

ROBUST PREDICTIONS IN INFINITE-HORIZON GAMES—AN UNREFINABLE FOLK THEOREM

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 played by a type t_i in a Bayesian Nash equilibrium, then there are perturbations of t_i for which a_i is the *only* rationalizable plan and whose unique rationalizable belief regarding the play of the game is arbitrarily close to the equilibrium belief of t_i . As an application to repeated games, we prove an unrefinable folk theorem: Any individually rational and feasible payoff is the unique rationalizable payoff vector for some perturbed type profile. This is true even if perturbed types are restricted to believe that the repeated-game payoff structure and the discount factor are common knowledge.

JEL Numbers: C72, C73.

1. INTRODUCTION

In the infinite-horizon dynamic games commonly used in economic applications, the set of equilibrium strategies is often very large. For example, the classic folk theorems for repeated games state that every individually-rational payoff profile can be achieved in a subgame-perfect equilibrium. A less transparent example is Rubinstein's (1982) bargaining game; although there is a unique subgame-perfect equilibrium, any outcome can occur in Nash equilibrium. Consequently, economists focus on strong refinements of equilibrium and ignore other equilibria. For instance, they might select the Rubinstein outcome in bargaining games or an efficient outcome in repeated games. All of these applications assume common knowledge of payoffs. The robustness program in game theory seeks to determine when strong predictions from equilibrium refinements can be maintained despite a slight relaxation

Date: First Version: March 2009; This Version: January 2012.

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. We are grateful to Gabriel Carroll, George Mailath, the audiences in several presentations, two anonymous referees, and an editor for very useful comments.

of common-knowledge assumptions.¹ Here, we show a lack of robustness of such predictions: *any* equilibrium outcome may become *uniquely* rationalizable when beliefs are perturbed, so that no equilibrium action can ever be ruled out without an extremely precise knowledge of players' beliefs.

Our work here builds on existing results, which show a similar lack of robustness in finite games (Weinstein and Yildiz (2007), Chen (2012)). Many important economic models, including those mentioned above, employ infinite-horizon dynamic games, so here we establish several different extensions which apply to such games. Our most notable application is an “unrefinable” folk theorem for infinite repeated games: For every payoff v in the interior of the individually rational and feasible set, and for sufficiently patient players, we construct a perturbation such that v is the *unique* rationalizable outcome. Moreover, in the situation described by the perturbation, all players anticipate that the payoffs are within an 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 the standard folk theorems suggests the need for a refinement to obtain clear predictions, the multiplicity in our unrefinable folk theorem shows the impossibility of a robust refinement. In the same vein, in Rubinstein’s bargaining model, we show that any bargaining outcome is the unique rationalizable outcome for some perturbation. Once again, no refinement can robustly rule out these outcomes.

These applications follow from our Proposition 2, which states: For any Bayesian Nash equilibrium and any type t_i , there exists a perturbed type \hat{t}_i for which the equilibrium action plan of t_i is the unique rationalizable plan. Furthermore, the unique rationalizable belief of \hat{t}_i regarding the outcome is arbitrarily close to the equilibrium belief of t_i . In particular, if the original game has complete information, then the perturbed type assigns probability nearly one to the equilibrium path (Corollary 1). Here the meaning of “perturbation” is that \hat{t}_i may be chosen such that 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. Hence, if a researcher has noisy information about the players’ beliefs up to a finite order but does not have any other information, then he cannot

¹See for example, Kreps, Milgrom, Roberts and Wilson (1982), Fudenberg, Kreps, and Levine (1988), Rubinstein (1989), Monderer and Samet (1989), and Kajii and Morris (1997).

distinguish some of the perturbations \hat{t}_i from the original type t_i . Consequently, he cannot verify a prediction about the behavior of t_i unless it is also true for \hat{t}_i . In particular, by Proposition 2, he cannot verify any prediction of an equilibrium refinement that does not follow from equilibrium alone.

In some applications, a researcher may believe that even if there is higher-order uncertainty about payoffs, there is common knowledge of some of the basic structure of payoffs and information. In particular, in a repeated game, he may wish to retain common knowledge that the players' payoffs in the repeated game are the discounted sum of the stage-game payoffs. The perturbations constructed in Proposition 2 would not maintain such common knowledge, and in general, restrictions on perturbations sometimes lead to sharper predictions. In the particular case of repeated games, however, we show (Proposition 5) that our conclusions remain intact: the perturbed types in the unrefinable folk theorem can be constructed while maintaining full common knowledge that we are playing a repeated game with commonly known discount factor, with uncertainty only concerning the stage-game payoffs.² In the same vein, Penta (2012a) describes robust predictions, under sequential rationality, when the fact that certain parameters are known to certain players is common knowledge. He shows that restrictions on information, combined with restricted payoff spaces, may lead to sharper predictions. In Section 6 we extend Penta's characterization to infinite-horizon games.

Our Proposition 2 applies more narrowly than the existing structure theorems, but with a stronger conclusion. It applies only to action plans played in some equilibrium, and not to all rationalizable plans. The stronger conclusion is that the perturbed types actually expect the selected equilibrium outcome to occur as the unique rationalizable play. Without this stronger conclusion, the selected outcome may be realized only by types who are surprised by their opponents' moves and play moves they did not expect to play (see Example 3). This would prevent one from applying the existing structure theorems to the analysis of equilibrium payoffs, so the stronger conclusion is important to our unrefinable folk theorem. We have also established the natural extension of previous results to all rationalizable actions in infinite dynamic games.

²This result also suggests that one may not need non-trivial commitment types for reputation formation; uncertainty about the stage payoffs may be enough when one allows more sophisticated information structures.

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 present extensions of our results to Penta's framework in Section 6. Section 7 concludes. The proofs of our general results are presented in the appendix.

2. BASIC DEFINITIONS

We suggest that the reader skim this 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 plan* (or simply 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, including histories of infinite length. We write $z(a)$ for the terminal history that is reached by profile a .

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 $u_i(\theta, a)$.

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.

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 $\Theta \times T_1 \times \cdots \times T_n$ 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 $T = T_1 \times \cdots \times T_n$ 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 $\kappa_{t_i} = p(\cdot | t_i)$ 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 $T^* = T_1^* \times \cdots \times T_n^*$ (Mertens and Zamir (1985); see also 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.”

Equivalence of Actions and 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 a game $(\Gamma, \Theta, T, \kappa)$ is *continuous at infinity*⁴ (first defined by Fudenberg and Levine (1983)) iff for any $\varepsilon > 0$, there exists $L < \infty$, such that

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

⁴ In that we do not allow L to depend on θ , this definition assumes that the possible payoff functions in the game are *equicontinuous* at infinity. This equicontinuity, as opposed to mere continuity of each θ , holds in all of our applications but is not needed in our propositions; it is useful for establishing certain properties of interim correlated rationalizability, noted in the next section. See Weinstein and Yildiz (2012) for more.

for all $i \in N$ and all terminal histories $h, \tilde{h} \in Z$.

