Order of Limits in Reputations*

Nuh Aygün Dalkıran†

October 30, 2015

Abstract

The fact that small departures from complete information might have large effects on the set of equilibrium payoffs draws interest in the adverse selection approach to study reputations in repeated games. It is well known that these large effects on the set of equilibrium payoffs rely on long-run players being arbitrarily patient. We study reputation games where a long-run player plays a fixed stage-game against an infinite sequence of short-run players under imperfect public monitoring. We show that in such games, introducing arbitrarily small incomplete information does not open the possibility of new equilibrium payoffs far from the complete information equilibrium payoff set. This holds true no matter how patient the long-run player is, as long as her discount factor is fixed. This result highlights the fact that the aforementioned large effects arise due to an order of limits argument, as anticipated.

JEL classification: C73, D82

Keywords: Reputations, Repeated Games with Short-Run and Long-Run Players, Continuity, Order of Limits.

*I am indebted to my advisors Mehmet Ekmekci and Ehud Kalai, as well as my committee members Nabil Al-Najjar and Asher Wolinsky for their guidance throughout this project. I would like to thank Alp Atakan, Johannes Hörner, Larry Samuelson, and Serdar Yüksel for their encouragement and helpful comments. I would like to also thank the participants of my talk at the annual meeting of the Society for Economic Dynamics (SED 2013, Seoul) and the seminar participants at Sabanci University for useful comments. This paper was one of the winners of 2014 Hakan Orbay Research Award, which was awarded by Sabanci University School of Management. An earlier version was circulated as "Limits of Reputations." Any remaining errors are mine.

†Northwestern University, Kellogg MEDS and Bilkent University, Department of Economics; email: dalkiran@bilkent.edu.tr; phone: +90 312 290 2006.

"Much of the interest in reputation models stems from the fact that seemingly quite small departures from perfect information about types can have large effects on the set of equilibrium payoffs." *Repeated Games and Reputations*, Mailath and Samuelson (2006), p 460.

---

1 Introduction

One of the most prominent results in the reputations literature is due to Fudenberg and Levine (1989, 1992), who studied infinitely repeated, reputation games where a long-run player faces an infinite sequence of short-run players. They showed that an *arbitrarily patient* strategic long-run player can guarantee herself a payoff close to her Stackelberg payoff when there is a small ex-ante probability that the long-run player is a commitment type who always plays the Stackelberg action. Their result implies that quite small perturbations of the complete information model might have large effects on the set of *limit* equilibrium payoffs.

This paper studies the set of equilibrium payoffs in repeated games with incomplete information as in Fudenberg and Levine (1992), when the long-lived player's discount factor is fixed. We show that even when the discount factor of the long-run player is very high, arbitrarily small perturbations cannot open the possibility of equilibrium payoffs far from the complete information equilibrium payoff set – as long as the discount factor of the long-run player is fixed.

Our main result might seem in stark contrast with the opening quotation of this paper, yet, it is indeed complementary to Fudenberg and Levine's (1992) result. Our main result highlights that, as anticipated, Fudenberg and Levine's (1992) result holds true due to a specific order of limits. That is, if the discount factor of the long-run player tends to 1 while holding the commitment type's ex-ante probability fixed, then the aforementioned reputation result à la Fudenberg and Levine

(1992) holds true; however, if the commitment type's ex-ante probability tends to 0 while holding the discount factor of the long-run player fixed, then the incomplete information equilibrium payoffs cannot be far from the complete information equilibrium payoff set. As far as we know, this is the first paper that explicitly points out the importance of the order of limits issue in these results.

From a technical point of view, our main result is an upper-hemi continuity result. We show that in reputation games of this type, the equilibrium payoff set is, for a fixed discount factor, upper-hemi continuous in the prior probability that the long-run player is a commitment type at zero when there is full-support imperfect public monitoring. We are aware that upper-hemi continuity results in the game theory literature are plenty. Yet, our result is the first result that explicitly provides a proof for the current upper-hemi continuity property, which highlights the order of limits issue in the reputations literature. Furthermore, given Bayesian updating, sequential rationality, and the dynamic structure of reputation games, our result is not a straightforward generalization of any such result in the literature. Other techniques might, of course, be used to prove similar results, yet our method of proof is relatively novel, employing the Abreu, Pearce, and Stacchetti (1990) techniques. We believe that this is another technical contribution of this paper because using such techniques to study repeated games with incomplete information is rare.[1] It is our hope that our proof will inspire other researchers to use similar techniques to tackle similar problems in the literature.

While Cripps and Thomas (2003) established both upper-hemi continuity and lower-hemi continuity of the equilibrium payoff set of repeated games with two long-lived players with equal discount factors when one-sided incomplete information vanishes, their results do not extend to our setting. Unfortunately, we fail to provide a proof (or a counter example) for the lower-hemi continuity counterpart of our result, thus this stays as a hard open problem. An affirmative conjecture for a necessary condition of the lower-hemi continuity counterpart of our result was done by Cripps, Mailath, and Samuelson (2004), but they also failed to provide a proof or a counter example for this necessary condition.[2] [3]

[1] The only other such paper we know of is Peski (2008).

