A necessary and sufficient condition for a PBE Folk Theorem

Pablo Schenone[*,†]

Northwestern University

November 8, 2011

Abstract

This paper studies repeated games of incomplete information where each player knows his own payoffs and where the unknown state of the world can be identified by the combined private information of all players. We obtain a condition that is both necessary and sufficient for a Perfect Bayesian Equilibrium (PBE) folk theorem to hold. This contrasts with the existing literature where, due to the difficulty in keeping track of beliefs as play evolves, the analysis has focused on either Nash equilibrium for one-sided incomplete information or has dealt with various ex-post solution concepts. Finally, we also show the condition obtained is also necessary and sufficient to obtain Ex-Post folk theorems.

[*]Corresponding Author: pschenone@u.northwestern.edu

[†]I am very thankful to Wojciech Olszewski, Todd Sarver and Marciano Siniscalchi for useful discussions and their detailed comments. I am particularly grateful to Eddie Dekel for his comments and encouragement when this project was in its formative stages and for his continuing guidance and support. I also wish to thank Jeff Ely, Ehud Kalai, Mehmet Ekmekci, Nenad Kos, Peter Klibanoff, Laura Doval, Ignacio Franceschelli, Toan Phan, Javier Donna, Gregory Veramendi and other participants of the Northwestern theory bag lunch. All remaining errors are, of course, my own.

1 Introduction

Early studies of infinitely repeated games offered the insight that, by considering agents who care about future payoffs, certain behaviors that seem irrational in a static environment are optimal if the agents interaction is repeated. This insight has influenced a range of applications, and has also led to a large number of papers proving "folk theorems" in general environments. By a "folk theorem" we mean a theorem along the following lines: "Approximately any payoff profile that can be generated in the game (feasibility condition) and that gives each player at least his min-max payoff (individual rationality condition) can be supported as the result of an equilibrium strategy". Thus, when analyzing the payoff predictions of a given game under a given solution concept, a natural first step is to ask whether or not a folk theorem holds. In particular, for games of incomplete information, one can also ask how the initial beliefs of the players affect the set of sustainable payoffs.

We study a class of incomplete-information games and answer the question: what conditions are both necessary and sufficient for a Perfect Bayesian Equilibrium (PBE) folk theorem to hold? To the best of our knowledge this is the first folk theorem to be proved for this solution concept. The underlying assumptions on the class of games are the following: (i) full dimensionality of the complete-information games, (ii) the state of the world can be uniquely identified by pooling together the information of all players, (iii) all players know their own payoffs, and (iv) perfect monitoring and recall. We also show that this condition is necessary and sufficient for a Belief-Free Equilibrium (BFE) folk theorem to hold. An interesting corollary of these two results is that a PBE folk theorem either holds or not independently of initial priors. Finally, we obtain a result on the structure of the supporting strategies; in the long run, all player-types receiving positive probability must play in the same way.

Before developing our general results further, consider the following example, grounded in the Hotelling (1929) product differentiation model, and analyzed further in the examples section. A monopolist produces a good at 0 marginal cost and each period his choice variable is the price p he charges. Furthermore, this good can have different features: for example, the good could be ice cream and the features could be different flavors, packagings, recipes, etc. There is a representative consumer that has private information about his preferences for different features and each period decides whether to buy or not at the specified price.

If we assume the monopolist can only choose a price each period but cannot change the features of the product, our results imply that, for all interesting parameter values and regardless of the beliefs the monopolist holds about the preferences of the consumer, a folk theorem will not hold. Indeed, a folk theorem will hold only in the uninteresting case where the consumer does not want to buy the product, even if it is free. Intuitively, the folk theorem fails to hold because the private information of the consumer acts on utility by shifting it rigidly, but does not change the ordinal ranking of his preference over action profiles: lower price is always better, regardless of the consumer's type.

For this case, where a folk theorem fails so some feasible and individually rational payoff profiles cannot be approximated in equilibrium, we characterize the set of payoffs that can be approximated in Belief-Free equilibrium. Knowing that a folk theorem fails tells us that the model makes non-trivial payoff predictions, and this result shows how tight these predictions are. In particular, our results imply that we will have one degree of freedom when choosing what payoff vectors can be sustained: given a type of the consumer, we can specify any feasible and individually rational payoff for that type, but we will then be constrained in what we can sustain for the other type(s). In contrast, there will be no restrictions on what we can sustain for the monopolist. Appendix $A.5$ discusses a possible extension of this result to PBE, but this is ongoing work.

In the case where the monopolist can also change the product features each period, there is a non-trivial set of parameter values such that the folk theorem holds. This is because different types of consumers value different aspects of the product, and therefore the ordinal ranking of preferences changes across the different types of consumer: some might prefer a high price if this is accompanied by one specific feature while another type of consumer may prefer a lower price and another feature.

In general, two properties must be satisfied for a folk theorem to hold. First, for any given state and any given subset of players, there must be a "separating" action that delivers a "good" payoff in that state for that subset of players but delivers a "bad" payoff in any other state and for all players not included in the given subset. Second, we need an action that, uniformly across players and states, yields payoffs less than or equal to the min max payoff.

While it should not be surprising that such a condition is sufficient to obtain a PBE (hence BFE) folk theorem, the important fact is that the condition is also necessary. Indeed, necessity of the condition implies that for many games the folk theorem fails. This is in contrast with the complete-information case under perfect monitoring and recall, where a mild full-dimensionality condition on payoffs is sufficient to have a folk theorem. Moreover, while full-dimensionality is a generic condition (i.e., given any game that does not satisfy it there is a game that satisfies the condition and is "arbitrarily close" to the original game) this is not true of the condition presented here.

The first result, that our condition is necessary and sufficient to obtain a Perfect Bayesian folk theorem, makes two important contributions. Perfect Bayesian equilibrium is a solution concept that introduces a degree of non-stationarity into the model, as discussed in Section 4, and this has made it hard to analyze the set of equilibrium payoffs that can be approximated in equilibrium. By connecting PBE to BFE through our conditions we are able to deal with this non-stationarity in a novel way and this allows us to obtain results that could not be obtained otherwise. Second, while most folk theorems for other solution concepts provide only sufficient conditions for the result to hold, our condition is also necessary; and this allows us to identify the cases where the folk theorem fails. To show that a given model under a given solution concept produces non-trivial payoff predictions one must first prove that the folk theorem fails, and this is achieved with our theorem.

The second result, that the necessary and sufficient condition for obtaining a PBE folk theorem is also necessary and sufficient to obtain a BFE folk theorem, also has two important implications. First, because BFE is a solution concept that does not depend on the beliefs a player holds about the type of his rivals, whether a PBE folk theorem holds or not is independent of the aforementioned beliefs. Second, if one adopts PBE as the solution concept for a game, but upon checking our conditions concludes that a folk theorem holds, strengthening the solution concept to BFE will not refine the payoff predictions made by the model.

The third and final result, that in the long run all player-types receiving positive probability must play in the same way, is useful for two reasons. First, as discussed in Section 4, it provides some intuition as to why a PBE folk theorem holds if and only if a BFE folk theorem holds. Second, it provides a basis with which to rule out certain strategy profiles as potential equilibrium profiles. A simple proof of this result is given, based on the results by Kalai and Lehrer (1993).

1.1 Related Literature

While the study of conditions under which a PBE folk theorem holds is new, there are a number of studies on the equilibrium payoff set in repeated games of incomplete information. As already mentioned, studying solution concepts like PBE that (a) depend non-trivially on the beliefs a player holds about his rivals and (b) satisfy sequential rationality, can lead to difficulties. Thus, most papers in the literature take one of two routes: they either consider solution concepts that are independent of the beliefs a player might have about his rivals, but retain sequential rationality, or they study solution concepts that depend on these beliefs but are not sequentially rational, such as Nash equilibrium, oftentimes also restricting the number of informed parties. Our paper, by considering a solution concept that satisfies both (a) and (b) above, lies between these two extremes.

To prove that a BFE folk theorem holds if and only if a certain condition on payoff functions holds, we use techniques developed in Fudenberg-Yamamoto (2010), henceforth FY. In that paper, the authors start by defining a novel solution concept, Perfect Type-Contingent Ex-Post Equilibrium (PTXE), that coincides with BFE when there is perfect monitoring but differs otherwise. Then they develop a linear programming technique for analyzing what payoff profiles can be sustained under this solution concept and, finally, they use the technique to show sufficient conditions for a PTXE folk theorem to hold. Though their setting varies from ours substantially, so that our BFE folk theorem is very different from theirs, the technique we use to prove our theorem follows their approach closely. Indeed, rather than assuming a product information structure, known-own-payoffs and perfect monitoring, FY assume a general partition information model, do not impose known-own-payoffs and assume unknown monitoring. That is, monitoring is not only imperfect but also depends, in a non-trivial way, on the unknown state of the world; hence, this assumption neither implies nor is implied by perfect monitoring. Moreover, the unknown monitoring assumption drives many of the FY results: by endowing each player with a signal correlated with the type of his opponent, it provides an exogenous channel through which a player can learn the private information of his rival, regardless of the rival's action. This is not true when we consider perfect monitoring: if a player wishes to hide his private information from his rivals he can always play a pooling strategy and prevent any sort of updating by his opponent.

We also borrowed insights from the literature on Nash equilibrium. Hart (1985) proves, for two-player games with no discounting and one-sided incomplete information, that a feasible and individually rational payoff vector can be sustained if and only if certain incentive compatibility constraints hold. Shalev (1994) shows stronger results for the special case of known-own-payoffs but still in a one-sided incomplete information, no-discounting framework. In particular, for Nash equilibrium and if priors have full support, Shalev shows the set of sustainable payoffs is independent of said priors. Cripps and Thomas (2003) and Peski (2008) analyze similar results for Nash equilibrium and one-sided incomplete information but now for the case with discounting. The common feature of these works is that the analysis of the set of sustainable payoffs is based on the analysis of *finitely revealing strategies*. These are strategies where the informed party reveals his information to the uninformed party, either through an explicitly modeled communication device or simply through his choice of actions, but only in a finite set of periods. These strategies are reminiscent of our asymptotically constant strategies, which play an important role in proving our results.

The paper is organized as follows. Section 2 introduces the main notation we will need. Section 3 presents the necessary and sufficient conditions for a game to have a BFE folk theorem in two theorems: Theorem 1 deals with the simple case of one informed player and many uninformed players, while Theorem 2 is the generalization to an arbitrary number of informed payers. The former is presented as it enables a simpler and less cumbersome presentation of the main ideas. These results, while interesting in their own right, are also the building blocks for the results on Perfect Bayesian Equilibrium (PBE), presented in Section 4. Theorem 4 states that if a folk theorem holds for the PBE solution concept, then the game must satisfy the condition that guarantees a BFE folk theorem. Schematically, we summarize these results as:

$$\text{Conditions} \underset{\text{Theorems 1,2}}{\Leftrightarrow} \text{BFE Folk Theorem} \underset{\text{Trivial}}{\Rightarrow} \text{PBE Folk Theorem} \underset{\text{Theorem 4}}{\Rightarrow} \text{Conditions}$$

Section 5 develops simple examples, including the product-choice game described above, that illustrate how these results are applied. Section 6 concludes.

2 Model: Notations and general assumptions

Information and static payoffs: There is a finite set of players, N. Each of these players is endowed with a finite action set A_n. An element $a_n \in A_n$ is an *action* for player

n and an element $a \in \prod_{n \in N} A_n$ is an *action profile*. For any set X, the set $\Delta(X)$ will denote all probability distributions over X. In particular, we refer to elements of $\Delta(A)$ as *mixed actions*. Information is modeled as a collection of finite sets $\{\Theta_n : n \in N\}$, one for each player. An element of $\theta_n \in \Theta_n$ is player n's *type* or, when the identity of the player is relevant, we refer to it as a *player-type*. An element $\theta \in \prod_{n \in N} \Theta_n$ is a *state*.

We assume private-values: each player $n \in N$ has a Bernoulli utility function $u_n : A \times \Theta_n \to \mathbb{R}$ that will be extended to payoffs over lotteries of action profiles in a standard expected-utility fashion. We denote the utility of player $n \in N$ from lottery $\alpha \in \Delta(A)$ when his type is θ_n as $u_n(\alpha|\theta_n)$. We also use $u_n(\alpha|\theta)$ to save notation, but it should be understood that only the n^{th} dimension of θ affects u_n.

Since we allow for the possibility that some players have no private information, it is useful to distinguish between the players that have private information and those that do not. Denote the set of players with private information as $I \subsetneq N$. Thus, players $n \notin I$ have type sets Θ_n that are singletons. Finally, we denote with $\underline{w}_n(\theta_n)$ the minimax payoff of player n when his type is θ_n; $\underline{w}_n(\theta_n) \equiv \min_{\alpha_{-n}} \max_{a_n} u_n(a_n, \alpha_{-n}|\theta_n)$.

Full dimensionality and maximal payoffs: We assume that each complete information game $G(\theta)$ has full dimensionality. Let $V(\theta)$ be the feasible, individually rational payoffs, $V(\theta) \equiv \{v \in \mathbb{R}^N : v \in co\{(u_n(a|\theta_n))_{n \in N} : a \in A\},\ v_n \geq \underline{w}_n(\theta_n)\}$; we assume that $\dim V(\theta) = N$ for all $\theta \in \Theta$. This guarantees that the standard folk theorem results apply in each of the complete-information games. Let $\hat{w}_n(\theta) = \max\{v_n \in \mathbb{R} : (\exists v_{-n} \in V_{-n}(\theta)) : (v_n, v_{-n}) \in V(\theta)\}$. As a consequence of the folk theorem for complete-information games, $\hat{w}_n(\theta)$ is the (tight) upper bound on the payoff player-type θ_n can obtain in any equilibrium. Finally, let $V \equiv \prod_{\theta \in \Theta} V(\theta)$.

