

Perfect Public Ex-Post Equilibria of Repeated Games with Uncertain Outcomes*

Drew Fudenberg[†] and Yuichi Yamamoto[‡]
Department of Economics, Harvard University[§]

First version: December 6, 2008

This version: May 2, 2009

Abstract

This paper studies repeated games with imperfect public monitoring where the players are uncertain both about the payoff functions and about the relationship between the distribution of signals and the actions played. To analyze these games, we introduce the concept of perfect public ex-post equilibrium (PPXE), and show that it can be characterized with an extension of the techniques used to study perfect public equilibria. We then develop identifiability conditions that are sufficient for a folk theorem; these conditions imply that there are PPXE in which the payoffs are approximately the same as if the monitoring structure and payoff functions were known.

Journal of Economic Literature Classification Numbers: C72, C73.

Keywords: repeated game, public monitoring, incomplete information, perfect public equilibrium, folk theorem, belief-free equilibrium, ex-post equilibrium.

*We thank Nageeb Ali, Susan Athey, Olivier Gossner, Michihiro Kandori, Bill Sandholm, Satoru Takahashi, Joel Watson, Thomas Wiseman, and seminar participants at École Polytechnique, Harvard/MIT, Hitotsubashi University, Kobe University, Northwestern, UCSD, University of Tokyo, University of Wisconsin-Madison, and 2009 Far East and South Asia Meeting of the Econometric Society for insightful comments, and NSF grant 0646816 for financial support.

[†]dfudenberg@harvard.edu

[‡]yamamot@fas.harvard.edu

[§]Littauer Center, Harvard University, 1805 Cambridge Street, Cambridge, MA 02138

1 Introduction

The role of repeated play in facilitating cooperation is one of the main themes of game theory. Past work has shown that reciprocation can lead to more cooperative equilibrium outcomes even if there is *imperfect public monitoring*, so that players do not directly observe their opponents' actions but instead observe noisy public signals whose distribution depends on the actions played. This work has covered a range of applications, from oligopoly pricing (e.g. Green and Porter (1984) and Athey and Bagwell (2001)), repeated partnerships (Radner, Myerson, and Maskin (1986)) and relational contracts (Levin (2003)). These applications are accompanied by a theoretical literature on the structure of the set of equilibrium payoffs and its characterization as the discount factor approaches 1, most notably Abreu, Pearce, and Stachetti (1986), Abreu, Pearce, and Stachetti (1990, hereafter APS), Fudenberg and Levine (1994, hereafter FL), Fudenberg, Levine, and Maskin (1994, hereafter FLM), and Fudenberg, Levine, and Takahashi (2007). All of these papers assume that the players know the distribution of public signals as a function of the actions played. In some cases this assumption seems too strong: For example, the players in a partnership may know that high effort makes good outcomes more likely, but not know the exact probability of a bad outcome when all agents work hard. This paper allows for such uncertainty, and also allows for uncertainty about the underlying payoff functions.

Specifically, we study repeated games in which the state of the world, chosen by Nature at the beginning of the play, influences the distribution of public signals and/or the payoff functions of the stage game. The effect of the state on the payoff functions can be direct, and can also be an indirect consequence of the effect of the state on the distribution of signals. For example, in a repeated partnership, the players will tend to have higher expected payoffs at a given action profile at states where high output is most likely, so even if the payoff to high output is known, uncertainty about the probability of high output leads to uncertainty about the expected payoffs of the stage game.

Because actions are imperfectly observed, the players' posterior beliefs need not coincide in later periods, even when they share a common prior on the distribution of states.¹ This complicates the verification of whether a given strategy

¹Cripps and Thomas (2003), Gossner and Vieille (2003), and Wiseman (2005) study

profile is an equilibrium, and thus makes it difficult to provide a characterization of the entire equilibrium set. Instead, we consider a subset of Nash equilibria, called *perfect public ex-post equilibria* or *PPXE*. A strategy profile is a PPXE if it is public- i.e. it depends only on publicly available information- and if its continuation strategy constitutes a Nash equilibrium given any state and given any history. In a PPXE, a player's best reply does not depend on her belief, so that the equilibrium set has a recursive structure and the analysis is greatly simplified. As with ex-post equilibrium, PPXE are robust to variations in beliefs about the underlying uncertainty- a PPXE for a given prior distribution is a PPXE for an arbitrary prior.²

Before launching into the general characterization of PPXE, we give a few examples to illustrate how these equilibria work. One important fact is that even though players start out not knowing the state, conditioning play on outcomes can indirectly allow the state to determine play and equilibrium payoffs. For example, if the outcome perfectly reveals the state, there can be PPXE where player 1's preferred PPE is played from period 2 on in state ω_1 and player 2's preferred PPE is played from period 2 on in state ω_2 . The first two examples consider special structures that make it easy to give explicit constructions of PPXE; the last two examples use our non-constructive characterization of the limits of PPXE payoffs. In particular, applying our linear programming characterization to Example 4 shows that payoffs can be bounded away from efficiency even though the distribution of outcomes can reveal the state and the folk theorem would hold in each state if the state were known.

To characterize the limit of the set of PPXE payoffs as the discount factor

symmetric-information settings where actions and payoffs are perfectly observed, so players always have the same beliefs, and this difficulty does not arise. In Aumann and Hart (1992), Aumann and Maschler (1995), Hörner and Lovo (2009), Wiseman (2008), and Hörner, Lovo, and Tomala (2008), players receive private signals about the payoff functions and so can have different beliefs. (In Wiseman (2008) the players privately observe their own realized payoff each period, in the other papers the players do not observe their own realized payoffs, and the private signals are the players' initial information or "type.")

²See Bergmann and Morris (2007) for a discussion of various definitions of ex-post equilibrium. Miller (2007) analyzes a different sort of ex-post equilibrium: he considers repeated games of adverse selection, where players report their types each period, as in Section 8 of FLM, and adds the restriction that announcing truthfully should be optimal regardless of the announcements of the other players.

goes to 1, we extend the linear programming characterization of the limit payoffs of PPE. That is, we show in Section 4 that the limit of the set of payoff vectors to PPXE as the discount factor goes to 1 is the intersection of the “maximal half-spaces” in various directions, where each component λ_i^ω of the direction vector λ corresponds to the weight attached to player i ’s payoff in state ω . The main new feature is that in a PPXE, the equilibrium payoffs are allowed to vary with the state, and can do so even if the state does not influence the expected payoffs to each action profile- for example there can be PPXE where player 1 does better in state ω_1 and player 2 does better in state ω_2 . Thus PPXE can involve a form of “utility transfer” across states. For this reason, the “maximal half space” in these “cross-state directions” can be the whole space, while in FL the maximal half space in each direction is bounded by the feasible set.

In Section 5, we use this characterization to prove an “ex-post” folk theorem: For any map from states to payoff vectors that are feasible and individually rational in that state, there is a PPXE whose payoffs in each state approximate the target map as the discount factor tends to 1. This theorem uses individual and pairwise full rank conditions as in FLM, and adds the assumption that for every pair (i, ω) and $(j, \tilde{\omega})$ of individuals and states, there is a profile α that has “state-wise full rank,” which means roughly that the observed signals reveal the state regardless of whether i or j (but not both!) unilaterally deviate from α .

As in FLM, a weaker, “static-threats,” version of the folk theorem holds under milder informational conditions. Section 6 shows that pairwise full rank can be replaced by the condition of “pairwise identifiability,” which can be satisfied with a smaller number of signals and that statewise full rank can be relaxed to “state-wise identifiability.” Both of these identifiability conditions are equivalent to their full-rank analogs when individual rank conditions are satisfied, but in general they are weaker and can be satisfied in models with fewer signals relative to the size of the action spaces. Even the statewise identifiability condition is stronger than needed, as shown in Section 6.2. In particular, when the individual full rank conditions are satisfied, statewise identifiability requires more signals than in FLM, but statewise distinguishability can be satisfied without a larger signal space. Very roughly speaking, the key is that for every pair of players i, j and pair of states $\omega, \tilde{\omega}$, there be a strategy profile whose signal distribution distinguishes between the two states regardless of the deviations of player j , and such that continuation

payoffs can give a large reward to player i in state ω without increasing player i 's incentive to deviate and without affecting player j 's payoff in state $\tilde{\omega}$.

While the study of uncertain monitoring structures is new, there is a substantial literature on repeated games with unknown payoff functions and perfectly observed actions, notably Aumann and Hart (1992), Aumann and Maschler (1995), Cripps and Thomas (2003), Gossner and Vieille (2003), Wiseman (2005), Hörner and Lovo (2009), Wiseman (2008), and Hörner, Lovo, and Tomala (2008).³ Our work makes two extensions to this literature- first to the case of unknown payoff functions and imperfectly observed actions but a known monitoring technology, and from there to the case where the monitoring structure is itself unknown.

PPXE is closely related to the “belief-free” equilibria used by Hörner and Lovo (2009) and Hörner, Lovo, and Tomala (2008). This equilibrium concept allows players to condition on their type, while PPXE does not, as it requires public strategies. Nevertheless, PPXE can be used to analyze incomplete-information games; here it amounts to a “pooling equilibrium” as players are not allowed to condition on their type. We say more about the comparison of these equilibrium concepts in Section 7. PPXE is also related to belief-free equilibria in repeated games with private monitoring, as in Piccione (2002), Ely and Välimäki (2002), Ely, Hörner, and Olszewski (2005), Yamamoto (2007), Kandori (2008), and Yamamoto (2009).⁴ However, unlike the belief-free equilibria in those papers, PPXE does not require that players be indifferent, and so it is not subject to the robustness critiques of Bhaskar, Mailath, and Morris (2008); this is what motivates our choice of a different name for the concept.

³Cripps and Thomas (2003), Gossner and Vieille (2003), and Wiseman (2005) study symmetric-information settings. In Aumann and Hart (1992), Aumann and Maschler (1995), Hörner and Lovo (2009), Wiseman (2008), and Hörner, Lovo, and Tomala (2008), players receive private signals about the payoff functions and so can have different beliefs. (In Wiseman (2008) the players privately observe their own realized payoff each period, in the other papers the players do not observe their own realized payoffs, and the private signals are the players' initial information or “type.”

⁴Belief-free equilibria and the use of indifference conditions have also been applied to repeated games with random matching (Takahashi (2008), Deb (2008)).

2 Unknown Signal Structure and Perfect Public Ex-Post Equilibria

2.1 Model

Let $I = \{1, \dots, I\}$ represent the set of players. At the beginning of the game, Nature chooses the state of the world ω from a finite set $\Omega = \{\omega_1, \dots, \omega_O\}$. Assume that players cannot observe the true state ω , and let $\mu \in \Delta\Omega$ denote the players' common prior over ω .⁵ For now we assume that the game begins with symmetric information: Each player's beliefs about ω correspond to the prior. We relax this assumption in Section 7.

Each period, players move simultaneously, and player $i \in I$ chooses an action a_i from a finite set A_i . Given an action profile $a = (a_i)_{i \in I} \in A \equiv \times_{i \in I} A_i$, players observe a public signal y from a finite set Y according to the probability function $\pi^\omega(a) \in \Delta Y$; we call the function π^ω the "monitoring technology." Player i 's realized payoff is $u_i^\omega(a_i, y)$, so that her expected payoff conditional on $\omega \in \Omega$ and on $a \in A$ is $g_i^\omega(a) = \sum_{y \in Y} \pi_y^\omega(a) u_i^\omega(a_i, y)$; $g^\omega(a)$ denotes the vector of expected payoffs associated with action profile a .

In the infinitely repeated game, players have a common discount factor $\delta \in (0, 1)$. Let (a_i^τ, y^τ) be the realized pure action and observed signal in period τ , and denote player i 's private history at the end of period $t \geq 1$ by $h_i^t = (a_i^\tau, y^\tau)_{\tau=1}^t$.⁶ Let $h_i^0 = \emptyset$, and for each $t \geq 1$, let H_i^t be the set of all h_i^t . Likewise, a public history up to period $t \geq 1$ is denoted by $h^t = (y^\tau)_{\tau=1}^t$, and H^t denotes the set of all h^t . A strategy for player i is defined to be a mapping $s_i : \bigcup_{t=0}^{\infty} H_i^t \rightarrow \Delta A_i$. Let S_i be the set of all strategies for player i , and let $S = \times_{i \in I} S_i$. Note that the case of a known public monitoring structure corresponds to a single possible state, $\Omega = \{\omega\}$.

⁵Because our arguments deal only with ex-post incentives, they extend to games without a common prior. However, as Dekel, Fudenberg, and Levine (2004) argue, the combination of equilibrium analysis and a non-common prior is hard to justify.

⁶As written, this formulation assumes that players do not observe their realized payoffs $u_i^\omega(a_i, y)$, unless the realized payoff does not depend on ω . Since we restrict attention to ex-post equilibria, where players' belief about the state do not matter, we do not need to impose this restriction, with the exception of Lemma 9, where the restriction is explicitly stated. If players observe the realized payoff, then player i 's private history after period t also includes $(u_i^\omega(a_i^\tau, y^\tau))_{\tau=1}^t$.

We define the set of feasible payoffs in a given state ω to be

$$V(\omega) \equiv \text{co}\{(g^\omega(a)) | a \in A\} = \{g^\omega(\eta) | \eta \in \Delta(A)\};$$

where $\Delta(A)$ is the set of all probability distributions over A : As in the standard case of a game with a known monitoring structure, the feasible set is both the set of feasible average discounted payoffs in the infinite-horizon game when players are sufficiently patient and the set of expected payoffs of the stage game that can be obtained when players use of a public randomizing device to implement distribution η over the action profiles.

Next we define the set of feasible payoffs of the overall game to be

$$V \equiv \times_{\omega \in \Omega} V(\omega),$$

so that a point $v \in V = (v^{\omega_1}, \dots, v^{\omega_o}) = ((v_1^{\omega_1}, \dots, v_I^{\omega_1}), \dots, (v_1^{\omega_o}, \dots, v_I^{\omega_o}))$.

Note that a given $v \in V$ may be generated using different action distributions $\eta(\omega)$ in each state ω . If players observe ω at the start of the game and are very patient, then any payoff in V can be obtained by a state-contingent strategy of the infinitely repeated game. Looking ahead, there will be equilibria that approximate payoffs in V if the state is *identified* by the signals, so that players learn it over time. Note also that, even if players have access to a public randomizing device, the set of feasible payoffs of the stage game is the smaller set

$$V^U = \{g^\omega(\eta) | \eta \in \Delta(A)\}_{\omega \in \Omega},$$

because play in the stage game must be a constant independent of ω .

2.2 Perfect Public Ex-Post Equilibria

This paper studies a special class of Nash equilibria called *perfect public ex-post equilibria* or PPXE; this is an extension of the concept of perfect public equilibrium that was introduced by FLM. Given a public strategy profile $s \in S$ and a public history $h^t \in H^t$, let $s|_{h^t}$ denote its continuation strategy profile after h^t .

Definition 1. A strategy $s_i \in S_i$ is *public* if it depends only on public information, i.e., for all $t \geq 1$, $h_i^t = (a_i^\tau, y^\tau)_{\tau=1}^t \in H_i^t$, and $\tilde{h}_i^t = (\tilde{a}_i^\tau, \tilde{y}^\tau)_{\tau=1}^t \in H_i^t$ satisfying $y^\tau = \tilde{y}^\tau$ for all $\tau \leq t$, $s_i(h_i^t) = s_i(\tilde{h}_i^t)$. A strategy profile $s \in S$ is *public* if s_i is public for all $i \in I$.

Definition 2. A strategy profile $s \in S$ is a *perfect public ex-post equilibrium* if for every $\omega \in \Omega$ the profile is a perfect public equilibrium of the game with known monitoring structure π^ω .⁷

Given a discount factor $\delta \in (0, 1)$, let $E(\delta)$ denote the set of PPXE payoffs, i.e., $E(\delta)$ is the set of all vectors $v = (v_i^\omega)_{(i,\omega) \in I \times \Omega} \in \mathbf{R}^{I \times |\Omega|}$ such that there is a PPXE $s \in S$ satisfying

$$(1 - \delta)E \left(\sum_{t=1}^{\infty} \delta^{t-1} g_i^\omega(a^t) \middle| s, \omega \right) = v_i^\omega$$

for all $i \in I$ and $\omega \in \Omega$. Note that $v \in E(\delta)$ specifies the equilibrium payoff for all players and for all possible states. Note also that the set of PPXE can be empty, in contrast to the case of perfect public equilibria of games with a known state.⁸ However, the conditions of our ex-post folk theorem imply that PPXE exist for sufficiently large discount factors.

The notion of minmax payoff extends to PPXE in a natural way. Let $\underline{v}_i^\omega = \min_{\alpha_{-i}} \max_{a_i} g_i^\omega(a_i, \alpha_{-i})$ be the minmax payoff for player i in state ω , and let

$$V^* \equiv \{v \in V \mid \forall i \in I, \forall \omega \in \Omega, v_i^\omega \geq \underline{v}_i^\omega\}$$

be the subset of the feasible payoff state where each player receives at least her minmax payoff in each state. Then $E(\delta) \subseteq V^*$, since any perfect public equilibrium of the game with known monitoring structure π must give each player i payoff at least \underline{v}_i^ω .

By definition, any continuation strategy of a PPXE is also a PPXE. Thus any PPXE specifies PPXE continuation play after each signal y , where the continuation payoffs $w(y) = (w_i^\omega(y))_{(i,\omega) \in I \times \Omega}$ corresponding to this signal specify payoffs for every player and every state. We will write $\pi^\omega(\alpha) \cdot w_i^\omega$ for player i 's expected

⁷That is, s is a public strategy, and for every $\omega \in \Omega$, and any public history $h^t \in H^t$, the continuation strategy profile $s|_{h^t}$ constitutes a Nash equilibrium of the ‘‘continuation game’’ corresponding to $\{h^t, \omega\}$. In this continuation game, players know that the state is ω , and because all opponents are using public strategies, each player can compute the expected payoff to any of their strategies (public or private) even though $\{h^t, \omega\}$ is not the root of a proper subgame.

