

A STRUCTURE THEOREM FOR RATIONALIZABILITY IN INFINITE-HORIZON GAMES

JONATHAN WEINSTEIN AND MUHAMET YILDIZ

ABSTRACT. We show that in any game that is continuous at infinity, if a plan of action a_i is rationalizable for a type t_i , then there are perturbations of t_i for which following a_i for an arbitrarily long future is the *only* rationalizable plan. One can pick the perturbation from a finite type space with common prior. Furthermore, for Bayesian Nash equilibria, the perturbation can be picked so that the unique rationalizable belief of the perturbed type regarding the play of the game is arbitrarily close to the equilibrium belief of t_i . As an application we prove an unusual folk theorem: Any individually rational and feasible payoff is the unique rationalizable payoff vector for some perturbed type profile, according to which the repeated-game payoff structure and the discount factor are common knowledge.

JEL Numbers: C72, C73.

1. INTRODUCTION

In economic applications with infinite-horizon dynamic games, the sets of equilibrium strategies and rationalizable strategies are often very large. For example, the literature on repeated games is filled with folk theorems, concluding that every individually rational payoff can be supported by a subgame-perfect equilibrium. For a less transparent example, in Rubinstein's (1982) bargaining game, although there is a unique subgame-perfect equilibrium, any outcome can be supported in Nash equilibrium. Consequently, economists focus on strong refinements of equilibrium and ignore other rationalizable strategies and equilibria. This is so common that we rarely think about rationalizable strategies in extensively-analyzed dynamic games. Of course, all these applications make strong common-knowledge assumptions.

Date: First Version: March 2009; This Version: April 2011.

This paper was written when the authors were members at the Institute for Advanced Study. We thank them for their generous financial support and hospitality.

In this paper, building on existing theorems for finite games, we prove structure theorems for rationalizability in infinite-horizon dynamic games that characterize the robust predictions of any refinement. The main attraction of our results is that they are readily applicable to most economic applications. Indeed, we provide two immediate applications, one in repeated games with sufficiently patient players and one in bargaining, showing that no refinement can robustly rule out any individually rational outcome in these games.

We consider an arbitrary dynamic game that is continuous at infinity, and has finitely many moves at each information set and a finite type space. Note that virtually all of the games analyzed in economics, such as repeated games with discounting and bargaining games, are continuous at infinity. For any type t_i in this game, consider a rationalizable plan of action a_i , which is a complete contingent plan that determines which move the type t_i will take at any given information set of i .¹ Fix some arbitrary integer L . We show that, by perturbing the interim beliefs of type t_i , we can find a new type \hat{t}_i who plays according to a_i in the first L information sets in *any* rationalizable action. The types t_i and \hat{t}_i have similar beliefs about the payoff functions, similar beliefs about the other players' beliefs about the payoff functions, similar beliefs about the other players' beliefs about the players' beliefs about the payoff functions, and so on, up to an arbitrarily chosen finite order. Moreover, we can pick \hat{t}_i from a finite model with a common prior, so that our perturbations do not rely on some esoteric large type space or the failure of the common-prior assumption.

In Weinstein and Yildiz (2007) we showed this result for finite-action games in normal form, under the assumption that the space of payoffs is rich enough so that any action can be dominant under some payoff specification. While this richness assumption holds when one relaxes all common-knowledge assumptions on payoff functions in a static game, it fails if one fixes a non-trivial dynamic game tree. This is because a plan of action cannot be strictly dominant when some information sets may not be reached. Chen (2008) has nonetheless extended the structure theorem to finite dynamic games, showing that the same result holds

¹The usual notation in dynamic games and games of incomplete information clash; action a_i stands for a move in dynamic games but for an entire contingent plan in incomplete-information games; t stands for time in dynamic games but type profile in incomplete-information games; h_i stands for history in dynamic games but hierarchy in incomplete-information games, etc. Following Chen, we will use the notation customary in incomplete information games, so a_i is a complete contingent plan of action. We will sometimes use “move” to distinguish an action at a single node.

under the weaker assumption that all payoff functions on the terminal histories are possible. This is an important extension, but the finite-horizon assumption rules out many major dynamic applications of game theory, such as repeated games and sequential bargaining. Since the equilibrium strategies can discontinuously expand when one switches from finite- to infinite-horizon, as in the repeated prisoners' dilemma game, it is not clear what the structure theorem for finite-horizon game implies in those applications. Here, we extend Chen's results further by allowing infinite-horizon games that are continuous at infinity, an assumption that is made in almost all applications. There is a challenge in this extension, because the construction employed by Weinstein and Yildiz (2007) and Chen (2008) relies on the assumption that there are finitely many actions. The finiteness (or countability) of the action space is used in a technical but crucial step of ensuring that the constructed type is well-defined, and there are counterexamples to that step when the action space is uncountable. Unfortunately, in infinite-horizon games, such as infinitely-repeated prisoners dilemma, there are uncountably many strategies, even in reduced form. However, continuity at infinity turns out to be enough to make infinite-horizon games behave well enough for the result to carry over.

We now briefly explain the implications of our structure theorem to robustness.² Imagine a researcher who subscribes to an arbitrary refinement of rationalizability, such as sequential equilibrium or proper equilibrium. Applying his refinement, he can make many predictions about the outcome of the game, describing which histories we may observe. Let us confine ourselves to predictions about finite-length (but arbitrarily long) outcome paths. For example, in the repeated prisoners' dilemma game, "players cooperate in the first round" and "player 1 plays tit-for-tat in the first $10^{1,000,000}$ periods" are such predictions, but "players always cooperate" and "players eventually defect" are not. Our result implies that any such prediction that can be obtained by a refinement, but not by mere rationalizability, relies crucially on assumptions about the infinite hierarchies of beliefs embedded in the model. Therefore, refinements cannot lead to any new prediction about finite-length outcome paths that is robust to misspecification of interim beliefs.

²For a more detailed discussion of the ideas in this paragraph, we refer to Weinstein and Yildiz (2007). In particular, there, we have extensively discussed the meaning of perturbing interim beliefs from the perspective of economic modelling and compared alternative formulations, such as the ex-ante perturbations of Kajii and Morris (1997).

One can formally derive this from our result by following the formulation in Weinstein and Yildiz (2007). Here, we will informally illustrate the basic intuition. Suppose that the above researcher observes a "noisy signal" about the players' first-order beliefs (which are about the payoff functions), the players' second-order beliefs (which are about the first-order beliefs), \dots , up to a finite order k , and does not have any information about the beliefs at order higher than k . Here, the researcher's information may be arbitrarily precise, in the sense that the noise in his signal may be arbitrarily small and k may be arbitrarily large. Suppose that he concludes that a particular type profile $t = (t_1, \dots, t_n)$ is consistent with his information, in that the interim beliefs of each type t_i could lead to a hierarchy of beliefs that is consistent with his information. Suppose that for this particular specification, his refinement leads to a sharper prediction about the finite-length outcome paths than rationalizability. That is, for type profile t , a particular path (or history) h of length L is possible under rationalizability but not possible under his refinement. But there are many other type profiles that are consistent with his information. In order to verify his prediction that h will not be observed under his refinement, he has to make sure that h is not possible under his refinement for any such type profile. Otherwise, his prediction would not follow from his information or solution concept; it would rather be based on his modeling choice of considering t but not the alternatives. Our result establishes that he cannot verify his prediction, and his prediction is indeed based on his choice of modeling: there exist a type profile \hat{t} that is also consistent with his information and, for \hat{t} , h is the *only* rationalizable outcome for the first L moves, in which case h is the only outcome for the first L moves *according to his refinement* as well.

Our structure theorem has two limitations. First, it only applies to finite-length outcomes. Second and more importantly, the perturbed types may find the unique rationalizable outcome unlikely at the beginning of play. In particular, a player may expect to play different moves in the future from what he actually plays according to the unique rationalizable plan. Focusing on Bayesian Nash equilibria, we prove a stronger structure theorem that does not have these limitations. For any Bayesian Nash equilibrium of any Bayesian game that is continuous at infinity, we show that for every type t_i in the Bayesian game there exists a perturbed type for which the equilibrium action of t_i is the unique rationalizable action and the unique rationalizable belief of the perturbed type is arbitrarily close to the equilibrium

belief of t_i . In particular, if the original game is of complete information, then the perturbed type assigns nearly probability one to the equilibrium path. (We also show that such perturbations can be found only for Bayesian Nash equilibria.)

As an application of this stronger result and the usual folk theorems, we show an unusual folk theorem. We show that every individually rational and feasible payoff v in the interior can be supported by the unique rationalizable outcome for some perturbation for sufficiently patient players. Moreover, in the actual situation described by the perturbation, all players anticipate that the payoffs are within ε neighborhood of v . That is, the complete-information game is surrounded by types with a unique solution, but the unique solution varies in such a way that it traces all individually rational and feasible payoffs. While the multiplicity in usual folk theorems may suggest a need for a refinement, the multiplicity in our unusual folk theorem emphasizes the impossibility of a robust refinement. In the same vein, in Rubinstein's bargaining model, we show that any bargaining outcome can be supported as a unique rationalizable outcome for some perturbation. Once again, no refinement can rule out these outcomes without imposing a common knowledge assumption.

In some applications, a researcher may want to retain certain structural assumptions regarding payoffs. In particular, in a repeated game, he may want to keep it common knowledge that the players' payoffs in the repeated game is the discounted sum of the stage-game payoffs. In general, such restrictions may lead to sharper predictions. In the particular case of repeated games, however, we show that our conclusions above remain intact: the perturbed types in the unusual folk theorem can be constructed while maintaining common knowledge of the repeated-game payoff structure and the discount factor.

After laying out the model in the next section, we present our general results in Section 3. We present our applications to repeated games and bargaining in Sections 4 and 5, respectively. We discuss the relation of our general results to broader literature on robustness in Section 6. The proofs of our general results are presented in the appendix.

2. BASIC DEFINITIONS

This section will be a bit tedious because of the notation involved in defining dynamic Bayesian games and hierarchies of beliefs. We suggest that the reader skim the section quickly and refer back as necessary; the main text is not very notation-heavy.

Extensive game forms. We consider standard n -player extensive-form games with possibly infinite horizon, as modeled in Osborne and Rubinstein (1994). In particular, we fix an extensive game form $\Gamma = (N, H, (\mathcal{I}_i)_{i \in N})$ with perfect recall where $N = \{1, 2, \dots, n\}$ is a finite set of players, H is a set of histories, and \mathcal{I}_i is the set of information sets at which player $i \in N$ moves. We use $i \in N$ and $h \in H$ to denote a generic player and history, respectively. We write $I_i(h)$ for the information set that contains history h , at which player i moves, i.e. the set of histories i finds possible when he moves. The set of available moves at $I_i(h)$ is denoted by $B_i(h)$. We have $B_i(h) = \{b_i : (h, b_i) \in H\}$, where (h, b_i) denotes the history in which h is followed by b_i . We assume that $B_i(h)$ is finite for each h . An action a_i of i is defined as any contingent plan that maps the information sets of i to the moves available at those information sets; i.e. $a_i : I_i(h) \mapsto a_i(h) \in B_i(h)$. We write $A = A_1 \times \dots \times A_n$ for the set of action profiles $a = (a_1, \dots, a_n)$.³ We write Z for the set of terminal nodes, at which no player moves. We write $z(a)$ for the terminal history that is reached by profile a . We say that actions a_i and a'_i are *equivalent* if $z(a_i, a_{-i}) = z(a'_i, a_{-i})$ for all $a_{-i} \in A_{-i}$.

Type spaces. Given an extensive game form, a Bayesian game is defined by specifying the belief structure about the payoffs. To this end, we write $\theta(z) = (\theta_1(z), \dots, \theta_n(z)) \in [0, 1]^n$ for the payoff vector at the terminal node $z \in Z$ and write Θ^* for the set of all payoff functions $\theta : Z \rightarrow [0, 1]^n$. The payoff of i from an action profile a is denoted by $u_i(\theta, a)$. Note that $u_i(\theta, a) = \theta_i(z(a))$. We endow Θ^* with the product topology (i.e. the topology of pointwise convergence). Note that Θ^* is compact and u_i is continuous in θ . Note, however, that Θ^* is not a metric space. We will use only finite type spaces, so by a *model*, we mean a finite set $\Theta \times T_1 \times \dots \times T_n$ associated with beliefs $\kappa_{t_i} \in \Delta(\Theta \times T_{-i})$ for each $t_i \in T_i$, where $\Theta \subseteq \Theta^*$. Here, t_i is called a type and $T = T_1 \times \dots \times T_n$ is called a type space. A model (Θ, T, κ) is said to be a *common-prior model (with full support)* if and only if there exists a probability distribution $p \in \Delta(\Theta \times T)$ with support $\Theta \times T$ and such that $\kappa_{t_i} = p(\cdot | t_i)$

³**Notation:** Given any list X_1, \dots, X_n of sets, write $X = X_1 \times \dots \times X_n$ with typical element x , $X_{-i} = \prod_{j \neq i} X_j$ with typical element x_{-i} , and $(x_i, x_{-i}) = (x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_n)$. Likewise, for any family of functions $f_j : X_j \rightarrow Y_j$, we define $f_{-i} : X_{-i} \rightarrow X_{-i}$ by $f_{-i}(x_{-i}) = (f_j(x_j))_{j \neq i}$. This is with the exception that h is a history as in dynamic games, rather than a profile of hierarchies (h_1, \dots, h_n) . Given any topological space X , we write $\Delta(X)$ for the space of probability distributions on X , endowed with Borel σ -algebra and the weak topology.

for each $t_i \in T_i$. Note that $(\Gamma, \Theta, T, \kappa)$ defines a Bayesian game. In this paper, we consider games that vary by their type spaces for a fixed game form Γ .

Hierarchies of Beliefs. Given any type t_i in a type space T , we can compute the first-order belief $h_i^1(t_i) \in \Delta(\Theta^*)$ of t_i (about θ), second-order belief $h_i^2(t_i) \in \Delta(\Theta^* \times \Delta(\Theta^*)^n)$ of t_i (about θ and the first-order beliefs), etc., using the joint distribution of the types and θ . Using the mapping $h_i : t_i \mapsto (h_i^1(t_i), h_i^2(t_i), \dots)$, we can embed all such models in the universal type space, denoted by $T^* = T_1^* \times \dots \times T_n^*$ (Mertens and Zamir (1985) and Brandenburger and Dekel (1993)). We endow the universal type space with the product topology of usual weak convergence. We say that a sequence of types $t_i(m)$ converges to a type t_i , denoted by $t_i(m) \rightarrow t_i$, if and only if $h_i^k(t_i(m)) \rightarrow h_i^k(t_i)$ for each k , where the latter convergence is in weak topology, i.e., “convergence in distribution.”