Histories, strategies and updating: For every $t \in \mathbb{N}$, $\mathcal{H}_t \equiv A^t$ will be the set of histories of length t, while $\mathcal{H} = \bigcup_t \mathcal{H}_t$ will be the set of all possible histories. We assume perfect monitoring and recall: at the beginning of every period $t+1$, the full history h^t of past play is known to the players. For any two histories of any two lengths, h^t and h'^{τ}, we can concatenate them into a history of length $t+\tau$ that starts off as h^t and continues as h'^{τ}. We use $h^t * h'^{\tau}$ to denote the concatenated history. We use S_n to denote the set of behavior strategies for player n, $S_n : \mathcal{H} \to \Delta(A_n)$, an element $\sigma_n \in S_n$ is a *strategy* for player n, and a vector $\sigma \equiv (\sigma_n^{\theta_n})_{n,\theta_n} \in \prod_{n\in N} S_n^{\Theta_n}$ is a *strategy profile.* Given a history h^t we define the strategy conditional on arriving at h^t as $\sigma_{h^t} \in S$; $\sigma_{h^t}(h'^{\tau}) \equiv \sigma(h^t * h'^{\tau})$. Since at the beginning of the game the full state θ might be unknown to the players, each player $n \in N$ has a prior over the information of his rivals. It is assumed that there is a common prior, denoted $\mu \in \Delta(\Theta)$, over the set Θ of states. The beliefs a player-type θ_n has about the types of his rivals are obtained by conditioning the common prior on his own type. We denote such beliefs by $\mu(\cdot|\theta_n) \in \Delta(\Theta_{-n})$. Finally, a strategy profile σ and a prior $\mu(\cdot|\theta_n)$ induce, for each player-type θ_n, a distribution over states and histories via Bayesian updating, denoted as $P(\cdot|\theta_n) \in \Delta(\mathcal{H} \times \Theta_{-n})$.

Dynamic Payoffs: Finally, given a discount factor $\delta \in (0,1)$, a history h^t (possibly empty) a player-type $\theta_n \in \Theta_n$ and a distribution $P \in \Delta(\mathcal{H} \times \Theta_{-n})$, we let $U_n(\sigma_{h^t}|\delta, \theta_n, P)$ denote the δ discounted expected payoff to player n from following the conditional strategy σ_{h^t} starting at history h^t. In symbols, $U_n(\sigma_{h^t}|\delta, \theta_n, P(\cdot|\theta_n, h^t)) = (1-\delta)E_{P(\cdot|\theta_n,h^t)}\{\sum_{\tau=0} \delta^{\tau} u_n(\sigma_{h^t}(h^{'\tau})|\theta_n)\}$

2.1 Equilibrium definitions

A BFE is a strategy profile such that, for each state, the state-contingent behaviors form a subgame-perfect Nash Equilibrium (SPNE). Formally,

Definition (Belief Free Equilibrium). *A strategy profile σ is a* Belief-Free Equilibrium *if, for each $\theta \in \Theta$, $(\sigma_n^{\theta_n})_{n\in N}$ is a SPNE of the complete-information game $G(\theta)$*

A PBE is a strategy profile σ and beliefs μ that is both sequentially rational given the beliefs each player holds about his rivals and where beliefs are updated using Bayes rule whenever possible. Formally, for a fixed $\mu \in \Delta(\theta)$, say a strategy profile $\sigma \in S$ is *sequentially rational for beliefs* $P \in \Delta(\Theta \times \mathcal{H})$ if, for every history $h \in \mathcal{H}$, period

$t \in \mathbb{N}$, player $n \in N$, state $\theta \in \Theta$, and deviation $\sigma'_n \in S_n$, $U_n(\sigma_{h^t}|\delta, \theta_n, P(\cdot|h^t)) \geq U_n(\sigma'_n, \sigma_{h^t,-n}|\delta, \theta_n, P(\cdot|h^t))$.

Definition (Perfect Bayesian Equilibrium). *Let $\mu \in \Delta(\Theta)$ be the common prior and $\delta \in (0,1)$ the common discount factor. A strategy profile $\sigma \in S$ and beliefs P are a (μ, δ) Perfect Bayesian Equilibrium if σ is sequentially rational for P and P is derived from μ and σ via Bayes' rule whenever possible.*

3 A folk theorem for Belief Free Equilibrium

3.1 One informed player

The case where there is only one informed player is presented first. The condition for a BFE folk theorem to hold is that a system of inequalities admits a solution, and this system is smaller when there is only one informed player. Therefore, without loss of generality, let $I = \{1\}$ and notice that θ and Θ are now isomorphic to θ_1 and Θ_1 respectively. Hence, we drop the subscript for notational convenience.

Condition A, which is necessary and sufficient for a BFE folk theorem to hold, is an identification condition. Given any state, there must be a collection of action profiles that bound the payoff player 1 receives in that state (with a bound no less than the min max payoff) while simultaneously capping his utility at the min max level for any other state. We present it as two sub-conditions: conditions $A.1$ and $A.2$. The main difference between $A.1$ and $A.2$ is that $A.1$ focuses solely on player 1 while $A.2$ focuses on all players simultaneously.

Condition A.1. *For any $\theta^+ \in \Theta$, there exists an action profile $\alpha_{\theta^+} \in \Delta(A)$ such that:*

- $u_1(\alpha_{\theta^+}|\theta^+) \geq \hat{w}_1(\theta^+)$
- $u_1(\alpha_{\theta^+}|\theta) \leq \underline{w}_1(\theta)$ *for all $\theta \neq \theta^+$*

Condition $A.1$ states that an action exists which rewards player 1 (with his maximum possible payoff) in a given state, while simultaneously punishes him (with his min max payoff) if not. In particular, condition $A.1$ requires conflict of interests across the different types of player 1: given any two states, at least one of the most preferred actions for player 1 in one state must not be individually rational in the other, and vice versa. For instance, a game where two types of player 1 have the same ordinal preferences over pure actions violates this condition.[1]

Before stating condition $A.2$ it is first necessary to define what a $J-$ *dominated* payoff profile is. In words, a payoff $v(\theta) \in V(\theta)$ is J-dominated by a payoff $v'(\theta) \in V(\theta)$ for some subset $J \subset N$ if they are different, $v'(\theta)$ weakly Pareto dominates $v(\theta)$ for all players in J, and $v(\theta)$ weakly Pareto dominates $v'(\theta)$ for the other players. Formally, we have the following definition.

Definition (J-domination). *Let $\theta \in \Theta$ and $J \subset N$. A payoff profile $v(\theta) \equiv (v_n(\theta))_{n \in N} \in V(\theta)$ J-dominates $v'(\theta) \equiv (v'_n(\theta))_{n \in N} \in V(\theta)$ if $v(\theta) \neq v'(\theta)$, $v_n(\theta) \geq v'_n(\theta)$ for all $n \in J$, and $v_n(\theta) \leq v'_n(\theta)$ for all $n \notin J$.*

Condition A.2. *For any $\theta^+ \in \Theta$, any $J \subset N$ and any J-undominated profile $(v_i(\theta^+)) \in V(\theta^+)$ there exists an action profile $\alpha_{J,\theta^+} \in \Delta(A)$ such that:*

- $u_n(\alpha_{J,\theta^+}|\theta^+) \geq v_n(\theta^+)$ *for all* $n \in J$
- $u_n(\alpha_{J,\theta^+}|\theta^+) \leq v_n(\theta^+)$ *for all* $n \notin J$
- $u_1(\alpha_{J,\theta^+}|\theta) \leq \underline{w}_1(\theta)$ *for all* $\theta \in \Theta$, $\theta \neq \theta^+$

Condition $A.2$ focuses on all players simultaneously. Formally, given a state θ^+, any extreme point of $V(\theta^+)$ can be approximated while simultaneously capping the utility of player 1 in any other state with his min max payoff. By convexity of $V(\theta^+)$, the

[1]That ordinal preferences over pure actions are independent of θ does not imply that the preferences over lotteries, or mixed actions, are independent of type.

same applies to any payoff in $V(\theta^+)$, since any such payoff is a convex combination of extreme points. Conceptually then, $A.2$ says that payoffs in game $G(\theta^+)$ can be analyzed independently of payoffs in any other game $G(\theta)$, $\theta \neq \theta^+$. Given any feasible and individually rational payoff in state θ^+, we can approximate that payoff with an action that, in any other state, is not individually rational for player 1; hence, no other type of player 1 has incentives to mimic this action.

Condition $A.2$ has another important implication. By setting $J = \emptyset$, $A.2$ states that there exists an action that delivers, for all player-types simultaneously, a payoff no greater than the min max.

Theorem 1. *Condition A holds if and only if for any strict, smooth subset $W \subset V$ there exists $\bar{\delta} \in (0,1)$ such that, for any $\delta \geq \bar{\delta}$ $W \subset E^{BFE}(\delta)$*

The following subsection offers a proof sketch of Theorem 1. Readers interested in the other main results may skip it without loss of continuity.

3.1.1 Theorem 1 proof sketch

To give a proof sketch of Theorem 1 it is important to sketch the basic argument behind FY first.[2] We seek conditions under which $V = \lim_{\delta\to 1} E^{BFE}(\delta)$, where $E^{BFE}(\delta)$ is the set of BFE payoff profiles when discount factor is δ. It is straightforward to show that $\lim_{\delta\to 1} E^{BFE}(\delta) \subset V$; to show the opposite inclusion requires more work. As in Abreu-Pearce-Stacchetti (1990), heceforth APS, or Fudenberg-Levine-Maskin (1994), heceforth FLM, any closed and convex set $W \subset \mathbb{R}^{|\Theta|\times N}$ satisfying recursive conditions analogous to the self-replication conditions in FLM will be such that $W \subset \lim_{\delta\to 1} E^{BFE}(\delta)$. Since these sets are convex, they can be expressed as the intersection of a collection of half-spaces; i.e., given a convex set $W \subset \mathbb{R}^{|\Theta|\times N}$, there is a collection of scores $(k(\lambda))_{\lambda\in\mathbb{R}^{|\Theta|\times N}}$, such that $W = \cap_{\lambda\in\mathbb{R}^{|\Theta|\times N}}\{v : v\cdot\lambda \leq k(\lambda)\}$. Formally, the vectors λ represent the normal to each of the supporting hyperplanes; intuitively, they represent the directions

[2] Complete proofs and more detail on FY are given in the appendix.

in which the hyperplanes increase. Suppose $Q \subset \mathbb{R}^{|\Theta| \times N}$ is closed, convex and such that (a) Q satisfies appropriate recursive conditions and (b) the corresponding scores, $(k^*(\lambda))_{\lambda \in \mathbb{R}^{|\Theta| \times N}}$, satisfy $v \cdot \lambda \leq k^*(\lambda)$ for all $v \in V$ and all $\lambda \in \mathbb{R}^{|\Theta| \times N}$. From (a) we conclude $Q \subset \lim_{\delta \to 1} E^{BFE}(\delta)$; from (b), $V \subset \cap_{\lambda \in \mathbb{R}^{|\Theta| \times N}} \{v : v \cdot \lambda \leq k\lambda\} = Q$. Hence, from (a) and (b), we conclude $V \subset \lim_{\delta \to 1} E^{BFE}(\delta)$. This, and the opposite inclusion, show $V = \lim_{\delta \to 1} E^{BFE}(\delta)$.

Fudenberg and Yamamoto find a set Q that satisfies (a) above and characterize the scores $\{k^*(\lambda)\}_{\lambda \in \mathbb{R}^{N \times \Theta}}$; we take these scores and impose $k^*(\lambda) \geq \max\{v \cdot \lambda : v \in V\}$ for each $\lambda \in \mathbb{R}^{N \times \Theta}$ to obtain condition A.

For most directions $\lambda \in \mathbb{R}^{N \times \Theta}$ either $k^*(\lambda) = +\infty$ or $k^*(\lambda) = \max\{\lambda \cdot v : v \in V\}$, so the inequality is satisfied automatically. Only three sets of directions, Λ^5, Λ^6, and Λ^7, allow for $k^*(\lambda) < \max\{\lambda \cdot v : v \in V\}$. Condition $A.1$ guarantees $k^*(\lambda) \geq \max\{\lambda \cdot v : v \in V\}$ for $\lambda \in \Lambda^6$ and condition $A.2$ guarantees $k^*(\lambda) < \max\{\lambda \cdot v : v \in V\}$ for $\lambda \in \Lambda^7$. A corollary of condition $A.2$ is that $k^*(\lambda) \geq \max\{\lambda \cdot v : v \in V\}$ for $\lambda \in \Lambda^5$.

The set Λ^6 is the set of directions $\lambda \in \mathbb{R}^{|\Theta| \times N} \setminus \{0\}$ such that (a) only the informed player has non zero weight, i.e, $\lambda_n^\theta = 0$ for all $n \neq 1$ and all λ; and (b) there is a unique "positive" state θ^+ such that $\lambda_1^{\theta^+} > 0$.[3] For these directions,

$$k^*(\lambda) = \max\{\sum_{n,\theta} \lambda_n^\theta u_n(\alpha|\theta) : \alpha \in \Delta(A)\}$$

We can rewrite this as

$$k^*(\lambda) = \lambda_1^{\theta^+} u_1(\alpha^*|\theta^+) + \sum_{\theta \neq \theta^+} \lambda_1^\theta u_1(\alpha^*|\theta)$$

On the other hand, it is simple to show that for these directions,

$$\max\{\lambda \cdot v : v \in V\} = \lambda_1^{\theta^+} \hat{w}_1(\theta^+) + \sum_{\theta \neq \theta^+} \lambda_1^\theta \underline{\mathrm{w}}_1(\theta)$$

[3]In general, when there is an arbitrary number of informed players, these are directions where (a) only one player has non-zero weights, (b) he has positive weight on at least one state (c) he cannot distinguish between two states in which his weights are positive (d) no other player can distinguish between two states that receive non-zero weight. Details are included in the appendix.

Comparing the expressions for k^* and $\max\{\lambda \cdot v : v \in V\}$ term-by-term shows that condition $A.1$ implies $k^*(\lambda) \geq \lambda \cdot v$ for all $v \in V$ and all $\lambda \in \Lambda^6$. Necessity is a consequence of the linearity of the problem.