⁸With a known state, repeated play of a static Nash equilibrium is a perfect public equilibrium of the repeated game. Similarly, repeated play of a static ex-post equilibrium is a PPXE, but static ex-post equilibria need not exist.

continuation payoff at state ω under action profile α , where w_i^ω is the vector $(w_i^\omega(y))_{y \in Y}$. This recursive structure of the equilibrium payoff set motivates the following definition.

Definition 3. For $\delta \in (0, 1)$ and $W \subseteq \mathbf{R}^{I \times |\Omega|}$, a pair $(\alpha, v) \in (\times_{i \in I} \Delta A_i) \times \mathbf{R}^{I \times |\Omega|}$ of an action profile and a payoff vector is *ex-post enforceable with respect to δ and W* if there is a function $w = (w^\omega)_{\omega \in \Omega} : Y \rightarrow W$ such that

$$v_i^\omega = (1 - \delta)g_i^\omega(\alpha) + \delta \pi^\omega(\alpha) \cdot w_i^\omega$$

for all $i \in I$ and $\omega \in \Omega$, and

$$v_i^\omega \geq (1 - \delta)g_i^\omega(a_i, \alpha_{-i}) + \delta \pi^\omega(a_i, \alpha_{-i}) \cdot w_i^\omega$$

for all $i \in I$, $\omega \in \Omega$, and $a_i \in A_i$.

For each $\delta \in (0, 1)$, $W \subseteq \mathbf{R}^{I \times |\Omega|}$, and $\alpha \in \times_{i \in I} \Delta A_i$, let $B(\delta, W, \alpha)$ denote the set of all payoff vectors $v \in \mathbf{R}^{I \times |\Omega|}$ such that (α, v) is ex-post enforceable with respect to δ and W . Let $B(\delta, W)$ be a union of $B(\delta, W, \alpha)$ over all $\alpha \in \times_{i \in I} \Delta A_i$.

To prove our main results, we will use the fact that various useful properties of PPE extend to PPXE.

Definition 4. A subset W of $\mathbf{R}^{I \times |\Omega|}$ is *ex-post self-generating with respect to δ* if $W \subseteq B(\delta, W)$.

Theorem 1. *If a subset W of $\mathbf{R}^{I \times |\Omega|}$ is bounded and ex-post self-generating with respect to δ , then $W \subseteq E(\delta)$.*

Proof. See Appendix. The proof is very similar to APS. The key is that when W is ex-post self-generating, the continuation payoffs $w(y)$ used to enforce $v \in V \subset \mathbf{R}^{I \times |\Omega|}$ have the property that for each $y \in Y$, the vector $w(y) \in \mathbf{R}^{I \times |\Omega|}$ can in turn be ex-post generated using a single next-period action α (independent of ω) so that the strategy profile constructed by “unpacking” the ex-post generation conditions does not directly depend on ω . *Q.E.D.*

Definition 5. A subset W of $\mathbf{R}^{I \times |\Omega|}$ is *locally ex-post generating* if for each $v \in W$, there exist $\delta_v \in (0, 1)$ and an open neighborhood U_v of v such that $W \cap U_v \subseteq B(\delta_v, W)$.

Theorem 2. *If a subset W of $\mathbf{R}^{I \times |\Omega|}$ is compact, convex, and locally ex-post generating, then there is $\bar{\delta} \in (0, 1)$ such that $W \subseteq E(\delta)$ for all $\delta \in (\bar{\delta}, 1)$.*

Proof. See Appendix; this is a straightforward generalization of FLM. *Q.E.D.*

3 Examples

Before proceeding with the general analysis, we present several examples to illustrate properties of PPXE. The first two examples make special assumptions that permit the explicit construction of PPXE strategies. The third and fourth examples are similar variants of a repeated partnership game. Here we use our linear programming bounds and folk theorem to show how incentive problems can lead to inefficiency even if there are action profiles that reveal the state and if the folk theorem would hold in each state if the state were known.

Example 1. There are two players, $I = \{1, 2\}$, and two possible states, $\Omega = \{\omega_1, \omega_2\}$. In every stage game, player 1 chooses an action from $A_1 = \{U, D\}$, while player 2 chooses an action from $A_2 = \{L, R\}$. Their expected payoffs $g_i^\omega(a)$ are as follows.

	<i>L</i>	<i>R</i>
<i>U</i>	2,2	0, 1
<i>D</i>	0,0	1, 1

	<i>L</i>	<i>R</i>
<i>U</i>	1,1	0, 0
<i>D</i>	1,0	2, 2

Here, the left table shows expected payoffs for state ω_1 , and the right table shows payoffs for state ω_2 . Suppose that the set of possible public signals is $Y = A \times \Omega$, and that the monitoring technology is such that $\pi_y^\omega(a) = \varepsilon > 0$ for $y \neq (a, \omega)$, and $\pi_y^\omega(a) = 1 - 7\varepsilon$ for $y = (a, \omega)$.

Note that (U, L) is a static Nash equilibrium for each state. Hence, playing (U, L) in every period is a PPXE, yielding the payoff vector $((2, 2), (1, 1))$. Likewise, playing (D, R) in every period is a PPXE, yielding the payoff vector $((1, 1), (2, 2))$. “Always (U, L) ” Pareto-dominates “always (D, R) ” for state ω_1 , but is dominated for state ω_2 . Note that these equilibrium payoff vectors are in the set V^U . Let $Y(\omega_1)$ be the set $\{y = (a, \omega) \in Y | \omega = \omega_1\}$, and $Y(\omega_2)$ be the set $\{y = (a, \omega) \in Y | \omega = \omega_2\}$. Consider the following strategy profile:

- In period one, play (U, L) .
- If $y \in Y(\omega_1)$ occurs in period one, then play (U, L) afterwards.
- If $y \in Y(\omega_2)$ occurs in period one, then play (D, R) afterwards.

After every one-period public history $h^1 \in H^1$, the continuation strategy profile is a PPXE. Also, given any state $\omega \in \Omega$, nobody wants to deviate in period one,

since (U, L) is a static Nash equilibrium and players cannot affect the distribution of the continuation play. Therefore, this strategy profile is a PPXE; its payoff vector converges to $v^* = ((2 - 4\varepsilon, 2 - 4\varepsilon), (2 - 4\varepsilon, 2 - 4\varepsilon))$ as $\delta \rightarrow 1$. Observe that $v^* \notin V^U$ if $\varepsilon \in (0, \frac{1}{8})$. In particular, this equilibrium approximates the efficient payoff vector $((2, 2), (2, 2))$ as the noise parameter ε goes to zero.

The idea of this construction is that continuation play depends on what players have learned about the state. When players observe $y \in Y(\omega_1)$ and learn that ω_1 is more likely, they choose “always (U, L) ,” which yields an efficient payoff $(2, 2)$ in state ω_1 , but gives an inefficient outcome $(1, 1)$ in ω_2 . Likewise, when players observe $y \in Y(\omega_2)$ and learn that ω_2 is more likely, they choose “always (D, R) ” to achieve an efficient payoff $(2, 2)$ in state ω_2 but an inefficient payoff in ω_1 . In this sense PPXE allows “utility transfers” across states.

Example 1 is misleadingly simple, because there is an ex-post equilibrium of the static game, and for this reason there is a PPXE for all discount factors. It is also very easy to construct equilibria that approximate efficient payoffs in this example: simply specify that (U, L) is played for T periods, and then either (U, L) or (D, R) is played forever afterwards, depending on which state is more likely. In the next example there is no static ex-post equilibrium, and hence no PPXE for a range of small discount factors, but where the folk theorem still applies.

Example 2. Suppose that there are two players, $I = \{1, 2\}$, and two states, $\Omega = \{\omega_1, \omega_2\}$. In every stage game, player 1 chooses U or D while player 2 chooses L or R . Players’ expected payoffs $g_i^\omega(a)$ are as follows.

ω_1	L	R
U	10, -4	1, 1
D	1, 1	0, 0

ω_2	L	R
U	0, 0	1, 1
D	1, 1	10, -4

Note that these are the payoff matrices in Example 2 of Hörner and Lovo (2009). Note also that the minimax payoffs $((v_1^{\omega_1}, v_2^{\omega_1}), (v_1^{\omega_2}, v_2^{\omega_2}))$ are $((1, \frac{1}{6}), (1, \frac{1}{6}))$, and that the set V^* has non-empty interior: $\dim V^* = \dim V^*(\omega_1) + \dim V^*(\omega_2) = 2 + 2 = 4$, where $V^*(\omega)$ is the set of feasible and individually rational payoffs given a state $\omega \in \Omega$.

This stage game does not have a static ex-post equilibrium, so we know that regardless of the monitoring structure it does not have a PPXE for a range of discount factors near 0.

We will now consider two repeated games with these payoff matrices and differing monitoring structures.

Example 2a. First suppose that actions are observable, but states (and rewards) are not, so that $Y = A$ and $\pi_y^\omega(a) = 1$ if $y = a$. In this case, it is easy to adapt the argument of Hörner and Lovo (2009) to show that there is no PPXE. In a PPXE, a player's equilibrium payoff conditional on ω cannot fall below the minimax payoff for that ω . In particular, player 2's equilibrium payoff conditional on ω_1 must be positive. This implies that the outcome $(10, -4)$ realizes at most the fifth of the time, and hence player 1's equilibrium payoff conditional on ω_1 is at most $\frac{14}{5}$. Likewise, player 1's equilibrium payoff conditional on ω_2 is at most $\frac{14}{5}$. However, if player 1 randomizes $(.5U, .5D)$ independent of the state, she earns at least 3 in one of the states. Therefore, there is no PPXE for any discount factor.

Note that with this monitoring structure, the outcome gives no information at all about the state, so it is impossible for the players to learn it.

Example 2b. Now suppose that the set of possible public signals is $Y = A \times \Omega$, and that the monitoring technology is perfect: $\pi_y^\omega(a) = 1$ if $y = (a, \omega)$, and $\pi_y^\omega(a) = 0$ otherwise. As we will see, this example satisfies all of the full-rank conditions of our general ex-post folk theorem, so in particular a PPXE exists, but our proof of the general folk theorem is not constructive. Because this example has perfect monitoring, it is easy to give an explicit construction of a PPXE whose payoffs converge to the efficient frontier in each state. Claim 7 in the appendix shows that the idea of this construction extends to any case where actions and states are perfectly observed; the idea is to wait one period, learn the state, and play a subgame-perfect equilibrium for the corresponding known-state game. However, the strategies used in the construction need to be a bit more complicated. In particular, as in Example 2b, the stage game need not have a static ex-post equilibrium, so the subgame-perfect equilibrium strategies for the known states must provide future incentives to prevent current deviations, and the recursive nature of PPXE requires that these future incentives correspond to equilibrium play in each state, including states that past signals have ruled out. (Note that in Example 1, there is no need for intertemporal incentives, so the PPXE we constructed there could prescribe play from period 2 on that is constant regardless of future signals about the state.)

Section 6.2 below presents a detailed analysis of two variants of a two-player partnership game; we summarize the main findings here to help preview and motivate our results. In these partnership games, each player i has the two actions $\{C_i, D_i\}$, and the stage game payoffs in each state make D_i a dominant strategy, so (D_1, D_2) is a static ex-post equilibrium. There are three possible outcomes, and two states, corresponding to differences in the productivity of effort, i.e. the probabilities of H and M if both players choose D are independent of the state. Finally, in each state the monitoring structure is additive: the change in probabilities induced by player i 's changing from C_i to D_i is the same regardless of the action of the other player.

In Example 3, the uncertainty is symmetric in the state and across players: In state ω_1 , if player 1 chooses C_1 instead of D_1 , then the probabilities of H and M increase by p_H and p_M , while player 2's choice of C_2 increases the probabilities by q_H and q_M . In state ω_2 , the roles are reversed: player 1's effort increase the probabilities by q_H and q_M , and player 2's effort increases the probabilities by p_H and p_M . In Example 4, the states only influences the productivity of player 2's effort: If player 1 chooses C_1 instead of D_1 , then the probabilities of H and M increase by p_H and p_M , independent of the state. In contrast, if player 2 chooses C_2 instead of D_2 , then the probabilities of H and M increase by q_H and q_M in state ω_1 , but they increase only by βq_H and βq_M in state ω_2 .

In each example, the conditions of FLM's Theorem 6.1 apply in each state considered in isolation, so if the state were known the folk theorem would apply. Moreover in each example there are action profiles that reveal the state, in the sense that the outcome distribution at that profile is different at state ω_1 than at state ω_2 . However, we will show that our ex-post-threats folk theorem applies to Example 3, while in Example 4 the folk theorem fails, and moreover PPXE payoffs can be bounded away from efficiency.

The key difference between these examples is that in Example 4, the two states are "entangled" in the sense that for any α_1 , the distribution $\pi^{\omega_2}(\alpha_1, C_2)$ is a convex combination of $\pi^{\omega_1}(\alpha_1, C_2)$ and $\pi^{\omega_2}(\alpha_1, D_2)$, while this is not the case in Example 3 provided that α_1 assigns positive probability to both actions. Hence in Example 4, lowering the expected value of the continuation payoffs under $\pi^{\omega_2}(\alpha_1, D_2)$ also lowers the continuation payoffs under $\pi^{\omega_2}(\alpha_1, C_2)$. Using this fact and our linear programming characterization. we can show that the sum

of player 1's PPXE payoff in state ω_1 and player 2's PPXE payoff in state ω_2 is strictly less than the sum $g_1^{\omega_1}(C_1, C_2) + g_2^{\omega_2}(C_1, C_2) - e$ where the efficiency gap $e > 0$ depends on β and the parameters of the payoff matrix. Because no player's payoff in either state can be less than the corresponding minmax level, this cross-state bound implies that for some parameter values player 2's PPXE payoff in state ω_2 is strictly less than $g_2^{\omega_2}(C_1, C_2)$.

4 Characterizing $E(\delta)$

4.1 Using Linear Programming to Bound $E(\delta)$

In this subsection, we provide a bound on the set of PPXE payoffs that holds for any discount factor; the next subsection shows that this bound is tight as the discount factor converges to one.

Consider the following linear programming problem. Let $\alpha \in \times_{i \in I} \Delta A_i$, $\lambda \in \mathbf{R}^{I \times |\Omega|}$, and $\delta \in (0, 1)$.

$$\begin{aligned}
 \text{(LP-Average)} \quad k^*(\alpha, \lambda, \delta) = & \max_{\substack{v \in \mathbf{R}^{I \times |\Omega|} \\ w: Y \rightarrow \mathbf{R}^{I \times |\Omega|}}} \lambda \cdot v \quad \text{subject to} \\
 \text{(i)} \quad v_i^\omega = & (1 - \delta)g_i^\omega(\alpha) + \delta\pi^\omega(\alpha) \cdot w_i^\omega \\
 & \text{for all } i \in I \text{ and } \omega \in \Omega, \\
 \text{(ii)} \quad v_i^\omega \geq & (1 - \delta)g_i^\omega(a_i, \alpha_{-i}) + \delta\pi^\omega(a_i, \alpha_{-i}) \cdot w_i^\omega \\
 & \text{for all } i \in I, \omega \in \Omega, \text{ and } a_i \in A_i, \\
 \text{(iii)} \quad \lambda \cdot v \geq & \lambda \cdot w(y) \quad \text{for all } y \in Y.
 \end{aligned}$$

If there is no (v, w) satisfying the constraints, let $k^*(\alpha, \lambda, \delta) = -\infty$. If for every $K > 0$ there is (v, w) satisfying all the constraints and $\lambda \cdot v > K$, then let $k^*(\alpha, \lambda, \delta) = \infty$.

Here condition (i) is the ‘‘adding-up’’ condition, condition (ii) is ex-post incentive compatibility, and condition (iii) requires that the continuation payoffs lie in half-space corresponding to direction vector λ and payoff vector v . Note that when $\lambda_i^\omega \neq 0$ and $\lambda_j^{\tilde{\omega}} \neq 0$ for some $\omega \neq \tilde{\omega}$, condition (iii) allows ‘‘utility transfer’’ across states. This utility transfer is the most significant way that LP-average differs from the linear program in FL, so we will discuss it in more detail below.

As we show in Lemma 1(a), the value $k^*(\alpha, \lambda, \delta)$ is independent of δ , so that we denote it by $k^*(\alpha, \lambda)$. Now let

$$k^*(\lambda) = \sup_{\alpha} k(\alpha, \lambda)$$

be the highest score that can be approximated in direction λ by any choice of α .

For each $\lambda \in \mathbf{R}^{I \times |\Omega|} \setminus \{0\}$ and $k \in \mathbf{R}$, let $H(\lambda, k) = \{v \in \mathbf{R}^{I \times |\Omega|} \mid \lambda \cdot v \leq k\}$. For $k = \infty$ or $\lambda = 0$, let $H(\lambda, k) = \mathbf{R}^{I \times |\Omega|}$. For $k = -\infty$ and $\lambda \neq 0$, let $H(\lambda, k) = \emptyset$. Then, let

$$H^*(\lambda) = H(\lambda, k^*(\lambda))$$

be the maximal half-space in direction λ , and set

$$Q = \bigcap_{\lambda \in \mathbf{R}^{I \times |\Omega|}} H^*(\lambda).$$

Lemma 1. (a) $k^*(\alpha, \lambda, \delta)$ is independent of δ .

(b) If $(\lambda_i^\omega)_{i \in I} \neq 0$ for some ω and $(\lambda_i^{\tilde{\omega}})_{i \in I} = 0$ for all $\tilde{\omega} \neq \omega$, then $k^*(\lambda) \leq \sup_{\alpha} \lambda \cdot g(\alpha)$.

(c) If $\lambda_i^\omega < 0$ for some (i, ω) and $\lambda_j^{\tilde{\omega}} = 0$ for all $(j, \tilde{\omega}) \neq (i, \omega)$ then $k^*(\lambda) \leq \lambda_i^\omega y_i^\omega$.

(d) Consequently $Q \subseteq V^*$.

