REPUTATION WITHOUT COMMITMENT

JONATHAN WEINSTEIN AND MUHAMET YILDIZ

Abstract. In the reputation literature, players have *commitment types* which represent the possibility that they do not have standard payoffs but instead are constrained to follow a particular plan. In this paper, we show that arbitrary commitment types can emerge from incomplete information about the stage payoffs. In particular, any finitely repeated game with commitment types is strategically equivalent to a standard finitely repeated game with incomplete information about the stage payoffs, such that the types with identical solutions have almost identical prior probability in two games. Then, classic reputation results can be achieved with uncertainty concerning only the stage payoffs.

JEL Numbers: C72, C73.

1. Introduction

The reputation literature relies on the existence of *commitment types*. These types are not strategic but are certain to follow a particular plan. Since the seminal work of Kreps, Milgrom, Roberts, and Wilson (1982) (henceforth, Gang of Four), it has been well-established that inclusion of commitment types may alter predicted outcomes dramatically, as this may entice the original "rational" types to imitate the commitment types, in order to form a reputation for playing according to the committed plan. Building on this insight, a large literature has emerged, with applications in a wide range of areas.[1]

Of course, commitment types can be modeled by using a payoff function that rewards a player who follows a specific plan. Nevertheless, such payoff functions often contradict the underlying structure of the original game. For example, the tit-for-tat types used by Gang

Date: First Version: August 2012; This Version: June 2014.

Weinstein: Washington University in St Louis; j.weinstein@wustl.edu. Yildiz: MIT Economics Department; myildiz@mit.edu. We thank the referees and the seminar participants at Harvard, Koc, Princeton, and Yale Universities, in particular, Dilip Abreu and Stephen Morris, for useful comments.

[1]We refer to the textbook of Mailath and Samuelson (2006) for a review.

of Four in the analysis of finitely repeated prisoners dilemma cannot be justified by a payoff function that is consistent with a repeated game, i.e. a (possibly discounted) sum of stage-game payoffs. The only commitment types that arise within such an additive structure are those who commit to playing the same action throughout the game, and these would not have any impact on behavior in a repeated prisoners dilemma game.

The justification for commitment types is important for the interpretation of the results. If one can derive an outcome (say cooperation in the finite prisoner's dilemma) using only types that arise from a natural payoff structure (such as a repeated game with different stage payoffs), one can then interpret the result as an outcome of incomplete information within a rational framework. On the other hand, some payoff functions, such as one that rewards a player only if he follows tit-for-tat, seem irrational in a repeated-game context. They may instead reflect psychological anomalies and super-game concerns (motivations that lie outside the game), such as maintaining reputation in a broader context. If one needs to introduce such payoff functions, then the result is probably best seen as reflecting the issues relating to the super-game concerns, psychological anomalies or irrationality that have been overlooked in the original model. (That may be why the commitment types are often referred to as "crazy types" and treated within the context of modeling irrationality, as in Kreps (1990).)

In this paper, we show that for any given plan, a commitment type who is *required* to follow this plan can be mimicked by a "twin" whose beliefs cause him to follow the same plan. Specifically, the twin knows it is common knowledge that payoffs follow a repeated-game strcuture, and his unique rationalizable action is to follow the given plan. Moreover, by embedding a collection of such twins into a single type space, every game with commitment types can be converted to a standard repeated game with incomplete information about the stage-game payoff function, such that the twins have prior probabilities almost identical to the commitment types. Therefore, any model of reputation formation in finitely repeated games, where players form a reputation for *commitment,* can be converted to a strategically equivalent model in which they form a reputation for certain *beliefs* about the stage-game payoffs. This is true provided that one allows sufficient variations in stage-game payoffs and considers a rich set of information structures.

Of course, one may also wish to restrict the the stage-game payoff functions. For example, in a standard prisoners' dilemma game, one might want to assume that it is common knowledge that cooperation is not dominant. Under such restrictions, twins may not exist for some commitment types. Indeed, we also prove an opposing benchmark, showing that one needs some amount of variations in the stage-game payoffs in order to have any reputational effect. We show that if the stage game is dominance-solvable and the stage game payoffs are restricted to be a sufficiently small neighborhood of the original stage-game payoff function, then the unique sequential equilibrium of the repeated game with incomplete information prescribes all players to repeat the stage-game solution throughout the game (as in the subgame-perfect equilibrium of the complete information version), regardless of the length of the game.

2. Preview of Results

In this section, we preview our result more carefully on the example analyzed by Gang of Four: the finitely-repeated prisoner's dilemma game in which player 1 may be committed to tit-for-tat, though this has small ex-ante probability.

Consider the repeated game in which the following prisoner's dilemma is repeated $\bar{t}$ times:

(PD)

	Cooperate	Defect
Cooperate	$5, 5$	$0, 6$
Defect	$6, 0$	$1, 1$

All the previous moves are publicly observable (perfect monitoring), and the payoff of a player in the repeated game is the sum of his payoffs in the stage game above. A simple application of backward induction in this game yields the play of (Defect, Defect) at every history. Indeed, it is well known that the only Nash equilibrium outcome is playing (Defect, Defect) at every period.

Gang of Four consider an incomplete information game G in which player 1 may be committed to playing tit-for tat. Player 1 has two types, a "rational" type τ_1^*, whose payoffs and available moves are as in the repeated-prisoners' dilemma game above, and a committed type τ_1^{T4T} which can only play tit-for-tat. That is, the latter type must play cooperate in the first round and imitate the last move of player 2 in the subsequent periods. The prior probability of τ_1^{T4T} is some small $\varepsilon > 0$. Player 2 still has one type τ_2^*, which is "rational" as in the original game. Gang of Four shows that in any sequential equilibrium of the new game each rational type τ_i^* must play Cooperate at all but few periods. (Interestingly, by varying the set of commitment types, one can generate a rich set of equilibrium behavior, obtaining a Folk Theorem for long games (Fudenberg and Maskin, 1986)).

As we mentioned in the introduction, one can replicate the above equilibrium behavior with payoff uncertainty by assigning the payoff function of τ_1^{T4T} as 1 at histories at which player 1 plays according to tit-for-tat and 0 at all other histories. Here, the solution concept is sequential equilibrium with the restriction that player 2 assigns probability 1 on τ_1^* off the path. Such a payoff function is incompatible with the repeated game payoff structure, and one cannot replicate the commitment to tit-for-tat by simply modifying the stage-game payoff function for τ_1^{T4T}. Indeed, such modifications can lead to only two commitment types: the type that plays Cooperate throughout and the type that plays Defect throughout. Commitment to cooperation can be justified by the stage-game payoff function

(CC)

	Cooperate	Defect
Cooperate	1	1
Defect	0	0

for example. The inclusion of such simple commitment types cannot affect the behavior of rational types in this game, though in other games such as Cournot duopoly, it could result in a player becoming a Stackelberg leader by convincing the other player he is committed to a certain action.

Fortunately, the austere information structure above is not the only structure we can consider. For any $\varepsilon' > \varepsilon$, applying our Proposition 1 to the game G in Gang of Four generates a game G' with the following properties.

Ex-ante Proximity: The prior probability of the rational type profile (τ_1^*, τ_2^*) is at least $1 - \varepsilon'$, and each τ_i^* knows that his stage-game payoffs are as in (PD).

Repeated-Game Structure: All types can play all strategies and maximize the sum of stage-game payoffs, which need not be as in (PD).

Strategic Equivalence: G and G' are strategically equivalent in the following sense.

(1) G' contains types τ_1^*, τ_2^*, a twin $\hat{\tau}_1^{T4T}$ of the tit-for-tat type τ_1^{T4T} in G, and a number of other new types (of both players) that we use to encode the beliefs of type $\hat{\tau}_1^{T4T}$.

(2) Though $\hat{\tau}_1^{T4T}$ is *allowed* to play any plan of action, tit-for-tat is his unique rationalizable plan.

(3) Rational type τ_1^* is certain that he faces the rational type τ_2^*, and the rational type τ_2^* in turn puts probability $1-\varepsilon$ on τ_1^* and probability ε on the twin $\hat{\tau}_1^{T4T}$ of τ_1^{T4T}.

By strategic equivalence property, the strategic situation the rational types face is the same as in G, except now τ_2^* thinks that $\hat{\tau}_1^{T4T}$ plays tit-for-tat as a result of some rational reasoning under incomplete information rather than as a result of commitment or an unconventional payoff function. Therefore, under the broad set of solution concepts that are invariant to such changes, the solution sets for rational types (τ_1^*, τ_2^*) are identical in G and G'. The conditional probabilities specified above are achieved by a prior distribution in G' putting probability $1-\varepsilon'$ on (τ_1^*, τ_2^*), $\varepsilon(1-\varepsilon')/(1-\varepsilon)$ on $\left(\hat{\tau}_1^{T4T}, \tau_2^*\right)$ and the remaining small probability $(\varepsilon'-\varepsilon)/(1-\varepsilon)$ on the newly constructed types.

Three points are worth emphasizing. First, when $\varepsilon'-\varepsilon$ is small compared to ε, the prior probabilities of (τ_1^*, τ_2^*) and $\left(\hat{\tau}_1^{T4T}, \tau_2^*\right)$ are approximately $1-\varepsilon$ and ε, respectively, with much smaller probability on the new types. Hence, the type spaces of G and G' are nearly identical, and the twin $\hat{\tau}_1^{T4T}$ assigns much larger probability to the standard type τ_2^* than to the new types. Despite this, $\hat{\tau}_1^{T4T}$ has a unique rationalizable plan because $\hat{\tau}_1^{T4T}$ believes that his own plan has non-negligible impact on his payoff only if he faces one of the newly constructed types. He finds these types unlikely, but they are likely enough to be his main concern. Second, the unique rationalizable plan emerges under intricate beliefs that require a large number of new types for encoding, especially when the game is long. Nonetheless, we are able to encode such beliefs by putting only a negligible amount of prior probability on the new types. Third, our proof is based on reward and punishment mechanisms typical in repeated games as well as theories of social learning that rely on the existence of payoff functions as in (CC). One could object that such stage-games are too different from the original prisoners' dilemma game to be conceivable.

Accordingly, our second benchmark (Proposition 5) allows only small variations in the stage-game payoffs. Since Nash equilibrium is robust to such perturbations and (Defect, Defect) is the only Nash equilibrium here, one cannot expect to have any discontinuity in that front. This is similar to the fact that for a given $\bar{t}$, (Defect, Defect) becomes the unique equilibrium outcome as $\varepsilon \to 0$ in the Gang of Four example. Nevertheless, Gang of Four demonstrates that such continuity is misleading: for any given ε, long cooperation is obtained when $\bar{t}$ is large enough. In this paper, we prove something stronger: a *uniform* continuity result with respect to the stage game payoffs. The distinction in these results arises because we impose common knowledge of the constraint on payoffs. According to our result, when the stage-game payoffs are sufficiently close to (PD) throughout the type space, (Defect, Defect) is the only sequential equilibrium outcome, no matter how large $\bar{t}$ is.

Therefore, one needs to allow some substantial amount of variation in stage-game payoffs in order to provide an incomplete-information foundation for the commitment types. While the amount of necessary variation may depend on the details of the game and the commitment types at hand, our main result shows that one can always provide such a foundation as long as there is enough variation in allowable stage-game payoff functions.

Our results here build on our previous work on non-robustness in repeated games. In Weinstein and Yildiz (2013) we showed that, in any infinitely repeated game, any individually rational and feasible outcome is the *unique* rationalizable outcome of an appropriately chosen perturbation which maintains common knowledge of the repeated-game structure and discounting criterion. A key lemma leading to this result showed that for any plan whatsoever, there is a type who follows this plan as a unique rationalizable action, although he believes in common knowledge of the repeated-game structure. An extension of this lemma to finitely repeated games plays an important role in our construction.

Our contribution to our previous work is twofold. First, extending the above lemma to finitely repeated games involves a considerably more nuanced construction, as providing effective future incentives is harder in finitely-repeated games. Second, and more importantly, perturbations considered in the two work are very different: here we use the ex-ante perturbations that are used in the traditional reputation literature (and also by Kajii and Morris (1997) in the context of robustness), while in Weinstein and Yildiz (2013) we use perturbations of interim beliefs in universal type space. In general, ex-ante perturbations are significantly more restrictive, and the two approaches often yield quite different results. One reason the results here can be achieved with ex-ante perturbations is that our construction centers around perturbing the commitment types, who do not have set beliefs. The main difficulty turns out to be embedding types constructed in the lemma into a common-prior model without affecting the types' rationalizable actions, while keeping the ex-ante probabilities of the new types arbitrarily small.

We introduce the basic definitions and formulations in Section 3. In Section 4, we present our construction of a new type space in which the commitment types are replaced by types for which the committed action plan is uniquely rationalizable. In Section 5, we show that, for the original rational types, the constructed game is strategically equivalent to the model with commitment types, under a very broad set of solution concepts. We present the most general version of our result in Section 6. After presenting our continuity result in Section 7, we conclude in Section 9. Some of the more complicated proofs are relegated to the Appendix.

