Learning from Private Information in Noisy Repeated Games

Citation

Fudenberg, Drew, and Yuichi Yamamoto. 2011. Learning from private information in noisy repeated games. *Journal of Economic Theory* 146(5): 1733-1769.

Published Version

doi:10.1016/j.jet.2011.03.003

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:9962008>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Learning from Private Information in Noisy Repeated Games*

Drew Fudenberg[†] and Yuichi Yamamoto[‡]

First version: February 18, 2009

This version: December 17, 2010

Abstract

We study the perfect type-contingently public ex-post equilibrium (PTXE) of repeated games where players observe imperfect public signals of the actions played, and both the payoff functions and the map from actions to signal distributions depends on an unknown state. The PTXE payoffs when players are patient are determined by the solutions to a family of linear programming problems. Using this characterization, we develop conditions under which play can be as if the players have learned the state. We provide a sufficient condition for the folk theorem, and a characterization of the PTXE payoffs in games with a known monitoring structure.

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 the editor and anonymous referees for helpful comments, and NSF grant 0646816 for financial support.

[†]Department of Economics, Harvard University, dfudenberg@harvard.edu

[‡]Department of Economics, Harvard University, yamamot@fas.harvard.edu

1 Introduction

The fact that repeated interactions can allow new and more efficient equilibrium outcomes is one of game theory's most important insights. It has been shown to apply in a range of settings, including games with imperfect public information about opponents' actions, and games where the monitoring structure- the map from actions to signal distributions- is itself unknown.¹ It has also been shown in games with private information about the payoff functions.² This paper studies games with the combination of these features: both the monitoring structure and the payoff functions can depend on an unknown state variable, and the players may have initial private information about the state. This describes, for example, a repeated partnership game where players observe group output but do not observe each other's effort, and each player has private information about the effect of her effort on the probability distribution of output.

Our main goal in this paper is to understand how the information structure of the game- meaning the combination of the monitoring structure and the initial private information- determines the extent to which the player's initial private information can be revealed in equilibrium. We address this question indirectly, by computing the limit of the equilibrium payoffs when players are patient. More specifically, we restrict attention to the *perfect type-contingently public ex-post equilibrium* or *PTXE* (Fudenberg and Yamamoto [17]). These are ex-post equilibria where each player's strategy depends only on the realized public outcomes and his initial private information (hence "type-contingent") but not on the player's private information about his own past actions.

PTXE generalizes several solution concepts: It reduces to the PPXE of [17] if players have no private information, the belief-free equilibria of [21] and [22] when actions are observed,³ and the perfect public equilibrium (PPE) of Fuden-

¹For repeated games with public monitoring, see Green and Porter [19], Radner [29], Abreu, Pearce, and Stacchetti [1,2], Fudenberg and Levine [14], Fudenberg, Levine, and Maskin [15], Athey and Bagwell [3], and Fudenberg, Levine, and Takahashi [16]. Fudenberg and Yamamoto [17] consider games where the monitoring structure is unknown.

²See Kohlberg [27], Forges [12], Sorin [32,33], Hart [20], Aumann and Maschler [5], Cripps and Thomas [8], Gossner and Vieille [18], Renault and Tomala [31], Wiseman [34,35], Hörner and Lovo [21], and Hörner, Lovo, and Tomala [22] for games with private information.

³These equilibria are different than the belief-free equilibria of repeated games with private monitoring (Piccione [28], Ely and Välimäki [10], Ely, Hörner, and Olszewski [9], Yamamoto [36,37], and Kandori [25]), which require that players be indifferent.

berg, Levine, and Maskin [15] (hereafter FLM) in complete-information games with a known monitoring structure. As with ex-post equilibria more generally, these equilibria are robust to the specification of the players' prior beliefs: a PTXE for a given prior distribution is a PTXE for an arbitrary prior.⁴

Any PPXE of the symmetric information game (where no player has initial private information about the state) induces a PPXE of the game where some players do have private information: these PPXE correspond to pooling equilibria of the incomplete-information game. Thus the folk theorems of [17] apply to games with private information. However, those theorems require that the distribution of signals vary with the state in a sufficiently rich way (essentially so that the state can be learned from the signals generated by some fixed action profile), and this is more restrictive than necessary when some players have private information. For example, if one player knows the state, he may be able to communicate it to the others using a strategy that conditions on the player's private information. This paper takes the possibility of such implicit communication into account, and so generates a larger set of equilibrium payoffs. In some cases, such as the partnership games we define in Section 3, in which a player's productivity is private information, there often exist asymptotically efficient equilibria, while equilibrium payoffs are bounded away from efficiency if the players ignore their private information.

Moreover, we can characterize the limit payoffs of PTXE with linear programming techniques. Specifically, the set of limit equilibrium payoffs is the intersection of maximal half-spaces in various directions, where the direction vectors λ assign weights on each player's payoff in each state, the maximal half-space in direction λ is all vectors v with $\lambda \cdot v$ no greater than the maximum score for λ , and this score is the highest weighted sum of payoffs that can be obtained with continuation payoffs that satisfy the incentive constraints and whose weighted sum is no higher than the sum they are supporting.

Roughly speaking, there are PTXE where players learn the state if the score is sufficiently large in "cross-state" directions that give non-zero weight to two or more states. For this to be the case, informed players must be willing to reveal their information, and uninformed players must not "jam" the information revelation of their informed opponents. A key point is that the relevant conditions

⁴See Bergemann and Morris [6] for a discussion of various definitions of ex-post equilibrium.

depend on whether the informed player's payoff in a given state is given positive or negative weight. With a positive weight the informed player wants to reveal the state, and our conditions imply that other players cannot prevent this; with a negative weight the informed player might prefer to hide the state, but under our conditions this is not possible.

We use these results to prove a folk theorem. While the exact conditions are complicated to state, the key assumption is that for each pair of players i and j (where possibly $i = j$) and each pair of states ω and $\omega' \neq \omega$, either (i) there is a player $l \neq i, j$ whose private information distinguishes ω and ω' , and player l can reveal this information regardless of the actions of i and j by choosing different actions in state ω and state ω' , (ii) player i or j (or both) can distinguish ω and ω' using initial private information, and the informed player is willing to reveal this information while the other one cannot interfere, or (iii) there is an action profile α (independent of the private information) that distinguishes (more formally, “statewise identifies”) ω from ω' . Conditions (i) and (ii) lead to a sort of “endogenous learning” where players transmit their private information to the opponents, while condition (iii) is a sort of “exogenous learning” based on the distribution of signals at a fixed action profile. Note that condition (i) does not require that player l is willing to reveal his information. This is because the condition can be used for directions where player l 's continuation payoff has zero weight and hence is unconstrained, and our “individual full rank” assumption ensures that there is some specification of the continuation payoffs that induces player l to play the specified actions in the two states. In contrast, if no third player can distinguish the states, then the incentives of the revealing player become relevant, as without additional conditions it may be that any continuation payoffs that induce player i to reveal his information must increase or decrease player j 's continuation payoff in a way that lowers the score.

We then consider a few cases with additional structure that simplifies our characterization. We begin with the case where the state space has one component that only influences payoffs and a second component that only influences the monitoring structure; here we show that when the full rank conditions are satisfied the limit set can be determined for each payoff function separately. Next we consider games with a product structure, where there is a separate and independent signal associated with each player's action, and moreover each player knows the effect

of his action on the signal distribution while the others do not. For example, in a game of bilateral production and exchange, the public signal might be the quality of a player's output, with each player having private information about the probability that she will make a high-quality good when she exerts high effort. Here we show that the scores for two classes of cross-state directions are high enough to be compatible with the folk theorem, but that the scores in the remaining class need not be. Finally, as an illustration of our characterization, we examine in detail a repeated partnership example where only group output is observed, and the state determines the productivity of player 2. We show that if player 1's private information reveals player 2's productivity while 2 has no private information (i.e. "1 knows 2's productivity"), then the folk theorem holds in general, while if only player 2 knows player 2's productivity, the folk theorem can fail, and moreover the limit equilibrium payoffs can be bounded away from efficiency. Intuitively, player 2 cannot be induced to reveal the state when doing so would lower his equilibrium payoff, and this leads to a bound on the extent to which equilibria can trade off player 2's payoffs between the two states; in some cases this bound is so strong that it rules out the efficient outcome.

Finally, we specialize to the case of a known monitoring structure, where we show that the set of limit equilibrium payoffs with imperfectly observed actions is the same as in the observed-action case studied by [21] and [22] provided that the monitoring structure satisfies a full-rank condition. [22] provide an equivalent characterization (for observed actions) that has a much different form; each characterization may be better suited for some applications. Our results show that their conclusions about limit payoffs extend to imperfectly observed actions; their work is complementary and more informative because it also explicitly constructs equilibrium strategies. The assumption of a known monitoring structure also lets us provide a sufficient condition for the folk theorem that is easier to verify: the key is that for every pair of states ω and ω' , there be at least three players whose private information distinguishes between ω and ω' ; [22] use this same condition to show in games with observed actions the set of ex-post perfect equilibria is non-empty. In the case of one-sided incomplete information, we are able to further extend and refine their results; for example, we find a simpler sufficient condition for the existence of PTXE.

2 Framework

2.1 Model

Let $I = \{1, \dots, I\}$ be the set of players. At the beginning of the game, Nature chooses the state of the world ω from a finite set $\Omega = \{\omega_1, \dots, \omega_O\}$. Then each player observes a private signal, which gives (possible imperfect) information about the true state ω . The set of player i 's private signals, Θ_i , is a partition of Ω , and given the true state $\omega \in \Omega$, he observes a private signal $\theta_i \in \Theta_i$ that contains ω . For notational convenience, let $\theta_i(\omega)$ denote this θ_i , i.e., $\omega \in \theta_i(\omega)$, and let $\theta(\omega) = (\theta_i(\omega))_{i \in I}$. Given $\theta_i \in \Theta_i$, player i forms a prior about the true state ω , which is denoted by $\mu_i(\theta_i) \in \Delta \Theta_i$.

Each period, players move simultaneously, and player $i \in I$ chooses an action a_i from a finite set A_i .⁵ 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 structure.” Player i 's realized payoff is $u_i^\omega(a_i, y)$, so that her expected payoff conditional on $\omega \in \Omega$ and $a \in A$ 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 . If there are $\omega' \neq \omega$ such that $\theta_i(\omega) = \theta_i(\omega')$ and $u_i^\omega(a_i, y) \neq u_i^{\omega'}(a_i, y)$ for some $a_i \in A_i$ and $y \in Y$, then we assume that player i does not observe the realized value of u_i as the game is played.⁶ If there are no such $\omega' \neq \omega$, it is immaterial whether or not u_i is observed, as player i can compute it from a_i , y , and θ_i .⁷

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 from period one to period $t \geq 1$ by $h_i^t = (a_i^\tau, y^\tau)_{\tau=1}^t$. Let $h_i^0 = \emptyset$, and for each $t \geq 0$, 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

⁵All of our results extend immediately to the case where A_i depends on θ_i .

⁶As we explain in the next section, the equilibria we consider remain equilibria when players are provided with additional channels of information about the state. Thus the assumption that players do not observe their realized payoffs has no role in the results; it allows us to generalize past work (such as most of the references in footnote 2) that did not require players observe their realized payoffs.

⁷We call this the case of known own payoffs; note that it does not imply that each player i knows the their stage-game payoff function g_i^ω as that payoff is an expected value with respect to the possibly unknown distribution π^ω .

of all h^t . A strategy for player i is defined to be a mapping $s_i : \Theta_i \times \bigcup_{t=0}^{\infty} H_i^t \rightarrow \Delta A_i$. Let S_i be the set of all strategies for player i , and let $S = \times_{i \in I} S_i$.

We define the feasible payoff set 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 ,⁸ and we define the set of feasible payoffs of the overall game to be

$$V \equiv \times_{\omega \in \Omega} V^\omega.$$

Note that a feasible payoff vector $v \in V$ may be generated using different action distributions η^ω 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 state-contingent strategy of the infinitely repeated game.

2.2 Preliminaries

Player i 's strategy $s_i \in S_i$ is *type-contingently public* if it depends only on $\theta_i \in \Theta_i$ and $h^t \in H^t$, that is, if $s_i(\theta_i, h_i^t) = s_i(\theta_i, \tilde{h}_i^t)$ whenever h_i^t and $\tilde{h}_i^t$ correspond to the same public history. A strategy profile $s \in S$ is type-contingently public if s_i is type-contingently public for each $i \in I$. Given a type-contingently public strategy profile $s \in S$, let $s_i|_{(\theta_i, h^t)}$ denote player i 's continuation strategy when his type is θ_i and the past public history is h^t , and let $s|_{(\theta, h^t)} = (s_i|_{(\theta_i, h^t)})_{i \in I}$.⁹ This paper studies a special class of Nash equilibria called *perfect type-contingently public ex-post equilibria* or *PTXE*.

Definition 1. A strategy profile $s \in S$ is a *perfect type-contingently public ex-post equilibrium (PTXE)* if s is type-contingently public, and if for any $\omega \in \Omega$ and $h^t \in H^t$, $s|_{(\theta(\omega), h^t)}$ is a Nash equilibrium of the infinitely repeated game with $\Omega = \{\omega\}$.

Remark 1. PTXE is an ex-post equilibrium concept in the sense that it requires each player's strategy is a best response irrespective of the true value of the state.

⁸As in the standard case of a game with a known monitoring structure, the feasible set V^ω 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.

⁹Here, the word "continuation strategy" is an abuse of language, because $s_i|_{(\theta_i, h^t)}$ is not a strategy for the entire game; it specifies a play for a given type θ_i but not for $\theta_i \neq \theta_i$.

For this reason, the set of PTXE is independent of the players' beliefs about the state, which makes the analysis of equilibria much simpler. The ex-post property also implies that a PTXE for a given partition Θ is also a PTXE for any finer partition; in particular a PTXE for the trivial partition (where players have no private information) remains a PTXE when Θ is informative. The PPXE we considered in [17] are the same as the PTXE with the trivial partition; the point of this paper is that finer partitions on Θ can support a larger set of PTXE, as in the examples of Section 6.3, where there are efficient limit PTXE but the PPXE are bounded away from efficiency uniformly in δ .

Remark 2. A second consequence of the ex-post nature of PTXE is that a PTXE of the game where players do not observe their realized stage game payoffs remains a PTXE if players do observe these realized payoffs and the payoffs reveal information about ω . That said, additional equilibrium outcomes could arise here under a relaxed equilibrium definition that allowed players to condition on this additional private information. We do not investigate that possibility in this paper.

Given a discount factor $\delta \in (0, 1)$, let $E(\delta)$ denote the set of PTXE 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 PTXE s satisfying $(1 - \delta)E[\sum_{t=1} \delta^{t-1} g_i^\omega(a^t) | s, \omega] = v_i^\omega$ for all i and ω . Note that $v \in E(\delta)$ specifies the equilibrium payoff for all players and all states.

Let $\vec{\alpha}_i = (\alpha_i^{\theta_i})_{\theta_i \in \Theta_i}$ where $\alpha_i^{\theta_i} \in \Delta A_i$ for each $\theta_i \in \Theta_i$, and let $\vec{\alpha} = (\vec{\alpha}_i)_{i \in I}$. Thus $\vec{\alpha}$ is an action profile contingent on private information; it specifies a mixed action α_i for each private signal θ_i of each player i . Let $g(\vec{\alpha}) = (g_i^\omega(\alpha^{\theta(\omega)}))_{(i,\omega)}$ denote the payoff vector of type-contingent profile $\vec{\alpha}$. If the action profile α is used independently of private information, we denote its payoff vector by $g(\alpha) = (g_i^\omega(\alpha))_{(i,\omega)}$.

By definition, any continuation strategy $s|_{h^t} = (s|_{\theta(\omega), h^t})_{\omega \in \Omega}$ of a PTXE is also a PTXE. Thus any PTXE specifies PTXE 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 the payoffs for every player and every state. We will write $\pi^\omega(\alpha) \cdot w_i^\omega$ for the expected continuation payoff at state ω under action profile α .

In [17], we showed that the limit of the equilibrium payoffs as $\delta \rightarrow 1$ is determined by the solutions $k^*(\vec{\alpha}, \lambda, \delta)$ to the following family of linear programming problems; for each type-contingent action profile $\vec{\alpha}$, direction $\lambda \in \mathbf{R}^{I \times |\Omega|} \setminus \{0\}$,

and $\delta \in (0, 1)$,

$$\begin{aligned}
k^*(\vec{\alpha}, \lambda, \delta) &= \max_{\substack{v \in \mathbf{R}^{I \times |\Omega|} \\ w: Y \rightarrow \mathbf{R}^{I \times |\Omega|}}} \lambda \cdot v \quad \text{subject to} \\
\text{(i)} \quad v_i^\omega &= (1 - \delta)g_i^\omega(\alpha^{\theta(\omega)}) + \delta\pi^\omega(\alpha^{\theta(\omega)}) \cdot w_i^\omega \quad (1) \\
&\quad \text{for all } i, \omega, \\
\text{(ii)} \quad v_i^\omega &\geq (1 - \delta)g_i^\omega(a_i, \alpha_{-i}^{\theta_{-i}(\omega)}) + \delta\pi^\omega(a_i, \alpha_{-i}^{\theta_{-i}(\omega)}) \cdot w_i^\omega \quad (2) \\
&\quad \text{for all } i, \omega, \text{ and } a_i \in A_i, \\
\text{(iii)} \quad \lambda \cdot v &\geq \lambda \cdot w(y) \quad \text{for all } y \in Y.
\end{aligned}$$

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

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 $\lambda_j^{\omega'} \neq 0$ for some $\omega \neq \omega'$, condition (iii) allows “utility transfer” across states.