Proof. As in past work, part (a) follows from the fact that the constraint set in (iii) is a half-space: Suppose that (v, w) satisfies constraints (i) through (iii) in LP-Average for $(\alpha, \lambda, \delta)$. For $\tilde{\delta} \in (0, 1)$, let

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

Then (v, \tilde{w}) satisfies constraints (i) through (iii) in LP-Average for $(\alpha, \lambda, \tilde{\delta})$, so that the set of feasible v in LP-Average is independent of δ , and thus so is $k^*(\alpha, \lambda, \delta)$.

Let Λ^* be the set of $\lambda \in \mathbf{R}^{I \times |\Omega|}$ such that $(\lambda_i^\omega)_{i \in I} \neq 0$ for some $\omega \in \Omega$ and $(\lambda_i^{\tilde{\omega}})_{i \in I} = 0$ for all $\tilde{\omega} \neq \omega$. Since parts (b) and (c) consider a single state ω they follow from FL Lemma 3.1. Thus $\bigcap_{\lambda \in \Lambda^*} H^*(\lambda) \subseteq V^*$, and part (d) follows from $Q \subseteq \bigcap_{\lambda \in \Lambda^*} H^*(\lambda)$. *Q.E.D.*

Since we already know that $E(\delta) \subseteq V^*$, part (d) of this lemma shows that Q is “not too big”: it doesn’t contain any payoff vector we can rule out on *a priori* grounds. The next lemma shows that Q is “big enough” to contain all the payoffs of PPXE.

Lemma 2. *For every $\delta \in (0, 1)$, $E(\delta) \subseteq E^*(\delta) \subseteq Q$, where $E^*(\delta)$ is the convex hull of $E(\delta)$.*

Proof. The proof is the same as in Fudenberg, Levine, and Takahashi (2007); we restate it in the Appendix to make it easy to see that the proof applies to the present setting. *Q.E.D.*

To help explain the role of cross-state utility transfers, we will show that the conclusion of Lemma 2 does not hold if constraint (iii) is replaced by the uniform-over-states version

$$(iii') \quad \sum_{i \in I} \lambda_i^\omega v_i^\omega \geq \sum_{i \in I} \lambda_i^\omega w_i^\omega(y) \quad \text{for all } \omega \in \Omega \text{ and } y \in Y.$$

The resulting “uniform” LP problem corresponds to a form of ex-post enforceability on half-spaces. This condition is too restrictive to capture all of the payoffs of PPXE, as shown by the combination of the following claim and the example that follows it.

Claim 1. *In the LP-Uniform problem formed by replacing (iii) in LP-Average with (iii'), the solution $k^U(\alpha, \lambda, \delta) \leq \lambda \cdot g(\alpha)$ for each α and λ . Therefore, $k^U(\lambda, \delta) \equiv \sup_\alpha k^U(\alpha, \lambda, \delta) \leq \sup_\alpha \lambda \cdot g(\alpha)$, and the computed set Q^U is a subset of payoffs V^U that can be attained with actions that are independent of the state.*

Proof. Inspection of the constraints in the LP-Uniform problem shows that it is equivalent to solving a separate LP problem for each state $\omega \in \Omega$ in isolation. As FL show, a solution to the LP problem for given (α, ω) cannot exceed $\sum_{i \in I} \lambda_i^\omega g_i^\omega(\alpha)$. Therefore, $k^U(\alpha, \lambda, \delta)$, the maximal score in LP-Uniform for a given α , is at most $\sum_{\omega \in \Omega} \sum_{i \in I} \lambda_i^\omega g_i^\omega(\alpha) = \lambda \cdot g(\alpha)$, so $\sup_\alpha k^U(\alpha, \lambda, \delta) \leq \sup_\alpha \lambda \cdot g(\alpha)$. *Q.E.D.*

In both Examples 1 and 2b, we constructed PPXE with payoffs outside of V^U .

4.2 Computing the Limit of $E(\delta)$ as Players Become Patient

Now we show that the set $E(\delta)$ of PPXE payoffs expands to equal all of Q as the players become sufficiently patient, provided that a full-dimensionality condition is satisfied. For each set B , let $\text{int}B$ denote the interior of B , and $\text{bd}B$ denote the boundary of B .

Definition 6. A subset W of $\mathbf{R}^{I \times |\Omega|}$ is *smooth* if it is closed and convex; it has a nonempty interior; and there is a unique unit normal for each point on $\text{bd}W$.⁹

Lemma 3. *If $\dim Q = I \times |\Omega|$, then for any smooth strict subset W of Q , there is $\bar{\delta} \in (0, 1)$ such that $W \subseteq E(\delta)$ for $\delta \in (\bar{\delta}, 1)$.*

Proof. From lemma 1(d), Q is bounded, and hence W is also bounded. Then, from Theorem 2, it suffices to show that W is locally ex-post generating, i.e., for each $v \in W$, there exist $\delta_v \in (0, 1)$ and an open neighborhood U_v of v such that $W \cap U_v \subseteq B(\delta_v, W)$.

First, consider $v \in \text{bd}W$. Let λ be normal to W at v , and let $k = \lambda \cdot v$. Since $W \subset Q \subseteq H^*(\lambda)$, there exist α, \tilde{v} , and $(\tilde{w}(y))_{y \in Y}$ such that $\lambda \cdot \tilde{v} > \lambda \cdot v = k$, (α, \tilde{v}) is enforced using continuation payoffs $(\tilde{w}(y))_{y \in Y}$ for some $\tilde{\delta} \in (0, 1)$, and $\tilde{w}(y) \in H(\lambda, \lambda \cdot \tilde{v})$ for all $y \in Y$. For each $\delta \in (\tilde{\delta}, 1)$ and $y \in Y$, let

$$w(y, \delta) = \frac{\delta - \tilde{\delta}}{\delta(1 - \tilde{\delta})}v + \frac{\tilde{\delta}(1 - \delta)}{\delta(1 - \tilde{\delta})} \left(\tilde{w}(y) - \frac{v - \tilde{v}}{\tilde{\delta}} \right).$$

By construction, (α, v) is enforced by $(w(y, \delta))_{y \in Y}$ for δ , and there is $\kappa > 0$ such that $|w(y, \delta) - v| < \kappa(1 - \delta)$. Also, since $\lambda \cdot \tilde{v} > \lambda \cdot v = k$ and $\tilde{w}(y) \in H(\lambda, \lambda \cdot \tilde{v})$ for all $y \in Y$, there is $\varepsilon > 0$ such that $\tilde{w}(y) - \frac{v - \tilde{v}}{\tilde{\delta}}$ is in $H(\lambda, k - \varepsilon)$ for all $y \in Y$, thereby

$$w(y, \delta) \in H \left(\lambda, k - \frac{\tilde{\delta}(1 - \delta)}{\delta(1 - \tilde{\delta})} \varepsilon \right)$$

for all $y \in Y$. Then, as in the proof of FL's Theorem 3.1, it follows from the smoothness of W that $w(y, \delta) \in \text{int}W$ for sufficiently large δ , i.e., (α, v) is enforced with respect to $\text{int}W$. To enforce u in the neighborhood of v , use α and a translate of $(w(y, \delta))_{y \in Y}$.

⁹A sufficient condition for each point on $\text{bd}W$ to have a unique unit normal is that $\text{bd}W$ is a C^2 -submanifold of $\mathbf{R}^{I \times |\Omega|}$.

Next, consider $v \in \text{int}W$. Choose λ arbitrarily, and let α and $(w(y, \delta))_{y \in Y}$ be as in the above argument. By construction, (α, v) is enforced by $(w(y, \delta))_{y \in Y}$. Also, $w(y, \delta) \in \text{int}W$ for sufficiently large δ , since $|w(y, \delta) - v| < \kappa(1 - \delta)$ for some $\kappa > 0$ and $v \in \text{int}W$. Thus, (α, v) is enforced with respect to $\text{int}W$ when δ is close to one. To enforce u in the neighborhood of v , use α and a translate of $(w(y, \delta))_{y \in Y}$, as before. *Q.E.D.*

These two lemmas establish the following theorem.

Theorem 3. *If $\dim Q = I \times |\Omega|$, then $\lim_{\delta \rightarrow 1} E(\delta) = Q$.*

It is possible that $\dim Q < I \times |\Omega|$, so that this theorem does not apply, but that $\lim_{\delta \rightarrow 1} E(\delta) \neq \emptyset$. A trivial example of this occurs when the state ω has no effect on either the monitoring structure or the payoffs, so that it cannot possibly be observed, but is simply a nuisance parameter. In this $E(\delta)$ is a subset of the space V^U of payoff that can be generated with actions that are independent of the state, so $Q \subseteq E(\delta)$ has dimension at most I . In this particular case, the solution is obviously to ignore the state and characterize the perfect public equilibria of the game where (any) ω is known; these equilibria correspond to the full set of PPXE of the game with the noise parameter added. More generally, the full-dimension conditions could fail due to the imperfect observability of ω , but ω might matter for the payoff functions. In this case one might be able to characterize $\lim_{\delta \rightarrow 1} E(\delta)$ using an extension of the iterative algorithm in Fudenberg, Levine, and Takahashi (2007), but this remains a topic for future research.

5 A Perfect Ex-Post Folk Theorem

In this section we give simple and easy-to verify sufficient conditions for a folk theorem to hold in PPXE. This theorem shows that any map from states of the world to payoffs that are feasible and individually rational in that state can be approximated by equilibrium payoffs as the discount factor goes to 1, and in particular by payoffs of a PPXE. More formally, our folk theorem gives conditions under which $\lim_{\delta \rightarrow 1} E(\delta) = V^*$.¹⁰ When this is true, so that efficient payoffs can be approximated by PPXE, the players have little reason to play other sorts of

¹⁰Recall that $V^* \equiv \{v \in V \mid \forall i \in I, \forall \omega \in \Omega, v_i^\omega \geq v_i^\omega\}$.

equilibria or to try to change the monitoring structure. Conversely, when the set of PPXE is empty, or when all PPXE are far from efficient but there are efficient sequential equilibria, the PPXE restriction might be less compelling.

Since we have already shown that $Q \subseteq V^*$ and that $\lim_{\delta \rightarrow 1} E(\delta) = Q$ under the full-dimension condition, it remains to show that $V^* \subseteq Q$, which is equivalent to showing that $k^*(\lambda) \geq \max_{v \in V^*} \lambda \cdot v$ for each direction λ . Our sufficient conditions are actually stronger than that: they will imply that $k^*(\lambda) = \infty$ for directions λ with non-zero components in two or more states. Conversely, the folk theorem fails if there is a λ such that $k^*(\lambda) < \max_{v \in V^*} \lambda \cdot v$; we use this fact in Example 4 below.

For each $i \in \mathbf{I}$, $\alpha \in \times_{i \in \mathbf{I}} \Delta A_i$, and $\omega \in \Omega$, let $\Pi_{(i, \omega)}(\alpha)$ be a matrix with rows $(\pi_y^\omega(a_i, \alpha_{-i}))_{y \in Y}$ for all $a_i \in A_i$.

Definition 7. Profile α has *individual full rank for* (i, ω) if $\Pi_{(i, \omega)}(\alpha)$ has rank equal to $|A_i|$. Profile α has *individual full rank* if it has individual full rank for all players and all states.

Individual full rank implies that at each state, every possible deviation of any one player leads to a statistically different distribution on outcomes; on this condition there are continuation payoffs that make every player indifferent between all actions. However, as we discuss in Section 6.2, many of our results hold under weaker but harder-to-verify conditions.

Let $\Pi_{(i, \omega)(j, \tilde{\omega})}(\alpha)$ be a matrix constructed by stacking matrices $\Pi_{(i, \omega)}(\alpha)$ and $\Pi_{(j, \tilde{\omega})}(\alpha)$.

Definition 8. For each $i \in \mathbf{I}$, $j \neq i$, and $\omega \in \Omega$, profile α has *pairwise full rank for* (i, ω) and (j, ω) if $\Pi_{(i, \omega)(j, \omega)}(\alpha)$ has rank equal to $|A_i| + |A_j| - 1$.

Pairwise full rank implies that deviations by player i can be distinguished from deviations by j .

Definition 9. For each $i \in \mathbf{I}$, $j \in \mathbf{I}$, $\omega \in \Omega$, and $\tilde{\omega} \neq \omega$, profile α has *statewise full rank for* (i, ω) and $(j, \tilde{\omega})$ if $\Pi_{(i, \omega)(j, \tilde{\omega})}(\alpha)$ has rank equal to $|A_i| + |A_j|$.

Note that both pairwise full rank and statewise full rank imply individual full rank. Note also that the pairwise full rank conditions require as many signals as in FLM, and the statewise full rank conditions require at most twice as many signals.

(Statewise full rank requires only one more signal than FLM if all players have the same number of actions; it requires twice as many signals if one player has more than two actions and all the other players have only two.)

The statewise full rank condition guarantees that the observed signals will reveal the state, regardless of the play of player i in state ω and the play of player j (possibly equal to i) in state $\tilde{\omega}$, assuming that everyone else plays according to α . This condition is more restrictive than necessary for the existence of a strategy that allows the players to learn the state: For that it would suffice that there be a single profile α where the distributions on signals are all distinct, which requires only two signals.¹¹ On the other hand, the condition is less restrictive than the requirement that the state is revealed to an outside observer even if a pair of players deviates. For example, statewise full rank is consistent with a signal structure where a joint deviation by players 1 and 2 could conceal the state from the outside observer, as in a two-player game with $A_1 = A_2 = \{L, R\}$ and $\pi_y^\omega(L, R) = \pi_y^{\tilde{\omega}}(R, L)$. Intuitively, since equilibrium conditions only test for unilateral deviations, the statewise full rank condition is sufficient for the existence of an equilibrium where the players eventually learn the state. In Section 6.2, we introduce the more complicated but substantially weaker condition of statewise distinguishability, and show that it is sufficient for a static-threat version of the folk theorem.

The following is an ex-post folk theorem. Note that the set of assumptions of this theorem is generically satisfied if $|Y| \geq 2|A_i|$ for all $i \in \mathbf{I}$.

Condition IFR. Every pure action profile has individual full rank.

Condition PFR. For each (i, ω) and (j, ω) satisfying $i \neq j$, there is an action profile α that has pairwise full rank for (i, ω) and (j, ω) .

Condition SFR. For each (i, ω) and $(j, \tilde{\omega})$ satisfying $\omega \neq \tilde{\omega}$, there is an action profile α that has statewise full rank.

Theorem 4. *Suppose that (IFR), (PFR), and (SFR) hold. Then, for any smooth strict subset W of V^* , there is $\bar{\delta} \in (0, 1)$ such that $W \subseteq E(\delta)$ for all $\delta \in (\bar{\delta}, 1)$.*

The following lemmas are useful in this proof.

¹¹Note that players only need to distinguish between a finite set of signal distributions, and not between all possible convex combinations of them.

Lemma 4. *Suppose that (PFR) holds. Then, there is an open and dense set of profiles each of which has pairwise full rank for all (i, ω) and (j, ω) satisfying $i \neq j$.*

Proof. Analogous to that of Lemma 6.2 of FLM. *Q.E.D.*

Lemma 5. *Suppose that (IFR) holds. Then, for any $i \in \mathbf{I}$, $\omega \in \Omega$, and $\varepsilon > 0$, there is a profile $\underline{\alpha}^\omega$ such that $\underline{\alpha}_i^\omega \in \arg \max_{\alpha_i} g_i^\omega(\alpha_i, \underline{\alpha}_{-i}^\omega)$; $|g_i^\omega(\underline{\alpha}^\omega) - v_i^\omega| < \varepsilon$; and $\underline{\alpha}^\omega$ has individual full rank for all $(j, \tilde{\omega}) \neq (i, \omega)$.*

Proof. Analogous to that of Lemma 6.3 of FLM. *Q.E.D.*

Lemma 6. *Suppose that a profile α has statewise full rank for (i, ω) and $(j, \tilde{\omega})$ satisfying $\omega \neq \tilde{\omega}$ and that α has individual full rank for all players and states. Then, $k^*(\alpha, \lambda) = \infty$ for direction λ such that $\lambda_i^\omega \neq 0$ and $\lambda_j^{\tilde{\omega}} \neq 0$.*

Remark 1. Because $k^*(\alpha, \lambda) \leq \lambda \cdot g(\alpha)$ in the known-monitoring-structure case of FL, this lemma shows a key difference between that setting and the uncertain monitoring structure case we consider here. The idea is that under statewise full rank, the continuation payoffs in such half-spaces can give player i a very large payoff in state ω by giving player j a very low payoff in that state, while reversing this transfer in state $\tilde{\omega}$.