3. Basic Definitions

We study, for simplicity, a standard 2-player finitely repeated game with perfect monitoring and normal-form stage games; see Section 6 for the general case. We write $N = \{1, 2\}$ for the set of players, $T = \{0, 1, \ldots, \bar{t}\}$ for the set of dates t, and fix a finite set $A = A_1 \times A_2$ of stage-game action profiles $a = (a_1, a_2)$.[2] Note that, since we have perfect monitoring, the non-initial histories in the repeated game are of the form $h^t = (a^0, \ldots, a^{t-1})$ where $a^{t'}$ denotes the stage-game strategy profile played at date $t' \in T$. We write h^0 for the empty initial history, and write H for the set of all non-terminal histories. An outcome path or a terminal history is a list $\left(a^0, \ldots, a^{\bar{t}}\right)$; the set of all terminal histories is denoted by Z.

The payoff vector from an outcome path $\left(a^0, a^1, \ldots, a^{\bar{t}}\right)$ in repeated game is simply the sum[3] of the stage game payoffs:

$$u\left(a^0, a^1, \ldots, a^{\bar{t}} | g\right) = g\left(a^0\right) + g\left(a^1\right) + \cdots + g\left(a^{\bar{t}}\right), \tag{3.1}$$

where the function $u = (u_1, u_2)$ denotes the payoffs from the repeated game and the function $g = (g_1, g_2)$ denotes the payoffs from the stage game. While the particular stage payoffs are not necessarily known, this formula will be common knowledge throughout the games we study here. That is, it is common knowledge that the stage payoff function g is fixed throughout the game and that the players simply maximize the sum of these payoffs.

[2]Following the convention in game theory, we write $-i$ for the player $j \neq i$ and drop the subscript to denote the profiles, e.g., $x = (x_1, x_2) \in X = X_1 \times X_2$ and $X_{-1} = X_2$.

[3]Discounting would not affect our results; setting the discount rate to 1 simplifies the expressions.

We write $\mathcal{G} = [0,1]^A$ for the set of all possible stage-game payoff functions $g_i : A \to [0,1]$. Here, we put a uniform bound on the stage game payoffs so that small variations of the probability distributions on stage payoffs lead to small variations in expected payoffs, as in the reputation literature. This restriction strengthens our results.

We fix a complete-information repeated game in which it is common knowledge that the stage-game payoffs are a fixed (g_1^*, g_2^*). The payoff function in the repeated game is $u(\cdot|g^*)$, given by the formula in (3.1). This could, for example, be the repeated prisoner's dilemma game, with g^* defined as in (PD).

In the complete-information game, a strategy of a player i is a mapping $s_i : H \to A_i$, which maps each non-terminal history to a strategy in the stage game. Because we analyze incomplete information games, however, we will avoid the word strategy for this mapping and call it instead an *action plan,* because we reserve the word *strategy* for mappings from types to action plans. (We refer to the strategies in the stage game as *moves.*) The set of all action plans is denoted by S_i. The outcome path induced by a profile (s_1, s_2) is denoted by $z(s_1, s_2)$. We also allow (behavioral) mixed strategies and write Σ_i for the set of mixed action plans $\sigma_i : H \to \Delta(A_i)$ for player i.

We consider two kinds of elaboration, corresponding to two distinct ways in which the common-knowledge assumption in the complete information game may be relaxed. The first notion of relaxation is the one considered in the reputation literature—through commitment types.

Definition 1. An *ε-elaboration with (one-sided) commitment types* is a Bayesian game such that

- the sets of types for players 1 and 2 are $\{\tau_1^*\} \cup C$ and $\{\tau_2^*\}$, respectively, where $C \subset S_1$;
- the probability of τ_1^* is $1 - \varepsilon$ and the probability of each $c \in C$ is some $\pi(c) \in [0,1]$ that add up to ε;
- the set of moves available to τ_i^* is as in the repeated game above, while the only available action plan for type $c \in C$ is c,
- the payoffs are as in the complete information game.

Such an elaboration can be denoted by a pair (C, π) of commitment types C and a vector of probabilities on C.

Here, each action plan $c \in C$ corresponds to a type of Player 1 who can only play c. The incomplete information is only about whether Player 1 can play all action plans or has committed to a particular action plan. The type τ_1^* that can play all plans is called the *rational* type while the types $c \in C$, who can play only according to one plan of action, are called *commitment* types. The commitment types are also sometimes called *crazy* types.

A commitment type c also be modeled by a payoff function that gives 1 if Player 1 plays according to c throughout and 0 if he ever deviates. Such a payoff function clearly violates the additive structure in (3.1). Indeed, except for the trivial commitment types who play the same stage-game strategy throughout the repeated game, commitment types cannot be justified by a new payoff function that has the additive structure above.

We will next introduce type spaces in which all payoff functions satisfy the additive structure of (3.1). While they enforce this restriction on the repeated structure of the game, these spaces are also "richer" in that they may have a larger variety of types. Our type spaces will have two-sided incomplete information with a common prior. Formally, by a type space, we mean a list $(\mathcal{T}, \pi, \boldsymbol{g})$ where $\mathcal{T} = \mathcal{T}_1 \times \mathcal{T}_2$ is the set of type profiles $\tau = (\tau_1, \tau_2)$, $\pi \in \Delta(\mathcal{T})$ is the common prior on type profiles, and $\boldsymbol{g} = (\boldsymbol{g}_1, \boldsymbol{g}_2)$ is the profile of type-dependent stage-game payoff functions $\boldsymbol{g}_i : \mathcal{T}_i \to \mathcal{G}$. [4] A *Bayesian repeated game* (without commitment types) is a list $(N, A, (\mathcal{T}, \pi, \boldsymbol{g}))$. We should emphasize that this notation suppresses many important common-knowledge assumptions, such as the fact that the game is repeated, all previous actions are publicly observable (i.e. perfect monitoring), and the payoffs in the repeated game are given by the formula (3.1). We also note that players do not observe either their own or other players' payoffs at each stage of the game. A strategy of a player i in a Bayesian repeated game $(N, A, (\mathcal{G}, \mathcal{T}, \pi))$ is a mapping $\sigma_i : \mathcal{T}_i \to \Sigma_i$.

The second notion of elaboration allows richer type spaces and two-sided incomplete information, but does not allow any payoff function outside of the additive structure in (3.1):

Definition 2. An *ε-elaboration without commitment types* is a Bayesian game $(N, A, (\mathcal{T}, \pi, \boldsymbol{g}))$ with distinguished types τ_1^*, τ_2^* where

(1) $\boldsymbol{g}_i(\tau_i^*) = g_i^*$ for each $i \in N$ and
(2) $\pi(\tau^*) = 1 - \varepsilon$.

[4] Here, $\Delta(X)$ denotes the set of all probability measures on the finite set X.

The first and second conditions state that the original complete information game is embedded in the elaboration and has a high ex-ante probability of $1-\varepsilon$. The last condition states that the rational types (τ_1^*, τ_2^*) know their payoffs, and their payoffs are as in the original complete information game. The novelty in this definition is that the payoffs under all possible specifications in the Bayesian game satisfy the additive repeated game structure of (3.1). That is, the formula (3.1) remains common knowledge. In that sense, all the types in an elaboration without commitment types are rational, although we reserve the term rational for types (τ_1^*, τ_2^*) as in the elaborations with commitment types.

Both elaborations above fall under the category of ε-elaborations as defined by Kajii and Morris (1997). An ε-elaboration without commitment types is a Kajii-Morris elaboration with the additional restriction that the formula (3.1) is common knowledge. While ε-elaborations with commitment types were presented above in terms of uncertainty about the strategies, they could also be represented as Kajii-Morris elaborations with specific simple type space in which the formula (3.1) fails.

Finally, we review two standard concepts in game theory. First, for any Bayesian game, interim correlated rationalizability (henceforth ICR) is the outcome of iterated elimination of action plans for types that are never a weak best response, as defined by Dekel, Fudenberg, and Morris (2007). We write $S_i^\infty[\tau_i|G]$ for the set of all interim correlated rationalizable action plans for type $\tau_i \in \mathcal{T}_i$ in game $G = (N, A, (\mathcal{G}, \mathcal{T}, \pi))$. We will give a more detailed definition of ICR later in the construction. We just note here that ICR is the weakest known rationalizability concept for Bayesian games, and all the action plans that are played by a type with positive probability in any equilibrium are ICR for that type.

Second, we say that action plans s_i and s_i' are *equivalent* if $z(s_i, s_{-i}) = z(s_i', s_{-i})$ for all action plans $s_{-i} \in S_{-i}$, i.e., they lead to the same outcome no matter what strategy the other player plays. Note that s_i and s_i' are equivalent iff $s_i(h^t) = s_i(h^t)$ for every history h^t in which i played according to s_i throughout; they may differ only in their prescriptions for histories that they preclude. Hence, in reduced form, action plans can be represented as mappings that maps the history of other players' play into own stage game actions. We write $\bar{S}_i$ for the set of reduced-form action plans $\bar{s}_i$; these map each $\left(a_{-i}^l\right)_{0\le l<t}$ to some action $a_i \in A_i$ in the stage game.

4. Irrelevance of Commitment Types

In this section, we show that one can replace each commitment type c with a regular type for which c is uniquely rationalizable with a slight perturbation of the prior distribution. Consequently, one can transform any elaboration with commitment types to an elaboration without commitment types, such that, from the point of view of rational types who believe in the ICR concept, the two elaborations are identical. Under ICR as well as a broader set of solution concepts, this will lead to the same set of solutions for the rational types in the two elaborations.

Proposition 1. *For any $\varepsilon, \varepsilon' \in (0,1)$ with $\varepsilon' > \varepsilon$ and for any ε-elaboration G with commitment types (C, π) there exists a ε'-elaboration $G' = (N, A, (\mathcal{G}, \mathcal{T}, \pi'))$ without commitment types in which the commitment types are replaced by types with unique rationalizable action plans:*

(1) $\pi'(g^*, \tau_2^* | \tau_1^*) = 1$ *and* $\pi'(g^*, \tau_1^* | \tau_2^*) = \pi(g^*, \tau_1^* | \tau_2^*) = 1 - \varepsilon$, *and*
(2) *for every $c \in C$ there exists $\tau_1^c \in \mathcal{T}_1$ such that all ICR action plans of τ_1^c are equivalent to c, and* $\pi'(\tau_1^c | \tau_2^*) = \pi(c | \tau_2^*) = \pi(c)$.

Here, the first condition establishes that the interim beliefs of rational types regarding their own payoffs and "rationality" of their opponents are identical in the two elaborations. The second condition establishes that each commitment type c is replaced by a type τ_1^c for which following c is uniquely rationalizable, and that the rational type of player 2 in G' assigns the same probability to the type τ_1^c as the rational type in G assigns to the commitment type c. Note that $\pi'(\tau_1^c | \tau_2^*) = \sum_{g_1} \pi'((g_1, g_2^*), \tau_1^c | \tau_2^*)$, meaning that type τ_2^* knows his own payoff function.

The proposition thereby establishes that any ε-elaboration G with commitment types can be converted into an ε'-elaboration by replacing commitment types c with types τ_1^c for whom c is the only rationalizable action plan in reduced form. These types follow c not because they are committed or have payoffs that are inconsistent with playing a repeated game but because their reasoning under their information leads them to do so. Moreover, from the point of view of the rational types these are the only types with positive probability, mirroring the elaboration with commitment types.

The equivalence is established despite the following constraints:

(1) The repeated-game payoff structure is maintained throughout G'. That is, it is common knowledge throughout that the payoff in the repeated game is the sum of the payoffs in the stage game, and that the stage game is fixed throughout the game. Type τ_1^c knows all this and yet follows c as its unique rationalizable plan.
(2) The ex-ante distribution π' in G' can be arbitrarily close to the distribution π in G, in that ε' can be arbitrarily close to ε.

Note that the commitment plans c can be arbitrarily complex. Hence, the types τ_1^c that replace the commitment types c may hold highly complicated beliefs, and the elaboration without commitment types may contain a large number of new types that are used to encode these beliefs within the standard common prior type space. Despite this, the second condition requires that the ex-ante probability of these types be made arbitrarily small.

Proof of Proposition 1. The first step in our construction is the following lemma, which establishes that any given action plan s_i is the only rationalizable action for a type $\tau_i^{s_i}$ from some common prior model. (The proof of Lemma 1 is the lengthiest step of the proposition and is relegated to the appendix.)