As argued in [17], the score $k^*(\vec{\alpha}, \lambda, \delta)$ is independent of δ , so we denote it by $k^*(\vec{\alpha}, \lambda)$. Let $k^*(\lambda) = \sup_{\vec{\alpha}} k^*(\vec{\alpha}, \lambda)$ be the highest score attainable in direction λ for any choice of $\vec{\alpha}$. For each $\lambda \in \mathbf{R}^{I \times |\Omega|} \setminus \{0\}$ and $k \in \mathbf{R}$, let $H(\lambda, k) = \{v \in \mathbf{R}^{I \times |\Omega|} \mid \lambda \cdot v \leq k\}$, with $H(\lambda, k) = \mathbf{R}^{I \times |\Omega|}$ for $k = \infty$ or $\lambda = 0$, and $H(\lambda, k) = \emptyset$ for $k = -\infty$ and $\lambda \neq 0$. Now let

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

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

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

The following proposition establishes that the intersection Q of the maximal half-spaces is equal to the limit set of PTXE payoffs as $\delta \rightarrow 1$. The proof is omitted, as it is similar to [17], which builds on the techniques of [14].

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

Our goal in this paper is to use this characterization to compute $\lim_{\delta \rightarrow 1} E(\delta)$ in some cases of interest. To do this we provide conditions under which the maximal half-spaces in the various directions are “large.”

3 Examples

Before developing our general results, we provide a few examples of PTXE to illustrate the ways that players can “learn the state” in equilibrium.

Example 1. Let $I = \{1, 2\}$ and $\Omega = \{\omega_1, \omega_2\}$, $\Theta_1 = \{(\omega_1), (\omega_2)\}$ and $\Theta_2 = \{(\omega_1, \omega_2)\}$. Player 1 chooses either U or D , and player 2 chooses either L or R . The payoffs for state ω_1 are in the left panel, and those for state ω_2 are in the right.

	L	R		L	R
U	2, 2	0, 1	U	1, 1	0, 1
D	1, 0	1, 1	D	1, 0	2, 2

In this example, both (U, L) and (D, R) are static ex-post equilibria.

Assume that $Y = A$ and $\pi_y^\omega(a) = \varepsilon$ if $y \neq a$. Note that the signal distribution does not depend on the state here, so that players cannot learn the state from state-independent actions. Instead, the efficient outcome $((2, 2), (2, 2))$ can be approximated if player 1 reveals his private information to player 2 through his actions. Specifically, consider the following three-phase automaton.

- Phase 1. Player 1 chooses U if $\theta_1 = (\omega_1)$, and D if $\theta_1 = (\omega_2)$. Player 2 chooses L . If the observed signal is $y = (U, L)$ or $y = (D, R)$, then go to Phase 2. If $y = (D, L)$, then go to Phase 3. If $y = (U, R)$, stay.
- Phase 2. Players choose (U, L) in the rest of the game.
- Phase 3. Players choose (D, R) in the rest of the game.

We claim that the strategy profile with initial state Phase 1 is a PTXE if δ is close to one and ε is close to zero. First, players do not want to deviate in Phase 2 or Phase 3, as (U, L) and (D, R) are static ex-post equilibria. Also, player 1 with $\theta_1 = (\omega_1)$ does not want to deviate in Phase 1. Indeed, if he deviates

to D , then players are likely to go to Phase 3 and play (D, R) forever, while if he does not deviate, then players are likely to go to Phase 2 so that (U, L) is played thereafter. Likewise, we can check that player 1 with $\theta_1 = (\omega_2)$ does not want to deviate in Phase 1. Player 2's prescribed play is always a static best response, and since 2's play has no effect on the transitions between stages 2 does not want to deviate either. Note that the payoffs of this equilibrium converge to $((2 - \varepsilon, 2 - \varepsilon), (2 - \varepsilon, 2 - \varepsilon))$.

Example 2. The next example is a two-player partnership game with two actions $\{C_i, D_i\}$ per player, three possible outcomes H, M, L , and two states.¹⁰ 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 $i \in I$, $\omega \in \Omega$, and $y \in Y$. We assume that 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. We impose restrictions on the realized payoffs 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.¹¹

Using our results, we will show that a folk theorem holds if player 1 knows the state and player 2 does not, but that PTXE payoffs are bounded away from efficiency for some parameters if player 2 knows the state and player 1 does not. The key is that player 2 can learn whether the true state is ω_1 or ω_2 by playing C_2 no matter what player 1 does, since player 2's marginal productivity is dependent on the state but not on player 1's action. Thus for the case in which player only 1

¹⁰If there were only two outcomes as in Radner, Myerson, and Maskin [30], then payoffs are bounded away from efficiency even if the state is known, while with three outcomes the folk theorem holds for generic signal distributions as FLM shows.

¹¹Specifically we assume $r_i(H) > r_i(M) > r_i(L)$; $e_1 > p_H(r_1(H) - r_1(L)) + p_M(r_1(M) - r_1(L))$; $e_2 > q_H(r_2(H) - r_2(L)) + q_M(r_2(M) - r_2(L))$; $e_1 < p_H(r_1(H) + r_2(H) - r_1(L) - r_2(L)) + p_M(r_1(M) + r_2(M) - r_1(L) - r_2(L))$; and $e_2 < \beta q_H(r_1(H) + r_2(H) - r_1(L) - r_2(L)) + \beta q_M(r_1(M) + r_2(M) - r_1(L) - r_2(L))$.

knows the state, even if player 1 tries to hide his private information, player 2 can learn the true state from the signal distribution. On the other hand, if only player 2 knows the state and he tries to hide it then player 1 cannot learn the true state. That is, if player 2 chooses D_2 , then for given any player 1's action, the signal distribution is the same for both states, and hence player 1 cannot learn from the observed signals. See Section 6.3 for details.

4 Sufficient Conditions for Efficient State Learning

In this section we develop “distinguishability” conditions that are sufficient for limit equilibria in which payoffs are as if players have learned the true state. In Section 4.2, we relate these conditions to the incentives and information of the players; roughly speaking, the distinguishability conditions are equivalent to assuming that if informed players are willing to reveal the state then uninformed players cannot prevent them from doing so. When the distinguishability conditions are satisfied, the maximal half-spaces in “cross-state” directions (those that give non-zero weights to payoffs in two or more states) are the whole space, so the cross-state directions impose no constraints on the limit equilibrium payoffs. The maximal half-spaces in directions that give non-zero weights to a single state are the same as in the known-state case considered by FLM, so combining FLM's assumptions, our distinguishability assumptions, and Proposition 1 establishes the existence of limit equilibria with the desired properties.

4.1 Statewise Full Rank and Statewise Distinguishability

We begin with the statewise full rank condition, which is sufficient for the maximal score to be infinity for all cross-state directions. For each $(i, \omega) \in \mathbf{I} \times \Omega$ and each type-contingent action profile $\vec{\alpha} \in \times_{i \in \mathbf{I}} \times_{\theta_i \in \Theta_i} \Delta A_i$, let $\Pi_{(i, \omega)}(\vec{\alpha})$ be a matrix with rows $(\pi_y^\omega(a_i, \alpha_{-i}^{\theta_{-i}(\omega)}))_{y \in Y}$ for all $a_i \in A_i$. Let $\Pi_{(i, \omega)(j, \omega')}(\vec{\alpha})$ be a matrix constructed by stacking two matrices, $\Pi_{(i, \omega)}(\vec{\alpha})$ and $\Pi_{(j, \omega')}(\vec{\alpha})$.

Definition 2. For each (i, ω) and (j, ω') satisfying $\omega \neq \omega'$, profile $\vec{\alpha}$ has *statewise full rank for (i, ω) and (j, ω')* if $\Pi_{(i, \omega)(j, \omega')}(\vec{\alpha})$ has rank $|A_i| + |A_j|$.

Statewise full rank implies that players can distinguish ω and ω' even if player

i at state ω or player j at state ω' deviates. For each pair (i, ω) and (j, ω') , there is more likely to be a type-contingent profile $\vec{\alpha}$ that has statewise full rank as the partitions Θ_l for player l become finer. The intuition is that if player l has more information, then it is easier for the players to learn the true state through inferences based on player l 's actions. Note that PPXE does not allow players to condition their play on their types, so it rules out this information channel. We say more about learning from other players' actions in Section 4.2.

The next lemma shows that statewise full rank is sufficient for the maximal score to be infinity for all cross-state directions. We say that a type-contingent profile $\vec{\alpha}$ is *ex-post enforceable* if there are $v \in \mathbf{R}^{I \times |\Omega|}$, $\delta \in (0, 1)$, and $w = (w^\omega)_{\omega \in \Omega} : Y \rightarrow \mathbf{R}^{I \times |\Omega|}$ such that (1) holds for all i and ω , and (2) holds for all i , ω , and a_i . The proof of the lemma is omitted, as it is very similar to Lemma 6 of [17].

Lemma 1. *Suppose profile $\vec{\alpha}$ is ex-post enforceable and has statewise full rank for (i, ω) and (j, ω') satisfying $\omega \neq \omega'$. Then, $k^*(\vec{\alpha}, \lambda) = \infty$ for λ such that $\lambda_i^\omega \neq 0$ and $\lambda_j^{\omega'} \neq 0$.*

While the statewise full rank condition is sufficient for efficient learning, it requires at least $|A_i| + |A_j|$ signals. The following condition, statewise distinguishability, can be satisfied with fewer signals and is sufficient for the maximal score to be infinity for all cross-state directions that have at least one positive component. We will soon relax this condition even further, but this definition is a useful expositional tool for explaining the more complicated definitions to come.

Definition 3. Profile $\vec{\alpha}$ *statewise distinguishes* (i, ω) from (j, ω') if there is $\xi = (\xi(y))_{y \in Y} \in \mathbf{R}^{|Y|}$ such that

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

To interpret this condition, without loss of generality we assume $\pi^{\omega'}(\alpha^{\theta(\omega')}) \cdot \xi = 0$. Clause (i) of this condition assures that the signals generated by $\vec{\alpha}$ statistically distinguish ω from ω' , and moreover picks out a direction ξ where the

difference has a particular sign. Clause (ii) says that changing player i 's continuation payoff function in state ω from $w_i^\omega(y)$ 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 ($\lambda_i^\omega \Delta w_i^\omega(y) + \lambda_j^{\omega'} \Delta w_j^{\omega'}(y) = 0$) without changing player j 's expected continuation payoff to any action. Since clause (i) implies $\pi^\omega(\alpha^{\theta(\omega)}) \cdot \xi > 0$, this change in the continuation payoffs increases player i 's expected continuation payoff at state ω , which implies an increase in the score for λ such that $\lambda_i^\omega > 0$. Note that this definition is not symmetric between i and j because condition (ii) is an inequality and condition (iii) is an equality. When this condition is satisfied, scaling up the vector ξ can generate arbitrarily large scores for all cross-state directions λ that have at least one positive component.

Our next step is to replace statewise distinguishability with an ensemble of three weaker conditions- this ensemble is weaker because it will allow different action profiles to be used in different directions.

Definition 4. Profile $\vec{\alpha}$ m -statewise distinguishes (i, ω) from (j, ω') if there is $\xi = (\xi(y))_{y \in Y} \in \mathbf{R}^{|Y|}$ such that

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

Note that this condition relaxes statewise distinguishability by replacing the last equality in (iii) with an inequality. Lemma 4(a) below shows that a profile that m -statewise distinguishes (i, ω) from (j, ω') can be used to generate an infinite score for all λ such that $\lambda_i^\omega > 0$ and $\lambda_j^{\omega'} < 0$; the “ m ” refers to the fact that positive and negative components are “mixed” in these directions.

Definition 5. Profile $\vec{\alpha}$ p -statewise distinguishes (i, ω) from (j, ω') if there is $\xi = (\xi(y))_{y \in Y} \in \mathbf{R}^{|Y|}$ such that

- (i) $\pi^\omega(\alpha^{\theta(\omega)}) \cdot \xi > \pi^{\omega'}(\alpha^{\theta(\omega')}) \cdot \xi$,

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

Lemma 4(b) below shows that a profile that p -statewise distinguishes (i, ω) from (j, ω') can be used to generate an infinite score for all “positive” directions λ such that $\lambda_i^\omega > 0$ and $\lambda_j^{\omega'} > 0$. As this suggests, this condition is symmetric:

Lemma 2. *Suppose $\vec{\alpha}$ p -statewise distinguishes (i, ω) from (j, ω') . Then $\vec{\alpha}$ p -statewise distinguishes (j, ω') from (i, ω) .*

Proof. Let ξ be a vector utilized to p -statewise distinguish (i, ω) from (j, ω') . Then the vector $-\xi$ satisfies all the conditions of p -statewise distinguishability of (j, ω') from (i, ω) . *Q.E.D.*

Note that if $\vec{\alpha}$ statewise distinguishes (i, ω) from (j, ω') , then it m -statewise distinguishes this pair and p -statewise distinguishes this pair.

As we will explain later, the combination of m - and p -statewise distinguishability is sufficient for a static-threat folk theorem. However, it is not sufficient for a perfect folk theorem, because the maximal score might not be high enough in cross-state directions where all the non-zero components are negative. The following condition is sufficient for the score to be infinitely large for these directions.

Definition 6. Profile $\vec{\alpha}$ n -statewise distinguishes (i, ω) from (j, ω') if there is $\xi = (\xi(y))_{y \in Y} \in \mathbf{R}^{|Y|}$ such that

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

Lemma 4(c) below shows that a profile that n -statewise distinguishes (i, ω) from (j, ω') can be used to generate an infinite score for all “negative” directions λ such that $\lambda_i^\omega < 0$ and $\lambda_j^{\omega'} < 0$. Also, n -statewise distinguishability is symmetric, as the next lemma shows. We omit the proof, since it is very similar to that of Lemma 2.

Lemma 3. *Suppose $\vec{\alpha}$ n -statewise distinguishes (i, ω) from (j, ω') . Then $\vec{\alpha}$ n -statewise distinguishes (j, ω') from (i, ω) .*

Now we state the main result of this section, which shows that the score for cross-state directions can be infinity if the corresponding statewise condition is satisfied. The proof can be found in the appendix.

Lemma 4.

- (a) *Suppose $\vec{\alpha}$ is ex-post enforceable and m -statewise distinguishes (i, ω) from (j, ω') . Then $k^*(\alpha, \lambda) = \infty$ for λ such that $\lambda_i^\omega > 0$ and $\lambda_j^{\omega'} < 0$.*
- (b) *Suppose $\vec{\alpha}$ is ex-post enforceable and p -statewise distinguishes (i, ω) from (j, ω') . Then $k^*(\alpha, \lambda) = \infty$ for λ such that $\lambda_i^\omega > 0$ and $\lambda_j^{\omega'} > 0$.*
- (c) *Suppose $\vec{\alpha}$ is ex-post enforceable and n -statewise distinguishes (i, ω) from (j, ω') . Then $k^*(\alpha, \lambda) = \infty$ for λ such that $\lambda_i^\omega < 0$ and $\lambda_j^{\omega'} < 0$.*

4.2 Sufficient Conditions for Statewise Distinguishability

In games with incomplete information, players have three possible sources of information about the state: (i) inference based on the public signals at a state-independent action profile; (ii) the information contained in their own types; and (iii) inferences based on the correlation between the opponents’ actions and the opponents’ types. The first information source is studied by [17]. The second information source is sufficient for perfect learning if every player can distinguish ω and ω' . (Note that this corresponds to assumption (ii) of condition (SFR) in Section 5.) Here we investigate the third information source: inferences based on the correlation between the opponents’ actions and the opponents’ types. For this information to generate large scores in cross-state directions, the informed player must be willing to reveal his information, and uninformed players must not “jam”

the information revelation of their informed opponents. We address these issues by providing simple sufficient conditions under which a type-contingent action profile satisfies the various distinguishability conditions.

Definition 7. *Player i can reveal whether ω or ω' if there are $a \in A$ and $a'_i \in A_i$ such that $\pi^\omega(a) \neq \pi^{\omega'}(a'_i, a_{-i})$.*

This says that player i can generate different signal distributions at ω and ω' , using a type-contingent action. Note that this is necessary for players $-i$ to learn the state from the correlation between player i 's actions and his types. As the next lemma shows, this condition is sufficient for p -statewise distinguishability for (i, ω) and (i, ω') .

Lemma 5. *Suppose $\theta_i(\omega) \neq \theta_i(\omega')$ and player i can reveal whether ω or ω' . Then there is $\vec{\alpha}$ that p -statewise distinguishes (i, ω) from (i, ω') .*

Proof. Let $a \in A$ and $a'_i \in A_i$ be such that $\pi^\omega(a) \neq \pi^{\omega'}(a'_i, a_{-i})$. Then there is $\xi \in \mathbf{R}^{|Y|}$ such that $\pi^\omega(a) \cdot \xi > \pi^{\omega'}(a'_i, a_{-i}) \cdot \xi$. Let $a_i^* \in \arg \max_{a_i''} \pi^\omega(a_i'', a_{-i}) \cdot \xi$ and $a_i^{**} \in \arg \min_{a_i''} \pi^{\omega'}(a_i'', a_{-i}) \cdot \xi$. Let $\vec{\alpha}$ be a type-contingent action profile such that players play (a_i^*, a_{-i}) at state ω and (a_i^{**}, a_{-i}) at state ω' . Then this $\vec{\alpha}$ p -statewise distinguishes (i, ω) from (i, ω') . Indeed, clause (i) follows from $\pi^\omega(a_i^*, a_{-i}) \cdot \xi \geq \pi^\omega(a) \cdot \xi > \pi^{\omega'}(a'_i, a_{-i}) \cdot \xi \geq \pi^{\omega'}(a_i^{**}, a_{-i}) \cdot \xi$. Also, clauses (ii) and (iii) hold, by definition of a_i^* and a_i^{**} . *Q.E.D.*

To get the intuition, recall that p -statewise distinguishability is relevant to directions λ that put positive weights on payoff for (i, ω) and (i, ω') . In these directions, player i 's payoffs at ω and ω' are both maximized, so she is willing to reveal her information at both states.