The set Λ^7 is the set of directions $\lambda \in \mathbb{R}^{N\times\Theta}$ such that (a) player 1 and at least one other player $n \neq 1$ have non-zero weights (b) there is at most one state $\theta^+ \in \Theta$ such that uninformed players receive non zero weight and (c) the informed player can only receive positive weight in that state θ^+.[4] For these directions

$$k^*(\lambda) = \max\{\sum_{n,\theta} \lambda_n^\theta u_n(\alpha|\theta) : \alpha \in \Delta(A)\}$$

Re-writing this expression, separating the states that receive positive weight from those that receive negative weight, and using that $\lambda \in \Lambda^7$ yields the following:

$$k^*(\lambda) = \sum_{\theta\neq\theta^+} \lambda_1(\theta)u_1(\alpha|\theta) + \sum_{\{n:\lambda_j(\theta^+)>0\}} \lambda_n(\theta^+)u_n(\alpha|\theta^+) + \sum_{\{n:\lambda_j(\theta^+)\leq 0\}} \lambda_n(\theta^+)u_n(\alpha|\theta^+)$$

The maximization problem $\max\{\lambda \cdot v : v \in V\}$ is solved in two steps. First, if v^* solves the above problem for some $\lambda \in \Lambda^7$, then $v_1^*(\theta) = \underline{\mathrm{w}}_1(\theta)$ for all $\theta \neq \theta^+$. Second, if v^* solves the above problem for some $\lambda \in \Lambda^7$ and $J = \{n : \lambda_n^{\theta^+} > 0\}$, $v^*(\theta^+)$ must be $J-$*undominated*. Thus,

$$\max\{\lambda\cdot v : v \in V\} = \sum_{\theta\neq\theta^+} \lambda_1(\theta)\underline{\mathrm{w}}_1(\theta) + \sum_{\{n:\lambda_j(\theta^+)>0\}} \lambda_n(\theta^+)v_n^*(\theta^+) + \sum_{\{n:\lambda_j(\theta^+)\leq 0\}} \lambda_n(\theta^+)v_n^*(\theta^+)$$

Comparing the expressions for $k^*(\lambda)$ and $\max\{\lambda \cdot v : v \in V\}$ term-by-term shows that condition $A.2$ implies $k^*(\lambda) \geq \lambda \cdot v$ for all $v \in V$ and all $\lambda \in \Lambda^7$. Necessity is, again, a consequence of the linearity of the problem.

The set Λ^5 is the set of directions $\lambda \in \mathbb{R}^{N\times\Theta}$ such that (a) only the informed player

[4] In general, these are directions where (a) at least two players receive non-zero weight, (b) at least two states receive non zero-weight, (c) if player n receives non-zero weight at state θ, and player n' (possibly $n = n'$) receives non-zero weight in state θ', no other player $n'' \neq n, n'$ can distinguish θ and θ' and (d) if a player n receives positive weight in state θ and a player $n' \neq n$ receives non zero weight in state θ', player n cannot distinguish θ from θ'. Again, see FY (2010) or Appendix A for details.

receives non zero weights and (b) the informed player receives non-positive weights in all states.[5] For these directions

$$k^*(\lambda) = -\min_{\alpha_{-1}} \max_{a_1} \sum_{n,\theta} |\lambda_1^\theta| u_1(a_1, \alpha_{-1}|\theta)$$

In contrast,

$$\max\{\lambda \cdot v : v \in V\} = -\sum_{\theta} |\lambda_1^\theta| \underline{\text{w}}_1(\theta)$$

Condition $A.2$, when $J = \emptyset$, implies that there is an action α such that, if α is played, all player-types receive payoffs no greater than their min max payoffs. This implies

$$k^*(\lambda) \geq \max\{\lambda \cdot v : v \in V\}$$

for all $\lambda \in \Lambda^5$.

Summarizing, condition A implies that $k^*(\lambda) \geq \max\{v \cdot \lambda : v \in V\}$ for all λ, this implies that $V \subset \lim_{\delta \to 1} E^{BFE}(\delta)$ and as a consequence the folk theorem holds. The converse is due to the linearity of the problem and details are given in the appendix.

3.2 Multiple informed players

Generalizing the case of one informed player to the case of many informed players is straightforward. Condition A requires that each type of the informed player can be separated from the other types via state-contingent payoffs that arise from state non-contingent actions. Since a type of the informed player is equivalent to a full state, this condition also implies we can separate any state from any other state. Condition B also requires this separation, but emphasizes the difference between a full state and a type of a particular player $i \in I$. Condition $B.1$ requires that for each informed player and for each of his types there exists an action that separates that type from the rest

[5] In general, this direction is one where (a) a unique player n has non-zero weights, (b) his weights are non-positive and (c) no other player can distinguish between two states that receive non-zero weight.

of his types; condition $B.2$ requires that for every state there has to be an action that separates that state from the other states.

Condition B.1. *For any $i \in I$ and $\theta_i^+ \in \Theta$ there exists an action profile $\alpha \in \Delta(A)$ such that:*

- $u_i(\alpha|\theta_i^+) - \hat{w}_i(\theta_i^+) \geq 0$
- $u_i(\alpha|\theta_i) - \underline{w}_i(\theta_i) \leq 0$ *for all* $\theta_i \neq \theta_i^+$

Condition B.2. *For any $\theta^+ \in \Theta$, for any $J \subset N$ and for any J-undominated profile $(v_n(\theta^+)) \in V(\theta^+)$ there exists an action profile $\alpha \in \Delta(A)$ such that:*

- $u_n(\alpha|\theta_n^+) - v_n(\theta^+) \geq 0$ *for all* $n \in J$
- $u_n(\alpha|\theta_n^+) - v_n(\theta^+) \leq 0$ *for all* $n \notin J$
- $u_i(\alpha|\theta_i) - \underline{w}_i(\theta_i) \leq 0$ *for all* $i \in I$, $\theta_i \in \Theta_i$, *and* $\theta_i \neq \theta_i^+$

As before, we let condition B stand for the conjunction of $B.1$ and $B.2$.

Theorem 2. *Condition B holds if and only if for any strict, smooth subset $W \subset V$ there exists $\bar{\delta} \in (0,1)$ such that, for any $\delta \geq \bar{\delta}$, $W \subset E^{PTXE}(\delta)$*

The proof of this theorem follows step by step proof of theorem 1; the only difference is the definition of the sets Λ^5, Λ^6, and Λ^7. The formal argument is presented in the appendix.

4 Perfect Bayesian Equilibrium

Although PBE has been the standard solution concept for incomplete-information dynamic games in general, and repeated games in particular, to the best of our knowledge this is the first folk theorem developed for this solution concept. It has been hard to obtain folk theorems for PBE because the existing techniques used to analyze the equilibrium payoff-set exploit that, both in terms of continuation behaviors and payoffs that can be supported in equilibrium, the game at any one node is indistinguishable from the game at any other node. The difficulty in dealing with Perfect Bayesian equilibrium is that belief updating, both on and off the path of play, breaks this indistinguishability: since, a priori, the set of continuation behaviors and payoffs that can be sustained in equilibrium depends on the initial beliefs that players hold, we cannot claim that the continuation games at nodes with different posteriors are indistinguishable Thus, the standard recursive techniques used in other works do not apply.

The main result in this section, Theorem 4, states that if a PBE folk theorem holds for a given game then condition B must be satisfied. As a corollary, a PBE folk theorem holds if and only if a BFE folk theorem holds. To understand, intuitively, why this result is true, note that the main difference between these solution concepts is the type of *incentive compatibility*, or *non-deviation*, condition imposed. In BFE we ask that players behave in such a way that they are best responding to the strategy of their rivals *for each of their possible types*. This is a "pointwise" constraint, because we are giving appropriate incentives in a type-by-type fashion. On the other hand, the incentive compatibility constraint implicit in PBE asks players to play a best response to their *beliefs* about the play of the rivals. This is an "'expected" or "on average" constraint. More formally, BFE imposes that, for each player-type θ_n,

$$U_n(\sigma_n^{\theta_n}, (\sigma_{n'}^{\theta_{n'}})_{n' \neq n, \theta_{n'} \in \Theta_{n'}} | \theta_n, \delta, P(\cdot|\theta_n)) \geq U_n(\hat{\sigma}_n, (\sigma_{n'}^{\theta_{n'}})_{n' \neq n, \theta_{n'} \in \Theta_{n'}} | \theta_n, \delta, P(\cdot|\theta_n))$$

for any $\hat{\sigma}_n$, $(\theta_{n'})_{n' \neq n} \in \Theta_{-n'}$. Note that here the beliefs $P(\cdot|\theta_n)$ held by player-type θ_n play no role whatsoever. In contrast, PBE imposes

$$U_n(\sigma_n^{\theta_n}, (\sigma_{n'})_{n' \neq n} | \theta_n, \delta, P(\cdot|\theta_n)) \geq U_n(\hat{\sigma}_n, (\sigma_{n'})_{n' \neq n} | \theta_n, \delta, P(\cdot|\theta_n))$$

for any $\hat{\sigma}_n$. Theorem 3, presented below, links the previous conditions through two important claims: first, as time goes to infinity, the beliefs a player holds about his rivals converge to a limiting distribution; second, two player types receiving positive probability under this limiting distribution must play in the same way. Thus, in the limit when $t \to \infty$, all player-types that affect the PBE incentive constraints are playing the same strategy; hence the pointwise and the average constraints coincide. Since beliefs converge, if δ is sufficiently high, the long run payoffs dominate the (overall) expected discounted utility for each player-type θ_n and this completes the argument.

Nevertheless, none of these asymptotic results are required for the case $I = \{1\}$. Since, by definition, the informed player knows the types of his rivals, pointwise and on average constraints coincide even in the short run, thus no asymptotic considerations have to be made.

A sketch of the argument for this simpler case follows. Assume a PBE folk theorem holds, pick any $\theta^+ \in \Theta$, and consider the payoff vector $v \in V$ such that (a) $v_1(\theta^+) = \hat{w}_1(\theta^+)$, (b) $v_1(\theta) = \underline{\mathrm{w}}_1(\theta)$ and (c) payoffs for uninformed players are arbitrarily chosen. Since a folk theorem holds, we can approximate this payoff with payoffs from PBE strategies. Concretely, for each $k \geq 1$, let $v^k \in V$ be such that (a) $v^k \to v$, (b) v^k is in the interior of V, so that there is (δ_k, σ^k) that supports v^k as a PBE. We use the fact that for each $n \neq 1$ the set Θ_n is a singleton to calculate

$$v_1^k(\theta) = E_h\{\sum_t (1-\delta_k)\delta_k^t u_1(\sigma_1^{k^\theta}(h^t), \sigma_{-1}^k(h^t)|\theta)\} = E_{h,\theta_{-1}}\{\sum_t (1-\delta_k)\delta_k^t u_1(\sigma_1^{k^\theta}(h^t), \sigma_{-1}^k(h^t)|\theta)\}$$

We then use incentive compatibility to calculate:

$$v_1^k(\theta) = E_h\{\sum_{t\geq 1} (1-\delta_k)\delta_k^t u_1(\sigma_1^{k^\theta}(h^t), \sigma_{-1}^k(h^t)|\theta)\} \geq E_h\{\sum_{t\geq 1} (1-\delta_k)\delta_k^t u_1(\sigma_1^{k^{\theta^+}}(h^t), \sigma_{-1}(h^t)|\theta)\}$$

where equality holds if $\theta = \theta^+$. Since Bernoulli utilities are linear in probabilities, the above expression can be re written as:

$$v_1^k(\theta) \geq u_1(E_h\{\sum (1-\delta_k)\delta_k^t (\sigma_1^{k^{\theta^+}}, \sigma_{-1}^k)(h^t)\}|\theta)$$

Let $\alpha_k \equiv E_h\{\sum (1-\delta_k)\delta_k^t (\sigma_1^{k^{\theta^+}}, \sigma_{-1}^k)(h^t)\}$. Since $\Delta(A)$ is compact, hence every sequence has a convergent sub-sequence, $\alpha_k \to \alpha$ for some $\alpha \in \Delta(A)$, up to a sub-sequence. Thus, taking limits as $k \to \infty$ and noting that equality holds when $\theta = \theta^+$ we get

$$v_1(\theta) \geq u_1(\alpha|\theta) \text{ with equality if } \theta = \theta^+$$

which, replacing $v_1(\theta)$, means

- $\hat{w}_1(\theta^+) = u_1(\alpha|\theta^+)$
- $\underline{\mathrm{w}}_1(\theta) \geq u_1(\alpha|\theta)$ for $\theta \neq \theta^+$

and this is precisely what we need for condition $A.1$. That condition $A.2$ holds is proved in an analogous manner.

As already mentioned, the proof for the case of multiple informed parties requires a previous result, Theorem 3. First, define the notion of an asymptotically constant strategy:

Definition (Asymptotically Constant Strategy). *A strategy profile σ is asymptotically constant if for every $\varepsilon > 0$, for any θ, θ' and for P-almost all histories, there exists $T \in \mathbb{N}$ such that, for all $\tau \geq T$, $\|\sigma^{\theta}_{h^\tau} - \sigma^{\theta'}_{h^\tau}\| < \varepsilon$ provided $P(\theta|h) > 0$ and $P(\theta'|h) > 0$.*

The important thing to note is that, if a strategy is asymptotically constant, then any two player-types receiving positive probability after some (infinite) history, must play the same strategy after that history.

Theorem 3. *Let σ be a PBE of the game $G(\mu, \delta)$. Then, σ is asymptotically constant.*

This result is used to prove Theorem 4 below:

Theorem 4. *Let $\mu \in \Delta(\Theta)$. If for any smooth, strict subset $W \subset V$ there exists $\bar{\delta} \in (0, 1)$ such that for all $\delta > \bar{\delta}$ $W \subset E^{PBE}(\mu, \delta)$ then condition B holds.*

5 Examples

5.1 A public good partnership game.

The following is an example where condition $A.1$ fails and thus the folk theorem fails. We also show which payoffs can, and cannot, be sustained in BFE equilibrium.