Remark 2. The proof of this lemma is complicated, so we illustrate it here with a simple example. Assume $A_i = \{a'_i, a''_i\}$ and $A_j = \{a'_j, a''_j\}$, and consider LP-Average problem for direction λ such that $\lambda_i^\omega = \lambda_j^{\tilde{\omega}} = 1$ and all other components of λ are zero. Constraints (i) and (ii) for $(l, \bar{\omega}) \in \mathbf{I} \times \Omega \setminus \{(i, \omega), (j, \tilde{\omega})\}$ can be satisfied by some choice of $(w_l^{\bar{\omega}}(y))_{y \in Y}$ because of individual full rank, and constraint (iii) is vacuous for these coordinates. So the LP problem reduces to finding $(w_i^\omega(y))_{y \in Y}$ and $(w_j^{\tilde{\omega}}(y))_{y \in Y}$ to solve

$$k^*(\alpha, \lambda, \delta) = \max_{v, w} v_i^\omega + v_j^{\tilde{\omega}}$$

subject to

$$\begin{aligned}
v_i^\omega &= (1 - \delta)g_i^\omega(\alpha) + \delta\pi^\omega(\alpha) \cdot w_i^\omega, \\
v_j^{\tilde{\omega}} &= (1 - \delta)g_j^{\tilde{\omega}}(\alpha) + \delta\pi^{\tilde{\omega}}(\alpha) \cdot w_j^{\tilde{\omega}}, \\
v_i^\omega &\geq (1 - \delta)g_i^\omega(a_i, \alpha_{-i}) + \delta\pi^\omega(a_i, \alpha_{-i}) \cdot w_i^\omega, \quad \forall a_i \in A_i \\
v_j^{\tilde{\omega}} &\geq (1 - \delta)g_j^{\tilde{\omega}}(a_j, \alpha_{-j}) + \delta\pi^{\tilde{\omega}}(a_j, \alpha_{-j}) \cdot w_j^{\tilde{\omega}}, \quad \forall a_j \in A_j \\
v_i^\omega + v_j^{\tilde{\omega}} &\geq w_i^\omega(y) + w_j^{\tilde{\omega}}(y), \quad \forall y \in Y.
\end{aligned}$$

We claim that $k^*(\alpha, \lambda, \delta) = \infty$ if α has statewise full rank. It suffices to show that for any sufficiently large v_i^ω and $v_j^{\tilde{\omega}}$, there exist $(w_i^\omega(y), w_j^{\tilde{\omega}}(y))_{y \in Y}$ that satisfy the first four constraints with equalities and

$$w_i^\omega(y) + w_j^{\tilde{\omega}}(y) = 0, \quad \forall y \in Y.$$

Eliminate this last equation by solving for $w_j^{\tilde{\omega}}(y)$. Then the coefficient matrix for the set of the remaining four equations is

$$\begin{pmatrix}
(\pi_y^\omega(a'_i, \alpha_{-i}))_{y \in Y} \\
(\pi_y^\omega(a''_i, \alpha_{-i}))_{y \in Y} \\
(\pi_y^{\tilde{\omega}}(a'_j, \alpha_{-j}))_{y \in Y} \\
(\pi_y^{\tilde{\omega}}(a''_j, \alpha_{-j}))_{y \in Y}
\end{pmatrix}$$

The statewise full rank condition guarantees that this matrix has rank four, so the system has a solution for any $(v_i^\omega, v_j^{\tilde{\omega}})$, and thus $k^*(\alpha, \lambda) = \infty$. Intuitively, this construction makes $w_i^\omega(y)$ large for signals y that are more likely under in state ω than in state $\tilde{\omega}$ and makes $w_i^\omega(y)$ negative for signals that are more likely under $\tilde{\omega}$, while keeping player i indifferent between all actions in state ω , and player j indifferent in state $\tilde{\omega}$. This would not be possible if the signal distribution were the same at the two states, or more generally if the above matrix were singular.

This example only explains why the k^* can be made arbitrarily large when exactly two components of λ are non-zero. And a similar idea applies even if λ has other nonzero components. For example, suppose that $\lambda_i^\omega = \lambda_j^{\tilde{\omega}} = \lambda_l^{\bar{\omega}} = 1$ and other components are zero. First, choose $(v_i^\omega, v_j^{\tilde{\omega}}, w_i^\omega, w_j^{\tilde{\omega}})$ as in the above example, so that constraints (i) and (ii) for (i, ω) and $(j, \tilde{\omega})$ are satisfied, v_i^ω and $v_j^{\tilde{\omega}}$ are large, and $w_i^\omega(y) + w_j^{\tilde{\omega}}(y) = 0$ for all $y \in Y$. What remains is to find $w_l^{\bar{\omega}}$ that satisfies constraints (i) and (ii) for $(l, \bar{\omega})$ and the feasibility constraint

$$v_i^\omega + v_j^{\tilde{\omega}} + v_l^{\bar{\omega}} \geq w_i^\omega(y) + w_j^{\tilde{\omega}}(y) + w_l^{\bar{\omega}}(y), \quad \forall y \in Y.$$

The individual full rank condition implies there is $w_i^{\bar{\omega}}(y)$ that satisfies constraints (i) and (ii), and since $w_i^{\omega}(y) + w_j^{\tilde{\omega}}(y) = 0$ and $v_i^{\omega} + v_j^{\tilde{\omega}}$ can be arbitrarily large, the feasibility constraint can be satisfied for any value of $w_i^{\bar{\omega}}(y)$.

Proof of Lemma 6. Let (i, ω) and $(j, \tilde{\omega})$ be such that $\lambda_i^{\omega} \neq 0$, $\lambda_j^{\tilde{\omega}} \neq 0$, and $\tilde{\omega} \neq \omega$. Let α be a profile that has statewise full rank for all (i, ω) and $(j, \tilde{\omega})$ satisfying $\omega \neq \tilde{\omega}$.

First, we claim that for every $K > 0$, there exist $z_i^{\omega} = (z_i^{\omega}(y))_{y \in Y}$ and $z_j^{\tilde{\omega}} = (z_j^{\tilde{\omega}}(y))_{y \in Y}$ such that

$$\pi^{\omega}(a_i, \alpha_{-i}) \cdot z_i^{\omega} = \frac{K}{\delta \lambda_i^{\omega}} \quad (1)$$

for all $a_i \in A_i$,

$$\pi^{\tilde{\omega}}(a_j, \alpha_{-j}) \cdot z_j^{\tilde{\omega}} = 0 \quad (2)$$

for all $a_j \in A_j$, and

$$\lambda_i^{\omega} z_i^{\omega}(y) + \lambda_j^{\tilde{\omega}} z_j^{\tilde{\omega}}(y) = 0 \quad (3)$$

for all $y \in Y$. To prove that this system of equations indeed has a solution, eliminate (3) by solving for $z_j^{\tilde{\omega}}(y)$. Then, there remain $|A_i| + |A_j|$ linear equations, and its coefficient matrix is $\Pi_{(i, \omega)(j, \tilde{\omega})}(\alpha)$. Since statewise full rank implies that this coefficient matrix has rank $|A_i| + |A_j|$, we can solve the system.

Next, for each $(l, \bar{\omega}) \in \mathbf{I} \times \mathbf{\Omega}$, we choose $(\tilde{w}_l^{\bar{\omega}}(y))_{y \in Y}$ so that

$$(1 - \delta)g_l^{\bar{\omega}}(a_l, \alpha_{-l}) + \delta \pi^{\bar{\omega}}(a_l, \alpha_{-l}) \cdot \tilde{w}_l^{\bar{\omega}} = 0 \quad (4)$$

for all $a_l \in A_l$. Note that this system has a solution, since α has individual full rank. Intuitively, continuation payoffs $\tilde{w}^{\bar{\omega}}$ are chosen so that players are indifferent over all actions and their payoffs are zero.

Let $K > \max_{y \in Y} \lambda \cdot \tilde{w}(y)$, and choose $(z_i^{\omega}(y))_{y \in Y}$ and $(z_j^{\tilde{\omega}}(y))_{y \in Y}$ to satisfy (1) through (3). Then, let

$$w_l^{\bar{\omega}}(y) = \begin{cases} \tilde{w}_i^{\omega}(y) + z_i^{\omega}(y) & \text{if } (l, \bar{\omega}) = (i, \omega) \\ \tilde{w}_j^{\tilde{\omega}}(y) + z_j^{\tilde{\omega}}(y) & \text{if } (l, \bar{\omega}) = (j, \tilde{\omega}) \\ \tilde{w}_l^{\bar{\omega}}(y) & \text{otherwise} \end{cases}$$

for each $y \in Y$. Also, let

$$v_l^{\bar{\omega}} = \begin{cases} \frac{K}{\lambda_i^{\omega}} & \text{if } (l, \bar{\omega}) = (i, \omega) \\ 0 & \text{otherwise} \end{cases}.$$

We claim that this (v, w) satisfies constraints (i) through (iii) in LP-Average. It follows from (4) that constraints (i) and (ii) are satisfied for all $(l, \bar{\omega}) \in (\mathbf{I} \times \Omega) \setminus \{(i, \omega), (j, \tilde{\omega})\}$. Also, using (1) and (4), we obtain

$$\begin{aligned} & (1 - \delta)g_i^\omega(a_i, \alpha_{-i}) + \delta\pi^\omega(a_i, \alpha_{-i}) \cdot w_i^\omega \\ &= (1 - \delta)g_i^\omega(a_i, \alpha_{-i}) + \delta\pi^\omega(a_i, \alpha_{-i}) \cdot (\tilde{w}_i^\omega + z_i^\omega) \\ &= \frac{K}{\lambda_i^\omega} \end{aligned}$$

for all $a_i \in A_i$. This shows that (v, w) satisfies constraints (i) and (ii) for (i, ω) . Likewise, from (2) and (4), (v, w) satisfies constraints (i) and (ii) for $(j, \tilde{\omega})$. Furthermore, using (3) and $K > \max_{y \in Y} \lambda \cdot \tilde{w}(y)$,

$$\begin{aligned} \lambda \cdot w(y) &= \lambda \cdot \tilde{w}(y) + \lambda_i^\omega z_i^\omega(y) + \lambda_j^{\tilde{\omega}} z_j^{\tilde{\omega}}(y) \\ &= \lambda \cdot \tilde{w}(y) < K = \lambda \cdot v \end{aligned}$$

for all $y \in Y$, and hence constraint (iii) holds.

Therefore, $k^*(\alpha, \lambda) \geq \lambda \cdot v = K$. Since K can be arbitrarily large, we conclude $k^*(\alpha, \lambda) = \infty$. *Q.E.D.*

Lemma 7. *Suppose that a profile α has pairwise full rank for all (i, ω) and (j, ω) satisfying $i \neq j$. Then, $k^*(\alpha, \lambda) = \lambda \cdot g(\alpha)$ for direction λ such that $(\lambda_i^\omega)_{i \in \mathbf{I}}$ has at least two non-zero components for some ω while $\lambda_j^{\tilde{\omega}} = 0$ for all $j \in \mathbf{I}$ and $\tilde{\omega} \neq \omega$.*

Proof. It follows from Lemma 1(b) that $k^*(\lambda, \alpha) \leq \lambda \cdot g(\alpha)$. Thus, in what follows, we establish that $k^*(\lambda, \alpha) \geq \lambda \cdot g(\alpha)$. To do so, we need to show that there exist continuation payoffs in $H(\lambda, \lambda \cdot g(\alpha))$ that enforce $(\alpha, g(\alpha))$.

As in the proof of Lemma 6, for each $i \in \mathbf{I}$ and $\tilde{\omega} \neq \omega$, there exist $(w_i^{\tilde{\omega}}(y))_{y \in Y}$ such that

$$v_i^{\tilde{\omega}} = (1 - \delta)g_i^{\tilde{\omega}}(a_i, \alpha_{-i}) + \delta\pi^{\tilde{\omega}}(a_i, \alpha_{-i}) \cdot w_i^{\tilde{\omega}}$$

for all $a_i \in A_i$. Moreover, it follows from Lemmas 4.3, 5.3, and 5.4 of FLM that there exist $(w_i^\omega(y))_{(i,y)}$ such that

$$v_i^\omega = (1 - \delta)g_i^\omega(a_i, \alpha_{-i}) + \delta\pi^\omega(a_i, \alpha_{-i}) \cdot w_i^\omega$$

for all $i \in \mathbf{I}$ and $a_i \in A_i$, and

$$\lambda \cdot w(y) = \sum_{i \in \mathbf{I}} \lambda_i^\omega w_i^\omega(y) = \sum_{i \in \mathbf{I}} \lambda_i^\omega v_i^\omega = \lambda \cdot v.$$

Obviously, the specified continuation payoffs are in $H(\lambda, \lambda \cdot g(\alpha))$ and enforce $(\alpha, g(\alpha))$, as desired. *Q.E.D.*

Lemma 8. *Suppose that α has individual full rank for all $(j, \tilde{\omega}) \neq (i, \omega)$ and has the best-response property for player i and for state ω . Then, $k^*(\alpha, \lambda) = \lambda \cdot g(\alpha)$ for direction λ such that $\lambda_i^\omega \neq 0$ and $\lambda_j^{\tilde{\omega}} = 0$ for all $(j, \tilde{\omega}) \neq (i, \omega)$.*

Proof. This is a straightforward generalization of Lemmas 5.1 and 5.2 of FLM. *Q.E.D.*

Proof of Theorem 4. From Lemma 3, it suffices to show that $Q = V^*$. To do so, we will compute the maximum score $k^*(\lambda)$ for each direction λ .

Case 1. Consider λ such that $\lambda_i^\omega \neq 0$ and $\lambda_j^{\tilde{\omega}} \neq 0$ for some $\tilde{\omega} \neq \omega$ and i possibly equal to j . In this case, players can transfer utilities across different states ω and $\tilde{\omega}$ while maintaining the feasibility constraint, and this construction allows $k^*(\alpha, \lambda, \delta) > \lambda \cdot g(\alpha)$, as Example 1 shows. In particular, from (SFR) and Lemma 6 we obtain $k^*(\lambda) = \infty$ for this direction λ .

Case 2. Consider λ such that $(\lambda_i^\omega)_{i \in I}$ has at least two non-zero components for some ω while $\lambda_i^{\tilde{\omega}} = 0$ for all $i \in I$ and $\tilde{\omega} \neq \omega$. Lemma 4 shows that every profile α can be approximated arbitrarily closely by a profile that has pairwise full rank for all players, and it follows from Lemma 7 that $k^*(\lambda) = \sup_\alpha k^*(\lambda, \alpha) = \max_{v \in V} \lambda \cdot v$.

Case 3. Consider λ such that $\lambda_i^\omega \neq 0$ for some (i, ω) and $\lambda_j^{\tilde{\omega}} = 0$ for all $(j, \tilde{\omega}) \neq (i, \omega)$. Suppose first that $\lambda_i^\omega > 0$. Since every pure action profile has individual full rank, $a^* \in \arg \max_{a \in A} g_i^\omega(a)$ also has individual full rank. Therefore, from Lemma 8,

$$k^*(\lambda) \geq k^*(a^*, \lambda) = \lambda_i^\omega g_i^\omega(a^*) = \max_{v \in V} \lambda \cdot v.$$

On the other hand, from Lemma 1(b), $k^*(\lambda) \leq \max_{v \in V} \lambda \cdot v$. Hence, we have $k^*(\lambda) = \max_{v \in V} \lambda \cdot v$.

Next, suppose that $\lambda_i^\omega < 0$. It follows from Lemmas 5 and 8 that for every $\varepsilon > 0$, there is a profile $\underline{\alpha}^\omega$ such that $|k^*(\underline{\alpha}^\omega, \lambda) - \lambda_i^\omega v_i^\omega| < \varepsilon$. Lemma 3.2 of FL shows that $k^*(\lambda) \leq \lambda_i^\omega v_i^\omega$, so $k^*(\lambda) = \lambda_i^\omega v_i^\omega$.

Combining these cases, we obtain $Q = V^*$. *Q.E.D.*

6 Weaker Sufficient Conditions for Folk Theorems

In this section we present a few alternative theorems that use weaker informational conditions to prove “static-threats” folk theorems, meaning that the theorems only ensure the attainability of payoffs that Pareto-dominate the payoffs of a static ex-post equilibrium. Consequently, these theorems assume that a static ex-post equilibrium exists. This is always true when the state only matters for the monitoring structure but has no impact on the expected payoffs (that is $g^\omega(a) = g(a)$), and it is also satisfied for generic payoff functions g when the state has a sufficiently small impact on the payoff function. Several of our other assumptions in this section seem more likely to be satisfied if the uncertainty is “small,” though that is not necessary, as shown by Example 3, and we have not tried to prove formal results along those lines.

6.1 Relaxing Full Rank to Identifiability

This subsection develops informational conditions that are analogous to the identifiability conditions of FLM. These conditions do not require individual full rank, so that a given player may have several actions that generate the same signal distributions, and not all actions need be enforceable, but deviations by different players can still be distinguished.

Definition 10. For each $i \in I$, $j \neq i$, and $\omega \in \Omega$, a profile α is *pairwise identifiable* for (i, ω) and (j, ω) if $\text{rank}\Pi_{(i,\omega)(j,\omega)}(\alpha) = \text{rank}\Pi_{(i,\omega)}(\alpha) + \text{rank}\Pi_{(j,\omega)}(\alpha) - 1$.

This is exactly the FLM definition of pairwise identifiability. (Recall that pairwise full rank is equivalent to the combination of individual full rank and pairwise identifiability.)

Definition 11. For each $i \in I$, $j \in I$, $\omega \in \Omega$, and $\tilde{\omega} \neq \omega$, a profile α is *statewise identifiable* for (i, ω) and $(j, \tilde{\omega})$ if $\text{rank}\Pi_{(i,\omega)(j,\tilde{\omega})}(\alpha) = \text{rank}\Pi_{(i,\omega)}(\alpha) + \text{rank}\Pi_{(j,\tilde{\omega})}(\alpha)$.

Note that statewise full rank is the combination of individual full rank and statewise identifiability. Thus when individual full rank is satisfied, statewise identifiability requires just as many signals as statewise full rank, in contrast to the statewise distinguishability condition in the next subsection.

We say that α is ex-post enforceable if it is ex-post enforceable with respect to $\mathbf{R}^{I \times |\Omega|}$ and δ for some $\delta \in (0, 1)$. This is equivalent to α being enforceable with respect to \mathbf{R}^I and δ for each information structure π^ω in isolation.

Condition X-Eff. If a pure action profile a gives a Pareto-efficient payoff vector for some $\omega \in \Omega$, then it is ex-post enforceable.

FLM show that any Pareto-efficient action profile is enforceable. (X-Eff) extends this to ex-post enforceability, so it is automatically satisfied when there is a single state.

Condition U-Eff. If a pure action profile a gives a Pareto-efficient payoff vector for some $\tilde{\omega} \in \Omega$, then it gives a Pareto-efficient payoff vector for every ω .

(U-Eff) says roughly that efficient actions are uniformly efficient. It is typically satisfied if $u_i^\omega(y, a_i)$ is independent of (or very insensitive to) ω and the various distributions π^ω are sufficiently similar.

Condition PID. For each $i \in I$, $j \neq i$, and $\omega \in \Omega$, every pure action profile is pairwise identifiable for (i, ω) and (j, ω) .