In contrast, even if player i can reveal whether ω or ω' there might be no profile that m -statewise distinguishes (i, ω) from (i, ω') . The reason is that if λ puts negative weight on payoffs at (i, ω') , player i 's payoffs at ω' is minimized in the corresponding LP problem so that he might not want to reveal the true state. However, the following condition is sufficient for m -statewise distinguishability; the idea is that player i at state ω' cannot conceal his private information if he cannot generate the same signal distribution as in state ω .

Definition 8. *Player i at ω' cannot hide state ω if there is $a \in A$ such that $\pi^\omega(a)$ is not in the convex hull of $\{\pi^{\omega'}(a'_i, a_{-i})\}_{a'_i \in A_i}$.*

Lemma 6. Suppose $\theta_i(\omega) \neq \theta_i(\omega')$ and player i at ω' cannot hide state ω . Then there is $\vec{\alpha}$ that m -statewise distinguishes (i, ω) from (i, ω') .

Proof. Let $a \in A$ be such that $\pi^\omega(a)$ is not in the convex hull of $\{\pi^{\omega'}(a'_i, a_{-i})\}_{a'_i \in A_i}$. Then from the separating hyperplane theorem, there is ξ such that $\pi^\omega(a) \cdot \xi > \pi^{\omega'}(a'_i, a_{-i}) \cdot \xi$ for all $a'_i \in A_i$. Let $a_i^* \in \arg \max_{a'_i} \pi^\omega(a'_i, a_{-i}) \cdot \xi$ and $a_i^{**} \in \arg \max_{a'_i} \pi^{\omega'}(a'_i, a_{-i}) \cdot \xi$. Let $\vec{\alpha}$ be a type-contingent action profile such that players play (a_i^*, a_{-i}) at state ω and (a_i^{**}, a_{-i}) at state ω . We claim that this $\vec{\alpha}$ m -statewise distinguishes (i, ω) from (i, ω') . Clause (i) follows from $\pi^\omega(a_i^*, a_{-i}) \cdot \xi \geq \pi^\omega(a) \cdot \xi > \pi^{\omega'}(a_i^{**}, a_{-i}) \cdot \xi$. Also, clauses (ii) and (iii) hold, by definition of a_i^* and a_i^{**} . *Q.E.D.*

A similar idea applies to n -statewise condition; here a relevant direction λ puts negative weights on payoffs at (i, ω) and (i, ω') , so we need to take into account player i 's incentive for information revelation at both states.

Definition 9. Player i cannot shuffle states ω and ω' if there is $a \in A$ such that the convex hull of $\{\pi^\omega(a'_i, a_{-i})\}_{a'_i \in A_i}$ and the convex hull of $\{\pi^{\omega'}(a'_i, a_{-i})\}_{a'_i \in A_i}$ do not intersect.

Lemma 7. Suppose $\theta_i(\omega) \neq \theta_i(\omega')$ and player i cannot shuffle states ω and ω' . Then there is $\vec{\alpha}$ that n -statewise distinguishes (i, ω) from (i, ω') .

Proof. Let $a \in A$ be such that the convex hull of $\{\pi^\omega(a'_i, a_{-i})\}_{a'_i \in A_i}$ and the convex hull of $\{\pi^{\omega'}(a'_i, a_{-i})\}_{a'_i \in A_i}$ do not intersect. Then from the separating hyperplane theorem, there is ξ such that $\pi^\omega(a'_i, a_{-i}) \cdot \xi > \pi^{\omega'}(a''_i, a_{-i}) \cdot \xi$ for all $a'_i \in A_i$ and $a''_i \in A_i$. Let $a_i^* \in \arg \min_{a'_i} \pi^\omega(a'_i, a_{-i}) \cdot \xi$ and $a_i^{**} \in \arg \max_{a'_i} \pi^{\omega'}(a'_i, a_{-i}) \cdot \xi$. Let $\vec{\alpha}$ be a type-contingent action profile such that players play (a_i^*, a_{-i}) at state ω and (a_i^{**}, a_{-i}) at state ω . Then as in the proof of the last lemma, we can show that this $\vec{\alpha}$ n -statewise distinguishes (i, ω) from (i, ω') . *Q.E.D.*

Next we consider statewise distinguishability for (i, ω) and (j, ω') where $i \neq j$ and only player i knows the state; i.e., $\theta_i(\omega) \neq \theta_i(\omega')$ and $\theta_j(\omega) = \theta_j(\omega')$.

Definition 10. Player j at state ω' is irrelevant for (i, ω) if there are $a \in A$ and $a'_i \in A_i$ such that $\pi^\omega(a)$ is not a linear combination of $\{\pi^{\omega'}(a'_i, a'_j, a_{-ij})\}_{a'_j \in A_j}$.

This says that there is an action profile a such that if player i wants to reveal whether the true state is ω or ω' by choosing a_i at ω and a'_i at ω' , the

uninformed player j cannot interfere, in the sense that a change in player j 's action at state ω' cannot result in the overall distribution when i plays a'_i being the same as the distribution in ω under a . For an example where this condition fails, suppose that there are two players with $A_i = \{a'_i, a''_i\}$ for each i , and that $\pi^\omega(a'_1, a'_2) = \pi^\omega(a''_1, a''_2) = \pi^{\omega'}(a'_1, a'_2) = \pi^{\omega'}(a''_1, a'_2)$ and $\pi^\omega(a'_1, a'_2) = \pi^\omega(a''_1, a'_2) = \pi^{\omega'}(a'_1, a'_2) = \pi^{\omega'}(a''_1, a''_2)$. Here player j at state ω' is not irrelevant for (i, ω) . On the other hand this condition can be satisfied even if player j has an option to “jam” player i 's information revelation: Suppose that there are two states, ω_1 and ω_2 , and that player 1 knows the state while other players do not. Let $A_1 = \{U, D\}$ and $A_2 = \{J, NJ\}$. Suppose that player 1's action is observable if player 2 chooses NJ , while it is unobservable if player 2 chooses J . Suppose that player l 's actions are always observable for each $l \neq 1$. Let a be an action profile such that $a_1 = U$ and $a_2 = NJ$, and let $a'_1 = D$. Then $\pi^{\omega_1}(a)$ is not a linear combination of $\{\pi^{\omega_2}(a'_1, a'_2, a_{-12})\}_{a'_2 \in A_2}$, so that player 2 at state ω_2 is irrelevant for $(1, \omega_1)$.

The following lemma shows that this irrelevance condition is sufficient for statewise distinguishability (and hence sufficient for p - and m -statewise distinguishability).

Lemma 8. *Suppose $\theta_i(\omega) \neq \theta_i(\omega')$ and player $j \neq i$ at ω' is irrelevant for (i, ω) . Then there is $\vec{\alpha}$ that statewise distinguishes (i, ω) from (j, ω') .*

Proof. Let $a \in A$ and $a'_i \in A_i$ be such that $\pi^\omega(a)$ is not a linear combination of $\{\pi^{\omega'}(a'_i, a'_j, a_{-ij})\}_{a'_j \in A_j}$. Then there is ξ such that $\pi^\omega(a) > 0$ and $\pi^{\omega'}(a'_i, a'_j, a_{-ij}) \cdot \xi = 0$ for all $a'_j \in A_j$. Let $a_i^* \in \arg \max_{a''_i \in A_i} \pi^\omega(a''_i, a_{-i}) \cdot \xi$, and let $\vec{\alpha}$ be a type-contingent action profile such that players play (a_i^*, a_{-i}) at state ω and (a'_i, a_{-i}) at state ω' . We claim that this $\vec{\alpha}$ statewise distinguishes (i, ω) from (j, ω') . Clause (ii) of statewise distinguishability follows from $a_i^* \in \arg \max_{a''_i \in A_i} \pi^\omega(a''_i, a_{-i}) \cdot \xi$. Also, since $\pi^\omega(a_i^*, a_{-i}) \geq \pi^\omega(a) > 0$ and $\pi^{\omega'}(a'_i, a'_j, a_{-ij}) \cdot \xi = 0$ for all $a'_j \in A_j$, clause (i) and (iii) hold. *Q.E.D.*

The intuition is as follows. Recall that statewise distinguishability is a combination of p - and m -statewise distinguishability, so the corresponding direction λ gives positive weight to player i 's payoff at state ω . Therefore, player i at state ω is willing to reveal his private information. Also, player j at state ω' cannot interfere with this information revelation.

For n -statewise distinguishability, we need a stronger assumption, since the corresponding λ puts negative weight on player i 's payoff at state ω , so that he might want to hide his private information.

Definition 11. *Player j at state ω' is strongly irrelevant for (i, ω) if there are $(a_i, a_{-i}) \in A$ such that any convex combination of $\{\pi^\omega(a'_i, a_{-i})\}_{a'_i \in A_i}$ is not a linear combination of $\{\pi^{\omega'}(a_i, a'_j, a_{-ij})\}_{a'_j \in A_j}$.*

This condition is a combination of “cannot hide” and “irrelevant” conditions: here player i at state ω cannot conceal his private information and player j at state ω' is irrelevant to player i 's information revelation.

Lemma 9. *Suppose $\theta_i(\omega) \neq \theta_i(\omega')$ and player $j \neq i$ at ω' is strongly irrelevant for (i, ω) . Then there is $\vec{\alpha}$ that n -statewise distinguishes (i, ω) from (j, ω') .*

Proof. Let $a \in A$ and $a'_i \in A_i$ be such that any convex combination of $\{\pi^\omega(a''_i, a_{-i})\}_{a''_i \in A_i}$ is not a linear combination of $\{\pi^{\omega'}(a'_i, a'_j, a_{-ij})\}_{a'_j \in A_j}$. Then there is ξ such that $\pi^\omega(a''_i, a_{-i}) > 0$ for all $a''_i \in A_i$ and $\pi^{\omega'}(a'_i, a'_j, a_{-ij}) \cdot \xi = 0$ for all $a'_j \in A_j$. Let $a_i^* \in \arg \min_{a''_i \in A_i} \pi^\omega(a''_i, a_{-i}) \cdot \xi$, and let $\vec{\alpha}$ be a type-contingent action profile such that players play (a_i^*, a_{-i}) at state ω and (a'_i, a_{-i}) at state ω' . Then as in Lemma 8, we can show that this $\vec{\alpha}$ n -statewise distinguishes (i, ω) from (j, ω') . *Q.E.D.*

When player j knows the state and player i does not (i.e., $\theta_j(\omega) \neq \theta_j(\omega')$ and $\theta_i(\omega) = \theta_i(\omega')$), the statewise conditions are satisfied under the same conditions as the case where player i knows the state. Recall that p - and n -statewise conditions are symmetric (Lemmas 2 and 3), so the sufficient conditions for the previous case apply. For the m -statewise condition, we obtain the following lemma.

Lemma 10. *Suppose $\theta_j(\omega) \neq \theta_j(\omega')$ and player $i \neq j$ at ω is strongly irrelevant for (j, ω') . Then there is $\vec{\alpha}$ that m -statewise distinguishes (i, ω) from (j, ω') .*

Proof. Let $a \in A$ and $a'_j \in A_j$ be such that any convex combination of $\{\pi^{\omega'}(a''_j, a_{-j})\}_{a''_j \in A_j}$ is not a linear combination of $\{\pi^\omega(a'_i, a'_j, a_{-ij})\}_{a'_i \in A_i}$. Then there is ξ such that $\pi^{\omega'}(a''_j, a_{-j}) < 0$ for all $a''_j \in A_j$ and $\pi^\omega(a'_i, a'_j, a_{-ij}) \cdot \xi = 0$ for all $a'_i \in A_i$. Let $a_j^* \in \arg \max_{a''_j \in A_j} \pi^{\omega'}(a''_j, a_{-j}) \cdot \xi$, and let $\vec{\alpha}$ be a type-contingent action profile such that players play (a'_j, a_{-j}) at state ω and (a_j^*, a_{-j}) at state ω' . Then we can show that this $\vec{\alpha}$ n -statewise distinguishes (i, ω) from (j, ω') . *Q.E.D.*

Finally we consider pairs (i, ω) and (j, ω') where there is a player $l \neq i, j$ who knows the state (here possibly $i = j$). If either player i or j can distinguish ω from ω' , then the previous lemmas still apply. Thus the interesting case is when both player i and j do not know the state.

Definition 12. Both player i at ω and player j at ω' are irrelevant for information revelation by player $l \neq i, j$ if there are $a \in A$ and $a'_l \in A_l$ such that any linear combination of $\{\pi^\omega(a'_i, a_l, a_{-il})\}_{a'_i \in A_i}$ is not a linear combination of $\{\pi^{\omega'}(a'_j, a'_l, a_{-jl})\}_{a'_j \in A_j}$.

This says that if player l wants to reveal his private information then neither player i at ω nor player j at ω' can interfere. The next lemma shows that this condition is sufficient for p -, m -, and n -statewise distinguishability.

Lemma 11. Suppose there is $i \in I$, $j \in I$, $l \neq i, j$ such that $\theta_l(\omega) \neq \theta_l(\omega')$ and that both player i at ω and player j at ω' are irrelevant for information revelation by player l . Then there is an $\vec{\alpha}$ that p -, m -, and n -statewise distinguishes (i, ω) from (j, ω') .

Proof. Let $a \in A$ and $a'_l \in A_l$ be such that any linear combination of $\{\pi^\omega(a'_i, a_l, a_{-il})\}_{a'_i \in A_i}$ is not a linear combination of $\{\pi^{\omega'}(a'_j, a'_l, a_{-jl})\}_{a'_j \in A_j}$. Then there are ξ and $\kappa > 0$ such that $\pi^\omega(a'_i, a_l, a_{-il}) = \kappa$ for all $a'_i \in A_i$ and $\pi^{\omega'}(a'_j, a'_l, a_{-jl}) \cdot \xi = 0$ for all $a'_j \in A_j$. Let $\vec{\alpha}$ be a type-contingent action profile such that players play a at state ω and (a'_l, a_{-l}) at state ω' . Then it is easy to check that this $\vec{\alpha}$ satisfies all the conditions of p -, m -, and n -statewise distinguishability. *Q.E.D.*

5 Ex-Post Folk Theorems

In this section we provide two sorts of folk theorem in PTXE: The first shows that all feasible individually rational payoffs can be approximated by payoffs of PTXE, and the second uses weaker conditions to obtain a “static-threats” version. In both cases, the key is finding the appropriate conditions on the combination of initial private information and the information revealed by the public outcomes.

Recall that $\Pi_{(i, \omega)}(\vec{\alpha})$ is a matrix with rows $(\pi_y^\omega(a_i, \alpha_{-i}^{\theta_{-i}(\omega)}))_{y \in Y}$ for all $a_i \in A_i$, and that $\Pi_{(i, \omega)(j, \omega')}(\vec{\alpha})$ is a matrix constructed by stacking two matrices, $\Pi_{(i, \omega)}(\vec{\alpha})$ and $\Pi_{(j, \omega')}(\vec{\alpha})$.

Definition 13. Profile $\vec{\alpha}$ has *individual full rank* for (i, ω) if $\Pi_{(i, \omega)}(\vec{\alpha})$ has rank $|A_i|$. Profile $\vec{\alpha}$ has *individual full rank* if it has individual full rank for all players and all states.

This condition implies that at each state, every possible deviation of any one player leads to a statistically different distribution on outcomes.

Definition 14. For each (i, ω) and (j, ω) satisfying $i \neq j$, profile $\vec{\alpha}$ has *pairwise full rank* for (i, ω) and (j, ω) if $\Pi_{(i, \omega)(j, \omega)}(\vec{\alpha})$ has rank $|A_i| + |A_j| - 1$.

Note that pairwise full rank implies individual full rank; it implies that deviations by one player can be distinguished from deviations by another.

Condition IFR. Every pure action profile $\vec{\alpha}$ has individual full rank.

Condition PFR. For each (i, ω) and (j, ω) satisfying $i \neq j$, there is a profile $\vec{\alpha}$ that has pairwise full rank for (i, ω) and (j, ω) .

Condition SFR. For each pair of states (ω, ω') satisfying $\omega \neq \omega'$, at least one of the following two conditions holds: (i) For each i and j (possibly $i = j$), there is a profile $\vec{\alpha}$ that has statewise full rank for (i, ω) and (j, ω') , or (ii) $\theta_l(\omega) \neq \theta_l(\omega')$ for all $l \in I$.

(SFR) requires that for each pair of states ω and $\omega' \neq \omega$, either (i) for every (i, j) there is a profile that lets players distinguish state ω from state ω' , regardless of whether player i deviates in state ω or player j deviates in state ω' , or (ii) players can distinguish these ω and ω' using their private information θ .

Note that (SFR) fails for (i, ω) and (i, ω') if π^ω is independent of ω (so that the monitoring structure is known) and $\theta_j(\omega) = \theta_j(\omega')$ for all j (so no player's private information distinguishes between ω and ω'). We say more about the case of a known monitoring structure in Section 7.

The next proposition establishes a general folk theorem in PTXE. Let $V^* \equiv \{v \in V \mid \forall i \in I \forall \omega \in \Omega \ v_i^\omega \geq \underline{v}_i^\omega\}$ where $\underline{v}_i^\omega = \min_{\alpha_{-i}} \max_{a_i} g_i^\omega(a_i, \alpha_{-i})$. A subset W of $\mathbf{R}^{I \times |\Omega|}$ is *smooth* if it is closed and convex; it has a non-empty interior; and there is a unique unit normal for each point on $\text{bd}W$.¹²

¹²A 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 $\mathbf{R}^{I \times |\Omega|}$.