We say that actions a_i and a'_i are *equivalent* if $z(a_i, a_{-i}) = z(a'_i, a_{-i})$ for all $a_{-i} \in A_{-i}$. For any integer L , we say that a_i and a'_i are *L -equivalent* if $z(a_i, a_{-i})^L = z(a'_i, a_{-i})^L$ 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 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 repeated games with discounting, games of sequential bargaining with discounting, all finite-horizon games, and so on. Games that are excluded include repeated games with a limit of averages criterion, or bargaining without discounting; generally, any case in which there can be a significant effect from the arbitrarily far future.

Interim Correlated Rationalizability. 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 . Note that BR_i is non-empty under continuity at infinity, because this implies continuity with respect to the product topology on histories, and that topology is compact by Tychonoff's theorem. A solution concept $\Sigma_i : t_i \mapsto \Sigma_i[t_i] \subseteq A_i$, $i \in N$, is said to have the *best-response property* 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 $\pi(a_{-i} \in \Sigma_{-i}[t_{-i}]) = 1$. We define *interim correlated rationalizability* (ICR), denoted by S^∞ , as the largest solution concept with best-response property. This largest set is well-defined because the set of solution concepts with best-response property is closed under coordinate-wise union, i.e., $S^\infty[t_i]$ is the union of the sets $\Sigma_i[t_i]$ over all solution concepts Σ with the best-response property.

Under our assumption of continuity at infinity, interim correlated rationalizability can be computed by the following elimination procedure⁵: 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

⁵See Weinstein and Yildiz (2012) for a proof of this claim.

$\pi(a_{-i} \in S_{-i}^{k-1}[t_{-i}]) = 1$. 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 $S^k[t] = S_1^k[t_1] \times \cdots \times S_n^k[t_n]$. Then,

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

This equality of the two concepts implies that the infinite intersection is non-empty.

Interim correlated rationalizability was 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.” They also show that ICR is upper-hemicontinuous for finite games. While this is not known to be true for all infinite games, we show that it is true under the present assumptions in Weinstein and Yildiz (2012).

ICR is the weakest rationalizability concept, and our main results such as Proposition 2 carry over to any stronger, non-empty concept by a very simple argument: If an action is uniquely ICR for a perturbed type, it is also uniquely selected by the stronger concept at that type. In particular, our result is true without modification for the interim sequential rationalizability (ISR) concept of Penta (2012a), if no further restriction on players’ information and beliefs is made. The concept of ISR does entail some modification to our arguments when combined with restrictions on players’ information; see Section 6.

Miscellaneous Definitions and Notation. We fix a set $\bar{A} = \bar{A}_1 \times \cdots \times \bar{A}_n$ of action profiles where \bar{A}_i selects one representative from each reduced-form equivalence class of action plans for player i . 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 $\Pr(\cdot | \pi, s_i)$ and $E[\cdot | \pi, a_i]$ 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 $E[\cdot | s, t_i]$.

It has recently become common in game theory to reserve the term *knowledge* for cases where truth axiom holds, and refer to *certainty* when players are certain but may be wrong⁶. In this paper, we use this nuanced language, although the distinction is not important for our analysis, and the general reader will not lose by reading “common certainty” and “common knowledge” as alike. In particular, for any type space (Θ, T, κ) , we say that type t_i *knows* an event (or a proposition) if the event (or the proposition) holds on $\Theta \times \{t_i\} \times T_{-i}$, and say that type t_i is *certain of* the event if κ_{t_i} assigns probability 1 to that event. Likewise, an event (or a proposition) is *common knowledge* (according to t_i) if it holds everywhere on $\Theta \times T$; the event is common certainty according to t_i if it holds on a “belief-closed subspace” that contains t_i . When types are embedded in a universal type space, knowledge is relaxed to certainty; we still say that an event is common knowledge according to $h_i(t_i)$ if it is common knowledge according to t_i . In our informal discussions, we say that one *drops* a common knowledge assumption if he allows perturbed types coming from spaces in which the assumption may fail, and we say that one *retains* a common knowledge assumption if he restricts the perturbations to the type spaces in which the assumption holds throughout.

3. STRUCTURE THEOREM

In this section we will present our main result for general infinite-horizon games. Given any game $(\Gamma, \Theta, T, \kappa)$ that is continuous at infinity and any Bayesian Nash equilibrium $s : T \rightarrow A$, we will show that there are perturbations \hat{t}_i of types t_i for which $s_i(\hat{t}_i)$ is the *only* rationalizable plan. Moreover, the unique rationalizable belief of \hat{t}_i regarding the outcomes is arbitrarily close to the belief of \hat{t}_i under s . The following structure theorem extends existing results and plays a crucial role in our construction.

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.*

⁶The truth axiom states that anything which is known is always true. The distinction of *knowledge* from *certainty* can be critical in dynamic games when a zero-probability event occurs. An event which was known is then still known, while that which was merely certain may be no longer certain.

In Weinstein and Yildiz (2007) we showed the structure theorem for finite-action games in normal form, under the assumption that the space of payoffs is rich enough that any action is 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 (2012) has nonetheless extended the structure theorem to finite dynamic games, showing that the same result holds under the weaker assumption that all payoff functions on the terminal histories are possible. Here, we extend Chen’s results further by allowing infinite-horizon games that are continuous at infinity.

The result can then be interpreted as follows: Consider a type t with a rationalizable action plan a . For some arbitrarily chosen k , suppose we find it impossible to distinguish types whose beliefs are similar up through order k . The lemma states that, for any L , there is a finite Bayesian game with type \hat{t} who we cannot distinguish from t and for which a is the unique rationalizable action plan through period L .

Proposition 1 has one important limitation. 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, as we establish next.

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, 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.

Then, by Proposition 1, 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 find this outcome unlikely. 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. Therefore, the perturbed type is nearly certain that he will play Defect in the second round.

The above example demonstrates that the beliefs of the perturbed types in Proposition 1 may drastically diverge from the unique rationalizable outcome. This prevents us from applying Proposition 1 to study the expected payoffs and the players' intended play. Our next result overcomes this limitation. For this, we need an outcome to be a Bayesian Nash equilibrium rather than merely rationalizable.

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. 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 limitation of Proposition 1 above does not apply. 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; in fact, Proposition 2 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}^* .

We will later apply this result to some popular complete-information games. In order to state the result for 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

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

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 . To see this, first note that for $\lambda = 0$, under $t_{i,m,0}$, it is m th-order mutual certainty 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 $t_j^{a^*}$, 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 UNREFINABLE FOLK THEOREM

In this section, we consider infinitely repeated games with complete information. Under the standard assumptions for the folk theorem, we prove an unrefinable 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 $B = B_1 \times \cdots \times B_n$ 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. Without loss of generality, we will assume that each player i has at least two moves in the stage game, i.e., $|B_i| \geq 2$. 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

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

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 1. *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 $v \in \text{int}V^*$ and $\varepsilon > 0$. By Lemma 1, 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 unrefinable 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 the standard folk theorems may suggest a need for a refinement, the multiplicity in our unrefinable folk theorem emphasizes the impossibility of a robust refinement.