For each $i \in N$ and for each belief $\pi \in \Delta(\Theta \times A_{-i})$, we write $BR_i(\pi)$ for the set of actions $a_i \in A_i$ that maximize the expected value of $u_i(\theta, a_i, a_{-i})$ under the probability distribution π . Moreover, a solution concept $\Sigma_i : t_i \mapsto \Sigma_i[t_i] \subseteq A_i$, $i \in N$, is said to be *closed under rational behavior* if and only if for each t_i and for each $a_i \in \Sigma_i[t_i]$, there exists a belief $\pi \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that $a_i \in BR_i(\text{marg}_{\Theta \times A_{-i}} \pi)$, $\text{marg}_{\Theta \times T_{-i}} \pi = \kappa_{t_i}$ and $\pi(a_{-i} \in \Sigma_{-i}[t_{-i}]) = 1$.

Interim Correlated Rationalizability. We define interim correlated rationalizability (ICR), denoted by S^∞ , as the largest solution concept that is closed under rational behavior. Under certain regularity conditions, e.g., in finite games, the interim correlated rationalizability can be computed by the following elimination procedure. For each i and t_i , set $S_i^0[t_i] = A_i$, and define sets $S_i^k[t_i]$ for $k > 0$ iteratively, by letting $a_i \in S_i^k[t_i]$ if and only if $a_i \in BR_i(\text{marg}_{\Theta \times A_{-i}} \pi)$ for some $\pi \in \Delta(\Theta \times T_{-i} \times A_{-i})$ such that $\text{marg}_{\Theta \times T_{-i}} \pi = \kappa_{t_i}$ and $\pi(a_{-i} \in S_{-i}^{k-1}[t_{-i}]) = 1$. That is, a_i is a best response to a belief of t_i that puts positive probability only to the actions that survive the elimination in round $k - 1$. We write $S_{-i}^{k-1}[t_{-i}] = \prod_{j \neq i} S_j^{k-1}[t_j]$ and $S^k[t] = S_1^k[t_1] \times \dots \times S_n^k[t_n]$. Then,⁴

$$S_i^\infty[t_i] = \bigcap_{k=0}^{\infty} S_i^k[t_i].$$

⁴In complete information games, this equality holds whenever the action spaces are compact and the utility functions are continuous (Bernheim (1984)). The equality may fail in other complete information games (Lipman (1994)).

Interim correlated rationalizability has been introduced by Dekel, Fudenberg, and Morris (2007) (see also Battigalli and Siniscalchi (2003) for a related concept). They show that the ICR set for a given type is completely determined by its hierarchy of beliefs, so we will sometimes refer to the ICR set of a hierarchy or “universal type.” ICR is the weakest rationalizability concept, and hence our results remain true under other notions of rationalizability.

Continuity at Infinity. We now turn to the details of the extensive game form. If a history $h = (b^l)_{l=1}^L$ is formed by L moves for some finite L , then h is said to be *finite* and have length L . If h contains infinitely many moves, then h is said to be *infinite*. A game form is said to have *finite horizon* if for some $L < \infty$ all histories have length at most L ; the game form is said to have *infinite horizon* otherwise. For any history $h = (b^l)_{l=1}^L$ and any L' , we write $h^{L'}$ for the subhistory of h that is truncated at length L' ; i.e. $h = (b^l)_{l=1}^{\min\{L, L'\}}$. We say that θ is *continuous at infinity* (first defined by Fudenberg and Levine (1983)) iff for any $\varepsilon > 0$, there exists $L < \infty$, such that

$$(2.1) \quad \left| \theta_i(h) - \theta_i(\tilde{h}) \right| < \varepsilon \text{ whenever } h^L = \tilde{h}^L$$

for all $i \in N$ and all terminal histories $h, \tilde{h} \in Z$. We say that a game $(\Gamma, \Theta, T, \kappa)$ is *continuous at infinity* if each $\theta \in \Theta$ is continuous at infinity.

We will confine ourselves to the games that are continuous at infinity throughout, including our perturbations. Note that most games analyzed in economics are continuous at infinity. This includes all finite-horizon games, repeated games with discounting, games of sequential bargaining, and so on. Of course, our assumption that $B_i(h)$ is finite restricts the games to finite stage games and finite set of possible offers in repeated games and bargaining, respectively.

3. STRUCTURE THEOREM

In this section we will present our main result, which shows that in a game that is continuous at infinity, if an action a_i is rationalizable for a type t_i , then there are perturbations of t_i for which following a_i for arbitrarily long future is the *only* rationalizable plan. As we will explain, we also prove a stronger version of the theorem for outcomes that occur in equilibrium.

Weinstein and Yildiz (2007) have proven a version of this structure theorem for finite action games under a richness assumption on Θ^* that is natural for static games but rules out fixing a dynamic extensive game form. Chen (2008) has proven this result for finite games under a weaker richness assumption that is satisfied in our formulation. The following result is implied by Chen's theorem.

Lemma 1 (Weinstein and Yildiz (2007) and Chen (2008)). *For any finite-horizon game $(\Gamma, \Theta, T, \kappa)$, for any type $t_i \in T_i$ of any player $i \in N$, any rationalizable action $a_i \in S_i^\infty [t_i]$ of t_i , and any neighborhood U_i of $h_i(t_i)$ in the universal type space T^* , there exists a hierarchy $h_i(\hat{t}_i) \in U_i$, such that for each $a'_i \in S_i^\infty [\hat{t}_i]$, a'_i is equivalent to a_i , and \hat{t}_i is a type in some finite, common-prior model.*

That is, if the game has finite horizon, then for any rationalizable action of a given type, we can make the given action uniquely rationalizable (in the reduced game) by perturbing the interim beliefs of the type. Moreover, we can do this by only considering perturbations that come from finite models with a common prior. In the constructions of Weinstein and Yildiz (2007) and Chen (2008), finiteness (or countability) of action space A is used in a technical but crucial step that ensures that the constructed type is indeed well-defined, having well-defined beliefs. The assumption ensures that a particular mapping is measurable, and there is no general condition that would ensure the measurability of the mapping when A is uncountable. Unfortunately, in infinite-horizon games, such as infinitely repeated games, there are uncountably many histories and actions. (Recall that an action here is a complete contingent plan of a type, not a move.) Our main result in this section extends the above structure theorem to infinite-horizon games. Towards stating the result, we need to introduce one more definition.

Definition 1. An action a_i is said to be L -equivalent to a'_i iff $z(a_i, a_{-i})^L = z(a'_i, a_{-i})^L$ for all $a_{-i} \in A_{-i}$.

That is, two actions are L -equivalent if both actions prescribe the same moves in the first L moves on the path against every action profile a_{-i} by others. For the first L moves a_i and a'_i can differ only at the information sets that they preclude. Of course this is the same as

the usual equivalence when the game has a finite horizon that is shorter than L . We are now ready to state our first main result.

Proposition 1. *For any game $(\Gamma, \Theta, T, \kappa)$ that is continuous at infinity, for any type $t_i \in T_i$ of any player $i \in N$, any rationalizable action $a_i \in S_i^\infty[t_i]$ of t_i , any neighborhood U_i of $h_i(t_i)$ in the universal type space T^* , and any L , there exists a hierarchy $h_i(\hat{t}_i) \in U_i$, such that for each $a'_i \in S_i^\infty[\hat{t}_i]$, a'_i is L -equivalent to a_i , and \hat{t}_i is a type in some finite, common-prior model.*

Imagine a researcher who wants to model a strategic situation with genuine incomplete information. He can somehow make some noisy observations about the players' (first-order) beliefs about the payoffs, their (second-order) beliefs about the other players' beliefs about the payoffs, \dots , up to a finite order. The noise in his observation can be arbitrarily small, and he can observe arbitrarily many orders of beliefs. Suppose that given his information, he concludes that his information is consistent with a type profile t that comes from a model that is continuous at infinity. Note that the set of hierarchies that is consistent with his information is an open subset $U = U_1 \times \dots \times U_n$ of the universal type space, and $(h_1(t_1), \dots, h_n(t_n)) \in U$. Hence, our proposition concludes that for every rationalizable action profile $a \in S^\infty[t]$ and any finite length L , the researcher cannot rule out the possibility that in the actual situation the first L moves have to be as in the outcome of a in *any* rationalizable outcome. That is, rationalizable outcomes can differ from a only after L moves. Since L is arbitrary, he cannot practically rule out any rationalizable outcome as the unique solution.

Notice that Proposition 1 differs from Lemma 1 only in two ways. First, instead of assuming that the game has a finite horizon, Proposition 1 assumes only that the game is continuous at infinity, allowing most games in economics. Second, it concludes that for the perturbed types all rationalizable actions are equivalent to a_i up to an arbitrarily long but finite horizon, instead of concluding that all rationalizable actions are equivalent to a_i . These two statements are, of course, equivalent in finite-horizon games.

A main step in our proof is indeed Lemma 1. There are, however, many involved steps that need to be spelled out carefully. Hence, we relegate the proof to the appendix. In order to illustrate the main idea, we now sketch out the proof for a simple but important case. Suppose that $\Theta = \{\bar{\theta}\}$ and $T = \{\bar{t}\}$, so that we have a complete information game, and a^* is a

Nash equilibrium of this game. For each m , perturb every history h at length m by assuming that thereafter the play will be according to a^* , which describes different continuations at different histories. Call the resulting history h^{m,a^*} . This can also be described as a payoff perturbation: define the perturbed payoff function θ^m by setting $\theta^m(h) = \bar{\theta}(h^{m,a^*})$ at every terminal history h . Now consider the complete-information game with perturbed model $\tilde{\Theta}^m = \{\theta^m\}$ and $T^m = \{\bar{t}^m\}$, where according to \bar{t}^m it is common knowledge that the payoff function is θ^m (essentially, players are forced to play according to a^* after the m th information set). We make three observations towards proving the proposition. We first observe that, since $\bar{\theta}$ is continuous at infinity, by construction, $\theta^m \rightarrow \bar{\theta}$, implying that $h_i(\bar{t}_i^m) \rightarrow h_i(\bar{t}_i)$. Hence, there exists $\bar{m} > L$ such that $h_i(\bar{t}_i^{\bar{m}}) \in U_i$. Second, there is a natural isomorphism between the payoff functions that do not depend on the moves after length \bar{m} , such as $\theta^{\bar{m}}$, and the payoff functions for the finite-horizon extensive game form that is created by truncating the moves at length \bar{m} . In particular, there is an isomorphism φ that maps the hierarchies in the universal type space $T^{\bar{m}*}$ for the truncated extensive game form to the types in universal type space T^* for the infinite-horizon game form that make the common-knowledge assumption that the moves after length \bar{m} are payoff-irrelevant. Moreover, the rationalizable moves for the first \bar{m} nodes do not change under the isomorphism, in that $a_i \in S_i^\infty[\varphi(t_i)]$ if and only if the restriction $a_i^{\bar{m}}$ of a_i to the truncated game is in $S_i^\infty[t_i]$ for any $t_i \in T^{\bar{m}*}$. We third observe that, since a^* is a Nash equilibrium, it remains a Nash equilibrium after the perturbation. This is because enforcing Nash equilibrium strategies after some histories does not give a new incentive to deviate. Therefore, a_i^* is a rationalizable strategy in the perturbed complete information game: $a_i^* \in S_i^\infty[\bar{t}_i^{\bar{m}}]$. Now, these three observations together imply that the hierarchy $\varphi^{-1}(h_i(\bar{t}_i^{\bar{m}}))$ for the finite-horizon game form is in an open neighborhood $\varphi^{-1}(U_i) \subset T_i^{\bar{m}*}$ and the restriction $a_i^{*\bar{m}}$ of a_i^* to the truncated game form is rationalizable for $\varphi^{-1}(h_i(\bar{t}_i^{\bar{m}}))$. Hence, by Lemma 1, there exists a type \tilde{t}_i such that (i) $h_i(\tilde{t}_i) \in \varphi^{-1}(U_i)$ and (ii) all rationalizable actions of \tilde{t}_i are \bar{m} -equivalent to $a_i^{*\bar{m}}$. Now consider a type \hat{t}_i with hierarchy $h_i(\hat{t}_i) \equiv \varphi(h_i(\tilde{t}_i))$, where \hat{t}_i can be picked from a finite, common-prior model because the isomorphic type \tilde{t}_i comes from such a type space. Type \hat{t}_i has all the properties in the proposition. First, by (i), $h_i(\hat{t}_i) \in U_i$ because

$$h_i(\hat{t}_i) = \varphi(h_i(\tilde{t}_i)) \in \varphi(\varphi^{-1}(U_i)) \subset U_i.$$

Second, by (ii) and the isomorphism in the second observation above, all rationalizable actions of \hat{t}_i are \bar{m} -equivalent to a_i^* .

There are two limitations of Proposition 1. First, it is silent about the tails. Given a rationalizable action a_i , it does not ensure that there is a perturbation under which a_i is the unique rationalizable plan—although it does ensure for an arbitrary L that there is a perturbation under which following a_i is the uniquely rationalizable plan up to L . The second limitation, which equally applies to Chen's (2008) result, is as follows. Given any rationalizable path $z(a)$ and L , Proposition 1 establishes that there is a profile $t = (t_1, \dots, t_n)$ of perturbed types for which $z^L(a)$ is the unique rationalizable path up to L . Nevertheless, these perturbed types may all find the path $z^L(a)$ unlikely at the start of play. This may lead to implausible-seeming outcomes such as in the following example – we use a two-stage game for simplicity, since the relevant idea is the same as for infinite games.

Cooperation in Twice-Repeated Prisoners' Dilemma. Consider a twice-repeated prisoners' dilemma game with complete information and with no discounting. We shall need the standard condition $u(C, D) + u(D, C) > 2u(D, D)$, where u is the payoff of player 1 in the stage game and C and D stand for the actions Cooperate and Defect, respectively. In the twice-repeated game, though of course there is a unique Nash equilibrium outcome, the following "tit-for-tat" strategy is rationalizable:

a^{T4T} : play Cooperate in the first round, and in the second round play what the other player played in the first round.

We show this rationalizability as follows. First, note that defection in every subgame, which we call a^{DD} , by both players is an equilibrium, so a^{DD} is rationalizable. Next, defection in the first period followed by tit-for-tat in the second period, which we call a^{DT} , is a best response to a^{DD} and therefore rationalizable. Finally, under the inequality above, a^{T4T} is a best response to a^{DT} and so is rationalizable. This tells us that cooperation in both rounds is possible under rationalizable play.

This counterintuitive sort of conclusion is one reason standard rationalizability is not ordinarily used for extensive-form games; it is extremely permissive. This makes the results of Chen (2008) more surprising. By his theorem, there exists a perturbation t^{T4T} of the common-knowledge type for which a^{T4T} is the unique rationalizable action. If both players have type t^{T4T} , the unique rationalizable action profile (a^{T4T}, a^{T4T}) leads to cooperation

in both rounds. However, we can deduce that the constructed type will necessarily have certain odd properties. Since t^{T4T} has a unique best reply, the player must assign positive probability to the event that the other player cooperates in the first round. Such cooperation must make him update his beliefs about the payoffs in such a way that Cooperate becomes a better response than Defect. Since the definition of perturbation requires that, ex ante, he believes with high probability the payoffs are similar to the repeated prisoner dilemma, under which Defect is dominant in the second round, this drastic updating implies that t^{T4T} finds it unlikely that the other player will play Cooperate in the first round. Hence, when both players have type t^{T4T} , the story must be as follows: they each cooperate in the first round even though they think they are playing Prisoners' Dilemma, motivated by a belief that the other player has plan a^{DT} . Then, when they see the other player cooperate, they drastically update their payoffs (which they believe to be correlated with the other player's type) and believe that it is optimal to cooperate in the second period.