(PID) is stronger than needed, it is sufficient that it applies to the pure action profiles that yield Pareto-efficient payoffs.

Lemma 9. *If $u_i^\omega(y, a_i)$ is independent of ω and (U-Eff) holds, then (X-Eff) holds.*

Proof. Because each player's payoff depends only on their own action and the realized signal, Lemma 6.1 of FLM applied to each state ω in isolation implies that profile a is enforceable for each ω . *Q.E.D.*

Condition SID. For each $i \in I$, $j \in I$, $\omega \in \Omega$, and $\tilde{\omega} \neq \omega$, there is a profile that is ex-post enforceable and statewise identifiable for (i, ω) and $(j, \tilde{\omega})$.

Intuitively, it is this condition that will allow the players to “learn the state” in a PPXE. It can be replaced by the less restrictive but harder to check condition of statewise distinguishability, as we show in the next subsection.

Theorem 5. *Suppose (PFR) holds or (X-Eff) and (PID) hold. Suppose also that (SID) holds. Assume that there is a static ex-post equilibrium α^0 , and let $V^0 \equiv \{v \in V \mid \forall i \in I \forall \omega \in \Omega \ v_i^\omega \geq g_i^\omega(\alpha^0)\}$. Then, for any smooth strict subset W of V^0 , there is $\bar{\delta} \in (0, 1)$ such that $W \subseteq E(\bar{\delta})$ for all $\delta \in (\bar{\delta}, 1)$.*

This theorem is established by the following lemmas that determine the maximal score k^* in various directions. The next lemma says that score of a static ex-post equilibrium can be enforced in any direction; this score will be used to generate the score in directions that minimize a player's payoff.

Lemma 10. *Suppose that there is a static ex-post equilibrium α^0 . Then, $k^*(\alpha^0, \lambda) \geq \lambda \cdot g(\alpha^0)$ for any direction λ .*

Proof. Let $v_i^\omega = w_i^\omega(y) = g_i^\omega(\alpha^0)$ for all $i \in \mathbf{I}$, $\omega \in \Omega$, and $y \in Y$. Then, this (v, w) satisfies constraints (i) through (iii) in LP-Average, and $\lambda \cdot v = \lambda \cdot g(\alpha^0)$. Hence, $k^*(\alpha^0, \lambda) \geq \lambda \cdot g(\alpha^0)$. *Q.E.D.*

The next lemma determines the maximal score for direction λ that considers a single state ω and has a positive component or at least two nonzero components, when (X-Eff) holds.

Lemma 11. (a) *Suppose that (PFR) or (X-Eff) and (PID) hold, and let a be a profile that gives a Pareto-efficient payoff for some $\omega \in \Omega$. Then, $k^*(a, \lambda) = \lambda \cdot g(a)$ for direction λ such that $(\lambda_i^\omega)_{i \in \mathbf{I}}$ has at least two non-zero components while $\lambda_j^{\tilde{\omega}} = 0$ for all $j \in \mathbf{I}$ and $\tilde{\omega} \neq \omega$.*

(b) *Suppose that (PFR) or (X-Eff) and (PID) hold. Then $k^*(\lambda) = \max_{v \in V} \lambda \cdot v$ for direction λ such that $\lambda_i^\omega > 0$ and $\lambda_j^{\tilde{\omega}} = 0$ for all $(j, \tilde{\omega}) \neq (i, \omega)$.*

Proof. Part (a). Lemma 1(b) shows that the maximum score in direction λ is at most $\lambda \cdot g(a)$. Because a is a pure action profile, and it is enforceable for all ω and pairwise identifiable from (X-Eff) and (PID), is enforceable on hyperplanes corresponding to λ from Theorem 5.1 of FLM, so the score $\lambda \cdot g(a)$ can be attained. If (PFR) holds this follows from Lemmas 4 and 7.

Part (b). Let a be a Pareto-efficient profile that maximizes player i 's payoff in state ω . If (X-Eff) holds, a is ex-post enforceable, and since the profile has the best-response property in state ω , Lemma 5.2 of FLM implies it is enforceable on λ . If (PFR) holds, this follows from Lemma 8. *Q.E.D.*

Lemma 12. *Suppose (SID) holds. Then, $k^*(\lambda) = \infty$ for direction λ such that there exist $i \in \mathbf{I}$, $j \in \mathbf{I}$, $\omega \in \Omega$, and $\tilde{\omega} \neq \omega$ such that $\lambda_i^\omega \neq 0$ and $\lambda_j^{\tilde{\omega}} \neq 0$.*

Proof. See Appendix. As in Lemma 5.5 of FLM, the idea is that if profile a is enforceable, then any action $a'_i \neq a_i$ leading to the same distribution of public signals as a_i cannot increase player i 's payoff and so can be ignored. Statewise identifiability implies that once we delete these redundant actions, the matrix $\Pi_{(i,\omega)(j,\tilde{\omega})}$ corresponding to the remaining actions satisfies statewise full rank. Therefore, as in Lemma 6, we can choose continuation payoffs to attain infinitely large maximal score while maintaining exact indifference among the remaining actions. These continuation payoffs also deter deviations to the deleted actions. *Q.E.D.*

6.2 Relaxing Statewise Identifiability

When individual full rank holds, statewise identifiability implies statewise full rank, which can require that there be twice as many signals as required by the FLM folk theorem. The following, more complex, condition can be satisfied with far fewer signals. In part, this condition is related to the fact that linear independence of the outcome distributions is not needed for an action profile to be enforceable, as linear independence tests all linear combinations of the distributions, while it is sufficient to rule out convex combinations.¹²

Definition 12. Profile α *statewise distinguishes* (i, ω) from $(j, \tilde{\omega})$ if there is a vector $\xi = (\xi(y))_{y \in Y} \in \mathbf{R}^{|Y|}$ such that

- (i) $\xi \cdot \pi^\omega(\alpha) > \xi \cdot \pi^{\tilde{\omega}}(\alpha)$,
- (ii) $\xi \cdot \pi^\omega(\alpha) = \xi \cdot \pi^\omega(a_i, \alpha_{-i}) \geq \xi \cdot \pi^\omega(\tilde{a}_i, \alpha_{-i})$ for all $a_i \in \text{supp} \alpha_i$ and $\tilde{a}_i \in A_i$,
- (iii) $\xi \cdot \pi^{\tilde{\omega}}(\alpha) = \xi \cdot \pi^{\tilde{\omega}}(a_j, \alpha_{-j})$ for all $a_j \in A_j$.

We illustrate these conditions in Figure ???. Clause (i) implies that the signals generated by α statistically distinguish ω from $\tilde{\omega}$. Clearly, there must be some such profile for there to be equilibria where the play varies with the state. Clause (ii) says that changing player i 's continuation payoff function in state ω from

¹²See Kandori and Matsushima (1998). In the study of mechanism design with transferable utility, Kosenok and Severinov (2008) and Rahman and Obara (2008) gave a weaker sufficient condition for budget-balanced implementation; the balanced-budget constraint roughly corresponds to directions λ where every component is strictly positive.

$w_i^\omega(y)$ to $w_i^\omega(y) + \xi(y)$ preserves incentive compatibility for player i , and clause (iii) says that the change in player i 's continuation payoff (of $\Delta w_i^\omega(y) \equiv \xi(y)$) can be offset to preserve the feasibility constraint ($\lambda_i^\omega \Delta w_i^\omega(y) + \lambda_j^{\tilde{\omega}} \Delta w_j^{\tilde{\omega}}(y) = 0$) without changing player j 's expected continuation payoff to any action. Note that this transfer scheme increases player i 's expected continuation payoff by $E[\Delta w_i^\omega | \alpha] \equiv \xi \cdot \pi^\omega(\alpha)$, so the maximal score for λ with $\lambda_i^\omega > 0$ can be made infinitely large by utility transfer between states ω from $\tilde{\omega}$.¹³

Condition SD. For each (i, ω) and $(j, \tilde{\omega})$ satisfying $\omega \neq \tilde{\omega}$, there is an ex-post enforceable action profile α that statewise distinguishes (i, ω) from $(j, \tilde{\omega})$.

(SD) is sufficient for the static-threat folk theorem, as it implies that profile α can generate an infinite score in all of the “cross-state” directions. The folk theorem holds with even weaker conditions, because different profiles can be used in different directions.

Definition 13. Profile α *n-statewise distinguishes* (i, ω) from $(j, \tilde{\omega})$ if there is a vector $\xi = (\xi(y))_{y \in Y} \in \mathbf{R}^{|Y|}$ such that

- (i) $\xi \cdot \pi^\omega(\alpha) > \xi \cdot \pi^{\tilde{\omega}}(\alpha)$,
- (ii) $\xi \cdot \pi^\omega(\alpha) = \xi \cdot \pi^\omega(a_i, \alpha_{-i}) \geq \xi \cdot \pi^\omega(\tilde{a}_i, \alpha_{-i})$ for all $a_i \in \text{supp} \alpha_i$ and $\tilde{a}_i \in A_i$,
- (iii) $\xi \cdot \pi^{\tilde{\omega}}(\alpha) = \xi \cdot \pi^{\tilde{\omega}}(a_j, \alpha_{-j}) \geq \xi \cdot \pi^{\tilde{\omega}}(\tilde{a}_j, \alpha_{-j})$ for all $a_j \in \text{supp} \alpha_j$ and $\tilde{a}_j \in A_j$.

Note that this condition relaxes statewise distinguishability by replacing the last equality in (iii) with an inequality. We will show that a profile n -statewise distinguishes (i, ω) from $(j, \tilde{\omega})$ can be used to generate an infinite score for all λ such that $\lambda_i^\omega > 0$ and $\lambda_j^{\tilde{\omega}} < 0$; the “ n ” in the conditions name refers to the fact that in these directions there is a negative trade-off between the utilities of the two players. Intuitively, if player i 's continuation payoff in state ω is increased by ξ , and $\lambda_i^\omega > 0$ and $\lambda_j^{\tilde{\omega}} < 0$, then to satisfy the feasibility constraint it is sufficient that $\lambda_i^\omega \Delta w_i^\omega(y) + \lambda_j^{\tilde{\omega}} \Delta w_j^{\tilde{\omega}}(y) = 0$, which implies that $\Delta w_i^\omega = \xi$ and $\Delta w_j^{\tilde{\omega}} = -\frac{\lambda_i^\omega}{\lambda_j^{\tilde{\omega}}} \xi$ are

¹³If $\lambda_i^\omega < 0$, then player i 's continuation payoff must be decreased to achieve a high score. This requires a different sort of transfer and in turn requires a different condition on the information structure, but this condition is not needed for a static-threats folk theorem.

colinear. When this is true clause (iii) implies that the change preserves incentive compatibility for player j in state $\tilde{\omega}$, as Figure ?? shows. Thus the maximal score can be made infinitely large by utility transfer between states ω from $\tilde{\omega}$.

The next condition is similar, but it is tailored to be useful “positive” directions λ such that $\lambda_i^\omega > 0$ and $\lambda_j^{\tilde{\omega}} > 0$.

Definition 14. Profile α *p-statewise distinguishes* (i, ω) from $(j, \tilde{\omega})$ if there is a vector $\xi = (\xi(y))_{y \in Y} \in \mathbf{R}^{|Y|}$ such that

- (i) $\xi \cdot \pi^\omega(\alpha) > \xi \cdot \pi^{\tilde{\omega}}(\alpha)$,
- (ii) $\xi \cdot \pi^\omega(\alpha) = \xi \cdot \pi^\omega(a_i, \alpha_{-i}) \geq \xi \cdot \pi^\omega(\tilde{a}_i, \alpha_{-i})$ for all $a_i \in \text{supp}\alpha_i$ and $\tilde{a}_i \in A_i$,
- (iii) $\xi \cdot \pi^{\tilde{\omega}}(\alpha) = \xi \cdot \pi^{\tilde{\omega}}(a_j, \alpha_{-j}) \leq \xi \cdot \pi^{\tilde{\omega}}(\tilde{a}_j, \alpha_{-j})$ for all $a_j \in \text{supp}\alpha_j$ and $\tilde{a}_j \in A_j$.

For λ such that $\lambda_i^\omega > 0$ and $\lambda_j^{\tilde{\omega}} > 0$, the feasibility constraint implies that Δw_i^ω and $\Delta w_j^{\tilde{\omega}}$ point in opposite directions, and clause (iii) assures that player j 's incentive compatibility is preserved when ξ is added to player i 's continuation payoffs. This is illustrated in Figure ??.

Note that if α statewise distinguishes (i, ω) from $(j, \tilde{\omega})$, then it n -statewise distinguishes this pair and p -statewise distinguishes this pair. So the following condition is weaker than (SD).

Condition Weak-SD. For each (i, ω) and $(j, \tilde{\omega})$ satisfying $\omega \neq \tilde{\omega}$, there is an ex-post enforceable action profile α that n -statewise distinguishes (i, ω) from $(j, \tilde{\omega})$, and an ex-post enforceable action profile $\tilde{\alpha}$ that p -statewise distinguishes (i, ω) from $(j, \tilde{\omega})$.

This is sufficient for the static-threat folk theorem as shown by the following lemmas. See the Appendix for their proofs.

Lemma 13. *Suppose that α is ex-post enforceable and n -statewise distinguishes (i, ω) from $(j, \tilde{\omega})$. Then $k^*(\alpha, \lambda) = \infty$ for direction λ such that $\lambda_i^\omega > 0$ and $\lambda_j^{\tilde{\omega}} < 0$.*

Lemma 14. *Suppose that α is ex-post enforceable and p -statewise distinguishes (i, ω) from $(j, \tilde{\omega})$. Then $k^*(\alpha, \lambda) = \infty$ for direction λ such that $\lambda_i^\omega > 0$ and $\lambda_j^{\tilde{\omega}} > 0$.*

Combining Lemmas 10, 11, 13 and 14 yields the following folk theorem:

Theorem 6. *Suppose (X-Eff) and (Weak-SD) hold. Assume that there is a static ex-post equilibrium α^0 , and let $V^0 \equiv \{v \in V \mid \forall i \in \mathbf{I} \forall \omega \in \Omega \ v_i^\omega \geq g_i^\omega(\alpha^0)\}$. Then, for any smooth strict subset W of V^0 , there is $\bar{\delta} \in (0, 1)$ such that $W \subseteq E(\delta)$ for all $\delta \in (\bar{\delta}, 1)$.*

To illustrate statewise distinguishability, we return to the two partnership games that we introduced in Section 3.

Example 3. There are two players, two states, and three outcomes $Y = \{H, M, L\}$. Each player i 's action space is $A_i = \{C_i, D_i\}$. The probabilities of H and M if both players choose D are independent of the state, and in each state the monitoring structure is additive: the change in probabilities induced by player i 's changing from C_i to D_i is the same regardless of the action of the other player. Moreover, in this example the uncertainty is symmetric in the state: In state ω_1 , if player 1 chooses C_1 instead of D_1 , then the probabilities of H and M increase by p_H and p_M , while player 2's choice of C_2 increases the probabilities by q_H and q_M ; in state ω_2 , the roles are reversed.

The realized payoff functions are independent of ω and given by

$$u_i(C_i, y) = r_i(y) - e_i \quad \text{and} \quad u_i(D_i, y) = r_i(y)$$

for each $i \in \mathbf{I}$, $\omega \in \Omega$, and $y \in Y$. We assume that for each $i \in \mathbf{I}$,

$$\begin{aligned} r_i(H) &> r_i(M) > r_i(L), \\ e_i &> p_H(r_i(H) - r_i(L)) + p_M(r_i(M) - r_i(L)), \\ e_i &> q_H(r_i(H) - r_i(L)) + q_M(r_i(M) - r_i(L)). \end{aligned}$$

Here the left-hand side of the second inequality is the cost of player 1's choice of C_1 for state ω_1 (or the cost of player 2's choice of C_2 for state ω_2), and the right-hand side is an increase in player 1's benefit from the project when he chooses C_1 instead of D_1 for state ω_1 (or an increase in player 2's benefit when he chooses C_2 for state ω_2). Since the left-hand side is greater than the right-hand side, we conclude that D_1 strictly dominates C_1 for state ω_1 , and D_2 strictly dominates C_2 for state ω_2 . Likewise, the third inequality asserts that D_1 strictly dominates C_1

for state ω_2 , and D_2 strictly dominates C_2 for state ω_1 . Thus, D_i strictly dominates C_i for each state. Moreover, we assume that for each $i \in I$,

$$e_i < p_H(r_1(H) + r_2(H) - r_1(L) - r_2(L)) + p_M(r_1(M) + r_2(M) - r_1(L) - r_2(L))$$

and

$$e_i < q_H(r_1(H) + r_2(H) - r_1(L) - r_2(L)) + q_M(r_1(M) + r_2(M) - r_1(L) - r_2(L)),$$

so that choosing C_i instead of D_i always increases the total surplus. Summing up, the payoff matrix of the stage game corresponds to a prisoner's dilemma for each state; hence, V^* has a non-empty interior and (D_1, D_2) is a static ex-post equilibrium.

Note that individual full rank is satisfied, and that pairwise full rank is satisfied at every profile and every state if the matrix

$$\begin{pmatrix} p_H & p_M \\ q_H & q_M \end{pmatrix}$$

has full rank. For example, the matrix $\Pi_{(1, \omega_1)(2, \omega_1)}(D_1, C_2)$ is represented by

$$\begin{pmatrix} o_H + q_H & o_M + q_M & 1 - (o_H + q_H + o_M + q_M) \\ o_H + p_H + q_H & o_M + p_M + q_M & 1 - (o_H + p_H + q_H + o_M + p_M + q_M) \\ o_H + q_H & o_M + q_M & 1 - (o_H + q_H + o_M + q_M) \\ o_H & o_M & 1 - (o_H + o_M) \end{pmatrix},$$

and this matrix has rank three if the above two-by-two matrix has full rank. Therefore, the profile (D_1, C_2) has pairwise full rank for $(1, \omega_1)$ and $(2, \omega_1)$. On the other hand, statewise identifiability is not satisfied at any profile, as there are only three signals, while four signals would be needed to satisfy statewise identifiability and individual full rank.