Proposition 2. *Suppose (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 $\delta \in (\bar{\delta}, 1)$.*

To prove this proposition, we compute the maximal scores for each direction. The key point is that (SFR) implies the maximal score for cross-state directions can be made large enough to establish the folk theorem. When the first condition in (SFR) holds, that conclusion comes from Lemma 1. When the second condition holds, the following lemma applies:

Lemma 12. *Suppose (PFR) and (IFR) hold. Let λ be such that $\theta_i(\omega) \neq \theta_i(\omega')$ for all $i \in I$, $\omega \in \Omega$ and $\omega' \neq \omega$ satisfying $(\lambda_j^\omega)_{j \in I} \neq 0$ and $(\lambda_j^{\omega'})_{j \in I} \neq 0$. Then $k^*(\lambda) \geq \max_{v \in V^*} \lambda \cdot v$.*

This lemma shows that the maximal score in cross-state directions doesn't exclude any feasible payoffs if all players know the state. The intuition behind the lemma is simple. If each player i can distinguish ω and ω' using private information θ_i , players can choose different action profiles contingent on whether the true state is ω or ω' . Therefore we expect that the score on state ω will not constrain the score on state ω' so that the maximal score for directions vectors that only weight these two states will be high enough to achieve the folk theorem. The formal proof is delegated to the appendix.

Combining this lemma and Lemma 1 shows that the maximum score in all cross-state directions is at least $\max_{v \in V^*} \lambda \cdot v$. This implies that the set Q is determined by λ that has non-zero components only for a single state. The following lemmas show that (IFR) and (PFR) imply that the maximal score for such directions is $\max_{v \in V^*} \lambda \cdot v$. The proofs are omitted, as they are straightforward generalizations of FLM.

Lemma 13. *Suppose (PFR) holds. Then $k^*(\lambda) = \max_{v \in V^*} \lambda \cdot v$ for all λ such that (i) $(\lambda_i^\omega)_{i \in I} \neq 0$ for some ω and $(\lambda_i^{\omega'})_{i \in I} = 0$ for all $\omega' \neq \omega$, and (ii) $(\lambda_i^\omega)_{i \in I}$ has at least two non-zero components or at least one positive component.*

Lemma 14. *Suppose (IFR) holds. Then $k^*(\lambda) = \max_{v \in V^*} \lambda \cdot v$ for all λ such that $\lambda_i^\omega < 0$ for some (i, ω) and $\lambda_j^{\omega'} = 0$ for all $(j, \omega') \neq (i, \omega)$.*

From these lemmas, we obtain $Q = V^*$ and hence Proposition 2 follows. Thus the folk theorem obtains if (IFR), (PFR), and (SFR) hold and if V^* is full dimensional.

As we have seen in Section 4.1, statewise full rank is stronger than needed for efficient learning, and can be replaced with statewise distinguishability.

Condition Pointwise-SD. For each ω and ω' satisfying $\omega \neq \omega'$, at least one of the following conditions holds: (i) For each i and j (possibly $i = j$), there is an ex-post enforceable action profile $\vec{\alpha}$ that m -statewise distinguishes (i, ω) from (j, ω') , there is an ex-post enforceable action profile $\vec{\alpha}'$ that p -statewise distinguishes (i, ω) from (j, ω') , and there is an ex-post enforceable action profile $\vec{\alpha}''$ that n -statewise distinguishes (i, ω) from (j, ω') , or (ii) $\theta_l(\omega) \neq \theta_l(\omega')$ for all $l \in I$.

This says that for each pair of states ω and $\omega' \neq \omega$, either (i) for every (i, j) there is a profile that lets players distinguish state ω from state ω' , regardless of whether player i deviates in state ω or player j deviates in state ω' , or (ii) players can distinguish these ω and ω' using their private information θ .

Note that (Pointwise-SD) is weaker than (SFR), since if $\vec{\alpha}$ has statewise full rank, then it satisfies the m -, p -, and n -statewise distinguishability conditions, but the converse is false, as the pointwise condition allows different profiles to be used for different directions. On the other hand, (Pointwise-SD) is a “strong” form of statewise distinguishability as it requires the n -statewise condition.

Proposition 3. *Suppose (IFR), (PFR), and (Pointwise-SD) 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 proof of this proposition parallels to that of Proposition 2, with the difference that Lemma 4 is used instead of Lemma 1 for the conclusion that the maximum scores in cross-state directions is infinite.

An even weaker condition is sufficient for a static-threat folk theorem: For that result it is sufficient that the m - and p -statewise conditions can each be satisfied for some profile.

Condition Pointwise-WeakSD. For each ω and ω' satisfying $\omega \neq \omega'$, at least one of the following conditions holds: (i) For each i and j (possibly $i = j$), there is an ex-post enforceable action profile $\vec{\alpha}$ that m -statewise distinguishes (i, ω) from (j, ω') and there is an ex-post enforceable action profile $\vec{\alpha}'$ that p -statewise distinguishes (i, ω) from (j, ω') , or (ii) $\theta_l(\omega) \neq \theta_l(\omega')$ for all $l \in I$.

Lemmas 5, 6, 8, and 11 imply that (Pointwise-WeakSD) holds if the condition (i), (ii), or (iii) we stated in the introduction is satisfied.

Proposition 4. *Suppose (PFR) and (Pointwise-WeakSD) hold. Assume that there is an ex-post equilibrium $\vec{\alpha}^0$, i.e., $\vec{\alpha}$ such that $\alpha_i^{\theta_i(\omega)} \in \arg \max_{\alpha_i} g_i^\omega(\alpha_i, \alpha_{-i}^{\theta_{-i}(\omega)})$ for all $i \in \mathbf{I}$ and $\omega \in \Omega$. 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 $\delta \in (\bar{\delta}, 1)$.*

The proof of this proposition is similar to that of Proposition 3, with the following differences. In this proposition, we do not assume (IFR) or (Pointwise-SD), so that Lemma 4(c) and Lemma 14 may not apply. Therefore, it might be that $k^*(\lambda) < \max_{v \in V^*} \lambda \cdot v$ for each $\lambda \neq 0$ such that $\lambda_i^\omega \leq 0$ for all (i, ω) . For these directions, we apply the next lemma to show that $k^*(\lambda) \geq \max_{v \in V^0} \lambda \cdot v$. The proof is straightforward and hence omitted.

Lemma 15. *Suppose there is a static ex-post equilibrium $\vec{\alpha}^0$. Then, for any direction λ , $k^*(\vec{\alpha}^0, \lambda) \geq \lambda \cdot g(\vec{\alpha}^0)$.*

Also, since Proposition 4 does not assume (IFR), Lemma 12 does not apply, so it might be that $k^*(\lambda) < \max_{v \in V^*} \lambda \cdot v$ for some cross-state directions λ . For these directions, we use the following lemma to show that $k^*(\lambda) \geq \max_{v \in V^0} \lambda \cdot v$. The proof of the lemma can be found in the appendix.

Lemma 16. *Suppose (PFR) holds. Let λ be such that $\theta_i(\omega) \neq \theta_i(\omega')$ for all $i \in \mathbf{I}$, $\omega \in \Omega$ and $\omega' \neq \omega$ satisfying $(\lambda_j^\omega)_{j \in \mathbf{I}} \neq 0$ and $(\lambda_j^{\omega'})_{j \in \mathbf{I}} \neq 0$. Then, $k^*(\lambda) \geq \max_{v \in V^0} \lambda \cdot v$.*

6 Applications and Examples

This section explores the effect of some plausible assumptions about the monitoring structure. The first two cases are fairly general; the third illustrates how to apply the general results by determining the limit payoffs in the partnership game described in Section 3.

6.1 Separable State Space

In general, the set of limit payoffs depends on the state's impact on both the monitoring structure and the payoff functions. When these dependencies are separable, the characterization of the limit set can be simplified. To see this, suppose that the state consists of two components, one that influences realized payoffs but not the monitoring structure and one that influences the monitoring structure but has no effect on the realized payoffs. That is, $\Omega = \Phi \times \Psi$, where $u_i^\omega(a_i, y) = u_i^{\omega'}(a_i, y)$ if $\phi = \phi'$, and $\pi_y^\omega(a) = \pi_y^{\omega'}(a)$ if $\psi = \psi'$. For example, this separability condition is consistent with a quantity-setting oligopoly game where players do not know the probability distribution of the market price, and each player knows their own cost function but not the cost function of the opponents: Here the price is the public signal y , and the state (cost parameters) has no effect on the distribution of prices at fixed output levels. Similar examples arise in partnership games where players know their own effort cost but not that of the opponents.

Condition ψ -SFR. For each (i, ω) and (j, ω') satisfying $\psi \neq \psi'$, there is an ex-post enforceable profile $\vec{\alpha}$ that has statewise full rank for (i, ω) and (j, ω') .

For each $\psi \in \Psi$, let $Q(\psi)$ denote the set Q for the known monitoring structure game corresponding to ψ , i.e., the game where the state space is restricted to $\Omega = \Phi \times \{\psi\}$ and the payoff functions u_i^ω and the monitoring structure π^ω for a given $\omega \in \Omega$ are the same as those of the original game. The next proposition shows that the equilibrium payoff set Q of the entire game is a product of $Q(\psi)$ over all ψ . Proposition 7 in Section 7 gives a formula to calculate each set $Q(\psi)$ when the monitoring structure π^ω has “strong full rank.”

Proposition 5. *Suppose the state space Ω is separable and (ψ -SFR) holds. Then $Q = \times_{\psi \in \Psi} Q(\psi)$.*

Proof. As Lemma 1 shows, if a profile $\vec{\alpha}$ is ex-post enforceable and has statewise full rank for (i, ω) and (j, ω') satisfying $\omega \neq \omega'$, then $k^*(\vec{\alpha}, \lambda) = \infty$ for direction λ such that $\lambda_i(\omega) \neq 0$ and $\lambda_j(\omega') \neq 0$. Thus from (ψ -SFR), $k^*(\lambda) = \infty$ for all λ such that $\lambda_i(\omega) \neq 0$ and $\lambda_j(\omega') \neq 0$ for (ω, ω') satisfying $\psi \neq \psi'$. This proves $Q = \times_{\psi \in \Psi} Q(\psi)$. Q.E.D.

6.2 Games with a Product Structure

In this section, we briefly discuss games with a product structure, in which there is a separate signal y_i associated with the action of each player i , each player i knows the distribution of “her” signal, and no player $j \neq i$ has any private information about the distribution of y_i . This case is of interest as a number of economic situations have this extra structure; it applies for example to bilateral production and exchange, where the public signal is the quality of a player’s output, and each player has private information about the probability that she will make a high-quality good when she exerts high effort. This structure on its own is not sufficient for the various folk theorems in Section 5, because m -statewise distinguishability can fail, but there is full statewise distinguishability in all of the other cross-state directions.

Formally, we assume that (i) $Y = \times_{i \in I} Y_i$; (ii) $\Omega = \times_{i \in I} \Omega_i$; (iii) $\sum_{y_{-i} \in Y_{-i}} \pi_y^\omega(a) = \sum_{y_{-i} \in Y_{-i}} \pi_y^{\omega'}(a)$ for each $i \in I$, $a \in A$, $y_i \in Y_i$, $\omega \in \Omega$, and $\omega' \in \Omega$ such that $\omega_i = \omega'_i$; and (iv) $\pi_y^\omega(a) = \prod_{i \in I} \pi_{y_{-i} \in Y_{-i}}^\omega(a)$ for each $a \in A$, $y \in Y$, and $\omega \in \Omega$. Note that the distribution of y_i depends only on a_i and ω_i here. We also assume that $\Theta_i = \{\theta_i^{\omega_i} | \omega_i \in \Omega_i\}$ where $\theta_i^{\omega_i} = \{\omega' | \omega'_i = \omega_i\}$; that is, player i knows the distribution of y_i but not the distribution of y_{-i} . We also assume that every state has some impact on the distribution of signals in the following sense: for each $\omega \in \Omega$ and $\omega' \neq \omega$, there is $a \in A$ such that $\pi^\omega(a) \neq \pi^{\omega'}(a)$. Note that this rules out the case where the signal distribution is known and the states refer only to the player’s payoffs.

Intuitively, in this setup each player i is able to signal his private information ω_i whenever he wants, as no other player’s action can be confused with his own. Thus we might expect that the main obstacle to information revelation comes when player i ’s information will be used to lower his payoff. We verify these intuitions in the on-line supplementary material. Specifically we provide an example where m -statewise distinguishability fails, and prove the following lemma.

Lemma 17.

- (a) *Let (i, ω) and (j, ω') be such that $i \neq j$ and $\omega_i \neq \omega'_i$. Then there is a profile $\vec{\alpha}$ that statewise distinguishes (i, ω) from (j, ω') .*
- (b) *Let (i, ω) and (j, ω') be such that $\omega_{-ij} \neq \omega'_{-ij}$ (and possibly $i = j$). Then*

there is a profile $\vec{\alpha}$ that p -, m -, and n -statewise distinguishes (i, ω) from (j, ω') .

- (c) Let (i, ω) and (j, ω') be such that $\omega_j \neq \omega'_j$ (and possibly $i = j$). Then there is a profile $\vec{\alpha}$ that p -statewise distinguishes (i, ω) from (j, ω') .

6.3 A Two-Player, Two-Actions Partnership

Now we apply our results to Example 2 from Section 3 to illustrate the effect of players knowing their own productivity. There are two players, two actions $A_i = \{C_i, D_i\}$, two states, and three outcomes $Y = \{H, M, L\}$, and the state only influences the productivity of player 2's effort. We will show that the PTXE folk theorem holds when player 1's private information distinguishes ω_1 from ω_2 . On the other hand, when player 2's private information distinguishes ω_1 from ω_2 but player 1's does not, PTXE cannot approximate the efficient outcome for a range of parameters. Intuitively, player 2 cannot be induced to reveal the state when doing so would lower his equilibrium payoff, and this leads to a bound on the extent to which equilibria can trade off player 2's payoffs between the two states; in some cases this bound is so strong that it rules out the efficient outcome.

If player 1 chooses C_1 instead of D_1 , 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 , 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 .¹³ We assume that the vectors (p_H, p_M) and (q_H, q_M) are linearly independent; this implies that individual full rank and pairwise full rank are satisfied at every profile and every state. However, as [17] show, no type-independent profile p -statewise distinguishes $(1, \omega_1)$ and $(2, \omega_2)$, and as a result, the set of PPXE payoffs is bounded away from efficiency uniformly in the discount factor.

6.3.1 The Case where Player 1 Knows the State

Suppose that player 1 knows the state and player 2 does not, that is, $\Theta_1 = \{(\omega_1), (\omega_2)\}$ and $\Theta_2 = \{(\omega_1, \omega_2)\}$. The following table shows whether statewise distinguishability conditions are satisfied or not in this case; see the on-line supplementary

¹³Thus while the state space has a product structure the signals do not.

material for the formal proof.

$(i, \omega), (j, \omega')$	p -statewise	m -statewise
$(1, \omega_1), (1, \omega_2)$	(C_1, C_2)	(C_1, C_2)
$(1, \omega_2), (1, \omega_1)$	(C_1, C_2)	(C_1, C_2)
$(2, \omega_1), (2, \omega_2)$	$((C_1, D_1), C_2)$	(C_1, C_2)
$(2, \omega_2), (2, \omega_1)$	$((C_1, D_1), C_2)$	$((C_1, D_1), C_2)$
$(1, \omega_1), (2, \omega_2)$	$((C_1, D_1), C_2)$	(C_1, C_2)
$(2, \omega_2), (1, \omega_1)$	$((C_1, D_1), C_2)$	Not satisfied
$(2, \omega_1), (1, \omega_2)$	(C_1, C_2)	(C_1, C_2)
$(1, \omega_2), (2, \omega_1)$	(C_1, C_2)	$((C_1, D_1), C_2)$

Since m -statewise distinguishability does not hold for $((2, \omega_2), (1, \omega_1))$, the maximal scores for the corresponding directions are not infinitely large. Nevertheless, as shown in the on-line supplementary material, these scores are high enough to achieve the ex-post folk theorem for any $\beta \in (0, 1)$. Very roughly speaking, this is because the state-independent profile (D_1, C_2) yields sufficiently high payoffs in the corresponding directions (i.e., $\lambda \cdot g(D_1, C_2) > \max_{v \in V^*} \lambda \cdot v$ for directions λ such that $\lambda_2^{\omega_2} > 0$, $\lambda_1^{\omega_1} < 0$, $\lambda_2^{\omega_2} = 0$, and $\lambda_1^{\omega_2} \leq 0$) and hence players need not to learn the state to obtain high scores for these directions. This example shows that the statewise conditions are sufficient but not necessary for the folk theorem.

6.3.2 The Case where Player 2 Knows the State

Suppose next that player 2 knows the state and player 1 does not. The following table shows whether statewise conditions are satisfied or not. Again, see the on-line supplementary material for the formal proof.