[2] There are well-known examples of lack of lower hemi-continuity in dynamic games with asymmetric information. See for example, Section 14.4.1 of Fudenberg and Tirole (1991).

[3] Cripps, Mailath, and Samuelson (2004) conjecture the following *affirmative* hypothesis in their paper, which appears as a presumption for their Theorem 3: There exists a particular equilibrium in the complete information

1.1 Related Literature

The first papers that introduced the adverse selection approach to study reputations are Kreps, Milgrom, Roberts, and Wilson (1982), Kreps and Wilson (1982), and Milgrom and Roberts (1982). They show that the intuitive expectation of cooperation in early rounds of the finitely repeated prisoners' dilemma and entry deterrence in early rounds of the chain store game can be rationalized due to "reputation effects."

As mentioned above, Fudenberg and Levine (1989, 1992) extended this idea to infinitely repeated games and showed that a patient long-run player facing infinitely many short-run players can guarantee herself a payoff close to her Stackelberg payoff when there is a slight probability that the long-run player is a commitment type who always plays the Stackelberg action. When compared to the folk theorem (see Fudenberg and Maskin (1986) and Fudenberg, Levine, and Maskin (1994)), their results imply another intuitive expectation: the equilibria with relatively high payoffs are more likely to arise due to reputation effects.

Fudenberg, Kreps, and Maskin (1990) provided an upper bound for the equilibrium payoffs of a long-run player facing infinitely many short-run players under imperfect public monitoring, which is independent of the discount factor of the long-run player and might be strictly less than the Stackelberg payoff. Hence, Fudenberg and Levine (1992)'s results imply that under imperfect public monitoring new equilibrium payoffs may arise with incomplete information when the discount factor of the long-run player is sufficiently high.

---

game and a bound such that for *any* commitment type prior that is less than this bound, there exists an equilibrium of the incomplete information game where the long-run player's payoff is arbitrarily close to her payoff from this particular equilibrium of the complete information game. This is not exactly the lower-hemi continuity counterpart of our main result, but it is a necessary condition for the lower-hemi continuity counterpart of our main result. In a footnote, Cripps, Mailath, and Samuelson (2004) writes: "We conjecture this hypothesis is redundant, given the other conditions of the theorem, but have not been able to prove it." Unfortunately, we also fail to provide such a proof. Yet, an immediate corollary to our main result implies that one can identify a particular equilibrium in the complete information game and a sequence of priors converging to zero such that each incomplete information game with those priors has an equilibrium with the long-run player payoff arbitrarily close to the payoff from the particular equilibrium of the complete information game. Clearly, in the very special case when the complete information equilibrium payoff set of the long-run player is a singleton, our result implies that the continuity hypothesis conjecture of Cripps, Mailath, and Samulson (2004) is true.

Even though the results of Fudenberg and Levine (1989, 1992) hold for both perfect and imperfect public monitoring, Cripps, Mailath, and Samuelson (2004) showed that reputation effects are not sustainable in the long-run when there is imperfect public monitoring. In other words, it is impossible to maintain a permanent reputation for playing a strategy that does not play an equilibrium of the complete information game under imperfect public monitoring.

Since Cripps, Mailath, and Samuelson's (2004) work, there has been a large literature which studies the possibility / impossibility of maintaining permanent reputations: Ekmekci (2011) showed that reputation can be sustained permanently in the steady state by using rating systems. Ekmekci, Gossner, and Wilson (2012) showed that impermanent types would lead to permanent reputations, as well. Atakan and Ekmekci (2012, 2013, 2014) provided positive and negative results on permanent reputations with long-lived players on both sides. Liu (2011) provided dynamics that explain accumulation, consumption, and restoration of reputation when the discovery of the past is costly. Liu and Skyrpacz (2014) provided similar dynamics for reputations when there is limited record-keeping.

To sum up, the adverse selection approach to study reputations in repeated games has been quite fruitful. This approach teaches us that reputational concerns can explain the emergence of intuitive equilibria in both finitely and infinitely repeated games. There has been considerable amount of work in the literature which focus on whether or not it is possible to maintain a permanent reputation and how reputation is accumulated, consumed, and restored.

The next section describes our model. Section 3 provides a motivating example, Section 4 presents our main result, and Section 5 concludes the paper.

2 The Model

Our model is a standard model of an infinitely repeated game with incomplete information under imperfect public monitoring.[4]

[4] The notation we employ is similar to that of Cripps, Mailath, and Samuelson (2004). Hence, we refer the interested reader toCripps, Mailath, and Samuelson (2004) for further technical details of the model. We also refer the reader to Chapter 2 of Mailath and Samuelson (2006) for definitions of basic concepts which are skipped here.