This sort of perturbation, in which the induced behavior can only occur on a path the players themselves assign low probability, is to some extent unconvincing.⁵ As mentioned above, this motivates our Proposition 2 which shows that equilibrium outcomes can be induced by perturbations without this property. This reinforces the natural view that rationalizability is a weak solution concept in a dynamic context.

Stronger Structure Theorem for Equilibrium Outcomes. These limitations of Proposition 1 are the motivation for our next proposition, a stronger version of the structure theorem for which we need an outcome to be a Bayesian Nash equilibrium rather than merely rationalizable. In order to state the result formally, we need to introduce some new formalism. We write \bar{A} for the set of reduced-form action profiles in which each equivalence class is represented by a unique action. We call a probability distribution $\pi \in \Delta(\Theta^* \times T_{-i}^* \times \bar{A}_{-i})$ a *rationalizable belief of type t_i* if $\text{marg}_{\Theta \times T_{-i}} \pi = \kappa_{h_i(t_i)}$ and $\pi(a_{-i} \in S_{-i}^\infty[t_{-i}]) = 1$. Given any strategy profile $s^* : T \rightarrow A$, we write $\pi^*(\cdot | t_i, s^*) \in \Delta(\Theta^* \times T_{-i}^* \times \bar{A}_{-i})$ for the belief of type t_i given that the other players play according to s_{-i}^* . We write $\text{Pr}(\cdot | \pi, s_i)$ and $E[\cdot | \pi, a_i]$

⁵The possibility of a player assigning small probability to the actual outcome arises under rationalizability whenever we do not have an equilibrium. In this dynamic example, the disconnect between the actual situation and the players' beliefs is more severe: their belief about their own future play differs from what they end up playing. They anticipate defecting in the second period while they cooperate in the actual realized type profile.

for the resulting probability measure and expectation operator from playing a_i against belief π , respectively. The expectation operator under $\pi^*(\cdot|t_i, s^*)$ is denoted by $E[\cdot|s^*, t_i]$. Recall that we consider the open neighborhoods of beliefs in the weak* topology as in the usual convergence in distribution. With this formalism, our result is stated as follows.

Proposition 2. *Let $G = (\Gamma, \Theta, T, \kappa)$ be a Bayesian game that is continuous at infinity, and $s^* : T \rightarrow A$ be a strategy profile in G . Then, the following are equivalent.*

(A): s^* is a Bayesian Nash equilibrium of G .

(B): For any $i \in N$, for any $t_i \in T_i$, for any neighborhood U_i of $h_i(t_i)$ in the universal type space T^* , and for any neighborhood V_i of the belief $\pi^*(\cdot|t_i, s^*)$ of type t_i under s^* , there exists a hierarchy $h_i(\hat{t}_i) \in U_i$, such that

- (1) $a_i \in S_i^\infty[\hat{t}_i]$ iff a_i is equivalent to $s_i^*(t_i)$, and
- (2) the unique rationalizable belief $\hat{\pi} \in \Delta(\Theta^* \times T_{-i}^* \times \bar{A}_{-i})$ of \hat{t}_i is in V_i .

Moreover, for every $\varepsilon > 0$, \hat{t}_i above can be chosen so that $|E[u_j(\theta, a)|\pi, a_i^*] - E[u_j(\theta, a)|s^*, t_i]| \leq \varepsilon$ for all $j \in N$.

Given a Bayesian Nash equilibrium s^* , the first conclusion states that the equilibrium action $s_i^*(t_i)$ is the only rationalizable action for the perturbed type in reduced form. Hence, the first limitation of Proposition 1 does not apply. The second conclusion states that the rationalizable belief of the perturbed type \hat{t}_i is approximately the same as the equilibrium belief of the original type t_i . Hence, the second limitation of Proposition 1 does not apply, either. Moreover, the second conclusion immediately implies that the interim expected payoffs according to the perturbed type \hat{t}_i under rationalizability are close to the equilibrium expected payoffs according to t_i . All in all, Proposition 2 establishes that no equilibrium outcome can be ruled out as the unique rationalizable outcome without knowledge of infinite hierarchy of beliefs, both in terms of actual realization and in terms of players' ex-ante expectations.

One may wonder if one can reach such a strong conclusion for other rationalizable strategies. The answer is a firm No, according to Proposition 2. In fact, the proposition establishes that the converse is also true: if for every type t_i one can find a perturbation under which the the players' interim beliefs are close to the beliefs under the original strategy profile s^*

(condition 2) and if the action $s_i^*(t_i)$ is uniquely rationalizable for the perturbed type (condition 1), then s^* is a Bayesian Nash equilibrium. This is simply because, by the Maximum Theorem, the two conditions imply that $s_i^*(t_i)$ is indeed a best reply for t_i against s_{-i}^* .

In our applications, we will explore the implications of this result for some important complete-information games in Economics. In order to state the result for the complete-information games, we fix a payoff function θ^* , and consider the game in which θ^* is common knowledge. This game is represented by type profile $t^{CK}(\theta^*)$ in the universal type space.

Corollary 1. *Let $(\Gamma, \{\theta^*\}, \{t^{CK}(\theta^*)\}, \kappa)$ be a complete-information game that is continuous at infinity, and a^* be a Nash equilibrium of this game. For any $i \in N$, for any neighborhood U_i of $h_i(t_i^{CK}(\theta^*))$ in the universal type space T^* , and any $\varepsilon > 0$, there exists a hierarchy $h_i(\hat{t}_i) \in U_i$, such that for every rationalizable belief π of \hat{t}_i ,*

- (1) $a_i \in S_i^\infty[\hat{t}_i]$ iff a_i is equivalent to a_i^* ;
- (2) $\Pr(z(a^*) | \pi, a_i^*) \geq 1 - \varepsilon$, and
- (3) $|E[u_j(\theta, a) | \pi, a_i^*] - u_j(\theta^*, a^*)| \leq \varepsilon$ for all $j \in N$.

For any Nash equilibrium a^* of any complete-information game, the corollary presents a profile \hat{t} of perturbations under which (1) the equilibrium a^* is the unique rationalizable action profile, (2) all players' rationalizable beliefs assign nearly probability one to the equilibrium outcome $z(a^*)$, and (3) the expected payoffs under these beliefs are nearly identical to the equilibrium payoffs. As established in Proposition 2, one can find such perturbations only for Nash equilibria.

The proof of Proposition 2 uses a contagion argument that is suitable for equilibrium. In order to illustrate the construction, we sketch the proof for the complete-information games considered in the corollary. Building on Proposition 1 we first show that for each action a_i there exists a type t^{a_i} for which a_i is uniquely rationalizable, extending a result of Chen to infinite-horizon games. For any Nash equilibrium a^* of any complete-information game $(\Gamma, \{\theta^*\}, \{t^{CK}(\theta^*)\}, \kappa)$, we construct a family of types $t_{j,m,\lambda}$, $j \in N$, $m \in \mathbb{N}$, $\lambda \in [0, 1]$, by

$$\begin{aligned} t_{j,0,\lambda} &= t^{a_j^*}, \\ \kappa_{t_{j,m,\lambda}} &= \lambda \kappa_{t^{a_j^*}} + (1 - \lambda) \delta_{(\theta^*, t_{-i,m-1,\lambda})} \quad \forall m > 0, \end{aligned}$$

where $\delta_{(\theta^*, t_{-i, m-1, \lambda})}$ is the Dirac measure that puts probability one on $(\theta^*, t_{-i, m-1, \lambda})$. For large m and small λ , $t_{i, m, \lambda}$ satisfies all the desired properties of \hat{t}_i . First note that for $\lambda = 0$, under $t_{i, m, 0}$, it is m th-order mutual knowledge that $\theta = \theta^*$. Hence, when m is large and λ is small, the belief hierarchy of $t_{i, m, 0}$ is close to the belief hierarchy of $t_i^{CK}(\theta^*)$, according to which it is common knowledge that $\theta = \theta^*$. Second, for $\lambda > 0$, a_j^* is uniquely rationalizable for $t_{j, m, \lambda}$ in reduced form. To see this, observing that it is true for $m = 0$ by definition of $t_{j, 0, \lambda}$, assume that it is true up to some $m - 1$. Then, any rationalizable belief of any type $t_{j, m, \lambda}$ must be a mixture of two beliefs. With probability λ , his belief is the same as that of $t^{a_j^*}$, to which a_j^* is the unique best response in reduced form actions. With probability $1 - \lambda$, the true state is θ^* and the other players play a_{-j}^* (in reduced form), in which case a_j^* is a best reply, as a^* is a Nash equilibrium under θ^* . Therefore, in reduced form a_j^* is the unique best response to any of his rationalizable beliefs, showing that a_j^* is uniquely rationalizable for $t_{j, m, \lambda}$ in reduced form. Finally, for any $m > 0$, under rationalizability type $t_{i, m, \lambda}$ must assign at least probability $1 - \lambda$ on (θ^*, a_{-i}^*) in reduced form because a_{-i}^* is uniquely rationalizable for $t_{-i, m-1, \lambda}$ in reduced form.

4. APPLICATION: AN UNUSUAL FOLK THEOREM

In this section, we consider infinitely repeated games with complete information. Under the standard assumptions for the folk theorem, we prove an unusual folk theorem, which concludes that for every individually rational and feasible payoff vector v , there exists a perturbation of beliefs under which there is a unique rationalizable outcome and players expect to enjoy approximately the payoff vector v under any rationalizable belief.

For simplicity, we consider a simultaneous-action stage game $G = (N, B, g)$ where $B = B_1 \times \dots \times B_n$ is the set of profiles $b = (b_1, \dots, b_n)$ of moves and $g^* : B \rightarrow [0, 1]^n$ is the vector of stage payoffs. We have perfect monitoring. Hence, a history is a sequence $h = (b^l)_{l \in \mathbb{N}}$ of profiles $b^l = (b_1^l, \dots, b_n^l)$. In the complete-information game, the players maximize the average discounted stage payoffs. That is, the payoff function is

$$\theta_\delta^*(h) = (1 - \delta) \sum_{l=0}^{\infty} \delta^l g^*(b^l) \quad (\forall h = (b^l)_{l \in \mathbb{N}})$$

where $\delta \in (0, 1)$ is the discount factor, which we will let vary. Denote the repeated game by $G_\delta = (\Gamma, \{\theta_\delta^*\}, \{t^{CK}(\theta_\delta^*)\}, \kappa)$.

Let $V = co(g(B))$ be the set of feasible payoff vectors (from correlated mixed action profiles), where co takes the convex hull. Define also the pure-action min-max payoff as

$$\underline{v}_i = \min_{b_{-i} \in B_{-i}} \max_{b_i \in B_i} g^*(b)$$

for each $i \in N$. We define the set of feasible and individually rational payoff vectors as

$$V^* = \{v \in V \mid v_i > \underline{v}_i \text{ for each } i \in N\}.$$

We denote the interior of V^* by $intV^*$. The interior will be non-empty when a weak form of full-rank assumption holds. The following lemma states a typical folk theorem (see Proposition 9.3.1 in Mailath and Samuelson (2006) and also Fudenberg and Maskin (1991)).

Lemma 2. *For every $v \in intV^*$, there exists $\bar{\delta} < 1$ such that for all $\delta \in (\bar{\delta}, 1)$, G_δ has a subgame-perfect equilibrium a^* in pure strategies, such that $u(\theta_\delta^*, a^*) = v$.*

The lemma states that every feasible and individually rational payoff vector in the interior can be supported as the subgame-perfect equilibrium payoff when the players are sufficiently patient. Given such a large multiplicity, both theoretical and applied researchers often focus on efficient equilibria (or extremal equilibria). Combining such a folk theorem with Corollary 1, our next result establishes that the multiplicity is irreducible.

Proposition 3. *For all $v \in intV^*$ and $\varepsilon > 0$, there exists $\bar{\delta} < 1$ such that for all $\delta \in (\bar{\delta}, 1)$, every open neighborhood U of $t^{CK}(\theta_\delta^*)$ contains a type profile $\hat{t} \in U$ such that*

- (1) *each \hat{t}_i has a unique rationalizable action a_i^* in reduced form, and*
- (2) *under every rationalizable belief π of \hat{t}_i , the expected payoffs are all within ε neighborhood of v :*

$$|E[u_j(\theta, a) \mid \pi, a_i^*] - v| \leq \varepsilon \quad \forall j \in N.$$

Proof. Fix any $v \in intV^*$ and $\varepsilon > 0$. By Lemma 2, there exists $\bar{\delta} < 1$ such that for all $\delta \in (\bar{\delta}, 1)$, G_δ has a subgame-perfect equilibrium a^* in pure strategies, such that $u(\theta_\delta^*, a^*) = v$. Then, by Corollary 1, for any $\delta \in (\bar{\delta}, 1)$ and any open neighborhood U of $t^{CK}(\theta_\delta^*)$, there exists a type profile $\hat{t} \in U$ such that each \hat{t}_i has a unique rationalizable action a_i^* in reduced form (Part 1 of Corollary 1), and under every rationalizable belief π of \hat{t}_i , the expected payoffs are all within ε neighborhood of $u(\theta_\delta^*, a^*) = v$ (Part 3 of Corollary 1). \square

Proposition 3 establishes an unusual folk theorem. It states that every individually rational and feasible payoff v in the interior can be supported by the unique rationalizable outcome for some perturbation. Moreover, in the actual situation described by the perturbation, all players play according to the subgame-perfect equilibrium that supports v and all players anticipate that the payoffs are within ε neighborhood of v . That is, the complete-information game is surrounded by types with a unique solution, but the unique solution varies in such a way that it traces all individually rational and feasible payoffs. While the multiplicity in usual folk theorems may suggest a need for a refinement, the multiplicity in our unusual folk theorem emphasizes the impossibility of a robust refinement.

Chassang and Takahashi (2011) examine the question of robustness in repeated games from an ex ante perspective. That is, following Kajii and Morris (1997), they define an equilibrium as robust if approximately the same outcome is possible in a class of elaborations. (An elaboration is an incomplete-information game in which each player believes with high probability that the original game is being played.) They consider specifically elaborations with serially independent types, so that the moves of players do not reveal any information about their payoffs and behavior in the future. They obtain a useful one-shot robustness result—to paraphrase, an equilibrium of the repeated game is robust if the equilibrium at each stage game, augmented with continuation values, is risk-dominant. There are two major distinctions. First, their perturbations are defined from an ex ante perspective, by what players believe before receiving information. Ours are from an interim perspective, based on what players believe just before play begins. This could be subsequent to receiving information, but our setup does not actually require reference to a particular information structure (type space with prior). For more on the distinction between these approaches, see our 2007 paper. Second, while they focus on serially independent types, whose moves do not reveal any information about the future payoffs, the moves of our perturbed types reveal information about both their own and the other players' payoffs in the future stage games.