For motivation, imagine a two player partnership game, with action sets $\{Work(W), Shirk(S)\}$. Working has a cost $c \in (0,1)$ for player 2, and this is common knowledge. The cost of working for player 1 is $\theta \in \{c, 2\} = \Theta$. The per capita gross payoffs of the project are 2 if both work, 1 if only one person works and 0 if no player works. Since we assume that, when one player works and the other shirks, both players get the gross payoff 1, we interpret this as a partnership that produces a public good. The payoff matrices are as follows:

c	W	S
W	$(2-c, 2-c)$	$(1-c, 1)$
S	$(1, 1-c)$	$(0,0)$

and

2	W	S
W	$(0, 2-c)$	$(-1, 1)$
S	$(1, 1-c)$	$(0,0)$

It is easy to see that part of condition $A.1$ is satisfied. If $\theta^+ = c$ then action (W, W) satisfies condition $A.1$: it yields $\hat{w}_1(c) = 2 - c$ and $\underline{w}_1(2) = 0$. Yet, condition $A.1$ fails when we set $\theta^+ = 2$: action profile (S, W) yields the best possible payoff for player 1 in state 2, but is still individually rational in state c. Thus, by virtue of theorems 1 and 4, the folk theorem (either for BFE or PBE) fails. Intuitively, the reason why condition $A.1$ fails is that payoffs are not "state contingent enough". Indeed, the payoff from working is state contingent because different states are equivalent to different working costs; in contrast, the payoff from shirking only depends on whether my rival worked or not, but not on the state. We could think that adding a state-contingent benefit to shirking, as we do in the next example, might solve this issue. This is only true if the state-contingent benefit from shirking outweighs the non-contingent payoff obtained when player 1 shirks and player 2 works. Otherwise, the non-contingent aspect of utilities dominates and makes condition $A.1$ fail.

Nonetheless, using the techniques used to prove Theorem 1, we can calculate what payoffs can still be approximated by BFE payoffs.[6] For instance, if we want to approximate a vector $v \in V$ such that $v(\theta = 2) = (1, 1 - c)$ with BFE payoffs, we must have $v_1(\theta = c) \geq 1$. More generally, a necessary condition for a payoff to be approximately sustained in BFE payoffs is to have $v_1(c) \geq v_1(2)$. It is also simple to check this conclusion extends to PBE. Indeed, since $u_1(a|c) \geq u_1(a|2)$ for all $a \in A$, if type c mimics the behavior of type 2, his payoffs must be weakly larger. Thus, any PBE payoff vector must also satisfy $v_1(c) \geq v_1(2)$.[7]

5.2 A private good partnership game.

Now we modify the partnership game so that condition $A.1$ is satisfied but condition $A.2$ might, depending on the cost c, be violated. Assume that if one player works and the other shirks, only the working player gets the gross payoff of 1.[8] As before, player 1 has two types: the first type pays a cost of c if he works and his second type pays a cost of 2 if he works, but now the second type also has a utility benefit from shirking. He gets y if he shirks alone and x if both players shirk together. For simplicity

[6] See appendix $A.5$ for details

[7] A necessary and sufficient condition for $v \in V$ to be approximately sustained in equilibrium is that $v_1(c) \geq v_1(2)$ and that, for some $\alpha \in \Delta(A)$, both $v_1(c) \geq u_1(\alpha|c)$ and $v_1(2) \geq u_1(\alpha|2)$.

[8] Working still has a coordination feature, in the sense that per capita payoff for the working player increases if both work together.

of calculations assume $0 < x < y < 2 - c$, so that x is player 1's min max payoff.[9] The payoff matrices are as follows:

c	W	S
W	$(2-c, 2-c)$	$(1-c, 0)$
S	$(0, 1-c)$	$(0, 0)$

and

2	W	S
W	$(0, 2-c)$	$(-1, 0)$
S	$(y, 1-c)$	$(x, 0)$

Condition $A.1$ is now satisfied because, for player 1, shirking when his rival works no longer guarantees the non-contingent payoff. Thus, the driving force behind the desire for shirking is no longer uniform across states.

Nonetheless, for certain parameter values, condition $A.2$ might not hold: for almost all $c \in (0, 1)$, θ^+, and $J \subset N$, the clauses of condition $A.2$ are satisfied. The exception is when $J = \{2\}$ and $\theta^+ = c$. In this case, the mixture between (W, W) and (S, W) that yields payoff $1 - c$ to player 1, is $\{2\}$ − *undominated.* For condition $A.2$ to hold this mixture must not be individually rational for player 1 in state 2. After some algebra, we conclude that this holds if and only if $\frac{y}{2-c} < x$. In other words, this particular clause of condition $A.2$ holds if working costs are sufficiently small in comparison to the relative utility between shirking alone and shirking when the other player is shirking.

As in condition $A.1$, what drives this result is that, comparing $\theta = c$ to $\theta = 2$, actions (W, W) and (S, W) swap their place in the *ordinal* preference ranking with respect to the min max benchmark. Yet, contrary to condition $A.1$, the *magnitude* of the change is also important. If (S, W) becomes "too attractive" or, alternatively, (W, W) in state

[9] Yes, player 1 type 2 is rather despicable: he not only has huge working costs, he also rather enjoys the prospect of shirking if he knows somebody else is working... even though that work does not benefit him!

c is not attractive enough, then condition $A.2$ will still fail.

Indeed, note that J domination, when J is a proper subset of N, is equivalent to maximizing a player's payoff subject to his rival obtaining at least some benchmark payoff. In particular, the maximum player 2 can get in state c, while simultaneously leaving player 1 on his individual rationality constraint, is to mix (W, W) and (S, W) with a weight that depends on c. While player 1 strictly prefers working to shirking, higher values of c mean that working is less attractive; thus, a higher weight must be assigned to (W, W) in order to compensate the bad payoff from (S, W). Now, assume that this action (the one that maximized player 2 utility in state c while keeping player 1 at his $\min\max$ payoff) is played in state 2. Since the high-cost type prefers shirking to working, as opposed to the low-cost type, higher working costs mean a high weight on the *bad* action. Thus, in general, we expect that the high-cost type will not mimic the behavior of the low-cost player and in this sense the states are *separated*: we can analyze extreme payoffs in the low-cost state knowing that the high cost player does not want to mimic this action. The only case where the argument fails is the case where y, the payoff from (S, W), is very high. In this case, even though obtaining this payoff is very unlikely, the reward is large enough that the high-cost type wants to take the risk. Thus, mimicking is profitable and we can no longer analyze one state independently of the other. Hence, if y is large enough with respect to c, condition $A.2$ is violated.

In short, this example shows that, for a folk theorem to hold, it is important that the structure of the game changes from state to state: "good" actions must become "bad" actions and vice versa, but the relative magnitude in which these changes relate to each other is equally important.

5.3 Monopoly product-choice game

A monopolist produces a good at 0 marginal cost and a consumer, who has private information about his preferences, decides whether or not to buy. This good has different features, indexed by points in the set $X = [0, 1]$, referred to as the set of features, or locations. We first analyze the case where, each period, the monopolist is constrained to produce at a given location $l \in [0, 1]$, so his only choice variable is price. We interpret this as a model where changing locations has prohibitively high costs. Then we look at the case where, each period, the monopolist is allowed to change his location, so his new choice set is both price and location. For consistency with the finite action sets assumption, the set of locations available for the monopolist to produce is the subset $L = \{\frac{m}{M} : 0 \leq m \leq M\}$ for some large $M \in \mathbb{N}$. We also assume prices increase in penny amounts and are bounded above by a large positive integer P.[10] If the monopolist does not sell to the consumer in a given period, his profit in that period is 0.

There are two types of the consumer, $x, x' \in L$, where $x < x'$. If consumer type x (resp. x') purchases good with features y at a price p, his utility is $u_1(y, p|\ x) = w - |x - y| - p$ (resp. $u_1(y, p|x') = w - |x' - y| - p$). The quantity w measures the utility of receiving the product with the features he values the most at a price of 0, and the term $|x - y|$ represents the disutility associated from consuming a good that does not have his most preferred features. Finally, if the consumer does not purchase the good at a given period, his utility for that period is 0. This implies that the min max payoffs of all player-types is 0.

5.3.1 Fixed location for the monopolist

Assume that the monopolist produces at location $l \in L$ where $x < l < x'$, $w > |x - l|$, and $w > |x' - l|$. Thus, if free, both consumer types would like to buy the specific good provided by the monopolist. Condition $A.1$ fails: indeed, if we set $\theta^+ = x$ then $\hat{w}_1(x) = w - |l - x|$. This utility can only be achieved if the monopolist sets price 0 and the consumer purchases, but this yields utility $w - |x' - l| > 0$ to type

[10] The assumption of penny increments will actually be irrelevant since utilities will be linear in prices and we are allowing the monopolist to play mixed strategies; i.e. to choose prices randomly.

x'. Condition $A.2$ is violated for the cases $J = \{1\}$ and $J = \{1, 2\}$. If $J = \{1\}$, the only undominated payoff profiles at a state $\theta \in \{x, x'\}$ are those where consumer type θ obtains his maximum utility, thus condition $A.2$ is violated for the same reason as is $A.1$. Similarly, if $J = \{1, 2\}$, the only undominated payoff profiles at a state $\theta \in \{x, x'\}$ are those of the form $(v_1, v_2) = (w - |\theta - l| - p, p)$, where $p \leq w - |\theta - l|$. Indeed, the only way to improve the consumer's (resp. the monopolist's) welfare is to reduce (resp. increase) the price, but this harms the monopolist (resp. consumer). Moreover, these are the only $\{1, 2\}$-undominated payoffs because any other payoff implies a probability that purchase does not take place, but these can be dominated by having a purchase happen. Thus, if p is small enough, condition $A.2$ will fail for the same reason as condition $A.1$ did. On the other hand, if $J = \{2\}$, condition $A.2$ will hold. Indeed, the only $\{2\}$-undominated profiles at a state $\theta \in \{x, x'\}$ are of the form $(0, |\theta - l|)$, achieved by setting $p = w - |\theta - l|$ and having the consumer make a purchase. We now show that condition $A.2$ is satisfied when $J = \{2\}$ and when $J = \emptyset$. Consider the action profile where the consumer buys the product and price is the maximum possible price P. The utility profile for playing this action in every state is $v = (w - |x - l| - P, P, w - |x' - l| - P, P)$. Thus, if P is high enough, this action profile satisfies condition $A.2$ for the case $J = \{2\}$ and for any choice of θ^+. For $J = \emptyset$ condition $A.2$ holds since any action profile where the consumer does not purchase yields the $\min\max$ payoff for all player-types.

Knowing that a folk theorem will not hold, the next step is to ask what payoffs can be obtained in PBE or BFE. For the latter, we apply the conditions in the appendix and conclude the following: we have one degree of freedom in specifying what to sustain in any state, but this imposes constraints on what we can sustain in the other state. Once we specify a payoff for a consumer type $\theta \in \{x, x'\}$, if that payoff is generated by a price p and a probability of purchase q, then we must give the other type of the consumer a payoff no less than what he can achieve is he purchases at price p with probability q. Concretely, consider any feasible, individually rational payoff vector $v = (v_1(x), v_2(x), v_1(x'), v_2(x'))$ and fix $\theta \in \{x, x'\}$. Then, there are $q \in [0, 1]$ and $p \in [0, P]$ such that $v_1(\theta) = q(w -$

$|\theta - l| - p)$, and v can be approximated in BFE if and only if $v_1(\theta') \geq q(w - |\theta' - l| - p)$ for $\theta' \neq \theta$. For the monopolist there are no bounds on what we can approximate state-wise. Finally, note that we can rewrite the above constraint on what the consumer might obtain as $v_1(\theta') \geq v_1(\theta) + q(|\theta - l| - |\theta' - l|)$. We now show that this constraint also applies to PBE; hence the set of payoffs sustainable in PBE and BFE coincide. Indeed, assume v is in the interior of V and that σ is a (δ, μ) PBE strategy profile that sustains v for some (δ, μ).[11] Simple calculation shows that, by mimicking the behavior of type θ, consumer type θ' can guarantee himself a utility level $v_1(\theta) + q(|\theta - l| - |\theta' - l|)$. Indeed, let σ be such a strategy. Then $\sigma_1^\theta : \mathcal{H} \to [0, 1]$ indicates, for type θ and after each history, the probability of purchase. Similarly, $\sigma_2 : \mathcal{H} \to \Delta\{0, ..., P\}$ indicates the pricing strategy of the monopolist. The utility of consumer type θ from following his prescribed strategy is $E_h\{\sum(1-\delta)\delta^t[\sigma_1^\theta(h^t)(w - |\theta - l| - \sigma_2(h^t))]\} = v_1(\theta)$. The utility of type θ' from mimicking type θ is $E_h\{\sum(1 - \delta)\delta^t[\sigma_1^\theta(h^t)(w - |\theta' - l| - \sigma_2(h^t))]\} \leq v_1(\theta')$; where the inequality follows form the "no profitable deviation" condition. Let $q = E_h\{\sum(1 - \delta)\delta^t\sigma_1^\theta(h^t)\}$. Rearranging terms shows that the constraint $v_1(\theta') \geq v_1(\theta) + q(|\theta - l| - |\theta' - l|)$ must hold. Thus, the above bound for BFE sustainability also applies to PBE. In particular, how far we are from a folk theorem, measured by the constraint on what payoffs we can sustain, is a function of how different utility functions are.

5.3.2 Endogenous location

Assume now that the monopolist can, at each period, also choose the location at which he will sell. Then, a PBE or BFE folk theorem holds if and only if $w \leq x' - x$. That is, the conditions will be satisfied only when, even if free, one type of the consumer is not willing to purchase the good corresponding to the other type. When the above inequality is not satisfied, the conditions fail for the same reason as in the fixed-location

[11] By considering limits we can also extend this result to payoffs in the boundary of V, but this adds unnecessary complications concerning limits.

model. Condition $A.1$, under $w \leq x' - x$, is satisfied because the best that consumer type x (resp. x') can hope for is to buy the product that the monopolist produces at location x (resp. x') at price 0. If $w \leq x' - x$, this action profile is not individually rational for type x'(resp. x) and thus condition $A.1$ holds. Condition $A.2$ holds for similar reasons.

As mentioned in the introduction, the main difference between the fixed-location model and the endogenous-location model is the effect that private information has on utilities. Indeed, returning to the fixed-location model, assume (without loss of generality) that $|x - l| < |x' - l|$. Then, we can rewrite the model in terms of a "high valuation consumer" (whose valuation is $w - |x - l|$) and a "low valuation consumer" whose valuation is $w - |x' - l|$. In this model, private information acts by shifting utility upward by the scalar amount $|x' - l| - |x - l|$. In contrast, the model where location is not fixed cannot be written in this way, and the private information of the consumer alters his ordinal preferences rather than simply shifting utility in a rigid manner.

