ENDOGENOUS REPUTATION IN REPEATED GAMES

PRISCILLA T. Y. MAN

ABSTRACT. Reputation is often modelled by a small but positive prior probability that a player is a "behavioral type" in repeated games. This paper investigates the source of this reputation by allowing two fully rational, equally patient players in a standard infinitely repeated game to commit privately to any behavioral strategy they wish at the beginning of the game. Even with forward induction restrictions on off-equilibrium beliefs, a folk theorem holds: Every subgame perfect equilibrium payoff vector of the standard game is a sequential equilibrium payoff vector in the game with reputation choice when a public randomization device is available. This result holds for all finite two-person stage games, including common and conflicting interests games, as well as for all discount factors. The result indicates that the refinement power of reputation effect in repeated games relies crucially on the reputation being exogenous.

1. INTRODUCTION

In repeated games, reputation is often modelled by a small but positive prior probability of a player being a "behavioral type"[1]. Reputation effect is said to be present if the folk theorem breaks down given such priors. While reputation effect can usually be obtained with a small prior probability on a behavioral type, the restrictions on the support of this prior distribution are not so innocuous, particularly in games with two long-run players. For example, Aumann and Sorin (1989) allow only automatons with finite memory. Schmidt (1993, Section 3) also offers an example of how reputation effect may fail when there is a "perverse" type. How restrictive are these assumptions? Where do these "reputations" come from? How should we interpret these reputations?

This paper attempts to answer these questions by allowing players to choose the reputation they wish to have. To be specific, we consider a canonical infinitely repeated game with two ex-ante fully rational and equally patient players. They may privately

Date: March 17, 2011.
Key words and phrases. reputation, repeated games, forward induction.

[1]See Mailath and Samuelson (2006, Chapter 15-16) for a survey of the literature.

committed to any strategy they wish before the start of the standard repeated game. Commitment is costless and perfect. Since a committed player would behave like a "behavioral type" in the ensuing repeated game, by examining the set of equilibrium payoffs in this augmented game, we can investigate the strength of reputation effect when reputations are "endogenous".

However, by assigning high probability on "bad" types (for example, a type that plays her min-max strategy forever after) off the equilibrium path, one can sustain any subgame perfect equilibrium payoff easily. To avoid these artificial equilibria, we impose a forward induction restriction on off-equilibrium beliefs. In particular, if a player observes a deviation of her opponent, she must put probability 1 on the set of continuing strategies that can make the deviation optimal as long as these strategies exist. This can be thought of as a variation of the Intuitive Criterion (Cho and Kreps, 1987) in the sense that if a player "signals" her "commitment type" by a costly deviation, the opponent cannot put positive probability on types that can never gain by sending this "off-equilibrium message".

Even with a forward induction restriction on off-equilibrium beliefs, our main theorem states that any subgame perfect equilibrium payoff vector of a standard two-person infinitely repeated game with a public randomization device is an equilibrium payoff vector of the game with reputation choice. The folk theorem remains.

Our result does not depend on the particular class of the stage game (as long as it is a finite two-person game). Thus it covers classes of stage games studied for reputation effects for two equally patient players, including games with common interests (Aumann and Sorin, 1989; Cripps and Thomas, 1997), strictly conflicting interests (Cripps, Dekel, and Pesendorfer, 2005) and strictly dominant action games (Chan, 2000). Since generic payoff is not required, the stage game can also be the reduced normal form of an extensive form game. Moreover, if the stage game exhibits *aligned interests* — that is, any player can achieve her highest feasible payoff while allowing all opponents to get strictly above their min-max payoff, then our result extends to N-person games when the discount factor is sufficiently high.

Since we do not preclude equilibria in which some players commit, our result does not render reputation effects non-existent or inconsequential. Nevertheless, it does mean that if reputation effect is to be an equilibrium refinement in repeated games, then the source of this reputation must be exogenous. Behavioral types that can be motivated by exogenous circumstances, such as bounded memories or finitely complex strategies, are likely to be more plausible than those motivated by "a player would most like her opponent to believe she is of this type".

Closest to this paper is the work of Gul and Pearce (1996), who argue that forward induction has no refinement power in finite stage games in the presence of a public randomization device[2]. On the other hand, both Osborne (1990) and Al-Najjar (1995) obtain some refinement in finitely repeated games with forward induction concepts — the former with Stable Sets (Kohlberg and Mertens, 1986) and the latter with Forward Induction Equilibrium (FIE). In particular, Al-Najjar (1995) demonstrates that FIE implies a version of the reputation effects in Fudenberg and Levine (1989) in games with one long-run and one short-run player. Our paper differs from these work on two aspects: First, they consider finitely repeated games while we analyze infinitely repeated games. Second, both Gul and Pearce (1996) and Al-Najjar (1995) require continuations after any deviation to be subgame perfect (in the "unperturbed" game) whereas we allow an opponent's belief to be inconsistent with the uncommitted player's contingent plan after a player's deviation (because the opponent believes with probability 1 that the deviation is coming from a committed type). This continuation is not even a Nash equilibrium in the standard game without reputation choices. In this sense our equilibrium can be thought of as a self-confirming equilibrium (Fudenberg and Levine, 1993), which in general is a consequence of forward induction (see Battigalli and Siniscalchi, 2002, 2003).

The rest of the paper is organized as follows: Section 2 defines the game and the equilibrium concept. Section 3 states and proves the main theorem. Section 4 discusses some variations of our model. Section 5 concludes.

---

[2]See also the comment by Govindan and Robson (1998). Our construction avoids using inadmissible strategies largely due to the fact that we deal with infinitely repeated games as opposed to finite general stage games.

2. Set Up

The following conventions will be maintained through out this paper unless otherwise stated: For any two sets S and S', $S \subset S'$ means S is a proper subset of S' while $S \subseteq S'$ means S is a subset of S' but not necessarily proper. All sets are endowed with the Borel σ-algebra. If S is a finite set, ΔS denotes the set of all probability distributions over S. If S is a subset of a vector space, the convex hull of S is denoted as $\operatorname{co} S$. If λ is a probability distribution, $\operatorname{supp} \lambda$ is the support of λ.

2.1. **The Underlying Game.**

Consider a 2-person finite stage game $G = (A_i, u_i)_{i=1,2}$. A typical player will be called player i (a "she") and her opponent player $-i$ (a "he"). Let A be the set of joint action profiles. The set of mixed strategies available to player i is ΔA_i.

A standard infinite repetition of G is denoted as $G^\infty(\delta)$ where $\delta \in (0, 1)$ is the common discount factor. When no confusion may arise we will suppress the dependence on δ. All past realized actions are perfectly and publicly observed but individual randomization devices are not observed. Players have perfect recall. A public randomization device generates a publicly and perfectly observable signal $\omega^t \in [0, 1]$ at time t. The signals are independently and uniformly distributed on $[0, 1]$.[3] Let f denote the probability distribution function of the uniform distribution.

A history at time 1, the beginning of the game, is the realization of the first signal, ω^1. A history at time t is a list of past signal realization and action profile, denoted by $h^t = (\omega^1, a^1, \ldots, a^{t-1}, \omega^t)$. The set of all time t histories is $H^t = [0, 1] \times (A \times [0, 1])^{t-1}$. The set of all histories is $H = \cup_{t \geq 1} H^t$. A behavioral strategy of player i is a measurable function $s_i : H \to \Delta A_i$. Let S_i be the set of all behavioral strategies of player i. Since all histories start and end with a realization of the correlation signal, if hh' denotes the concatenation of the two histories h and h', it is assumed that the last signal in h is replaced by the first signal in h'. For any history h, let $s_i|h$ be the continuation strategy derived from s_i defined by $s_i|h(h') = s_i(hh')$.

[3]As in many standard settings, we can allow for any signal distribution that is absolutely continuous with respect to the Lebesgue measure. We assume the uniform distribution for simplicity.

A strategy s_i is said to *exclude* a history h if h can never be realized as long as player i adopts s_i. Consider a history $h^t = (\omega^1, a^1, \ldots, \omega^t)$. For any behavioral strategy of player i, $s_i \in S_i$, define the probability that s_i does not exclude h^t as

$$\beta_i(h^t, s_i) = \prod_{k=1}^{t-1} s_i(h^{tk})(a_i^k),$$

where h^{tk} is the first k-period history extracted from h^t. The density for nature not excluding h^t is given by $\beta_0(h^t) = \prod_{k=1}^{t} f(\omega^k) = 1$. Thus, given a strategy pair $s = (s_1, s_2)$, the density that h^t is reached can be written as

$$\beta(h^t, s) = \beta_0(h^t) \prod_{i=1,2} \beta_i(h^t, s_i) = \prod_{i=1,2} \beta_i(h^t, s_i).$$

The average discounted expected payoff of player i given the joint strategy profile s is then

$$U_i(s) = (1-\delta) \int_H \sum_t \delta^{t-1} u_i(a^t) \beta(h, s) dh.$$

Continuation payoffs at history h is similarly defined to be

$$U_i(s|h) = (1-\delta) \int_H \sum_t \delta^{t-1} u_i(a^t) \beta(h', s|h) dh'.$$

Let $\mathcal{E}(\delta)$ be the set of subgame perfect equilibrium payoff vectors of the game $G^\infty(\delta)$.

Before proceeding, it would be useful to remark on probability distributions over (behavioral) strategies of a player. The classic result of the realization equivalence of behavioral strategies and mixed strategies (Kuhn, 1953) continues to hold in our set up given the appropriate definition of mixed strategies (Mertens, Sorin, and Zamir, 2000, Theorem 1.8, p. 55). Thus a mixture of some behavioral strategies is considered to be the same mixture over their realization equivalent mixed strategies[4]. The "convex hull" of a subset of S_i should be similarly interpreted.