Structure Theorem with Uncertainty only about the Stage Payoffs. An important drawback of the structure theorems is that they may rely on the 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 (by considering only the type spaces in which these structures are true throughout), those theorems may become inapplicable. Indeed, recent papers (e.g. Weinstein and Yildiz (2011) and Penta (2012a, 2012b)) 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 any fixed discount factor $\delta \in (0, 1)$, we define

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

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 predictions which are robust against 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 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 plan 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 unrefinable folk theorem above.

Proposition 4 is proved in the Appendix. The proof first 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 $\theta \in \Theta_\delta^*$. The construction of these types is rather involved, and uses ideas from learning and incentives in repeated games. Using the existence of these types, one then constructs the nearby types in the proposition following the ideas sketched in illustrating the proof of Corollary 1 above. In the following example we illustrate the gist of the idea on the twice-repeated prisoners' dilemma.

Example 1. Consider again the twice-repeated prisoners' dilemma 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 $\delta = 1$. Given a type according to which the payoffs g^{PD} are common knowledge, we will construct a nearby type for which tit-for-tat 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 assign zero probability to some history. But now consider a type $t_{i, b_i, 1}$ that puts probability 1/2 on $\left(\theta_{\delta, g^{(b_i, b_{-i})}}, t_{-i, C, 0}\right)$

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 construct types $t_{i, k}$ with approximate k th-order mutual certainty of prisoners' dilemma payoffs who Defect at all histories in their unique rationalizable plan. Type $t_{i, 1}$ 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). Proceeding inductively on k , 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 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 certainty 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 certainty of the prisoners' dilemma payoffs. 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 types $\hat{t}_{i, k}$, which are close to common-knowledge types, defect in period 1 and play tit-for-tat in period 2. 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 believe that their opponent probably plays defection followed by tit-for-tat, so they cooperate in the first period. In the second period, if they saw D , they still think they are playing prisoner's dilemma, so they defect. If they saw C , they think they are playing $g^{(C, C)}$, so they cooperate. That is, their unique rationalizable action is tit-for-tat with cooperation at the initial node.

Early literature identified two mechanisms through which a 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 in Example 1, our perturbed types in the proof of Proposition 4 generalize this idea: 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 (2012) combine the two mechanisms in order to obtain a very strong conclusion: any rationalizable action can be made uniquely rationalizable under some perturbation.

At another level, however, Propositions 4 and 5 make a stronger point than the previous reputation and contagion literatures, in the following sense: The existing models mainly rely on behavioral commitment types (or “crazy” types) that follow a complete plan of action throughout the game, suggesting that non-robustness may be due to psychological/behavioral concerns that are overlooked in game-theoretical analyses. By proving the unrefinable folk theorem while allowing uncertainty only about the stage payoffs⁷, Propositions 4 and 5 show that informational concerns can lead to the non-robustness results, even without a full range of crazy types.

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 from our work here. 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

⁷Of course, this allows for “crazy” types who always play the same action – but not for those who play any more complicated plan, say tit-for-tat.

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 future payoffs, the moves of our perturbed types reveal information about both their own and the other players' stage-game payoffs, which are assumed to be constant over time.

Some other papers have also restricted attention to perturbations which keep some payoff structure common knowledge. In Weinstein and Yildiz (2011), we dealt with nice games, which are static games with unidimensional action spaces and strictly concave utility functions. We obtained a characterization for sensitivity of Bayesian Nash equilibria in terms of a local version of ICR, allowing arbitrary common-knowledge restrictions on payoffs.⁸ 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. Most generally, Penta (2012b) proved a version of the structure theorem under arbitrarily given common-knowledge restrictions on payoffs, identifying a set of actions that can be made uniquely rationalizable by perturbing the interim beliefs under the given common knowledge restrictions on payoffs.

We establish that equilibrium refinements are not upper hemicontinuous, even if one imposes common-knowledge restrictions on the payoff structure. This results in the lack of robustness above. One may, however, raise the same criticism for unrefined solution concepts, such as Bayesian Nash equilibrium and ICR. Extending the results of Dekel, Fudenberg, and Morris (2007) for finite games to the infinite games we analyze here, we show in Weinstein and Yildiz (2012) that ICR is upper hemicontinuous under the usual continuity and compactness properties, provided that the space of payoffs can be embedded into a compact metric space. In particular, we show that ICR is upper hemicontinuous whenever the payoffs are restricted to be in Θ_δ^* , imposing common knowledge of the repeated-game payoff structure and the discount factor. In fact, it suffices for the discount factor to be unknown but bounded away from 1. Therefore, the predictions of ICR are robust under the above restriction. This further implies that the unique solution for the perturbed types remains robust with respect to further perturbations. We should note, however, that the set Θ^* of *all*

⁸Weinstein 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.

payoff functions in infinite-horizon games is not metrizable (or sequentially compact), and so we do not know if ICR is upper hemicontinuous in the entire universal type space. We should also note that, while Bayesian Nash equilibrium is not upper hemicontinuous in general (by Proposition 1), it is robust with respect to perturbations that assign high probability on the unique outcome (as in Proposition 2).

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.

The existing literature already illustrates that the subgame-perfect equilibrium is sensitive to incomplete information. For example, for high δ , the literature on the 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. The 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. The 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 already a point of discontinuity for the complete-information model. 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. INFORMATION AND SEQUENTIAL RATIONALITY

We have discussed earlier that when analyzing robustness, one may want to consider only perturbations which retain some structural common-knowledge assumptions, such as the additive payoff structure in a repeated game, or the fact that a player knows the true value of a certain parameter. When the set of possible payoff functions is the same from the point of view of every player, our formalism suffices for this. If each player may have his own

information, and furthermore this information (unlike mere *beliefs*) is never doubted even when probability-zero events occur, a slightly different setup, introduced by Penta (2012a), is necessary. This setup is needed, for instance, to analyze a case in which it is common knowledge that players know (and never doubt) their own utility functions. When the underlying set of payoff parameters is sufficiently rich (e.g. when all possible payoff functions are available as in our model above), retaining such assumptions does not lead to any change, and the original characterization in Proposition 1 remains intact. In restricted parameter sets, retaining the informational assumption may lead to somewhat sharper predictions. For example, in private value environments, this allows one round of elimination of weakly dominated actions in addition to rationalizability. In this section, building on an extension of the result of Penta (2012a) to infinite horizon games, we will extend our results to Penta’s setup. Note also that Penta’s framework is related to that of Battigalli and Siniscalchi (2003), who introduced a version of rationalizability for extensive-form games which allowed for restrictions on players’ beliefs about their opponents’ behavior.