Lemma 1. *For any $s_i \in S_i$, there exists a Bayesian repeated game $G^{s_i} = (N, A, (\mathcal{G}^{s_i}, \mathcal{T}^{s_i}, \pi^{s_i}))$ with a type $\tau_i^{s_i} \in \mathcal{T}_i^{s_i}$ such that*

(1) *$\pi^{s_i}(g, \tau) > 0$ for every $(g, \tau) \in \mathcal{G}^{s_i} \times \mathcal{T}^{s_i}$ and*
(2) *for every action plan s_i', $s_i' \in S_i^{\infty}[\tau_i^{s_i} | G^{s_i}]$ if and only if s_i' is equivalent to s_i.*

By relabeling if necessary, we take all of the types above to be distinct from each other and from τ^*, fixing also a unique type $\tau_i^{s_i}$ for each s_i. We construct $G' = (N, A, (\mathcal{G}', \mathcal{T}', \pi'))$ by setting

$$
\begin{aligned}
\mathcal{G}' &= \{g^*, (\mathbf{0}, g_2^*)\} \cup \bigcup_{c \in C} \mathcal{G}^c \\
\mathcal{T}_i' &= \{\tau_i^*\} \cup \bigcup_{c \in C} \mathcal{T}_i^c \qquad (\forall i \in N) \\
\pi'(g, \tau) &= \begin{cases} 1 - \varepsilon' & \text{if } (g, \tau) = (g^*, \tau^*), \\ \frac{1-\varepsilon'}{1-\varepsilon} \pi(c) & \text{if } (g, \tau) = ((\mathbf{0}, g_2^*), (\tau_1^c, \tau_2^*)), \\ \frac{\varepsilon' - \varepsilon}{(1-\varepsilon)|C|} \pi^c(g, \tau) & \text{if } (g, \tau) \in \mathcal{G}^c \times \mathcal{T}^c, \\ 0 & \text{otherwise,} \end{cases}
\end{aligned}
$$

where

$$\mathbf{0}\,(a) = 0 \qquad (\forall a \in A)\,.$$

We now observe that G' satisfies the properties in the proposition. Indeed, rational type τ_1^* of Player 1 assigns probability 1 on (g^*, τ_2^*). Likewise, we have

$$\pi'\left(\mathcal{G}' \times \{\tau_2^*\}\right) = 1 - \varepsilon' + \frac{1-\varepsilon'}{1-\varepsilon} \sum_{c \in C} \pi(c) = 1 - \varepsilon' + \frac{1-\varepsilon'}{1-\varepsilon}\varepsilon = (1-\varepsilon')\,/\,(1-\varepsilon)$$

and therefore, in the interim, τ_2^* assigns probability $1-\varepsilon$ to (g^*, τ_1^*) and probability $\pi\,(c)$ to τ_1^c for each c. On the other hand, since the beliefs of type τ_1^c altered substantially when $(\mathcal{G}^c, \mathcal{T}^c, \pi^c)$ was incorporated in G', it is not clear that τ_1^c follows c as the unique ICR action. The next lemma states that this is indeed the case.

Lemma 2. *For any $c \in C$, $i \in N$, and any $\tau_i \in T_i^c$, $S_i^\infty\,[\tau_i|G'] = S_i^\infty\,[\tau_i|G^c]$; in particular, $S_1^\infty\,[\tau_1^c|G'] = c$.*

This lemma completes the proof of the proposition; its proof is in the appendix. □

Our proof has two main steps. The first, found in Lemma 1, is to construct a type space in which a given action plan is uniquely rationalizable for a type. We constructed such a type space in Weinstein and Yildiz (2013) for infinite-horizon repeated games, but without requiring that the constructed type space has a common prior, a property that is essential for our proposition here. In this paper, using the ideas in that construction, we first construct such a type space for finite-horizon games without common prior and then convert it to a common-prior type space, using this time the ideas and the results developed by Lipman (2003) and Weinstein and Yildiz (2007).

The main economic ideas involved in these constructions come from social learning and reward and punishment mechanisms in repeated games. First, for the class of action plans that is consistent with learning in a single decision problem, we construct games in which players whose payoffs depend only on their own action update their optimal stage game strategies as they observe the past behavior of the other players and learn about their own payoff from the other players' moves, as in social learning. We show that such processes can generate any action plan that is consistent with individual learning as the unique ICR action plan. The generated action plans can be arbitrarily complex because individual learning only puts a relatively mild restriction on behavior, similar to the sure-thing principle. Accordingly, the type space generating these behavior can be highly complex, leading players to update their beliefs about own payoffs, other players' beliefs, and other players' beliefs about beliefs and so on, as they observe the other players' moves. These ideas apply to the finite and infinite repeated games, and the proof is nearly identical for both cases. For infinitely repeated games, in Weinstein and Yildiz (2013) we extend the result to all action plans (including plans that contradicts the condition for individual learning), using a reward and punishment mechanism. Unfortunately, it is harder to come up with effective reward and punishment mechanisms for finite horizon games. After all, one cannot provide any future incentive in the last period. Hence, here, we use a more nuanced construction that combines social learning with a reward and punishment mechanism to extend the result to all action plans in finitely repeated games.

The second main step is to incorporate the above type spaces in one common prior model, replacing each commitment type with one of these type spaces. One must do this in such a way that (i) the original complete-information game still has high prior probability $(1-\varepsilon')$, (ii) the interim beliefs of the rational types are as in the original elaboration with commitment types, and (iii) the types' rationalizable behavior in the constructed type space remain the same after incorporating them into common prior model. The conditions (ii) and (iii) oppose each other, making the construction more difficult. To see this, note that (i) and (ii) require that the common prior π' puts a high probability on τ_1^c, requiring that probability to be $\frac{1-\varepsilon'}{1-\varepsilon}\pi(c)$ as in our proof. When ε and ε' are close, this probability is approximately $\pi(c)$. When ε and ε' are close, this also requires that π' puts a very small probability on $\mathcal{T}^c$, the original type profiles in the constructed type space in the first step. That probability can be at most $(\varepsilon'-\varepsilon)/(1-\varepsilon)$, which is negligible with respect to $\frac{1-\varepsilon'}{1-\varepsilon}\pi(c)$ when ε and ε' are close. These constraints make the belief of type τ_1^c in game G' substantially different from the belief of the type τ_1^c in game $\mathcal{G}^c$. In our construction, type τ_1^c in game G' assigns probability

$$p^c = \frac{|C|(1-\varepsilon')\pi(c)}{|C|(1-\varepsilon')\pi(c)+(\varepsilon'-\varepsilon)\pi^c(\tau)} \tag{4.1}$$

on type τ_2^*. Note that, for fixed $\pi(c)$, when $\varepsilon'-\varepsilon$ approaches 0, p^c approaches 1.[5] In contrast, τ_1^c in game G^c assigns zero probability on τ_2^*. Consequently, the belief hierarchies

[5]Note also that the technique we use in transforming the model without common-prior to the one with common prior also renders $\pi^c(\tau_1^c)$ small, bringing p^c near 1 even when ε and ε' are far apart.

of the types in G' can be quite different from the belief hierarchies of the types in G^c with the same label, which could lead to distinct set of ICR actions. We circumvent this problem with the following trick. We set the beliefs such that, whenever player 2 has type τ_2^*, the payoff of type τ_1^c is 0 for every move in the stage game, making him indifferent among all outcomes. Since $p^c < 1$, his best responses are identical to his best responses conditional on the type of player 2 being other than τ_2^*, thereby replicating the best responses of his twin in G^c. Since this was the only difference between the two type spaces, the rationalizable actions turn out to be identical in games G^c and G', as shown formally by Lemma 2.

Roughly speaking, from the point of view of rational types, Proposition 1 replaces commitment types by types who follow the same plans as their unique rationalizable plan. Hence, under any rationalizable solution concept, the rational types face the same strategic uncertainty in both games leading the same set of possible behavior. We will next establish such strategic equivalence formally.

5. Strategic Equivalence

In this section, we show that, in Proposition 1, the elaborations G with commitment types and G' without commitment types are "strategically equivalent" for rational types. By this we mean that, for a broad set of solution concepts to be delineated, the set of solutions for each rational type are identical in games G and G'. Therefore, the same set of behavior can be supported by reputational models regardless of whether one allows commitment types. In other words, the same set of behavior is supported whether one allows payoff functions that are inconsistent with the repeated-game structure or imposes this structure throughout. This will further imply, as we detail in Section 5.3, that the predictions of models with commitment types are nearly indistinguishable from that of those without commitment types.

Clearly, our result here applies to any solution concept that is invariant to replacing commitment types with types that have unique rationalizable action plans (in reduced form). In general Bayesian games, this invariance condition is somewhat stronger than elimination of non-rationalizable strategies, because the new game contains some new types, encoding the beliefs of the types with unique rationalizable plans. We first establish our result for a general class of such invariant solution concepts. We also establish the same strategic equivalence for sequential equilibrium; this requires an additional off-path belief restriction commonly imposed in the reputation literature.

5.1. **Strategic Equivalence under Invariant Solutions.** The following definitions are standard: A solution concept Σ maps every Bayesian game G to a set $\Sigma(G)$ of mixed strategies in game G. For any type spaces $\mathcal{T}$ and $\mathcal{T}'$ with $\mathcal{T} \subset \mathcal{T}'$ and any strategy profile σ on $\mathcal{T}'$, σ_T denotes the restriction of σ to $\mathcal{T}$. In the following definitions, we also use the convention that two probability distributions that have common support and agree on this support are identical, ignoring any difference in domains.

Definition 3. A solution concept Σ is said to be *invariant to elimination of non-rationalizable strategies* if and only if

$$\Sigma(G) = \Sigma(G')$$

for any two games G and G' with identical type spaces such that (i) if an action plan s_i is available for a type τ_i in game G then s_i is available for τ_i in G' and (ii) if s_i is not available for τ_i in G then $s_i \notin S_i^\infty[\tau_i|G']$.

Definition 4. A solution concept Σ is said to be *invariant to trivial enrichments of the type spaces* if and only if

$$\Sigma(G) = \{\sigma_{\mathcal{T}}|\sigma \in \Sigma(G')\}$$

for any two games G and G' with type spaces $\mathcal{T}$ and $\mathcal{T}'$ such that (i) $\mathcal{T} \subset \mathcal{T}'$, (ii) every type in $\mathcal{T}$ has identical set of available action plans in games G and G', and (iii) any type in $\mathcal{T}$ with multiple action plans has identical interim beliefs in games G and G'.

Note that the transformation in the first definition allows only elimination of non-rationalizable actions and the transformation in the second definition allows only inclusion of new types such that the types who put positive probability to the new types are trivial in that they can play only according to one plan. Proposition 1 implies that under any solution concept that is invariant to the above transformations, elaborations with or without commitment types have the same strategic implications for rational types. Due to its importance, we state this corollary as a proposition:

Proposition 2. *Let Σ be a solution concept that is invariant to elimination of non-rationalizable strategies and to trivial enrichment of the type spaces. Then, for any $\varepsilon, \varepsilon' \in (0,1)$ with $\varepsilon' > \varepsilon$ and for any ε-elaboration G with commitment types, there exists an ε'-elaboration $G' = (N, A, (\mathcal{G}, \mathcal{T}, \pi'))$ without commitment types such that*

$$\{\sigma(\tau^*) \,|\sigma \in \Sigma(G)\} = \{\sigma(\tau^*) \,|\sigma \in \Sigma(G')\},$$

i.e., the set of solutions for rational types are identical in games G and G'.

Proof. Note that, in Proposition 1, the elaboration G' can be obtained from G by (1) introducing new types such that only committed types believe in the new types, and (2) allowing commitment types to play any action plan in the repeated game. The first step is a trivial enrichment as in Definition 3 and the second undoes an elimination covered by Definition 4, so the conclusion follows. □

5.2. **Strategic Equivalence under Sequential Equilibrium.** We will next establish the same strategic equivalence under sequential equilibrium, which is defined as follows. Given any Bayesian repeated game with a type space $(\mathcal{G}, \mathcal{T}, \pi)$, a *belief structure* is a list $\mu = (\mu_{i,\tau_i,h})_{i\in N, \tau_i \in \mathcal{T}_i, h\in H}$ of type specific beliefs $\mu_{i,\tau_i,h} \in \Delta(\mathcal{G} \times \mathcal{T}_{-i})$ regarding the underlying payoffs and the other player's types, beliefs that vary with the history of play.[6] An assessment is a pair $(\tilde{\sigma}, \mu)$ of strategy profile $\tilde{\sigma} : \mathcal{T} \to \Sigma$ and a belief structure μ. An *assessment* $(\tilde{\sigma}, \mu)$ is said to be *sequentially rational* if $\tilde{\sigma}_i(\cdot|\tau_i)$ is a sequential best response to $\mu_{i,\tau_i,h}$ and $\tilde{\sigma}_{-i}$, i.e., the restriction of $\tilde{\sigma}_i(\cdot|\tau_i)$ to the continuation game after every history h is a best response to $\tilde{\sigma}_{-i}$ and the beliefs $\mu_{i,\tau_i,h}$ in the continuation game. An assessment $(\tilde{\sigma}, \mu)$ is said to be *consistent* if there exists a sequence $(\tilde{\sigma}^n, \mu^n) \to (\tilde{\sigma}, \mu)$ such that $\tilde{\sigma}^n$ assigns positive probability to each available move at every history and μ^n is derived from Bayes' rule and $\tilde{\sigma}^n$. An assessment $(\tilde{\sigma}, \mu)$ is said to be a *sequential equilibrium* if it is sequentially rational and consistent.