2.2. **The Game with Reputation Choice.** We augment $G^\infty(\delta)$ to a game $G^*(\delta)$ by giving players the choice to commit to any behavioral strategy they wish before G^∞

[4]While the infinite dimension of the history space creates technical inconveniences in defining mixed strategies, they can be well-defined by modelling explicitly a private randomization device. See Mertens, Sorin, and Zamir (2000, p. 55-56).

begins. At the start of the game, each player i can choose from the set $\Theta_i = S_i \cup \{*_i\}$ where $*_i$ represents the choice of not committing to any particular strategy. The commitment choice of player i is denoted by π_i, which is a probability distribution on Θ_i. Players choose their commitment types simultaneously. The commitment choice will not be revealed to the other player at any point in the game. After the commitment stage, players proceed to play G^∞. If a player has committed, she has no choice in G^∞. The strategy she has committed to will be perfectly implemented for her. We will refer to her as a "behavioral type" in G^∞. If a player chooses not to commit, she then play the ensuing game. We will refer to her as a "normal type".

The initial history (before the type choices) of G^* is denoted by $\emptyset$. A non-null history of G^* is of the form $((\theta_1, \theta_2), h)$ where θ_i is the (realized) type choice of player i. Given any $\theta_i \in \Theta_i$ and $h \in H$, an information set of player i is a collection of all histories $((\theta_i, \theta_{-i}), h)$ such that for all players $j = 1, 2$, either $\theta_j = *_j$ or $\beta_j(h, \theta_j) > 0$. Denote this information set as $h_i = (\theta_i, h)$. This h_i is also called a private history. Private histories after player i has committed to a behavioral type are uninteresting since player i has no choice there. Thus we will focus on the set of all private histories of player i of the form $(*_i, h)$, denoted by $\tilde{H}_i$. A behavioral strategy of player i in the play of G^∞ embedded in G^* is therefore a mapping from $\tilde{H}_i$ to $\Delta\, A_i$. Since there is a one-to-one relationship between $\tilde{H}_i$ and H, we will refer to a typical behavioral strategy (in the play of G^∞ embedded in G^*) by s_i, the same as in G^∞. A complete strategy of player i in G^* is therefore a pair (π_i, s_i)

Next we will define beliefs. Again, since player i has no choice after committing to a behavioral type, we are not interested in her beliefs at such information sets. We will define beliefs only on the set of all "relevant" information sets of player i, given by $H_i^* = \tilde{H}_i \cup \emptyset$. Even with this restriction, the infinite dimension introduced by the public randomization device still makes it technically hard to define beliefs as distributions over nodes within an information set. However, since player i's continuation payoff is not affected by the types of player $-i$ beyond the actions those types would take, player i's optimal response at an information set is determined by her opponent's *continuation strategy*. We will therefore define the belief of player i at her private history $h_i \in H_i^*$, denoted by $\mu_i(h_i)$, as a continuation strategy of player $-i$. As the game is infinitely repeated, $\mu_i(h_i) \in S_{-i}$ for all $h_i \in H_i^*$.

An assessment of G^* is a belief system μ and a strategy profile (π, s). Given a type θ_i of player i and her behavioral strategy s_i, the continuation strategy at her information set $h_i = (\theta_i, h)$ is defined as $\theta_i|h(h') = \theta_i(hh')$ for all $h' \in H$ (where $*_i|h$ is taken as $s_i|h$, the same definition applies hereafter).

2.3. **Equilibrium.** An equilibrium of G^* is an assessment $\langle \mu, (\pi, s)\rangle$. We require sequential rationality: the continuation strategy chosen at each information set of player i maximizes her continuation payoffs.

Definition 2.1 (Sequential Rationality)**.** An assessment $\langle \mu, (\pi, s)\rangle$ is *sequentially rational* if for all players i,

(1) At the initial history $\emptyset$, for all $\theta_i \in \operatorname{supp} \pi_i$,

$$\theta_i \in \arg\max_{\theta_i' \in \Theta_i} U_i(\theta_i', \mu_i(\emptyset));$$

(2) For all private histories $h_i = (*_i, h) \in H_i^*$,

$$s_i \in \arg\max_{s_i' \in S_i} U_i(s_i'|h, \mu_i(h_i)).$$

The infinite dimension of histories renders it impossible to define consistency of beliefs in the conventional manner. Instead, we define consistency as follows:

Definition 2.2 (Continuation Consistency)**.** An assessment $\langle \mu, (\pi, s)\rangle$ is *continuation consistent* if for all players i,

(1) For all histories $h \in H$, $\beta_{-i}(h, \mu_i(\emptyset)) = \int_{\Theta_{-i}} \beta_{-i}(h, \theta_{-i})\pi_i(\theta_{-i})d\theta_{-i}$.
(2) For all private histories $h_i \in H_i^*$, all histories $h \in H$, if $\beta_{-i}(h, \mu_i(h_i)) > 0$, then $\mu_i(h_i h) = \mu_i(h_i)|h$;[5]

The first condition requires initial beliefs to be "correct" up to realization equivalence. The second condition requires player i to have a "time consistent" belief over her opponent's continuation strategy — if she had believed her opponent to play s_{-i} before and

[5]If $h_i = \emptyset$, $h_i h$ is taken to be $(*_i, h)$.

no contradictory evidence has appeared yet, she should continue to believe that player $-i$ will play according to s_{-i}.

We show in Appendix A that in a finite version of this game (with a finite signal space for the public randomization device and repeated for finitely many periods), then conventional consistency coincides with our continuation consistency. In this sense, continuation consistency is a substitute for consistency in the definition of a sequential equilibrium in the particular class of games we analyze.

Continuation consistency leaves a lot of freedom on the off-equilibrium beliefs. By manipulating such off-equilibrium beliefs — for example, letting player i believe player $-i$ will play his min-max strategy forever after any of his deviations — one can sustain a large number of equilibrium outcomes. To preclude these artificial equilibria, we require player i to believe that player $-i$ is attempting to obtain a continuation payoff that makes $-i$'s deviation weakly profitable whenever such a belief is possible.

To define it formally we need some notations. Take player i. For any $v_i \in \mathbb{R}$, define the set of all behavioral strategies of i that can potentially achieve a payoff of at least v_i by

$$S_i^1(v_i) = \{s_i \in S_i : \exists s_{-i} \in S_{-i} \text{ s.t. } U_i(s_i, s_{-i}) \geq v_i\}.$$

This definition ignores the possibility that player $-i$ may never find it optimal to play the s_{-i} required. In light of this, we can also define

$$R_{-i}^1(v_i) = \left\{ s_{-i} \in S_{-i} : \exists s_i \in \operatorname{co} S_i^1(v_i) \text{ s.t. } s_{-i} \in \arg\max_{s'_{-i} \in S_{-i}} U_{-i}(s_i, s'_{-i}) \right\}$$

to be the set of optimal responses of player $-i$ against a belief over player i's continuation strategies concentrated on $S_i^1(v_i)$. Now define inductively, for any integer $k > 1$, given $S_i^{k-1}(v_i)$ and $R_{-i}^{k-1}(v_i)$

$$S_i^k(v_i) = \left\{ s_i \in S_i^{k-1}(v_i) : \exists s_{-i} \in R_{-i}^{k-1}(v_i) \text{ s.t. } U_i(s_i, s_{-i}) \geq v_i \right\}; \text{ and}$$
$$R_{-i}^k(v_i) = \left\{ s_{-i} \in R_{-i}^{k-1}(v_i) : \exists s_i \in \operatorname{co} S_i^k(v_i) \text{ s.t. } s_{-i} \in \arg\max_{s'_{-i} \in S_{-i}} U_{-i}(s_i, s'_{-i}) \right\}.$$

With a slight abuse of notations, define also

$$S_i^\infty(v_i) = \bigcap_{k\geq 1} S_i^k(v_i); \text{ and}$$

$$R_{-i}^\infty(v_i) = \bigcap_{k\geq 1} R_{-i}^k(v_i).$$

For a given v_i, all the above sets can be empty (they certainly are if v_i is above the highest feasible payoff to i). Nonetheless, for any k, there must be a v_i such that $S_i^k(v_i)$ is non-empty since Nash equilibrium strategy pairs can never be eliminated. Continuity of the payoff functions and compactness of the strategy space ensure that these sets are compact. The set of v_i such that $S_i^k(v_i)$ is non-empty is also compact for any k.

Now consider player i's private history $h_i = (*_i, h)$. Let s_i be player i's equilibrium strategy and $\mu_i(h_i) = s'_{-i}$. Then player i expects player $-i$ to get the continuation payoff $U_{-i}(s_i|h, s'_{-i})$. Suppose i observe an action $a'_{-i} \notin \operatorname{supp} s'_{-i}$. The lowest possible expected continuation payoff of player $-i$ such that the deviation is optimal is

$$v_{-i} = \frac{1}{\delta}\left[U_{-i}(s_i|h, s'_{-i}) - (1-\delta)\sum_{a_i\in A_i} u_{-i}(a'_{-i}, a_i)s_i(h)(a_i)\right]. \tag{1}$$

Notice that $S_{-i}^1(v_{-i})$ is the set of all possible continuation strategies player $-i$ can adopt after the deviation for his deviation to be optimal. If player i believes in the rationality of player $-i$ as long as it is possible, she would expect player $-i$ to continue with some strategies in $S_{-i}^1(v_{-i})$. Her best response to this belief would be in $R_i^1(v_{-i})$. Of course, if player $-i$ is more sophisticated, he should believed that player i would respond using only strategies in $R_i^1(v_{-i})$. Hence his set of "conditional rationalizable continuation strategies" is actually $S_{-i}^2(v_{-i})$, and so on. We can keep iterating until we get to $S_{-i}^\infty(v_{-i})$.