Consider a compact set $C = C_0 \times C_1 \times \dots \times C_n$ of payoff parameters $c = (c_0, c_1, \dots, c_n)$ where the underlying payoff functions θ depends on the payoff parameters c : $\theta = f(c)$ for some continuous and one-to-one mapping $f : C \rightarrow \Theta^*$. We will assume it is common knowledge that θ lies in the subspace $f(C) \subseteq \Theta^*$. It will also be assumed to be common knowledge throughout the section that the true value of the parameter c_i is known by player i . For any type t_i , we will write $\mathbf{c}_i(t_i)$ for the true value of c_i , which is known by t_i . Note that this formulation subsumes our model above, by simply letting C_1, \dots, C_n be trivial (singletons) so that $\Theta^* = C_0$. We will write $T^{C^*} \subset T^*$ for the subspace of the universal type space in which it is common knowledge that $\theta \in f(C)$ and each player i knows the true value of c_i . As in Penta (2012a), we will restrict perturbations to lie in T^{C^*} . Following Penta, we will further focus on multistage games in which all previous moves are publicly observable.

A *conjecture* of a player i is a conditional probability system $\mu_i = (\mu_{i,h})_{h \in H}$ that is consistent with Bayes’ rule (on positive probability events), where $\mu_{i,h} \in \Delta(C_0 \times T_{-i} \times A_{-i})$ for each $h \in H$. Here, it is implicitly assumed that it remains common knowledge throughout the game that $(c_1, \dots, c_n) = (\mathbf{c}_1(t_1), \dots, \mathbf{c}_n(t_n))$. In particular, player i assigns probability 1 to $\mathbf{c}_i(t_i)$ throughout the game. For each conjecture μ_i of type t_i , we write $SBR_i(\mu_i|t_i)$ for the set of actions $a_i \in A_i$ that remain a best response to μ_i at all information sets

that are not precluded by a_i ; we refer to $a_i \in SBR_i(\mu_i|t_i)$ as a *sequential best response*. A solution concept $\Sigma_i : t_i \mapsto \Sigma_i[t_i] \subseteq A_i$, $i \in N$, is said to have the *sequential best-reply property* if and only if for each t_i and for each $a_i \in \Sigma_i[t_i]$, there exists a conjecture μ_i of t_i such that $a_i \in SBR_i(\mu_i|t_i)$, the beliefs about (θ, t_{-i}) according to $\mu_{i,\emptyset}$ agree with κ_{t_i} and $\mu_{i,\emptyset}(a_{-i} \in \Sigma_{-i}[t_{-i}]) = 1$, where \emptyset denotes the initial node of the game. We define *interim sequential rationalizability* (ISR), denoted by ISR^∞ , as the largest solution concept that has the sequential best-reply property. In finite games this is equivalent to the result of an iterative elimination process similar to iterative elimination of strictly dominated actions (see Penta (2012a) for that alternative definition). Note that ISR differs from ICR only in requiring sequential rationality, rather than normal-form rationality. The only restriction here comes from the common knowledge assumption that the player i does not change his belief about c_i , since the players' conjectures off the path are otherwise unrestricted. The resulting solution concept is relatively weak (e.g. weaker than extensive form rationalizability) and equal to ICR in rich environments.⁹ We use the following richness assumption of Penta (2012a).

Assumption 1. *For every $a_i \in A_i$ there exists c^{a_i} such that a_i is conditionally dominant under c^{a_i} , i.e., at every history that is consistent with a_i , following a_i is better than deviating from a_i .*

Lemma 2 (Penta (2012a)). *Under Assumption 1, for any finite-horizon multistage game $(\Gamma, \Theta, T, \kappa)$ with $\Theta \subset f(C)$, for any type $t_i \in T_i$ of any player $i \in N$, any ISR action $a_i \in ISR_i^\infty[t_i]$ of t_i , and any neighborhood U_i of $h_i(t_i)$ in the universal type space T^* , there exists a hierarchy $h_i(\hat{t}_i) \in U_i \cap T_i^{C^*}$, such that for each $a'_i \in ISR_i^\infty[\hat{t}_i]$, a'_i is equivalent to a_i .*

Note that the above model maintains two common-knowledge assumptions throughout the perturbations: (i) each player i assigns probability one on the true value of c_i (in defining the interim beliefs at the beginning of the game) and (ii) the players never doubt this fact throughout the play of the game (in defining the conjectures μ). Lemma 7 establishes that, under Assumption 1, maintaining the common knowledge of (i) has no bite because ISR is equal to ICR in static games. It also establishes, however, that maintaining common

⁹For example, ISR is equal to ICR if for every a_i and c_i , there exists (c_0, c_{-i}) such that a_i is conditionally dominant under (c_0, c_i, c_{-i}) (cf. Assumption 1). ISR is equal to ICR also when no player has any information. See Penta (2009) for further details.

knowledge of (ii) leads to potentially sharper predictions in dynamic games, as ISR may be a strict refinement of ICR in such games. The next result extends Penta’s result to infinite horizon games.

Proposition 7. *Under Assumption 1, consider any multistage game $(\Gamma, \Theta, T, \kappa)$ that is continuous at infinity and $\Theta \subset f(C)$ is such that each $\theta = f(c) \in \Theta$ is in the interior of $f(C_0 \times \{(c_1, \dots, c_n)\})$. For any type $t_i \in T_i$ of any player $i \in N$, any ISR action $a_i \in \text{ISR}_i^\infty [t_i]$ 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 $h_i(\hat{t}_i) \in U_i \cap T_i^{C^*}$, such that for each $a'_i \in \text{ISR}_i^\infty [\hat{t}_i]$, a'_i is L -equivalent to a_i , and \hat{t}_i is a type in some finite, common-prior model.*

Penta (2012a) proves this result for finite games without requiring that θ is in the interior of $f(C_0 \times \{(c_1, \dots, c_n)\})$. Here, we extend this result to infinite-horizon games, with the above requirement that one can make slight payoff perturbations in payoffs by changing c_0 alone. This is required only for uniformly small perturbations, in that there exists $\varepsilon > 0$ such that if $|\theta(z) - \theta'(z)| \leq \varepsilon$ for all $z \in Z$, then there exists a c'_0 that leads to θ' instead of θ .¹⁰ Roughly speaking, Proposition 7 characterizes the robust prediction of common knowledge of sequential rationality and the informational assumptions, such as the true value of each c_i is known by player i , who never updates his beliefs regarding c_i . These are the predictions that can be made by interim sequential rationality alone. One cannot obtain a sharper robust prediction than those of interim sequential rationalizability by considering its refinements, even if one is willing to retain common knowledge assumptions regarding players’ information.

Using Proposition 7, one can also extend our other results to this framework. Here, we will only formally present the extension of Corollary 1, our structure theorem for equilibrium in the case of complete information; the proof is relegated to the Appendix.

Proposition 8. *Under Assumption 1, let $(\Gamma, \{\theta^*\}, \{t^{CK}(\theta^*)\}, \kappa)$ be a multi-stage complete-information game that is continuous at infinity, with $\theta^* = f(c^*) \in \Theta$ in the interior of $f(C_0 \times \{(c_1^*, \dots, c_n^*)\})$. Let also $a^* \in \text{ISR}^\infty [t^{CK}(\theta^*)]$ be a Nash equilibrium of this game.*

¹⁰While this assumption rules out pure private value environments in which $|C_0| = 1$, it allows approximate private value environments in which the players know their payoff functions up to an arbitrarily small error ε .