2.1 The Complete Information game

A long-run player (Player 1) plays an infinitely repeated stage-game with a sequence of different short-run players (Player 2). The stage-game is a finite simultaneous-move game of imperfect public monitoring. The action sets of Player 1 and Player 2 in the stage-game are denoted by I and J, respectively. The public signal, y, is drawn from a finite set, Y. The probability that y is realized under the action profile (i,j) is given by ρ^y_{ij}.

The ex ante stage-game payoffs are given by $u_1(i,j)$ and $u_2(i,j)$.

Player 1 ("she") is a long-run player with a fixed discount factor $\delta < 1$. Her payoff in the infinitely repeated game is the average discounted sum of stage-game payoffs, $(1-\delta)\sum_{t=0}^{\infty}\delta^t u_1(i_t, j_t)$. Player 2 ("he"), on the other hand, denotes a sequence of short-run players, each of whom plays the stage-game only once.

Player 1's actions are private. Hence, Player 1 in period t has a private history, consisting of public signals and her past actions, denoted by $h_{1t} \equiv ((i_0,y_0),(i_1,y_1),\ldots,(i_{t-1},y_{t-1})) \in H_{1t} \equiv (I \times Y)^t$. Player 2, the short-lived players, only observes the public history, i.e., $(y_0, y_1, \ldots, y_{t-1}) \in Y^t$.

A behavioral strategy for Player 1 is denoted by $\sigma_1 : \bigcup\limits_{t=0}^{\infty} H_{1t} \rightarrow \Delta(I)$ whereas a behavioral strategy for Player 2 is denoted by $\sigma_2 : \bigcup\limits_{t=0}^{\infty} H_t \rightarrow \Delta(J)$. A strategy profile $\sigma = (\sigma_1, \sigma_2)$ induces a probability distribution P^σ over $(I \times J \times Y)^\infty$. Let $\{\mathcal{H}_{1t}\}_{t=0}^{\infty}$ denote the filtration on $(I \times J \times Y)^\infty$ induced by the private histories of Player 1 and $\{\mathcal{H}_t\}_{t=0}^{\infty}$ denote the filtration induced by the public histories. $E^\sigma[\cdot|\mathcal{H}_{it}\ ]$ denotes Player i's expectations with respect to P^σ conditional on $\mathcal{H}_{it}$, where $\mathcal{H}_{2t} = \mathcal{H}_t$.

In equilibrium, the short-run player plays a best-response after every equilibrium history. Player 2's strategy σ_2 is a best-response to σ_1 if, for all t,

$$E^\sigma[u_2(i_t, j_t)|\mathcal{H}_t] \geq E^\sigma[u_2(i_t, j)|\mathcal{H}_t] \text{ for all } j \in J$$

The set of such best-responses are denoted by $BR_2(\sigma_1)$.

We continue with the definition of a Nash equilibrium in the complete information game:

Definition 1 *A Nash equilibrium of the complete-information game is a strategy profile* $\sigma^* = (\sigma_1^*, \sigma_2^*)$ *with* $\sigma_2^* \in BR_2(\sigma_1^*)$ *such that for all* σ_1

$$E^{\sigma^*}[(1-\delta)\sum_{t=0}^{\infty}\delta^t u_1(i_t,j_t)] \geq E^{(\sigma_1,\sigma_2^*)}[(1-\delta)\sum_{t=0}^{\infty}\delta^t u_1(i_t,j_t)].$$

We assume that the monitoring structure has full support. That is, every signal y is possible after any action profile.

Assumption 1 *(Full Support): $\rho_{ij}^y > 0$ for all $(i,j) \in I \times J$ and $y \in Y$.*

Remark 1 *The full support monitoring assumption ensures that all finite sequences of public signals occur with positive probability, hence must be followed by optimal behavior in any Nash equilibrium. Therefore, any Nash equilibrium outcome is also a perfect Bayesian equilibrium outcome. Furthermore, since there is only one long-run and one short-run player, Nash equilibrium outcomes coincide with perfect public equilibrium outcomes.*[5]

2.2 The Incomplete Information Game

There is incomplete information regarding the type of the long-run player 1. At time $t = -1$, Player 1's type is selected. With probability $1 - p_0 > 0$, Player 1 ("she") is a "(n)ormal" type long-run player with a fixed discount factor $\delta < 1$. Her payoff in the infinitely repeated game is the average discounted sum of stage-game payoffs, $(1-\delta)\sum_{t=0}^{\infty}\delta^t u_1(i_t,j_t)$. With probability $p_0 > 0$, she is a "(c)ommitment" type who always, independent of history, plays the same (possibly mixed) action $s_1 \in \Delta(I)$ in each period.[6]

A state of the world is now a type for Player 1 and sequence of actions and signals. The set of states is $\Omega = \{n, c\} \times (I \times J \times Y)^\infty$ with generic outcome w. The prior p_0, commitment strategy, and the strategy profile of the normal players $\tilde{\sigma} = (\tilde{\sigma}_1, \sigma_2)$ induce a probability measure P over Ω, which describes how an uninformed player expects play to evolve.[7]

The strategy profile $\tilde{\sigma} = (\tilde{\sigma}_1, \sigma_2)$ determines a probability measure $\widetilde{P}$ over Ω, which describes how play evolves when Player 1 is the normal type. Let $E[\cdot]$ denote unconditional expectations taken

[5] For technical details we refer the interested reader to Kandori and Matsushima (1998, Appendix) or Sekiguchi (1997, Proposition 3).