We will not take a particular stand on how many times we should iterate the above procedures. Instead, fix any $k > 0$, we define formally the above restrictions as follows:

Definition 2.3 (kth-order Forward Induction)**.** An assessment $\langle \mu, (\pi, s)\rangle$ satisfies *kth-order forward induction* if for all players i, all private histories $h_i = (*_i, h)$, if $\mu_i(h_i) = s'_{-i}$, then for all $a_i \in \operatorname{supp} s_i(h)$, all $a_{-i} \notin \operatorname{supp} s'_{-i}(h)$,

$$\mu_i((h_i, (a_{-i}, a_i))) \in \operatorname{co} S^k_{-i}(v_{-i}),$$

where v_{-i} is defined by Equation (1), whenever $S^k_{-i}(v_{-i})$ is non-empty.

Two points are worth noting. First, forward induction applies only when a unilateral deviation is detected. It puts no restriction on continuation beliefs when both players have deviated. Second, this restriction does not have to hold ω by ω. It suffices as long as the continuation strategy is ex-ante optimal.

Last but not least, we will focus on equilibria in which players do not commit. This leads us to our equilibrium concept:

Definition 2.4 (Strong Sequential Equilibrium)**.** An assessment $\langle \mu, (\pi, s) \rangle$ is a *kth-order strong sequential equilibrium* of the augmented repeated game G^* if it is sequentially rational, continuation consistent, satisfies kth-order forward induction and π_i is a degenerate distribution putting all probability mass on $*_i$ for both $i = 1, 2$.

Does a kth-order strong sequential equilibrium exists (for some k at least)? If they do, what are the equilibrium payoffs? Would the possibility of commitment types create reputation effects? We answer these questions in the next section.

3. Main Theorem

Theorem 3.1. *For all $\delta \in (0, 1)$, all $k > 0$ and all subgame perfect equilibrium payoff vectors $v \in \mathcal{E}(\delta)$ of the standard repeated game $G^\infty(\delta)$, there exists a kth-order strong sequential equilibrium of $G^*(\delta)$ with equilibrium payoffs equal to v.*

The rest of this section contains the proof. The proof consists of two parts. The first part defines a self-generation operator similar to that of Abreu, Pearce, and Stacchetti (1990) and shows that a self-generating set of payoffs is a set of kth-order strong sequential equilibrium payoffs. The second step shows $\mathcal{E}(\delta)$ is self-generating.

3.1. Self-Generation.

Recall that a self-generation operator (for a 2-person game) in Abreu, Pearce, and Stacchetti (1990) is an operator B_δ on sets $W \subseteq \mathbb{R}^2$ defined as follows: For any $w \in \mathbb{R}^2$, $w \in B_\delta(W)$ if there exists an $\alpha \in \Delta A_1 \times \Delta A_2$ and a function $\gamma : A \to \operatorname{co} W$ such that, for $i = 1, 2$,

$$w_i \geq \sum_{a_{-i} \in A_{-i}} \left[(1-\delta) u_i(a_i, a_{-i}) + \delta \gamma_i(a_i, a_{-i})\right] \alpha_{-i}(a_{-i}) \quad \text{for all } a_i \in A_i;$$

and if the inequality is strict for a_i, then $a_i \notin \operatorname{supp} \alpha_i$. A bounded set $W \subset \mathbb{R}^2$ is *self-generating* if $W \subseteq B_\delta(W)$.

Consider a compact set $W \subseteq \mathbb{R}^2$. For $i = 1, 2$, let

$$\overline{w}_i = \max_{w \in W} w_i$$

$$\underline{w}_i = \min_{w \in W} w_i.$$

Also, let $\overline{w}^i$ and $\underline{w}^i$ be the payoff vectors that achieve the maximum and minimum above, respectively.

Define an operator B^*_δ on compact sets $W \subseteq \mathbb{R}^2$ as follows: For any $w \in \mathbb{R}^2$, $w \in B^*_\delta(W)$ if there exists an $\alpha^* \in \Delta A_1 \times \Delta A_2$ and a function $\gamma^* : A \to \operatorname{co} W$ such that, for all $i = 1, 2$,

$$w_i \geq \sum_{a_{-i} \in A_{-i}} \left[(1-\delta) u_i(a_i, a_{-i}) + \delta \gamma^*_i(a_i, a_{-i})\right] \alpha^*_{-i}(a_{-i}) \quad \text{for all } a_i \in A_i; \tag{2}$$

and if the inequality is strict for a_i, then

(1) $a_i \notin \operatorname{supp} \alpha^*_i$; and

(2) $w_i \geq (1-\delta) \sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i}) \alpha^*_{-i}(a_{-i}) + \delta \overline{w}_i$.

Compared with B_δ, the only extra condition of B^*_δ is the requirement (2) when the inequality is strict for a_i.[6] This extra requirement says that when a_i is strictly suboptimal, it is not profitable even if player i obtains the highest possible continuation payoff in W after this deviation.

Definition 3.2. A compact set of payoffs $W \subset \mathbb{R}^2$ is B^*_δ-*self-generating* if $W \subseteq B^*_\delta(W)$.

[6]There is also the additional requirement that W is compact. This is inconsequential as we will focus on the set of subgame perfect equilibrium payoffs of G^∞, which is compact.

Remark 3.3. It is obvious that if W is B^*_δ-self-generating then it is self-generating. Recall that the set of subgame perfect equilibrium payoffs $\mathcal{E}(\delta)$ is the largest self-generating set (Abreu, Pearce, and Stacchetti, 1990). Hence if W is B^*_δ-self-generating, $W \subseteq \mathcal{E}(\delta)$.

Proposition 3.4. *Take $\delta \in (0,1)$. Let $W \subseteq \mathbb{R}^2$ be a compact, B^*_δ-self-generating set. Then for all $k > 0$, W is a set of kth-order strong sequential equilibrium payoffs of $G^*(\delta)$.*

Proof. The proof is separated into two parts. The first part constructs an auxiliary strategy pair. In the second part we use this auxiliary strategy pair to construct the kth-order strong sequential equilibrium desired.

Part 1: Auxiliary Strategies and Beliefs

Fixing k, define

$$v_i^k = \max\left\{v_i \in \mathbb{R} : S_i^k(v_i) \neq \emptyset\right\}$$

to be the highest achievable payoff for player i given kth-order of sophistication of players. The definition of $S_i^k(v_i^k)$ allows the possibility that for some $s_i \in S_i^k(v_i^k)$ and some $s_{-i} \in R_{-i}^{k-1}(v_i^k)$, $U_i(s_i, s_{-i}) > v_i^k$. However, this higher payoff is not possible as long as player $-i$ is best responding to player i, as the following lemma shows.

Lemma 3.5. *For all $s_i \in S_i$, if s_{-i} is a best response to s_i, then $U_i(s_i, s_{-i}) \leq v_i^k$.*

Proof. Suppose not. Then there exists an $v_i' > v_i^k$ such that $s_i \in S_i^1(v_i')$ (due to the existence of s_{-i}). Moreover, since s_{-i} is a best response to s_i, for any j, if $s_i \in S_i^j(v_i')$, then $s_{-i} \in R_{-i}^j(v_i')$. Hence $s_i \in S_i^{j+1}(v_i')$. By induction $s_i \in S_i^k(v_i')$ but this contradicts the fact that v_i^k is the highest payoff to i such that $S_i^k(v_i)$ is non-empty. □

If $v_i^k = \overline{w}_i$, define $\hat{s}$ to be the subgame perfect equilibrium strategy achieving $\overline{w}_i$ for player i. Define also $s_i' = \hat{s}_i$.

If instead $v_i^k > \overline{w}_i$, consider an infinitely repeated game $\hat{G}$ obtained from G^∞ by restricting player i's behavioral strategies to those in $S_i^k(v_i^k)$. There is no restriction on player $-i$'s strategies. The game $\hat{G}$ has a Nash equilibrium[7], call it $\hat{s}$.

[7]The existence of the Nash equilibrium can be proved as follows: first take a T-period version of the game and discretize the signal space of the correlation device. Take limit of the equilibria of these finite games as the discretized grid gets finer and show that there exists a limit that is also a Nash equilibrium of the T-period game with a continuous signal space. Finally, take limit as T goes to infinity and use the standard argument for repeated games with discounting to complete the argument.

Notice that we have some freedom of choosing the continuation play $\hat{s}|h$ at history h excluded by $\hat{s}_i$. This is because if $s_i \in S_i^k(v_i^k)$, so are all s_i' that differs from s_i at histories excluded by s_i. Also, since any unilateral deviation by player $-i$ from $\hat{s}_{-i}$ cannot reach such histories, changing the continuation play creates no profitable deviations for player $-i$. We will pick $\hat{s}$ such that for all h not excluded by $\hat{s}_i$, all $h' = (h, a, \omega)$ where $a_i \notin \operatorname{supp} \hat{s}_i(h)$, players continue by playing the subgame perfect equilibrium of G^∞ that achieves $\underline{w}_i$ for player i. (This subgame perfect equilibrium exists since $W \subseteq \mathcal{E}(\delta)$, see Remark 3.3.) Similarly, we can pick freely the continuation strategies of player $-i$ prescribed by $\hat{s}_{-i}$ at histories excluded by it.

By construction, $\hat{s}_i \in \operatorname{co} S_i^k(v_i^k)$. Since $\hat{s}_{-i}$ is a best response to $\hat{s}_i$, by Lemma 3.5 $U_i(\hat{s}) \leq v_i^k$. Moreover, no strategy $s_i \in S_i$ can achieve strictly higher than v_i^k against $\hat{s}_{-i}$. Clearly no strategy in $S_i^k(v_i^k)$ can achieve strictly higher than $U_i(\hat{s})$ by the definition of an equilibrium. Any strategy not in $S_i^k(v_i^k)$ achieves strictly less than v_i^k against all strategies in $R_{-i}^{k-1}(v_i^k)$ and $\hat{s}_{-i}$ is in this set. Let $s_i' \in S_i$ be a strategy that is a best response to $\hat{s}_{-i}$. Note that s_i' may not be in $S_i^k(v_i^k)$.