Then, for any $i \in N$, for any $L < \infty$, 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 $h_i(\hat{t}_i) \in U_i \cap T_i^{C*}$, such that for every ISR belief π of \hat{t}_i ,

- (1) $a_i \in ISR_i^\infty[\hat{t}_i]$ iff a_i is L -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 $j \in N$.

Like Proposition 4, this result remains silent for the tail behavior, establishing uniqueness of ISR only up to an arbitrary finite horizon. The result is stronger than Corollary 1, in that the perturbed types are in T_i^{C*} , retaining common knowledge of informational assumptions. Note that the result also assumes that a^* is ISR, putting a weak restriction on equilibrium. Since subgame-perfect equilibria of a repeated game are ISR, the unrefinable folk theorem in Proposition 3 also extends to the current setup.

7. CONCLUSION

In economic models there are often a multitude of equilibria. This problem is especially acute in infinite-horizon games, such as repeated games, in which the folk theorem applies, establishing that any feasible payoff vector can be supported by an equilibrium. In response to such multiplicity, economists often focus on refinements. In this paper, we develop a structure theorem for infinite-horizon games that can be readily used in applications. Our result establishes that without any common-knowledge assumption regarding payoffs and information structure, one cannot obtain any robust prediction that is not implied by Bayesian Nash equilibrium alone. As an application, we prove an unrefinable folk theorem, showing that no feasible payoff vector can be excluded if there is noise in our knowledge of players' beliefs. Our construction allows uncertainty only about the stage payoffs. This shows that, even without the large set of commitment types used in the reputation literature, the uncertainty behind the structure theorem can operate with full force.

APPENDIX A. PROOF OF PROPOSITION 1

A.1. Preliminaries. We start by describing some 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 pre-image of the mapping $(w, y) \mapsto (f(w), g(y))$.

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. 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 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' .¹¹ 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 and Virtually 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

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

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

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

¹¹Proof: 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 $\mu_i(a'_i) \in \mu_i(\bar{A}_i)$.

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, so we refer to them as “virtually truncated.” We formalize this 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 $A_i^m = \{a_i^m | a_i \in A_i\}$.

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) $T_i = \tau_i^m(T_i^m)$ 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 types, 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 $t_i \in T_i^{*m}$.¹² First two statements immediately follow from (A.4). Part 3 follows from (A.4) and Lemma 6. \square

In Weinstein and Yildiz (2007), we proved a version of Proposition 1 for finite action games. We used a richness assumption on Θ^* that is natural for static games but rules out fixing a dynamic extensive game form. Chen (2012) has proven this result for finite dynamic games, under a weaker richness assumption that is satisfied in our formulation. Our proof of Proposition 1 will take advantage of these earlier results. In particular, we will use this lemma, which is implied by Chen's theorem:

Lemma 8 (Weinstein and Yildiz (2007) and Chen (2012)). *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 $h_i(\hat{t}_i) \in U$, 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.*

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

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

¹²If 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 $k > 1$.

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

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

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 ‘‘virtually truncated’’ Bayesian 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

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

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

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

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

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

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 ,

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

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 8, 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 8, there exists a hierarchy $h_i(\bar{t}_i^m)$ in $(\tau_i^m)^{-1}(U_i)$ as in the statement above.

Please note that the unique ICR action in this perturbation will be robust to further small perturbations, just as in the original structure theorem of Weinstein and Yildiz (2007), *so long as* we confine attention to the truncated game form Γ^m , since here the game is finite and the results of Dekel, Fudenberg, and Morris (2007) apply. However, once we apply the following step to pull back the constructed type to lie in the original, infinite game-form, this statement is known to be true only for perturbations that retain common knowledge of $\hat{\Theta}^m$. The statement is not necessarily true for the perturbations that lie outside the image of the embedding.

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, $\hat{a}_i^m \in S_i^\infty[\bar{t}_i^m]$. 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 9. *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 plan 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

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

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 9. 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 $t_j^{s_j^*(t_j)}$. 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

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

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 $\pi(t_i, \bar{m}, 0) \in V_i$. 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 5,¹³ 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 $\pi(t_i, \bar{m}, \bar{\lambda}) \in V_i$. 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

¹³To 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) .

APPENDIX C. PROOF OF PROPOSITION 4

We write $T_i^{CK(\Theta_\delta^*)}$ for the set of types of player i according to which it is common knowledge that $\theta \in \Theta_\delta^*$, i.e. that we are playing a repeated game with discount factor δ . In order to harness our previous constructions, in Lemma 11 we will construct, for *every* possible plan a_i and finite time horizon L , a type in $T_i^{CK(\Theta_\delta^*)}$ for which all rationalizable plans are L -equivalent to a_i . These types then play the role that dominant-action types would play in richer environments. Our first step towards this, Lemma 10, constructs types who do not believe the other players' actions ever affect them directly, but who find others' actions informative about their own payoffs. They are further constructed so as to always choose the “myopic” action, optimizing the expected payoff in the current period. This construction will not work on all plans, but only on those satisfying this version of the sure-thing principle:

Definition 1. A plan a_i is said to be *sure-thing compliant* if and only if there is no partial history h and move $b_i \in B_i$ such that $a_i(h, (a_i(h), b_{-i})) = b_i$ for every b_{-i} but $a_i(h) \neq b_i$.

In other words, a plan is sure-thing compliant if whenever the player plays b_i in all possible continuations next period, he also plays b_i this period. This is of course equivalent to the sure-thing principle of Savage if the player has the same preferences over his moves in both periods. Given that, in our construction in the next proof, player i is actually facing a single-player decision problem with unknown payoffs, it is not hard to see that this particular construction can *only* work for sure-thing compliant plans. Of course the necessity of the condition is not relevant to later results, and our further construction in Lemma 11 extends the result to all plans.

Lemma 10. *For any δ , any L , and any sure-thing compliant action plan a_i , there exists a type $t^{a_i, L} \in T_i^{CK(\Theta_\delta^*)}$ for which all rationalizable plans are L -equivalent to a_i .*

Proof. We will induct on L . When $L = 1$, it suffices to consider a type $t^{a_i, 1}$ who is certain that in the stage game, $a_i(\emptyset)$ yields payoff 1 while all other actions yield payoff 0. Now fix L, a_i and assume the result is true for all players and 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 uninformative — this ensures that he always chooses the myopically best action.

Formally: Let \hat{H} be the set of histories of length $L - 1$ in which player i always follows the plan a_i , so that $|\hat{H}| = |B_{-i}|^{L-1}$, where B_{-i} is the set of profiles of static moves for the other players. For each history $h \in \hat{H}$, we construct a pair (t_{-i}^h, θ^h) , and our constructed type $t^{a_i, L}$ assigns equal weight to each of $|B_{-i}|^{L-1}$ such pairs. Each type t_{-i}^h is constructed by applying the inductive hypothesis to a plan a_{-i}^h which plays according to history h so long as i follows a_i , and simply repeats the previous move forever if player i deviates. Such plans are sure-thing compliant for the player(s) $-i$ because at every history, the current action is repeated on at least one branch.