[6] $\Delta(I)$ denotes the set of all possible probability distributions over I.

[7] The filtrations $\{\mathcal{H}_{1t}\}_{t=0}^{\infty}$ and $\{\mathcal{H}_t\}_{t=0}^{\infty}$ on $(I \times J \times Y)^\infty$ can also be viewed as filtrations on Ω in the obvious way.

with respect to the measure P and $\widetilde{E}[\cdot]$ denote the conditional expectations taken with respect to the measure $\widetilde{P}$.[8]

Given the strategy σ_2, the normal type Player 1 has the same objective function as in the complete-information game. Player 2, on the other hand, is maximizing $E[u_2(i_t, j)|\mathcal{H}_t]$, so that after any history h_t, he is updating his beliefs over the type of Player 1 that he is facing. The profile $(\tilde{\sigma}_1, \sigma_2)$ is a Nash equilibrium of the incomplete-information game if each player is playing a best-response. At any equilibrium, player 2's posterior belief in period t that player 1 is the commitment type is given by the $\mathcal{H}_t$-measurable random variable $p_t : \Omega \to [0, 1]$. By Assumption 1, Bayes' rule determines this posterior after all sequences of signals. Thus, in period t, player 2 is maximizing

$$p_t u_2(s_1, j) + (1 - p_t)\widetilde{E}[u_2(i_t, j)|\mathcal{H}_t].$$

The reputation of Player 1 is modeled as the belief of short-lived Player 2s regarding Player 1's type. Hence, in period t, it is quantified as Player2's posterior belief p_t.

Let $V(p_0, \delta)$ denote the equilibrium payoff set of the (normal type) long-run player 1 when the ex-ante commitment prior is p_0 and her discount factor is δ. In particular, $V(0, \delta)$ denotes the equilibrium payoff set of the long-run player 1 with discount factor δ in the repeated game under *complete information.*

3 A Motivating Example

To motivate our main result and to show how it compares to Fudenberg and Levine (1992), we provide an example of a moral-hazard mixing game (see Fudenberg and Levine (1994)). There is a long-lived seller (Player 1) who faces an infinite sequence of buyers (Player 2) who only plays the stage-game once. There are two actions available to the seller:, $A_1 = \{H, L\}$, where H and L denote producing a *high-quality* and a *low-quality* product respectively. Each buyer also has two possible actions: *buying* the product (B) and *not buying* the product (N), i.e., $A_2 = \{B, N\}$. Player 1 (the

[8] Note that $\sigma_1 : \bigcup_{t=0}^{\infty} H_{1t} \to \Delta(I)$ can be viewed as the sequence of functions $(\sigma_{10}, \sigma_{11}, \ldots, \sigma_{1t}, \ldots)$ with $\sigma_{1t} : H_{1t} \to \Delta(I)$. Hence, it can be extended from H_{1t} to Ω so that $\sigma_{1t}(w) \equiv \sigma_{1t}(h_{1t}(w))$, where $h_{1t}(w)$ is Player 1's t-period history under w. The same applies to σ_2 as well.

seller) is denoted as the row player, and Player 2 (each buyer) is denoted as the column player in the stage-game, with the following payoff matrix:

	B	N
H	$1,2$	$-1,0$
L	$2,-2$	$0,0$

Note that there is a *unique* Nash equilibrium of this stage-game, and in this equilibrium the row player plays L (producing *low* quality) and the column player plays N (*not buying* the product).[9] Note also that a rational buyer (Player 2) would play B only if he anticipates that the seller (Player 1) plays H with a probability of at least $\frac{1}{2}$.

Player 1's discount factor is $\delta < 1$. The actions of Player 1 are not observed by the buyers. However, every period an informative public signal about player 1's action is observed. Let $Y = \{h, l\}$ be the set of signals. Let $q > \frac{2}{3}$ be the probability of h occurring if Player 1 plays action H, and again for simplicity, q is the probability of l occurring if Player 1 plays L.[10][11] There is a positive probability $p_0 > 0$ that the seller is an honorable firm (a *commitment type*) who always produces *high-quality* product, that is, she plays action H at every period of the repeated game independent of the history.

By employing techniques introduced by Abreu, Pearce, and Stacchetti (1990), it can be shown that in the repeated game with complete information, i.e., when there is *no* commitment type, any $v \in [0, 1 - \frac{1-q}{2q-1}]$ is a subgame perfect equilibrium payoff of Player 1 for any $\delta > \overline{\delta}$ for some $\overline{\delta} < 1$, and no value outside of this range is an equilibrium payoff of Player 1 for any discount factor. That is, $V(0, \delta) = [0, 1 - \frac{1-q}{2q-1}]$ for any $\delta > \overline{\delta}$.[12]

Below, we show that for any $\delta > \overline{\delta}$ there exists an $\eta > 0$ such that any perfect Bayesian Nash equilibrium payoff of the long-lived seller (Player 1) in the incomplete information game is close to the set of her subgame perfect equilibrium payoffs of the complete information game when $p_0 < \eta$.