We will use $\hat{s}_i$, $\hat{s}_{-i}$ and s_i' in the construction to follow.

Part 2: Constructing the Equilibrium Assessment

Now suppose $W \subseteq B_\delta^*(W)$. Take any $w^* \in W$. We will construct a kth-order strong sequential equilibrium with payoffs w^*.

We will first construct a strategy (π_i, s_i) for each player i. The states in strategy are $W \cup \{\emptyset, \overline{v}_1^k, \overline{v}_2^k\}$. The $\overline{v}_i^k$'s are special states that, once these states are reached, continuation strategies depend on history since last entrance into the state[8]. The actions and transition rules are as follows:

(1) At the initial history $\emptyset$, both players choose $*_i$ with probability 1. Transit to state w^*.

(2) The current state is $w \in W$. Player i plays α_i^* in the definition of B_δ^* in the current period. After observing the realized action profile $a = (a_i, a_{-i})$, transit between states according to the following rules:

[8]As one will see from the construction, if $v_i^k = \overline{w}_i$, the state $\overline{v}_i^k$ will never be reached. In such cases one can construct the assessment without using this special state.

(a) If Inequalities (2) for a_i and a_{-i} are both weak or both strict, transit to states within W according to $\gamma^*(a)$ using the public randomization device.

(b) If Inequality (2) is strict for a_i but not a_{-i}, calculate

$$v_i = \frac{1}{\delta}\left[w_i - (1-\delta)\sum_{a'_{-i}\in A_{-i}} u_i(a_i, a'_{-i})\alpha^*_{-i}(a'_{-i})\right].$$

By Condition (2) of the operator B^*_δ, $v_i > \overline{w}_i$. If $v_i > v_i^k$, transit according to $\gamma^*(a)$ using the correlation device. If $v_i \in (\overline{w}_i, v_i^k]$, find $\lambda \in (0,1]$ such that $v_i = \lambda v_i^k + (1-\lambda)\overline{w}_i$. Using the public randomization device, transit to state $\overline{v}_i^k$ with probability λ and to $\overline{w}^i \in W$ that achieves $\overline{w}_i$ with probability $1-\lambda$.

(3) The current state is $\overline{v}_i^k$. Let h be the history since entrance into the state $\overline{v}_i^k$. Player i plays $s'_i(h)$ (defined in Part 1 above) in the current state. Player $-i$ plays $\hat{s}_{-i}(h)$ in the current state. Upon observing the realized joint action profile a, transit between states according to the following rules:

(a) If $a_j \in \operatorname{supp}\hat{s}_j(h)$ for both $j = 1, 2$, stay in $\overline{v}_i^k$ and update the history.

(b) For $j = 1, 2$, if $a_j \notin \operatorname{supp}\hat{s}_j(h)$ but $a_{-j} \in \operatorname{supp}\hat{s}_{-j}(h)$, calculate

$$v_j = \frac{1}{\delta}\left[U_j(\hat{s}|h) - (1-\delta)\sum_{a'_{-j}\in A_{-j}} u_j(a_j, a'_{-j})\hat{s}_{-j}(h)(a'_{-j})\right].$$

Depending on v_j, transit as follows:

(i) If $v_j < \underline{w}_j$ or $v_j > v_j^k$, transit to state $\underline{w}^j \in W$.

(ii) If $v_j \in [\underline{w}_j, \overline{w}_j]$, find $\lambda \in [0,1]$ such that $v_j = \lambda\overline{w}_j + (1-\lambda)\underline{w}_j$. Using the public randomization device, transit to $\overline{w}^j$ with probability λ and to $\underline{w}^j$ with probability $1-\lambda$.

(iii) If $v_j \in (\overline{w}_j, v_j^k]$, find $\lambda \in (0,1]$ such that $v_j = \lambda v_j^k + (1-\lambda)\overline{w}_j$. Using the public randomization device, transit to $\overline{v}_j^k$ with probability λ and to $\overline{w}^j$ with probability $1-\lambda$. (If $j = i$, transition to $\overline{v}_j^k$ means restarting the history since entrance into the state.)

(c) If $a_j \notin \operatorname{supp}\hat{s}_j(h)$ for both $j = 1, 2$, transit to state $\underline{w}^i \in W$.

The beliefs of the players are given as follows:

(1) At the initial history, $\mu_i(\emptyset)$ is the strategy of $-i$ constructed above (but believing that $-i$ will play $\hat{s}_{-i}$ if state $\overline{v}^k_{-i}$ is reached).

(2) At histories h where the current state is $w \in W$, set $\mu_i(h)$ to be the continuation strategy $s_{-i}|h$ according to the construction above.

(3) If the current state is $\overline{v}^k_i$ and h is the history since entrance into the current state, that is, the current history is some $h' = h''h$. Then $\mu_{-i}(h') = \hat{s}_i|h$ while $\mu_i(h') = \hat{s}_{-i}|h$.

To show that $\langle \mu, (\pi, s) \rangle$ constructed above is a kth-order strong sequential equilibrium, we first verify sequential rationality. Since the one-step deviation property holds in G^*, we check that there is no one-step deviation at all histories.

First note that there is a profitable one-step deviation at the initial history if and only if there is a profitable deviation in the ensuing game, so we will only check the other histories. If the current state is $w \in W$ and Inequality (2) holds as equality for a_i, player i cannot gain by changing the mixed strategy weight put on a_i. If Inequality (2) is strict for a_i, notice that upon entry to state $\overline{v}^k_i$, the expected payoff to player i is $U_i(s'_i, \hat{s}_{-i}) \leq v^k_i$ (the way we change the continuation payoffs does not create profitable deviation). However, λ is chosen such that

$$\begin{aligned} v_i &= \lambda v^k_i + (1-\lambda)\overline{w}_i \\ &\geq \lambda U_i(s'_i, \hat{s}_{-i}) + (1-\lambda)\overline{w}_i \end{aligned}$$

where v_i is the continuation payoff required to make the current period deviation to a_i weakly profitable. Thus playing $a_i \notin \operatorname{supp} \alpha^*_i$ is suboptimal. Finally, consider the case when the current state is $\overline{v}^k_i$. Given that player $-i$ believes (possibly incorrectly) that player i will play according to $\hat{s}_i$, playing $\hat{s}_{-i}$ is optimal by construction. Given that player i believes (correctly) that player $-i$ will play according to $\hat{s}_{-i}$, playing s'_i is optimal as well.

Next we verify continuation consistency. Condition (1) is satisfied because the only possible divergence between the belief of player i and the behavioral strategy of a normal player $-i$ occurs in state $\overline{v}^k_{-i}$, which is off the equilibrium path. Condition (2) is also satisfied since since player i never switches her beliefs.

Finally we demonstrate that the assessment satisfies kth-order forward induction. Forward induction puts no restriction on the equilibrium path. Similarly there is no restriction if both players have deviated. Now suppose we start in state $w \in W$ and a deviation by player i (to $a_i \notin \operatorname{supp} \alpha_i^*$) is observed. If Inequality (2) holds at equality for a_i, then the continuation strategy pair, say $\tilde{s}$, is a subgame perfect equilibrium in G^∞ in which the payoff to player i equals v_i in Definition 2.3. Thus $\tilde{s}_i \in S_i^k(v_{-i})$ for all k. If instead Inequality (2) is strict for a_i, recall that $\hat{s}_i \in \operatorname{co} S_i^k(v_i^k)$. Also, let $(\overline{s}_i, \overline{s}_{-i})$ be the subgame perfect equilibrium that achieves $\overline{w}^i$. Then $\overline{s}_i \in S_i^k(\overline{w}_i)$ and $\overline{s}_{-i} \in R_{-i}^{k-1}(\overline{w}_i)$. Thus if $v_i = \lambda v_i^k + (1-\lambda)\overline{w}_i$, then the strategy that plays $\hat{s}_i$ with probability λ and $\overline{s}_i$ with $1-\lambda$ using the correlation device is in $\operatorname{co} S_i^k(v_i)$. Hence $-i$'s off-equilibrium belief satisfies the restriction. With a similar argument, off-equilibrium beliefs after deviations from state $\overline{v}_i^k$ also satisfy kth-order forward induction.

Therefore $\langle \mu, (\pi, s)\rangle$ is a kth-order strong sequential equilibrium of $G^*(\delta)$ with payoffs $w^* \in W$. □

3.2. **Generating Subgame Perfect Equilibrium Payoffs.**

Proposition 3.6. *For all $\delta \in (0,1)$, the set of subgame perfect equilibrium payoffs $\mathcal{E}(\delta)$ is B_δ^*-self-generating.*

Proof. Fix $\delta \in (0,1)$. First note that $\mathcal{E}(\delta)$ is compact and convex in $\mathbb{R}^2$ (Abreu, Pearce, and Stacchetti, 1990). Hence we can denote the highest and lowest subgame perfect equilibrium payoffs to player i by

$$\overline{v}_i = \max_{v \in \mathcal{E}(\delta)} v_i$$

$$\underline{v}_i = \min_{v \in \mathcal{E}(\delta)} v_i.$$

By Abreu, Pearce, and Stacchetti (1990), $\mathcal{E}(\delta)$ is self-generating under the B_δ operator. Thus for any $v \in \mathcal{E}(\delta)$, there exists an $\alpha \in \Delta A_1 \times \Delta A_2$ and a function $\gamma : A \to \operatorname{co} \mathcal{E}(\delta)$

such that, for $i = 1, 2$,

$$v_i \geq \sum_{a_{-i} \in A_{-i}} \left[(1-\delta)u_i(a_i, a_{-i}) + \delta\gamma(a_i, a_{-i})\right] \alpha_{-i}(a_{-i}) \quad \text{for all } a_i \in A_i;$$

and if the inequality is strict for a_i, then $a_i \notin \operatorname{supp} \alpha_i$.