$(i, \omega), (j, \omega')$	p -statewise	m -statewise
$(1, \omega_1), (1, \omega_2)$	(C_1, C_2)	(C_1, C_2)
$(1, \omega_2), (1, \omega_1)$	(C_1, C_2)	(C_1, C_2)
$(2, \omega_1), (2, \omega_2)$	$(C_1, (C_1, D_2))$	(C_1, C_2)
$(2, \omega_2), (2, \omega_1)$	$(C_1, (C_1, D_2))$	Not satisfied
$(1, \omega_1), (2, \omega_2)$	$(C_1, (C_1, D_2))$	(C_1, C_2)
$(2, \omega_2), (1, \omega_1)$	$(C_1, (C_1, D_2))$	$(C_1, (C_1, D_2))$
$(2, \omega_1), (1, \omega_2)$	(C_1, C_2)	(C_1, C_2)
$(1, \omega_2), (2, \omega_1)$	(C_1, C_2)	Not satisfied

In this case the on-line supplementary material shows that the folk theorem fails because the maximum score in direction $\lambda = ((0, -1), (0, 1))$ is too low. The reason is that state-contingent play is needed to generate high payoffs toward this direction and the failure of m -statewise distinguishability implies that learning the true state is somewhat costly. Moreover, if the cost of effort is high, then for $\lambda = ((0, -\varepsilon), (1, 0))$ the maximal score can be so low that it rules out equilibrium with the payoffs of the efficient action profile (C_1, C_2) . Specifically, this is the case if player 1's effort cost is high enough so that $g_1^{\omega_2}(D_1, C_2) - g_1^{\omega_2}(C_1, C_2)$ is close to zero.¹⁴ Intuitively, player 2 cannot be induced to reveal the state when doing so would lower his equilibrium payoff, and as a result the maximal score for direction λ with $\lambda_2^{\omega_1} < 0$ is lower than $\lambda \cdot g(C_1, C_2)$.

7 Known Monitoring Structure

So far we have studied a general model, where both payoffs and monitoring structure can depend on the state of the world, and provided sufficient conditions for the folk theorems. However, these sufficient conditions may not be satisfied in some games. One notable example is the case of a known monitoring structure; here a state-independent profile α cannot induce different signal distributions for different states, so for players to distinguish the states they must have “enough” private information. In this section we provide conditions for the limit equilibrium payoffs of games with a known monitoring structure to coincide with the limit equilibrium payoffs of the game with observed actions. This shows that the results of [21] and [22] are robust to imperfect monitoring, as PTXE reduces to the belief-free equilibria of those papers for perfect-monitoring games. Also we provide sufficient condition for folk theorems with a known monitoring structure.

Formally, the *monitoring structure is known* if $\pi_y^\omega(a) = \pi_y^{\omega'}(a)$ for all $y \in Y$, $a \in A$, $\omega \in \Omega$, and $\omega' \neq \omega$. We maintain this assumption throughout this section. Since π^ω does not depend on ω , we denote it by π .

In this section, we often impose the following strong full rank condition. As we will see, under this condition the case of a known but imperfect monitoring

¹⁴The derivation of this bound on the maximal score is very similar to the proof of Claims 12 and 13 which are used to prove Proposition 12; all of these proofs are in the supplementary materials.

structure is very similar to that where actions are perfectly observed. Let $\Pi_i(a)$ denote the matrix with rows $(\pi_y(a'_i, a_{-i}))_{y \in Y}$ for all $a'_i \in A_i$. Also, for each $i \in \mathbf{I}$, $j \in \mathbf{I}$, $a \in A$, and $a' \in A$, let $\Pi_{(i,a)(j,a')}$ denote the matrix constructed by stacking two matrices $\Pi_i(a)$ and $\Pi_j(a')$.

Definition 15. The monitoring structure π has *strong full rank* if

- (i) $\Pi_{(i,a)(j,a)}$ has rank $|A_i| + |A_j| - 1$ for all $i, j \in \mathbf{I}$ and $a \in A$; and
- (ii) for any $i, j \in \mathbf{I}$, if there is $l \neq i, j$, then $\Pi_{(i,a)(j,(a'_l, a_{-l}))}$ has rank $|A_i| + |A_j|$ for all $l \neq i, j$, $a \in A$, and $a'_l \neq a_l$.

Note that we allow $i = j$ in this definition, and hence the second clause is not vacuous even in two-player games. The first clause imposes FLM's pairwise full rank condition on every action profile. The second clause implies that the state-wise full rank condition holds for (i, ω) and (j, ω') if player l can distinguish the states ω and ω' .¹⁵ The strong full rank condition is obviously satisfied for games with perfectly observable actions. It is also satisfied if the signals are isomorphic to the actions and players observe the intended action with a small noise, i.e. $Y = A$ and $\pi_y(a) < \varepsilon$ for all $a \in A$ and $y \neq a$ where ε is close to zero.

7.1 Known Monitoring Structure and Strong Full Rank

In this subsection, we show that games with a known monitoring structure and strong full rank have the same set of limit equilibrium payoffs as games with observed actions. Specifically, we have the following proposition.

Proposition 6. *Suppose that the monitoring structure is known and has strong full rank. Suppose also that Q is full dimensional. Then the limit PTXE payoff set for this game, $\lim_{\delta \rightarrow 1} E(\delta) = Q$, is equal to the limit set of PTXE payoffs (or belief-free equilibrium payoffs) for the game that has the same information structure $(\Omega, (\Theta_i)_{i \in \mathbf{I}})$ and the same expected payoffs $(g_i)_{i \in \mathbf{I}}$ but with perfectly observable actions.*

¹⁵To see this, let $\tilde{\alpha}$ be such that $\alpha^{\theta(\omega)} = a$ and $\alpha^{\theta(\omega')} = (a'_l, a_{-l})$. Then this $\tilde{\alpha}$ has statewise full rank for (i, ω) and (j, ω') , as the corresponding matrix has rank $|A_i| + |A_j|$.

This shows that with a known monitoring structure and strong full rank, the analysis of the observed-action case carries over in the obvious way. When strong full rank fails, the known-monitoring-structure game can have a strictly smaller set of limit equilibrium payoffs than when actions are perfectly observable, for much the same reason that this can occur when the structure of the game is known.

To prove this proposition, we compute the maximal score for each direction λ and show that the score does not depend on the monitoring imperfection. For this, it is helpful to classify the directions so that the maximal score can be computed in the same way for all directions in a given class. As shown by FLM, when the state is known, there are three sorts of directions to consider: (i) maximizing the payoff of some player i ($\lambda_i > 0$ and $\lambda_j = 0$ for all $j \neq i$), (ii) minimizing the payoff of some player i ($\lambda_i < 0$ and $\lambda_j = 0$ for all $j \neq i$), or (iii) trading off the payoffs of two or more players ($\lambda_i \neq 0$ for at least two players i .) In our analysis here we combine all three sorts of directions together into the class Λ^1 of “single-state” directions; the maximal scores for such λ can be computed as in FLM. Λ^2 through Λ^4 are the sets of cross-state directions that satisfy the relevant version of statewise distinguishability, so the maximal scores in these directions are infinitely large. Λ^5 and Λ^6 are directions that (i) weight only the payoffs of a single player i but do so in more than one state and (ii) do not satisfy the relevant version of statewise distinguishability. Λ^7 is all of the directions that do not fit into classes 1 through 6. The maximal scores for Λ^5 , Λ^6 , and Λ^7 are not necessarily high enough for the folk theorem.

Let Λ^1 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 $(\lambda_i^{\omega'})_{i \in I} = 0$ for all $\omega' \neq \omega$. Since these directions consider only a single state, Lemmas 5.2 and 5.4 of FLM show that the maximum score is the maximum feasible score. As a result we obtain the following lemma.

Lemma 18. *Suppose the monitoring structure is known and has strong full rank. Then for each $\lambda \in \Lambda^1$, $k^*(\lambda) = \max_{v \in V^*} \lambda \cdot v$.*

Let Λ^2 be the set of λ such that there are $i \in I$, $j \in I$, $l \neq i, j$, $\omega \in \Omega$, and $\omega' \in \Omega$ such that $\lambda_i^\omega \neq 0$, $\lambda_j^{\omega'} \neq 0$, and $\theta_l(\omega) \neq \theta_l(\omega')$. Here player l can distinguish between ω and ω' , and the strong full rank condition implies that if player l tries to reveal this information by a state-contingent action, both player i at state ω and player j at state ω' are irrelevant to the information revelation. Thus player l 's

private information can be fully revealed, and as a result the maximal scores for these directions are infinity.

Lemma 19. *Suppose the monitoring structure is known and has strong full rank. Then for each $\lambda \in \Lambda^2$, $k^*(\lambda) = \infty$.*

Proof. Let $\lambda \in \Lambda^2$, and let $i \in \mathbf{I}$, $j \in \mathbf{I}$, $l \neq i, j$, $\omega \in \Omega$, and $\omega' \in \Omega$ be such that $\lambda_i^\omega \neq 0$, $\lambda_j^{\omega'} \neq 0$, and $\theta_l(\omega) \neq \theta_l(\omega')$. Since monitoring structure has strong full rank, both player i at ω and player j at ω' are irrelevant for player l 's information revelation. Then from Lemmas 11 there is $\vec{\alpha}$ that p , m , and n -statewise distinguishes (i, ω) from (j, ω') and from Lemma 4, we have $k^*(\lambda) = \infty$. *Q.E.D.*

Let Λ^3 be the set of λ such that there are $i \in \mathbf{I}$, $j \neq i$, $\omega \in \Omega$, and $\omega' \neq \omega$ such that $\lambda_i^\omega > 0$, $\lambda_j^{\omega'} \neq 0$, and $\theta_i(\omega) \neq \theta_i(\omega')$. Here player i can distinguish between ω and ω' , and the score is increasing in player i 's payoff in state ω . Since the strong full rank condition implies that player j at ω' is irrelevant to (i, ω) , player i 's private information can be fully revealed and the maximal scores for these directions are infinity as well.¹⁶

Lemma 20. *Suppose the monitoring structure is known and has strong full rank. Then for each $\lambda \in \Lambda^3$, $k^*(\lambda) = \infty$.*

Proof. Let $\lambda \in \Lambda^3$, and let $i \in \mathbf{I}$, $j \neq i$, $\omega \in \Omega$, and $\omega' \in \Omega$ be such that $\lambda_i^\omega > 0$, $\lambda_j^{\omega'} \neq 0$, and $\theta_i(\omega) \neq \theta_i(\omega')$. Since the monitoring structure has strong full rank, player j at ω' is irrelevant for (i, ω) . Then from Lemma 8, there is $\vec{\alpha}$ that statewise distinguishes (i, ω) from (j, ω') . Since $\lambda_i^\omega > 0$, Lemma 4 applies. *Q.E.D.*

Let Λ^4 be the set of λ such that there are $i \in \mathbf{I}$, $\omega' \in \Omega$, and $\omega'' \neq \omega'$ such that $\lambda_i^{\omega'} > 0$, $\lambda_i^{\omega''} > 0$, $(\lambda_j^\omega)_{\omega \in \Omega} = 0$ for all $j \neq i$, and $\theta_i(\omega') \neq \theta_i(\omega'')$. Here only player i 's payoffs matter, the score is increasing in i 's payoff in ω and ω' , and player i can distinguish between these two states. Once again, the maximal scores for these directions are infinity, as strong full rank implies that player i can reveal whether ω or ω' .

Lemma 21. *Suppose the monitoring structure is known and has strong full rank. Then for each $\lambda \in \Lambda^4$, $k^*(\lambda) = \infty$.*

¹⁶The intersection of Λ^2 and Λ^3 might be non-empty but this is irrelevant as the maximal score is infinity for either case.

Proof. Let $\lambda \in \Lambda^4$, and let $i \in \mathbf{I}$, $\omega' \in \Omega$, and $\omega'' \in \Omega$ be such that $\lambda_i^{\omega'} > 0$, $\lambda_i^{\omega''} > 0$, and $\theta_i(\omega') \neq \theta_i(\omega'')$. Since the monitoring structure has strong full rank, player i can reveal whether the state is ω or ω' . Then from Lemma 5, there is $\vec{\alpha}$ that p -statewise distinguishes (i, ω') from (i, ω'') . Since $\lambda_i^{\omega'} > 0$ and $\lambda_i^{\omega''} > 0$, Lemma 4(b) applies. *Q.E.D.*

Let $\Lambda^5(i)$ be the set of λ such that $(\lambda_i^\omega)_{\omega \in \Omega} \leq 0$, $(\lambda_i^\omega)_{\omega \in \Omega} \neq 0$, $(\lambda_j^\omega)_{\omega \in \Omega} = 0$ for all $j \neq i$, and $\theta_j(\omega) = \theta_j(\omega')$ for all $j \neq i$, $\omega \in \Omega$, and $\omega' \neq \omega$ satisfying $\lambda_i^\omega \neq 0$ and $\lambda_i^{\omega'} \neq 0$. Here only player i 's payoffs matter, the score is decreasing in i 's payoff, and no other player can distinguish between the states; these directions determine the minmax payoff for player i , taking into account a trade-off between the minmax level in one state and the payoffs in other states. Let $\Lambda^5 = \bigcup_{i \in \mathbf{I}} \Lambda^5(i)$.

Lemma 22. *Suppose the monitoring structure is known and has strong full rank. Then for each i and $\lambda \in \Lambda^5(i)$, $k^*(\lambda) = \max_{\alpha_{-i}} \min_{a_i} \sum_{\omega \in \Omega} \lambda_i^\omega(\omega) g_i^\omega(a_i, \alpha_{-i})$, that is, $k^*(\lambda) = -\min_{\alpha_{-i}} \max_{a_i} \sum_{\omega \in \Omega} -\lambda_i^\omega g_i^\omega(a_i, \alpha_{-i})$.*

The proof is delegated to the appendix. The intuition is as follows: Strong full rank implies that constraints (i) and (ii) can be satisfied for all $j \neq i$, and because $(\lambda_j^\omega)_{\omega \in \Omega} = 0$ the continuation payoffs assigned to $j \neq i$ are irrelevant. Thus we only need to consider continuation payoffs for player i that satisfy (i) and (ii) for ω such that $\lambda_i^\omega \neq 0$, and the feasibility constraint (iii). Note also that player $j \neq i$ has to use the same action α_j for all states ω with $\lambda_i^\omega \neq 0$, as he cannot distinguish these states by definition of $\Lambda^5(i)$. Summing the incentive-compatibility constraints over the states ω (taking into account that $\lambda_i^\omega \leq 0$) yields a “weaker aggregate incentive condition,” which corresponds to a game with a known state where player i 's payoff is $\sum_{\omega \in \Omega} -\lambda_i^\omega g_i^\omega(a)$. Using this analogy, we can show that the maximal score in the direction of minimizing this payoff (that is, maximizing $-\sum_{\omega \in \Omega} -\lambda_i^\omega v_i^\omega$) is at most the corresponding minmax payoff, namely $-\min_{\alpha_{-i}} \max_{a_i} \sum_{\omega \in \Omega} -\lambda_i^\omega g_i^\omega(a_i, \alpha_{-i})$. We then use the strong full rank assumption to show that this bound is attained.

Let Λ^6 be the set of λ such that there is $i \in \mathbf{I}$ such that $\lambda_i^\omega > 0$ for some $\omega \in \Omega$, $(\lambda_j^\omega)_{\omega \in \Omega} = 0$ for all $j \neq i$, $\theta_i(\omega') = \theta_i(\omega'')$ for all $\omega' \in \Omega$ and $\omega'' \neq \omega'$ satisfying $\lambda_i^{\omega'} > 0$ and $\lambda_i^{\omega''} > 0$, and $\theta_j(\omega') = \theta_j(\omega'')$ for all $j \neq i$, $\omega' \in \Omega$, and $\omega'' \neq \omega'$ satisfying $\lambda_i^{\omega'} \neq 0$ and $\lambda_i^{\omega''} \neq 0$. In words, this says that only player i 's payoff has non-zero weights, that player i cannot distinguish between any two states where

his utility gets positive weight, and no other player can distinguish between any two states where player i 's utility gets non-zero weight.

Finally, we construct a set Λ^7 that we show contains all directions that do not belong to one of the preceding sets. We define Λ^7 to be the set of all λ satisfying the following properties.

- (i) $(\lambda_i^\omega)_{\omega \in \Omega} \neq 0$ and $(\lambda_j^\omega)_{\omega \in \Omega} \neq 0$ for some $i \in \mathbf{I}$ and $j \neq i$.
- (ii) $(\lambda_l^{\omega'})_{l \in \mathbf{I}} \neq 0$ and $(\lambda_l^{\omega''})_{l \in \mathbf{I}} \neq 0$ for some $\omega' \in \Omega$ and $\omega'' \neq \omega'$.
- (iii) $\theta_l(\omega''') = \theta_l(\omega''')$ for $l \in \mathbf{I}$, $\omega''' \in \Omega$, and $\omega''' \neq \omega''$, if $\lambda_{l'}^{\omega'''} \neq 0$ for some $l' \neq l$ and $\lambda_{l''}^{\omega'''} \neq 0$ for some $l'' \neq l$.
- (iv) $\theta_l(\omega''') = \theta_l(\omega''')$ for $l \in \mathbf{I}$, $\omega''' \in \Omega$, and $\omega''' \neq \omega''$ if $\lambda_{l'}^{\omega'''} > 0$ and $\lambda_{l''}^{\omega'''} \neq 0$ for some $l' \neq l$.

In words, this is the set of directions where the score depends on the payoffs of players i and j in some state ω , and where it also depends on the payoff of some player l (possibly i or j) in two other states ω' and ω'' , but this player l cannot distinguish between any states ω''' and ω'' if either (condition (iii)) in each of these states there is at least one other player whose payoff matters or (condition (iv)) the score is increasing in l 's payoff in state ω''' and depends on the payoff of some l' in state ω''' .