Structure Theorem with Uncertainty about the Stage Payoffs. An important drawback of the structure theorems is that they may rely on existence of types who are far from the payoff and information structure assumed in the original model. If a researcher is willing to make common knowledge assumptions regarding these structures, those structure theorems may become inapplicable. Indeed, recent papers (e.g. Weinstein and Yildiz (2011) and

Penta (2008)) study the robust predictions when some common knowledge assumptions are retained.

In repeated games, one may wish to maintain common knowledge of the repeated-game payoff structure. Unfortunately, in our proofs of the propositions above, the types we construct do not preserve common knowledge of such a structure — they may depend on the entire history in ways which are not additively separable across stages. It is more difficult to construct types with unique rationalizable action when we restrict the perturbations to preserve common knowledge of the repeated-game structure, but in our next two propositions we are able to do this. The proofs (deferred to the Appendix) are somewhat lengthy and require the use of incentive structures similar to those in the repeated-game literature.

For simplicity, we exclude the trivial cases by assuming that each player has at least two moves. For any fixed discount factor $\delta \in (0, 1)$, we define

$$(4.1) \quad \Theta_\delta^* = \left\{ \theta_{\delta,g}(h) \equiv (1 - \delta) \sum_{l=0}^{\infty} \delta^l g(b^l) \mid g : B \rightarrow [0, 1]^n \right\}$$

as the set of repeated games with discount factor δ . Here, Θ_δ^* allows uncertainty about the stage payoffs g , but fixes all the other aspects of the repeated game, including the discount factor. For a fixed complete information repeated game with stage-payoff function g^* , we are interested in the robust predictions against the perturbations in which it remains common knowledge that the payoffs come from Θ_δ^* , allowing only uncertainty about the stage payoffs. The complete information game is represented by type profile $t^{CK}(\theta_{\delta,g^*})$ in the universal type space. The next result extends the structure theorem in Corollary 1 to this case.

Proposition 4. *For any $\delta \in (0, 1)$, let $(\Gamma, \{\theta^*\}, \{t^{CK}(\theta_{\delta,g^*})\}, \kappa)$ be a complete-information repeated game and a^* be a Nash equilibrium of this game. For any $i \in N$, for any neighborhood U_i of $h_i(t_i^{CK}(\theta_{\delta,g^*}))$ in the universal type space T^* , any $\varepsilon > 0$ and any L , there exists a hierarchy $h_i(\hat{t}_i) \in U_i$, such that*

- (1) $a_i \in S_i^\infty[\hat{t}_i]$ iff a_i is L -equivalent to a_i^* ;
- (2) $|E[u_j(\theta, a) \mid \pi] - u_j(\theta^*, a^*)| \leq \varepsilon$ for all $j \in N$ and for all rationalizable belief π of \hat{t}_i on (θ, a) , and
- (3) according to \hat{t}_i it is common knowledge that $\theta \in \Theta_\delta^*$.

Proposition 4 strengthens Corollary 1 by adding the last condition that the perturbed type still finds it common knowledge that he is playing a repeated game that is identical to the original complete-information game in all aspects except for the stage payoffs. The conclusion is weakened only by being silent about the tails, which will be immaterial to our conclusions. Indeed, using Proposition 4 instead of Corollary 1 in the proof of Proposition 3, which is the main result in this application, one can easily extend that folk theorem to the world in which a researcher is willing to retain common knowledge of the repeated game structure:

Proposition 5. *For all $v \in \text{int}V^*$, there exists $\bar{\delta} < 1$ such that for all $\delta \in (\bar{\delta}, 1)$, for all $\varepsilon > 0$ and all $L < \infty$, every open neighborhood U of $t^{CK}(\theta_\delta^*)$ contains a type profile $\hat{t} \in U$ such that*

- (1) *each \hat{t}_i has a unique rationalizable action a_i^* up to date L in reduced form;*
- (2) *under every rationalizable belief π of \hat{t}_i , the expected payoffs are all within ε neighborhood of v :*

$$|E[u_j(\theta, a) | \pi] - v| \leq \varepsilon \quad \forall j \in N,$$

- (3) *and it is common knowledge according to \hat{t} that $\theta \in \Theta_\delta^*$.*

That is, even if a researcher is willing to assume the repeated game payoff structure, for high discount factors, he cannot rule out any feasible payoff vector as the approximate outcome of the unique rationalizable belief for some nearby type. Hence, allowing uncertainty about the stage payoffs is sufficient to reach the conclusion of the unusual folk theorem above.

Proposition 4 is proved in the Appendix. The proof involves showing that each action plan is uniquely rationalizable up to an arbitrarily long finite horizon for a type for which it is common knowledge that $\theta \in \Theta_\delta^*$. Using this fact one then constructs the nearby types in the proposition following the ideas sketched to illustrate the proof of Corollary 1 above. The construction of the former types is quite involved using the ideas from learning and incentives in repeated games. In the next example we illustrate the gist of the idea on the twice-repeated prisoners' dilemma.

Example 1. Consider the twice-repeated prisoners' dilemma above with $g_1^{PD}(C, D) + g_1^{PD}(D, C) > 2g_1^{PD}(D, D)$, where g_1^{PD} is the payoff of player 1 in the stage game and C and D stand for

the actions Cooperate and Defect, respectively. Recall that D dominates C and $\delta = 1$. We will construct a nearby type for which tit-for-tat strategy is uniquely rationalizable. To this end, we first construct some types (not necessarily nearby) for which certain action plans are uniquely rationalizable. For any strategy profile $b \in \{C, D\}^2$ in the stage game, consider the payoff function g^b where $g_i^b(b'_1, b'_2) = 1$ if $b'_i = b_i$ and $g_i^b(b'_1, b'_2) = 0$ otherwise. For a type $t_{i,b_i,0}$ that puts probability 1 on $\theta_{\delta,g^{(b_i,b_{-i})}}$ for some b_{-i} , playing b_i in the first round is uniquely rationalizable. Such a type may have multiple rationalizable actions in the second round as he may put zero probability at some history. But now consider type $t_{i,b_i,1}$ that puts probability $1/2$ on $(\theta_{\delta,g^{(b_i,b_{-i})}}, t_{-i,C,0})$ and probability $1/2$ on $(\theta_{\delta,g^{(b_i,b_{-i})}}, t_{-i,D,0})$ for some b_{-i} . Since types $t_{-i,C,0}$ and $t_{-i,D,0}$ play C and D , respectively, as their unique rationalizable move in the first round, type $t_{i,b_i,1}$ puts positive probability at all histories at the beginning of the second period that are not precluded by his own action. Hence, his unique rationalizable action plan is to play b_i at all histories. We next consider types $t_{i,k}$ with approximate k th-order mutual knowledge of prisoners' dilemma payoffs who Defect at all histories in their unique rationalizable plan. We consider $t_{i,1}$ who puts probability $1/2$ on each of $(\theta_{\delta,g^{PD}}, t_{-i,C,1})$ and $(\theta_{\delta,g^{PD}}, t_{-i,D,1})$. Since the other player does not react to the moves of player i and i is certain that he plays a prisoners' dilemma game, his unique rationalizable plan is to defect everywhere (as he assigns positive probabilities to both moves). For any small ε and $k > 1$, consider the type $t_{i,k}$ who puts probability $1 - \varepsilon$ on $(\theta_{\delta,g^{PD}}, t_{-i,k-1})$ and probability ε on $(\theta_{\delta,g^{PD}}, t_{-i,C,1})$. By the inductive application of the previous argument, type $t_{i,k}$ also defects at all histories as the unique rationalizable plan. Moreover, when ε is small, there is approximate k th-order mutual knowledge of prisoners' dilemma. Now for arbitrary $k > 1$ and small $\varepsilon > 0$, consider the type $\hat{t}_{i,k}$ that puts probability $1 - \varepsilon$ on $(\theta_{\delta,g^{PD}}, t_{-i,k-1})$ and probability ε on $(\theta_{\delta,g^{(C,C)}}, t_{-i,C,1})$. He has approximate k th-order mutual knowledge of prisoners' dilemma. Moreover, since his opponent does not react to his moves and ε is small, his unique rationalizable move at the first period is D . In the second period, if he observes that his opponent played D in the first period, he becomes sure that they play prisoners' dilemma and plays D as his unique rationalizable move. If he observes that his opponent played C , however, he updates his belief and put probability 1 on $g^{(C,C)}$ according to which C dominates D . In that case, he too plays C in the second period. The nearby types $\hat{t}_{i,k}$ play tit-for-tat while defecting at the beginning. Now consider the nearby types $\tilde{t}_{i,k}$ that put

probability $1 - \varepsilon$ on $(\theta_{\delta, g^{PD}}, \hat{t}_{-i, k-1})$ and probability ε on $(\theta_{\delta, g^{(C, C)}}, t_{-i, C, 1})$. These types play tit-for-tat with cooperation at the initial node as their unique rationalizable action.

5. APPLICATION: INCOMPLETE INFORMATION IN BARGAINING

In a model of bilateral bargaining with complete information, Rubinstein (1982) shows that there exists a unique subgame-perfect equilibrium. Subsequent research illustrates that the equilibrium result is sensitive to incomplete information. In this section, using Proposition 2, we show quite generally that the equilibrium must be highly sensitive: every bargaining outcome can be supported as the unique rationalizable outcome for a nearby model.

We consider Rubinstein's alternating-offer model with finite set of divisions. There are two players, $N = \{1, 2\}$, who want to divide a dollar. The set of possible shares is $X = \{0, 1/m, 2/m, \dots, 1\}$ for some $m > 1$. At date 0, Player 1 offers a division $(x, 1 - x)$, where $x \in X$ is the share of Player 1 and $1 - x$ is the share of Player 2. Player 2 decides whether to accept or reject the offer. If he accepts, the game ends with division $(x, 1 - x)$. Otherwise, we proceed to the next date. At date 1, Player 2 offers a division $(y, 1 - y)$, and Player 1 accepts or rejects the offer. In this fashion, players make offers back and forth until an offer is accepted. We denote the bargaining outcome by (x, l) if players reach an agreement on division $(x, 1 - x)$ at date l . In the complete-information game, the payoff function is

$$\theta^* = \begin{cases} \delta^l(x, 1 - x) & \text{if the outcome is } (x, l) \\ 0 & \text{if players never agree} \end{cases}$$

for some $\delta \in (0, 1)$.

When $X = [0, 1]$, in the complete information game $G^* = (\Gamma, \{\theta^*\}, \{t^{CK}(\theta^*)\}, \kappa)$, there is a unique subgame perfect equilibrium, and the bargaining outcome in the unique subgame-perfect equilibrium is

$$(x^*, 0) = (1/(1 + \delta), 0).$$

That is the players immediately agree on division $(x^*, 1 - x^*)$. When $X = \{0, 1/m, \dots, 1\}$ as in here, there are more subgame-perfect equilibria due to multiple equilibrium behavior in the case of indifference. Nevertheless, the bargaining outcomes of these equilibria all converge to $(x^*, 0)$ as $m \rightarrow \infty$.

In contrast with the unique subgame-perfect equilibrium, there is a large multiplicity of non-subgame-perfect Nash equilibria, but these equilibria are ignored as they rely on incredible threats or sequentially irrational moves off the path. Building on such non-subgame-perfect Nash equilibria and Proposition 2, the next result shows that each bargaining outcome is the outcome of unique rationalizable action plan under some perturbation.

Proposition 6. *For any bargaining outcome $(x, l) \in X \times \mathbb{N}$ and any $\varepsilon > 0$, every open neighborhood U of $t^{CK}(\theta_\delta^*)$ contains a type profile $\hat{t} \in U$ such that*

- (1) *each \hat{t}_i has a unique rationalizable action a_i^* in reduced form;*
- (2) *the bargaining outcome under a^* is (x, l) , and*
- (3) *every rationalizable belief of \hat{t}_i assigns at least probability $1 - \varepsilon$ on (x, l) .*

Proof. We will show that the complete-information game has a Nash equilibrium a^* with bargaining outcome (x, l) . Proposition 2 then establishes the existence of type profile \hat{t} as in the statement of the proposition. Consider the case of even l , at which Player 1 makes an offer; the other case is identical. Define a^* in reduced-form as

- (a_1^*) at any date $l' \neq l$, offer only $(1, 0)$ and reject all the offers; offer $(x, 1 - x)$ at date l ;
- (a_2^*) at any date $l' \neq l$, offer only $(0, 1)$ and reject all the offers; accept only $(x, 1 - x)$ at l .

It is clear that a^* is a Nash equilibrium, and the bargaining outcome under a^* is (x, l) . \square

That is, for every bargaining outcome (x, l) , one can introduce a small amount of incomplete information in such a way that the resulting type profile has a unique rationalizable action profile and it leads to the bargaining outcome (x, l) . Moreover, in the perturbed type profile, players are all nearly certain that (x, l) will be realized. Unlike in the case of non-subgame-perfect equilibria, one cannot rule out these outcomes by refinement because there is a unique rationalizable outcome. In order to rule out these outcomes, one either needs to introduce irrational behavior or rule out the information structure that leads to the perturbed type profile by fiat (as he cannot rule out these structures by observation of finite-order beliefs without ruling out the original model). Therefore, despite the unique subgame-perfect outcome in the original model, and despite the fact that this outcome has generated many important and intuitive insights, one cannot make any prediction on the

outcome without introducing irrational behavior or making informational assumptions that cannot be verified by observing finite-order beliefs.

Existing literature illustrates already that the subgame-perfect equilibrium is sensitive to incomplete information. For example, for high δ , the literature on Coase conjecture establishes that if one party has a private information about his own valuation, then he gets everything—in contrast to the nearly equal sharing in the complete information game. This further leads to delay due to reputation formation in bargaining with two-sided incomplete information on payoffs (Abreu and Gul (2000)) or on players’ second-order beliefs (Feinberg and Skrzypacz (2005)).

Proposition 6 differs from these results in many ways. First difference is in the scope of sensitivity: while the existing results show that another outcome may occur under a perturbation, Proposition 6 shows that any outcome can be supported by a perturbation. Second difference is in the solution concept: while the existing result show sensitivity with respect to a sequential equilibrium or all sequential equilibria, there is a unique rationalizable outcome in Proposition 6, ruling out reinstating the original outcome by a refinement. Third, the existing results often consider the limit $\delta \rightarrow 0$, which is a point of discontinuity for the complete-information model already. In contrast, δ is fixed in Proposition 6. Finally, existing results consider simple perturbations, and these perturbations may correspond the specification of economic parameters, such as valuation, or may be commitment types. In contrast, given the generality of the results, the types constructed in our paper are complicated, and it is not easy to interpret how they are related to the economic parameters. (In specific examples, the same results could be obtained using simple types that correspond to economic parameters, as in Izmalkov and Yildiz (2010)).