Take $v \in \mathcal{E}(\delta)$. Construct for it $\alpha^* = \alpha$. For each $a = (a_i, a_{-i}) \in A$, define $\gamma^*(a) = \gamma(a)$ if either $a_i \in \operatorname{supp} \alpha_i$ for all i or $a_i \notin \operatorname{supp} \alpha_i$ for all i. If $a_i \notin \operatorname{supp} \alpha_i$ but $a_{-i} \in \operatorname{supp} \alpha_{-i}$, let

$$\hat{\gamma}_i(a_i) = \frac{1}{\delta}\left[v_i - (1-\delta)\sum_{a'_{-i} \in A_{-i}} u_i(a_i, a'_{-i})\alpha_{-i}(a'_{-i})\right].$$

Pick $\gamma^*(a) = \gamma(a)$ if $\hat{\gamma}_i(a_i) > \overline{v}_i$; and pick $\gamma^*(a)$ such that $\gamma^*_i(a) = \hat{\gamma}_i(a_i)$ if $\hat{\gamma}_i(a_i) \leq \overline{v}_i$.

We first show that we can find such $\gamma^*(a) \in \mathcal{E}(\delta)$. This is obvious if $\gamma^*(a) = \gamma(a)$. In the only case when they differ, notice that self-generation under the B_δ operator guarantees

$$v_i \geq \sum_{a'_{-i} \in A_{-i}} \left[(1-\delta)u_i(a_i, a'_{-i}) + \delta\gamma(a_i, a'_{-i})\right] \alpha_{-i}(a'_{-i}).$$

Hence

$$\hat{\gamma}_i(a_i) = \frac{1}{\delta}\left[v_i - (1-\delta)\sum_{a'_{-i} \in A_{-i}} u_i(a_i, a'_{-i})\alpha_{-i}(a'_{-i})\right] \geq \sum_{a'_{-i} \in A_{-i}} \gamma_i(a_i, a'_{-i})\alpha_{-i}(a'_{-i}).$$

Since $\gamma(a) \in \mathcal{E}(\delta)$ for all $a \in A$, the right hand side of the above equation is no less than $\underline{v}_i$. Therefore $\overline{v}_i \geq \hat{\gamma}_i(a_i) \geq \underline{v}_i$. So choose $\lambda \in [0, 1]$ such that $\hat{\gamma}_i(a_i) = \lambda\overline{v}_i + (1-\lambda)\underline{v}_i$ Let $\overline{v}^i$ and $\underline{v}^i$ be the two payoff vectors in $\mathcal{E}(\delta)$ achieving $\overline{v}_i$ and $\underline{v}_i$ respectively. We can set

$$\gamma^*(a) = \lambda\overline{v}^i + (1-\lambda)\underline{v}^i,$$

which is obviously in $\mathcal{E}(\delta)$ since $\mathcal{E}(\delta)$ is convex.

Lastly, we need to show that the conditions of the B^*_δ operator are satisfied. Take $v \in \mathcal{E}(\delta)$. For all i, all $a_i \in A_i$,

$$\sum_{a'_{-i} \in A_{-i}} \left[(1-\delta)u_i(a_i, a'_{-i}) + \delta\gamma^*_i(a_i, a'_{-i})\right] \alpha^*_{-i}(a'_{-i})$$

$$
= \begin{cases} \sum_{a'_{-i} \in A_{-i}} \left[(1-\delta)U_i(a_i, a'_{-i}) + \delta\gamma_i(a_i, a'_{-i})\right] \alpha_{-i}(a'_{-i}) & \text{if } a_i \in \operatorname{supp}\alpha_i; \\ \sum_{a'_{-i} \in A_{-i}} (1-\delta)u_i(a_i, a'_{-i})\alpha_{-i}(a'_{-i}) + \delta\hat{\gamma}_i(a_i) & \text{if } a_i \notin \operatorname{supp}\alpha_i \text{ and } \hat{\gamma}_i(a_i) \le \overline{v}_i; \\ \sum_{a'_{-i} \in A_{-i}} \left[(1-\delta)U_i(a_i, a'_{-i}) + \delta\gamma_i(a_i, a'_{-i})\right] \alpha_{-i}(a'_{-i}) & \text{otherwise.} \end{cases}
$$

$$
= \begin{cases} v_i & \text{if } a_i \in \operatorname{supp}\alpha_i; \\ v_i & \text{if } a_i \notin \operatorname{supp}\alpha_i \text{ and } \hat{\gamma}_i(a_i) \le \overline{v}_i; \\ \sum_{a'_{-i} \in A_{-i}} \left[(1-\delta)U_i(a_i, a'_{-i}) + \delta\gamma_i(a_i, a'_{-i})\right] \alpha_{-i}(a'_{-i}) & \text{otherwise.} \end{cases}
$$

$$
\le v_i
$$

by the definition of the B_δ self-generation operator.

Thus Inequality (2) is satisfied for all a_i, with strict inequality only when $a_i \notin \operatorname{supp}\alpha_i$ but $\hat{\gamma}_i(a_i) > \overline{v}_i$. It remains to check that in this case, requirement (2) of the B^*_δ operator is satisfied. But this is almost immediate since

$$
\begin{aligned}
\overline{v}_i &< \hat{\gamma}_i(a_i) \\
\overline{v}_i &< \frac{1}{\delta}\left[v_i - (1-\delta)\sum_{a'_{-i} \in A_{-i}} u_i(a_i, a'_{-i})\alpha_{-i}(a'_{-i})\right] \\
v_i &> (1-\delta)\sum_{a'_{-i} \in A_{-i}} u_i(a_i, a'_{-i})\alpha_{-i}(a'_{-i}) + \delta\overline{v}_i.
\end{aligned}
$$

Therefore $\mathcal{E}(\delta)$ is B^*_δ-self-generating. □

Theorem 3.1 now follows from Propositions 3.4 and 3.6.

4. Alternative Specifications

4.1. **N-Player Games.** Extending the result to games with more than 2 players involves conceptual issues regarding the definition of forward induction. In particular, since our S^k_i and R^k_{-i} operators implicitly *eliminates* strategies, a seemingly natural extension of the definition to a N-player version would assumes more than (common strong) beliefs in player i's rationality — it implies all players render any joint profile involving the eliminated strategy as infinitely unlikely than any remaining strategy profiles[9]. Therefore, instead of giving a precise definition of forward induction for N-player repeated games with reputation choice, we provide two minimal requirements on the definition of forward induction. We then introduce a class of N-player games to which our main theorem extends as long as these two minimal requirements are satisfied and the discount factor is sufficiently high.

Let $G = (A_i, u_i)_{i \in N}$ be a finite stage game. Let $\overline{u}_i$ and v_i^m be the highest feasible payoff and the min-max payoff to player i in G respectively. Denote the infinite reputation of G as G^∞ and the augmentation of G^∞ with reputation choice as G^*. For a moment suppose $S_i(v_i)$ is defined for all players $i \in N$ and all $v_i \in \mathbb{R}$. Then the following definition of forward induction follows from Definition 2.3:

Definition 4.1. An assessment $\langle \mu, (\pi, s) \rangle$ of G^* satisfies forward induction if for all players $i \in N$, all private histories $h_i = (*_i, h)$ where $\mu_i(h_i) = s'_{-i} \in \times_{j \neq i} S_j$, if there is exactly one $j \neq i$ such that $a_j \notin \operatorname{supp} s'_j(h)$ and $a_k \in \operatorname{supp} s'_k(h)$ for all $k \neq j$, then

$$\operatorname{Proj}_j \mu_i((h_i, a)) \in \operatorname{co} S_j(v_j),$$

where Proj_j denotes the projection onto S_j, and

$$v_j = \frac{1}{\delta}\left[U_j(s_i | h, s'_{-i}) - (1-\delta) \sum_{a'_{-j} \in A_{-j}} u_j(a_j, a'_{-j}) s'_{-j}(h)(a'_{-j})\right];$$

whenever $S_j(v_j)$ is nonempty.

We can now give two minimal requirements on forward induction in terms of $S_i(v_I)$:

Definition 4.2. A forward induction restriction on G^* satisfies

Feasibility: if for all $i \in N$, $S_i(v_i) = \emptyset$ whenever $v_i > \overline{u}_i$.

Nash Axiom: if for all $i \in N$, all $v_i \in \mathbb{R}$, if $s = (s_1, \ldots, s_N)$ is a Nash equilibrium of G^∞ such that $U_i(s) \geq v_i$, then $s_i \in S_i(v_i)$.

A strong sequential equilibrium of G^* is again an assessment that is sequentially rational, continuation consistent, satisfies forward induction and in which no player commits.

[9] See Man (2010, Figure 11) for an example.

Next we define the class of games to which our main theorem extends.

Definition 4.3. A finite stage game $G = (A_i, u_i)_{i \in N}$ exhibits *aligned interests* if for each players $i \in N$, there exists an $\overline{a}^i \in \arg\max_{a \in A} u_i(a)$ such that at least one of the following is true:

(1) $\overline{a}^i$ is a Nash equilibrium of G; or
(2) For all $j \neq i$, $u_j(\overline{a}^i) > v_j^m$.

Obviously, all common interests games exhibit aligned interests. Battle of the Sexes is also an aligned interest game. Prisoner's Dilemma is not. It is somewhat unfortunate that games labelled as "conflicting interests games" or even "strictly conflicting interests games" in the literature may still be aligned interests games according to our definition. This is because the definitions in the literature consider the Stackelberg payoff while ours considers the highest feasible payoff.