Claim 2. *In Example 3, (D_1, C_2) statewise distinguishes (i, ω) from $(j, \tilde{\omega})$ satisfying $\omega \neq \tilde{\omega}$.*

Proof. First, consider $((i, \omega), (j, \tilde{\omega})) = ((1, \omega_1), (2, \omega_2))$. In this case, let $\xi = (\xi(y))_{y \in Y}$ be a solution to the system

$$\begin{aligned} \xi(H)p_H + \xi(M)p_M + \xi(L)(1 - p_H - p_M) &= 0 \\ \xi(H)q_H + \xi(M)q_M + \xi(L)(1 - q_H - q_M) &= K \end{aligned}$$

for some $K > 0$. This system has a solution, since the matrix

$$\begin{pmatrix} p_H & p_M \\ q_H & q_M \end{pmatrix}$$

has full rank. Then, we can check that

$$\xi \cdot \pi^{\omega_1}(C_1, C_2) = \xi \cdot \pi^{\omega_1}(D_1, C_2) = \xi \cdot \pi^{\omega_2}(D_1, C_2) + K = \xi \cdot \pi^{\omega_2}(D_1, D_2) + K$$

so that all the desired conditions are satisfied. For $((i, \omega), (j, \tilde{\omega})) = ((2, \omega_1), (2, \omega_2))$, we can use the same ξ .

Likewise, for $((i, \omega), (j, \tilde{\omega})) = ((1, \omega_2), (1, \omega_1))$ or $((i, \omega), (j, \tilde{\omega})) = ((2, \omega_2), (1, \omega_1))$, use ξ that solves

$$\begin{aligned} \xi(H)p_H + \xi(M)p_M + \xi(L)(1 - p_H - p_M) &= 0 \\ \xi(H)q_H + \xi(M)q_M + \xi(L)(1 - q_H - q_M) &= -K \end{aligned}$$

for some $K > 0$. For $((i, \omega), (j, \tilde{\omega})) = ((1, \omega_1), (1, \omega_2))$ or $((i, \omega), (j, \tilde{\omega})) = ((2, \omega_1), (1, \omega_2))$, use ξ that solves

$$\begin{aligned} \xi(H)p_H + \xi(M)p_M + \xi(L)(1 - p_H - p_M) &= -K \\ \xi(H)q_H + \xi(M)q_M + \xi(L)(1 - q_H - q_M) &= 0 \end{aligned}$$

for some $K > 0$. Finally, for $((i, \omega), (j, \tilde{\omega})) = ((1, \omega_2), (2, \omega_1))$ or $((i, \omega), (j, \tilde{\omega})) = ((2, \omega_2), (2, \omega_1))$, use ξ that solves

$$\begin{aligned} \xi(H)p_H + \xi(M)p_M + \xi(L)(1 - p_H - p_M) &= K \\ \xi(H)q_H + \xi(M)q_M + \xi(L)(1 - q_H - q_M) &= 0 \end{aligned}$$

for some $K > 0$. *Q.E.D.*

We conclude that the static-threat folk theorem applies to Example 3. In contrast, payoffs are bounded away from efficiency in Example 4, which is a related partnership game. As we remarked earlier, this is because the states are “entangled”:

Definition 15. Profile α entangles states ω and $\tilde{\omega}$ for player j if there is $\bar{\pi} \in \text{co}\{\pi^{\tilde{\omega}}(a_j, \alpha_{-j}) | a_j \in A_j\}$ such that $\pi^{\tilde{\omega}}(\alpha) = \kappa \pi^{\omega}(\alpha) + (1 - \kappa) \bar{\pi}$ for some $\kappa \in (0, 1]$.

Lemma 15. *If profile α p -statewise distinguishes (i, ω) from $(j, \tilde{\omega})$ then α does not entangle ω and $\tilde{\omega}$ for player j .*

Proof. If α entangles states ω and $\tilde{\omega}$ for player j then for any ξ such that $\xi \cdot \pi^{\tilde{\omega}}(\alpha) \leq \xi \cdot \pi^{\tilde{\omega}}(a_j, \alpha_{-j})$ for all $a_j \in A_j$, we have $\xi \cdot \pi^{\tilde{\omega}}(\alpha) \leq \xi \cdot \bar{\pi}$ for all $\bar{\pi} \in \text{co}\{\pi^{\tilde{\omega}}(a_j, \alpha_{-j}) | a_j \in A_j\}$, so that $\xi \cdot \pi^{\tilde{\omega}}(\alpha) \geq \xi \cdot \pi^{\omega}(\alpha)$. Thus α does not p -statewise distinguish (i, ω) from $(j, \tilde{\omega})$. *Q.E.D.*

Example 4. As in the previous example, there are two players, two actions $A_i = \{C_i, D_i\}$, two states, and three outcomes $Y = \{H, M, L\}$. The state only influences the productivity of player 2's effort: If player 1 chooses C_1 instead of D_1 , then the probabilities of H and M increase by p_H and p_M , independent of the state. In contrast, if player 2 chooses C_2 instead of D_2 , then the probabilities of H and M increase by q_H and q_M in state ω_1 , but they increase only by βq_H and βq_M in state ω_2 . If $\beta < 1$, the states have different outcome distributions, so can be identified by repeated observation. However, every profile where player 2 plays C_2 with positive probability entangles the states for player 2, and if player 2 plays D_2 the states cannot be distinguished, so no profile statewise distinguishes $(1, \omega_1)$ and $(2, \omega_2)$.

As in Example 3, the payoffs are

$$u_i(C_i, y) = r_i(y) - e_i \quad \text{and} \quad u_i(D_i, y) = r_i(y)$$

for each $i \in I$ and $y \in Y$. We again make assumptions on the r_i so that the stage game payoffs in each state correspond to a prisoner's dilemma: D_i is a dominant strategy, so (D_1, D_2) is a static ex-post equilibrium, (C_1, C_2) is efficient, and V^* has a non-empty interior.¹⁴

Individual full rank and the pairwise full rank are satisfied at every profile and every state, if the matrix

$$\begin{pmatrix} p_H & p_M \\ q_H & q_M \end{pmatrix}$$

has full rank, as in Example 3. However, no profile can statewise distinguish $(1, \omega_1)$ from $(2, \omega_2)$, as the following claim shows.

¹⁴The conditions on the payoffs are: $r_i(H) > r_i(M) > r_i(L)$; $e_1 > p_H(r_1(H) - r_1(L)) + p_M(r_1(M) - r_1(L))$; $e_2 > q_H(r_2(H) - r_2(L)) + q_M(r_2(M) - r_2(L))$; $e_1 < p_H(r_1(H) + r_2(H) - r_1(L) - r_2(L)) + p_M(r_1(M) + r_2(M) - r_1(L) - r_2(L))$; and $e_2 < \beta q_H(r_1(H) + r_2(H) - r_1(L) - r_2(L)) + \beta q_M(r_1(M) + r_2(M) - r_1(L) - r_2(L))$.

In what follows, we prove that the folk theorem fails in this example. Specifically, we show that the maximal score $k^*(\lambda)$ in direction $\lambda' = ((1,0), (0,1))$ is strictly less than the value $\lambda' \cdot g(C_1, C_2) < \max_{v \in V^*} \lambda' \cdot v$. To do this, we use the fact that the monitoring technology has an additive form, so that it suffices to consider only the pure action profiles, as in Lemma 4.1 of FL.¹⁵

Claim 3. For $\alpha = (C_1, C_2)$,

$$k^*(\alpha, \lambda') \leq \lambda' \cdot g(C_1, C_2) - \frac{1-\beta}{\beta} (g_2^{\omega_2}(C_1, D_2) - g_2^{\omega_2}(C_1, C_2)).$$

Proof. See Appendix. As we mentioned in Section 3, this inefficiency result comes from the fact that the two states are entangled for player 2 and hence the profile (C_1, C_2) does not p -statewise distinguish $(1, \omega_1)$ from $(2, \omega_2)$. *Q.E.D.*

Claim 4. For $\alpha = (D_1, C_2)$, $k^*(\alpha, \lambda') \leq \lambda' \cdot g(D_1, C_2) - \frac{1-\beta}{\beta} (g_2^{\omega_2}(D_1, D_2) - g_2^{\omega_2}(D_1, C_2))$.

Proof. The same as in the previous claim. *Q.E.D.*

Claim 5. For $\alpha = (C_1, D_2)$, $k^*(\alpha, \lambda') \leq \lambda' \cdot g(C_1, D_2)$.

Proof. Since $\pi^{\omega_1}(C_1, D_2) = \pi^{\omega_2}(C_1, D_2)$ and $\pi^{\omega_1}(D_1, D_2) = \pi^{\omega_2}(D_1, D_2)$, the set of the constraints in the LP-Average problem for λ' is isomorphic with that for $\lambda'' = ((0,0), (1,1))$. Then the maximal score for λ' equals that for λ'' , and the statement follows from Lemma 1(b). *Q.E.D.*

Claim 6. For $\alpha = (D_1, D_2)$, $k^*(\alpha, \lambda') \leq \lambda' \cdot g(D_1, D_2)$.

Proof. The same as in the last claim. *Q.E.D.*

Now we combine these claims to show that $k^*(\lambda') < \lambda' \cdot g(C_1, C_2)$. Since $g_1^{\omega_1}(C_1, D_2) = g_1^{\omega_2}(C_1, D_2)$, we have

$$\begin{aligned} \lambda' \cdot g(C_1, D_2) &= g_1^{\omega_1}(C_1, D_2) + g_2^{\omega_2}(C_1, D_2) = g_1^{\omega_2}(C_1, D_2) + g_2^{\omega_2}(C_1, D_2) \\ &< g_1^{\omega_2}(C_1, C_2) + g_2^{\omega_2}(C_1, C_2) \leq g_1^{\omega_1}(C_1, C_2) + g_2^{\omega_2}(C_1, C_2) = \lambda' \cdot g(C_1, C_2). \end{aligned}$$

¹⁵FL used a more restrictive definition of “additive monitoring structure,” but the proof of their Lemma 4.1 applies to any case where the effect of one player’s action on the distribution of signals is independent of the action of the other player.

Also,

$$\begin{aligned}
\lambda' \cdot g(D_1, C_2) &= g_1^{\omega_1}(D_1, C_2) + g_2^{\omega_2}(D_1, C_2) \\
&= g_1^{\omega_1}(C_1, C_2) + g_2^{\omega_2}(C_1, C_2) \\
&\quad + (g_1^{\omega_1}(D_1, C_2) + g_2^{\omega_1}(D_1, C_2) - g_1^{\omega_1}(C_1, C_2) - g_2^{\omega_1}(C_1, C_2)) \\
&< g_1^{\omega_1}(C_1, C_2) + g_2^{\omega_2}(C_1, C_2) \\
&= \lambda' \cdot g(C_1, C_2).
\end{aligned}$$

Here, the second equality comes from the additive structure, which implies that $g_2^{\omega_1}(D_1, C_2) - g_2^{\omega_1}(C_1, C_2) = g_2^{\omega_2}(D_1, C_2) - g_2^{\omega_2}(C_1, C_2)$. Combined with the previous claims, it follows that $k^*(\lambda') < \lambda' \cdot g(C_1, C_2)$, so that the folk theorem fails. Moreover, because the player's equilibrium payoffs cannot be below their min-max level in any state, this bound implies that for some parameter values player 2's PPXE payoff in state ω_2 is strictly less than $g_2^{\omega_2}(C_1, C_2)$.¹⁶

7 Incomplete Information and Belief-Free Equilibria

7.1 PPXE of Incomplete-Information Games

So far we have assumed that the players have symmetric information about the state. Now suppose that each player i observes a private signal $\theta_i \in \Theta_i$ at the beginning of the game, where Θ_i is a partition of Ω . Any public strategy s_i of the game where player i has a trivial partition, $\Theta_i = \{(\Omega)\}$ induces a public strategy for any non-trivial partition Θ_i : player i simply ignores the private information and sets $s'_i(h, \theta_i) = s_i(h)$ for all h and all θ_i . Since by definition play in a PPXE is optimal regardless of the state, any PPXE for the symmetric-information game (where all players have the trivial partition) induces a PPXE for any incomplete-information game (any partitions Θ_i) with the same payoff functions and prior. That is, if strategy profile s is a PPXE of the symmetric-information game, then the profile s' where $s'_i(h, \theta_i) = s_i(h)$ for all players i , types θ_i , and histories h is

¹⁶For example, suppose that $p_H = .5$, $p_M = 0$, $q_H = 0$, $q_M = .5$, $\beta = .8$, $r_i(H) = 100$, $r_i(M) = 99$, $r_i(L) = 0$, $e_1 = 99$, and $e_2 = 79$. Then the minmax payoffs are 0 for all players and all states, $g_1^{\omega_1}(C_1, C_2) = 0.5$, $g_2^{\omega_2}(C_1, C_2) = 10.6$, and $g_2^{\omega_2}(C_1, D_2) = 50$. Using Claim 3, we have $v_1^1 + v_2^2 < 1.25$, and since $v_1^1 \geq 0$, $v_2^2 < 1.25$, it cannot achieve payoff $g_2^{\omega_2}(C_1, C_2) = 10.6$.

a PPXE of the incomplete-information game. Moreover, since θ_i is private information, any strategy that conditions on θ_i will not be a function of only the public information. Thus the PPXE of the incomplete-information games are isomorphic to the PPXE of the associated symmetric-information game, so the limit PPXE payoffs can be computed using LP-average, and our sufficient conditions for the folk theorem still apply.

However, we would expect the folk theorem to hold under weaker conditions if players are allowed to condition their play on their private information. Players *can* condition on their private information in the “belief-free equilibrium” studied by Hörner and Lovo (2009) and Hörner, Lovo, and Tomala (2008). These papers define a belief-free equilibrium for games with observable actions and incomplete information to be a strategy profile s such that for each state ω , profile s is a subgame-perfect equilibrium of the game where all players know the state is ω .¹⁷

When the information partitions are trivial (and actions are perfectly observed) belief-free equilibrium is equivalent to PPXE. In this case the game is one of complete information, and players have no way to learn the state, so one way to study the game is to replace the payoff functions in each state with their expected value, and apply subgame-perfect equilibrium to the resulting standard game. It may be that the folk theorem holds in this game, but the set of PPXE is empty, which might raise some questions about the strength of the robustness argument for ex-post equilibria; we are agnostic on the status of PPXE when the folk theorem fails but efficient payoffs can be supported by other sorts of equilibria.

When the information partitions are non-trivial, belief-free equilibrium allows a larger set of strategies than does symmetric-information PPXE, so the limit PPXE payoffs must be a weak or strict subset of the limit payoffs of belief-free equilibria. In the next subsection we study games with the monitoring structure of Hörner and Lovo (2009) and Hörner, Lovo, and Tomala (2008), and show that the inclusion is strict: some limit payoffs of belief-free equilibria are not limit payoffs of PPXE. In ongoing work Fudenberg and Yamamoto (2009), we define

¹⁷Hörner and Lovo (2009) study two-player games where the information partition has a product structure; Hörner, Lovo, and Tomala (2008) extends the analysis to general partitions and N -player games. These papers assume that players do not observe their realized payoffs as the game is played: The players’ only information is their initial private signal θ_i and the sequence of realized actions.

the notion of a “type-contingent perfect public ex-post equilibrium,” which allows players to condition on their initial private information in addition to the public history. This equilibrium concept reduces to the belief-free equilibrium of Hörner and Lovo (2009) and Hörner, Lovo, and Tomala (2008) when actions are perfectly observed. We then develop the appropriate linear programming characterization of limit equilibrium payoffs, which we hope to use to extend the results of Hörner, Lovo, and Tomala (2008) to games where actions are imperfectly observed and the monitoring structure is unknown.

7.2 Incomplete Information and Perfectly Observed Actions

Consider the following example from Hörner and Lovo (2009). There are two players, $I = \{1, 2\}$, and two states, $\Omega = \{\omega_1, \omega_2\}$. Player 1 knows the state, but player 2 does not: $\Theta_1 = \{(\omega_1), (\omega_2)\}$ and $\Theta_2 = \{(\Omega)\}$. Player $i \in I$ chooses actions $a_i \in A_i = \{T, B\}$, and observes a public signal $y \in A$. Assume that $\pi_y^\omega(a) = 1$ for $y = a$, so that actions are perfectly observable, and players cannot learn the state from the signals. The payoff matrix conditional on ω_1 is

	T	B
T	1, 1	$-L, 1 + G$
B	$1 + G, -L$	0, 0

where $0 < L - G < 1$. This game can be regarded as prisoner’s dilemma where T is cooperation and B is defection. On the other hand, the payoff matrix conditional on an ω_2 is

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

Note that this game is also prisoner’s dilemma, but now the role of each action is reversed; B is cooperation and T is defection.

Hörner and Lovo (2009) show that player 1’s best limit payoff in belief-free equilibrium is $1 + \frac{G}{1+L}$ in each state, which is the highest payoff consistent with individual rationality for player 2 in the games where the state is known. We will show that PPXE cannot attain as high a limit payoff. Intuitively, this is because

(a) the public signals do not directly reveal the state, so with trivial partitions ($\Theta_i = \{(\Omega)\}$) learning the state is impossible, and (b) the same conclusion obtains if player 1 does start out knowing the state but we restrict attention to equilibria in which player 1's play doesn't depend on his prior information.

Because the PPXE payoff set for games with asymmetric information is identical with that for the corresponding symmetric-information game, we can compute the limit set of PPXE payoffs for asymmetric-information games by using LP-Average.