Lemma 23. $\bigcup_{n=1}^7 \Lambda^n = \mathbf{R}^{I \times |\Omega|} \setminus \{(0, \dots, 0)\}$

Proof. Let λ be such that $\lambda \neq (0, \dots, 0)$ and $\lambda \notin \Lambda^7$. It suffices to show that $\lambda \in \bigcup_{n=1}^6 \Lambda^n$. If λ does not satisfy the clause (ii) of the definition of Λ^7 , then $\lambda \in \Lambda^1$. If λ does not satisfy (iii), then $\lambda \in \Lambda^2$. If λ does not satisfy (iv), then $\lambda \in \Lambda^3$. If λ satisfies (iii) and (iv) but not (i), then $\lambda \in \Lambda^4 \cup \Lambda^5 \cup \Lambda^6$. *Q.E.D.*

Lemma 24. Suppose the monitoring structure is known and has strong full rank. Then for each $\lambda \in \Lambda^6 \cup \Lambda^7$, $k^*(\lambda) = \max_{\alpha} \lambda \cdot g(\alpha)$.

The proof is given in the appendix. The first step of the proof is to show that for each $\lambda \in \Lambda^6 \cup \Lambda^7$, there is a single “type” θ_i^* that is relevant; we use this to show that the upper and lower bounds on the score are both $\max_{\alpha} \lambda \cdot g(\alpha)$.

Combining the above lemmas yields the following characterization of the maximal scores in each direction and thus of the set Q .

Proposition 7. *Suppose the monitoring structure is known and has strong full rank. Then*

$$k^*(\lambda) = \begin{cases} \max_{v \in V^*} \lambda \cdot v & \text{if } \lambda \in \Lambda^1. \\ \infty & \text{if } \lambda \in \Lambda^2 \cup \Lambda^3 \cup \Lambda^4 \\ \max_{\alpha_{-i}} \min_{a_i} \sum_{\omega \in \Omega} \lambda_i^\omega g_i^\omega(a_i, \alpha_{-i}) & \text{if } \lambda \in \Lambda^5(i) \\ \max_{\alpha} \lambda \cdot g(\alpha) & \text{if } \lambda \in \Lambda^6 \cup \Lambda^7 \end{cases},$$

and $Q = \bigcap_{i \in \{1, \dots, 7\}, \lambda \in \Lambda^i} H^*(\lambda)$.

This proposition shows that the monitoring imperfection does not affect the maximal score. Then the set Q does not depend on the monitoring imperfection as well, and hence Proposition 6 follows.

7.2 One-Sided Incomplete Information

In this subsection we consider the case where only player 1's payoff function is uncertain, and he knows his own payoff function while the other players do not. Formally, we say the game has one-sided incomplete information if $g_i^\omega(a) = g_i^{\omega'}(a)$ for all $i \neq 1$, $a \in A$, $\omega \in \Omega$, and $\omega' \neq \omega$, and that $\theta_1(\omega) = (\omega)$ for all ω , and $\Theta_i = \{(\Omega)\}$ for all $i \neq 1$. This is the assumption made in Hörner and Lovo (2009, Section 4) and Hörner, Lovo and Tomala (2009, Section 6) analysis of reputations, so once again our results can be seen as extending theirs.

Section 5 of [22] derives several sufficient conditions for Q (denoted by V^* in their paper) to be non-empty,¹⁷ which implies that there is a PTXE in the undiscounted case. However these conditions do not assure the existence of PTXE with the discounted payoff criterion used in this paper, because Q might not be full dimensional and in that case their existence result and our Proposition 1 would not apply. In this subsection, we give a simple sufficient condition for Q to be full dimensional; under this condition, Q equals the set of limit PTXE payoffs so that there are PTXE for sufficiently large δ .

Let V^U be the set of feasible payoffs of the stage game with public randomization, that is, $V^U = \text{co}\{g(a) | a \in A\}$. Note that $\dim V^U$ is at most $|\Omega| + I - 1$, since

¹⁷They also give tight conditions for Q to be non-empty by imposing restrictions on the payoff functions as well as on the information structure.

$g_i^\omega(a) = g_i^{\omega'}(a)$ for all $i \neq 1$, $a \in A$, $\omega \in \Omega$, and $\omega' \neq \omega$. Let

$$V^{U*} = \{v \in V^U \mid \forall i \exists \alpha_{-i} \forall \omega, v_i^\omega \geq \max_{a_i} g_i^\omega(a_i, \alpha_{-i})\}.$$

Condition Non-E. The set V^{U*} has dimension $|\Omega| + I - 1$.

This condition is likely to be satisfied if there is an action α_{-1} that gives low payoffs to player 1 for every state ω . The next proposition shows that (Non-E) is sufficient for the set Q to be full dimensional; then Proposition 1 applies, and so (Non-E) is a sufficient condition for the existence of PTXE.

Proposition 8. *Suppose that the monitoring structure is known and has strong full rank, and that there is one-sided incomplete information. Suppose also that (Non-E) holds. Then $\dim Q = I \times |\Omega|$.*

Proof. Let v be in the relative interior of V^{U*} . It suffices to show that $k^*(\lambda) > \lambda \cdot v$ for all λ .

First, consider $\lambda \in \Lambda^1$. Since $V^{U*} \subseteq V^*$, v is an interior point of V^* . Then $\lambda \cdot v < \max_{v' \in V^*} \lambda \cdot v' = k^*(\lambda)$ for $\lambda \in \Lambda^1$. Likewise, since $V^{U*} \subseteq V^U$, v is an interior point of V^U . Then, $\lambda \cdot v < \max_{v' \in V^U} \lambda \cdot v' = \max_{\alpha} \lambda \cdot g(\alpha) = k^*(\lambda)$ for $\lambda \in \Lambda^6$ and $\lambda \in \Lambda^7$.

Since $k^*(\lambda) = \infty$ for $\lambda \in \Lambda^2 \cup \Lambda^3 \cup \Lambda^4$, it remains to consider $\lambda \in \Lambda^5$. By the definition of Λ^5 , $(\lambda_1^\omega)_{\omega \in \Omega} \neq 0$, $\lambda_1^\omega \leq 0$ for all $\omega \in \Omega$, and $(\lambda_j^\omega)_{\omega \in \Omega} = 0$ for all $j \neq 1$. Also, since v is in the relative interior of V^{U*} , there is α_{-1} such that $v_1^\omega > \max_{a_1} g_1^\omega(a_1, \alpha_{-1})$ for all $\omega \in \Omega$. Taken together, we obtain

$$\begin{aligned} \lambda \cdot v &= \sum_{\omega \in \Omega} \lambda_1^\omega v_1^\omega < \sum_{\omega \in \Omega} \lambda_1^\omega \max_{a_1} g_1^\omega(a_1, \alpha_{-1}) = \sum_{\omega \in \Omega} \min_{a_1} \lambda_1^\omega g_1^\omega(a_1, \alpha_{-1}) \\ &< \max_{\alpha'_{-1}} \min_{a_1} \sum_{\omega \in \Omega} \lambda_1^\omega g_1^\omega(a_1, \alpha'_{-1}) = k^*(\lambda), \end{aligned}$$

as desired. Here, the equality in the third line comes from

$$\begin{aligned} \lambda_1^\omega \max_{a_1} g_1^\omega(a_1, \alpha_{-1}) &= -\lambda_1^\omega \min_{a_1} (-g_1^\omega(a_1, \alpha_{-1})) = |\lambda_1^\omega| \min_{a_1} (-g_1^\omega(a_1, \alpha_{-1})) \\ &= \min_{a_1} |\lambda_1^\omega| (-g_1^\omega(a_1, \alpha_{-1})) = \min_{a_1} \lambda_1^\omega g_1^\omega(a_1, \alpha_{-1}). \end{aligned}$$

Q.E.D.

Remark 3. If there is a “commitment type” ω^* , for which there is some $a_1^* \in A_1$ such that $g_1^{\omega^*}(a_1^*, a_{-1})$ is independent of a_{-1} and $g_1^{\omega^*}(a_1^*, a_{-1}) \geq g_1^{\omega^*}(a)$ for all $a_1 \in A_1$, the minimax payoff of this commitment type equals his best payoff $g_1^{\omega^*}(a_1^*, a_{-i})$. In this case the set Q does not have full dimension, and our results do not apply.¹⁸ Moreover, in this case the set of PTXE is often empty. Suppose that there are two players, and player 2 has a unique best reply against a_1^* , and call it a_2^* . In a PTXE, player 1 in state ω^* always play a_1^* , so that player 2 must play a_2^* after every history, independently of the state. Then player 1’s optimal strategy for state $\omega \neq \omega^*$ is to choose $a_1^\omega \in \arg \max_{a_1 \in A_1} g_1^\omega(a_1, a_2^*)$ after every history. For this strategy profile to be a PTXE, a_2^* must be a best reply to a_1^ω for all $\omega \neq \omega^*$, but such a condition is not satisfied in general. Thus we conclude that there is no PTXE for any discount factor.¹⁹

7.3 The Folk Theorem with Known Monitoring Structure

Our general folk theorem uses (SFR) or (Pointwise-SD), which require either that all players can distinguish every pair of states, or that there are profiles $\vec{\alpha}$ that satisfy various full rank conditions. With a known monitoring structure (and strong full rank) the following simpler condition is sufficient.

Proposition 9. *Suppose that the monitoring structure is known and has strong full rank. Suppose also that for each (ω, ω') satisfying $\omega \neq \omega'$, there are at least three players who can distinguish ω and ω' , i.e., there are $i \in I$, $j \neq i$, and $l \neq i, j$ such that $\theta_i(\omega) \neq \theta_i(\omega')$, $\theta_j(\omega) \neq \theta_j(\omega')$, and $\theta_l(\omega) \neq \theta_l(\omega')$. Then, for any smooth strict subset W of V^* , there exists $\bar{\delta} \in (0, 1)$ such that $W \subseteq E(\delta)$ for all $\delta \in (\bar{\delta}, 1)$.*

Proof. Since there are at least three players who can distinguish ω and ω' , any cross-state direction λ is an element of Λ^2 . Then, from Proposition 7, we have $k^*(\lambda) = \infty$. Since $k^*(\lambda) = \max_{v \in V^*} \lambda \cdot v$ for any $\lambda \in \Lambda^1$, we obtain $Q = V^*$.
Q.E.D.

Theorem 5.3 of [22] shows that Q is non-empty for games with perfect monitoring, if there are at least three players who can distinguish ω and ω'

¹⁸[21] make essentially this point on page 475.

¹⁹If there are observed actions, these same assumptions imply that there is not a belief-free equilibrium. [21] note that there is a belief-free equilibrium with a commitment type in strictly dominant action games with a unique Stackelberg type.

for each (ω, ω') satisfying $\omega \neq \omega'$; our result shows that the assumptions of that proposition are in fact sufficient for a folk theorem.

In the next proposition, we consider the case in which there are at least two players who can distinguish states. Recall that V^U is the set of feasible payoffs of the stage game with public randomization, that is, $V^U = \text{co}\{g(a) | a \in A\}$.

Proposition 10. *Suppose that the monitoring structure is known and satisfies strong full rank. Suppose also that for each (ω, ω') satisfying $\omega \neq \omega'$, there are at least two players who can distinguish ω and ω' , i.e., there are $i \in \mathbf{I}$ and $j \neq i$ such that $\theta_i(\omega) \neq \theta_i(\omega')$ and $\theta_j(\omega) \neq \theta_j(\omega')$. Let $V^{**} \equiv \{v \in V^* | \exists \tilde{v} \in V^U \forall i \in \mathbf{I} \forall \omega \in \Omega, v_i^\omega \geq \tilde{v}_i^\omega\}$. Then, for any smooth strict subset W of V^{**} , there is $\bar{\delta} \in (0, 1)$ such that $W \subseteq E(\delta)$ for all $\delta \in (\bar{\delta}, 1)$.*

Note that if there is a “bad outcome” $\alpha \in \Delta A$ such that $g_i^\omega(\alpha) \leq v_i^\omega$ for all $i \in \mathbf{I}$ and $\omega \in \Omega$, then we have $V^{**} = V^*$, so that the folk theorem obtains. Theorem 5.11 of [22] shows that Q is non-empty for games with perfect monitoring and a bad outcome, if there are at least two players who can distinguish ω and ω' for each (ω, ω') satisfying $\omega \neq \omega'$. Again our result shows that the assumptions of the proposition are sufficient for a folk theorem.

Proof. It suffices to show that $V^{**} \subseteq Q$. To do so, we compute the maximal score $k^*(\lambda)$ for every direction, using Proposition 7.

First, consider $\lambda \in \Lambda^1$. It follows from Proposition 7 that $k^*(\lambda) = \max_{v \in V^*} \lambda \cdot v$ for this direction. Next, consider λ such that $\lambda_i^{\omega'} \neq 0$ and $\lambda_i^{\omega''} \neq 0$ for some $i \in \mathbf{I}$, $\omega' \in \Omega$, and $\omega'' \neq \omega'$, and $(\lambda_j^\omega)_{\omega \in \Omega} = 0$ for all $j \neq i$. Since there are at least two players who can distinguish ω' and ω'' , there is $l \neq i$ such that $\theta_l(\omega') \neq \theta_l(\omega'')$. Thus $\lambda \in \Lambda^2$, and hence $k^*(\lambda) = \infty$ for this direction.

Consider λ such that $\lambda_i^{\omega'} \neq 0$ and $\lambda_j^{\omega''} \neq 0$ for some $i \in \mathbf{I}$, $j \neq i$, $\omega' \in \Omega$, and $\omega'' \neq \omega'$, and $\theta_l(\omega') \neq \theta_l(\omega'')$ for some $l \neq i, j$. Again, $\lambda \in \Lambda^2$ in this case, so that $k^*(\lambda) = \infty$. Consider λ such that $\lambda_i^{\omega'} > 0$ and $\lambda_j^{\omega''} \neq 0$ for some $i \in \mathbf{I}$, $j \neq i$, $\omega' \in \Omega$, and $\omega'' \neq \omega'$, and $\theta_l(\omega') = \theta_l(\omega'')$ for all $l \neq i, j$. Since there are at least two players who can distinguish ω' and ω'' , it must be that $\theta_i(\omega') \neq \theta_i(\omega'')$ and $\theta_j(\omega') \neq \theta_j(\omega'')$. This implies that $\lambda \in \Lambda^3$, and hence $k^*(\lambda) = \infty$.

Finally, consider λ such that $\lambda \leq 0$, $\lambda_i^{\omega'} < 0$ and $\lambda_j^{\omega''} < 0$ for some $i \in \mathbf{I}$, $j \neq i$, $\omega' \in \Omega$, and $\omega'' \neq \omega'$, and for any pair (i, ω''') and (j, ω''') satisfying $\omega''' \neq \omega''$,

$\lambda_i^{\omega'''} < 0$, and $\lambda_j^{\omega'''} < 0$, and for any $l \neq i, j$, $\theta_l(\omega''') = \theta_l(\omega''''')$. By definition, $\lambda \in \Lambda^7$ in this case, so that $k^*(\lambda) = \max_{\alpha} \lambda \cdot g(\alpha) = \max_{v \in V^U} \lambda \cdot v$.

From the above arguments, obviously we have $V^{**} \subset H^*(\lambda)$ for all λ . Therefore, $V^{**} \subseteq Q$. *Q.E.D.*

8 Conclusion

This paper shows how to extend the insights and techniques of the repeated games literature to games with imperfectly observed actions, an unknown monitoring structure, and private information. Our analysis is based on the fact that the set of PTXE payoffs has a recursive structure, and says little about the entire set of equilibrium payoffs. When the folk theorem holds in PTXE, or more generally when there are asymptotically efficient PTXE, the restriction to PTXE may be of less concern, especially given their desirable robustness properties. When the set of PTXE is small or empty, it would be nice to know more about the entire set of sequential equilibrium payoffs; that more difficult problem is still unresolved. Another open question is to extend the analysis of PTXE to other settings where repeated play has been shown to support more efficient outcomes, such as games with long-run and short-run players (Fudenberg, Kreps, and Maskin [13] and [14]), games with overlapping generations of players (Kandori [24]), community enforcement (Kandori [23] and Ellison [11]), games with imperfect private monitoring (Compte [7] and Kandori and Matsushima [26]), and games where the state evolves according to a finite Markov chain (Athey and Bagwell [4]).

Appendix

A.1 Proof of Lemma 4

Lemma 4.

- (a) Suppose $\vec{\alpha}$ is ex-post enforceable and m -statewise distinguishes (i, ω) from (j, ω') . Then $k^*(\alpha, \lambda) = \infty$ for λ such that $\lambda_i^{\omega} > 0$ and $\lambda_j^{\omega'} < 0$.
- (b) Suppose $\vec{\alpha}$ is ex-post enforceable and p -statewise distinguishes (i, ω) from (j, ω') . Then $k^*(\alpha, \lambda) = \infty$ for λ such that $\lambda_i^{\omega} > 0$ and $\lambda_j^{\omega'} > 0$.