Corollary 4.4. *Let G be an aligned interests game. Suppose the definition of forward induction satisfies feasibility and the Nash Axiom. Then there exists a $\underline{\delta} \in (0,1)$ such that for all $\delta \in (\underline{\delta}, 1)$, every subgame perfect equilibrium payoff vector of $G^\infty(\delta)$ is a strong sequential equilibrium payoff vector of $G^*(\delta)$.*

Proof. Since G exhibits aligned interests, by the standard folk theorem there exists a discount factor $\underline{\delta} \in (0,1)$ such that for all $\delta \in (\underline{\delta}, 1)$ and for all i, the payoff vector $u(\overline{a}^i)$ is a subgame perfect equilibrium payoff vector of the infinitely repeated game $G^\infty(\delta)$. Feasibility and Nash Axiom then imply $v_i^k = \overline{u}_i$. One can complete the proof by B_δ^*-generate $\mathcal{E}(\delta)$ without using the special states (so that all continuations are subgame perfect). □

4.2. **Definitions of Forward Induction.** In their work on forward induction in finitely repeated games, Al-Najjar (1995) and Gul and Pearce (1996) adopt definitions of forward induction equilibrium concept requiring the continuation play at any history to be a subgame perfect equilibrium. In our framework, we can let $S^*(\delta)$ to be the set of all subgame perfect equilibria strategy profile of G^∞ and define

$$S_i(v_i) = \{s_i \in S_i : \exists\, s^* \in S^*(\delta) \text{ s.t. } u_i(s^*) \geq v_i \text{ and } s_i^* = s_i\}\,.$$

We can then replace the S_{-i}^k in Definition 2.3 by $S_{-i}(v_{-i})$ to obtain a new definition of forward induction and hence strong sequential equilibrium. Under this definition, our main theorem continue to hold:

Corollary 4.5. *Let G be a finite two person game. For all $\delta \in (0,1)$ and all subgame perfect equilibrium payoff vectors $v \in \mathcal{E}(\delta)$ of $G^\infty(\delta)$, there exists a strong sequential equilibrium of $G^*(\delta)$ with equilibrium payoffs equal to v.*

Proof. The definition of $S_i(v_i)$ implies $v_i^k = \overline{v}_i$ for both $i = 1, 2$. Theorem 3.1 can be proved by B_δ^*-generating $\mathcal{E}(\delta)$ without using the special states. □

Our S^k set is not exactly the set of rationalizable or iteratively admissible strategies as in Bernheim (1984) and Pearce (1984). This is because strategies eliminated in our iterations may not be iteratively dominated (they may be a best response against other opponent strategies). Since there are usually not many dominated strategies in a typical infinitely repeated games, sequential equilibrium in rationalizable strategies is likely to be a weaker requirement than our strong sequential equilibrium.

Finally, some other existing forward induction concepts simply do not extend well to infinitely repeated games. Refinements for signaling games such as the Intuitive Criterion (Cho and Kreps, 1987), D1, D2 and their iterative versions (Banks and Sobel, 1987) do not take care of continuation strategies of the "sender" of the message and hence are inappropriate in infinitely repeated games. Equilibrium refinements based on strategic stability for general finite extensive form games (see Kohlberg and Mertens, 1986; Govindan and Wilson, 2009; Man, 2010), on the other hand, are technically hard to extend to infinite games.

4.3. **Public Randomization.** Loosely speaking, the intuition behind our result is a signaling idea. If commitment is known, then a player would like to commit to her Stackelberg strategy. However, when the commitment is unknown, the opponent can rationally choose some strategy that renders her Stackelberg strategy suboptimal. Thus it is crucial for player i to inform her opponent of her commitment. She can send this message by taking a costly deviation. If player i can potentially get a payoff strictly higher than that in the equilibrium and the deviation makes a "normal type" strictly worse off, forward induction requires player $-i$ to believe that player i is "crazy". The best response to this belief can make player i's deviation strictly profitable.

Our construction eliminates such signaling opportunities mostly by making a player indifferent between the equilibrium continuation and deviation. Our proof relies on the public randomization device to achieve this exact indifference. Is the public randomization device indispensable? No if there is another method to convexify the payoff set. However, the usual technique of replacing a public randomization device with sequences of action profiles (see Fudenberg and Maskin, 1991) may fail as we need to eliminate also the signaling opportunities along the sequence. It is unclear whether this is feasible.

5. Conclusion

This paper analyzes repeated games with endogenous reputation choices. We show that even with a forward induction restriction on the off-equilibrium beliefs, every subgame perfect equilibrium payoff of a standard two-person repeated game with a public randomization device can be sustained in a sequential equilibrium in which no players would find it optimal to establish a reputation at the first place.

Our result implies that the refinement power of "reputation effect" in the literature relies critically on the prior distribution being exogenous. While our result does not rule out the possibility of equilibria in which a player commits with positive probability, it does state that for every target payoff vector, there is a strategy deterring any commitment at the first place. In light of this, reputations motivated by exogenous constraints are likely to be better justified than reputations that are results of previous choices.

Alternatively, since no player commits in the equilibrium we have constructed, our result can be interpreted as the failure of forward induction to refine the set of equilibria in infinitely repeated games with a correlation device. This statement is not entirely true. While no subgame perfect equilibrium payoff is eliminated, not every equilibrium survives.

For example, equilibria sustained by the grim trigger strategy are ruled out. Nevertheless, the observation that forward induction weakens when the payoff set is convexified is fairly general. In the "burning money" game, Ben-Porath and Dekel (1992) give an example in their Remark 2.2 that signaling by burning a dollar may fail when the game without burning has a continuum of strategies for each player. The payoff set in that particular example is again convex.

Appendix A. Alternative Definitions of Beliefs and Consistency

In this appendix we consider a finite version of G^*. We will first define the finite game. There are two ways of describing beliefs in this game: the conventional method of defining beliefs as probability distributions over nodes in an information set; and the method we use in the main text that defines beliefs as an expected continuation strategy of the opponent. We then give give two propositions establishing the equivalence between consistency of conventional assessment and continuation consistency of the alternative assessment. We take this equivalence as a justification of substituting consistency with continuation consistency in the infinite game in our main text.

A.1. **The Game.** Let G^T be a standard T-fold repetition of a two-person finite stage game $G = (A_i, u_i)_{i=1,2}$. To keep the game finite, we discretize the signal space of the public randomization device to a finite set Ω. Signals are independent across periods and identically distributed according to the probability mass function f. Histories in G^T are defined as usual. We will call the set of all histories H. A behavioral strategy of player i is a mapping from histories to a stage game mixed action, $s_i : H \to \Delta A_i$. The set of all behavioral strategies is S_i. Concatenation of histories and continuation strategies are defined as in the main text.

As in the main text, the probability that player i's behavioral strategy s_i does not exclude history h is denoted as $\beta_i(h, s_i)$. Let Z be the set of all terminal nodes in G^T. Two pure (behavioral) strategies $s_i, s_i' \in S_i$ are *realization equivalent* if $\beta_i(z, s_i) = \beta_i(z, s_i')$ for all $z \in Z$. Let S_i^* be the collection of equivalence classes of the pure strategies defined by realization equivalence. Notice that if ρ_i is a mixed strategy (i.e., a distribution over the pure strategies) of player i in G^T, then given player $-i$'s strategy and nature's choices, the outcome of the game depends only on the probability ρ_i puts on each equivalence class in S_i^*. Thus we will represent mixed strategies as distributions over S_i^*. A behavioral strategy $s_i \in S_i$ and a mixed strategy ρ_i are *realization equivalent* if

$$\beta_i(z, s_i) = \sum_{s_i' \in S_i^*} \beta_i(z, s_i')\rho_i(s_i') \quad \text{for all } z \in Z.$$

The probability of nature not excluding a t-length history h^t is given by $\beta_0(h^t) = \prod_{k=1}^{t} f(\omega^k)$.

As in the main text we augmented G^T by allowing reputation choice at the beginning of the game. Player i can choose at the initial history a "reputation" among the set $\Theta_i = S_i^* \cup \{*_i\}$. The commitment choice is denoted by $\pi_i \in \Delta\,\Theta_i$. Call the augmented game G^{T*}. Histories, information sets and strategies in G^{T*} are as defined in G^*.

There will be two ways of describing beliefs. The first way describes beliefs as distribution over nodes in information sets. We call such a belief system a "conventional belief system", typically denoted by μ. We will use $\mu_i(\theta_{-i}|h_i)$ to denote the probability player i assigns to player $-i$ being of type θ_{-i} at her information set h_i according to the belief system μ. The second way — the one we adopted in the main text — describes beliefs as an expected continuation strategy. Specifically, player i belief at $h_i = (*_i, h)$,[10] denoted by $\hat{\mu}_i(h_i)$, is a continuation strategy of player $-i$ at the history h. We call such a belief system a "continuation belief system".

A conventional assessment of $\hat{G}$ is a conventional belief system and a behavioral strategy profile, $\langle \mu, (\pi, s)\rangle$. A continuation assessment is a continuation belief system and a behavioral strategy profile $\langle \hat{\mu}, (\pi, s)\rangle$.

A.2. **Equivalence between Conventional and Continuation Belief Consistency.** It would be useful to introduce a couple notations and definitions before we discuss the relationship between conventional and continuation belief systems. First of all, as a notation convention, given (π_i, s_i), for any histories $h \in H$, $\beta_i(h, *_i)$ is defined to be $\beta_i(h, s_i)$.

[10]We define beliefs only at information sets after player i has chosen not to commit since she has no choice at other information sets.