6. CONCLUDING REMARKS AND LITERATURE REVIEW

Early literature identified two mechanisms through which small amount of incomplete information can have a large effect: reputation formation (Kreps, Milgrom, Roberts, and Wilson (1982)) and contagion (Rubinstein (1989)). In reputation formation, one learns about the other players’ payoffs from their unexpected moves. As we have seen in the twice-repeated prisoners’ dilemma game, our perturbed types exhibit this property at a higher level: they learn not only about the other players’ payoffs but also about their own payoffs

from the others' unexpected moves. Moreover, our perturbations are explicitly constructed using a generalized contagion argument. Hence, the perturbations here and in Chen (2008) combine the two mechanisms in order to obtain a very strong conclusion: any rationalizable action can be made uniquely rationalizable under some perturbation.

As the above example illustrates in our perturbations a player i may learn about the payoffs of j from the moves of another player k . This is reasonable in many contexts due to interdependence of preferences, and we are mainly dropping the common knowledge aspect of the assumption that k does not know the payoffs of j . Nevertheless, it may also be reasonable to keep it common knowledge that some parameters are only known by some players. For example, one may want to keep common knowledge that a firm's cost is its own private information, so that one does not update his beliefs about the firm's cost by observing some other player's move. Penta (2008) solves this problem. Keeping it common knowledge that some arbitrarily given parameters are known only by some arbitrarily given players, he obtains an identical structure theorem for what he calls interim sequential rationalizability (ISR) instead of interim correlated rationalizability (ICR). Of course, ISR depends on what is kept common knowledge and is equal to ICR when nothing is kept common knowledge. We believe that one can extend Penta's result to infinite-horizon games by modifying our construction, obtaining a more general result. We have not considered that extension for clarity because ICR is a more transparent solution concept, and for we believe that the case of dropping all common knowledge assumptions is an important benchmark.

Some other papers considered perturbations by keeping some payoffs structure common knowledge. For nice games, which are static games with unidimensional action spaces and strictly concave utility functions, Weinstein and Yildiz (2008) obtain a characterization for sensitivity of Bayesian Nash equilibria in terms of a local version of ICR, keeping arbitrary common knowledge restrictions on payoffs.⁶ In the same vein, Oury and Tercieux (2007) allow arbitrarily small perturbations on payoffs to obtain an equivalence between continuous partial implementation in Bayesian Nash equilibria and full implementation in rationalizable strategies.

⁶Weinstein and Yildiz (2011) also solve the problem of uncountable action spaces within the important class of nice games using a special structure of those games, which is clearly different from the structure in infinite-horizon games that allowed our characterization.

APPENDIX A. PROOF OF STRUCTURE THEOREM

We start with describing the notation we use in the appendix.

Notation 1. For any belief $\pi \in \Delta(\Theta \times A_{-i})$ and action a_i and for any history h , write $E[\cdot|h, a_i, \pi]$ for the expectation operator induced by action a_i and π conditional on reaching history h . For any strategy profile $s : T \rightarrow A$ and any type t_i , we write $\pi(\cdot|t_i, s_{-i}) \in \Delta(\Theta \times T_{-i} \times A_{-i})$ for the belief induced by t_i and s_{-i} . Given any functions $f : W \rightarrow X$ and $g : Y \rightarrow Z$, we write $(f, g)^{-1}$ for the preimage of the mapping $(w, y) \mapsto (f(w), g(y))$.

A.1. Preliminaries. We now define some basic concepts and present some preliminary results. By a *Bayesian game in normal form*, we mean a tuple $(N, A, u, \Theta, T, \kappa)$ where N is the set of players, A is the set of action profiles, (Θ, T, κ) is a model, and $u : \Theta \times A \rightarrow [0, 1]^n$ is the payoff function. While this notation is consistent with our formulation, we will also define some auxiliary Bayesian games with different action spaces, payoff functions and parameter spaces. For any $G = (N, A, u, \Theta, T, \kappa)$, we say that a_i and a'_i are *G-equivalent* if

$$u(\theta, a_i, a_{-i}) = u(\theta, a'_i, a_{-i}) \quad (\forall \theta \in \Theta, a_{-i} \in A_{-i}).$$

By a *reduced-form game*, we mean a game $G_R = (N, \bar{A}, u, \Theta, T, \kappa)$ where \bar{A}_i contains at least one representative action from each G -equivalence class for each i . Rationalizability depends only on the reduced form:

Lemma 3. *Given any game G and a reduced form G_R for G , for any type t_i , the set $S_i^\infty[t_i]$ of rationalizable actions in G is the set of all actions that are G -equivalent to some rationalizable action of t_i in G_R .*

The lemma follows from the fact that in the elimination process, all members of an equivalence class are eliminated at the same time; i.e. one eliminates, at each stage, a union of equivalence classes. It implies the following isomorphism for rationalizability.

Lemma 4. *Let $G = (N, A, u, \Theta, T, \kappa)$ and $G' = (N, A', u', \Theta', T', \kappa)$ be Bayesian games in normal form, $\mu_i : A_i \rightarrow A'_i$, $i \in N$, be onto mappings, and $\varphi : \Theta \rightarrow \Theta'$ and $\tau_i : T_i \rightarrow T'_i$, $i \in N$, be bijections. Assume (i) $\kappa_{\tau_i(t_i)} = \kappa_{t_i} \circ (\varphi, \tau_{-i})^{-1}$ for all t_i and (ii) $u'(\varphi(\theta), \mu(a)) = u(\theta, a)$ for all (θ, a) . Then, for any t_i and a_i ,*

$$(A.1) \quad a_i \in S_i^\infty[t_i] \iff \mu_i(a_i) \in S_i^\infty[\tau_i(t_i)].$$

Note that the bijections φ and τ are a renaming, and (i) ensures that the beliefs do not change under the renaming. On the other hand, μ_i can map many actions to one action, but (ii) ensures that all those actions are G -equivalent. The lemma concludes that rationalizability is invariant to such a transformation.

Proof. First note that (ii) implies that for any $a_i, a'_i \in A_i$,

$$(A.2) \quad a_i \text{ is } G\text{-equivalent to } a'_i \iff \mu_i(a_i) \text{ is } G'\text{-equivalent to } \mu_i(a'_i).$$

In particular, if $\mu_i(a_i) = \mu_i(a'_i)$, then a_i is G -equivalent to a'_i . Hence, there exists a reduced-form game $G_R = (N, \bar{A}, u, \Theta, T, \kappa)$ for G , such that μ is a bijection on \bar{A} , which is formed by picking a unique representative from each $\mu^{-1}(\mu(a))$. Then, by (A.2) again, $G'_R = (N, \mu(\bar{A}), u', \Theta', T', \kappa)$ is a reduced form for G' .⁷ Note that G_R and G'_R are isomorphic up to the renaming of actions, parameters, and types by μ , φ , and τ , respectively. Therefore, for any $a'_i \in \bar{A}_i$ and t_i , a'_i is rationalizable for t_i in G_R iff $\mu_i(a'_i)$ is rationalizable for $\tau_i(t_i)$ in G'_R . Then, Lemma 3 and (A.2) immediately yields (A.1). \square

We will also apply a Lemma from Mertens-Zamir (1985) stating that the mapping from types in any type space to their hierarchies is continuous, provided the belief mapping κ defining the type space is continuous.

Lemma 5 (Mertens and Zamir (1985)). *Let (Θ, T, κ) be any model, endowed with any topology, such that $\Theta \times T$ is compact and κ_{t_i} is a continuous function of t_i . Then, h is continuous.*

A.2. Truncated Games. We now formally introduce an equivalence between finitely-truncated games and payoff functions that implicitly assume such a truncation. For any positive integer m , define a truncated extensive game form $\Gamma^m = (N, H^m, (\mathcal{I}_i)_{i \in N})$ by

$$H^m = \{h^m | h \in H\}.$$

The set of terminal histories in H^m is

$$Z^m = \{z^m | z \in Z\}.$$

We define

$$\bar{\Theta}^m = ([0, 1]^{Z^m})^n$$

⁷Proof: Since μ_i is onto, $A'_i = \mu_i(A_i)$. Moreover, for any $\mu_i(a_i) \in A'_i$, there exists $a'_i \in \bar{A}_i$ that is G -equivalent to a_i . By (A.2), $\mu_i(a_i)$ is G' -equivalent to $\mu_i(a'_i) \in \mu_i(\bar{A}_i)$.

as the set of payoff functions for truncated game forms. Since Z^m is not necessarily a subset of Z , $\bar{\Theta}^m$ is not necessarily a subset of Θ^* . We will now embed $\bar{\Theta}^m$ into Θ^* through an isomorphism to a subset of Θ^* . Define the subset

$$\hat{\Theta}^m = \{\theta \in \Theta^* \mid \theta(h) = \theta(\bar{h}) \text{ for all } h \text{ and } \bar{h} \text{ with } h^m = \bar{h}^m\}.$$

This is the set of payoff functions for which moves after period m are irrelevant. Games with such payoffs are nominally infinite but inherently finite, as we formalize via the isomorphism $\varphi_m : \bar{\Theta}^m \rightarrow \hat{\Theta}^m$ defined by setting

$$(A.3) \quad \varphi_m(\theta)(h) = \theta(h^m)$$

for all $\theta \in \bar{\Theta}^m$ and $h \in Z$, where $h^m \in H^m$ is the truncation of h at length m . Clearly, under the product topologies, φ_m is an isomorphism, in the sense that it is one-to-one, onto, and both φ_m and φ_m^{-1} are continuous. For each $a_i \in A_i$, let a_i^m be the restriction of action a_i to the histories with length less than or equal to m . The set of actions in the truncated game form is $A_i^m = \{a_i^m \mid a_i \in A_i\}$.

Lemma 6. *Let $G = (\Gamma, \Theta, T, \kappa)$ and $G^m = (\Gamma^m, \Theta^m, T^m, \kappa)$ be such that (i) $\Theta^m \subset \bar{\Theta}^m$, (ii) $\Theta = \varphi_m(\Theta^m)$ and (iii) $T_i = \tau_i(T_i^m)$ for some bijection τ_i^m and such that $\kappa_{\tau_i^m(t_i^m)} = \kappa_{t_i^m} \circ (\varphi_m, \tau_{-i}^m)^{-1}$ for each $t_i^m \in T_i^m$. Then, the set of rationalizable actions are m -equivalent in G and G^m :*

$$a_i \in S_i^\infty[\tau_i^m(t_i^m)] \iff a_i^m \in S_i^\infty[t_i^m] \quad (\forall i, t_i^m, a_i).$$

Proof. In Lemma 4, take $\varphi = \varphi_m^{-1}$, $\tau_i = (\tau_i^m)^{-1}$, and $\mu : a_i \mapsto a_i^m$. We only need to check that $u^m(\varphi_m^{-1}(\theta), a^m) = u(\theta, a)$ for all (θ, a) where u^m denotes the utility function in the truncated game G^m . Indeed, writing $z^m(a^m)$ for the outcome of a^m in G^m , we obtain

$$\begin{aligned} u^m(\varphi_m^{-1}(\theta), a^m) &= \varphi_m^{-1}(\theta)(z^m(a^m)) = \varphi_m^{-1}(\theta)(z(a)^m) \\ &= \varphi_m(\varphi_m^{-1}(\theta))(z(a)) = \theta(z(a)) = u(\theta, a). \end{aligned}$$

Here, the first and the last equalities are by definition; the second equality is by definition of a^m , and the third equality is by definition (A.3) of φ_m . \square

Let T^{*m} be the $\bar{\Theta}^m$ -based universal type space, which is the universal type space generated by the truncated extensive game form. This space is distinct from the universal type space, T^* , for the original infinite-horizon extensive form. We will now define an embedding between the two type spaces, which will be continuous and one-to-one and preserve the rationalizable actions in the sense of Lemma 6.

Lemma 7. *For any m , there exists a continuous, one-to-one mapping $\tau^m : T^{*m} \rightarrow T^*$ with $\tau^m(t) = (\tau_1^m(t_1), \dots, \tau_n^m(t_n))$ such that for all $i \in N$ and $t_i \in T_i^{*m}$,*

- (1) t_i is a hierarchy for a type from a finite model if and only if $\tau_i^m(t_i)$ is a hierarchy for a type from a finite model;
- (2) t_i is a hierarchy for a type from a common-prior model if and only if $\tau_i^m(t_i)$ is a hierarchy for a type from a common-prior model, and
- (3) for all a_i , $a_i \in S_i^\infty[\tau_i^m(t_i)]$ if and only if $a_i^m \in S_i^\infty[t_i]$.

Proof. Since T^{*m} and T^* do not have any redundant type, by the analysis of Mertens and Zamir (1985), there exists a continuous and one-to-one mapping τ^m such that

$$(A.4) \quad \kappa_{\tau_i^m(t_i)} = \kappa_{t_i} \circ (\varphi_m, \tau_{-i}^m)^{-1}$$

for all i and $t_i \in T_i^{*m}$.⁸ First two statements immediately follow from (A.4). Part 3 follows from (A.4) and Lemma 6. \square

A.3. Proof of Proposition 1. We will prove the proposition in several steps.

Step 1. Fix any positive integer m . We will construct a perturbed incomplete information game with an enriched type space and truncated time horizon at m under which each rationalizable action of each original type remains rationalizable for some perturbed type. For each rationalizable action $a_i \in S_i^\infty[t_i]$, let

$$X[a_i, t_i] = \{a'_i \in S_i^\infty[t_i] \mid a'_i \text{ is } m\text{-equivalent to } a_i\}$$

and pick a representative action $r_{t_i}(a_i)$ from each set $X[a_i, t_i]$. We will consider the type space $\tilde{T}^m = \tilde{T}_1^m \times \dots \times \tilde{T}_n^m$ with

$$\tilde{T}_i^m = \{(t_i, r_{t_i}(a_i), m) \mid t_i \in T_i, a_i \in S_i^\infty[t_i]\}.$$

Note that each type here has two dimensions, one corresponding to the original type the second corresponding to an action. Note also that \tilde{T}^m is finite because there are finitely many equivalence classes $X[a_i, t_i]$, allowing only finitely many representative actions $r_{t_i}(a_i)$. Towards defining the beliefs, recall that for each $(t_i, r_{t_i}(a_i), m)$, since $r_{t_i}(a_i) \in S_i^\infty[t_i]$, there exists a belief $\pi^{t_i, r_{t_i}(a_i)} \in \Delta(\Theta \times T_{-i} \times A_{-i})$ under which $r_{t_i}(a_i)$ is a best reply for t_i and $\text{marg}_{\Theta \times T_{-i}}(\pi^{t_i, r_{t_i}(a_i)}) = \kappa_{t_i}$. Define a mapping $\phi_{t_i, r_{t_i}(a_i), m} : \Theta^* \rightarrow \Theta^*$ between the payoff functions by setting

$$(A.5) \quad \phi_{t_i, r_{t_i}(a_i), m}(\theta)(h) = E\left[\theta(h) \mid h^m, r_{t_i}(a_i), \pi^{t_i, r_{t_i}(a_i)}\right]$$

⁸If one writes $t_i = (t_i^1, t_i^2, \dots)$ and $\tau_i^m(t_i) = (\tau_i^{m,1}(t_i^1), \tau_i^{m,2}(t_i^2), \dots)$ as a hierarchies, we define τ_i^m inductively by setting $\tau_i^{m,1}(t_i^1) = t_i^1 \circ \varphi_m^{-1}$ and $\tau_i^{m,k}(t_i^k) = t_i^k \circ (\varphi_m, \tau_{-i}^{m,1}, \dots, \tau_{-i}^{m,k-1})^{-1}$ for $k > 1$.

at each $\theta \in \Theta^*$ and $h \in Z$. Define a joint mapping

$$(A.6) \quad \bar{\phi}_{t_i, r_{t_i}(a_i), m} : (\theta, t_{-i}, a_{-i}) \mapsto \left(\phi_{t_i, r_{t_i}(a_i), m}(\theta), (t_{-i}, r_{t_{-i}}(a_{-i}), m) \right)$$

on tuples for which $a_{-i} \in S_{-i}^\infty[t_{-i}]$. We define the belief of each type $(t_i, r_{t_i}(a_i), m)$ by

$$(A.7) \quad \kappa_{t_i, r_{t_i}(a_i), m} = \pi^{t_i, r_{t_i}(a_i)} \circ \bar{\phi}_{t_i, r_{t_i}(a_i), m}^{-1}.$$

Note that $\kappa_{t_i, r_{t_i}(a_i), m}$ has a natural meaning. Imagine a type t_i who wants to play $r_{t_i}(a_i)$ under a belief $\pi^{t_i, r_{t_i}(a_i)}$ about (θ, t_{-i}, a_{-i}) . Suppose he assumes that payoffs are fixed as if after m the continuation will be according to him playing $r_{t_i}(a_i)$ and the others playing according to what is implied by his belief $\pi^{t_i, r_{t_i}(a_i)}$. Now he considers the outcome paths up to length m in conjunction with (θ, t_{-i}) . His belief is then $\kappa_{t_i, r_{t_i}(a_i), m}$. Let $\tilde{\Theta}^m = \cup_{t_i, r_{t_i}(a_i)} \phi_{t_i, r_{t_i}(a_i), m}(\Theta)$. The perturbed model is $(\tilde{\Theta}^m, \tilde{T}^m, \kappa)$. We write $\tilde{G}^m = (\Gamma, \tilde{\Theta}^m, \tilde{T}^m, \kappa)$ for the resulting Bayesian game, which we will sometimes refer to as a normal-form game.