In an ε-elaboration without commitment types, sequential equilibria are defined as above. In an ε-elaboration with commitment types, the definition of course depends on how one formalizes the commitment types. In particular, the definition above implies that Player 2 puts probability 1 on the rational type of Player 1 if the history is not consistent with

[6] A more general definition of a belief structure would also specify the beliefs regarding past actions, but those beliefs are trivial because of perfect monitoring.

any commitment type—even when the history is also inconsistent with the strategy of the rational type. This is because the commitment types have only one action, so that only the rational types may tremble. This is an additional assumption when the commitment types are represented by payoff perturbations (violating the additive repeated game structure). In general, the possible off-the-path beliefs can vary depending on the way the commitment types are formulated, but the above assumption is usually maintained. We will keep this additional assumption in our definition for sequential equilibrium without commitment types:

Assumption 1. *For every history $h = (a^0, \ldots, a^{t-1})$,*

$$\mu_{2,\tau_2^*,h}\left(g^*, \tau_1^*\right) = 1$$

whenever h has zero probability under every type $\tau_1 \neq \tau_1^*$.

In our analysis we will focus on the behavior of the rational types under sequential equilibrium, which is formally defined as follows.

Definition 5. For any elaboration G (with or without commitment types), we write

$$SE^*\left(G\right) = \left\{\sigma\left(\cdot|\tau^*\right) | \left(\sigma, \mu\right) \text{ is a sequential equilibrium of } G \text{ that satisfies Assumption 1}\right\}$$

for the set of sequential equilibrium action plans for the rational types in G.

We are now ready to state the strategic equivalence result for sequential equilibrium.

Proposition 3. *For any $\varepsilon, \varepsilon' \in (0,1)$ with $\varepsilon' > \varepsilon$ and for any ε-elaboration G with commitment types (C, π) there exists an ε'-elaboration $G' = (N, A, (\mathcal{G}, \mathcal{T}, \pi'))$ without commitment types such that*

$$SE^*\left(G\right) = SE^*\left(G'\right),$$

i.e., under Assumption 1, the set of sequential equilibrium action plans for the rational types is same in games G and G'.

Proof. We will show that both conditions $\sigma\left(\cdot|\tau^*\right) \in SE^*\left(G\right)$ and $\sigma\left(\cdot|\tau^*\right) \in SE^*\left(G'\right)$ are characterized by the following conditions, (SR1) and (SR2). First, (σ, μ) is a sequential equilibrium of G if and only if the following three conditions are satisfied. The consistency condition for τ_2^* is

$$\text{(C)} \qquad \mu_{2,\tau_2^*,h}(c) = \mu_h^{\sigma\left(\cdot|\tau_1^*\right)}(c) \equiv \begin{cases} \frac{\pi(c)}{\Pr\left(h|\sigma\left(\cdot|\tau_1^*\right)\right)(1-\varepsilon)+\sum_{c'\in C^h}\pi(c')} & \text{if } c\in C^h \\ 0 & \text{otherwise} \end{cases} \qquad (\forall h, c)$$

where C^h is the set of commitment plans $c \in C$ that is consistent with history h. Of course, $\mu_h^{\sigma\left(\cdot|\tau_1^*\right)}(\tau_1^*) = 1 - \sum_{c\in C}\mu_h^{\sigma\left(\cdot|\tau_1^*\right)}(c)$. The consistency condition for player 1 is trivial, as player 2 has only one type. Note that $\mu_h^{\sigma\left(\cdot|\tau_1^*\right)}$ is a function of $\sigma(\cdot|\tau_1^*)$, and hence the following sequential rationality conditions are solely on $\sigma(\cdot|\tau^*)$. The sequential rationality conditions are

(SR1): $\sigma(\cdot|\tau_1^*)$ is a sequential best response to $\sigma(\cdot|\tau_2^*)$ under g_1^*, and

(SR2): at each history h, $\sigma(\cdot|\tau_2^*)$ is conditional best response to the mixed strategy

$$\tilde{\sigma} \equiv \mu_h^{\sigma\left(\cdot|\tau_1^*\right)}(\tau_1^*)\,\sigma(\cdot|\tau_1^*) + \sum_{c\in C}\mu_h^{\sigma\left(\cdot|\tau_1^*\right)}(c)\,c$$

under g_2^*.

Since all the other types are committed to a single plan, there are no other conditions. This shows that $\sigma(\cdot|\tau^*) \in SE^*(G)$ if and only if (SR1) and (SR2) are satisfied.

To show that $\sigma(\cdot|\tau^*) \in SE^*(G')$ implies the conditions (SR1) and (SR2), consider any sequential equilibrium (σ, μ') of G' that satisfies Assumption 1. Firstly, since type τ_1^* puts probability one on (g^*, τ_2^*), the sequential rationality condition for that type is (SR1). Secondly, since c is the unique rationalizable action plan of τ_1^c in G' (by Lemma 2) on all histories h consistent with c,

$$\text{(5.1)} \qquad \sigma(c(h)|h, \tau_1^c) = 1 \qquad \left(\forall c \in C^h, \forall h\right).$$

Hence, by Assumption 1 and consistency,

$$\text{(5.2)} \qquad \mu'_{2,\tau_2^*,h}(\tau_1^c) = \mu_h^{\sigma\left(\cdot|\tau_1^*\right)}(c) \qquad (\forall h, c),$$

which of course also implies that $\mu'_{2,\tau_2^*,h}(\tau_1^*) = \mu_h^{\sigma\left(\cdot|\tau_1^*\right)}(\tau_1^*)$. By (5.1) and (5.2), under the belief of type τ_2^*, player 1 plays according to $\tilde{\sigma}$ above, and the sequential rationality condition for type τ_2^* is (SR2).

To show that (SR1) and (SR2) are sufficient for $\sigma(\cdot|\tau^*) \in SE^*(G')$, take any $\sigma(\cdot|\tau^*)$ that satisfies (SR1) and (SR2). We will construct a sequential equilibrium (σ, μ') of G' that satisfies Assumption 1. Set $\mu'_{1,\tau_1^*,h}(g^*, \tau_2^*) = 1$ and $\mu'_{2,\tau_2^*,h} = \mu_h^{\sigma(\cdot|\tau_1^*)}$. For each $c \in C$, consider a sequential equilibrium (σ^c, μ^c) of the game in which the action plan of type τ_2^* is fixed as $\sigma(\cdot|\tau_2^*)$—as moves of nature, and the type space is $\mathcal{T}^c$ with the interim beliefs in G'. Set $\sigma(\cdot|\tau_i) = \sigma^c(\cdot|\tau_i)$ and $\mu'_{i,\tau_i,h} \equiv \mu^c_{i,\tau_i,h}$ for every $\tau_i \in \mathcal{T}_i^c$ and $c \in C$. We now show that (σ, μ') is a sequential equilibrium of G' and satisfies Assumption 1. Since Lemma 2 applies to the case $g_2^* = 0$, in which case $\sigma(\cdot|\tau_2^*)$ is rationalizable for type τ_2^*,

$$\sigma^c(c(h)|h, \tau_1^c) = 1 \qquad \left(\forall c \in C^h, \forall h\right). \tag{5.3}$$

Hence, $\mu'_{2,\tau_2^*,h}$ is consistent and satisfies Assumption 1. The sequential rationality conditions for rational types are (SR1) and (SR2) by construction and (5.3). The sequential rationality and consistency for types in $\mathcal{T}^c$ immediately follows from the construction and the fact that (σ^c, μ^c) is a sequential equilibrium in the auxiliary game. $\square$

The strategic equivalence under sequential equilibrium is somewhat subtle, requiring the lengthy proof above. This is because of the issues relating to the off-the-path beliefs, which play a central role in sequential equilibrium while not being relevant for ICR. If a type τ_1^c, who plans to follow c, deviates from c, then his subsequent behavior may be different from c as ICR cannot restrict the behavior at the contingencies that are precluded by one's own strategy. In that case, off the path beliefs of player 2 at the histories that are not consistent with any type could be different. Moreover, consistency may result in unforeseen restrictions on those beliefs as it is applied for types in $\mathcal{T}^c$ and τ_2^* simultaneously. Assumption 1 ensures that Player 2 assigns zero probability to τ_1^c whenever Player 1 deviates from c, resulting in beliefs that are identical to those with commitment types, as we show in the proof. Of course, at the histories that are consistent with commitment types, the rational types in the games G and G' face the same uncertainty regarding all relevant aspects, such as whether the other player is rational and which $c \in C$ he is playing if he is not rational. This leads to the same set of solutions for rational types in both games.

5.3. **Indistinguishability of Testable Predictions.** The strategic equivalence above implies that the testable predictions with or without commitment types are nearly indistinguishable. Imagine that an empirical or experimental researcher observes outcomes of games that essentially look like a fixed repeated game, as in g^*, but she does not know the players' beliefs about possible commitments or payoff variations. Using the data, she can obtain an empirical distribution on outcome paths. Because of sampling variation, there is some noise regarding the actual equilibrium distribution of the outcomes. The above strategic equivalence implies that the equilibrium distributions for elaborations with or without commitment types can be arbitrarily close, making it impossible to rule out one model without ruling out the the other given the sampling noise.

Towards stating this result formally, let Σ^* be the set of solution concepts that are (1) invariant to the elimination of non-rationalizable plans, (2) invariant to trivial enrichments of the type spaces, and (3) include all solutions generated by the sequential equilibria that satisfy Assumption 1. Given any solution concept $\Sigma \in \Sigma^*$ and any Bayesian game G, a solution σ leads to a probability distribution $\boldsymbol{z}(\cdot|\sigma) \in \Delta(Z)$ on the set Z of outcome paths, such that

$$\boldsymbol{z}(z|\sigma) = \sum_{\tau \in \mathcal{T}} \sum_{s \in S_z} \sigma(s|\tau)\,\pi(\tau) \qquad (\forall z \in Z),$$

where $\mathcal{T}$ is the sets of type profiles in G, $S_z = \{s \in S | z(s) = z\}$ is the set of profiles of action plans that lead to z, π is the (induced) common prior on $\mathcal{T}$, and $\sigma(s|\tau)$ is the probability of action plan s in equilibrium σ when the type profile is τ. A solution concept Σ yields a set

$$\mathcal{Z}(\Sigma, G) = \{\boldsymbol{z}(\cdot|\sigma) \,|\sigma \in \Sigma(G)\}$$

of probability distributions on outcome paths. Towards comparing the distance between such sets, we endow the set $2^{\Delta(Z)}$ of such subsets with the Hausdorff metric d, the standard metric for sets[7]. For any $X, Y \in 2^{\Delta(Z)}$,

$$d(X, Y) \leq \lambda$$

if and only if for each $x \in X$, there exist $y \in Y$ and $p \in \Delta(Z)$ with $x = (1-\lambda)y + \lambda p$, and for each $y \in Y$, there exist $x \in X$ and $p \in \Delta(Z)$ with $x = (1-\lambda)y + \lambda p$.

Our first corollary states that the set of distributions on the outcome paths are nearly identical with or without commitment types.

[7]More specifically, we use the Hausdorff metric induced by the "total variation" metric on $\Delta(Z)$, but since we only use the metric on sets, we will simply define the Hausdorff metric directly.

Corollary 1. *For any $\Sigma \in \Sigma^*$, any ε-elaboration G with commitment types, and any $\varepsilon' \in (\varepsilon, 1)$, there exists an ε'-elaboration G' without commitment types such that $d\left(\mathcal{Z}\left(\Sigma, G\right), \mathcal{Z}\left(\Sigma, G'\right)\right) \leq \left(\varepsilon' - \varepsilon\right) / \left(1 - \varepsilon\right)$.*

Proof. Define $\lambda = \left(\varepsilon' - \varepsilon\right) / \left(1 - \varepsilon\right)$. Consider the ε'-elaboration G' in the previous results. Recall that any type profile $\left(\tau_1, \tau_2^*\right)$ in G has identical solutions to a type profile $\left(f\left(\tau_1\right), \tau_2^*\right)$ in G' where $f\left(\tau_1^*\right) = \tau_1^*$ and $f\left(c\right) = \tau_1^c$. Moreover, $\pi'\left(f\left(\tau_1\right), \tau_2^*\right) = \left(1 - \lambda\right) \pi\left(\tau_1, \tau_2^*\right)$. Hence, $y \in \mathcal{Z}\left(\Sigma, G'\right)$ if and only if there exists $\sigma' \in \Sigma\left(G'\right)$ such that $y = \left(1 - \lambda\right) x + \lambda p$ for x and p where

$$x\left(z\right) = \sum_{\tau \in \mathcal{T}} \sum_{\{s \in S | z(s) = z\}} \sigma'\left(s | f\left(\tau_1\right), \tau_2^*\right) \pi\left(f\left(\tau_1\right), \tau_2^*\right)$$

and $p\left(z\right) = \sum_{\tau_2 \neq \tau_2^*} \sum_{\{s \in S | z(s) = z\}} \sigma'\left(s | \tau\right) \pi'\left(\tau\right)$. Since the set of solutions for $\left(\tau_1, \tau_2^*\right)$ and $\left(f\left(\tau_1\right), \tau_2^*\right)$ are identical, $x \in \mathcal{Z}\left(\Sigma, G\right)$, and the converse is also true in that there exists a $\sigma' \in \Sigma\left(G'\right)$ as above for every $x \in \mathcal{Z}\left(\Sigma, G\right)$. □

Suppose that one wants to restrict G' to be an ε-elaboration, so that the prior probability of rational types are identical. The results in the reputation literature are often continuous with respect to ε when the set and the relative probability of the commitment types are fixed. In that case, such a restriction would not make a difference, as established in the next corollary.