6 Concluding Remarks

The existing literature on repeated games under incomplete information has looked at what payoffs can be approximated when different solution concepts are assumed. Of these, the most prominently studied are Nash equilibrium, generally with one-sided incomplete information, and Belief Free equilibrium. For Perfect Bayesian Equilibrium, which is arguably the most widely used notion in applied theory, there are (to the best of our knowledge) no known results. The difficulty in analyzing Perfect Bayesian Equilibrium is that belief updating, both on and off the path of play, coupled with sequential rationality, adds a non-stationarity to the model that cannot be addressed with the existing techniques.

This paper considers the class of incomplete-information games that satisfy full dimensionality, known-own-payoffs, perfect monitoring and recall, and where the state of the world can be identified by pooling the information of all informed players together. For this class of games we provide a necessary and sufficient condition for obtaining a Perfect Bayesian Equilibrium folk theorem. Moreover, we show that this condition is necessary and sufficient for obtaining a Belief-Free Equilibrium folk theorem. We first show that the condition implies a Belief-Free folk theorem and then we show that if a Perfect Bayesian Folk theorem holds so must the condition. The (trivial) observation that a Belief Free folk theorem implies Perfect Bayesian folk theorem concludes the argument.

Finally, two questions are left open for future research. First, when our folk theorem fails, using the techniques from Fudenberg-Yamamoto we derive necessary and sufficient conditions for a specific payoff vector to be approximately sustainable in Belief Free Equilibrium as players become more patient. Since the argument used to prove Theorem 4 only relies on local approximations, it is natural to conjecture that an analogue to Theorem 4 will hold: If a payoff $v \in V$ can be approximated in Perfect Bayesian Equilibrium, then the condition for that payoff v to be approximated in Belief-Free equilibrium must also hold. Thus, a payoff can be sustained in Perfect Bayesian Equilibrium if and only if it can be sustained in Belief-Free equilibrium. For the examples provided in Section 5 we see that this conjecture is true but a formal argument is yet to be obtained. Second, the condition for obtaining a folk theorem asks for the existence of very specific action profiles. It is natural to conjecture that these will play an important role in constructing the equilibrium strategies. Having understood what payoffs can be sustained in a given game, understanding the structure of the behaviors that generate such payoffs would give us a more complete view of the predictions made by the model.

A Appendix

A.1 FY (2010) further details

The purpose of this appendix is to provide further details on the techniques employed in FY (2010) since most proofs build upon them. Following FY (2010), the set of payoffs that can be sustained in a PTXE is related to the linear programs below. Moreover, it is proved in FY that the solution to the following program is independent of δ.

Linear Program.

$$k^*(\overline{\alpha},\lambda) = \max_{\substack{v\in\mathbb{R}^{N\times|\Theta|}\\ w:A\to\mathbb{R}^{N\times|\Theta|}}} \lambda\cdot v$$

$$s.t. \begin{cases} v_n = (1-\delta)u_n(\alpha^{\theta_n}|\theta_n) + \delta w_n(\alpha^{\theta_n}) \textit{ for all } n,\ \theta \\ v_n \geq (1-\delta)u_n(a_n,\alpha_{-n}^{\theta_{-n}}|\theta_n) + \delta w_n(a_n,\alpha_{-n}^{\theta_{-n}}|\theta_n) \textit{ for all } i,\ \theta \textit{ and } a_i\in A_i \\ \lambda\cdot v \geq \lambda\cdot w(a) \textit{ for all } a \end{cases}$$

where $\overline{\alpha} \equiv (\alpha^\theta)_{\theta\in\Theta} \in \Delta(A)^\Theta$ is a type contingent mixed action measurable with respect to the player's information. It is clear that there is no loss of generality in assuming that the directions λ are picked from the set $\Lambda \equiv \{\lambda \in \mathbb{R}^{N\times|\Theta|} : \|\lambda\| \leq 1\}$. Henceforth, we make this normalization. Letting $k^*(\lambda) \equiv \sup_{\overline{\alpha}} k^*(\lambda,\overline{\alpha})$ we construct the half-spaces $K(\lambda) \equiv \{v \in \mathbb{R}^{N\times|\Theta|} : \lambda\cdot v \leq k^*(\lambda)\}$, and the corresponding set supported by them $Q \equiv \cap_\lambda K(\lambda)$. Letting $E^{PTXE}(\delta)$ be the set of equilibrium payoffs that can be achieved in a PTXE when the discount factor is δ, we can state Proposition 1 from FY (2010):

FY(2010) Proposition 1. *If* $\dim Q = N\times|\Theta|$ *then* $\lim_{\delta\to1} E^{PTXE}(\delta) = Q$

An important remark regarding this proposition is that $\lim_{\delta\to1} E^{PTXE}(\delta) \subset Q$ always holds, regardless of the dimensionality condition. This will be useful in proving our characterization theorem.

Finally, Proposition 6 from FY characterizes the value $k^*(\lambda)$ for each λ. To state this Proposition, we first define an appropriate partition of theset Λ. Note that we state the partition in terms of our information model. The generalization to an arbitrary information model can be found in FY.

- $\Lambda^1 \equiv \{\lambda \in \Lambda : \exists\theta$ such that $(\lambda_n^\theta)_n \neq 0$ and $(\forall\theta' \neq \theta)(\lambda_n^{\theta'})_n = 0\}$
- $\Lambda^2 \equiv \{\lambda \in \Lambda : (\exists n, n' \in N)(i \in I)(\exists\theta, \theta')(\theta_i \neq \theta_i)$ such that $\lambda_n^\theta \neq 0$ and $\lambda_{n'}^{\theta'} \neq 0\}$
- $\Lambda^3 \equiv \{\lambda \in \Lambda : (\exists i \in I, n \in N)(\exists\theta, \theta' \in \Theta)(\theta_i \neq \theta_i')$ such that $\lambda_i^\theta > 0, \lambda_n^{\theta'} \neq 0\}$
- $\Lambda^4 \equiv \{\lambda \in \Lambda : (\exists i \in I)(\theta', \theta'')(\theta_i' \neq \theta_i'')$ such that $\lambda_i^{\theta'}, \lambda_i^{\theta''} > 0$ and $(\lambda_n^\theta)_{\theta\in\Theta} = 0$ for all $n \neq i\}$
- For $i \in I$ $\Lambda^5(i) \equiv \{\lambda \in \Lambda : (\lambda_i^\theta)_{\theta\in\Theta} \leq 0, (\lambda_i^\theta)_{\theta\in\Theta} \neq 0(\lambda_n^\theta)_{\theta\in\theta} = 0$ for all $n \neq i$ and $(\forall n \neq i)(\theta, \theta' \in \Theta)(\lambda_i^\theta \neq 0$ and $\lambda_i^{\theta'} \neq 0$ implies $\theta_n = \theta_n')\}$
- $\Lambda^5 \equiv \cup_{i\in I}\Lambda^5(i)$
- $\Lambda^6 \equiv \{\lambda \in \Lambda : (\exists i \in I)(\exists\theta' \in \Theta)$ such that $\lambda_i^{\theta'} \neq 0(\lambda_n^\theta)_{\theta\in\Theta} = 0$ for all $n \neq i(\lambda_i^{\theta''} > 0, \lambda_i^{\theta'''} > 0$ imply $\theta_i'' = \theta_i''')$ and $(\lambda_i^{\theta''} \neq 0, \lambda_i^{\theta'''} \neq 0$ imply $\theta_n'' = \theta_n''')$ for all $n \neq i\}$
- $\Lambda^7 \equiv \{\lambda \in \Lambda : (\exists n', n'' \in N)(\theta', \theta'' \in \Theta)$ such that $((\lambda_{n'}^\theta)_{\theta\in\Theta} \neq 0, (\lambda_{n''}^\theta)_{\theta\in\Theta} \neq 0)((\lambda_n^{\theta'})_{n\in N}, (\lambda_n^{\theta''})_{n\in N} \neq 0)(\lambda_l^{\theta'''} \neq 0, \lambda_{l'}^{\theta''''} \neq 0, n \neq l, n \neq l', \theta''' \neq \theta''''$ implies $\theta_n''' = \theta_n'''')(\lambda_n^{\theta'''} > 0, \lambda_l^{\theta''''} \neq 0, n \neq l, \theta''' \neq \theta''''$ implies $\theta_n''' = \theta_n''''))\}$
- $\Lambda^0 = \cup_{j\in\{1,\ldots,4\}}\Lambda^j$

It is simple to check that $\cup_{t=1,\ldots,7}\Lambda^t = \Lambda$.

FY(2010) Proposition 7.

$$k^*(\lambda) = \begin{cases} \max\{\lambda \cdot v : v \in V\} & \textit{if } \lambda \in \Lambda^1 \\ \infty & \textit{if } \lambda \in \Lambda^2 \cup \Lambda^3 \cup \Lambda^4 \\ \max_{\alpha_{-i}} \min_{\alpha_i} \sum_{\theta \in \Theta} \lambda_i^\theta u_i(\alpha_i, \alpha_{-i}|\theta) & \textit{if } \lambda \in \Lambda^5(i) \\ \max_\alpha \{\lambda \cdot u(\alpha)\} & \textit{if } \lambda \in \Lambda^6 \cup \Lambda^7 \end{cases}$$

Notice the score is always large enough in directions Λ^0, in the sense that $\lambda \cdot v \geq \max\{\lambda \cdot v : v \in V\}$ for all $\lambda \in \Lambda^0$. Hence, our interests lie in the remainder $\Lambda \setminus \Lambda^0$ directions. Finally, for the case of one sided incomplete information, we will need FY proposition 8. Let $V^U = co\{u(a|\cdot) : a \in A\}$ and $V^{*U} = \{v \in V^U : (\forall n)(\exists \alpha_{-n})(\forall \theta) v_n(\theta) \geq \max_{a_n} u_n(a_n, \alpha_{-n}|\theta)\}$.

FY(2010) Proposition 8. *If* V^{*U} *has dimension* $|\Theta| + N - 1$ *then* $\dim Q = N$

These propositions, adapted to our information model, are the building blocks we need to construct the proofs of our main theorems. One can also generalize the information model and prove the following: if the joint information of all players does not identify the state, then achieving a folk theorem is impossible. Indeed, consider a generalized information model where states are elements ω of a set Ω. Each player n is endowed with a partition, possibly a trivial one if the player is uninformed, of Ω, denoted π_n. Let $\theta_n(\omega) \in \pi_n$ denote all states that belong to the same cell of π_n as ω. These represent all states that, when the true state is ω, cannot be distinguished from ω. Alternatively, states ω, ω' cannot be distinguished by a player n if and only if $\theta_n(\omega) = \theta_n(\omega')$. Our information model is the special case where Ω is a subset of an N-dimensional space, the partitions π_n are sections along the n^{th} dimension and the fact that player n cannot distinguish ω from ω' is the equality $\omega_n = \omega'_n$. The known-own-payoffs condition is simply that $\theta_n(\omega') = \theta_n(\omega)$ implies $u_n(\cdot|\omega) = u_n(\cdot|\omega')$. We can now prove that, under the known own payoffs condition, if a game admits a folk theorem then, for every $\omega \in \Omega$, $\{\omega\} = \cap_{n\in N}\theta_n(\omega)$. Proceed by contradiction. Assume there is a state $\omega \in \Omega$ such that $\{\omega, \omega'\} \subset \cap_n \theta_n(\omega)$, where $\omega' \neq \omega$. Then, there exists some informed player $i \in I$, such that $\{\omega, \omega'\} \subset \theta_i(\omega) \cap_{j\neq i} \theta_j(\omega)$. Pick $\lambda \in \mathbb{R}^{|\Omega|\times N}$ such that $\lambda_i(\omega) = p > 0 > q_i = \lambda_i(\omega')$, $\lambda_j(\omega') = q_j$ for some $j \neq i$ and zero otherwise. Then, $\lambda \in \Lambda^7$. By adjusting the relative sizes of p and q_i one can check that a necessary condition for $k^*(\lambda) \geq \max\{\lambda \cdot v : v \in V\}$ to hold for all p, q_i and q_j is that, for some $\alpha \in \Delta(A)$, $u_i(\alpha|\omega) \geq \hat{w}_i(\omega)$ and $u_i(\alpha|\omega') \leq \underline{w}_i(\omega')$. The known-own-payoffs condition implies that $u_i(\alpha|\omega) = u_i(\alpha|\omega')$ and $\underline{w}_i(\omega) = \underline{w}_i(\omega')$. Thus $\hat{w}_i(\omega) \leq \underline{w}_i(\omega)$. This a contradiction with the full dimensionality of $V(\omega)$.

A.2 Proof of Theorem 1

The proof of theorem 1 will build upon two lemmas. The first, shows that condition $A.1$ is equivalent to having a high score in directions $\Lambda^6 \cup \Lambda^7$. The second, shows that condition $A.2$ implies a high score in direction Λ^5.