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

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

$$z_i^\omega(y) = \frac{K}{\delta \lambda_i^\omega \pi^\omega(\alpha) \cdot \xi} \xi(y) \quad \text{and} \quad z_j^{\omega'}(y) = -\frac{K}{\delta \lambda_j^{\omega'} \pi^\omega(\alpha) \cdot \xi} \xi(y)$$

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

$$\pi^\omega(a_i, \alpha_{-i}) \cdot z_i^\omega = \frac{K}{\delta \lambda_i^\omega \pi^\omega(\alpha) \cdot \xi} \pi^\omega(a_i, \alpha_{-i}) \cdot \xi = \frac{K}{\delta \lambda_i^\omega} \quad (3)$$

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

$$\pi^\omega(a_i, \alpha_{-i}) \cdot z_i^\omega = \frac{K}{\delta \lambda_i^\omega \pi^\omega(\alpha) \cdot \xi} \pi^\omega(a_i, \alpha_{-i}) \cdot \xi \leq \frac{K}{\delta \lambda_i^\omega} \quad (4)$$

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

$$\pi^{\omega'}(a_j, \alpha_{-j}) \cdot z_j^{\omega'} = -\frac{K}{\delta \lambda_j^{\omega'} \pi^\omega(\alpha) \cdot \xi} \pi^{\omega'}(a_j, \alpha_{-j}) \cdot \xi(y) = 0 \quad (5)$$

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

$$\pi^{\omega'}(a_j, \alpha_{-j}) \cdot z_j^{\omega'} = -\frac{K}{\delta \lambda_j^{\omega'} \pi^\omega(\alpha) \cdot \xi} \pi^{\omega'}(a_j, \alpha_{-j}) \cdot \xi \leq 0 \quad (6)$$

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

$$\lambda_i^\omega z_i^\omega(y) + \lambda_j^{\omega'} z_j^{\omega'}(y) = 0 \quad (7)$$

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 $(\tilde{v}, \alpha)$. Let $K > \max_{y \in Y} \lambda \cdot \tilde{w}(y) - \lambda \cdot \tilde{v}$. Then, let

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

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

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

We claim that this (v, w) satisfies all the constraints in the LP problem. Obviously, constraints (i) and (ii) are satisfied for all $(l, \bar{\omega}) \in (\mathbf{I} \times \Omega) \setminus \{(i, \omega), (j, \omega')\}$, as $v_l^{\omega''} = \tilde{v}_i^{\omega''}$ and $w_l^{\bar{\omega}}(y) = \tilde{w}_l^{\omega''}(y)$. Also, since (3) and (4) hold and $\tilde{w}$ enforces $(\alpha, \tilde{v})$, 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 (5) and (6) that (v, w) satisfies constraints (i) and (ii) for (j, ω') . Furthermore, using (7) and $K > \max_{y \in Y} \lambda \cdot \tilde{w}(y) - \lambda \cdot \tilde{v}$,

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

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 $k^*(\alpha, \lambda) = \infty$, which proves part (a) of the lemma.

For parts (b) and (c), let $\xi = (\xi(y))_{y \in Y}$ be as in the definition of p - or n -statewise distinguishability. Without loss of generality, assume $\pi^{\omega'}(\alpha) \cdot \xi = 0$. The rest of the proof is the same as in part (a). Q.E.D.

A.2 Proof of Lemma 12

Lemma 12. *Suppose (PFR) and (IFR) hold. Let λ be such that $\theta_i(\omega) \neq \theta_i(\omega')$ for all $i \in \mathbf{I}$, $\omega \in \Omega$ and $\omega' \neq \omega$ satisfying $(\lambda_j^{\omega})_{j \in \mathbf{I}} \neq 0$ and $(\lambda_j^{\omega'})_{j \in \mathbf{I}} \neq 0$. Then, $k^*(\lambda) \geq \max_{v \in V^*} \lambda \cdot v$.*

Proof. For each $\omega \in \Omega$, let $\lambda(\omega) = (\lambda_i^{\omega'}(\omega))_{(i,\omega')}$ be such that $(\lambda_i^{\omega}(\omega))_{i \in I} = (\lambda_i^{\omega})_{i \in I}$ and $(\lambda_i^{\omega'}(\omega))_{i \in I} = 0$ for all $\omega' \neq \omega$. Let Ω^* be the set of all ω such that $\lambda(\omega) \neq 0$. We claim

$$k^*(\vec{\alpha}, \lambda) \geq \sum_{\omega \in \Omega^*} k^*(\vec{\alpha}, \lambda(\omega)) \quad (8)$$

for each $\vec{\alpha}$. In words, $k^*(\vec{\alpha}, \lambda)$ is at least the sum of the maximal scores when we solve the LP problem for each state ω in isolation. To prove this, consider the LP problem for $(\vec{\alpha}, \lambda)$ but constraint (iii) is replaced with a more restrictive condition

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

Let $k^U(\vec{\alpha}, \lambda)$ denote the solution to this new problem. Since condition (iii') does not allow utility transfer across different states, considering this new LP problem is equivalent to solving a separate LP problem for each state $\omega \in \Omega^*$ in isolation. Thus we have $k^U(\vec{\alpha}, \lambda) = \sum_{\omega \in \Omega^*} k^*(\vec{\alpha}, \lambda(\omega))$. Since $k^*(\vec{\alpha}, \lambda) \geq k^U(\vec{\alpha}, \lambda)$, (8) follows.

Recall that $\lambda(\omega)$ considers only a single state ω . Thus the maximal score $k^*(\vec{\alpha}, \lambda(\omega))$ depends on $\alpha^{\theta(\omega)}$ but not on $\alpha^{\theta'}$ for other θ' . This observation, together with the fact that all players can distinguish any state in the set Ω^* , implies that

$$\sup_{\vec{\alpha}} \sum_{\omega \in \Omega^*} k^*(\vec{\alpha}, \lambda(\omega)) = \sum_{\omega \in \Omega^*} \sup_{\vec{\alpha}} k^*(\vec{\alpha}, \lambda(\omega)).$$

It follows from Lemmas 13 and 14 that $\sup_{\vec{\alpha}} k^*(\vec{\alpha}, \lambda(\omega)) = \max_{v \in V^*} \lambda(\omega) \cdot v$. Therefore,

$$\sup_{\vec{\alpha}} \sum_{\omega \in \Omega^*} k^*(\vec{\alpha}, \lambda(\omega)) = \sum_{\omega \in \Omega^*} \max_{v \in V^*} \lambda(\omega) \cdot v = \max_{v \in V^*} \lambda \cdot v.$$

Using (8), we obtain the desired result. *Q.E.D.*

A.3 Proof of Lemma 16

Lemma 16. *Suppose (PFR) holds. Let λ be such that $\theta_i(\omega) \neq \theta_i(\omega')$ for all $i \in I$, $\omega \in \Omega$ and $\omega' \neq \omega$ satisfying $(\lambda_j^{\omega})_{j \in I} \neq 0$ and $(\lambda_j^{\omega'})_{j \in I} \neq 0$. Then, $k^*(\lambda) \geq \max_{v \in V^0} \lambda \cdot v$.*

Proof. The proof is very similar to Lemma 12. The only difference is that in the last step of the proof, we may not have $\sup_{\vec{\alpha}} k^*(\vec{\alpha}, \lambda(\omega)) = \max_{v \in V^*} \lambda(\omega) \cdot v$, since (IFR) might fail. Instead, we use Lemmas 13 and 15 to show that $\sup_{\vec{\alpha}} k^*(\vec{\alpha}, \lambda(\omega)) \geq \max_{v \in V^0} \lambda(\omega) \cdot v$. Q.E.D.

A.4 Proof of Lemma 22

Lemma 22. *Suppose the monitoring structure is known and has strong full rank. Then for each i and $\lambda \in \Lambda^5(i)$, $k^*(\lambda) = \max_{\alpha_{-i}} \min_{a_i} \sum_{\omega \in \Omega} \lambda_i^\omega(\omega) g_i^\omega(a_i, \alpha_{-i})$, that is, $k^*(\lambda) = -\min_{\alpha_{-i}} \max_{a_i} \sum_{\omega \in \Omega} -\lambda_i^\omega g_i^\omega(a_i, \alpha_{-i})$.*

To prove this lemma, we use the following claims.

Claim 1. *Let $\lambda \in \Lambda^5(i)$. Then for each $j \neq i$, there is $\theta_j^* \in \Theta_j$ that contains all ω such that $\lambda_i^\omega \neq 0$.*

Proof. Suppose not, so that there are $\omega \in \Omega$ and $\omega' \neq \omega$ such that $\theta_j(\omega) \neq \theta_j(\omega')$, $\lambda_i^\omega \neq 0$, and $\lambda_i^{\omega'} \neq 0$. Then $\lambda \notin \Lambda^5(i)$, since it does not satisfy the last condition of the definition of $\Lambda^5(i)$. A contradiction. Q.E.D.

Claim 2. *Suppose the monitoring structure is known. Let $\lambda \in \Lambda^5(i)$. Then for each $\vec{\alpha} = ((\alpha_i^{\theta_i})_{\theta_i \in \Theta_i})_{i \in I}$, $k^*(\vec{\alpha}, \lambda) \leq \min_{a_i} \lambda \cdot g(a_i, \alpha_{-i}^{\theta_{-i}^*})$ where θ_{-i}^* is chosen as in Claim 1 and $\alpha_{-i}^{\theta_{-i}^*} = (\alpha_j^{\theta_j^*})_{j \neq i}$.*

Proof. Let $a'_i \in \arg \min_{a_i} \lambda \cdot g(a_i, \alpha_{-i}^{\theta_{-i}^*})$. If $k^*(\vec{\alpha}, \lambda) = -\infty$ then the result is obvious. If $k^*(\vec{\alpha}, \lambda) > -\infty$, we can choose (v, w) to satisfy constraints (i) through (iii) in the LP problem associated with $(\vec{\alpha}, \lambda, \delta)$ for some $\delta \in (0, 1)$. It follows from constraint (ii) that

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

for all i , a_i , and ω such that $\lambda_i^\omega \neq 0$, since $\theta_j(\omega) = \theta_j^*$ for $j \neq i$ for such ω . Multiplying both sides by λ_i^ω , summing over all ω , and using the fact that $\lambda_j^\omega = 0$ for all $j \neq i$, we have

$$\begin{aligned} \lambda \cdot v &= \sum_{\omega \in \Omega} \lambda_i^\omega v_i^\omega \leq (1 - \delta) \sum_{\omega \in \Omega} \lambda_i^\omega g_i^\omega(a'_i, \alpha_{-i}^{\theta_{-i}^*}) + \delta \sum_{\omega \in \Omega} \sum_{y \in Y} \pi_y(a'_i, \alpha_{-i}^{\theta_{-i}^*}) \lambda_i^\omega w_i^\omega(y) \\ &= (1 - \delta) \lambda \cdot g(a'_i, \alpha_{-i}^{\theta_{-i}^*}) + \delta \sum_{y \in Y} \pi_y(a'_i, \alpha_{-i}^{\theta_{-i}^*}) \lambda \cdot w(y), \end{aligned}$$

so from (iii),

$$\lambda \cdot v \leq (1 - \delta)\lambda \cdot g(a'_i, \alpha_{-i}^{\theta^*}) + \delta\lambda \cdot v.$$

Subtracting $\delta\lambda \cdot v$ from both sides and dividing by $(1 - \delta)$, we get $\lambda \cdot v \leq \lambda \cdot g(a'_i, \alpha_{-i}^{\theta^*})$. Therefore, $k^*(\vec{\alpha}, \lambda, \delta) \leq g(a'_i, \alpha_{-i}^{\theta^*})$. *Q.E.D.*

For each $j \in \mathbf{I}$, let $\tilde{g}_j(a) = -\sum_{\omega \in \Omega} \lambda_j^\omega g_j^\omega(a)$. Let $\tilde{\lambda} \in \mathbf{R}^I$ be such that $\tilde{\lambda}_i = -1$ and $\tilde{\lambda}_j = 0$ for all $j \neq i$. Consider the following LP problem:

$$\begin{aligned} \tilde{k}^*(\alpha, \tilde{\lambda}, \delta) = \max_{\substack{\tilde{v} \in \mathbf{R}^I \\ \tilde{w}: Y \rightarrow \mathbf{R}^I}} \tilde{\lambda} \cdot \tilde{v} \quad \text{subject to} \\ \text{(i)} \quad \tilde{v}_j = (1 - \delta)\tilde{g}_j(\alpha) + \delta\pi(\alpha) \cdot \tilde{w}_j \quad \text{for all } j, \\ \text{(ii)} \quad \tilde{v}_j = (1 - \delta)\tilde{g}_j(a_j, \alpha_{-j}) + \delta\pi(a_j, \alpha_{-j}) \cdot \tilde{w}_j \quad \text{for all } j \text{ and } a_j, \\ \text{(iii)} \quad \tilde{\lambda} \cdot \tilde{v} \geq \tilde{\lambda} \cdot \tilde{w}(y) \quad \text{for all } y. \end{aligned}$$

This is the problem of finding the maximum score for a known-state game (i.e., $|\Omega| = 1$) for direction $\tilde{\lambda}$, so its value (which does not depend on δ) follows from past work:

Claim 3. *Suppose the monitoring structure is known and has strong full rank. Then $\sup_{\alpha} \tilde{k}^*(\alpha, \tilde{\lambda}) = -\min_{\alpha_{-i}} \max_{a_i} \tilde{g}_i(a_i, \alpha_{-i})$*

Proof. Strong full rank implies that every pure action profile has individual full rank. Then from FLM Lemma 6.3, the maximal score for direction $\tilde{\lambda}$ is given by player i 's minimax score. Therefore, $\tilde{k}^*(\alpha, \tilde{\lambda}) = -\min_{\alpha_{-i}} \max_{a_i} \tilde{g}_i(a_i, \alpha_{-i})$. *Q.E.D.*

Claim 4. *Suppose the monitoring structure is known and has strong full rank. Let $\lambda \in \Lambda^5(i)$. Then $k^*(\vec{\alpha}, \lambda) = \tilde{k}^*(\alpha, \tilde{\lambda})$ if $\vec{\alpha}$ is a state-independent action α .*

Proof. First, we show $k^*(\vec{\alpha}, \lambda) \leq \tilde{k}^*(\alpha, \tilde{\lambda})$. When $k^*(\vec{\alpha}, \lambda) = -\infty$, then this inequality obviously follows. So assume $k^*(\vec{\alpha}, \lambda) > -\infty$. Choose (v, w) to satisfy constraints (i) through (iii) in the LP problem for $(\vec{\alpha}, \lambda, \delta)$, and let $\tilde{v}_j = -\sum_{\omega \in \Omega} \lambda_j^\omega v_j^\omega$ and $\tilde{w}_j = -\sum_{\omega \in \Omega} \lambda_j^\omega w_j^\omega(y)$ for all $j \in \mathbf{I}$ and $y \in Y$. Then this $(\tilde{v}, \tilde{w})$ satisfies all the constraints of the LP problem for $(\alpha, \tilde{\lambda}, \delta)$, and $\lambda \cdot v = \tilde{\lambda} \cdot \tilde{v}$. This shows that $k^*(\vec{\alpha}, \lambda) \leq \tilde{k}^*(\alpha, \tilde{\lambda})$.

Next, we show $k^*(\vec{\alpha}, \lambda) \geq \tilde{k}^*(\alpha, \tilde{\lambda})$. As before we restrict attention to the case of $\tilde{k}^*(\alpha, \tilde{\lambda}) > -\infty$.

We claim there are $(z_i^\omega(y))_{(\omega,y)}$ such that

$$(1 - \delta) \left(-\frac{\tilde{g}_i(a_i, \alpha_{-i})}{\sum_{\omega \in \Omega} \lambda_i^\omega} - g_i^\omega(a_i, \alpha_{-i}) \right) = \delta \pi(a_i, \alpha_{-i}) \cdot z_i^\omega \quad (9)$$

for all $\omega \in \Omega$ and $a_i \in A_i$, and

$$\sum_{\omega \in \Omega} \lambda_i^\omega z_i^\omega(y) = 0 \quad (10)$$

for all $y \in Y$. To see that this system has a solution, choose ω' such that $\lambda_i^{\omega'} \neq 0$, and eliminate $z_i^{\omega'}$ using (10). Then we can check that (9) for ω' are redundant equations; that is, (9) for ω' automatically holds if (9) holds for all $\omega \neq \omega'$. This leaves $(|\Omega| - 1) \times |A_i|$ equations and $(|\Omega| - 1) \times |A_i|$ unknowns, and strong full rank assures that the coefficient matrix has full rank. Therefore, the system has a solution.

Choose $(\tilde{v}, \tilde{w})$ to satisfy all the constraints of the LP problem for $(\alpha, \tilde{\lambda}, \delta)$, let $v_i^\omega = -\frac{\tilde{v}_i}{\sum_{\omega \in \Omega} \lambda_i^\omega}$, and $w_i^\omega(y) = -\frac{\tilde{w}_i(y)}{\sum_{\omega \in \Omega} \lambda_i^\omega} + z_i^\omega(y)$. Since $\lambda \cdot v = \tilde{\lambda} \cdot \tilde{v}$, it suffices to show that this (v, w) satisfies all the constraints of the LP problem for $(\alpha, \lambda, \delta)$. (We can ignore the adding-up constraint and the incentive compatibility constraint for player $j \neq i$, as strong full rank holds.) Note that

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

for all $a_i \in A_i$ with equality if $a_i \in \text{supp} \alpha_i$. Here, the second equality comes from (9), and the inequality comes from the fact that $(\tilde{v}, \tilde{w})$ satisfies the constraints of the LP problem for $(\alpha, \tilde{\lambda}, \delta)$. Therefore, this (v, w) satisfies constraints (i) and (ii). Also,

$$\lambda \cdot w(y) = \sum_{\omega \in \Omega} \lambda_i^\omega w_i^\omega(y) = \sum_{\omega \in \Omega} \lambda_i^\omega \left(z_i^\omega(y) - \frac{\tilde{w}_i(y)}{\sum_{\omega \in \Omega} \lambda_i^\omega} \right) = -\tilde{w}_i(y) \leq -\tilde{v}_i = \lambda \cdot v.$$

Here, the third equality comes from (10) and the inequality comes from the fact that $(\tilde{v}, \tilde{w})$ satisfies the constraints of the LP problem for $(\alpha, \tilde{\lambda}, \delta)$. Therefore, this (v, w) satisfies constraint (iii). Q.E.D.