Step 2. For each t_i and $a_i \in S_i^\infty[t_i]$, the hierarchies $h_i(t_i, r_{t_i}(a_i), m)$ converge to $h_i(t_i)$.

Proof: Let $\tilde{T}^\infty = \bigcup_{m=1}^\infty \tilde{T}^m \cup T$ be a type space with beliefs as in each component of the union, and topology defined by the basic open sets being singletons $\{(t_i, r_{t_i}(a_i), m)\}$ together with sets $\{(t_i, r_{t_i}(a_i), m) : a_i \in S_i^\infty[t_i], m > k\} \cup \{t_i\}$ for each $t_i \in T$ and integer k . That is, the topology is almost discrete, except that there is non-trivial convergence of sequences $(t_i, r_{t_i}(a_i), m) \rightarrow t_i$. Since \tilde{T}^∞ is compact under this topology, Lemma 5 will now give the desired result, once we prove that the map κ from types to beliefs is continuous. This continuity is the substance of the proof – if not for the need to prove this, our definition of the topology would have made the result true by fiat.

At types $(t_i, r_{t_i}(a_i), m)$ the topology is discrete and continuity is trivial, so it suffices to show continuity at types t_i . Since Θ is finite, by continuity at infinity, for any ε we can pick an m such that for all $\theta \in \Theta$, $|\theta_i(h) - \theta_i(\tilde{h})| < \varepsilon$ whenever $h^m = \tilde{h}^m$. Hence, by (A.5),

$$\begin{aligned} \left| \phi_{t_i, r_{t_i}(a_i), m}(\theta)(h) - \theta(h) \right| &= \left| E \left[\theta(\tilde{h}) \mid \tilde{h}^m = h^m, r_{t_i}(a_i), \pi^{t_i, r_{t_i}(a_i)} \right] - \theta(h) \right| \\ &\leq E \left[\left| \theta(\tilde{h}) - \theta(h) \right| \mid \tilde{h}^m = h^m, r_{t_i}(a_i), \pi^{t_i, r_{t_i}(a_i)} \right] < \varepsilon. \end{aligned}$$

Thus, $\phi_{t_i, r_{t_i}(a_i), m}(\theta)(h) \rightarrow \theta(h)$ for each h , showing that $\phi_{t_i, r_{t_i}(a_i), m}(\theta) \rightarrow \theta$. From the definition (A.6) we see that this implies $\bar{\phi}_{t_i, r_{t_i}(a_i), m}(\theta, t_{-i}, a_{-i}) \rightarrow (\theta, t_{-i})$ as $m \rightarrow \infty$. (Recall that $(t_{-i}, r_{t_{-i}}(a_{-i}), m) \rightarrow t_{-i}$.) Therefore, by (A.7), as $m \rightarrow \infty$,

$$\kappa_{t_i, r_{t_i}(a_i), m} \rightarrow \pi^{t_i, r_{t_i}(a_i)} \circ \text{proj}_{\Theta \times T_{-i}}^{-1} = \text{marg}_{\Theta \times T_{-i}}(\pi^{t_i, r_{t_i}(a_i)}) = \kappa_{t_i},$$

which is the desired result.

Step 3. The strategy profile $s^* : \tilde{T}^m \rightarrow A$ with $s_i^*(t_i, r_{t_i}(a_i), m) = r_{t_i}(a_i)$ is a Bayesian Nash equilibrium in \tilde{G}^m .

Proof: Towards defining the belief of a type $(t_i, r_{t_i}(a_i), m)$ under s_{-i}^* , define mapping

$$\gamma : (\theta, t_{-i}, r_{t_{-i}}(a_{-i}), m) \mapsto (\theta, t_{-i}, r_{t_{-i}}(a_{-i}), m, r_{t_{-i}}(a_{-i})),$$

which describes s_{-i}^* . Then, given s_{-i}^* , his beliefs about $\Theta \times \tilde{T}_{-i} \times A_{-i}$ is

$$\pi(\cdot | t_i, r_{t_i}(a_i), m, s_{-i}^*) = \kappa_{t_i, r_{t_i}(a_i), m} \circ \gamma^{-1} = \pi^{t_i, r_{t_i}(a_i)} \circ \bar{\phi}_{t_i, r_{t_i}(a_i), m}^{-1} \circ \gamma^{-1},$$

where the second equality is by (A.7). His induced belief about $\Theta \times A_{-i}$ is

$$\begin{aligned} \text{marg}_{\Theta \times A_{-i}} \pi(\cdot | t_i, r_{t_i}(a_i), m, s_{-i}^*) &= \pi^{t_i, r_{t_i}(a_i)} \circ \bar{\phi}_{t_i, r_{t_i}(a_i), m}^{-1} \circ \gamma^{-1} \circ \text{proj}_{\Theta \times A_{-i}}^{-1} \\ \text{(A.8)} \qquad \qquad \qquad &= \pi^{t_i, r_{t_i}(a_i)} \circ \left(\phi_{t_i, r_{t_i}(a_i), m}, r_{-i} \right)^{-1} \end{aligned}$$

where $r_{-i} : (t_{-i}, a_{-i}) \mapsto r_{t_{-i}}(a_{-i})$. To see this, note that

$$\text{proj}_{\Theta \times A_{-i}} \circ \gamma \circ \bar{\phi}_{t_i, r_{t_i}(a_i), m} : (\theta, t_{-i}, a_{-i}) \mapsto \left(\phi_{t_i, r_{t_i}(a_i), m}(\theta), r_{t_{-i}}(a_{-i}) \right).$$

Now consider any deviation a'_i such that $a'_i(h) = r_{t_i}(a_i)(h)$ for every history longer than m . It suffices to focus on such deviations because the moves after length m are payoff-irrelevant under $\tilde{\Theta}^m$ by (A.5). The expected payoff vector from any such a'_i is

$$\begin{aligned} E \left[u(\theta, a'_i, s_{-i}^*) | \kappa_{t_i, r_{t_i}(a_i), m} \right] &= E \left[u \left(\phi_{t_i, r_{t_i}(a_i), m}(\theta), a'_i, r_{t_{-i}}(a_{-i}) \right) | \pi^{t_i, r_{t_i}(a_i)} \right] \\ &= E \left[\phi_{t_i, r_{t_i}(a_i), m}(\theta) (z(a'_i, r_{t_{-i}}(a_{-i}))) | \pi^{t_i, r_{t_i}(a_i)} \right] \\ &= E \left[E \left[\theta (z(a'_i, r_{t_{-i}}(a_{-i}))) | z(a'_i, r_{t_{-i}}(a_{-i}))^m, r_{t_i}(a_i), \pi^{t_i, r_{t_i}(a_i)} \right] | \pi^{t_i, r_{t_i}(a_i)} \right] \\ &= E \left[E \left[\theta (z(a'_i, r_{t_{-i}}(a_{-i}))) | z(a'_i, r_{t_{-i}}(a_{-i}))^m, a'_i, \pi^{t_i, r_{t_i}(a_i)} \right] | \pi^{t_i, r_{t_i}(a_i)} \right] \\ &= E \left[\theta (z(a'_i, r_{t_{-i}}(a_{-i}))) | \pi^{t_i, r_{t_i}(a_i)} \right], \end{aligned}$$

where the first equality is by (A.8); the second equality is by definition of u ; the third equality is by definition of $\phi_{t_i, r_{t_i}(a_i), m}$, which is (A.5); the fourth equality is by the fact that a'_i is equal to $r_{t_i}(a_i)$ conditional on history $z(a'_i, r_{t_{-i}}(a_{-i}))^m$, and the fifth equality is by the law of iterated

expectations. Hence, for any such a'_i ,

$$\begin{aligned} E \left[u_i \left(\theta, r_{t_i}(a_i), s_{-i}^* \right) \mid \kappa_{t_i, r_{t_i}(a_i), m} \right] &= E \left[\theta_i \left(z \left(r_{t_i}(a_i), r_{t_{-i}}(a_{-i}) \right) \mid \pi^{t_i, r_{t_i}(a_i)} \right) \right] \\ &\geq E \left[\theta_i \left(z \left(a'_i, r_{t_{-i}}(a_{-i}) \right) \mid \pi^{t_i, r_{t_i}(a_i)} \right) \right] \\ &= E \left[u_i \left(\theta, a'_i, s_{-i}^* \right) \mid \kappa_{t_i, r_{t_i}(a_i), m} \right], \end{aligned}$$

where the inequality is by the fact that $r_{t_i}(a_i)$ is a best reply to $\pi^{t_i, r_{t_i}(a_i)}$, by definition of $\pi^{t_i, r_{t_i}(a_i)}$. Therefore, $r_{t_i}(a_i)$ is a best reply for type $(t_i, r_{t_i}(a_i), m)$, and hence s^* is a Bayesian Nash equilibrium.

Step 4. Referring back to the statement of the proposition, by Step 2, pick m , t_i , and a_i such that $m > L$ and $h_i((t_i, r_{t_i}(a_i), m)) \in U_i$. By Step 3, a_i is rationalizable for type $(t_i, r_{t_i}(a_i), m)$.

Proof: Since $h_i((t_i, r_{t_i}(a_i), m)) \rightarrow h_i(t_i)$ and U_i is an open neighborhood of t_i , $h_i((t_i, r_{t_i}(a_i), m)) \in U_i$ for sufficiently large m . Hence, we can pick m as in the statement. Moreover, by Step 3, $r_{t_i}(a_i)$ is rationalizable for type $(t_i, r_{t_i}(a_i), m)$ (because it is played in an equilibrium). This implies also that a_i is rationalizable for type $(t_i, r_{t_i}(a_i), m)$, because m -equivalent actions are payoff-equivalent for type $(t_i, r_{t_i}(a_i), m)$.

The remaining steps will show that a further perturbation makes a_i uniquely rationalizable.

Step 5. Define hierarchy $h_i(\tilde{t}_i) \in T_i^{*m}$ for the finite-horizon game form Γ^m by

$$h_i(\tilde{t}_i) = (\tau_i^m)^{-1} (h_i((t_i, r_{t_i}(a_i), m))),$$

where τ_i^m is as defined in Lemma 7 of Section A.2. Type \tilde{t}_i comes from a finite game $G^m = (\Gamma^m, \Theta^m, T^m, \kappa)$ and $a_i^m \in S_i^\infty[\tilde{t}_i]$.

Proof: By Lemma 7, since type $(t_i, r_{t_i}(a_i), m)$ is from a finite model, so is \tilde{t}_i . Since a_i is rationalizable for type $(t_i, r_{t_i}(a_i), m)$, by Lemma 7, a_i^m is rationalizable for $h_i(\tilde{t}_i)$ and hence for type \tilde{t}_i in G^m .

Step 6. By Step 5 and Lemma 1, there exists a hierarchy $h_i(\bar{t}_i^m)$ in open neighborhood $(\tau_i^m)^{-1}(U_i)$ of $h_i(\tilde{t}_i)$ such that each element of $S_i^\infty[\bar{t}_i^m]$ is m -equivalent to a_i^m , and \bar{t}_i^m is a type in a finite, common-prior model.

Proof: By the definition of $h_i(\tilde{t}_i)$ in Step 5, $h_i(\tilde{t}_i) \in (\tau_i^m)^{-1}(U_i)$. Since U_i is open and τ_i^m is continuous, $(\tau_i^m)^{-1}(U_i)$ is open. Moreover, \tilde{t}_i comes from a finite game, and a_i^m is rationalizable for \tilde{t}_i . Therefore, by Lemma 1, there exists a hierarchy $h_i(\bar{t}_i^m)$ in $(\tau_i^m)^{-1}(U_i)$ as in the statement above.

Step 7. Define the hierarchy $h_i(\hat{t}_i)$ by

$$h_i(\hat{t}_i) = \tau_i^m(h_i(\bar{t}_i^m)).$$

The conclusion of the proposition is satisfied by \hat{t}_i .

Proof: Since $h_i(\bar{t}_i^m) \in (\tau_i^m)^{-1}(U_i)$,

$$h_i(\hat{t}_i) = \tau_i^m(h_i(\bar{t}_i^m)) \in \tau_i^m\left((\tau_i^m)^{-1}(U_i)\right) \subseteq U_i.$$

Since \bar{t}_i^m is a type from a finite, common-prior model, by Lemma 7, \hat{t}_i can also be picked from a finite, common-prior model. Finally, take any $\hat{a}_i \in S_i^\infty[\hat{t}_i]$. By Lemma 7, $\hat{a}_i^m \in S_i^\infty[\hat{t}_i]$. Hence, by Step 6, \hat{a}_i^m is m -equivalent to a_i^m . It then follows that \hat{a}_i is m -equivalent to a_i . Since $m > L$, \hat{a}_i is also L -equivalent to a_i .

APPENDIX B. PROOF OF PROPOSITION 2

Using Proposition 1, we first establish that every action can be made rationalizable for some type. This extends the lemma of Chen from equivalence at histories of bounded length to equivalence at histories of unbounded length.