Corollary 2. *Consider any ε-elaboration G with commitment types $\left(C, \pi\right)$ and a solution concept $\Sigma \in \Sigma^*$ such that $\Sigma\left(G^\alpha\right)$ is continuous with respect to α, where G^α is an $\alpha\varepsilon$-elaboration with commitment types $\left(C, \pi/\alpha\right)$ for $\alpha \geq 1$. Then, for any $\lambda > 0$, there exists an ε-elaboration G' without commitment types such that $d\left(\mathcal{Z}\left(\Sigma, G\right), \mathcal{Z}\left(\Sigma, G'\right)\right) \leq \lambda$.*

Proof. Apply the previous result starting from G^α for some $\alpha > 1$ that is sufficiently close to 1, in particular where $\alpha\varepsilon \leq \lambda(1 - \varepsilon) + \varepsilon$, and then apply continuity. □

6. General Case

In this section, we will present the result for the n-player case, allowing commitment types for all players. The definitions for the n-player case mirror the case of $n = 2$, and we will not repeat them here. Since we will allow commitment types for all players, an ε-elaboration with commitment types is now defined as a Bayesian game such that the set of types for each player i is $\{\tau_i^*\} \cup C_i$ where $C_i \subset S_i$ can be empty, type τ_i can play any action plan while a type $c_i \in C_i$ can play only c_i, and the probability $\pi(\tau^*)$ of rational type profile is $1-\varepsilon$. Note that when $\varepsilon < 1$, there some C_i is non-empty. Note also that the distribution of commitment type is not restricted; they can be correlated for example. Such a Bayesian game can be denoted by $(C_1, \ldots, C_n, \pi)$ where $\pi \in \Delta((\{\tau_1\} \cup C_1) \times \cdots \times (\{\tau_n\} \cup C_n))$ is the prior on the type profiles. The result is generalized to this case as follows.

Proposition 4. *For any $\varepsilon, \varepsilon' \in (0,1)$ with $\varepsilon' > \varepsilon$ and for any ε-elaboration G with commitment types $(C_1, \ldots, C_n, \pi)$ there exists a strategically-equivalent ε'-elaboration $G' = (N, A, (\mathcal{G}, \mathcal{T}, \pi'))$ without commitment types in which the commitment types are replaced by types with unique rationalizable action plans:*

(1) *for every $i \in N$, $\pi'\left(g^*, \tau_{-i}^* | \tau_i^*\right) = \pi\left(\tau_{-i}^* | \tau_i^*\right)$;*
(2) *for every $i \in N$ and $c_i \in C_i$, there exists $\tau_i^{c_i} \in \mathcal{T}_i$ such that all ICR action plans of $\tau_i^{c_i}$ are equivalent to c_i, and $\pi'\left(\tau_i^{c_i} | \tau_j^*\right) = \pi\left(c_i | \tau_j^*\right)$ for every $j \neq i$;*
(3) *for every $\Sigma \in \Sigma^*$,*

$$\{\sigma(\tau^*) \,|\, \sigma \in \Sigma(G)\} = \{\sigma(\tau^*) \,|\, \sigma \in \Sigma(G')\}.$$

The first two conditions all together state that each commitment type is replaced by a type that follows the committed action profile as his uniquely rationalizable plan, and the interim beliefs of the rational types remain intact under rationalizability. The last condition states that the two games are strategically equivalent for rational types under any invariant solution concept, including sequential equilibria that puts probability one on rational types off the path. An outline of the proof for this result can be found in the appendix.

7. Necessity of Commitment under CK of Approximate Payoffs

In the previous sections, while we imposed the constraint that it is always common knowledge that the payoffs are the sum of identical stage-game payoffs, we allowed those payoffs to lie anywhere in the interval $[0,1]$. In this section, by contrast, we make the stricter requirement that it is common knowledge that payoffs lie within ε of those in the complete-information game. Under this stricter requirement, we show that commitment types are not dispensable in reputation models. When the stage game is dominance solvable, there is a unique sequential Nash equilibrium outcome, in which the unique rationalizable strategy profile of the stage game is played throughout. Here, ε is uniform over all type spaces and the number of repetitions. For example, in the repeated prisoners' dilemma, one cannot have any cooperation without commitment types when it is common knowledge that the payoffs are approximately those in the prisoner's dilemma.

Define the distance between two stage-game payoff functions via the sup norm:

$$d(g', g) = \max_a |g'(a) - g(a)|$$

Proposition 5. *Fix a complete information stage game g^* which has unique rationalizable profile a^*. Then, there exists $\varepsilon > 0$ such that for any $\varepsilon' > 0$ and any $\bar{t}$, every ε'-elaboration $(N, A, (\mathcal{G}, \mathcal{T}, \pi))$ without commitment types, satisfying the additional requirement that $d(g, g^*) < \varepsilon$ for all $g \in \mathcal{G}$, has a unique sequential equilibrium in which a^* is played by all types at all histories.*

Proof. The elimination process for the finite stage game g^* is finite. Each time an action is eliminated (again by finiteness) it must be that for some δ it is never within δ of being a best reply. Choose ε so that 2ε is smaller than the minimum of these δ.

Now suppose there is a sequential equilibrium strategy profile s^* which contradicts the result. Consider one of the latest histories at which any violation of the profile a^* occurs, and of the violations at this history, consider an action a_i' which is eliminated first in the elimination process for g^*, say at stage k. When player i takes this action, he must believe that (a) the profile a^* is played at all future dates regardless of his action and (b) no action eliminated at stage $k - 1$ or earlier is played at the current history. But then by (b), the fact that a_i' is eliminated at stage k, and the choice of ε, his action is suboptimal in the stage game; and by (a) his action cannot affect future play. This contradicts the concept of sequential equilibrium. □

For example, in a repeated prisoners' dilemma game, if it is common knowledge that payoffs are close to the prisoners' dilemma, then in any sequential equilibrium the players defect throughout the game regardless of the number of repetitions. At some level this is a reflection of general continuity properties of Bayesian Nash equilibrium payoffs with respect to the perturbations of payoffs. Indeed, it is well known that, for any given $\bar{t}$, as $\varepsilon \to 0$, the Bayesian equilibrium payoffs in ε-elaborations of repeated prisoners' dilemma with commitment types approach the payoffs from defection throughout the game. This is in line with the continuity results for Nash equilibrium payoffs with respect to the prior distributions. Hence, for a given $\bar{t}$, our result here differs from the existing continuity results only in terms of the perturbations it considers, making the stage payoffs approach to the original game instead of making the probability of "crazy" types with unrelated payoffs to go to zero. Our result has a major strength however: ε is uniform with respect to the number of repetitions. In contrast, for any ε probability of a tit-for-tat type, cooperation prevails whenever the number of repetitions are sufficiently large, as famously established by Gang of Four.

8. Remarks

Continuity and Robustness of Equivalence. Since interim correlated rationalizability is upperhemicontinuous (Dekel, Fudenberg, and Morris, 2007), each type $\tau_i^{c_i}$ with unique rationalizable action c_i has the same unique rationalizable action on a open neighborhood of parameters and beliefs. Hence, when we fix the interim beliefs of rational types to obtain strategic equivalence, we can perturb the other parameters such as the stage-game payoff functions and beliefs for the newly constructed types, and obtain an open set of ε'-elaborations G' without commitment types that are strategically equivalent to G. In particular, type $\tau_i^{c_i}$ need not be exactly indifferent between his actions conditional on meeting a rational type; this was only a simplifying aspect in our construction.

On the other hand, our result is silent about continuity with respect to variation of the beliefs of the rational types. Such a continuity is directly tied to the continuity properties of the solution concept in the original game G, by our strategic equivalence result. Of course, since ε' must be larger than ε (albeit being arbitrarily close), our result and the reputation result that it is applied to are relevant only when the solution concept on the original model G is continuous with respect to the small variations in ε when the commitment types and their relative probabilities with respect to each other are fixed. This is indeed the case for most existing models.[8] Finally, since we also include many more types in our construction, the

[8]For example, sequential equilibrium would be continuous with respect to such scaling of the probabilities of commitment types. Bayesian Nash equilibrium behavior of the rational types would be continuous with respect to all variations of priors (with possibly varying commitment types and relative probabilities) because such variations could be represented as an ex-ante payoff perturbation.

new game G' allows many more perturbations of beliefs than those that are allowed in G. For example, the rational types can now put positive probability on newly constructed types that are meant to encode the beliefs of type $\tau_i^{c_i}$. Continuity with respect to such perturbations is related to the continuity properties of the solution concept in the original game with respect to the introduction of new commitment types. Here, while putting positive probability on types with unique rationalizable actions is equivalent to adding a new commitment type under our invariance conditions and strategic equivalence result, putting positive probability on some newly constricted types with multiple rationalizable actions could have possibly different implications under some solution concepts.

Commitment to Mixed Strategies. In some reputation models, the commitment types are allowed to play a mixed action plan. For the natural case that only the realized moves are observable, such mixed commitment types are incorporated in our paper as follows. A mixed commitment type σ_i induces a probability distribution μ^{σ_i} on pure action plans of the player in reduced form. From the point of view of the rational type τ_j^* of the other player, the commitment type σ_i can be replaced by pure commitment types in the support of μ^{σ_i}, by putting probability $\pi(\sigma_i)\,\mu^{\sigma_i}(s_i)$ on each s_i in the support of μ^{σ_i}, where $\pi(\sigma_i)$ is the probability of σ_i in the original elaboration G and $\mu^{\sigma_i}(s_i)$ is the probability of s_i under μ^{σ_i}. Application of Proposition 4 to the resulting elaboration with pure commitment types yields an elaboration G' without commitment types that is strategically equivalent for the rational types.

Infinitely Repeated Games. Here, we focus on finitely repeated games. It is actually easier to construct types that are committed to a particular plan of action up to an arbitrary finite horizon as the unique rationalizable plan in infinite-horizon games if one does not insist on the common-prior assumption. Indeed we provided such a result in Weinstein and Yildiz (2013) in another context as we discussed before. It also seems feasible to extend our construction within common-prior assumption to infinitely-repeated games, using finite-horizon truncations. Hence, it seems feasible to obtain a similar result for infinite-horizon games allowing only arbitrarily long but finite-horizon commitments. We do not pursue such results here mainly because the most major results in infinite-horizon reputation literature, such as the results of Fudenberg and Levine (1989), are based on types that commit to a playing a fixed move, and such types can easily be justified within the repeated-game framework.[9]

Other games. The applications in reputation formation are not confined to the repeated-games framework. Indeed, an important strand of literature explores the role of reputation in bargaining considering types that commit to dynamic plans (see for example Abreu and Gul (2000), Abreu and Pearce (2007), and Wolitzky (2012)). Of course, understanding the scope of reputation within the structural assumptions of those models is also very important. Here, as a first step, we established a strategic equivalence result for finitely repeated games.

How much variation in stage-game payoffs and type spaces do we need to support a commitment? Our main result establishes that arbitrary commitment types can result from ICR without any restriction on the stage-game payoffs and the type spaces. Moreover, our construction uses only a couple of simple stage-game payoff functions.[10] Feasibility of such stage-game payoff functions is sufficient for supporting arbitrary commitment types by introducing uncertainty on stage game payoffs. On the other hand, Proposition 5 shows that in the limit as the maximal variation in stage-game payoffs is taken to 0, commitment types cannot be generated.

Our construction (in the first part of the proof of Lemma 1) makes fundamental use of players who do not know their own payoffs. Some of the literature has focused on models with common knowledge that each player knows his own payoffs; Fudenberg, Kreps, and Levine (1988) call this a model with "personal types"˙ We do not know precisely to what extent our results can be recovered in a model with personal types. There are multiple difficulties. The first is the construction of a type with unique rationalizable plan, as in Lemma 1. This is considerably more difficult when using personal types, and while it is possible to generate commitment types for some non-trivial plans, we do not know if it is possible for all plans. The second difficulty arises when putting the types into a common-prior type space. The technique we used for Lemma 2 relied on the commitment types of Player 1 believing that their payoff is always identically zero when Player 2 is a normal type. With commitment

[9]When the commitment type plays a mixed move, the resulting pure action plans involve commitment to dynamically varying action plans. An extension of our results could be useful in that case.

