > v_j$. We will show that each \mathcal{U}_{ij} is open and dense in \mathcal{U}_0 , and so is $\mathcal{U} = \bigcap_{(i,j): i \neq j} \mathcal{U}_{ij}$ by the Baire category theorem. (Note that \mathcal{U}_0 is open in $\mathbf{R}^{n|A|}$, and hence the Baire category theorem applies to \mathcal{U}_0 .) Clearly, \mathcal{U}_{ij} is open. To show that \mathcal{U}_{ij} is dense in \mathcal{U}_0 , fix any $u \in \mathcal{U}_0 \setminus \mathcal{U}_{ij}$. Then there exists a vertex $v \in F_+$ where long minmaxing of player i is required (and hence

$v_i = \underline{v}_i$), $v_j = \underline{v}_j$, and there exists no $v' \in F_+$ such that $v'_i = v_i$ and $v'_j > v_j$. Then we have $\{(v'_i, v'_j) \in \mathbf{R}^2 : v' \in F_+, v'_i = v_i\} = \{(v_i, v_j)\}$.

Let $\underline{\alpha}^i$ be a minmax action profile of player i . Suppose that $\underline{\alpha}^i$ is pure. Then, since long minmaxing of player i is required at v , player i is indifferent among all actions against $\underline{\alpha}^i$. Since $|A_i| \geq 2$, this case is nowhere dense. Thus, we can assume that $\underline{\alpha}^i$ is not pure.

Suppose that $v_j > \underline{v}_j$. Since F_+ is full-dimensional, we have $\underline{v}_i = \min_{a \in A} u_i(a)$. Since $|A_i| \geq 2$, this case is nowhere dense. Thus, we can assume that $v_j = \underline{v}_j$ and hence (v_i, v_j) is the vector of minmax payoffs.

Since (v_i, v_j) is on the boundary of the projection of $\text{conv } u(A)$ on the ij -plane, we have either (i) $(v_i, v_j) = (u_i(\bar{a}), u_j(\bar{a}))$ for some $\bar{a} \in A$, or (ii) (v_i, v_j) is a strict convex combination of $(u_i(\hat{a}), u_j(\hat{a}))$ and $(u_i(\tilde{a}), u_j(\tilde{a}))$ for some $\hat{a}, \tilde{a} \in A$ with $\hat{a} \neq \tilde{a}$.

For Case (i), since $\underline{\alpha}^i$ is not pure, for each $a_i \in A_i$, there exists $a_{-i} \in \text{supp } \underline{\alpha}^i$ such that $(a_i, a_{-i}) \neq \bar{a}$. Consider $u' \in \mathcal{U}_0$ such that $u'_i(\bar{a}) = u_i(\bar{a})$ and $u'_i(a) < u_i(a)$ for any $a \neq \bar{a}$. Then, player i 's minmax payoff \underline{v}'_i under u'_i satisfies

$$\underline{v}'_i \leq \max_{a_i \in A_i} u'_i(a_i, \alpha^i_{-i}) < \max_{a_i \in A_i} u_i(a_i, \alpha^i_{-i}) = \underline{v}_i = u_i(\bar{a}) = u'_i(\bar{a}).$$

Thus, this case is nowhere dense.

For Case (ii), for each such u , there exists $\lambda \in (0, 1)$ such that

$$\begin{aligned} \underline{v}_i &= \lambda u_i(\hat{a}) + (1 - \lambda) u_i(\tilde{a}), \\ \underline{v}_j &= \lambda u_j(\hat{a}) + (1 - \lambda) u_j(\tilde{a}). \end{aligned}$$

Since \underline{v}_i is independent of u_{-i} , λ is also independent of u_{-i} generically (as long as $u_i(\hat{a}) \neq u_i(\tilde{a})$). Similarly, since \underline{v}_j is independent of u_{-j} , λ is also independent of u_{-j} generically. Thus, we can treat λ as a constant.

We assume without loss of generality that $\hat{a}_i \neq \tilde{a}_i$. Since α^i_{-i} is not pure and $\hat{a}_i \neq \tilde{a}_i$, for any $a_i \in A_i$, there exists $a_{-i} \in \text{supp } \alpha^i_{-i}$ such that $(a_i, a_{-i}) \neq \hat{a}, \tilde{a}$. Similarly to Case (i), consider $u'' \in \mathcal{U}_0$ such that $u''_i(\hat{a}) = u_i(\hat{a})$, $u''_i(\tilde{a}) = u_i(\tilde{a})$, and $u''_i(a) < u_i(a)$ for any $a \neq \hat{a}, \tilde{a}$. Then, player i 's minmax payoff \underline{v}''_i under u''_i satisfies

$$\underline{v}''_i \leq \max_{a_i \in A_i} u''_i(a_i, \alpha^i_{-i}) < \max_{a_i \in A_i} u_i(a_i, \alpha^i_{-i}) = \underline{v}_i = \lambda u_i(\hat{a}) + (1 - \lambda) u_i(\tilde{a}) = \lambda u''_i(\hat{a}) + (1 - \lambda) u''_i(\tilde{a}).$$

Thus, this case is also nowhere dense.