[9] Notice also that the unique Nash equilibrium of the stage game is not efficient.

[10] Note that the signals are independent of Player 2's choice of action here.

[11] It can be shown that when $q \leq \frac{2}{3}$ the complete information equilibrium payoff set is the singleton $\{0\}$.

[12] For the details of this argument we refer the reader to Mailath and Samuelson (2006, section 3.6).

Claim 1 *For any (fixed) $\delta > \bar{\delta}$, given any $\zeta > 0$, there exists an $\eta > 0$ such that if $p_0 < \eta$ then Player 1's any equilibrium payoff is not more than $1 - \frac{1-q}{2q-1} + \zeta$.*[13]

Proof. Suppose the probability that Player 1 is the commitment type is p_0 and she is expected to play H with probability $\alpha > p_0$.[14] Let $p_1(y)$ be the posterior probability that Player 1 is a commitment type after a signal y is observed. Bayes' rule yields:

$$
\begin{aligned}
p_1(h) &= \frac{p_0 q}{\alpha q + (1-\alpha)(1-q)} \\
p_1(l) &= \frac{p_0(1-q)}{\alpha(1-q) + (1-\alpha)q}
\end{aligned}
$$

Let $\kappa = \frac{q}{1-q}$ and observe that $\max_{\alpha \in [p_0,1]} \max\{\frac{p_1(h)}{p_0}, \frac{p_1(l)}{p_0}\} < \kappa$ since $q > \frac{2}{3}$.
For a given $\zeta > 0$, let t^* be such that $\delta^{t^*} < \frac{\zeta}{2}$, and let $\eta = \frac{\frac{1}{2}}{\kappa^{t^*}}$, and $p_0 < \eta$.
We have two observations that are true at any period $t < t^*$:

(i) At any period t after any history, the posterior probability with which Player 1 is a commitment type is less than $\frac{1}{2}$. This is simply because for any $t \leq t^*$ the posterior will be at most κ times the prior and $p_0 < \eta = \frac{\frac{1}{2}}{\kappa^{t^*}}$.

(ii) In any equilibrium, if Player 2's action is B after some public history, h_t, then H should be in the support of Player 1's strategy at time t. This is because, as mentioned before, to induce Player 2 to play B, the overall probability assigned to H should be at least $\frac{1}{2}$ and by **(i)** the posterior at t that Player 1 is a commitment type is less than $\frac{1}{2}$.[15]

Let

$$v' = \sup\{v \in V(\mu, \delta) \text{ for some } \mu \leq p_0\kappa\},$$

where $V(\mu, \delta)$ is the set of equilibrium payoffs of Player 1 for the commitment prior probability μ. Hence, v' is an upper bound for the continuation payoffs for Player 1 in the incomplete information game when $t \leq t^*$.

[13] It is possible to replicate Claim 1 for $\delta \leq \bar{\delta}$ but this is omitted since it adds no further insight to the main result.
[14] Note that here we refer to Player 1 considering both the normal and commitment types. That is, $\alpha = p_0 + (1 - p_0)\alpha_1(H)$, where $\alpha_1(H)$ is the probability with which the normal type of Player 1 chooses to play H.
[15] To be more precise, $\sigma_2(h_t) = B$ implies $H \in supp(\sigma_1(h_{1t}))$ for some h_{1t} compatible with h_t.

Suppose $p_0 < \eta$ and $q > \frac{2}{3}$; following **(i)** and **(ii)**, if Player 2's action is B at h_t, then H should be in the support of Player 1's action. Hence, Player 1's payoff is no more than

$$1(1-\delta) + \delta(qv_h + (1-q)v_l), \tag{1}$$

where v_h and v_l are the continuation payoffs for signals h and l respectively.

The incentive constraint that induces Player 1 to put a positive probability on H is given by

$$\delta(v_h - v_l) \geq (2q-1)(1-\delta). \tag{2}$$

Let $v(p_0, \delta)$ be any equilibrium payoff of Player 1 in the incomplete information game where the commitment prior is p_0 and her discount factor is δ.