[10]More precisely, it uses the family $g_i^{a_i^*,a_{-i}^*,\lambda} = \lambda 1_{a_i^*} + (1-\lambda)\, 1_{a_{-i}^*}$, $a_i^* \in A_i$, $a_{-i}^* \in A_{-i}^*$, $\lambda \in [0,1]$, where $1_{a_j^*}$ is the characteristic function of a_j^*, taking the value of 1 when a_j^* is played and 0 otherwise.

types, this technique cannot be used, as payoffs cannot be correlated with the opponent's type. It is an open question whether some other technique would be successful. Note that this second difficulty only arises if we assume personal types *and* a common prior.

9. Conclusion

The reputation literature, one of the main accomplishments of game theory, relies on the existence of commitment types. It is important for the interpretation of the results in this literature whether one can obtain the same results within a rationalistic framework in which all types can follow the plans that are available to rational types and all types' payoffs satisfy the structural payoff assumptions of the underlying model. If one can obtain the same results within such a framework, we can interpret the results as coming from incomplete information about payoffs. Otherwise, the result must be interpreted as stemming from the factors that are outside of the model, such as irrationality, psychological anomalies, and super-game concerns. In this paper, within the context of finitely repeated games, we have established that one can obtain all results within a rationalistic framework, allowing an interpretation based on incomplete information. This is the case when all stage-game payoffs are allowed. On the other hand, for games with dominance-solvable stage games, we show that reputation cannot have an impact when the stage game payoffs are sufficiently restricted. Hence, the scope of reputation within a rationalistic framework depends on the severity of the additional structural assumptions imposed when there are such assumptions.

Appendix A. Omitted Proofs

A.1. **Preliminary Definitions.** In the appendix, we will also consider type spaces without a common prior. Such a type space is a list $(\mathcal{G}, \mathcal{T}, \pi(\cdot|\cdot))$ where $\pi(\cdot|\tau_i) \in \Delta(\mathcal{G} \times \mathcal{T}_{-i})$ is the probability distribution of τ_i. Here, there need not be a single $\pi \in \Delta(\mathcal{G} \times \mathcal{T})$ that leads to these interim beliefs by Bayes' rule. Fix any $G = (N, A, (\mathcal{G}, \mathcal{T}, \pi(\cdot|\cdot)))$. For each $i \in N$ and for each belief $\beta \in \Delta(\mathcal{G} \times S_{-i})$, we write $BR_i(\beta)$ for the set of actions $s_i \in S_i$ that maximize the expected value of $u_i(z(s_i, s_{-i})|g)$ under the probability distribution β.

Interim correlated rationalizability (ICR) is computed by the following elimination procedure: For each i and τ_i, set $S_i^0[\tau_i|G] = S_i$, and define sets $S_i^k[\tau_i|G]$ for $k > 0$ iteratively, by letting $s_i \in S_i^k[\tau_i|G]$ if and only if $s_i \in BR_i\left(\text{marg}_{\mathcal{G} \times S_{-i}}\beta\right)$ for some $\beta \in \Delta(\mathcal{G} \times \mathcal{T}_{-i} \times S_{-i})$ such that

$\text{marg}_{\mathcal{G}\times\mathcal{T}_{-i}}\beta = \pi(\cdot|\tau_i)$ and $\beta\left(s_{-i} \in S_{-i}^{k-1}[\tau_{-i}|G]\right) = 1$. That is, s_i is a best response to a belief of τ_i that puts positive probability only to the actions that survive the elimination in round $k-1$. We write $S^k[\tau|G] = S_1^k[\tau_1|G] \times S_2^k[\tau_2|G]$. Then,

$$S_i^\infty[\tau_i|G] = \bigcap_{k=0}^{\infty} S_i^k[\tau_i|G].$$

The following class of action plans will play an important role in our construction:

Definition 6. A plan s_i is said to be *sure-thing compliant* if and only if there is no partial history h and move $a_i \in A_i$ such that $s_i(h, (s_i(h), a_{-i})) = a_i$ for every a_{-i} but $s_i(h) \neq a_i$.

In other words, a plan is sure-thing compliant if whenever the player plays a_i in all possible continuations next period, he also plays a_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.

Finally, we introduce the following notation regarding the sets with equivalent action plans. Given any two sets X, Y of action plans, we write $X \simeq Y$ if for every $x \in X$ there exists $y \in Y$ that is equivalent to x, and for every $y \in Y$ there exists $x \in X$ that is equivalent to y.

A.2. **Proof of Lemma 1.** Our proof has three main steps. First, we will prove it for sure-thing compliant action plans, without requiring the type space to have common prior or the full-support property (property 1 in the statement of the lemma). We then extend this result to all action plans, without requiring the properties on type space once again. Finally, we convert the latter type space to a type space with common prior and full support assumptions without altering the rationalizable actions, proving the lemma. The first step is the following lemma; Weinstein and Yildiz (2013) proved this lemma for infinite-horizon games, and the proof carries over to finitely repeated games with minor modifications. (The proof can be found in the online appendix.)

Lemma 3 (Weinstein-Yildiz 2013)**.** *For any sure-thing compliant action plan s_i, there exists a game $\tilde{G} = \left(N, A, \left(\tilde{\mathcal{G}}, \tilde{\mathcal{T}}, \tilde{\pi}(\cdot|\cdot)\right)\right)$ with a type $\tau_i^{s_i}$ such that $S_i^\infty\left[\tau_i^{s_i}|\tilde{G}\right] \simeq \{s_i\}$. (The type space does not necessarily have a common prior.)*

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

Lemma 4. *For any action plan s_i, there exists a game $\tilde{G} = \left(N, A, \left(\tilde{\mathcal{G}}, \tilde{\mathcal{T}}, \tilde{\pi}(\cdot|\cdot)\right)\right)$ with a type $\tau_i^{s_i}$ such that $S_i^\infty\left[\tau_i^{s_i}|\tilde{G}\right] \simeq \{s_i\}$. (The type space does not necessarily have a common prior.)*

Proof. Fix some $a^*_{-i} \in A_{-i}$, and define a function $v_{-i} : A_{-i} \to [0,1]$ by

$$v_{-i}(a_{-i}) = \begin{cases} 1 & \text{if } a_{-i} = a^*_{-i}, \\ 0 & \text{otherwise.} \end{cases}$$

For every $\hat{a}_i \in A_i$, define also a function $v_i^{\hat{a}_i} : A_i \to [0,1]$ by

$$v_i^{\hat{a}_i}(a_i) = \begin{cases} 1 & \text{if } a_i = \hat{a}_i, \\ 0 & \text{otherwise.} \end{cases}$$

Finally, consider the class of stage-game payoff functions $g_i^{\hat{a}_i} : A \to [0,1]$ for player i where

$$g_i^{\hat{a}_i}(a_i, a_{-i}) = \lambda v_i^{\hat{a}_i}(a_i) + (1-\lambda) v_{-i}(a_{-i})$$

for some $\lambda \in (0, 1/(2\bar{t}+1))$. Let also $\hat{S}_{-i}$ be the set of action profiles s_{-i} for which there exists a function $\rho : H \times A_{-i} \to A_i$ such that

(1) for any $t < \bar{t}$, any history h^t and any $(a_i, a_{-i}) \in A_i$,
$$s_{-i}\left(h^t, (a_i, a_{-i})\right) = a^*_{-i} \iff a_i = \rho\left(h^t, a_{-i}\right);$$
(2) $\rho\left(h^{\bar{t}-1}, a_{-i}\right) \equiv s_i\left(h^{\bar{t}-1}\right)$ for all those $h^{\bar{t}-1}$ such that player i has played according to s_i throughout.

At any history $(h, (a_i, a_{-i}))$, player $-i$ rewards i by playing a^*_{-i} if $a_i = \rho(h, a_{-i})$. The only restriction on which move is rewarded occurs at date $\bar{t}-1$ and in the contingency that i has followed s_i up to $\bar{t}-1$: he will be rewarded at $\bar{t}$ if he continues to follow s_i at $\bar{t}-1$. The set $\hat{S}_{-i}$ is symmetric in all other ways. Note also that player $j \neq i$ reacts differently to the rewarded move of player i. Hence, the actions in $\hat{S}_{-i}$ are all sure-thing compliant. Thus, by Lemma 3, for each $s_{-i} \in \hat{S}_{-i}$, there exists a game $G^{s_{-i}} = (N, A, (\mathcal{G}^{s_{-i}}, \mathcal{T}^{s_{-i}}, \pi^{s_{-i}}(\cdot|\cdot)))$ with a type $\tau_{-i}^{s_{-i}}$ such that $S^\infty_{-i}\left[\tau_{-i}^{s_{-i}} | G^{s_{-i}}\right] \simeq \{s_{-i}\}$.

Define the game $\tilde{G} = \left(N, A, \left(\tilde{\mathcal{G}}, \tilde{\mathcal{T}}, \tilde{\pi}(\cdot|\cdot)\right)\right)$ by

$$\begin{aligned}
\tilde{\mathcal{G}} &= \left\{g_i^{\hat{a}_i} | \hat{a}_i \in A_i\right\} \cup \bigcup_{s_{-i} \in \hat{S}_{-i}} \mathcal{G}^{s_{-i}}; \\
\tilde{\mathcal{T}}_i &= \{\tau_i^{s_i}\} \cup \bigcup_{s_{-i} \in \hat{S}_{-i}} \mathcal{T}_i^{s_{-i}}; \quad \tilde{\mathcal{T}}_{-i} = \bigcup_{s_{-i} \in \hat{S}_{-i}} \mathcal{T}_{-i}^{s_{-i}}; \\
\tilde{\pi}(\cdot|\tau_j) &= \pi^{s_{-i}}(\cdot|\tau_j) \qquad \left(\forall \tau_j \in \mathcal{T}_j^{s_{-i}}, j \in N, s_{-i} \in \hat{S}_{-i}\right); \\
\tilde{\pi}\left(g^{s_i\left(z(s_i, s_{-i})^{\bar{t}}\right)}, \tau_{-i}^{s_{-i}} | \tau_i^{s_i}\right) &= 1/\left|\hat{S}_{-i}\right| \qquad \left(\forall s_{-i} \in \hat{S}_{-i}\right).
\end{aligned} \tag{A.1}$$

The main attraction of the definition is the belief of the newly introduced type $\tau_i^{s_i}$ in (A.1). He assigns equal probabilities on types $\tau_{-i}^{s_{-i}}$ but also believes that there is a perfect correlation between the types $\tau_{-i}^{s_{-i}}$ and the way his own action affects his payoff. If he follows s_i throughout $h^{\bar{t}}$ and observes the moves of the other player, he learns what action $\hat{a}_i$ is best for him, which happens to be the action that he would have played at that history according to s_i. Note that each of the types other than $\tau_i^{s_i}$ has a unique rationalizable action in reduced form. Hence, $\tau_i^{s_i}$ has a unique rationalizable belief regarding the payoff functions and the outcomes.

In the rest of the proof, we will show that s_i is uniquely rationalizable for $\tau_i^{s_i}$ in reduced form, i.e., $S_i^\infty\left[\tau_i^{s_i}|\tilde{G}\right] \simeq \{s_i\}$. We now introduce couple of notations. For any history h^t, write $P_t^*\left(h^t\right)$ for the probability that a_{-i}^* is played at date t conditional on h according to the rationalizable belief of $\tau_i^{s_i}$. As noted above, by symmetry,

$$P_t^*\left(h^t\right) = \begin{cases} 1 & \text{if } t = \bar{t} \text{ and } i \text{ follows } s_i \text{ throughout } h^t; \\ 0 & \text{if } t = \bar{t} \text{ and } i \text{ follows } s_i \text{ up to } \bar{t}-1 \text{ but deviates at } \bar{t}-1 \text{ in } h^t; \\ 1/\left|A_{-i}\right| & \text{otherwise.} \end{cases} \tag{A.2}$$

Write $U_i\left(s_i'|h\right)$ for the expected payoff of i from playing s_i' under the rationalizable belief of type $\tau_i^{s_i}$ conditional on history h. Write also $U_i\left(h^t\right)$ for the realized expected payoff of $\tau_i^{s_i}$ up to date t at history h^t.

We now show that $U_i\left(s_i|h^t\right) > U_i\left(s_i'|h^t\right)$ for every history h^t and action plan s_i' such that i follows s_i throughout h^t and $s_i'\left(h^t\right) \neq s_i\left(h^t\right)$. So long as he follows s_i, every such history h^t is reached with positive probability under the rationalizable belief of $\tau_i^{s_i}$. This therefore shows that the expected payoff from s_i is strictly higher than any s_i' that is not equivalent to s_i. Therefore, $S_i^\infty\left[\tau_i^{s_i}|\tilde{G}\right] \simeq \{s_i\}$.