Lemma 16. *Suppose that $Y = A$ and $\pi_y^\omega(a) = 1$ for $y = a$. Then $k^*(\alpha, \lambda) \leq \lambda \cdot g(\alpha)$ for all α and λ .*

Proof. Let (v, w) be a solution to LP-Average associated with $(\lambda, \alpha, \delta)$. By definition, (v, w) satisfies all the constraints in LP-Average, and since $Y = A$ we can treat the continuation payoffs as a function of the realized actions. Then,

$$\begin{aligned}
k^*(\alpha, \lambda) &= \sum_{i \in I} \sum_{\omega \in \Omega} \lambda_i^\omega \cdot v_i^\omega \\
&= \sum_{i \in I} \sum_{\omega \in \Omega} \lambda_i^\omega \left((1 - \delta) g_i^\omega(\alpha) + \delta \sum_{a \in A} \alpha(a) w_i^\omega(a) \right) \\
&= (1 - \delta) \lambda \cdot g(\alpha) + \delta \sum_{a \in A} \alpha(a) \lambda \cdot w(a) \\
&\leq (1 - \delta) \lambda \cdot g(\alpha) + \delta \sum_{a \in A} \alpha(a) k^*(\alpha, \lambda) \\
&= (1 - \delta) \lambda \cdot g(\alpha) + \delta k^*(\alpha, \lambda).
\end{aligned}$$

Here, the inequality follows from constraint (iii) in LP-Average. Subtracting $\delta k^*(\alpha, \lambda)$ from both sides and dividing by $(1 - \delta)$, we obtain $k^*(\alpha, \lambda) \leq \lambda \cdot g(\alpha)$, as desired. *Q.E.D.*

Consider λ such that $\lambda_1^{\omega_1} = \lambda_1^{\omega_2} = 1$ and $\lambda_2^{\omega_1} = \lambda_2^{\omega_2} = 0$. It follows from the above lemma that for any α ,

$$k^*(\alpha, \lambda) \leq \lambda \cdot g(\alpha) = g_1^{\omega_1}(\alpha) + g_1^{\omega_2}(\alpha).$$

Note that the value $g_1^{\omega_1}(\alpha) + g_1^{\omega_2}(\alpha)$ is maximized by $\alpha = (T, T)$ or $\alpha = (B, B)$, and its value is 1. Hence,

$$k^*(\lambda) = \sup_{\alpha} k^*(\alpha, \lambda) = 1.$$

This result shows that Q is contained in the hyperplane $H(\lambda, 1) = \{v \in \mathbf{R}^{I \times |\Omega|} \mid v_1^{\omega_1} + v_1^{\omega_2} \leq 1\}$, so that $v_1^{\omega_1} + v_1^{\omega_2} \leq 1$ for any $v \in \lim_{\delta \rightarrow 1} E(\delta)$. In words, the sum of player 1's equilibrium payoffs for state ω_1 and for ω_2 cannot exceed 1. On the other hand, since the equilibrium payoff must be above the minimax payoff for each state, we have $v_1^{\omega_1} \geq 0$ and $v_1^{\omega_2} \geq 0$ for all $v \in \lim_{\delta \rightarrow 1} E(\delta)$. Therefore, we obtain

$$0 \leq v_1^{\omega_1} \leq 1 \quad \text{and} \quad 0 \leq v_1^{\omega_2} \leq 1$$

for all $v \in \lim_{\delta \rightarrow 1} E(\delta)$. Obviously, this value is less than the best belief-free equilibrium payoff, $1 + \frac{G}{1+L}$.

8 Concluding Remarks

This paper has restricted attention to the set of PPXE, and analyzed them with extensions of the techniques used to analyze PPE in games where the monitoring structure is known. When the statewise full rank conditions hold, along with the standard individual and pairwise full rank conditions, the set of PPXE satisfies an ex-post folk theorem, even if the set of static ex-post equilibria is empty. When a static ex-post equilibrium does exist, there is an ex-post PPXE folk theorem under even milder informational conditions.

Of course for a given discount factor the full set of sequential equilibria of these games is larger than the set of PPXE, and can permit a larger set of payoffs. In particular, because the game has finitely many actions and signals per period and is continuous at infinity, sequential equilibria exist for any discount factor, even if the set of PPXE is empty,¹⁸ so PPXE is not well-adapted to the study of games with uncertain monitoring structures and very impatient players. Conversely, when players are patient and mostly concerned with their long-run payoff, our informational conditions imply that there are PPXE where players eventually learn what the state is, and obtain the same payoffs as if the state was publicly observed.¹⁹

¹⁸This follows from the facts that sequential equilibria exist in the finite-horizon truncations (Kreps and Wilson (1982)) and that the set of equilibrium strategies is compact in the product topology (Fudenberg and Levine (1983)).

¹⁹When the ex-post folk theorem holds, and a feasible payoff vector v is not a limit payoff of PPXE, then for some player/state pair the payoff v_i^ω is not ex post individually rational, so that

9 Appendix

9.1 Proof of Theorem 1

Theorem 1. *If a subset W of $\mathbf{R}^{I \times |\Omega|}$ is bounded and ex-post self-generating with respect to δ , then $W \subseteq E(\delta)$.*

Proof. Let $v \in W$. We will construct a PPXE that yields v . Since $v \in B(\delta, W)$, there exist a profile α and a function $w : Y \rightarrow W$ such that (α, v) is ex-post enforced by w . Set the action profile in period one to be $s|_{h^0} = \alpha$ and for each $h^1 = y^1 \in Y$, set $v|_{h^1} = w(h^1) \in W$. The play in later periods is determined recursively, using $v|_{h^t}$ as a state variable. Specifically, for each $t \geq 2$ and for each $h^{t-1} = (y^\tau)_{\tau=1}^{t-1} \in H^{t-1}$, given a $v|_{h^{t-1}} \in W$, let $\alpha|_{h^{t-1}}$ and $w|_{h^{t-1}} : Y \rightarrow W$ be such that $(\alpha|_{h^{t-1}}, v|_{h^{t-1}})$ is ex-post enforced by $w|_{h^{t-1}}$. Then, set the action profile after history h^{t-1} to be $s|_{h^{t-1}} = \alpha|_{h^{t-1}}$, and for each $y^t \in Y$, set $v|_{h^t = (h^{t-1}, y^t)} = w|_{h^{t-1}}(y^t) \in W$.

Because W is bounded and $\delta \in (0, 1)$, payoffs are continuous at infinity so finite approximations show that the specified strategy profile $s \in S$ generates v as an average payoff, and its continuation strategy $s|_{h^t}$ yields $v|_{h^t}$ for each $h^t \in H^t$. Also, by construction, nobody wants to deviate at any moment of time, given any state $\omega \in \Omega$. Because payoffs are continuous at infinity, the one-shot deviation principle applies, and we conclude that s is a PPXE, as desired. *Q.E.D.*

9.2 Proof of Theorem 2

Theorem 2. *If a subset W of $\mathbf{R}^{I \times |\Omega|}$ is compact, convex, and locally ex-post generating, then there is $\bar{\delta} \in (0, 1)$ such that $W \subseteq E(\delta)$ for all $\delta \in (\bar{\delta}, 1)$.*

Proof. Suppose that W is locally ex-post generating. Since $\{U_v\}_{v \in W}$ is an open cover of the compact set W , there is a subcover $\{U_{v^m}\}_m$ of W . Let $\bar{\delta} = \max_m \delta_{v^m}$. Choose $u \in W$ arbitrarily, and let U_{v^m} be such that $u \in U_{v^m}$. Since $W \cap U_{v^m} \subseteq B(\delta_{v^m}, W)$, there exist α_u and $w_u : Y \rightarrow W$ such that (α_u, u) is ex-post enforced by

the payoff vector v cannot pointwise dominate any point $v' \in V^*$. However, as Olivier Gossner pointed out, there may be priors such that the expected payoffs to a sequential (but not ex-post) equilibrium Pareto-dominates all of the ex-post equilibrium payoffs, essentially because revealing information destroys opportunities for insurance, as in Hirshleiffer (1971).

w_u for δ_{y^m} . Given a $\delta \in (\bar{\delta}, 1)$, let

$$w(y) = \frac{\delta - \delta_u}{\delta(1 - \delta_u)}u + \frac{\delta_u(1 - \delta)}{\delta(1 - \delta_u)}w_u(y)$$

for all $y \in Y$. Then, it is straightforward that (α_u, u) is enforced by $(w(y))_{y \in Y}$ for δ . Also, $w(y) \in W$ for all $y \in Y$, since u and $w(y)$ are in W and W is convex. Therefore, $u \in B(\delta, W)$, meaning that $W \subseteq B(\delta, W)$ for all $\delta \in (\bar{\delta}, 1)$. (Recall that u and δ are arbitrarily chosen from W and $(\bar{\delta}, 1)$.) Then, from Theorem 1, $W \subseteq E(\delta)$ for $\delta \in (\bar{\delta}, 1)$, as desired. *Q.E.D.*

9.3 Proof of Lemma 2

Lemma 2. *For every $\delta \in (0, 1)$, $E(\delta) \subseteq E^*(\delta) \subseteq Q$, where $E^*(\delta)$ is the convex hull of $E(\delta)$.*

Proof. It is obvious that $E(\delta) \subseteq E^*(\delta)$. Suppose $E^*(\delta) \not\subseteq Q$. Then, since the score is a linear function, there is $v \in E(\delta)$ and λ such that $\lambda \cdot v > k^*(\lambda)$. In particular, since $E(\delta)$ is compact, there exist $v^* \in E(\delta)$ and λ such that $\lambda \cdot v^* > k^*(\lambda)$ and $\lambda \cdot v^* \geq \lambda \cdot \tilde{v}$ for all $\tilde{v} \in E^*(\delta)$. By definition, v^* is enforced by $(w(y))_{y \in Y}$ such that $w(y) \in E(\delta) \subseteq E^*(\delta) \subseteq H(\lambda, \lambda \cdot v^*)$ for all $y \in Y$. But this implies that $k^*(\lambda)$ is not the maximum score for direction λ , a contradiction. *Q.E.D.*

9.4 An Ex-Post Folk Theorem with Perfect Monitoring

Claim 7. *Suppose that monitoring is perfect, that is, $Y = A \times \Omega$ and $\pi_y^\omega(a) = 1$ if $y = (a, \omega)$. Fix a payoff vector $v \in \text{int}V^*$. Then there is $\bar{\delta}$ such that for all $\delta \in (\bar{\delta}, 1)$ there is a PPXE where players play a pure action profile α in period one and then along the equilibrium path play $s^\omega(\delta)$ from period two, where $s^\omega(\delta)$ is a subgame-perfect equilibrium for state ω and discount factor δ with payoff v^ω .*

Proof. Let $v = (v^\omega)_{\omega \in \Omega} \in \text{int}V^*$, and let $\varepsilon > 0$ be such that for each ω , any payoff vector within ε of v^ω is in the set $V^*(\omega)$. Then let $\delta \in (0, 1)$ be such that (i) $\varepsilon > \frac{1-\delta}{\delta} \sum_{i \in I} (\max_{a \in A} g_i(a) - \min_{a \in A} g_i(a))$, (ii) for each ω , there is a subgame-perfect equilibrium s^{ω, v^ω} for state ω and discount factor δ with payoff v^ω , and (iii) for each $\omega \in \Omega$ and for any payoff vector \tilde{v}^ω within ε of v^ω , there is a subgame-perfect equilibrium $s^{\omega, \tilde{v}^\omega}$ for state ω and discount factor δ . Note that these three

conditions hold if δ is close to one; the last condition (iii) comes from Theorem 6.2 of FLM.

Consider the following strategy profile:

Phase 1 : Play a pure action profile a in period one. If there is no unilateral deviator from a and ω is observed, then go to Phase (ω, v^ω) . If player i unilaterally deviates from a and ω is observed, then go to Phase $(\omega, (v_i - \frac{1-\delta}{\delta}(\max_{\tilde{a} \in A} g_i(\tilde{a}) - \min_{\tilde{a} \in A} g_i(\tilde{a})), (v_j^\omega)_{j \neq i}))$.

Phase $(\omega, \tilde{v}^\omega)$ (Here, $\omega \in \Omega$ and \tilde{v}^ω is within ε of v^ω .) : Play a subgame perfect equilibrium $s^{\omega, \tilde{v}^\omega}$ in the remaining periods, as long as ω is observed in every period of this phase. (Recall that $s^{\omega, \tilde{v}^\omega}$ is a subgame-perfect equilibrium for state ω with payoffs \tilde{v}^ω .) If in any period t , $\omega^t \neq \omega^{t-1}$ then go to phase $(\omega^t, w^{\omega^t}(a^t))$ in the next period, where $w^{\omega^t}(a^t) = (w_i^{\omega^t}(a^t))_{i \in I}$ is chosen so that

$$w_i^{\omega^t}(a^t) = v_i^{\omega^t} + \frac{1-\delta}{\delta}(v_i^{\omega^t} - g_i^{\omega^t}(a^t))$$

for all $i \in I$.

This strategy profile is well-defined, as $w^{\omega^t}(a^t)$ is within ε of v^ω by construction, and it is easy to see that this strategy profile is a PPXE. *Q.E.D.*

9.5 Proof of Lemma 12

Lemma 12. *Suppose that SID holds. Then, $k^*(\lambda) = \infty$ for direction λ such that there exist $i \in I$, $j \in I$, $\omega \in \Omega$, and $\tilde{\omega} \neq \omega$ such that $\lambda_i^\omega \neq 0$ and $\lambda_j^{\tilde{\omega}} \neq 0$.*

Proof. Let (i, ω) and $(j, \tilde{\omega})$ be such that $\lambda_i^\omega \neq 0$, $\lambda_j^{\tilde{\omega}} \neq 0$, and $\tilde{\omega} \neq \omega$. Let α be a profile that is ex-post enforceable and statewise identifiable for (i, ω) and $(j, \tilde{\omega})$. In what follows, we show that $k^*(\alpha, \lambda) = \infty$.

First, we claim that for every $K > 0$, there exist $z_i^\omega = (z_i^\omega(y))_{y \in Y}$ and $z_j^{\tilde{\omega}} = (z_j^{\tilde{\omega}}(y))_{y \in Y}$ such that (1) holds for all $a_i \in A_i$, (2) holds for all $a_j \in A_j$, and (3) holds for all $y \in Y$. To prove that this system of equations indeed has a solution, let $A'_i \subseteq A_i$ provide a basis for the space spanned by $(\pi_y^\omega(a'_i, \alpha_{-i}))_{y \in Y}$, meaning that the set $\{(\pi_y^\omega(a'_i, \alpha_{-i}))_{y \in Y}\}_{a'_i \in A'_i}$ is a basis for the space, so that $\text{rank} \Pi'_{(i, \omega)}(\alpha) = \text{rank} \Pi_{(i, \omega)}(\alpha) = |A'_i|$. Then if (1) holds for all $a'_i \in A'_i$, then (1) for $a'_i \notin A'_i$ is satisfied as well. Likewise, let $A'_j \subseteq A_j$ provide a basis for the space spanned by $(\pi_y^{\tilde{\omega}}(a'_j, \alpha_{-j}))_{y \in Y}$ for all $a'_j \in A'_j$; if (2) holds for all $a'_j \in A'_j$, then (2) for $a'_j \notin A'_j$

is satisfied. Thus, for the above system to have a solution, it suffices that there exist z_i^ω and $z_j^{\tilde{\omega}}$ such that (1) holds for all $a_i \in A'_i$, (2) holds for all $a_j \in A'_j$, and (3) holds for all $y \in Y$. Eliminate (3) by solving for $z_j^{\tilde{\omega}}(y)$. Then, there remain $|A'_i| + |A'_j|$ linear equations, and its coefficient matrix is $\Pi'_{(i,\omega)(j,\tilde{\omega})}(\alpha)$, which is constructed by stacking two matrices $\Pi'_{(i,\omega)}(\alpha)$ and $\Pi'_{(j,\tilde{\omega})}(\alpha)$. It follows from statewise identifiability that

$$\text{rank}\Pi'_{(i,\omega)(j,\tilde{\omega})}(\alpha) = \text{rank}\Pi'_{(i,\omega)}(\alpha) + \text{rank}\Pi'_{(j,\tilde{\omega})}(\alpha) = |A'_i| + |A'_j|.$$

Therefore, we can indeed solve the system.