Lemma 8. *For all plans of action a_i , there is a type t^{a_i} of player i such that a_i is the unique rationalizable action for t^{a_i} , up to reduced-form equivalence.*

Proof. The set of non-terminal histories is countable, as each of them has finite length. Index the set of histories where it is i 's move and the history thus far is consistent with a_i as $\{h(k) : k \in Z^+\}$. By Proposition 1, for each k there is a type t_{-i}^k whose rationalizable actions are always consistent with history $h(k)$. We construct type t^{a_i} as follows: his belief about t_{-i} assigns probability 2^{-k} to type t_{-i}^k . His belief about θ is a point-mass on the function θ_{a_i} , defined as 1 if all of i 's actions were consistent with a_i and $1 - 2^{-k}$ if his first inconsistent move was at history $h(k)$. Now, if type t^{a_i} plays action a_i he receives a certain payoff of 1. If his plan b_i is not reduced-form equivalent to a_i , let $h(k)$ be the shortest history in the set $\{h(k) : k \in Z^+\}$ where $b_i(h(k)) \neq a_i(h(k))$. By construction, there is probability at least 2^{-k} of reaching this history if he believes the other player's action is rationalizable, so his expected payoff is at most $1 - 2^{-2k}$. This completes the proof. \square

Proof of Proposition 2. We first show that (A) implies (B). Assume that s^* is a Bayesian Nash equilibrium of G . Construct a family of types $\tau_j(t_j, m, \lambda)$, $j \in N$, $t_j \in T_j$, $m \in \mathbb{N}$, $\lambda \in [0, 1]$, as

follows

$$\begin{aligned}\tau_j(t_j, 0, \lambda) &= t^{s_j^*(t_j)}, \\ \kappa_{\tau_j(t_j, m, \lambda)} &= \lambda \kappa_{t^{s_j^*(t_j)}} + (1 - \lambda) \beta_{t_j, m, \lambda} \quad \forall m > 0\end{aligned}$$

where

$$\beta_{t_j, m, \lambda}(\theta, \tau_{-j}(t_{-j}, m - 1, \lambda)) = \kappa_{t_j}(\theta, t_{-j}) \quad \forall (\theta, t_{-j}) \in \Theta \times T_{-j}.$$

For large m and small λ , $\tau_j(t_j, m, \lambda)$ satisfies all the properties of \hat{t}_i , as we establish below.

Now, we use mathematical induction on m to show that for all $\lambda > 0$ and for all m and t_j , $a_j \in S_j^\infty[\tau_j(t_j, m, \lambda)]$ if and only if a_j is equivalent to $s_j^*(t_j)$, establishing the first conclusion in (B). This statement is true for $m = 0$ by definition of $\tau_j(t_j, 0, \lambda)$ and Lemma 8. Now assume that it is true up to some $m - 1$. Consider any rationalizable belief of any type $\tau_j(t_j, m, \lambda)$. With probability λ , his belief is the same as that of $t^{s_j^*(t_j)}$. By definition, $s_j^*(t_j)$ is the unique best response to this belief in reduced form actions. With probability $1 - \lambda$, his belief on $\Theta^* \times A_{-j}$ is the same as the equilibrium belief of t_j on $\Theta^* \times A_{-j}$. The action $s_j^*(t_j)$ is also a best reply to this belief because s^* is a Bayesian Nash equilibrium in the original game. Therefore, $s_j^*(t_j)$ is the unique best response to the rationalizable belief of type $\tau_j(t_j, m, \lambda)$ in reduced form. Since type $\tau_j(t_j, m, \lambda)$ and his rationalizable belief are picked arbitrarily, this proves the statement.

Note that by the preceding paragraph, for any $\lambda > 0$ and $m > 0$, $\tau_j(t_j, m, \lambda)$ has a unique rationalizable belief

$$\pi(t_j, m, \lambda) = \kappa_{\tau_j(t_j, m, \lambda)} \circ \gamma_{j, m, \lambda}^{-1}$$

where

$$\gamma_{j, m, \lambda} : (\theta, h_{-j}(t_{-j}, m, \lambda)) \mapsto (\theta, h_{-j}(t_{-j}, m, \lambda), s_{-j}^*(t_{-j})).$$

Here, the mapping $\gamma_{j, m, \lambda}$ corresponds to the fact that the newly constructed types play according to the equilibrium strategy of the original types. We leave the actions of the other types unassigned as their actions are not relevant for our proof. For $\lambda = 0$, we define $\pi(t_j, m, \lambda)$ by the same equation, although the type $\tau_j(t_j, m, \lambda)$ may also have other rationalizable beliefs.

In order to show that for large m and small λ , the beliefs of $\tau_j(t_j, m, \lambda)$ are as in the proposition, note that for $\lambda = 0$, the m th-order belief of $\tau_j(t_j, m, 0)$ is equal to the m th-order belief of t_j . Hence, as $m \rightarrow \infty$, $h_j(\tau_j(t_j, m, 0)) \rightarrow h_j(t_j)$ for each j . Consequently, for each j , as $m \rightarrow \infty$, $\pi(t_j, m, 0)$ converges to

$$\pi_{t_j}^* = \kappa_{t_j} \circ (\gamma_j^*)^{-1} \quad \text{with } \gamma_j^* : (\theta, t_{-j}) \mapsto (\theta, t_{-j}, s_{-j}^*(t_{-j})).$$

Note that $\pi_{t_j}^*$ is the equilibrium belief of type t_j under s^* . Therefore, there exists $\bar{m} > 0$ such that $h_i(\tau_i(t_i, \bar{m}, 0)) \in U_i$ and $\pi(t_i, m, 0) \in V_i$. Moreover, for $j \in N$, $m \leq \bar{m}$, and $\lambda \in [0, 1]$, beliefs of

$\tau_j(t_j, m, 0)$ are continuous in λ . Hence, by Lemma 5,⁹ for each t_j , as $\lambda \rightarrow 0$, $h_j(\tau_j(t_j, \bar{m}, \lambda)) \rightarrow h_j(\tau_j(t_j, \bar{m}, 0))$ and (thereby) $\pi(t_j, m, \lambda) \rightarrow \pi(t_j, m, 0)$. Thus, there exists $\bar{\lambda} > 0$ such that $h_i(\tau_i(t_i, \bar{m}, \bar{\lambda})) \in U_i$ and $\pi(t_i, m, \bar{\lambda}) \in V_i$. Therefore, the type $\hat{t}_i = \tau_i(t_i, \bar{m}, \bar{\lambda})$ satisfies all the properties in (B).

In order to show the converse (i.e. that (B) implies (A)), take any type t_i and assume (B). Then, there exists a sequence of types $\hat{t}_i(m)$ with unique rationalizable beliefs $\hat{\pi}_m \in \Delta(\Theta^* \times T_{-i}^* \times A_{-i})$ and unique rationalizable action $s_i^*(t_i)$ where $\hat{\pi}_m$ converges to the belief $\pi_{t_i}^*$ of type t_i under s^* . Since $s_i^*(t_i) \in S_i^\infty[\hat{t}_i(m)]$, $s_i^*(t_i) \in BR(\text{marg}_{\Theta^* \times A_{-i}} \hat{\pi}_m)$ for each m . Since u_i is continuous and $\hat{\pi}_m \rightarrow \pi_{t_i}^*$, together with the Maximum Theorem, this implies that $s_i^*(t_i) \in BR(\text{marg}_{\Theta^* \times A_{-i}} \pi_{t_i}^*)$, showing that $s_i^*(t_i)$ is a best reply to s_{-i}^* for type t_i . Since t_i is arbitrary, this proves that s^* is a Bayesian Nash equilibrium. \square

APPENDIX C. PROOF OF PROPOSITION 4

In our proof we will need a couple of lemmas and definitions. First, we write $T_i^{CK(\Theta_\delta^*)}$ for the set of types of player i according to which it is common knowledge that $\theta \in \Theta_\delta^*$. Second, we consider the following action plans.

Definition 2. A plan a_i in reduced form is said to be *Bayes consistent* if and only if it never happens that for a partial history h and move $b_i \in B_i$, $a_i(h, (a_i(h), b_{-i})) = b_i$ for every b_{-i} but $a_i(h) \neq b_i$.

This concept is important in the next lemma because the construction we use in the proof is based on information received by player i rather than punishments and rewards. Consequently, he must follow this form of the sure-thing principle for the construction to work.

Lemma 9. *For any δ , any L and any Bayes consistent action plan a_i , there exists a type $t^{ai,L} \in T_i^{CK(\Theta_\delta^*)}$ for which playing according to a_i until L is uniquely rationalizable in reduced form.*

Proof. We will induct on L . The result is vacuous for $L = -1$. Fix L, a_i and assume the result is true for $L - 1$. In outline, the type we construct will have payoffs which are completely insensitive to the actions of the other players, but will find those actions informative about his own payoffs. He also will believe that if he ever deviates from a_i , the other players' subsequent actions are

⁹To ensure compactness, put all of the types in construction of types $t^{s_j^*}(t_j)$ together and for $\tau(t_j, m, \lambda)$ with $t_j \in T_j$, $j \in N$, $m \in \{0, 1, \dots, \bar{m}\}$, $\lambda \in [0, 1]$, use the usual topology for (t_j, m, λ) .

uninformative. Let $n = |B_{-i}|$ be the number of static moves for the other players. Let \hat{H} be the set histories of length $L - 1$ in which player i always follows the plan a_i , so that $|\hat{H}| = n^{L-1}$. In our construction $t^{a_i, L}$ will assign equal weight to each of n^{L-1} pairs (t_h^{-i}, θ_h) , one for each history in \hat{H} , where t_h^{-i} is assumed to play according to this history so long as i follows a_i , and to simply repeat his last move through period $L - 1$ if player i deviates. (Note that this plan is Bayes consistent because at every history there is at least one branch where the current action is repeated.) We will define each θ_h via an iterative process where we describe the average payoff function of player i conditional on reaching each node, starting at time 0. In particular, we will define a payoff $f : \hat{H} \times B_i \rightarrow \mathbb{R}$, representing i 's expected value of his stage-game payoffs conditional on reaching that history.

To construct $t^{a_i, L}$, we must first construct f , as follows: Fix $\varepsilon > 0$. Let $f(\emptyset, a_i(\emptyset)) = 1$ and $f(\emptyset, b) = 0$ for all $b \neq a_i(\emptyset)$, where \emptyset is the initial node. Next, assume $f(h, \cdot)$ has been defined and proceed for the relevant one-step continuations of h as follows:

Case 1: If $a_i(h, (a_i(h), b_{-i})) = a_i(h)$ for all b_{-i} , then let $f((h, b), \cdot) = f(h, \cdot)$ for every b .

Case 2: Otherwise, by Bayes consistency, at least two different actions are prescribed for continuations $(h, (a_i(h), b_{-i}))$. For each action $b_i \in B_i$, let $S_{b_i} = \{b_{-i} : a_i(h, (a_i(h), b_{-i})) = b_i\}$ be the set of continuations where b_i is prescribed. Then let

$$f((h, (a_i(h), b_{-i})), b_i) = \begin{cases} f(h, a_i(h)) + \varepsilon & \text{if } b_{-i} \in S_{b_i} \\ \frac{nf(h, b_i) - |S_{b_i}|(f(h, a_i(h)) + \varepsilon)}{n - |S_{b_i}|} & \text{if } b_{-i} \notin S_{b_i} \end{cases}$$

where the last denominator is non-zero by the observation that at least two different actions are prescribed.

These payoffs are chosen so that

$$(C.1) \quad f(h, b_i) = \frac{1}{n} \sum_{b_{-i}} f((h, (a_i(h), b_{-i})), b_i)$$

and so that $f(h, a_i(h)) \geq f(h, b_i) + \varepsilon$ for every h and $b_i \neq a_i(h)$. Define $g_h(b) = f(h, b_i)$ for each history h of length $L - 1$ and define θ_h accordingly, as in (4.1). Let $t^{a_i, L}$ assign equal weight to each of n^{L-1} pairs (t_h^{-i}, θ_h) , one for each history h of length $L - 1$ which is consistent with a_i , where the types t_h^{-i} are assumed (by induction), under rationalizable play, to always play consistently with h through stage $L - 1$ when a_i is followed, or to repeat their last move through stage $L - 1$ if not. We claim that under rationalizable play, from the perspective of type $t^{a_i, L}$, when he has followed a_i and reaches history h , $f(h, \cdot)$ is the expected value of the stage-game payoff g . This is true by definition for length- $L - 1$ histories. Since type $t^{a_i, L}$ always thinks his opponents' actions

are distributed uniformly over b_{-i} , the recursive relation (C.1) implies backwards-inductively that the claim is true. Note also that if he follows a_i through period L , player i always learns his true payoff. Let \bar{a}_i be the plan which follows a_i through period L , then plays the known optimal action from period $L + 1$ onward. We claim that \bar{a}_i strictly outperforms any plan which deviates by period L . The intuitive argument is as follows. Because type $t^{a_i, L}$ has stage-game payoffs which are insensitive to the other players' moves, he only has two possible incentives at each stage: to maximize his average stage-game payoffs at the current stage, and to receive further information about his payoffs. The former goal is strictly satisfied by the move prescribed by \bar{a}_i , and the latter is at least weakly satisfied by this move, since after a deviation he receives no further information.

Formally, we must show that for any fixed plan a'_i not L -equivalent to a_i and any rationalizable belief of $t^{a_i, L}$, the plan \bar{a}_i gives a better expected payoff. Given a rationalizable belief on opponents' actions, any initial deviation at or before L is reached with positive probability. Let h be a random variable equal to the earliest realized history at which a'_i differs from a_i , or ∞ if they do not differ by period L . Conditional on any non-infinite value of h , \bar{a}_i outperforms a'_i on average. In fact this is weakly true stage-by-stage, and strictly true at the first deviation, because:

At stage $|h| + 1$: The average payoff $f(h, b_i)$ is strictly optimized by $\bar{a}_i(h)$.

At stages $|h| + 2, \dots, L$: The plan \bar{a}_i optimizes stage-game payoffs relative to its information, which comes from a finer information partition than that available under plan a'_i (because the opponents' play is uninformative subsequent to a deviation.) Hence, even if it plays conditionally optimally, a'_i will never perform better on average than \bar{a}_i .

At stages $L + 1, \dots$: Under plan \bar{a}_i , player i now has complete information about his payoff and optimizes perfectly, so a'_i cannot do better.

If $h = \infty$, again \bar{a}_i cannot be outperformed because he optimizes based on complete information after L .

Finally, since there are only finitely many histories and types in the construction, the payoffs are bounded and so can be normalized to lie in $[0, 1]$. \square

This lemma implies the following result on repeated games.