Combining (1) and (2) with the fact that $v_h \leq v'$ and $v_l \leq v'$ gives us

$$v(p_0, \delta) \leq (1-\delta)(1 - \frac{1-q}{2q-1}) + \delta v'. \tag{3}$$

On the other hand, if H is not in the support of Player 1's action,

$$v(p_0, \delta) \leq (1-\delta)0 + \delta v' \leq (1-\delta)(1 - \frac{1-q}{2q-1}) + \delta v'. \tag{4}$$

An interpretation of inequality (3) is as follows: even though playing H gives Player 1 a current payoff of 1 when player 2 plays B, she bears an informational current payoff loss of $\frac{1-q}{2q-1}$ due to imperfect monitoring. Iterating forward gives:

$$v(p_0, \delta) \leq \sum_{s=0}^{t^*}(1-\delta)\delta^s(1 - \frac{1-q}{2q-1}) + \delta^{t^*+1} \sup_{\mu \in [0,1]} v(\mu, \delta). \tag{5}$$

Since $v(\mu, \delta) \leq 2$ for all $\mu \in [0,1]$ – since 2 is the highest payoff that Player 1 can get in the stage game– and since $\delta^{t^*} < \frac{\zeta}{2}$, we get

$$v(p_0, \delta) \leq 1 - \frac{1-q}{2q-1} + \zeta. \tag{6}$$

That is, whenever $p_0 \leq \eta$ we have $v(p_0, \delta) \leq 1 - \frac{1-q}{2q-1} + \zeta$. ■

Since Player 1's equilibrium payoff is bounded below by 0, this means any equilibrium payoff of the incomplete information game is close to the set of equilibrium payoffs of the complete information game when $p_0 < \eta$ as claimed. Technically, we had shown for a fixed $\delta > \overline{\delta}$, given any ζ there exists an η such that when $p_0 < \eta$, $V(p_0, \delta)$ is in the ζ neighborhood of $V(0, \delta)$.

In the incomplete information game, for every t^* we can choose a prior p_0 small enough such that for every $t \leq t^*$ Player 1's reputation level is less than $\frac{1}{2}$, *irrespective* of her strategy.[16] Hence, at any of these periods inducing Player 2 to play B bears the same cost, $\frac{1-q}{2q-1}$, on Player 1. For a fixed discount factor δ, if t^* is large enough, payoffs after t^* have almost no effect on Player 1's average discounted payoff in the repeated game.

Next, let us note what the main result of Fudenberg and Levine (1992) implies for this example: Let $\underline{v}(p_0, \delta) = \inf V(p_0, \delta)$ and $\overline{v}(p_0, \delta) = \sup V(p_0, \delta)$ for some (fixed) $p_0 \in (0, 1)$.

Claim 2 (Fudenberg and Levine (1992)) $\lim_{\delta \to 1} \underline{v}(p_0, \delta) = 1$ *for any (fixed)* $p_0 \in (0, 1)$.

Proof. See Corollary 3.2 of Fudenberg and Levine (1992). ■

Therefore, Fudenberg and Levine (1992) implies that when the long-lived seller (Player 1) becomes *arbitrarily patient*, i.e., as $\delta \to 1$, she guarantees herself a payoff close to 1 as long as $p_0 > 0$. The intuition behind their result is that by mimicking the commitment type often enough, a strategic long-run player can make her short-run opponents believe that she is a commitment type with sufficiently high probability. This will induce the short-run players to best respond to the commitment action except for a finite number of periods. But, when δ tends to 1 this finite number of periods will not matter, hence a lower bound for the equilibrium payoff of the arbitrarily patient long-run player will be the payoff that he can get by publicly committing herself to the action of the commitment type.[17]

On the other hand, our Claim 1 implies that for every (fixed) $\delta > \overline{\delta}$ given any ζ there exists an η such that when $p_0 \in (0, \eta)$, $\overline{v}(p_0, \delta) \leq 1 - \frac{1-q}{2q-1} + \zeta$. The intuition behind Claim 1 is that no matter how high the discount factor δ is, as long as it is fixed, there will come a time period t^* such that the effect of periods after t^* on the average discounted sum of payoffs will be negligible. Therefore, if commitment prior is so small that it takes for the long-run Player 1 longer than t^* to convince Player 2s that he is the commitment type with sufficiently high probability (greater than $\frac{1}{2}$) then the incomplete information equilibrium payoffs seems not far away from the complete information payoffs.

[16] This is where the full-support imperfect monitoring assumption bites.

[17] Observe that in the motivating example, the action of the commitment type is H and if Player 2s know that Player 1 is committed to play H then their best response would be B which will induce a payoff of 1 to Player 1.

To note the difference numerically, let $q = \frac{3}{4}$, $\delta = 0.99$, and $\zeta = 0.001$; if p_0 is positive but less than the corresponding η, then even though $\lim_{\delta\to 1,} \underline{v}(p_0, \delta) = 1$ our main result implies that $\overline{v}(p_0, 0.99) \leq \frac{1}{2} + 0.001$.[18] That is, no matter how high the discount factor is, as long as it is fixed, *the largest equilibrium payoff* to the long-lived seller is less than (or equal to) 0.501 for *arbitrarily small* commitment priors. On the other hand, no matter how small the commitment prior is, as long as it is fixed, *the smallest equilibrium payoff* to the long-lived seller will converge to 1 for *arbitrarily large* discount factors. These two results together clarify the importance of the order of limits in the standard reputation result.

Formally, the role of order of limits in terms of upper and lower bounds on equilibrium payoffs for the motivating example can be summarized by the following corollary:[19]

Corollary 1 $\lim_{\delta\to 1} \limsup_{p_0\to 0} \overline{v}(p_0, \delta) < \lim_{p_0\to 0} \lim_{\delta\to 1} \underline{v}(p_0, \delta)$.