Lemma A.1. *Condition A holds if and only if, for every $\lambda \in \Lambda^6 \cup \Lambda^7$, $k^*(\lambda) \geq \max_{v\in V} \lambda \cdot v$*

Proof. ($\Rightarrow$) Take $\lambda \in \Lambda^7$. In particular, there is $\theta^* \in \Theta$ such that (a) if $\lambda_n^\theta > 0$ for some $n \in N$ and $\theta \in \Theta$ then $\theta = \theta^*$ and (b) if $\lambda_n^\theta \neq 0$ for some $n \neq 1$ and $\theta \in \Theta$ then $\theta = \theta^*$. Let $J = \{n : \lambda_n^{\theta^*} > 0\}$. Pick α_J as in condition $A.2$. Note that the problem $\max_{v\in V} \lambda \cdot v$ will admit a solution v^* such that $v^*(\theta) = \underline{w}_1(\theta)$ for all $\theta \neq \theta^*$ and $(v_n^*(\theta^*))_n$ will be J-undominated. Let v^* be such a solution. Then:

$$
\begin{aligned}
k^*(\lambda) &= \max\left\{\sum_{\theta\neq\theta^*} \lambda_1(\theta) u_1(\alpha'|\theta) + \sum_{n\in J} \lambda_n(\theta^*) u_n(\alpha'|\theta^*) + \sum_{n\notin J} \lambda_n(\theta^*) u_n(\alpha'|\theta^*) : \alpha' \in \Delta(A)\right\} \\
&\geq \sum_{\theta\neq\theta^*} \lambda_1(\theta) \underline{w}_1(\theta) + \sum_{n\in J} \lambda_n(\theta^*) v_n(\theta^*) + \sum_{n\notin J} \lambda_n(\theta^*) v_n(\theta^*) \\
&= \max\{\lambda \cdot v : v \in V\}
\end{aligned}
$$

Hence, for all $\lambda \in \Lambda^7$ we get $k^*(\lambda) \geq \max\{\lambda \cdot v : v \in V^*\}$

Now take $\lambda \in \Lambda^6$. That is, $\lambda_n = 0$, for $n \neq 1$, $\lambda_1 \neq 0$, and there are θ, θ' such that $\lambda_1(\theta) > 0 > \lambda_1(\theta')$.[12] Let $\theta^+ = \theta$. Pick $\alpha_{\theta+}$ as in condition $A.1$. Then:

$$\begin{aligned}
k^*(\lambda) &= \max\{\sum_{\theta'' \in \Theta} \lambda_1(\theta'') u_1(\alpha|\theta'') : \alpha \in \Delta(A)\} \\
&\geq \lambda_1((\theta^+))\hat{w}_1(\theta^+) + \sum_{\theta'' \neq \theta^+} \lambda_1((\theta''))\underline{\mathrm{w}}_1(\theta'') \\
&= \max\{\lambda \cdot v : v \in V\}
\end{aligned}$$

Hence, for all $\lambda \in \Lambda^6$ we get $k^*(\lambda) \geq \max\{\lambda \cdot v : v \in V^*\}$. Hence, we get $k^*(\lambda) \geq \max\{\lambda \cdot v : v \in V\}$ for all $\lambda \in \Lambda^6 \cup \Lambda^7$

($\Leftarrow$) First, note that for any $v \in V$ the function $g_v : \Lambda \times \Delta(A) \to \mathbb{R}$ defined by $g(\lambda, \alpha) = \sum_i \sum_\theta \lambda_i(\theta)[u_i(\alpha|\theta) - v_i(\theta)]$ is linear in each component and almost periodic over $\Lambda' \times \Delta(A)$ for any convex $\Lambda' \subset \Lambda$. Thus, by Fan's minimax theorem, $\inf_{\lambda \in \Lambda'} \max_\alpha \{g(\lambda, \alpha)\} = \max_\alpha \inf_{\lambda \in \Lambda'}\{g(\lambda, \alpha)\}$ for any convex $\Lambda' \subset \Lambda$. We now use this to show condition A must hold. Pick $\theta^+ \in \Theta$ arbitrarily and $v \in V$ such that $v_1(\theta^+) = \hat{w}_1(\theta^+)$, $v_1(\theta) = \underline{\mathrm{w}}_1(\theta)$ for $\theta \neq \theta^+$. Let $\lambda \in \Lambda^6$ be such that $(\lambda_n)_{n \neq 1} = 0$, $\lambda_1(\theta^+) = p > 0$, $\lambda_1(\theta) = q_\theta < 0$ for all $\theta \neq \theta^+$. Then, for any $\alpha \in \Delta(A)$:

$$\begin{aligned}
\sum_\theta \lambda_1(\theta)[u_1(\alpha|\theta) - v(\theta)] &= p[u_1(\alpha|\theta^+) - v_1(\theta^+)] + \sum_{\theta \neq \theta^+} q_\theta [u_1(\alpha|\theta) - v_1(\theta)] \\
\max_{\alpha \in \Delta(A)} \sum_\theta \lambda_1(\theta)[u_1(\alpha|\theta) - v_1(\theta)] &= k^*(\lambda(p,q)) - \lambda(p,q) \cdot v \geq 0 \\
\inf_{p,q}\{k^*(\lambda(p,q)) - \lambda(p,q) \cdot v\} &= \max_{\alpha \in \Delta(A)} \inf_{(p,q)} \sum_\theta \lambda_1(\theta)[u_1(\alpha|\theta) - v(\theta)] \geq 0
\end{aligned}$$

where the second line follows from the characterization of k^* and the fact that $k^*(\lambda) - \max_{v \in V} \lambda \cdot v \geq 0$ and the last line follows from Fan's min max theorem. Thus, there must be an α^* such that $\sum_\theta \lambda_1(p,q)(\theta)[u_1(\alpha^*|\theta) - v(\theta)] \geq 0$ for all p, q. Letting $p \to \infty$,

[12] If there were no negative-weight states, this woud be a direction in Λ^1.

$q_\theta \to 0$ for all θ we conclude $u_1(\alpha^*|\theta^+) - \hat{w}_1(\theta^+) \geq 0$ and letting $q_\theta \to -\infty$ for some θ, $q_{\theta'} \to 0$ for all $\theta \neq \theta'$ and $p \to 0$ we conclude $u_1(\alpha^*|\theta^+) - \underline{w}_1(\theta^+) \leq 0$. Let $\alpha_{\theta^+} \equiv \alpha^*$. This action satisfies condition $A.1$.

Let $\theta^* \in \Theta$, $J \subset N$ and $v(\theta^*)$ a J-undominated payoff vector. Extend this vector to $\mathbb{R}^{N\times\Theta}$ by setting $v_1(\theta) = \underline{w}_1(\theta)$ for $\theta \neq \theta^*$ and arbitrarily for the other players. Choose $\lambda \in \Lambda^7$ such that $\lambda_n(\theta^*) = p(n) > 0$ for $n \in J$, $\lambda_n(\theta^*) = q(n) < 0$ for $n \notin J$, $\lambda_n(\theta) = 0$ for all $n \neq 1$ and $\theta \neq \theta^*$. Furthermore, let $\lambda_1(\theta) = q(1,\theta) < 0$ for all $\theta \neq \theta^*$. Following the same logic as before:

$$\sum_\theta \sum_n \lambda_n(\theta)[u_n(\alpha|\theta) - v_n(\theta)] = \sum_{n\in J} p(n)[u_n(\alpha) - v_n(\theta^*)] + \sum_{n\notin J} q(n)[u_n(\alpha) - v_n(\theta^*)] + \sum_{\theta\neq\theta^*} q(1,\theta)[u_1(\alpha|\theta) - v_1(\theta)]$$

$$k^*(\lambda(p,q)) - \lambda(p,q)\cdot v \geq 0$$

$$\inf_{p,q}\{k^*(\lambda(p,q)) - \lambda(p,q)\cdot v\} = \max_{\alpha\in\Delta(A)} \inf_{(p,q)} \sum_\theta \sum_n \lambda_n(\theta)[u_n(\alpha|\theta) - v(\theta)] \geq 0$$

Thus, there must be α^* such that for all (p,q), $\sum_n \lambda_n(\theta)[u_n(\alpha^*|\theta) - v(\theta)] \geq 0$. By adjusting the sizes of p, q in turns we obtain $[u_n(\alpha^*) - v_n(\theta^*)] \geq 0$ for $n \in J$, $[u_n(\alpha^*) - v_n(\theta^*)] \leq 0$ for $n \notin J$ and $[u_1(\alpha^*|\theta) - \underline{w}_1(\theta)] \leq 0$ for all $\theta \neq \theta^*$. □

Lemma A.2. *If condition A.2 holds, then $k^*(\lambda) \geq \lambda \cdot v$ for all $v \in V$ and all $\lambda \in \Lambda^5$*

Proof. Let condition $A.2$ hold, and pick $J = \emptyset$. Thus, there must exists α^* such that, for all $\theta \in \Theta$ and all $n \in N$, $u_n(\alpha^*|\theta) \leq \underline{w}_n(\theta)$. Then:

$$\begin{aligned} k^*(\lambda) &= \max_{\alpha_2} \min_{\alpha_1} \{\sum_\theta \lambda_1(\theta) u_1(\alpha_1, \alpha_2|\theta)\} \\ &= -\min_{\alpha_2} \max_{\alpha_1} \{\sum_\theta |\lambda_1(\theta)| u_1(\alpha_1, \alpha_2|\theta)\} \\ &= -\sum_\theta |\lambda_1(\theta)| \min_{\alpha_2} \max_{\alpha_1} \{u_1(\alpha_1, \alpha_2|\theta)\} \\ &= \max\{\lambda \cdot v : v \in V\} \end{aligned}$$

where the next to last equality follows from the existence of the uniform min max action.

□

Armed with these lemmas, we can easily prove the main theorem in this section. Before proceeding recall proposition 8 from FY.

FY(2010) Proposition 8. *If V^{*U} has dimension $|\Theta| + N - 1$ then* $\dim Q = N$

In condition $A.2$, set θ^* arbitrarily and $J = \emptyset$. Then, condition A implies the existence of a uniform min max action. This, and the fact that each $V(\theta)$ is fully dimensional, implies V^{*U} is fully dimensional.

Theorem 1. *Condition A holds if and only if for any strict, smooth subset $W \subset V$ there exists $\bar{\delta} \in (0,1)$ such that, for any $\delta \geq \bar{\delta}$ $W \subset E^{PTXE}(\delta)$*

Proof. $\Rightarrow$ If condition A holds then V^{*U} has dimension $|\Theta| + N - 1$ and thus $\dim Q = N$. Moreover, for all $\lambda \in \Lambda$, $k^*(\lambda) \geq \max_{v \in V} \lambda \cdot v$. Hence, by virtue of FY proposition 1, for any strict, smooth, compact subset $W \subset V$ there exists $\delta \in (0,1)$ such that $W \subset E^{PTXE}(\delta)$.

$\Leftarrow$ Assume that for any strict, smooth subset $W \subset V$ there exists $\delta \in (0,1)$ such that $W \subset E^{PTXE}(\delta)$. Denote with V° the relative interior of V. For any $v \in V^\circ$ let $\bar{B} \subset V$ be a closed ball containing v. B is a smooth strict subset of V. Thus, $v \in \bar{B} \subset \lim_{\delta \to 1} E^{PTXE}(\delta) \subset Q$ by assumption. Since this holds for any $v \in V^\circ$, it follows that $V^\circ \subset Q$ and thus $\sup_{v\in V} v\cdot\lambda \leq k^*(\lambda)$. By virtue of lemma A, condition A must then hold. □

A.3 Proof of Theorem 2

Lemma $B.1$ below is analogous to Lemma $A.1$ and shows that condition B is equivalent to a high enough score in directions $\Lambda^6 \cup \Lambda^7$.

Lemma B.1. *Condition B holds if and only if, for every* $\lambda \in \Lambda^6 \cup \Lambda^7$, $k^*(\lambda) \geq \max_{v\in V} \lambda \cdot v$

Proof. ($\Rightarrow$) Take $\lambda \in \Lambda^7$. Then, there is at most one state $\theta^+ \in \Theta$ in which players are allowed positive weights. Let $J = \{n : \lambda_n(\theta^+) > 0\}$.[13] Pick α_J as in condition $B.2$. Note that the problem $\max_{v\in V}\{\lambda\cdot v\}$ will admit a solution v^* such that $v_n^*(\theta) = \underline{\mathrm{w}}_n(\theta_n)$ for all $\theta_n \neq \theta_n^*$ and $(v_n^*(\theta^+))_n$ will be J-undominated. Let v^* be such a solution. Then:

$$\begin{aligned}
k^*(\lambda) =& \max\{\sum_{n\in J} \lambda_n(\theta^+)u_n(\alpha'|\theta^+) + \sum_{n\notin J} \lambda_n(\theta^+)u_n(\alpha'|\theta^+) + \sum_{i\in I}\sum_{\theta_i\neq\theta_i^+} \lambda(\theta_i,\theta_{-i}^+)u_i(\alpha'|\theta_i) : \alpha' \in \Delta(A)\} \\
\geq& \sum_i \sum_{\theta_i\neq\theta_i^*} \lambda_i(\theta_i,\theta_{-i}^*)\underline{\mathrm{w}}_i(\theta_i) + \sum_{n\in J}\lambda_n(\theta^*)v_n^*(\theta^*) + \sum_{n\notin J}\lambda_n(\theta^*)v_n^*(\theta^*) \\
=& \max\{\lambda\cdot v : v\in V\}
\end{aligned}$$

Hence, for all $\lambda \in \Lambda^7$ we get
$k^*(\lambda) \geq \max\{\lambda\cdot v : v \in V^*\}$
Now take $\lambda \in \Lambda^6$. That is, there exists $i \in I$ such that $\lambda_i^\theta > 0 > \lambda_i^{\theta'}$ for some $\theta, \theta' \in \Theta$,

[13] Recall that possibly $\lambda \leq 0$ and thus $J = \emptyset$.

and $\lambda_n = 0$, for $n \neq i$. Let $\theta^+ = \theta$. Pick α as in condition $B.1$. Then:

$$\begin{aligned} k^*(\lambda) =& \max\{\sum_{\theta'' \in \Theta} \lambda_i((\theta''))u_i(\alpha'|\theta'') : \alpha' \in \Delta(A)\} \\ \geq& \lambda_i(\theta^+)\hat{w}_i(\theta_i^+) + \sum_{\theta'' \neq \theta^+} \lambda_i((\theta''))\underline{w}_i(\theta_i'') \\ =& \max\{\lambda \cdot v : v \in V\} \end{aligned}$$

Hence, for all $\lambda \in \Lambda^6$ we get $k^*(\lambda) \geq \max\{\lambda \cdot v : v \in V^*\}$ and thus, we get $k^*(\lambda) \geq \max\{\lambda \cdot v : v \in V\}$ for all $\lambda \in \Lambda^6 \cup \Lambda^7$