Lemma 10. *For any $\delta \in (0, 1)$, any L and any action plan a_i , there exists a type $t_i^{a_i, L} \in T_i^{CK(\Theta_\delta^*)}$ for which playing according to a_i until L is uniquely rationalizable in reduced form.*

Proof. For some $b_{-i}^* \in B_{-i}$, consider a stage payoff function g_i with $g_i(b_i, b_{-i}^*) = 1$ and $g_i(b_i, b_{-i}) = 0$ for all $b_{-i} \neq b_{-i}^*$. Note that player i 's payoff does not depend on his own action, and the other

players may reward him by playing b_{-i}^* . Write $\hat{\theta} \in \Theta_\delta^*$ for a payoff function resulting from g_i , i.e., $\hat{\theta}_i(h) = (1 - \delta) \sum_l \delta^l g_i(b^l)$. Fix a large M with $\delta^M < \delta^L (|B_i| - 1) / (2|B_i| - 1)$. Let \hat{A}_{-i} be the set of action profiles a_{-i} such that

- (1) for any $l \leq L + 1$, any b_{-i} , and any h^{l-1} , there exists a unique $b_i [h^{l-1}, b_{-i}] \in B_i$ such that $a_{-i}(h^{l-1}, (b_i, b_{-i})) = b_{-i}^*$ if $b_i = b_i [h^{l-1}, b_{-i}]$ and $a_j(h^{l-1}, (b_i, b_{-i})) \neq b_j^*$ for every $j \neq i$ otherwise;¹⁰
- (2) $b_i [h^L, b_{-i}] = a_i(h^L)$ if player i has played according to a_i throughout the history h^L , and
- (3) for any $l \in \{L + 2, \dots, M\}$ and any h at the beginning of l , $a_{-i}(h) = a_{-i}(h^{L+1})$.

Note that the other players reward a unique move of i at any history, with the only restriction that player i is rewarded for sure at dates $L + 1, \dots, M$ if he sticks to a_i throughout $l = 0, \dots, L$. This implies that if player i sticks to a_i up to $L - 1$ and deviates at L , then he will not be rewarded at dates $L + 1, \dots, M$. This is the only restriction, and the set \hat{A}_{-i} is symmetric in all other ways. Note also that at any $l \leq M$, a player j either reacts differently to different moves of player i or repeats his previous move regardless. Hence, the actions in \hat{A}_{-i} are all Bayes-consistent up to date M , and thus for each $a_{-i} \in \hat{A}_{-i}$, there exists a Bayes-consistent action \hat{a}_{-i} that is M -equivalent to a_{-i} . Let \hat{A}_{-i}^M be a finite subset of \hat{A}_{-i} that consists of one Bayes-consistent element from each M -equivalence class in \hat{A}_{-i} . By Lemma 9, for each $a_{-i} \in \hat{A}_{-i}^M$, there exists $t^{a_{-i}, M}$ for which all rationalizable action profiles are M -equivalent to a_{-i} . Consider a type $t_i^{a_i, L}$ that assigns probability $1/|\hat{A}_{-i}^m|$ to each $(\hat{\theta}, t^{a_{-i}, M})$ with $a_{-i} \in \hat{A}_{-i}^m$. Note that, according to $t_i^{a_i, L}$ the rewarded actions up to $l = L - 1$ are independently and identically distributed with uniform distribution over his moves. This leads to the formulas for the probability of reward in the next paragraph.

For any history h at the beginning of any date l , write $P_l^*(h)$ for the probability that b_{-i}^* is played at t conditional on h according to the rationalizable belief of $t_i^{a_i, L}$. As noted above, by symmetry,

$$(C.2) \quad P_l^*(h) = 1/|B_i| \quad \forall l \leq L,$$

and

$$(C.3) \quad P_{L+1}^*(h) = \begin{cases} 1 & \text{if } i \text{ follows } a_i \text{ until } L \\ 0 & \text{if } i \text{ follows } a_i \text{ until } L - 1 \text{ and deviates at } L \\ 1/|B_i| & \text{otherwise.} \end{cases}$$

¹⁰Note that h^{l-1} is the list of moves played at dates $0, 1, \dots, l - 2$, and $a_j(h^{l-1}, b)$ is the move of player j at date l if players play b at $l - 1$ after history h^{l-1} .

Note that the expected payoff of type $t_i^{a_i, L}$ under any action a'_i is

$$(C.4) \quad U_i(a'_i) = \sum_l (1 - \delta) \delta^l E [P_l^* | a'_i].$$

Using the above formulas, we will now show that type $t_i^{a_i, L}$ does not have a best response that differs from a_i at some history of length $l \leq L$. Consider such an action plan a'_i . Define also a_i^* , by setting

$$a_i^*(h^l) = \begin{cases} a_i(h^l) & \text{if } l \leq L \\ a'_i(h^l) & \text{if } l > L \end{cases}$$

at each history h^l at the beginning of l . We will show that a_i^* yields a higher expected payoff than a'_i . To this end, for each history h , define $\tau(h)$ as the smallest date l such that the play of player i is in accordance with both a_i and a'_i throughout history h^l , $a_i(h^l) \neq a'_i(h^l)$, and player i plays $a'_i(h^l)$ at date l according to h . (Here, τ can be infinite.) Note also that, by (C.2) and (C.3), a_i^* always yields

$$U_i(a_i^*) = (1 - \delta^{L+1}) / |B_i| + (\delta^{L+1} - \delta^{M+1}) \cdot 1 + \sum_{l > M} (1 - \delta) \delta^l E [P_l^* | a_i^*].$$

On the event $\tau > L$, a_i^* and a'_i are identical and hence yield the same payoff. On the event $\tau = L$, by (C.2) and (C.3), a'_i yields the payoff of

$$U_i(a'_i | \tau = L) = (1 - \delta^{L+1}) / |B_i| + \sum_{l > M} (1 - \delta) \delta^l E [P_l^* | a'_i, \tau = L].$$

On the event $\tau < L$, by (C.2) and (C.3), a'_i yields the payoff of

$$U_i(a'_i | \tau < L) = (1 - \delta^{L+1}) / |B_i| + (\delta^{L+1} - \delta^{M+1}) \cdot 1 / |B_i| + \sum_{l > M} (1 - \delta) \delta^l E [P_l^* | a'_i, \tau < L].$$

Hence,

$$\begin{aligned} U_i(a_i^*) - U_i(a'_i) &\geq \Pr(\tau \leq L) \left[(\delta^{L+1} - \delta^{M+1}) (1 - 1/|B_i|) + \sum_{l > M} (1 - \delta) \delta^l [E [P_l^* | a_i^*, \tau \leq L] - E [P_l^* | a'_i, \tau \leq L]] \right] \\ &\geq \Pr(\tau \leq L) [(\delta^{L+1} - \delta^{M+1}) (1 - 1/|B_i|) - \delta^{M+1}] > 0, \end{aligned}$$

where the first inequality follows from the previous three displayed equations, the next inequality is by the fact that $P_l^* \in [0, 1]$, and the strict inequality follows from the fact that $\Pr(\tau \leq L) > 0$ by definition (as $t_i^{a_i, L}$ puts positive probability at all histories up to date L and a'_i differs from a_i at some such history) and the fact that $\delta^M < \delta^L (|B_i| - 1) / (2|B_i| - 1)$. \square

This lemma establishes that any action can be made uniquely rationalizable for arbitrarily long horizon even within the restricted class of repeated game payoffs with the given discount factor δ . Using this lemma, we will next prove Proposition 4.

Proof of Proposition 4. First, note that by continuity at infinity there exist $\bar{\lambda} \in (0, 1)$ and $l^* < \infty$ such that if a player i assigns at least probability $1 - \bar{\lambda}$ on the event that $\theta = \theta_{\delta, g^*}$ and everybody follows a^* up to date l^* , then the expected payoff vector under his belief will be within ε neighborhood of $u(\theta_{\delta, g^*}, a^*)$.

We construct a family of types $t_{j, m, l, \lambda}$, $j \in N$, $m, l \in \mathbb{N}$, $\lambda \in [0, \bar{\lambda}]$, by

$$\begin{aligned} t_{j, 0, l, \lambda} &= t_j^{a_j^*, l}, \\ \kappa_{t_{j, m, l, \lambda}} &= \lambda \kappa_{t_j^{a_j^*, l}} + (1 - \lambda) \delta_{(\theta_{\delta, g^*}, t_{-i, m-1, l', \lambda})} \quad \forall m > 0, \end{aligned}$$

where $t_j^{a_j^*, l} \in T_j^{CK}(\Theta_\delta^*)$ is the type for whom a_j^* is uniquely rationalizable up to date l , $\delta_{(\theta_{\delta, g^*}, t_{-i, m-1, l', \lambda})}$ is the Dirac measure that puts probability one on $(\theta_{\delta, g^*}, t_{-i, m-1, l', \lambda})$ and l' will be defined momentarily. The types $t_{j, m, l, \lambda}$ will be constructed in such a way that under any rationalizable plan they will follow a_j^* up to date l and the first m orders of beliefs will be within λ neighborhood of $t^{CK}(\theta_{\delta, g^*})$. Note that under $\kappa_{t_j^{a_j^*, l}}$ it is a unique best reply to follow a_j^* up to date l , and if the other players follow a_{-j}^* forever then it is a best response under θ_{δ, g^*} to follow a_j^* up to date l . In that case, it would be a unique best response to follow a_j^* up to date l if one puts probability λ on $\kappa_{t_j^{a_j^*, l}}$ and $(1 - \lambda)$ on the latter scenario with $\theta = \theta_{\delta, g^*}$. Since there are only finitely many plans to follow up to date l and the game is continuous at infinity, there exists a finite $l' \geq l^*$ such that it is still the unique best response under θ_{δ, g^*} to follow a_j^* up to date l if the other players played a_{-j}^* only up to date l' . We pick such an $l' \geq l^*$.

We now show that for large m and l and small λ , $t_{i, m, l, \lambda}$ satisfies all the desired properties of \hat{t}_i . First note that for $\lambda = 0$, under $t_{i, m, l, 0}$, it is m th-order mutual knowledge that $\theta = \theta_{\delta, g^*}$. Hence, there exist m^* and $\lambda^* > 0$ such that when $m \geq m^*$ and $\lambda \leq \lambda^*$, the belief hierarchy of $t_{i, m, l, \lambda}$ is within the neighborhood U_i of the belief hierarchy of $t_i^{CK}(\theta_{\delta, g^*})$, according to which it is common knowledge that $\theta = \theta_{\delta, g^*}$. Second, for $\lambda > 0$, a_j^* is uniquely rationalizable up to date l for $t_{j, m, l, \lambda}$ in reduced form. To see this, observing that it is true for $m = 0$ by definition of $t_{j, 0, l, \lambda}$, assume that it is true up to some $m - 1$. Then, any rationalizable belief of any type $t_{j, m, l, \lambda}$ must be a mixture of two beliefs. With probability λ , his belief is the same as that of $t_j^{a_j^*, l}$, and with probability $1 - \lambda$, he believes that the true state is θ_{δ, g^*} and the other players play a_{-j}^* (in reduced form) up to date l' . But we have chosen l' so that following a_j^* up to date l is a unique best response under that belief. Therefore, any rationalizable action of $t_{j, m, l, \lambda}$ is l -equivalent to a_j^* . Third, for any $m > 0$ and $l \geq l^*$, the expected payoffs are within ε neighborhood of $u(\theta_{\delta, g^*}, a^*)$. Indeed, under rationalizability, type $t_{i, m, l, \lambda}$ must assign at least probability $1 - \lambda \geq 1 - \bar{\lambda}$ on $\theta = \theta_{\delta, g^*}$ and that the other players follow a_{-i}^* up to date $l' \geq l^*$ while he himself follows a_i^* up to date $l \geq l^*$. The expected payoff vector is ε neighborhood of $u(\theta_{\delta, g^*}, a^*)$ under such a belief by definition of $\bar{\lambda}$ and l^* .

Finally, each $t_{j,m,l,\lambda}$ is in $T_j^{CK(\Theta_\delta^*)}$ because all possible types in the construction assigns probability 1 on $\theta \in \Theta_\delta^*$. We complete our proof by picking $\hat{t}_i = t_{i,m,l,\lambda}$ for some $m > m^*$, $l \geq \max\{L, l^*\}$, and $\lambda \in (0, \min\{\bar{\lambda}, \lambda^*\})$. \square

REFERENCES

- [1] Abreu, D. and F. Gul (2000): “Bargaining and Reputation” *Econometrica* 68, 85-117.
- [2] Battigalli, P. and M. Siniscalchi (2003): “Rationalization and Incomplete Information,” *Advances in Theoretical Economics* 3-1, Article 3.
- [3] Brandenburger, A. and E. Dekel (1993): “Hierarchies of Beliefs and Common Knowledge,” *Journal of Economic Theory*, 59, 189-198.
- [4] Chassang, S. and S. Takahashi (20011): “Robustness to Incomplete Information in Repeated Games,” *Theoretical Economics* 6, 49-93.
- [5] Chen, Y. (2008): “A Structure Theorem for Rationalizability in Dynamic Games”, Northwestern University Working Paper.
- [6] Dekel, E. D. Fudenberg, S. Morris (2007): “Interim Correlated Rationalizability,” *Theoretical Economics*, 2, 15-40.
- [7] Feinberg, Y. and A. Skrzypacz (2005) “Uncertainty about Uncertainty and Delay in Bargaining” *Econometrica* 73, 69-91.
- [8] Fudenberg, D. and D. Levine (1983): “Subgame-Perfect Equilibria of Finite and Infinite Horizon Games,” *Journal of Economic Theory* 31, 251-268.
- [9] Fudenberg, D. and E. Maskin (1991): “On the Dispensability of Public Randomization in Discounted Repeated Games,” *Journal of Economic Theory* 53, 428-438.
- [10] Kajii, A. and S. Morris (1997): “The Robustness of Equilibria to Incomplete Information,” *Econometrica*, 65, 1283-1309.
- [11] Kreps, D., P. Milgrom, J. Roberts and R. Wilson (1982): “Rational Cooperation in the Finitely-Repeated Prisoners’ Dilemma,” *Journal of Economic Theory*, 27, 245-52.
- [12] Mailath, G. and L. Samuelson (2006): *Repeated Games and Reputations*, Oxford University Press.
- [13] Mertens and Zamir (1985): “Formulation of Bayesian Analysis for Games with Incomplete Information,” *International Journal of Game Theory*, 10, 619-632.
- [14] Oury, M. and O. Tercieux (2007): “Continuous Implementation,” PSE Working Paper.
- [15] Penta, A. (2008): “Higher Order Beliefs in Dynamic Environments,” University of Pennsylvania Working Paper.
- [16] Rubinstein, A. (1982): “Perfect Equilibrium in a Bargaining Model,” *Econometrica* 50, 97-109.
- [17] Rubinstein, A. (1989): “The Electronic Mail Game: Strategic Behavior Under ‘Almost Common Knowledge,’” *The American Economic Review*, Vol. 79, No. 3, 385-391.

- [18] Weinstein, J. and M. Yildiz (2007): “A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinements,” *Econometrica*, 75, 365-400.
- [19] Weinstein, J. and M. Yildiz (2011): “Sensitivity of Equilibrium Behavior to Higher-order Beliefs in Nice Games,” *Games and Economic Behavior* 72, 288–300.

WEINSTEIN: NORTHWESTERN UNIVERSITY; YILDIZ: MIT