To define the payoff functions θ^h for all $h \in \hat{H}$, we will need to define an auxiliary function $f : \tilde{H} \times B_i \rightarrow \mathbb{R}$, where \tilde{H} is the set of prefixes of histories in \hat{H} . The motive behind the construction is that $f(h, \cdot)$ represents i 's expected value of his stage-game payoffs conditional on reaching the history h . The function f is defined iteratively on histories of increasing length. Specifically, define 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 empty history. 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 sure-thing compliance, at least two different actions are prescribed for continuations $(h, (a_i(h), b_{-i}))$ as we vary 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

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

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

These payoffs were chosen to satisfy the constraints

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

$$(C.2) \quad f(h, a_i(h)) \geq f(h, b_i) + \varepsilon \quad (\forall h, b_i \neq a_i(h)).$$

as can be verified algebraically.

For each history $h \in \hat{H}$, define the stage-game payoff function $g^h : B \rightarrow [0, 1]^n$ by setting $g_i^h(b) = f(h, b_i)$ and $g_j^h(b) = 0$ at each b and $j \neq i$. Define θ^h accordingly, by

$$\theta^h(b^0, b^1, \dots) \equiv (1 - \delta) \sum_{l=0}^{\infty} \delta^l g^h(b^l),$$

as in (4.1). Define $t^{a_i, L}$ as mentioned above, by assigning equal weight to each pair (t_{-i}^h, θ^h) .

We claim that under rationalizable play, from the perspective of type $t^{a_i, L}$, when he has followed a_i and reaches history $h \in \tilde{H}$, $f(h, \cdot)$ is his expected value of the stage-game payoff g_i . We show this by induction on the length of histories, backwards. When a history $h \in \hat{H}$ is reached, player i becomes certain (assuming rationalizable play) that the opposing types must be t_{-i}^h and thus the payoffs must be θ^h , which is the desired result for this case. Suppose the claim is true for all histories in \tilde{H} of length M . Note that type $t^{a_i, L}$ puts equal weight on all sequences of play for his opponent. Therefore, for a history $h \in \tilde{H}$ of length $M - 1$, the expected payoffs are given by the right-hand-side of (C.1) which proves the claim.

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: the myopic goal of maximizing his average stage-game payoffs at the current stage, and the desire 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, player i has a uniform belief on the other players' actions as long as he follows a_i . Let \hat{h} be a random variable equal to the shortest realized history at which a'_i differs from a_i before period L , or ∞ if they do not differ by period L . Note that the uniform belief on others' actions implies that $\hat{h} \neq \infty$ with positive probability. We show that conditional on any non-infinite value of \hat{h} , \bar{a}_i strictly outperforms a'_i on average. In fact this is weakly true stage-by-stage, and strictly true at the first deviation, because:

At stages $1, \dots, |\hat{h}|$: The plans are identical.

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

At stages $|\hat{h}| + 2, \dots, L$: Along the path observed by a player following a'_i , the other players are known to repeat their stage- $|\hat{h}| + 1$ move at stages $|\hat{h}| + 2, \dots, L$. So at these stages, the plan a'_i cannot do better than to optimize with respect to the history truncated at length $|\hat{h}| + 1$. The plan \bar{a}_i optimizes the expected stage-game payoffs with respect to a longer history, under which opposing moves are identical through stage $|\hat{h}| + 1$. Since he is therefore solving a less-constrained optimization problem, he must perform better than a'_i at each stage $|\hat{h}| + 2, \dots, L$.

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 $\hat{h} = \infty$, again \bar{a}_i cannot be outperformed because he optimizes based on complete information after L , and \bar{a}_i and a'_i prescribe the same behavior before L .

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

The next lemma builds on this result to generalize to all action plans.

Lemma 11. *For any $\delta \in (0, 1)$, any L and any action plan a_i , there exists a type $t_i^{a_i, L} \in T_i^{CK(\Theta_i^*)}$ for which playing according to a_i until L is uniquely rationalizable in reduced form.*

Proof. For some $b_{-i}^* \in B_{-i}$, which will be fixed throughout the proof, 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 $b_{-i} \neq b_{-i}^*$. That is, player i 's payoff does not depend on his own action, but only on whether the other players reward him by playing b_{-i}^* . Write $\hat{\theta} \in \Theta_i^*$ for a payoff function resulting from g_i , i.e., $\hat{\theta}_i(b^0, b^1, \dots) = (1 - \delta) \sum_l \delta^l g_i(b^l)$. This $\hat{\theta}$ will be fixed throughout the proof and the constructed type $t_i^{a_i, L}$ will be certain that payoffs are given by $\hat{\theta}$. Fix an M large enough that $\delta^M < \delta^L (|B_i| - 1) / (2|B_i| - 1) < \delta^L / 2$. Let \hat{H} be the set of all histories of length L or less. Let \hat{A}_{-i} be the set of action profiles a_{-i} for which there exists a function $\rho : \hat{H} \times B_{-i} \rightarrow B_i$ such that

- (1) for any $l \leq L + 1$, any history h^{l-1} of length $l - 1$ and any $(b_i, b_{-i}) \in B_i$, $a_{-i}(h^{l-1}, (b_i, b_{-i})) = b_{-i}^*$ if $b_i = \rho(h^{l-1}, b_{-i})$ and $a_j(h^{l-1}, (b_i, b_{-i})) \neq b_j^*$ for every $j \neq i$ otherwise;¹⁴
- (2) $\rho(h^L, b_{-i}) \equiv a_i(h^L)$ for all those h^L such that player i has played according to a_i throughout, and
- (3) for any $l \in \{L + 2, \dots, M\}$ and any h at the beginning of l , $a_{-i}(h) = a_{-i}(h^{L+1})$.