($\Leftarrow$) As before, note that for any $v \in V$ the function $g_v : \Lambda \times \Delta(A) \to \mathbb{R}$ defined by $g(\lambda, \alpha) = \sum_i \sum_\theta \lambda_i(\theta)[u_i(\alpha|\theta) - v_i(\theta)]$ is linear in each component and almost periodic over $\Lambda' \times \Delta(A)$ for any convex $\Lambda' \subset \Lambda$. Thus, by Fan's minimax theorem, $\inf_{\lambda \in \Lambda'} \max_\alpha \{g(\lambda, \alpha)\} = \max_\alpha \inf_{\lambda \in \Lambda'} \{g(\lambda, \alpha)\}$ for any $\Lambda' \subset \Lambda$. We now use this to show the two items of condition B must hold. Pick $\theta^+ \in \Theta$, $i \in I$ arbitrarily and $v \in V$ such that $v_i(\theta^+) = \hat{w}_i(\theta^+)$, $v_i(\theta_i, \theta_{-i}^+) = \underline{w}_i(\theta_i)$ for $\theta_i \neq \theta_i^+$. Let $\lambda \in \Lambda^6$ be such that $(\lambda_n)_{n \neq i} = 0$, $\lambda_i(\theta^+) = p > 0$, $\lambda_i(\theta_i, \theta_{-i}^+) = q_{\theta_i} < 0$ for all $\theta_i \neq \theta_i^+$. Then, for any $\alpha \in \Delta(A)$:

$$\begin{aligned} \sum_\theta \lambda_i(\theta)[u_i(\alpha|\theta) - v_i(\theta)] =& p[u_i(\alpha|\theta^+) - v_i(\theta^+)] + \sum_{\theta_i \neq \theta_i^+} q_{\theta_i}[u_i(\alpha|\theta_i) - v_i(\theta_i, \theta_{-i}^+)] \\ \max_{\alpha \in \Delta(A)} \sum_\theta \lambda_i(\theta)[u_i(\alpha|\theta) - v_i(\theta)] =& k^*(\lambda(p,q)) - \lambda(p,q) \cdot v \geq 0 \\ \inf_{p,q}\{k^*(\lambda(p,q)) - \lambda(p,q) \cdot v\} =& \max_{\alpha \in \Delta(A)} \inf_{(p,q)} \sum_\theta \lambda_i(\theta)[u_i(\alpha|\theta) - v_i(\theta)] \geq 0 \end{aligned}$$

where the second line follows from the characterization of k^* and the fact that $k^*(\lambda) - \max_{v \in V} \lambda \cdot v \geq 0$ and the last line follows from Fan's min max theorem. Thus, there must be an α^* such that $\sum_\theta \lambda_i(p,q)(\theta)[u_i(\alpha^*|\theta) - v_i(\theta)] \geq 0$ for all p, q. Letting $p \to \infty$, $q_{\theta_i} \to 0$ for all θ we conclude $u_i(\alpha^*|\theta^+) - \hat{w}_i(\theta^+) \geq 0$ and letting $q_{\theta_i} \to -\infty$ for some θ_i, $q_{\theta_i'} \to 0$ for all $\theta_i' \neq \theta_i$ and $p \to 0$ we conclude $u_i(\alpha^*|\theta_i) - \underline{w}_i(\theta_i) \leq 0$. Let $\alpha \equiv \alpha^*$ and obtain condition $B.1$.

Let $\theta^+ \in \Theta$, $J \subset N$ and $v(\theta^*)$ a J-undominated payoff vector. Extend this vector to $\mathbb{R}^{N\times|\Theta|}$ by setting $v_n(\theta) = \underline{w}_n(\theta)$ for $\theta \neq \theta^*$. Choose $\lambda \in \Lambda^7$ such that $\lambda_n(\theta^*) = p(n) > 0$ for $n \in J$, $\lambda_n(\theta^*) = q(n) < 0$ for $n \notin J$, $\lambda_n(\theta) = 0$ for all $n \notin I$ and $\theta \neq \theta^*$. Furthermore, let $\lambda_i(\theta_i, \theta^*_{-i}) = q(i, \theta_i) < 0$ for all $\theta_i \neq \theta^*_i$ and all $i \in I$. Following the same logic as before:

$$\begin{aligned}\sum_\theta \sum_n \lambda_n(\theta)[u_n(\alpha|\theta) - v_n(\theta)] =& \sum_{n\in J} p(n)[u_n(\alpha) - v_n(\theta^*)] + \sum_{n\notin J} q(n)[u_n(\alpha) - v_n(\theta^*)] \\ &+ \sum_{i\in I}\sum_{\theta_i\neq\theta^*_i} q(i,\theta_i)[u_i(\alpha|\theta_i,\theta^+_{-i}) - v_1(\theta_i,\theta^+_{-i})]\end{aligned}$$

$$k^*(\lambda(p,q)) - \lambda(p,q)\cdot v \geq 0$$

$$\inf_{p,q}\{k^*(\lambda(p,q)) - \lambda(p,q)\cdot v\} = \max_{\alpha\in\Delta(A)} \inf_{(p,q)} \sum_\theta \sum_n \lambda_n(\theta)[u_n(\alpha|\theta) - v(\theta)] \geq 0$$

Thus, there must be α^* such that for all (p,q), $\sum_n \lambda_n(\theta)[u_n(\alpha^*|\theta) - v(\theta)] \geq 0$. By adjusting the sizes of p, q in turns we obtain $[u_n(\alpha^*|\theta^*) - v_n(\theta^*)] \geq 0$ for $n \in J$, $[u_n(\alpha^*|\theta^*) - v_n(\theta^*)] \leq 0$ for $n \notin J$ and $[u_i(\alpha^*|\theta_i) - \underline{w}_i(\theta_i)] \leq 0$ for all $\theta_i \neq \theta^+_i$. □

From condition $B.2$ we get a similar result as lemma $A.2$. Notice that, in $\Lambda^5(i)$, if $\lambda^\theta_i < 0, \lambda^{\theta'}_i < 0$ then $\theta_{-i} = \theta'_{-i}$. Thus, directions in Λ^5 are those where there is a unique section $\theta_{-i} \in \Theta_{-i}$ where weights are non zero.

Lemma B.2. *If condition $B2$ holds then, for all $i \in I$, $k^*(\lambda) \geq \max_{v\in V} \lambda \cdot v$ for all $\lambda \in \Lambda^5(i)$.*

Proof. Let condition $B.2$ holds, and pick $J = \emptyset$. Thus, there must exist α^* such that, for all $\theta_i \in \Theta_i$ and all $i \in I$, $u_i(\alpha^*|\theta_i) \leq \underline{w}_i(\theta_i)$. Then, for all $\lambda \in \Lambda^5(i)$:

$$\begin{aligned} k^*(\lambda) &= \max_{\alpha_{-i}} \min_{\alpha_i} \{ \sum_{\theta} \lambda_i(\theta) u_i(\alpha_i, \alpha_{-i}|\theta) \} \\ &= - \min_{\alpha_{-i}} \max_{\alpha_i} \{ \sum_{\theta} |\lambda_i(\theta)| u_i(\alpha_i, \alpha_{-i}|\theta) \} \\ &= - \sum_{\theta} |\lambda_i(\theta)| \min_{\alpha_{-i}} \max_{\alpha_i} \{ u_i(\alpha_i, \alpha_{-i}|\theta) \} \\ &= \max\{ \lambda \cdot v : v \in V \} \end{aligned}$$

where the next to last equality follows from the existence of the uniform (across Θ_i) min max action. □

Theorem 2. *Condition B holds if and only if for any strict, smooth subset $W \subset V$ there exists $\bar{\delta} \in (0,1)$ such that, for any $\delta \geq \bar{\delta}$, $W \subset E^{PTXE}(\delta)$*

Proof. $\Rightarrow$ If condition B holds, by virtue of Lemma $B.1$ and $B.2$, $V \subset Q$. Since V is fully dimensional, so is Q. Hence, by virtue of FY proposition 1, for any strict, smooth, compact subset $W \subset V$ there exists $\delta \in (0,1)$ such that $W \subset E^{PTXE}(\delta)$.
$\Leftarrow$ Assume that for any strict, smooth subset $W \subset V$ there exists $\delta \in (0,1)$ such that $W \subset E^{PTXE}(\delta)$. Denote with V° the relative interior of V. For any $v \in V^\circ$ let $\bar{B} \subset V$ be a closed ball containing v. The set $\bar{B}$ is a smooth strict subset of V. Thus, $v \in \bar{B} \subset \lim_{\delta\to1} E^{PTXE}(\delta) \subset Q$ by assumption. Since this holds for any $v \in V^\circ$, it follows that $V^\circ \subset Q$ and thus $\sup_{v\in V} v \cdot \lambda \leq k^*(\lambda)$. By virtue of lemma $B.1$, condition B must then hold. □

A.4 Proof of Theorem 3 and Theorem 4

Theorem 3. *Let σ be a PBE of the game $G(\mu, \delta)$. Then, σ is asymptotically constant.*

Proof. Fix $\theta_n \in \Theta_n$, $j \neq n$, $\theta_j, \theta'_j \in \Theta_j$, and $h \in \mathcal{H}$ such that $P(\theta_j|h) > 0$, $P(\theta'_j|h) > 0$. Since $P(\theta_j|h) > 0$, $P(\theta'_j|h) > 0$ then $P(\theta_j|h^t) > 0$, $P(\theta'_j|h^t) > 0$ for large enough t. Fix

$\varepsilon > 0$. By Kalai-Lehrer (1993), there exists T such that, if $t > T$, then $\sigma|_{h_j^t}^{\theta_j}$ plays ε-like $\sum_{\theta_j''} P(\theta_j|h^t, \theta_i)\sigma|_{h_j^t}^{\theta_j''}$. That is, the realized strategy plays ε-like the prediction any other player-type makes. The same holds for $\sigma|_{h_j^t}^{\theta_j'}$. Thus, $\sigma|_{h_j^t}^{\theta_j}$ plays ε-like $\sigma|_{h_j^t}^{\theta_j'}$. In particular, for any $a_j \in A_j$ $|\sigma|_{h_j^t}^{\theta_j}(a_i) - \sigma|_{h_j^t}^{\theta_j'}(a_i)| \leq \varepsilon$. Taking $\varepsilon \to 0$ (so that $h^t \to h$) completes the proof. □

Theorem 4. *Let $\mu \in \Delta(\Theta)$. If for any smooth, strict subset $W \subset V$ there exists $\bar{\delta} \in (0, 1)$ such that for all $\delta > \bar{\delta}$ $W \subset E^{PBE}(\mu, \delta)$ then condition B must hold.*

The proof is a straightforward application of the following two lemmas. Since condition $B.1$ only deals with one "relevant" player, the argument is much simpler than for condition $B.2$. Thus the separation of the proof into two distinct lemmas.

Lemma 4.1. *Let $\mu \in \Delta(\Theta)$. If for any smooth, strict subset $W \subset V$ there exists $\bar{\delta} \in (0, 1)$ such that for all $\delta > \bar{\delta}$ $W \subset E^{PBE}(\mu, \delta)$ then condition $B.1$ must hold.*

Proof. Let $\theta^+ \in \Theta$, $i \in I$ and $K > 0$ be arbitrarily selected. Pick $v^* \in V$ such that $v_i^*(\theta_i^+, \theta_{-i}) = \hat{w}_i(\theta_i^+)$ for all $\theta_{-i} \in \Theta_{-i}$, $v_i^*(\theta_i, \theta_{-i}) = \underline{w}_i(\theta_i)$ for all $\theta_i \neq \theta_i^+$, $\theta_{-i} \in \Theta_{-i}$ and arbitrarily otherwise. Let $v \in V^\circ$ be such that $\|v - v^*\| \leq \frac{1}{K}$ and $v_i(\theta_i, \theta_{-i}) = v_i(\theta_i, \theta_{-i}')$ for all $\theta_{-i}, \theta_{-i}' \in \Theta_{-i}$. Then, v can be sustained as a (μ, δ) PBE for high enough δ. Let σ be the sustaining strategy.
Define $\alpha = E_{h,\theta_{-i}}\{\sum_{t\geq 0}(1-\delta)\delta^t(\sigma_i^{\theta^+}, (\sigma_j^{\theta_j})_{j\neq i})(h^t)\}$. By linearity of Bernoulli utility functions, $u_i(\alpha|\theta_i^+) = v_i(\theta_i^+)$. Also, by incentive compatibility, $u_i(\alpha|\theta_i) \leq v_i(\theta_i)$ for any $\theta_i \neq \theta_i^+$. Taking $K \to \infty$ concludes the proof. □

Lemma 4.2. *Let $\mu \in \Delta(\Theta)$. If for any smooth, strict subset $W \subset V$ there exists $\bar{\delta} \in (0, 1)$ such that for all $\delta > \bar{\delta}$ $W \subset E^{PBE}(\mu, \delta)$ then condition $B.2$ must hold.*

The proof of this lemma is similar to the proof of the previous one. We approximate an appropriate boundary payoff with payoffs in the interior (these can be sustained as PBEs); we then show that the dynamic strategies employed induce certain "static actions" and, finally, we show the existence of the desired action by considering the limit case as the interior payoffs approach the boundary. On the other hand, and contrary to the previous lemma, since many informed players are involved in the relevant incentive compatibility constraints, we need a slightly more notational-intensive proof.

Proof. Definitions and notations Consider $\theta^+ \in \Theta$, $J \subset N$ and $v(\theta^+) \in V(\theta^+)$ a J-undominated payoff profile. Extend $v(\theta^+)$ to $v \in V$ by setting (a) $v_n(\theta_n^+, \theta_{-n}) = v_n(\theta^+)$ for all $\theta_{-n} \in \Theta_{-n}$ (b) $v_i(\theta_i, \theta_{-i}) = \underline{\mathrm{w}}_i$ for all $\theta_i \neq \theta_i^+$ and all $\theta_{-i} \in \Theta_{-i}$. For each $K \in \mathbb{N}$ let $v^K \in V^\circ$ such that $\|v^K - v\| \leq \frac{1}{K}$ and such that (i) $v_n^K(\theta_n^+, \theta_{-n}) = v_n^K(\theta^+)$ for all $\theta_{-n} \in \Theta_{-n}$ (ii) $v_n^K(\theta_n, \theta_{-n}) = v_n^K(\theta_n, \theta'_{-n})$ for all $\theta_n \neq \theta_n^+$ and all $\theta_{-n}, \theta'_{-n} \in \Theta_{-n}$. Since this payoff is independent of the rival's type, we abbreviate it $v_n(\theta_n)$. Since each $v^K \in V^\circ$, there exist (σ^K, δ_K) such that σ^K sustains v^K as a μ, δ_K PBE. Note that w.l.o.g. $\delta_K \to 1$ as $K \to \infty$. Moreover, let $P^K \in \Delta(\Theta \times \mathcal{H})$ be the probability induced by σ^K and E^K denote expectation with respect to said measure. Finally, note that Tychonoff's theorem guarantees S is compact, and thus σ^K converges to some σ (up to a subsequence) and thus P^K converges (up to subsequence) to some P.