Proof. From Claim 1, it follows that given any $\zeta > 0$ there exits $\eta > 0$ such that whenever $p_0 \leq \eta$ we have $v(p_0, \delta) \leq 1 - \frac{1-q}{2q-1} + \zeta$ for any $\delta > \bar{\delta}$. Furthermore, $q > \frac{2}{3}$ implies $1 - \frac{1-q}{2q-1} < 1$, hence $v(p_0, \delta) < 1 + \zeta$ for any $p_0 \leq \eta$. Therefore, $\zeta > 0$ being arbitrarily small implies $\limsup_{p_0\to 0} \overline{v}(p_0, \delta) < 1$. This is true for any $\delta > \bar{\delta}$ which then implies $\lim_{\delta\to 1} \limsup_{p_0\to 0} \overline{v}(p_0, \delta) < 1$

By Claim 2, we have $\lim_{\delta\to 1} \underline{v}(p_0, \delta) = 1$ for any $p_0 \in (0, 1)$, hence it follows that $\lim_{p_0\to 0} \lim_{\delta\to 1} \underline{v}(p_0, \delta) = 1$. ■

4 Main Result

We are ready to provide our main result. We note once again that $V(p_0, \delta)$ denotes the equilibrium payoff set of the long-run player with the fixed discount factor δ in the incomplete information repeated game with the ex-ante commitment prior p_0, and $V(0, \delta)$ denotes the equilibrium payoff set of the long-run player with the fixed discount factor δ in the repeated game under *complete information.*

[18] For $q = \frac{2}{3}$, $\delta = 0.99$, and $\zeta = 0.001$, the corresponding η can be easily calculated as $\frac{1/2}{3^{757}}$.

[19] Note that there is no known algorithm yet to compute the exact incomplete information equilibrium payoff set $V(p_0, \delta)$. Hence, the order of limits result provided here is just about the lower bounds and upper bounds of equilibrium payoff sets. In the subsection 4.5, it will be further clarified why a general order of limits result cannot be obtained in the form $\lim_{\delta\to 1} \lim_{p_0\to 0} V(p_0, \delta) \neq \lim_{p_0\to 0} \lim_{\delta\to 1} V(p_0, \delta)$ technically.

Our main result is the following:

Theorem 1 *Suppose the monitoring distribution ρ satisfies Assumption 1. For any fixed $\delta < 1$, given any $\zeta > 0$, there exists an $\eta > 0$ such that for any prior $p_0 \in (0, \eta)$, any equilibrium payoff of the long-run player in the incomplete information repeated game with the commitment prior p_0, i.e., any $v \in V(p_0, \delta)$, is in the ζ neighborhood of $V(0, \delta)$.*

In words, our main result says that introducing arbitrarily small incomplete information does not open the possibility of new equilibrium payoffs that are far from the complete information equilibrium payoff set, even when the long-run player's discount factor is very high but fixed.

4.1 Outline of the Proof

We proceed as follows: first, we introduce the standard set operator à la Abreu, Pearce, and Stacchetti (1990) particular to our setting, which gives us decomposable payoffs for the long-run player in a given set. When applied repeatedly to a compact set that includes all stage-game payoffs of the long-run player, this operator converges to the complete information equilibrium payoff set of the long-run player. Then, we slightly modify this operator to introduce a new set operator. The modification is that Player 2, any of the short-run players, is not restricted to best-respond to the enforcing action of Player 1, but is allowed to best-respond to some (possibly mixed) action that is close (in the Euclidean metric) to the enforcing action.

In our first lemma we show that there exists a distance $\bar{\varepsilon} > 0$ such that all best-responses to this particular action are also best-responses to *some* (possibly mixed) action whose support is within the support of the action of the normal type of the long-run player. The essence of the argument is that the $\bar{\varepsilon}$ in Lemma 1 is uniform over all possible supports. In Lemma 2, we show that the two operators coincide for any distance smaller than $\bar{\varepsilon}$. The essence of Lemma 2 is that the operators coincide uniformly over all possible subsets of the real line. Lemma 3 extends the result of Lemma 2 to arbitrary iterations using these uniformities.

The rest of the proof makes use of the fact that the discount factor is fixed and hence a finite number of iterations, t^*, suffices to approximate the set of all complete information equilibrium payoffs of the long-run player. When the commitment prior in the incomplete information game is small enough, the posterior after the first t^* periods stays below a certain threshold due to full-support monitoring. Employing Lemma 3 allows us to show that one can identify a bound over the commitment prior so that the equilibrium payoffs of the long-run player in the incomplete information game with a prior less than this bound cannot be too far from her complete information equilibrium payoffs.