It follows from Claims 3 and 4 and $\tilde{g}_j(a) = -\sum_{\omega \in \Omega} \lambda_i^\omega g_j^\omega(a)$ that

$$\begin{aligned} k^*(\lambda) &\geq \sup_{\alpha} \tilde{k}^*(\alpha, \tilde{\lambda}) = -\min_{\alpha_{-i}} \max_{a_i} \tilde{g}_i(a_i, \alpha_{-i}) \\ &= -\min_{\alpha_{-i}} \max_{a_i} -\lambda \cdot g(a_i, \alpha_{-i}) = \max_{\alpha_{-i}} \min_{a_i} \lambda \cdot g(a_i, \alpha_{-i}). \end{aligned}$$

On the other hand, Claim 2 shows that $k^*(\lambda) \leq \max_{\alpha_{-i}} \min_{a_i} \lambda \cdot g(a_i, \alpha_{-i})$. Therefore, $k^*(\lambda) = \max_{\alpha_{-i}} \min_{a_i} \lambda \cdot g(a_i, \alpha_{-i})$.

A.5 Proof of Lemma 24

Lemma 24. *Suppose the monitoring structure is known and has strong full rank. Then for each $\lambda \in \Lambda^6 \cup \Lambda^7$, $k^*(\lambda) = \max_{\alpha} \lambda \cdot g(\alpha)$.*

The proof consists of a series of claims.

Claim 5. *Let $\lambda \in \Lambda^6$, and let $i \in \mathbf{I}$ be such that $(\lambda_i^\omega)_{\omega \in \Omega} \neq 0$. Then*

- (a) *there is $\theta_i^* \in \Theta_i$ such that $\theta_i(\omega) = \theta_i^*$ for all ω such that $\lambda_i^\omega > 0$; and*
- (b) *for each $j \neq i$, there is $\theta_j^* \in \Theta_j$ that contains all ω such that $\lambda_i^\omega \neq 0$.*

Let $\lambda \in \Lambda^7$. Then

- (c) *for each $i \in \mathbf{I}$, there is $\theta_i^* \in \Theta_i$ that contains all ω such that $\lambda_j^\omega \neq 0$ for some $j \neq i$; and*
- (d) *this θ_i^* contains all ω such that $\lambda_i^\omega > 0$.*

Proof. For part (a), suppose not, so that there are $\omega' \in \Omega$ and $\omega'' \neq \omega$ such that $\theta_i(\omega') \neq \theta_i(\omega'')$, $\lambda_i^{\omega'} > 0$, and $\lambda_i^{\omega''} > 0$. Then $\lambda \notin \Lambda^6$, as for λ to be in Λ^6 , $\theta_i(\omega') = \theta_i(\omega'')$ for all $\omega' \in \Omega$ and $\omega'' \neq \omega$ satisfying $\lambda_i^{\omega'} > 0$ and $\lambda_i^{\omega''} > 0$. A contradiction.

For part (b), suppose that there are $\omega' \in \Omega$ and $\omega'' \neq \omega'$ such that $\theta_j(\omega') \neq \theta_j(\omega'')$, $\lambda_i^{\omega'} \neq 0$, and $\lambda_i^{\omega''} \neq 0$. Then $\lambda \notin \Lambda^6$, as for λ to be in Λ^6 , $\theta_j(\omega') = \theta_j(\omega'')$ for all $j \neq i$, $\omega' \in \Omega$, and $\omega'' \neq \omega'$ satisfying $\lambda_i^{\omega'} \neq 0$ and $\lambda_i^{\omega''} \neq 0$. A contradiction.

For part (c), suppose that there are (j, ω') and (l, ω'') such that $j \neq i$, $l \neq i$, $\theta_i(\omega') \neq \theta_i(\omega'')$, $\lambda_j^{\omega'} \neq 0$, and $\lambda_l^{\omega''} \neq 0$. Then $\lambda \notin \Lambda^7$, as the last condition of the

definition of Λ^7 requires that $\theta_i(\omega') = \theta_i(\omega'')$ for all $i \in \mathbf{I}$, $\omega' \in \Omega$, and $\omega'' \neq \omega'$ such that $\lambda_j^{\omega'} \neq 0$ for some $j \neq i$ and $\lambda_l^{\omega''} \neq 0$ for some $l \neq i$. A contradiction.

For part (d), suppose that there are $i \in \mathbf{I}$ and $\omega' \in \Omega$ such that $\lambda_i^{\omega'} > 0$ and $\omega \notin \theta_i^*$. Let (j, ω'') be such that $j \neq i$ and $\lambda_j^{\omega''} \neq 0$. Then from part (c), $\omega'' \in \theta_i^*$, so that $\theta_i(\omega'') = \theta_i^* \neq \theta_i(\omega')$. This implies that $\lambda \notin \Lambda^7$, as the last condition of the definition of Λ^7 requires that $\theta_i(\omega') = \theta_i(\omega'')$ for all $i \in \mathbf{I}$, $\omega' \in \Omega$, and $\omega'' \neq \omega'$ such that $\lambda_i^{\omega'} > 0$ and $\lambda_j^{\omega''} \neq 0$ for some $j \neq i$. A contradiction. *Q.E.D.*

Claim 6. Suppose the monitoring structure is known, and let $\lambda \in \Lambda^6 \cup \Lambda^7$. Then for each $\vec{\alpha} = ((\alpha_i^{\theta_i})_{\theta_i \in \Theta_i})_{i \in \mathbf{I}}$, $k^*(\vec{\alpha}, \lambda) \leq \lambda \cdot g(\alpha^{\theta^*})$ where θ^* is chosen as in Claim 5 and $\alpha^{\theta^*} = (\alpha_i^{\theta_i^*})_{i \in \mathbf{I}}$.

Proof. Choose (v, w) to satisfy constraints (i) through (iii) in the LP problem associated with $(\vec{\alpha}, \lambda, \delta)$ for some $\delta \in (0, 1)$. It follows from constraint (ii) that

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

for all $i \in \mathbf{I}$, $a_i \in A_i$, and $\omega \in \Omega$ such that $\lambda_i^\omega \neq 0$, since $\theta_j(\omega) = \theta_j^*$ for $j \neq i$ for such ω . In particular, we have

$$v_i^\omega \geq (1 - \delta)g_i^\omega(\alpha^{\theta^*}) + \delta\pi(\alpha^{\theta^*}) \cdot w_i^\omega \quad (11)$$

for all $i \in \mathbf{I}$ and $\omega \in \Omega$ such that $\lambda_i^\omega \neq 0$. Also, from constraint (i), we obtain

$$v_i^\omega = (1 - \delta)g_i^\omega(\alpha^{\theta^*}) + \delta\pi(\alpha^{\theta^*}) \cdot w_i^\omega \quad (12)$$

for all $i \in \mathbf{I}$ and $\omega \in \Omega$ such that $\lambda_i^\omega > 0$, since $\theta(\omega) = \theta^*$ for such $\omega \in \Omega$. It follows from (11) and (12) that

$$\begin{aligned} \lambda \cdot v &\leq \sum_{i \in \mathbf{I}} \sum_{\omega \in \Omega} \lambda_i^\omega \left[(1 - \delta)g_i^\omega(\alpha^{\theta^*}) + \delta\pi(\alpha^{\theta^*}) \cdot w_i^\omega \right] \\ &= (1 - \delta)\lambda \cdot g(\alpha^{\theta^*}) + \delta \sum_{y \in Y} \pi_y(\alpha^{\theta^*}) \lambda \cdot w(y). \end{aligned}$$

Using constraint (iii),

$$\lambda \cdot v \leq (1 - \delta)\lambda \cdot g(\alpha^{\theta^*}) + \delta \sum_{y \in Y} \pi_y(\alpha^{\theta^*}) \lambda \cdot v = (1 - \delta)\lambda \cdot g(\alpha^{\theta^*}) + \delta \lambda \cdot v.$$

Subtracting $\delta \lambda \cdot v$ from both sides and dividing by $(1 - \delta)$, we get $\lambda \cdot v \leq \lambda \cdot g(\alpha^{\theta^*})$. Therefore, $k^*(\vec{\alpha}, \lambda, \delta) \leq g(\alpha^{\theta^*})$. *Q.E.D.*

Claim 7. Suppose the monitoring structure is known and has strong full rank. Let $\lambda \in \Lambda^6$. Then $k^*(\lambda) \geq \max_{\alpha} \lambda \cdot g(\alpha)$.

Proof. Let $\alpha \in \arg \max_{\alpha'} \lambda \cdot g(\alpha')$. Without loss of generality we can assume that α is a pure action profile, so that we denote it by a . In what follows, we show that $k^*(a, \lambda) \geq \lambda \cdot g(a)$.

Let $\lambda \in \Lambda^6$, and let (i, ω') be such that $\lambda_i^{\omega'} > 0$. Consider the LP problem associated with (a, λ, δ) . Note that we can ignore constraints (i) and (ii) for $j \neq i$, as $(\lambda_j^{\omega})_{\omega \in \Omega} = 0$.

Let $v_i^{\omega} = g_i^{\omega}(a)$ for each $\omega \in \Omega$. For $\omega \neq \omega'$, let $(w_i^{\omega}(y))_{y \in Y}$ be such that

$$g_i^{\omega}(a) = (1 - \delta)g_i^{\omega}(a'_i, a_{-i}) + \delta\pi(a'_i, a_{-i}) \cdot w_i^{\omega}(y) \quad (13)$$

for all $a'_i \in A_i$. Also, let

$$w_i^{\omega'}(y) = \frac{1}{\lambda_i^{\omega'}} \left(\lambda \cdot g(a) - \sum_{\omega \neq \omega'} \lambda_i^{\omega} w_i^{\omega}(y) \right) \quad (14)$$

for all $y \in Y$.

We claim that this (v, w) satisfies constraints (i) through (iii) in the LP problem. First, constraints (i) and (ii) hold for $\omega \neq \omega'$, since (13) holds. Also, as in the proof of Claim 4, we have

$$\begin{aligned} & (1 - \delta)g_i^{\omega'}(a'_i, a_{-i}) + \delta\pi(a'_i, a_{-i}) \cdot w_i^{\omega'}(y) \\ &= g_i^{\omega'}(a) + (1 - \delta) \frac{\lambda \cdot g(a'_i, a_{-i}) - \lambda \cdot g(a)}{\lambda_i^{\omega'}} \leq g_i^{\omega'}(a) \end{aligned}$$

for all $a'_i \in A_i$ with equality if $a'_i = a_i$. Here, the inequality is from the fact that a maximizes $\arg \max \lambda \cdot g(a')$ and $\lambda_i^{\omega'} > 0$. This shows that constraints (i) and (ii) hold for ω' . Finally, constraint (iii) follows from (14). Thus we conclude $k^*(a, \lambda) \geq \sum_{\omega \in \Omega} \lambda_i^{\omega} v_i^{\omega} = \lambda \cdot g(a)$, as desired. *Q.E.D.*

Claim 8. Suppose the monitoring structure is known and has strong full rank. Let $\lambda \in \Lambda^7$. Then for each α , $k^*(\alpha, \lambda) \geq \lambda \cdot g(\alpha)$.

Proof. Let $\lambda \in \Lambda^7$, and given this λ , let $\lambda_{(i, \omega)(j, \omega')}$ be a direction such that the components for (i, ω) and (j, ω') are equal to those of λ and the remaining components are zero. (Thus the direction $\lambda_{(i, \omega)(j, \omega')}$ has at most two non-zero components.) In order to prove the claim, it suffices to show that α is enforceable with

respect to the hyperplane orthogonal to λ at $g(\alpha)$. This enforceability follows from the following two facts: (i) If the monitoring structure has strong full rank, then α is enforceable with respect to the hyperplane orthogonal to $\lambda_{(i,\omega)(j,\omega')}$ at $g(\alpha)$ for each (i, ω) and (j, ω') such that $i \neq j$ (but possibly $\omega = \omega'$), $\lambda_i^\omega \neq 0$, and $\lambda_j^{\omega'} \neq 0$. (ii) α is enforceable with respect to the hyperplane orthogonal to λ at $g(\alpha)$ if α is enforceable with respect to the hyperplane orthogonal to $\lambda_{(i,\omega)(j,\omega')}$ at $g(\alpha)$ for each (i, ω) and (j, ω') such that $i \neq j$, $\lambda_i^\omega \neq 0$, and $\lambda_j^{\omega'} \neq 0$. Note that (i) follows from Lemma 5.4 of FLM, since here we assume that the monitoring structure does not depend on ω . Likewise, (ii) follows from Lemma 5.3 of FLM, since $\lambda \in \Lambda^7$ implies that for each (i, ω) such that $\lambda_i^\omega \neq 0$, there is (j, ω') such that $i \neq j$ and $\lambda_j^{\omega'} \neq 0$. *Q.E.D.*

The statement of Lemma 24 follows from Claims 6, 7, and 8.

References

- [1] D. Abreu, D. Pearce, E. Stacchetti, Optimal cartel equilibria with imperfect monitoring, J. Econ. Theory 39 (1986), 251-269.
- [2] D. Abreu, D. Pearce, E. Stacchetti, Toward a theory of discounted repeated games with imperfect monitoring, Econometrica 58 (1990), 1041-1063.
- [3] S. Athey, K. Bagwell, Optimal collusion with private information, RAND J. Econ. 32 (2001), 428-465.
- [4] S. Athey, K. Bagwell, Collusion with persistent cost shocks, Econometrica 76 (2008), 493-540.
- [5] R. Aumann, M. Maschler, Repeated Games with Incomplete Information, MIT Press, Cambridge, MA, 1995. With the collaboration of R.E. Stearns.
- [6] D. Bergemann, S. Morris, Belief-free incomplete information games, mimeo, 2007.
- [7] O. Compte, Communication in repeated games with imperfect private monitoring, Econometrica 66 (1998), 597-626.
- [8] M. Cripps, J. Thomas, Some asymptotic results in discounted repeated games of one-side incomplete information, Math. Oper. Res. 28 (2003), 433-462.

- [9] J. Ely, J. Hörner, W. Olszewski, Belief-free equilibria in repeated games, *Econometrica* 73 (2005), 377-415.
- [10] J. Ely, J. Välimäki, A robust folk theorem for the prisoner's dilemma, *J. Econ. Theory* 102 (2002), 84-105.
- [11] G. Ellison, Cooperation in the prisoner's dilemma with anonymous random matching, *Rev. Econ. Stud.* 61 (1993), 567-588.
- [12] F. Forges, Note on nash equilibria in infinitely repeated games with incomplete information, *Int. J. Game Theory* 13 (1984), 179-187.
- [13] D. Fudenberg, D. Kreps, E. Maskin Repeated games with long-run and short-run players, *Rev. Econ. Stud.* 57 (1990), 555-573.
- [14] D. Fudenberg, D.K. Levine, Efficiency and observability in games with long-run and short-run players, *J. Econ. Theory* 62 (1994), 103-135.
- [15] D. Fudenberg, D.K. Levine, E. Maskin, The folk theorem with imperfect public information, *Econometrica* 62 (1994), 997-1040.
- [16] D. Fudenberg, D.K. Levine, S. Takahashi, Perfect public equilibrium when players are patient, *Games Econ. Behav.* 61 (2007), 27-49.
- [17] D. Fudenberg, Y. Yamamoto, Repeated games where the payoffs and monitoring structure are unknown, *Econometrica* 78 (2010), 1673-1710.
- [18] O. Gossner, N. Vieille, Strategic learning in games with symmetric information, *Games Econ. Behav.* 42 (2003), 25-47.
- [19] E.J. Green, R.H. Porter, Noncooperative collusion under imperfect price information, *Econometrica* 52 (1984), 87-100.
- [20] S. Hart, Nonzero-sum two-person repeated games with incomplete information, *Math. Oper. Res.* 10 (1985), 117-153.
- [21] J. Hörner, S. Lovo, Belief-free equilibria in games with incomplete information, *Econometrica* 77 (2009), 453-487.
- [22] J. Hörner, S. Lovo, T. Tomala, Belief-free equilibria in games with incomplete information: characterization and existence, mimeo, 2010.
- [23] M. Kandori, Social norms and community enforcement, *Rev. Econ. Stud.* 59 (1992), 63-80.
- [24] M. Kandori, Repeated games played by overlapping generations of players, *Rev. Econ. Stud.* 59 (1992), 81-92.

- [25] M. Kandori, Weakly belief-free equilibria in repeated games with private monitoring, forthcoming in *Econometrica*, 2010.
- [26] M. Kandori, H. Matsushima, Private observation, communication and collusion, *Econometrica* 66 (1998), 627-652.
- [27] E. Kohlberg, Optimal strategies in repeated games with incomplete information, *Int. J. Game Theory* 4 (1975), 7-24.
- [28] M. Piccione, The repeated prisoner's dilemma with imperfect private monitoring, *J. Econ. Theory* 102 (2002), 70-83.
- [29] R. Radner, Repeated principal-agent games with discounting, *Econometrica* 53 (1985), 43-58.
- [30] R. Radner, R. Myerson, E. Maskin, An example of a repeated partnership game with discounting and with uniformly inefficient equilibria, *Rev. Econ. Stud.* 53 (1986), 59-70.
- [31] J. Renault, T. Tomala, Learning the state of nature in repeated games with incomplete information and signals, *Games Econ. Behav.* 47 (2004), 124-156.
- [32] S. Sorin, Big match with lack of information on one side (part I), *Int. J. Game Theory* 13 (1984), 201-255.
- [33] S. Sorin, Big match with lack of information on one side (part II), *Int. J. Game Theory* 14 (1985), 173-204.
- [34] T. Wiseman, A partial folk theorem for games with unknown payoff distributions, *Econometrica* 73 (2005), 629-645.
- [35] T. Wiseman, A partial folk theorem for games with private learning, mimeo, 2010.
- [36] Y. Yamamoto, Efficiency results in n player games with imperfect private monitoring, *J. Econ. Theory* 135 (2007), 382-413.
- [37] Y. Yamamoto, A limit characterization of belief-free equilibrium payoffs in repeated games, *J. Econ. Theory* 144 (2009), 802-824.