Let (\tilde{v}, \tilde{w}) be a pair of a payoff vector and a function such that \tilde{w} enforces (\tilde{v}, α) . Let $K > \max_{y \in Y} \lambda \cdot \tilde{w}(y) - \lambda \cdot \tilde{v}$, and choose z_i^ω and $z_j^{\tilde{\omega}}$ to satisfy (1) through (3). Then, let

$$w_l^{\bar{\omega}}(y) = \begin{cases} \tilde{w}_i^\omega(y) + z_i^\omega(y) & \text{if } (l, \bar{\omega}) = (i, \omega) \\ \tilde{w}_j^{\tilde{\omega}}(y) + z_j^{\tilde{\omega}}(y) & \text{if } (l, \bar{\omega}) = (j, \tilde{\omega}) \\ \tilde{w}_l^{\bar{\omega}}(y) & \text{otherwise} \end{cases}$$

for each $y \in Y$. Also, let

$$v_l^{\bar{\omega}} = \begin{cases} \tilde{v}_i + \frac{K}{\lambda_i^{\bar{\omega}}} & \text{if } (l, \bar{\omega}) = (i, \omega) \\ \tilde{v}_l^{\bar{\omega}} & \text{otherwise} \end{cases}.$$

Then, as in the proof of Lemma 6, this (v, w) satisfies constraints (i) through (iii) in LP-Average. Therefore, $k^*(\alpha, \lambda) \geq \lambda \cdot v = \lambda \cdot \tilde{v} + K$. Since K can be arbitrarily large, we conclude $k^*(\alpha, \lambda) = \infty$. *Q.E.D.*

9.6 Proof of Lemma 13

Lemma 13. *Suppose that α is ex-post enforceable and n -statewise distinguishes (i, ω) from $(j, \tilde{\omega})$. Then $k^*(\alpha, \lambda) = \infty$ for direction λ such that $\lambda_i^\omega > 0$ and $\lambda_j^{\tilde{\omega}} < 0$.*

Proof. Let $\xi = (\xi(y))_{y \in Y}$ be as in the definition of n -statewise distinguishability. Without loss of generality, assume $\xi \cdot \pi^{\tilde{\omega}}(\alpha) = 0$. Let $z_i^\omega = (z_i^\omega(y))_{y \in Y}$ and $z_j^{\tilde{\omega}} = (z_j^{\tilde{\omega}}(y))_{y \in Y}$ be such that

$$z_i^\omega(y) = \frac{K}{\delta \lambda_i^\omega \xi \cdot \pi^\omega(\alpha)} \xi(y)$$

and

$$z_j^{\tilde{\omega}}(y) = -\frac{K}{\delta\lambda_j^{\tilde{\omega}}\xi \cdot \pi^{\omega}(\alpha)}\xi(y)$$

for all $y \in Y$. Then, since $\xi \cdot \pi^{\omega}(\alpha) = \xi \cdot \pi^{\omega}(a_i, \alpha_{-i}) > 0$ for $a_i \in \text{supp}\alpha_i$, we have

$$\pi^{\omega}(a_i, \alpha_{-i}) \cdot z_i^{\omega} = \frac{K}{\delta\lambda_i^{\omega}\xi \cdot \pi^{\omega}(\alpha)}\pi^{\omega}(a_i, \alpha_{-i}) \cdot \xi = \frac{K}{\delta\lambda_i^{\omega}} \quad (5)$$

for all $a_i \in \text{supp}\alpha_i$. Also, since $\xi \cdot \pi^{\omega}(\alpha) > 0$ and $\xi \cdot \pi^{\omega}(\alpha) \geq \xi \cdot \pi^{\omega}(a_i, \alpha_{-i})$ for $a_i \notin \text{supp}\alpha_i$, we have

$$\pi^{\omega}(a_i, \alpha_{-i}) \cdot z_i^{\omega} = \frac{K}{\delta\lambda_i^{\omega}\xi \cdot \pi^{\omega}(\alpha)}\pi^{\omega}(a_i, \alpha_{-i}) \cdot \xi \leq \frac{K}{\delta\lambda_i^{\omega}} \quad (6)$$

for all $a_i \notin \text{supp}\alpha_i$. Likewise, since $\xi \cdot \pi^{\omega}(\alpha) > 0$, $\xi \cdot \pi^{\tilde{\omega}}(a_j, \alpha_{-j}) = 0$ for all $a_j \in \text{supp}\alpha_j$, and $\xi \cdot \pi^{\tilde{\omega}}(a_j, \alpha_{-j}) \leq 0$ for all $a_j \notin \text{supp}\alpha_j$,

$$\pi^{\tilde{\omega}}(a_j, \alpha_{-j}) \cdot z_j^{\tilde{\omega}} = -\frac{K}{\delta\lambda_j^{\tilde{\omega}}\xi \cdot \pi^{\omega}(\alpha)}\pi^{\tilde{\omega}}(a_j, \alpha_{-j}) \cdot \xi(y) = 0 \quad (7)$$

for all $a_j \in \text{supp}\alpha_j$, and

$$\pi^{\tilde{\omega}}(a_j, \alpha_{-j}) \cdot z_j^{\tilde{\omega}} = -\frac{K}{\delta\lambda_j^{\tilde{\omega}}\xi \cdot \pi^{\omega}(\alpha)}\pi^{\tilde{\omega}}(a_j, \alpha_{-j}) \cdot \xi \leq 0 \quad (8)$$

for all $a_j \notin \text{supp}\alpha_j$. Finally, it is obvious that

$$\lambda_i^{\omega}z_i^{\omega}(y) + \lambda_j^{\tilde{\omega}}z_j^{\tilde{\omega}}(y) = 0 \quad (9)$$

for all $y \in Y$.

Let (\tilde{v}, \tilde{w}) be a pair of a payoff vector and a function such that \tilde{w} enforces (\tilde{v}, α) . Let $K > \max_{y \in Y} \lambda \cdot \tilde{w}(y) - \lambda \cdot \tilde{v}$. Then, let

$$w_l^{\bar{\omega}}(y) = \begin{cases} \tilde{w}_i^{\omega}(y) + z_i^{\omega}(y) & \text{if } (l, \bar{\omega}) = (i, \omega) \\ \tilde{w}_j^{\tilde{\omega}}(y) + z_j^{\tilde{\omega}}(y) & \text{if } (l, \bar{\omega}) = (j, \tilde{\omega}) \\ \tilde{w}_l^{\tilde{\omega}}(y) & \text{otherwise} \end{cases}$$

for each $y \in Y$. Also, let

$$v_l^{\bar{\omega}} = \begin{cases} \tilde{v}_i^{\omega} + \frac{K}{\lambda_i^{\omega}} & \text{if } (l, \bar{\omega}) = (i, \omega) \\ \tilde{v}_l^{\tilde{\omega}} & \text{otherwise} \end{cases}.$$

We claim that this (v, w) satisfies all the constraints in LP-Average. Obviously, constraints (i) and (ii) are satisfied for all $(l, \bar{\omega}) \in (\mathbf{I} \times \Omega) \setminus \{(i, \omega), (j, \tilde{\omega})\}$, as $v_l^{\bar{\omega}} = \tilde{v}_i^{\bar{\omega}}$ and $w_l^{\bar{\omega}}(y) = \tilde{w}_l^{\bar{\omega}}(y)$. Also, since (5) and (6) hold and \tilde{w} enforces (α, \tilde{v}) , we obtain

$$\begin{aligned} & (1 - \delta)g_i^\omega(a_i, \alpha_{-i}) + \delta\pi^\omega(a_i, \alpha_{-i}) \cdot w_i^\omega \\ &= (1 - \delta)g_i^\omega(a_i, \alpha_{-i}) + \delta\pi^\omega(a_i, \alpha_{-i}) \cdot (\tilde{w}_i^\omega + z_i^\omega) \\ &= \tilde{v}_i^\omega + \frac{K}{\lambda_i^\omega} \\ &= v_i^\omega \end{aligned}$$

for all $a_i \in \text{supp}\alpha_i$, and

$$\begin{aligned} & (1 - \delta)g_i^\omega(a_i, \alpha_{-i}) + \delta\pi^\omega(a_i, \alpha_{-i}) \cdot w_i^\omega \\ &= (1 - \delta)g_i^\omega(a_i, \alpha_{-i}) + \delta\pi^\omega(a_i, \alpha_{-i}) \cdot (\tilde{w}_i^\omega + z_i^\omega) \\ &\leq \tilde{v}_i^\omega + \frac{K}{\lambda_i^\omega} \\ &= v_i^\omega \end{aligned}$$

for all $a_i \notin \text{supp}\alpha_i$. Hence, (v, w) satisfies constraints (i) and (ii) for (i, ω) . Likewise, it follows from (7) and (8) that (v, w) satisfies constraints (i) and (ii) for $(j, \tilde{\omega})$. Furthermore, using (9) and $K > \max_{y \in Y} \lambda \cdot \tilde{w}(y) - \lambda \cdot \tilde{v}$,

$$\begin{aligned} \lambda \cdot w(y) &= \lambda \cdot \tilde{w}(y) + \lambda_i^\omega z_i^\omega(y) + \lambda_j^{\tilde{\omega}} z_j^{\tilde{\omega}}(y) \\ &= \lambda \cdot \tilde{w}(y) \\ &< \lambda \cdot \tilde{v} + K \\ &= \lambda \cdot v \end{aligned}$$

for all $y \in Y$, and hence constraint (iii) holds.

Therefore, $k^*(\alpha, \lambda) \geq \lambda \cdot v = \lambda \cdot \tilde{v} + K$. Since K can be arbitrarily large, we conclude $k^*(\alpha, \lambda) = \infty$. *Q.E.D.*

9.7 Proof of Lemma 14

Lemma 14. *Suppose that α is ex-post enforceable and p -statewise distinguishes (i, ω) from $(j, \tilde{\omega})$. Then $k^*(\alpha, \lambda) = \infty$ for direction λ such that $\lambda_i^\omega > 0$ and $\lambda_j^{\tilde{\omega}} > 0$.*

Proof. Let $\xi = (\xi(y))_{y \in Y}$ be as in the definition of p -statewise distinguishability. Without loss of generality, assume $\xi \cdot \pi^{\tilde{\omega}}(\alpha) = 0$. Let

$$z_i^\omega(y) = \frac{K}{\delta \lambda_i^\omega \xi \cdot \pi^\omega(\alpha)} \xi(y)$$

and

$$z_j^{\tilde{\omega}}(y) = -\frac{K}{\delta \lambda_j^{\tilde{\omega}} \xi \cdot \pi^\omega(\alpha)} \xi(y)$$

for all $y \in Y$. Then, as in the proof of Lemma 13, we have (5) for all $a_i \in \text{supp}\alpha_i$, (6) for all $a_i \notin \text{supp}\alpha_i$, (7) for all $a_j \in \text{supp}\alpha_j$, (8) for all $a_j \notin \text{supp}\alpha_j$, and (9) for all $y \in Y$.

Let (\tilde{v}, \tilde{w}) be a pair of a payoff vector and a function such that \tilde{w} enforces (\tilde{v}, α) . Let $K > \max_{y \in Y} \lambda \cdot \tilde{w}(y) - \lambda \cdot \tilde{v}$, and choose (v, w) as in the proof of Lemma 13. Then, this (v, w) satisfies all the constraints in LP-Average, so that $k^*(\alpha, \lambda) \geq \lambda \cdot v = \lambda \cdot \tilde{v} + K$. Letting $K \rightarrow \infty$, we get $k^*(\alpha, \lambda) = \infty$. *Q.E.D.*

9.8 Proof of Claim 3

Claim 3. For $\alpha = (C_1, C_2)$,

$$k^*(\alpha, \lambda') \leq \lambda' \cdot g(C_1, C_2) - \frac{1-\beta}{\beta} (g_2^{\omega_2}(C_1, D_2) - g_2^{\omega_2}(C_1, C_2)).$$

Proof. Consider the associated LP-Average problem, and choose (v, w) to satisfy constraints (i) through (iii) of this problem. From player 2's IC constraint for state ω_2 , we have

$$\begin{aligned} \beta (q_H (w_2^{\omega_2}(H) - w_2^{\omega_2}(L)) + q_M (w_2^{\omega_2}(M) - w_2^{\omega_2}(L))) \\ \geq \frac{1-\delta}{\delta} (g_2^{\omega_2}(C_1, D_2) - g_2^{\omega_2}(C_1, C_2)). \end{aligned}$$

Then,

$$\begin{aligned}
v_1^{\omega_1} + v_2^{\omega_2} &= (1 - \delta)(g_1^{\omega_1}(C_1, C_2) + g_2^{\omega_2}(C_1, C_2)) \\
&\quad + \delta(\pi^{\omega_1}(C_1, C_2) \cdot w_1^{\omega_1} + \pi^{\omega_2}(C_1, C_2) \cdot w_2^{\omega_2}) \\
&= (1 - \delta)(g_1^{\omega_1}(C_1, C_2) + g_2^{\omega_2}(C_1, C_2)) + \delta\pi^{\omega_1}(C_1, C_2) \cdot (w_1^{\omega_1} + w_2^{\omega_2}) \\
&\quad - \delta(1 - \beta)(q_H(w_2^{\omega_2}(H) - w_2^{\omega_2}(L)) + q_M(w_2^{\omega_2}(M) - w_2^{\omega_2}(L))) \\
&\leq (1 - \delta)(g_1^{\omega_1}(C_1, C_2) + g_2^{\omega_2}(C_1, C_2)) + \delta(v_1^{\omega_1} + v_2^{\omega_2}) \\
&\quad - \frac{(1 - \delta)(1 - \beta)}{\beta}(g_2^{\omega_2}(C_1, D_2) - g_2^{\omega_2}(C_1, C_2))
\end{aligned}$$

Arranging,

$$v_1^{\omega_1} + v_2^{\omega_2} \leq g_1^{\omega_1}(C_1, C_2) + g_2^{\omega_2}(C_1, C_2) - \frac{1 - \beta}{\beta}(g_2^{\omega_2}(C_1, D_2) - g_2^{\omega_2}(C_1, C_2)).$$

So we have

$$\lambda \cdot v \leq \lambda \cdot g(C_1, C_2) - \frac{1 - \beta}{\beta}(g_2^{\omega_2}(C_1, D_2) - g_2^{\omega_2}(C_1, C_2)).$$

This proves the desired result. *Q.E.D.*

References

- Abreu, D., D. Pearce, and E. Stacchetti (1986): “Optimal Cartel Equilibria with Imperfect Monitoring,” *Journal of Economic Theory* 39, 251-269.
- Abreu, D., D. Pearce, and E. Stacchetti (1990): “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica* 58, 1041-1063.
- Athey, S., and K. Bagwell (2001): “Optimal Collusion with Private Information,” *RAND Journal of Economics* 32, 428-465.
- Aumann, R., and S. Hart (1992): *Handbook of Game Theory with Economic Applications, Volume 1*. North Holland, New York, NY.
- Aumann, R., and M. Maschler (1995): *Repeated Games with Incomplete Information*. MIT Press, Cambridge, MA. With the collaboration of R.E. Stearns.
- Bhaskar, V., G.J. Mailath, and S. Morris (2008): “Purification in the Infinitely Repeated Prisoner’s Dilemma,” *Review of Economic Dynamics* 11, 515-528.
- Bergemann, D., and S. Morris (2007): “Belief-Free Incomplete Information Games,” mimeo.
- Cripps, M., and J. Thomas (2003): “Some Asymptotic Results in Discounted Repeated Games of One-Side Incomplete Information,” *Mathematics of Operations Research* 28, 433-462.
- Deb, J. (2008): “Cooperation and Community Responsibility: A Folk Theorem for Repeated Random Matching Games,” mimeo.
- Dekel, E., D. Fudenberg, and D.K. Levine (2004): “Learning to Play Bayesian Games,” *Games and Economic Behavior* 46, 282-303.
- Ely, J., J. Hörner, and W. Olszewski (2005): “Belief-Free Equilibria in Repeated Games,” *Econometrica* 73, 377-415.
- Ely, J., and J. Välimäki (2002): “A Robust Folk Theorem for the Prisoner’s Dilemma,” *Journal of Economic Theory* 102, 84-105.

- Fudenberg, D., and D.K. Levine (1983): “Subgame-Perfect Equilibria of Finite and Infinite Horizon Games,” *Journal of Economic Theory* 31, 251-268
- Fudenberg, D., and D.K. Levine (1994): “Efficiency and Observability in Games with Long-Run and Short-Run Players,” *Journal of Economic Theory* 62, 103-135.
- Fudenberg, D., D.K. Levine, and E. Maskin (1994): “The Folk Theorem with imperfect public information,” *Econometrica* 62, 997-1040.
- Fudenberg, D., D.K. Levine, and S. Takahashi (2007): “Perfect Public Equilibrium when Players are Patient,” *Games and Economic Behavior* 61, 27-49.
- Fudenberg, D., and E. Maskin (1986): “The Folk Theorem in Repeated Games with Discounting and with Incomplete Information,” *Econometrica* 54, 533-554.
- Fudenberg, D., and Y. Yamamoto (2009): “Type-Contingent Perfect Public Ex-Post Equilibria,” mimeo.
- Gossner, O., and N. Vieille (2003): “Strategic Learning in Games with Symmetric Information,” *Games and Economic Behavior* 42, 25-47.
- Green, E.J., and R.H. Porter (1984): “Noncooperative Collusion under Imperfect Price Information,” *Econometrica* 52, 87-100.
- Hart, S. (1985): “Nonzero-Sum Two-Person Repeated Games with Incomplete Information,” *Mathematics of Operations Research* 10, 117-153
- Hirshleiffer, J. (1971): “The Private and Social Value of Information and the Reward to Inventive Activity,” *American Economic Review* 61, 561-574.
- Hörner, J., and S. Lovo (2009): “Belief-Free Equilibria in Games with Incomplete Information,” *Econometrica* 77, 453-487.
- Hörner, J., S. Lovo, and T. Tomala (2008): “Belief-Free Equilibria in Games with Incomplete Information: the N-Player Case,” mimeo.
- Kandori, M. (1992): “The Use of Information in Repeated Games with Imperfect Monitoring,” *Review of Economic Studies* 59, 581-593.

- Kandori, M. (2008): “Weakly Belief-Free Equilibria in Repeated Games with Private Monitoring,” mimeo.
- Kandori, M., and H. Matsushima (1998): “Private Observation, Communication and Collusion,” *Econometrica* 66, 627-652.
- Kosenok, G., and S. Severinov (2008): “Individually rational, budget-balanced mechanisms and allocation of surplus,” *Journal of Economic Theory* 140, 126-161.
- Kreps, D., and R. Wilson (1982): “Sequential Equilibria,” *Econometrica* 50, 863-894.
- Levin, J. (2003): “Relational Incentive Contracts,” *American Economic Review* 93, 835-857.
- Miller, D. (2007): “The Dynamic Cost of Ex Post Incentive Compatibility in Repeated Games of Private Information,” mimeo.
- Piccione, M. (2002): “The Repeated Prisoner’s Dilemma with Imperfect Private Monitoring,” *Journal of Economic Theory* 102, 70-83.
- Radner, R., R. Myerson, and E. Maskin (1986): “An Example of a Repeated Partnership Game with Discounting and with Uniformly Inefficient Equilibria,” *Review of Economic Studies* 53, 59-70.
- Rahman, D., and I. Obara (2008): “Secret Contracts for Efficient Partnerships,” mimeo.
- Takahashi, S. (2008): “Community Enforcement when Players Observe Partners’ Past Play,” mimeo.
- Wiseman, T. (2005): “A Partial Folk Theorem for Games with Unknown Payoff Distributions,” *Econometrica* 73, 629-645.
- Wiseman, T. (2008) “A Partial Folk Theorem for Games with Private Learning,” mimeo.
- Yamamoto, Y. (2007): “Efficiency Results in N Player Games with Imperfect Private Monitoring,” *Journal of Economic Theory* 135, 382-413.

Yamamoto, Y. (2009): "A Limit Characterization of Belief-Free Equilibrium Payoffs in Repeated Games," *Journal of Economic Theory* 144, 802-824.