First consider the case $t = \bar{t}$. Conditional on h^t, $\tau_i^{s_i}$ assigns probability 1 on $g^{s_i\left(h^t\right)}$. If he follows s_i and plays $s_i\left(h^t\right)$ at h^t, then his own action contributes λ to his payoff, while his own action contributes zero to his payoff otherwise. Moreover, since he has followed s_i throughout h^t, he will be rewarded for sure by the other player at $\bar{t}$, contributing $1-\lambda$ to his payoff regardless of his own move at h^t. Hence,

$$U_i\left(s_i|h^t\right) = U_i\left(h^t\right) + 1,$$

and

$$U_i\left(s_i'|h^t\right) = U_i\left(h^t\right) + 1 - \lambda,$$

yielding

$$U_i\left(s_i|h^t\right) - U_i\left(s_i'|h^t\right) = \lambda > 0.$$

Now consider the case $t < \bar{t}$. His payoff from following s_i is

$$\begin{aligned} U_i\left(s_i|h^t\right) &= U_i\left(h^t\right) + \lambda \sum_{t'=t}^{\bar{t}} E\left[v_i^{\hat{a}_i}\left(s_i\left(h^{t'}\right)\right)|h^t, s_i\right] + (1-\lambda) \sum_{t'=t}^{\bar{t}} E\left[P_{t'}^*\left(h^{t'}\right)|h^t, s_i\right] \\ &\geq U_i\left(h^t\right) + (1-\lambda)\frac{\bar{t}-t}{|A_{-i}|} + 1. \end{aligned}$$

To see the lower bound, note that, so long as he follows s_i, he gets 1 at date $\bar{t}$ (as in the previous case) and at least $(1-\lambda)/|A_{-i}|$ at each $t' < \bar{t}$. (At any $t' < \bar{t}$, $v_i^{\hat{a}_i} \geq 0$ and $P_{t'}^* \geq 1/|A_{-i}|$ when he follows s_i.) On the other hand, his payoff from following s_i' is

$$\begin{aligned} U_i\left(s_i'|h^t\right) &= U_i\left(h^t\right) + \lambda \sum_{t'=t}^{\bar{t}} E\left[v_i^{\hat{a}_i}\left(s_i'\left(h^{t'}\right)\right)|h^t, s_i'\right] + (1-\lambda) \sum_{t'=t}^{\bar{t}} E\left[P_{t'}^*\left(h^{t'}\right)|h^t, s_i'\right] \\ &\leq U_i\left(h^t\right) + \lambda\left(\bar{t}-t+1\right) + (1-\lambda)\frac{\bar{t}-t+1}{|A_{-i}|}. \end{aligned}$$

The upper bound comes from the fact that $v_i^{\hat{a}_i} \leq 1$ throughout and $P_{t'}^* \leq 1/|A_{-i}|$ after a deviation from s_i (by A.2). Combining the two inequalities, we obtain

$$U_i\left(s_i|h^t\right) - U_i\left(s_i'|h^t\right) \geq (1-\lambda)\left(1 - \frac{1}{|A_{-i}|}\right) - \lambda\left(\bar{t}-t\right) > 0,$$

where the strict inequality follows from $\lambda < 1/\left(2\bar{t}+1\right)$ and $|A_{-i}| \geq 2$. □

Proof of Lemma 1. By Lemma 4, there exists a game $\tilde{G} = \left(N, A, \left(\tilde{\mathcal{G}}, \tilde{\mathcal{T}}, \tilde{\pi}\left(\cdot|\cdot\right)\right)\right)$ with a type $\tilde{\tau}_i$ such that $S_i^\infty\left[\tilde{\tau}_i|\tilde{G}\right] \simeq \{s_i\}$. This falls short of the conditions of Lemma 1 in that $\left(\tilde{\mathcal{G}}, \tilde{\mathcal{T}}, \tilde{\pi}\left(\cdot|\cdot\right)\right)$ does not necessarily admit a common prior and the prior could not have a full support (Condition 1) even if existed. Here, we remedy this problem by converting $\tilde{G}$ to a common prior game $G^s = (N, A, (\mathcal{G}^{s_i}, \mathcal{T}^{s_i}, \pi^{s_i}))$ with the desired properties. First, for every $\lambda \in (0,1)$, define $G^\lambda = \left(N, A, \left(\tilde{\mathcal{G}}, \tilde{\mathcal{T}}, \pi^\lambda\left(\cdot|\cdot\right)\right)\right)$ by setting

$$\pi^\lambda\left(g, \tau_{-j}|\tau_j\right) = \frac{\lambda}{\left|\tilde{\mathcal{G}} \times \tilde{\mathcal{T}}_{-j}\right|} + (1-\lambda)\tilde{\pi}\left(g, \tau_{-j}|\tau_j\right)$$

at each $(g, \tau_j, \tau_{-j}) \in \tilde{\mathcal{G}} \times \tilde{\mathcal{T}}$. Now, as $\lambda \to 0$, $\pi^\lambda\left(g, \tau_{-j}|\tau_j\right) \to \tilde{\pi}\left(g, \tau_{-j}|\tau_j\right)$ everywhere. Together with a continuity result for belief hierarchies by Mertens and Zamir (1985), this implies that the belief hierarchy of type $\tilde{\tau}_i$ in game G^λ converges to the belief hierarchy of $\tilde{\tau}_i$ in game $\tilde{G}$. Thus, by upperhemicontinuity of ICR (Dekel, Fudenberg, and Morris, 2006), there exists $\bar{\lambda} > 0$ such that

$$S_i^\infty\left[\tilde{\tau}_i|G^{\bar{\lambda}}\right] \subseteq S_i^\infty\left[\tilde{\tau}_i|\tilde{G}\right] \simeq \{s_i\}.$$

Since $S_i^{\infty}\left[\tilde{\tau}_i|G^{\bar{\lambda}}\right]$ is non-empty, this implies that

$$S_i^{\infty}\left[\tilde{\tau}_i|G^{\bar{\lambda}}\right] \simeq \{s_i\}\,. \tag{A.3}$$

Moreover, since $\tilde{\mathcal{G}} \times \tilde{\mathcal{T}} \times S$ is finite, there exists some finite k such that

$$S_i^{\infty}\left[\tilde{\tau}_i|G^{\bar{\lambda}}\right] = S_i^{k}\left[\tilde{\tau}_i|G^{\bar{\lambda}}\right]. \tag{A.4}$$

Now, since $\pi^{\bar{\lambda}}(g, \tau_{-j}|\tau_j) > 0$ everywhere, by the main result of Lipman (2003), there exists a common-prior game $G^{s_i} = (N, A, (\mathcal{G}^{s_i}, \mathcal{T}^{s_i}, \pi^{s_i}))$ such that the common prior π^{s_i} is positive everywhere and there exists a type $\tau_i^{s_i} \in \mathcal{T}_i^{s_i}$ whose first k orders of beliefs are identical to that of type $\tilde{\tau}_i$ in game $G^{\bar{\lambda}}$. But by Dekel, Fudenberg, and Morris (2007), S^k is a function of the first k orders of beliefs, yielding

$$S_i^{k}\left[\tau_i^{s}|G^{s}\right] = S_i^{k}\left[\tilde{\tau}_i|G^{\bar{\lambda}}\right]. \tag{A.5}$$

Combining (A.3), (A.4) and (A.5), we obtain

$$S_i^{\infty}\left[\tau_i^{s}|G^{s}\right] \subseteq S_i^{k}\left[\tau_i^{s}|G^{s}\right] = S_i^{k}\left[\tilde{\tau}_i|G^{\bar{\lambda}}\right] \simeq \{s_i\}\,.$$

Since $S_i^{\infty}\left[\tau_i^{s}|G^{s}\right] \neq \varnothing$, this further implies that

$$S_i^{\infty}\left[\tau_i^{s}|G^{s}\right] \simeq \{s_i\}\,,$$

as desired. □

A.3. **Proof of Lemma 2.** Using induction on k, we will show that $S_i^k[\tau_i|G'] = S_i^k[\tau_i|G^c]$ for every k, $\tau_i \in \mathcal{T}_i^c$, and $i \in N$. This is true for $k = 0$ by definition. Towards an induction, assume that

$$S_{-i}^{k-1}\left[\tau_{-i}|G'\right] = S_{-i}^{k-1}\left[\tau_{-i}|G^c\right] \qquad \left(\forall \tau_{-i} \in T_{-i}^c\right). \tag{A.6}$$

Take any $\tau_i \in \mathcal{T}_i^c$ and write $B(\tau_i|G)$ for the set of all beliefs β of type τ_i after round $k-1$ in game G for any $G \in \{G', G^c\}$, where $\mathrm{marg}_{\mathcal{G}\times\mathcal{T}_{-i}}\beta = \pi(\cdot|\tau_i)$ and $\beta\left(s_{-i} \in S_{-i}^{k-1}\left[\tau_{-i}|G\right]\right) = 1$. First consider the case $\tau_i \neq \tau_1^c$. In that case, by definition, $\pi'(\cdot|\tau_i) = \pi^c(\cdot|\tau_i)$. Together with the inductive hypothesis (A.6), this implies that $B(\tau_i|G') = B(\tau_i|G^c)$. Therefore, $s_i \in S_i^k[\tau_i|G']$ if and only if $s_i \in BR_i\left(\mathrm{marg}_{\mathcal{G}'\times S_{-i}}\beta\right)$ for some $\beta \in B(\tau_i|G') = B(\tau_i|G^c)$, and this is the case if and only if $s_i \in S_i^k[\tau_i|G^c]$, showing that $S_i^k[\tau_i|G'] = S_i^k[\tau_i|G^c]$.

Now consider the case $\tau_i = \tau_1^c$. Then,

$$\pi'(\cdot|\tau_i) = p^c\delta_{\left((\mathbf{0},g_2^*),\tau_2^*\right)} + (1-p^c)\,\pi^c(\cdot|\tau_i) \tag{A.7}$$

where the probability $p^c \in (0,1)$ is defined in (4.1), and δ_x is the Dirac's measure on x, putting probability 1 on $\{x\}$. Hence, by the inductive hypothesis (A.6), $\beta \in B(\tau_i|G')$ if and only if

$$\beta = p^c \beta(\cdot|\tau_2^*) + (1-p^c)\,\beta(\cdot|\mathcal{T}_2^c) \tag{A.8}$$

for some conditional beliefs $\beta(\cdot|\tau_2^*) \in \Delta\left(\{((\mathbf{0}, g_2^*), \tau_2^*)\} \times S_2^{k-1}[\tau_2^*|G']\right)$ and

$$\beta(\cdot|\mathcal{T}_2^c) \in B(\tau_i|G^c).$$

Now, take any $s_i \in S_i^k[\tau_i|G']$. Then, $s_i \in BR_i\left(\mathrm{marg}_{\mathcal{G}'\times S_{-i}}\beta\right)$ for some $\beta \in B(\tau_i|G')$. By (A.8), for any s_i',

$$\begin{aligned} p^c \cdot 0 + (1-p^c)\,E_{\beta(\cdot|\mathcal{T}_2^c)}[u_i(s_i, s_{-i}|g)] &= E_\beta[u_i(s_i, s_{-i}|g)] \\ &\geq E_\beta\left[u_i\left(s_i', s_{-i}|g\right)\right] \\ &= p^c \cdot 0 + (1-p^c)\,E_{\beta(\cdot|\mathcal{T}_2^c)}\left[u_i\left(s_i', s_{-i}|g\right)\right] \end{aligned}$$

where $\beta(\cdot|\mathcal{T}_2^c) \in B(\tau_i|G^c)$. (Here, the inequality follows from s_i being a best response, and the equalities follow from (A.8).) Since $p^c < 1$, this further implies that

$$E_{\beta(\cdot|\mathcal{T}_2^c)}[u_i(s_i, s_{-i}|g)] \geq E_{\beta(\cdot|\mathcal{T}_2^c)}\left[u_i\left(s_i', s_{-i}|g\right)\right],$$

showing that $s_i \in BR_i\left(\mathrm{marg}_{\mathcal{G}'\times S_{-i}}\beta(\cdot|\mathcal{T}_2^c)\right)$. Therefore, $s_i \in S_i^k[\tau_i|G^c]$.

Conversely, take any $s_i \in S_i^k[\tau_i|G^c]$. By definition, $s_i \in BR_i\left(\mathrm{marg}_{\mathcal{G}'\times S_{-i}}\beta(\cdot|\mathcal{T}_2^c)\right)$ for some $\beta(\cdot|\mathcal{T}_2^c) \in B(\tau_i|G^c)$. Pick any $\beta(\cdot|\tau_2^*) \in \Delta\left(\{((\mathbf{0}, g_2^*), \tau_2^*)\} \times S_2^{k-1}[\tau_2^*|G']\right)$, and define $\beta \in B(\tau_i|G^c)$ by (A.8). Now, for any s_i',

$$\begin{aligned} E_\beta[u_i(s_i, s_{-i}|g)] &= p^c \cdot 0 + (1-p^c)\,E_{\beta(\cdot|\mathcal{T}_2^c)}[u_i(s_i, s_{-i}|g)] \\ &\geq p^c \cdot 0 + (1-p^c)\,E_{\beta(\cdot|\mathcal{T}_2^c)}\left[u_i\left(s_i', s_{-i}|g\right)\right] \\ &= E_\beta\left[u_i\left(s_i', s_{-i}|g\right)\right], \end{aligned}$$

where the inequality follows from s_i being a best response, and the equalities follow from (A.8). That is, $s_i \in BR_i\left(\mathrm{marg}_{\mathcal{G}'\times S_{-i}}\beta\right)$. Therefore, $s_i \in S_i^k[\tau_i|G']$.