Next we give two definitions of belief consistency, one for each way of describing beliefs. Consistency of conventional assessment is standard:

Definition A.1 (Consistency)**.** A conventional assessment $\langle \mu, (\pi, s) \rangle$ is *consistent* if there exists a sequence of completely mixed behavioral strategy profiles $\{(\pi^n, s^n)\}$ converging to (π, s) and sequence of belief systems $\{\mu^n\}$ derived from (π^n, s^n) using Bayes rule such that $\{\mu^n\}$ converges to μ.

Consistency of continuation assessment has been given in Definition 2.2.

The next definition describes when we can consider a conventional belief system and a continuation belief system "equivalent":

Definition A.2 (Compatibility)**.** A conventional belief system μ and a continuation belief system $\hat{\mu}$ are *compatible* if for all players $i = 1, 2$, all private histories $h_i = (*_i, h) \in H_i^*$, all actions of the opponent $a_{-i} \in A_{-i}$,

$$\hat{\mu}_i(h_i)(a_{-i}) = \sum_{\theta_{-i} \text{ s.t. } \beta_{-i}(h, \theta_{-i}) > 0} \mu_i(\theta_{-i}|h_i)\theta_{-i}(h)(a_{-i}).$$

The next two propositions are the "main theorems" in this appendix. They together establish a notion of equivalence between consistency of conventional assessments and consistency of continuation assessments in the class of games we focus on.

Proposition A.3. *Let $\langle \mu, (\pi, s) \rangle$ be a consistent conventional assessment. Then there exists a continuation belief system $\hat{\mu}$ compatible with μ such that $\langle \hat{\mu}, (\pi, s) \rangle$ is a continuation consistent assessment.*

Proof. Since $\langle \mu, (\pi, s) \rangle$ is consistent, let $\{(\pi^n, s^n)\} \to (\pi, s)$ and μ^n be the Bayes' Rule derived sequence of conventional belief systems converging to μ. Abuse notations and adopt the shorthand $\mu_i^n(\cdot|h) = \mu_i^n(\cdot|(*_i, h))$. Define for each n, for all players i, a behavioral strategy of her opponent $\hat{s}_{-i}^n \in S_{-i}$ by specifying: for all $h \in H$, all $a_{-i} \in A_{-i}$,

$$\hat{s}_{-i}^n(h)(a_{-i}) = \sum_{\theta_{-i} \text{ s.t. } \beta_{-i}(h, \theta_{-i}) > 0} \mu_i(\theta_{-i}|h)\theta_{-i}(h)(a_{-i}).$$

Now construct, for each n and each player i, her continuation belief by defining: for all $h_i \in H_i^*$,

$$\hat{\mu}_i^n(h_i) = \begin{cases} \hat{s}_{-i}^n & \text{if } h_i = \emptyset; \\ \hat{s}_{-i}^n | h & \text{if } h_i = (*_i, h). \end{cases}$$

Let $\hat{\mu} = \lim_{n\to\infty} \hat{\mu}^n$. This limit is well-defined since $\hat{s}^n$ is continuous in μ^n and s^n. We claim that $\hat{\mu}$ satisfies the requirement in the proposition.

Condition (2) of continuation consistency and compatibility between μ and $\hat{\mu}$ are satisfied by construction. It remains to show that Condition (1) of continuation consistency is satisfied. We will show a stronger result: that for all n, for all players i and all histories h,

$$\beta_{-i}(h, \hat{\mu}_i^n(\emptyset)) = \sum_{\theta_{-i}\in\Theta_{-i}} \beta_{-i}(h, \theta_{-i})\pi_i^n(\theta_{-i}). \tag{3}$$

The desired result will then follow by continuity.

We will prove this by an induction on the length of histories. Take any n. Consider a period 1 history ω^1. Since no strategy can ever exclude this history, $\beta_{-i}(\omega^1, s'_{-i}) = 1$ for all players i and all strategies $s'_{-i} \in S_{-i}$. Thus,

$$\beta_{-i}(\omega^1, \hat{s}_{-i}^n) = 1 = \sum_{\theta_{-i}\in\Theta_{-i}} \beta_{-i}(\omega^1, \theta_{-i})\pi_{-i}^n(\theta_{-i}).$$

Now suppose for all period $t-1$ histories h^{t-1}, Equation (3) holds. Consider a period t history $h^t = (h^{t-1}, a, \omega)$. We have

$$\begin{aligned}
\beta_{-i}(h^t, \hat{s}_{-i}^n) &= \beta_{-i}(h^{t-1}, \hat{s}_{-i}^n)\hat{s}_{-i}^n(h^{t-1})(a_{-i}) \\
&= \beta_{-i}(h^{t-1}, \hat{s}_{-i}^n) \sum_{\theta_{-i} \text{ s.t. } \beta_{-i}(h^{t-1},\theta_{-i})>0} \mu_i^n(\theta_{-i}|(*_i, h^{t-1}))\theta_{-i}(h^{t-1})(a_{-i}) \\
&= \beta_{-i}(h^{t-1}, \hat{s}_{-i}^n) \sum_{\theta_{-i}\in\Theta_{-i}} \frac{\beta_{-i}(h^{t-1}, \theta_{-i})\pi_{-i}^n(\theta_{-i})}{\sum_{\theta'_{-i}\in\Theta_{-i}} \beta_{-i}(h^{t-1}, \theta_{-i})\pi_{-i}^n(\theta_{-i})}\theta_{-i}(h^{t-1})(a_{-i}).
\end{aligned}$$

By the induction hypothesis, this is equal to

$$\sum_{\theta_{-i}\in\Theta_{-i}} \beta_{-i}(h^{t-1}, \theta_{-i})\pi_{-i}^n(\theta_{-i})\theta_{-i}(h^{t-1})(a_{-i})$$

$$= \sum_{\theta_{-i}\in\Theta_{-i}} \beta_{-i}(h^t, \theta_{-i})\pi^n_{-i}(\theta_{-i}). \qquad \square$$

Proposition A.4. *Let $\langle\hat{\mu}, (\pi, s)\rangle$ be a continuation consistent assessment. Then there exists a conventional belief system μ compatible with $\hat{\mu}$ such that $\langle\mu, (\pi, s)\rangle$ is consistent.*

Proof. Let $\langle\hat{\mu}, (\pi, s)\rangle$ be a continuation consistent assessment. For each player i, construct a behavioral strategy $\hat{s}_i$ using $-i$'s beliefs by specifying for all histories $h \in H$, all actions $a_i \in A_i$,

$$\hat{s}_i(h)(a_i) = \hat{\mu}_{-i}((*_{-i}, h))(a_i).$$

(In other words, $\hat{s}_i$ is the behavioral strategy constructed from $\hat{\mu}_{-i}$ after replacing continuation strategies at excluded histories with the "correct" continuation.) Now construct from $\hat{s}_i$ a sequence by defining, for all histories $h \in H$, all actions $a_i \in A_i$,

$$\overline{s}^n_i(h)(a_i) = \begin{cases} \overline{s}_i(h)(a_i) & \text{if } \overline{s}_i(h)(a_i) > 0; \\ \frac{1}{n} & \text{otherwise.} \end{cases}$$

Scale the $\overline{s}^n_i$ so that the probabilities add up to 1:

$$\hat{s}^n_i(h)(a_i) = \frac{\overline{s}^n_i(h)(a_i)}{\sum_{a'_i\in A_i} \overline{s}^n_i(h)(a'_i)}.$$

Thus for all i and all n, $\hat{s}^n_i$ is a completely mixed behavioral strategy of player i in G^T. By construction $\hat{s}^n_i \to \hat{s}_i$ as n goes to infinity.

For each player i and each n, let $\hat{\rho}^n_i$ be a mixed strategy realization equivalent to $\hat{s}^n_i$ and has full support on S^*_i. This completely mixed strategy exists by the theorem of Kuhn (1953).

Define the set of equivalence classes of pure strategies of player i that enables histories excluded by $\hat{\mu}_{-i}(\emptyset)$ (and hence by $\hat{s}_i$) as

$$\hat{S}_i = \{s_i \in S^*_i : \exists\, h \in H \text{ s.t. } \beta_i(h, \hat{\mu}_{-i}(\emptyset)) = 0 \text{ and } \beta_i(h, s_i) > 0\}.$$

Clearly, $\pi_i(\hat{S}_i) = 0$ (or else $\beta_i(h, \hat{\mu}_{-i}(\emptyset)) > 0$ for some of these histories h). Since $\hat{\rho}^n_i$ is realization equivalent to $\hat{s}^n_i$, $\hat{\rho}^n_i(\hat{S}_i) \to 0$. Now for each n, find M^n_i as the highest order in terms of $\frac{1}{n}$ of $\hat{\rho}_i^n(s_i)$ among all $s_i \in \hat{S}_i$. Since G^T is finite, we can choose $\hat{\rho}_i^n$ such that M_i^n is finite for all n.