That is: At any history in $h \in \hat{H}$, the other players reward a unique move $\rho(h, b_{-i})$ of i at each history (h, b) . The only restriction on which move is rewarded occurs at stage $L + 1$, when if player i has followed a_i so far, he will be rewarded if he continues to do so. Furthermore, at stages $\{L + 2, \dots, M\}$ the other players simply repeat their move from stage $L + 1$. 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 sure-thing compliant up to date M , and thus for each $a_{-i} \in \hat{A}_{-i}$, there exists a sure-thing

¹⁴Note 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} .

compliant action \hat{a}_{-i} that is M -equivalent to a_{-i} . Let \hat{A}_{-i}^M be a finite subset of A_{-i} that consists of one sure-thing compliant element from each M -equivalence class in \hat{A}_{-i} . By Lemma 10, 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 of length l , write $P_l^*(h)$ for the probability that b_{-i}^* is played at date l conditional on h according to the rationalizable belief of $t_i^{a_i, L}$. As noted above, by symmetry,

$$(C.3) \quad P_l^*(h) = 1/|B_i| \quad \forall l \leq L,$$

and

$$(C.4) \quad P_{L+1}^*(h) = \begin{cases} 1 & \text{if } i \text{ follows } a_i \text{ until } L \\ 0 & \text{if } i \text{ follows } a_i \text{ until } L - 1 \text{ and deviates at } L \\ 1/|B_i| & \text{otherwise.} \end{cases}$$

Denote the expected payoff of type $t_i^{a_i, L}$ under any action a'_i by $U_i(a'_i)$, and note that

$$(C.5) \quad U_i(a'_i) \equiv \sum_l (1 - \delta) \delta^l E [P_l^* | a'_i].$$

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 of length l . We will show that a_i^* yields a strictly 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. It equals the first *realized* difference in moves; note that even if the two plans are not equivalent, they may not differ on a particular history.) Conditioned on the event $\tau > L$, we know $\tau = \infty$, that is, a_i^* and a'_i play identical moves and hence yield the same payoff. We will show that a_i^* has a strictly higher expected payoff than a'_i when conditioned on each of the events $\tau = L$ and $\tau < L$. On the event $\tau = L$, by (C.3) and (C.4), a'_i yields a payoff of

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

while, by (C.4), a_i^* yields

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

Hence,

$$\begin{aligned} U_i(a_i^*|\tau = L) - U_i(a_i'|\tau = L) &= (\delta^{L+1} - \delta^{M+1}) + \sum_{l>M} (1 - \delta) \delta^l [E [P_l^* | a_i^*, \tau = L] - E [P_l^* | a_i', \tau = L]] \\ &\geq [\delta^{L+1} - \delta^{M+1} - \delta^{M+1}] > 0, \end{aligned}$$

where the first inequality holds because $P_l^* \in [0, 1]$ and the strict inequality follows from our defining assumption on M .

Similarly, on the event $\tau < L$, by (C.3), a_i' yields a payoff of

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

while a_i^* yields

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

Hence,

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

where the first inequality is by the fact that $P_l^* \in [0, 1]$ and the strict inequality follows from our defining assumption on M .

Finally, note that $\Pr(\tau \leq L) > 0$ (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), so we can conclude that a_i^* yields a strictly higher expected payoff than a_i' and hence a_i' is not optimal. \square

This lemma establishes that any action can be made uniquely rationalizable for an arbitrarily long horizon, even within the restricted class of repeated game payoffs with the given discount factor δ . Using this lemma, we can now 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 $t_j^{CK}(\theta_{\delta,g^*})$. Note that under $\kappa_{t_j^{a_j^*,l}}$ it is a unique best reply to follow a_j^* up to date l . Moreover, if $\theta = \theta_{\delta,g^*}$ and the other players follow a_{-j}^* forever, then it is a best response to follow a_j^* up to date l . Hence, it is 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

APPENDIX D. PROOFS OF PROPOSITIONS 7 AND 8

Here we show how to modify the proofs of Propositions 1 and 2 in order to retain the informational common-knowledge assumptions described in Propositions 7 and 8 and satisfy sequential rationality. Note that here, a Bayesian game also assigns a “payoff type” $\mathbf{c}_i(t_i) \in C_i$ for each type t_i , and hence a Bayesian game is a list $G = (\Gamma, \Theta, T, \mathbf{c}, \kappa)$.

Proof of Proposition 7. Note that as in Lemma 4, ISR^∞ depends only on the reduced form of a plan, and, as in Lemma 6, the ISR actions of “virtually truncated” games are equivalent to the ISR actions of truncated games. In light of these facts, we now describe the major modifications to each step of the proof of Proposition 1.

In Step 1, we observe that, by the definition of ISR, each $r_{t_i}(a_i)$ is a *sequential* best response to a conjecture $\mu^{t_i, r_{t_i}(a_i)}$ of t_i such that $\mu^{t_i, r_{t_i}(a_i)}$ agrees with κ_{t_i} and puts probability one on ISR actions. We define types $(t_i, r_{t_i}(a_i), m)$ by setting $\mathbf{c}_i(t_i, r_{t_i}(a_i), m) = \mathbf{c}_i(t_i)$, so that the private information does not change, and setting

$$(D.1) \quad \kappa_{t_i, r_{t_i}(a_i), m} = \mu_{\emptyset}^{t_i, r_{t_i}(a_i)} \circ \bar{\phi}_{t_i, r_{t_i}(a_i), m}^{-1}$$

where $\bar{\phi}_{t_i, r_{t_i}(a_i), m}$ is now defined as

$$(D.2) \quad \bar{\phi}_{t_i, r_{t_i}(a_i), m} : (c_0, t_{-i}, a_{-i}) \mapsto \left(\phi_{t_i, r_{t_i}(a_i), m}(f(c_0, \mathbf{c}_i(t_i), \mathbf{c}_{-i}(t_i))), (t_{-i}, r_{t_{-i}}(a_{-i}), m) \right).$$

Since $\phi_{t_i, r_{t_i}(a_i), m}(f(c_0, \mathbf{c}_i(t_i), \mathbf{c}_{-i}(t_i))) \rightarrow f(c_0, \mathbf{c}_i(t_i), \mathbf{c}_{-i}(t_i))$ as in the proof of Proposition 1, by the interior assumption in the hypothesis, there exists \bar{m} such that for every $m > \bar{m}$, $\phi_{t_i, r_{t_i}(a_i), m}(\theta) = f(G_{t_i, r_{t_i}(a_i), m}(c_0, \mathbf{c}_{-i}(t_{-i})), \mathbf{c}_i(t_i), \mathbf{c}_{-i}(t_{-i}))$ for some $G_{t_i, r_{t_i}(a_i), m}(c_0, \mathbf{c}_{-i}(t_{-i})) \in C$, ensuring that the newly constructed types are in T^{C^*} . In Step 2, we prove that $\kappa_{t_i, r_{t_i}(a_i), m} \rightarrow \kappa_{t_i}$, by observing that $\bar{\phi}_{t_i, r_{t_i}(a_i), m}(c_0, t_{-i}, a_{-i}) \rightarrow (f(c_0, \mathbf{c}_i(t_i), \mathbf{c}_{-i}(t_i)), t_{-i})$.