A.4. **Proof of Proposition 4.** Here we outline the proof of the general proposition. Note first that Lemma 1 applies to general case as well: for each $i \in N$ and $c_i \in C_i$ there exists a exists a Bayesian repeated game $G^{c_i} = (N, A, (\mathcal{G}^{c_i}, \mathcal{T}^{c_i}, \pi^{c_i}))$ with π^{c_i} positive everywhere and with a type

$\tau_i^{c_i} \in \mathcal{T}_i^{c_i}$ for which all ICR actions are equivalent to c_i. Again all those types can be taken unique and distinct from each other. Write $\mathcal{T}_i = \{\tau_i^*\} \cup C_i$, and define mapping $\tilde{\tau}_i$ on $\mathcal{T}_i$ by

$$\tilde{\tau}_i(\tau_i) = \begin{cases} \tau_i^* & \text{if } \tau_i = \tau_i^* \\ \tau_i^{c_i} & \text{otherwise,} \end{cases}$$

and mapping γ_i from $\tilde{\mathcal{T}}_i = \tilde{\tau}_i(\mathcal{T}_i)$ to stage-game payoff functions by

$$\gamma_i(\tau_i) = \begin{cases} g_i^* & \text{if } \tau_i = \tau_i^* \\ \mathbf{0} & \text{otherwise.} \end{cases}$$

Write also $\gamma(\tau) = (\gamma_1(\tau_1), \ldots, \gamma_n(\tau_n))$ for $\tau \in \tilde{\mathcal{T}} = \tilde{\mathcal{T}}_1 \times \cdots \times \tilde{\mathcal{T}}_n$. We construct $G' = (N, A, (\mathcal{G}', \mathcal{T}', \pi'))$ by setting

$$\begin{aligned} \mathcal{G}' &= \prod_{i \in N} \{g_i^*, \mathbf{0}\} \cup \bigcup_{i \in N, c_i \in C_i} \mathcal{G}^{c_i} \\ \mathcal{T}_j' &= \left\{\tau_j^*\right\} \cup \bigcup_{i \in N, c_i \in C_i} \mathcal{T}_j^{c_i} \qquad (\forall j \in N) \\ \pi'(g, \tau) &= \begin{cases} 1 - \varepsilon' & \text{if } (g, \tau) = (g^*, \tau^*), \\ \frac{1-\varepsilon'}{1-\varepsilon} \pi(\tau') & \text{if } \tau = \tilde{\tau}(\tau') \text{ and } g = \gamma(\tau) \text{ for some } \tau' \in \mathcal{T} \setminus \{\tau^*\}, \\ \frac{\varepsilon' - \varepsilon}{(1-\varepsilon)(|C_1| + \cdots + |C_n|)} \pi^{c_i}(g, \tau) & \text{if } (g, \tau) \in \mathcal{G}^{c_i} \times \mathcal{T}^{c_i} \text{ for some } c_i \in C_i \\ 0 & \text{otherwise.} \end{cases} \end{aligned}$$

Observe that G' satisfies the properties in the proposition. For any rational type τ_j^*,

$$\pi'\left(\tau_j^*\right) = \frac{1-\varepsilon'}{1-\varepsilon} \pi\left(\tau_j^*\right),$$

and hence $\pi'\left(\tilde{\tau}_{-j}(\tau_{-i}) | \tau_j^*\right) = \pi\left(\tau_{-i} | \tau_j^*\right)$ for every $\tau_{-j} \in \mathcal{T}_{-j}$. On the other hand, for type $\tau_i^{c_i}$, his payoffs vary only when the other types are in $\mathcal{T}_{-i}^{c_i}$. Hence, as in Lemma 2, $S_i^\infty[\tau_i^{c_i} | G'] = c_i$.

A.5. **Proof of Lemma 3.** We first introduce a more general notion of equivalence. Recall that $z(s)$ denotes the outcome of a profile s of action plans. In line with our notation for histories, we will write $z(s)^t$ for the truncation of $z(s)$ at the beginning date t; i.e., if $z(s) = \left(a^0, a^1, \ldots, a^{\bar{t}}\right)$, then $z(s)^t = \left(a^0, a^1, \ldots, a^{t-1}\right)$. Recall also that action plans s_i and s_i' are *equivalent* if $z(s_i, s_{-i}) = z(s_i', s_{-i})$ for all action plans $s_{-i} \in S_{-i}$, i.e., they lead to the same outcome no matter what strategy the other player plays. Note that s_i and s_i' are equivalent iff $s_i\left(h^t\right) = s_i'\left(h^t\right)$ for every history h^t in which i played according to s_i throughout; they may differ only in their prescriptions for histories that they preclude. Hence, in reduced form, action plans can be represented as mappings that maps the history of other players' play into own stage game actions. Similarly, action plans s_i and

s_i' are said to be *t-equivalent* if $z(s_i, s_{-i})^t = z(s_i', s_{-i})^t$ for all action plans $s_{-i} \in S_{-i}$, they lead to the same history up to date t no matter what strategy the other player plays. Because we have a finite horizon $\bar{t}$, equivalence is the same as $\bar{t}+1$-equivalence. Given any two sets X, Y of action plans, we write $X \simeq^t Y$ if for every $x \in X$ there exists $y \in Y$ that is equivalent to x, and for every $y \in Y$ there exists $x \in X$ that is t-equivalent to y. We prove the following more general version of Lemma 3 for t-equivalence. Note that the construction in this proof relies on the fact that players do not know their own stage-game payoffs and do not observe them at each stage, but can learn them from other players' actions.

Lemma 5 (Weinstein-Yildiz 2013)**.** *For any sure-thing compliant action plan s_i and any $t \in T$, there exists a game $\tilde{G} = \left(N, A, \left(\tilde{\mathcal{G}}, \tilde{\mathcal{T}}, \tilde{\pi}(\cdot|\cdot)\right)\right)$ with a type $\tau_i^{s_i,t}$ such that $S_i^\infty\left[\tau_i^{s_i,t}|\tilde{G}\right] \simeq^t \{s_i\}$. (The type space does not necessarily have a common prior.)*

Proof. We will induct on t. When $t = 1$, it suffices to consider a type $\tau^{s_i,t}$ who is certain that in the stage game, $s_i(\varnothing)$ yields payoff 1 while all other actions yield payoff 0. Now fix t, s_i and assume the result is true for all players and for $t-1$. 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 s_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 $t-1$ in which player i always follows the plan s_i, so that $|\hat{H}| = |A_{-i}|^{t-1}$, where A_{-i} is the set of profiles of static moves for the other players. For each history $h \in \hat{H}$, we construct a pair (τ_{-i}^h, g^h), and our constructed type $\tau^{s_i,t}$ assigns equal weight to each of $|A_{-i}|^{t-1}$ such pairs. Each type τ_{-i}^h is constructed by applying the inductive hypothesis to a plan s_{-i}^h which plays according to history h so long as i follows s_i, and simply repeats the previous move forever if player i deviates. Such plans are sure-thing compliant for the player $-i$ 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 $f : \tilde{H} \times A_i \to \mathbb{R}$, 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 $f(\varnothing, s_i(\varnothing)) = 1$ and $f(\varnothing, a) = 0$ for all $a \neq s_i(\varnothing)$, where $\varnothing$ 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 $s_i(h, (s_i(h), a_{-i})) = s_i(h)$ for all a_{-i}, then let $f((h, a), \cdot) = f(h, \cdot)$ for every a.

Case 2: Otherwise, by sure-thing compliance, at least two different actions are prescribed for continuations $(h,(s_i(h),a_{-i}))$ as we vary a_{-i}. For each action $a_i \in A_i$, let $S_{a_i} = \{a_{-i} : s_i(h,(s_i(h),a_{-i})) = a_i\}$ be the set of continuations where a_i is prescribed. Then let

$$f((h,(s_i(h),a_{-i})),a_i) = \begin{cases} f(h,s_i(h))+\varepsilon & \text{if} \quad a_{-i} \in S_{a_i} \\ \frac{|A_{-i}|f(h,a_i)-|S_{a_i}|(f(h,s_i(h))+\varepsilon)}{|A_{-i}|-|S_{a_i}|} & \text{if} \quad a_{-i} \notin S_{a_i} \end{cases}$$

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

$$\text{(A.9)} \qquad f(h,a_i) \;=\; \frac{1}{|A_{-i}|}\sum_{a_{-i}} f((h,(s_i(h),a_{-i})),a_i)$$

$$\text{(A.10)} \qquad f(h,s_i(h)) \;\geq\; f(h,a_i)+\varepsilon \qquad (\forall h, a_i \neq s_i(h)).$$

as can be verified algebraically.

For each history $h \in \hat{H}$, define the stage-game payoff function $g^h : A \to [0,1]^n$ by setting $g_i^h(a) = f(h,a_i)$ and $g_j^h(a) = 0$ at each a and $j \neq i$. Define $\tau^{s_i,t}$ as mentioned above, by assigning equal weight to each pair (τ_{-i}^h, θ^h).

We claim that under rationalizable play, from the perspective of type $\tau^{s_i,t}$, when he has followed s_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 knows (assuming rationalizable play) the opposing types must be τ_{-i}^h and thus the stage-game payoff function must be g^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 $\tau^{s_i,t}$ 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 (A.9) which proves the claim.

Note also that if he follows s_i through period t, player i always learns his true payoff. Let $\bar{s}_i$ be the plan which follows s_i through period t, then plays the known optimal action from period $t+1$ onward. We claim that $\bar{s}_i$ strictly outperforms any plan which deviates by period t. The intuitive argument is as follows. Because type $\tau^{s_i,t}$ has stage-game payoffs which are insensitive to the other players' moves, he only has two possible incentives at each date: the myopic goal of maximizing his average stage-game payoffs at the current date, and the desire to receive further information about his payoffs. The former goal is strictly satisfied by the move prescribed by $\bar{s}_i$, 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 s_i' not t-equivalent to s_i and any rationalizable belief of $\tau^{s_i,t}$, the plan $\bar{s}_i$ 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 s_i. Let $\hat{h}$ be a random variable equal to the shortest realized history at which s_i' differs from s_i before period t, or ∞ if they do not differ by period t. 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 $\hat{h}$, $\bar{s}_i$ strictly outperforms s_i' on average. In fact this is weakly true date-by-date, and strictly true at the first deviation, because:

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

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

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

At dates $t + 1, ...$: Under plan $\bar{s}_i$, player i now has complete information about his payoff and optimizes perfectly, so s_i' cannot do better.

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

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 $[0, 1]$. $\square$

References

[1] Abreu, D. and F. Gul (2000): "Bargaining and Reputation," *Econometrica* 68, 85-117.

[2] Abreu, D. and D. Pearce (2007): "Bargaining, Reputation, and Equilibrium Selection in Repeated Games with Contracts," *Econometrica*, 75: 653–710.

[3] Chen, Y. (2012): "A structure theorem for rationalizability in the normal form of dynamic games," *Games and Economic Behavior*, 75, 587-597.

[4] Dekel, E. D. Fudenberg, S. Morris (2007): "Interim Correlated Rationalizability," *Theoretical Economics*, 2, 15-40.

[5] Fudenberg, D., D. Kreps, and D. Levine (1988): "On the Robustness of Equilibrium Refinements," *Journal of Economic Theory* 44, 354-380.

[6] Fudenberg, D. and D. Levine (1989): "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, 57, 759-778.

[7] Fudenberg, D. and E. Maskin (1986): "The Folk Theorem in Repeated Games with Discounting or Incomplete Information," *Econometrica*, 54, 533-557.

[8] Kajii, A. and S. Morris (1997): "The Robustness of Equilibria to Incomplete Information," *Econometrica*, 65, 1283-1309.

[9] Kreps, D. (1990): *A Course in Microeconomic Theory*, Princeton University Press, Princeton.

[10] 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.

[11] Lipman, B. (2003): "Finite Order Implications of Common Prior," *Econometrica*, 71, 1255-1267.

[12] Mailath, G. and L. Samuelson (2006): *Repeated Games and Reputations*, Oxford University Press.

[13] Weinstein, J. and M. Yildiz (2007): "A Structure Theorem for Rationalizability with Application to Robust Predictions of Refinements," *Econometrica*, 75, 365-400.

[14] Weinstein, J. and M. Yildiz (2013): "Robust predictions in infinite-horizon games—An unrefinable folk theorem," *Review of Economic Studies*

[15] Wolitzky, A. (2012): "Reputational Bargaining with Minimal Knowledge of Rationality," *Econometrica*, 80, 2047-2088.