Step 1: Take any history h^t (possibly with $t = \infty$ thus denoting a complete history), any $\theta_n \in \Theta_n$, and any $\bar{\sigma} \in S$. Let $\alpha(\bar{\sigma}, \theta_n, h^t) = E^K_{h', \theta_{-n}|h^t}\{\sum_\tau (1-\delta_K)\delta_K^\tau (\bar{\sigma}^{\theta_n}_{h^t_n}, \bar{\sigma}_{h^t_{-n}})(h'^\tau)\}$. In words: given a dynamic strategy $\bar{\sigma}$, this action that summarizes the dynamic play a player n can expect, conditional on having arrived at history h^t, if he starts playing like player-type θ_n. By linearity of Bernoulli utility functions, it follows that $u_n(\alpha(\bar{\sigma}, \theta_n^+, h^t)|\theta_n) = E^K_{h', \theta_{-n}|h^t}\{\sum (1-\delta_K)\delta_K^\tau u_n((\bar{\sigma}_{h^t_n}\theta_n^+, \bar{\sigma}_{h^t_{-n}})(h'^\tau)|\theta_n)\}$. In words: the static payoff of action $\alpha(\bar{\sigma}, \theta_n^+, h^t)$ summarizes, conditional on reaching history h^t, the continuation payoff player-type θ_n can expect if he starts mimicking θ_n^+. For the case $\bar{\sigma} = \sigma^K$ note that, for $t = \infty$, so $h^t = h \in \mathcal{H}$ is an infinite history, if h is such that $P^K(\theta^+|h) > 0$ then $u_n(\alpha(\sigma^K, \theta_n^+, h)|\theta_n) = u_n(\alpha((\sigma^K_n, \sigma^{K^{\theta^+_{-n}}}_{-n}), \theta_n^+, h^t)|\theta_n)$. This follows from the fact -proved in theorem 3- that at infinity, all players in the support of the asymptotic distribution play the same strategy. Thus, at an infinite history where θ^+ receives positive probability, without loss of generality all player-types of the rival play as if they were θ^+_{-n}. Also, if h is such that $P^K(\theta^+|h) = 0$ then $u_n(\alpha(\sigma^K, \theta_n, h)|\theta_n) \geq \underline{\text{w}}_n(\theta_n)$ by individual rationality.

Step 2: Consider a player-type $\theta_i \neq \theta_i^+$. For each K, we construct a possible deviation from σ^K, denoted $\hat{\sigma}_i^K$, as follows: for $\theta_i' \neq \theta_i$ $\hat{\sigma}_i^{K^{\theta_i'}} = \sigma_i^{K^{\theta_i'}}$ and for each h^t such that $P^K(\theta^+|h^t) > 0$, $\hat{\sigma}_{h_i^t}^{K^{\theta_i}} = \sigma_{h_i^t}^{K^{\theta_i^+}}$, and $\hat{\sigma}_{h_i^t}^{K^{\theta_i}} = \sigma_{h_i^t}^{K^{\theta_i}}$ otherwise. That is, only player-type θ_i deviates and, for each history h^t, he only deviates if there is a positive probability that his rivals are θ^+_{-i}; in which case he deviates by mimicking θ_i^+. Fix an arbitrary period $T \in \mathbb{N}$; then the payoff from this deviation can be calculated as $U_i(\hat{\sigma}_i^{K^{\theta_i}}, \sigma_{-i}^K|\theta_i) = E_{h,\theta_{-i}}\{\delta_K^T u_i(\alpha(\hat{\sigma}, \theta_i, h^T)|\theta_i)\} + E_{h,\theta_{-i}}\{\sum_{t=0}^{T-1}(1-\delta_K)\delta_K^t u_i(\hat{\sigma}(h^t)|\theta_i)\}$ By equilibrium, this payoff is no greater than the compliance payoff $v_i^K(\theta_i)$. Moreover, let $\mathcal{H}_K^+ = \{h : P^K(\theta^+|h) > 0\}$ and $\mathcal{H}^+ = \{h : P(\theta^+|h) > 0\}$. Then, as $K \to \infty$ and using convergence results from Serfozo (1982), we get:

$$\underline{\text{w}}_i(\theta_i) \geq$$

$$U_i(\hat{\sigma}|\theta_i) = P(\mathcal{H}^+)E_{h|\mathcal{H}^+}\{u_i(\alpha(\hat{\sigma}, \theta_i, h^T)|\theta_i)\} + (1 - P(\mathcal{H}^+))E_{h|\complement\mathcal{H}^+}\{u_i(\alpha(\hat{\sigma}, \theta_i, h^T)|\theta_i)\}$$

$$\geq P(\mathcal{H}^+)E_{h|\mathcal{H}^+}\{u_i(\alpha(\hat{\sigma}, \theta_i, h^T)|\theta_i)\} + (1 - P(\mathcal{H}^+))\underline{\text{w}}_i(\theta_i)$$

where the last inequality follows from sequential rationality after every history.

Since this holds for all $T \in \mathbb{N}$, taking limits as $T \to \infty$ and noting that $\alpha(\hat{\sigma}, \theta_i, h) = \alpha((\sigma_i, \sigma_{-i}), \theta_i^+, h)$ for all $h \in H^+$, we get $u_i(\alpha^*|\theta_i) \leq \underline{\text{w}}_i(\theta_i)$ where $\alpha^* = E_{h|H^+}\{\alpha((\sigma_i, \sigma_{-i}^{\theta^+_{-i}}), \theta_i^+, h)\}$. Since player-type $\theta_i \neq \theta_i^+$ was arbitrarily selected, this holds for all such player types. The next step shows that $u_n(\alpha^*|\theta_n^+) = v_n(\theta^+)$ thus concluding the proof.

Step 3: By definition, strategy $\sigma^{K^{\theta^+}}$ yields playoffs $v_n^K(\theta^+)$ to player types θ_n^+. Since under strategy $\sigma^{K^{\theta^+}}$ the set $\mathcal{H}_K^+$ receives probability 1 we get:

$$v_n^K(\theta^+) = U_n(\sigma^{K^{\theta^+}}|\theta_n^+) =$$
$$E_{h,\theta_{-i}|\mathcal{H}_K^+}\{\delta_K^T u_i(\alpha(\sigma^{K^{\theta^+}}, \theta_i, h^T)|\theta_i^+)\} + E_{h,\theta_{-i}|\mathcal{H}_K^+}\{\textstyle\sum_{t=0}^{T-1}(1-\delta_K)\delta_K^t u_i(\sigma^{K^{\theta^+}}|\theta_i^+)\}$$

Taking limits as $K \to \infty$ first and $T \to \infty$ next yields

$$v_n(\theta^+) = u_n(\alpha^*|\theta_n^+)$$

□

A.5 Dropping Condition B

If a given game does not satisfy condition B then we know that a PTXE or PBE folk theorem cannot hold for that game. Yet, to the extent that the set of PTXE (and thus PBE) payoffs is not empty, there are some payoff profiles that can be sustained in equilibrium. This section deals with characterizing these. To this effect, consider an arbitrary payoff vector $v \in V$ and conditions $C.1$ and $C.2$ below:

Condition C.1. *For every $i \in I$ $\theta^+ \in \Theta$ there exists α such that*

- $u_i(\alpha|\theta_i^+) - v_i(\theta^+) \geq 0$
- $u_i(\alpha|\theta_i) - v_i(\theta_i, \theta_{-i}^+) \leq 0$ *for all* $\theta_i \neq \theta_i^+$

Condition C.2. *For every $\theta^+ \in \Theta$ and any $J \subset N$ there exists α such that*

- $u_n(\alpha|\theta_n^+) - v_n(\theta^+) \geq 0$ *for all* $n \in J$
- $u_n(\alpha|\theta_n^+) - v_n(\theta^+) \leq 0$ *for all* $n \notin J$
- $u_i(\alpha|\theta_i) - v_n(\theta_i, \theta_{-i}^+) \leq 0$ *for all* $\theta_i \neq \theta_i^+$

Conditions $C.1$ and $C.2$ are a straightforward generalization of conditions $B.1$ and $B.2$ respectively. Thus, by mimicking the proofs of Lemmas $B.1$ and $B.2$, it is straightforward to see that a payoff vector $v \in V$ satisfies $\lambda \cdot v \leq k^*(\lambda)$ for all $\lambda \in \Lambda^6 \cup \Lambda^7$ if and only if it satisfies conditions $C.1$ and $C.2$.

Checking that a specific vector $v \in V$ satisfies $k^*(\lambda) \geq \lambda \cdot v$ for all $\lambda \in \Lambda^5$ requires an extra definition. For $n \in N$ let $\mathcal{L}(n) = \{l \in \mathbb{R}^{\Theta} : (\exists\ a^*_{-n} : \Delta(A_n) \to A_{-n})$ such that $l(\theta) = \max_{\alpha_n} u_n(\alpha_n, a^*_{-n}(\alpha_n)|\theta_n)\}$.[14] Then, $k^*(\lambda) - \lambda \cdot v \geq 0$ for all $\lambda \in \Lambda^5$ if and only if for every i there exists some $l \in \mathcal{L}(i)$ such that $v_i(\theta) \geq l(\theta)$. Before proving this result, notice that when there is an action that minimaxes player i uniformly over his types then $l(\theta) \equiv \underline{\mathrm{w}}_i(\theta) \in \mathcal{L}(i)$ for all $v \in V$ and $n \in N$. Thus, when condition $B.2$ applies to the special case $J = \emptyset$, for each $n \in N$ there is an $l \in \mathcal{L}(n)$ such that, for every $v \in V$ $l(\theta) \geq v_n(\theta)$. In this sense, the sets $\mathcal{L}(i)$ generalize the notion of "lower bound" when the action profile that minimaxes all players cannot be sustained.

To show the result note:

$k^*(\lambda) - \lambda \cdot v \geq 0$ for all $\lambda \in \Lambda^5(i) \Leftrightarrow$
$(\forall \lambda \in \Lambda^5)\ \max_{\alpha_{-i}} \min_{a_i}\{\sum_\theta \lambda_i(\theta)[u_i(a_i, \alpha_{-i}|\theta) - v_i(\theta)]\} \geq 0 \Leftrightarrow$
$(\forall \lambda \in \Lambda^5(i))\ \max_{\alpha_i} \min a_{-i}\{\sum_\theta |\lambda_i(\theta)|[u_i(\alpha_i, a_{-i}|\theta) - v_i(\theta)]\} \leq 0 \Leftrightarrow$
$\max_{\lambda \in \Lambda^5(i)} \max_{\alpha_i} \min a_{-i}\{\sum_\theta |\lambda_i(\theta)|[u_i(\alpha_i, a_{-i}|\theta) - v_i(\theta)]\} \leq 0 \Leftrightarrow$
$\max_{\alpha_i} \min a_{-i} \max_{\theta \in \Theta}\{\sum_\theta |\lambda_i(\theta)|[u_i(\alpha_i, a_{-i}|\theta) - v_i(\theta)]\} \leq 0 \Leftrightarrow$
$(\forall \alpha_i)(\exists a_{-i}) : (\forall \theta)[u_i(\alpha_i, a_{-i}|\theta) - v_i(\theta)] \leq 0 \Leftrightarrow$
$(\exists a^*_{-i} : \Delta(A_i) \to A_{-i}) : (\forall \alpha_i)(\forall \theta)[u_i(\alpha_i, a^*_{-i}(\alpha_i)|\theta) - v_i(\theta)] \leq 0 \Leftrightarrow$
$(\exists l \in \mathcal{L}(i)) : (\forall \theta)[l(\theta) - v_i(\theta)] \leq 0$

In short, a payoff profile $v \in V$ is such that $v \in Q$, and thus $v \in \lim_{\delta \to 1} E^{PTXE}(\delta)$ under the dimensionality condition, if and only if it satisfies $C.1$, $C.2$ and for each i there exists $l \in \mathcal{L}(i)$ such that $l(\theta) \leq v_i(\theta)$ for each $\theta \in \Theta$.

References

Abreu, D,; Pearce, D. and Stacchetti, E. (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring", Econometrica.

[14] Of course, the real bite of this set is when $n \in I$. That is, when player n has some private information.

Aumann, R. and Maschler, M. (1995): "Repeated Games with Incomplete Information", The MIT Press.

Cripps, M., and Thomas, J. (2003): "Some asymptotic results in discounted repeated games of one sided incomplete information" Mathematics of Operations Research.

Cripps, M., Dekel, E., and Pesendorfer, W. (2005): "Reputation with equal discounting in repeated games with strictly conflicting interests" Journal of Economic Theory.

Cripps, M., Mailath, G., and Samuelson, L. (2007): "Disappearing private reputations in long-run relationships" Journal of Economic Theory.

Cripps, M., Mailath, G., and Samuelson, L. (2004): "Imperfect monitoring and impermanent reputations" Econometrica.

Escobar, J. and Toikka, J.,(2011): "A Folk Theorem with Markovian Private Information".

Fudenberg, D., and Y. Yamamoto (2010): "Learning from Private Information in Noisy Repeated Games," working paper.

Hörner, J., and Lovo, S. (2009): "Belief-Free Equilibria in games with incomplete information" Econometrica

Hörner, J., Lovo, S., and Tomala, T. (2011): "Belief-Free Equilibria in games with incomplete information:Characterization and Existence" Journal of Economic Theory.

Kalai, E., Lehrer E. (1995) "Rational learning leads to Nash Equilibrium" Econometrica.

Kalai, E., Lehrer E. (1995) "Subjective games and equilibria" Games and Economic behavior.

Lehrer, E., Smorodinsky R. (1997) "Repeated Large Games with incomplete information" Games and Economic behavior.

Mailath, G. (1988) ”On the behavior of separating Equilibria of signaling games with a finite set of types as the set of types becomes dense in an interval” Journal of Economic Theory.

Peski, Marcin (2008): “Repeated Games with incomplete information on one side”, Theoretical Economics.

Renault, J., and Tomala, T. (2004) ”Learning the state of nature in repeated games with incomplete information and signals” Games and Economic Behavior.

Serfozo, R. (1982) ”Convergence of Lebesgue integrals with varying measures” Sankhy: The Indian Journal of Statistics.

Shalev, J. (1994) ”Nonzero-Sum Two-Person Repeated Games with Incomplete Information and Known-Own-Payoffs” Games and Economic Behavior.