In Step 3, we prove that $\Sigma : (t_i, r_{t_i}(a_i), m) \mapsto \{r_{t_i}(a_i)\}$ is closed under sequentially rational behavior in \tilde{G}^m , so that $r_{t_i}(a_i) \in ISR_i^\infty[t_i, r_{t_i}(a_i), m]$. To this end, for each $(t_i, r_{t_i}(a_i), m)$, we construct a conjecture $\tilde{\mu}$ of type $(t_i, r_{t_i}(a_i), m)$ against which $r_{t_i}(a_i)$ is a sequential best response and $\tilde{\mu}_{\emptyset}$ puts probability 1 on the graph of Σ , by setting

$$\tilde{\mu}_h = \mu_h^{t_i, r_{t_i}(a_i)} \circ \tilde{\phi}_{t_i, r_{t_i}(a_i), m}^{-1} \circ \gamma^{-1}$$

where

$$\gamma : (c_0, t_{-i}, r_{t_{-i}}(a_{-i}), m) \mapsto (c_0, t_{-i}, r_{t_{-i}}(a_{-i}), m, r_{t_{-i}}(a_{-i}))$$

stipulates that the types play according to Σ , and the mapping

$$(D.3) \quad \tilde{\phi}_{t_i, r_{t_i}(a_i), m} : (c_0, t_{-i}, a_{-i}) \mapsto \left(G_{t_i, r_{t_i}(a_i), m}(f(c_0, \mathbf{c}_i(t_i), \mathbf{c}_{-i}(t_i))), (t_{-i}, r_{t_{-i}}(a_{-i}), m) \right)$$

incorporates the transformation of c_0 . By construction, $\mu_{\emptyset}^{t_i, r_{t_i}(a_i)}$ puts probability 1 on the graph of Σ , and the belief induced on $\tilde{\Theta}^m \times \tilde{T}_{-i}$ by $\tilde{\mu}_{\emptyset}$ is $\kappa_{t_i, r_{t_i}(a_i), m}$. Towards showing that $r_{t_i}(a_i)$ is a sequential best response to $\tilde{\mu}$, we also observe that each $\tilde{\mu}_h$ induces probability distribution $\mu_h^{t_i, r_{t_i}(a_i)} \circ \left(\phi_{t_i, r_{t_i}(a_i), m, r_{-i}} \right)^{-1}$ on $\Theta \times A_{-i}$ —as in the proof of Proposition 1, where that belief was $\pi^{t_i, r_{t_i}(a_i)} \circ \left(\phi_{t_i, r_{t_i}(a_i), m, r_{-i}} \right)^{-1}$. One can then simply replace $\pi^{t_i, r_{t_i}(a_i)}$ with $\mu_h^{t_i, r_{t_i}(a_i)}$ in the remainder of the proof of that step, to show that $r_{t_i}(a_i)$ is a best response to $\tilde{\mu}_h$ at each history h that is not precluded by $r_{t_i}(a_i)$, showing that $r_{t_i}(a_i)$ is a sequential best response to $\tilde{\mu}$ for type $(t_i, r_{t_i}(a_i), m)$.

In Step 6, we use Lemma 2 instead of Lemma 8, to obtain a hierarchy $h_i(\bar{t}_i^m)$ in open neighborhood $(\tau_i^m)^{-1}(U_i)$ of $h_i(\bar{t}_i)$ such that each element of $ISR_i^\infty[\bar{t}_i^m]$ is m -equivalent to a_i^m and $h_i(\bar{t}_i^m) \in T_i^{Cm*}$, which is the subspace of T_i^{*m} in which it is common knowledge that $\theta \in f(C)$ and the true value of c_j is known by player j for each j . This leads to the type \hat{t}_i constructed in Step 7 to remain in T_i^{C*} and have a_i as the unique ISR action up to m -equivalence. \square

Though Proposition 8 is a close analogue of the general result on equilibria, Proposition 2, its proof is more closely analogous to the final steps in our result specific to repeated games, Proposition 4. (It is in the lemmas preceding that proof that the steps specific to repeated games occur.)

Proof of Proposition 8. In the proof of Proposition 4, modify the types $t_{j,m,l,\lambda}$ by substituting θ^* for θ_{δ,g^*} and taking $t^{a_j^*,l}$ to be the type in T_j^{C*} for whom a_j^* is uniquely ISR up to date l and $\mathbf{c}_j(t^{a_j^*,l}) = c_j^*$ (by Proposition 7). Take also $\mathbf{c}_j(t_{j,m,l,\lambda}) = c_j^*$, so that $h_j(t_{j,0,l,\lambda}) \in T_j^{C*}$. Moreover, as in the proof of Proposition 4, since playing according to a_j^* up to l is the unique sequential best response for type $t_{j,m,l,\lambda}$ when the others follow a_{-j}^* forever, we can take l' sufficiently large so that following a_j^* remains the unique sequential best response up to l when the others follow a_{-j}^* up to l' . As in the proof of Proposition 4, this shows that a_j^* is the unique ISR plan for type $t_{j,m,l,\lambda}$. Finally, as in the proof of Proposition 4, one can select m, l , and λ to satisfy the other properties in the proposition. \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 (2011): “Robustness to Incomplete Information in Repeated Games,” *Theoretical Economics* 6, 49-93.
- [5] Chen, Y. (2012): “A structure theorem for rationalizability in the normal form of dynamic games,” *Games and Economic Behavior*, 75, 587-597.
- [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., D. Kreps, and D. Levine (1988): “On the Robustness of Equilibrium Refinements,” *Journal of Economic Theory* 44, 354-380.
- [10] Fudenberg, D. and E. Maskin (1991): “On the Dispensability of Public Randomization in Discounted Repeated Games,” *Journal of Economic Theory* 53, 428-438.
- [11] Kajii, A. and S. Morris (1997): “The Robustness of Equilibria to Incomplete Information,” *Econometrica*, 65, 1283-1309.
- [12] 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.
- [13] Mailath, G. and L. Samuelson (2006): *Repeated Games and Reputations*, Oxford University Press.
- [14] Mertens, J. and S. Zamir (1985): “Formulation of Bayesian Analysis for Games with Incomplete Information,” *International Journal of Game Theory*, 10, 619-632.
- [15] Monderer, D. and D. Samet (1989): “Approximating Common Knowledge with Common Beliefs,” *Games and Economic Behavior*, 1, 170-190.
- [16] Oury, M. and O. Tercieux (2007): “Continuous Implementation,” *Econometrica*, forthcoming.
- [17] Penta, A. (2012a): “Higher Order Uncertainty and Information: Static and Dynamic Games,” *Econometrica*, 80, 631-660.
- [18] Penta, A. (2012b): “On the Structure of Rationalizability on Arbitrary Spaces of Uncertainty,” *Theoretical Economics*, forthcoming.
- [19] Rubinstein, A. (1982): “Perfect Equilibrium in a Bargaining Model,” *Econometrica*, 50, 97-109.
- [20] Rubinstein, A. (1989): “The Electronic Mail Game: Strategic Behavior Under ‘Almost Common Knowledge,’” *The American Economic Review*, 79, No. 3, 385-391.
- [21] Weinstein, J. and M. Yildiz (2007): “A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinements,” *Econometrica*, 75, 365-400.
- [22] Weinstein, J. and M. Yildiz (2011): “Sensitivity of Equilibrium Behavior to Higher-order Beliefs in Nice Games,” *Games and Economic Behavior* 72, 288–300.
- [23] Weinstein, J. and M. Yildiz (2012): “Properties of Interim Correlated Rationalizability in Infinite Games,” mimeo.

WEINSTEIN: NORTHWESTERN UNIVERSITY; YILDIZ: MIT