We will now construct the sequence of behavioral strategy profiles for the consistent conventional belief system of G^{T*}, the game with reputation choice. For all players i, all $\theta_i \in \Theta_i$, define

$$\tilde{\pi}_i^n(\theta_i) = \begin{cases} \pi_i(\theta_i) & \text{if } \theta_i \in \operatorname{supp} \pi_i \\ \hat{\rho}_i^n(\theta_i) & \text{if } \theta_i \in \hat{S}_i \\ \left(\frac{1}{n}\right)^{M_i^n+1} & \text{if } \theta_i = *_i \text{ and } \pi_i(*_i) = 0 \\ \frac{1}{n} & \text{otherwise.} \end{cases}$$

And for all histories $h \in H$, all actions $a_i \in A_i$,

$$\tilde{s}_i^n(h)(a_i) = \begin{cases} s_i(h)(a_i) & \text{if } s_i(h)(a_i) > 0 \\ \frac{1}{n} & \text{if } s_i(h)(a_i) = 0 \text{ and } \pi_i(*_i) = 0 \\ \left(\frac{1}{n}\right)^{M_i^n+1} & \text{otherwise.} \end{cases}$$

Next, scale these "probabilities" so that they add up to one. That is,

$$\pi_i^n(\theta_i) = \frac{\tilde{\pi}_i^n(\theta_i)}{\sum_{\theta_i' \in \Theta_i} \tilde{\pi}_i^n(\theta_i')} \quad \text{for all } \theta_i \in \Theta_i.$$

Similarly,

$$s_i^n(h)(a_i) = \frac{\tilde{s}_i^n(h)(a_i)}{\sum_{a_i' \in A_i} \tilde{s}_i^n(h)(a_i')} \quad \text{for all } h \in H, \text{ all } a_i \in A_i.$$

It is straight-forward to check that $\{(\pi^n, s^n)\} \to (\pi, s)$. Using Bayes' rule, for all players i, all private histories $h_i = (\theta_i, h)$, all opponent types $\theta_{-i} \in \Theta_{-i}$, we get the conventional belief by

$$\begin{aligned} \mu_i^n(\theta_{-i}|h_i) &= \frac{\beta_0(h)\beta_i(h,\theta_i)\beta_{-i}(h,\theta_{-i})\pi_{-i}^n(\theta_{-i})}{\sum_{\theta'_{-i} \in \Theta_{-i}} \beta_0(h)\beta_i(h,\theta_i)\beta_{-i}(h,\theta'_{-i})\pi_{-i}^n(\theta'_{-i})} \\ &= \frac{\beta_{-i}(h,\theta_{-i})\pi_{-i}^n(\theta_{-i})}{\sum_{\theta'_{-i} \in \Theta_{-i}} \beta_{-i}(h,\theta'_{-i})\pi_{-i}^n(\theta'_{-i})}. \end{aligned} \tag{4}$$

Let $\mu = \lim_n \mu^n$. Then $\langle \mu, (\pi, s) \rangle$ is consistent.

It remains to show that μ is compatible with $\hat{\mu}$. First notice that for all players i, all behavioral strategies $s_i \in S_i$, all histories $h \in H$, all actions $a_i \in A_i$,

$$\beta_i((h,(a_i,a_{-i}),\omega),s_i) = \beta_i(h,s_i)s_i(h)(a_i) \quad \text{for all } a_{-i} \in A_{-i}, \text{ all } \omega \in \Omega.$$

Thus adopt the shorthand $\beta_i((h,a_i),s_i) = \beta_i(h,s_i)s_i(h)(a_i)$. Next, using Equation (4) and given n, we can write for all players i, all private histories $h_i = (*_i, h) \in H_i^*$ and all actions $a_{-i} \in A_{-i}$,

$$\begin{aligned}\sum_{\substack{\theta_{-i}\in\Theta_{-i}\\ \beta_{-i}(h,\theta_{-i})>0}} \mu_i^n(\theta_{-i}|h_i)\theta_{-i}(h)(a_{-i}) &= \sum_{\theta_{-i}\in\Theta_{-i}} \frac{\beta_{-i}(h,\theta_{-i})\pi_{-i}^n(\theta_{-i})}{\sum_{\theta'_{-i}\in\Theta_{-i}}\beta_{-i}(h,\theta'_{-i})\pi_{-i}^n(\theta'_{-i})}\theta_{-i}(h)(a_{-i}) \\ &= \frac{\sum_{\theta_{-i}\in\Theta_{-i}}\beta_{-i}((h,a_{-i}),\theta_{-i})\pi_{-i}^n(\theta_{-i})}{\sum_{\theta'_{-i}\in\Theta_{-i}}\beta_{-i}(h,\theta'_{-i})\pi_{-i}^n(\theta'_{-i})}. \qquad (5)\end{aligned}$$

There are two cases possible. First consider h such that $\beta_{-i}(h,\hat{\mu}_i(\emptyset)) > 0$. Then h is reached by some types θ_{-i} such that $\pi_{-i}(\theta_{-i}) > 0$. As n goes to infinity, the expression in Equation (5) converges to

$$\frac{\sum_{\theta_{-i}\in\operatorname{supp}\pi_{-i}}\beta_{-i}((h,a_{-i}),\theta_{-i})\pi_{-i}(\theta_{-i})}{\sum_{\theta'_{-i}\in\operatorname{supp}\pi_{-i}}\beta_{-i}(h,\theta'_{-i})\pi_{-i}(\theta'_{-i})} = \hat{\mu}_i(h)(a_{-i}),$$

with the equality following from strong consistency.

If instead $\beta_{-i}(h,\hat{\mu}_i(\emptyset)) = 0$, then h can only be reached by strategies in $\hat{S}_{-i}$ or by s_i if $\pi_{-i}(*_{-i}) = 0$. However, since the highest order of $\frac{1}{n}$ for $\hat{\rho}_{-i}^n$ among all pure strategies $\hat{s}_{-i} \in \hat{S}_{-i}$ is M_{-i}^n and $\pi_{-i}^n(*_{-i}) = n^{-(M_{-i}^n+1)}$ when $\pi_{-i}(*_{-i}) = 0$, the limit of the expression in Equation (5) is the same as

$$\begin{aligned}\lim_{n\to\infty}\frac{\sum_{\theta_{-i}\in\hat{S}_{-i}}\beta_{-i}((h,a_{-i}),\theta_{-i})\hat{\rho}_{-i}^n(\theta_{-i})}{\sum_{\theta_{-i}\in\hat{S}_{-i}}\beta_{-i}(h,\theta_{-i})\hat{\rho}_{-i}^n(\theta_{-i})} &= \lim_{n\to\infty}\frac{\beta_{-i}((h,a_{-i}),\hat{s}_{-i}^n)}{\beta_{-i}(h,\hat{s}_{-i}^n)} \\ &= \lim_{n\to\infty}\hat{s}_{-i}^n(h)(a_{-i}) \\ &= \hat{s}_{-i}(h)(a_{-i}),\end{aligned}$$

where the first equality follows from the realization equivalence between $\hat{\rho}_{-i}^n$ and $\hat{s}_{-i}^n$. Recall that $\hat{s}_{-i}$ is constructed to match $\hat{\mu}_i$ at every private histories in H_i^*. Thus $\hat{\mu}$ and μ are compatible. □

References

Abreu, D., D. G. Pearce, and E. Stacchetti (1990): "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica*, 58(5), 1041–1063.

Al-Najjar, N. (1995): "A Theory of Forward Induction in Finitely Repeated Games," *Theory and Decision*, 38, 173–193.

Aumann, R. J., and S. Sorin (1989): "Cooperation and Bounded Recall," *Games and Economic Behavior*, 1, 5–39.

Banks, J. S., and J. Sobel (1987): "Equilibrium Selection in Signaling Games," *Econometrica*, 55(3), 647–661.

Battigalli, P., and M. Siniscalchi (2002): "Strong Belief and Forward Induction Reasoning," *Journal of Economic Theory*, 106(2), 356–391.

——— (2003): "Rationalization and Incomplete Information," *Advances in Theoretical Economics*, 3(1), Article 3.

Ben-Porath, E., and E. Dekel (1992): "Signaling Future Actions and the Potential for Sacrifice," *Journal of Economic Theory*, 57, 36–51.

Bernheim, B. D. (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52(4), 1007–1028.

Chan, J. (2000): "On the Non-Existence of Reputation Effects in Two-Person Infinitely-Repeated Games," Working Paper, Johns Hopkins University.

Cho, I.-K., and D. M. Kreps (1987): "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102(2), 179–221.

Cripps, M. W., E. Dekel, and W. Pesendorfer (2005): "Reputation with Equal Discounting in Repeated Games with Strictly Conflicting Interests," *Journal of Economic Theory*, 121(2), 259–272.

Cripps, M. W., and J. P. Thomas (1997): "Reputation and Perfection in Repeated Common Interest Games," *Games and Economic Behavior*, 18(2), 141–158.

Fudenberg, D., and D. K. Levine (1989): "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, 57(4), 759–778.

——— (1993): "Self-Confirming Equilibrium," *Econometrica*, 61(3), 523–545.

Fudenberg, D., and E. Maskin (1991): "On the Dispensability of Public Randomization in Discounted Repeated Games," *Journal of Economic Theory*, 53(2), 428–438.

Govindan, S., and A. J. Robson (1998): "Forward Induction, Public Randomization and Admissibility," *Journal of Economic Theory*, 82, 451–457.

Govindan, S., and R. Wilson (2009): "On Forward Induction," *Econometrica*, 77(1), 1–28.

Gul, F., and D. G. Pearce (1996): "Forward Induction and Public Randomization," *Journal of Economic Theory*, 70, 43–64.

Kohlberg, E., and J.-F. Mertens (1986): "On the Strategic Stability of Equilibria," *Econometrica*, 54(5), 1003–1038.

Kuhn, H. W. (1953): "Extensive Games and the Problem of Perfection," in *Contributions to the Theory of Games*, ed. by H. W. Kuhn, and A. W. Tucker, vol. 2, pp. 193–216. Princeton University Press, Princeton, New Jersey.

Mailath, G. J., and L. Samuelson (2006): *Repeated Games and Reputations: Long-Run Relationships.* Oxford University Press, New York.

Man, P. (2010): "Forward Induction Equilibrium," Working Paper, University of Queensland.

Mertens, J.-F., S. Sorin, and S. Zamir (2000): "Repeated Games," CORE Discussion Paper, Université Catholique de Louvain.

Osborne, M. J. (1990): "Signaling, Forward Induction, and Stability in Finitely Repeated Games," *Journal of Economic Theory*, 50, 22–36.

Pearce, D. G. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52(4), 1029–1050.

Schmidt, K. M. (1993): "Reputation and Equilibrium Characterization in Repeated Games with Conflicting Interests," *Econometrica*, 61(2), 325–351.

School of Economics, University of Queensland
E-mail address: t.man@uq.edu.au