4.2 The Set Operator

We start by introducing the standard set operator, T, of Abreu, Pearce, and Stacchetti (1990):

Consider the operator $T : 2^{\mathbb{R}} \backslash \{\emptyset\} \to 2^{\mathbb{R}} \backslash \{\emptyset\}$ defined as follows:

$v \in T(W)$ if and only if there exists a nonempty $I_v \subseteq I$, $\alpha_2 \in \Delta(J)$ and $w = (w_1, w_2, \ldots, w_Y) \in W^{|Y|}$ such that:

$$(i)\ v = (1-\delta)u_1(i, \alpha_2) + \delta(\sum_{y,j} w_y \rho_{ij}^y \alpha_2(j)) \text{ for each } i \in I_v$$

$$(ii)\ v \geq (1-\delta)u_1(i, \alpha_2) + \delta(\sum_{y,j} w_y \rho_{ij}^y \alpha_2(j)) \text{ for each } i \in I$$

$$(iii)\ \alpha_2 \in BR_2(\alpha_1) \text{ for some } \alpha_1 \in \Delta(I_v)$$

This operator identifies decomposable payoffs of the long-run player for a given set, $W \in \mathbb{R}$. Here, (i) corresponds to feasibility, (ii) corresponds to incentive compatibility conditions, and (iii) simply says that Player 2 (a short run player) is best-responding to the enforcing action of the long-run player.

Let $M = \max_{i,j} |u_1(i,j)|$, and let $W_0 = [-M, M]$. It follows from the techniques introduced by Abreu, Pearce, and Stacchetti (1990) that $T^\infty(W_0) = \bigcap_{t=0}^{\infty} T^t(W_0) = V(0, \delta)$ where $T^t(W_0)$ is defined recursively as $T^1(W_0) = T(W_0)$; $T^k(W_0) = T(T^{k-1}(W_0))$ for all $k \in \mathbb{N}$.

Next, consider the incomplete information game: Recall that in period t Player 2 is best-responding to α_1', where $\alpha_1' = p_t s_1 + (1-p_t)\alpha_1$ and $\sigma_1(h_{1t}) = \alpha_1$, and p_t is the posterior at time t that Player 1 is a commitment type. Observe that when p_t happens to be arbitrarily small so is the Euclidean distance $||\alpha_1' - \alpha_1||$.[20]

Utilizing this observation, we next define our set operator by relaxing condition (iii) of the operator T as follows:

[20] $||\alpha_1' - \alpha_1|| = \sqrt{\sum_{i \in I}(\alpha_1'(i) - \alpha_1(i))^2}$.

For any $\varepsilon > 0$, let $T_\varepsilon : 2^{\mathbb{R}}\backslash\{\emptyset\} \to 2^{\mathbb{R}}\backslash\{\emptyset\}$ be such that:

$v \in T_\varepsilon(W)$ if and only if there exists a nonempty $I_v \subseteq I$, $\alpha_2 \in \Delta(J)$ and $w = (w_1, w_2, \ldots, w_Y) \in W^{|Y|}$ such that:

$$(i)\ v = (1-\delta)u_1(i, \alpha_2) + \delta(\sum_{y,j} w_y \rho_{ij}^y \alpha_2(j)) \text{ for each } i \in I_v$$

$$(ii)\ v \geq (1-\delta)u_1(i, \alpha_2) + \delta(\sum_{y,j} w_y \rho_{ij}^y \alpha_2(j)) \text{ for each } i \in I$$

$$(iii_\varepsilon)\ \alpha_2 \in BR_2(\alpha_1') \text{ for some } \alpha_1' : ||\alpha_1' - \tilde{\alpha}_1|| \leq \varepsilon \text{ for some } \tilde{\alpha}_1 \in \Delta(I_v)$$

T_ε is slightly more permissive than the T of Abreu, Pearce, and Stacchetti (1990), inasmuch as it only requires the short-run player to best-respond to an action that is close, in Euclidean metric, to the one played by the long-run player.

The motivation for our operator is as follows: In the game of incomplete information, the mixed action to which the short-run player best-responds is a weighted average of the action taken by the normal type of the long-run player and the commitment type's action. Provided that the latter type is very unlikely, this means that the short-run player is taking a best-response to an action that is nearly the normal type's action. The key to our main result will be then to show that, if the short-run player best-responds to an action that is close to the probability distributions over a set of actions, then he is actually also playing a best-response to an action within that set if the distance between the original action and the set of probability distributions over this set of actions is sufficiently small.

It is clear that both operators T and T_ε are monotone. That is, if $W_1 \subseteq W_2$, then $T(W_1) \subseteq T(W_2)$ and $T_\varepsilon(W_1) \subseteq T_\varepsilon(W_2)$. Moreover, for any $\varepsilon_1 > \varepsilon_2$, $T_{\varepsilon_2}(W) \subseteq T_{\varepsilon_1}(W)$.

4.3 Lemmata

We provide 3 lemmata, which will be used in the proof of our main result. All of the proofs of these lemmata are provided in the Appendix.

We start with a technical lemma that is key to our main result. In words, Lemma 1 says that if a short-run player best-responds to an action that is close to the set of probability distributions over a set of actions, then he is also playing a best-response to an action within this set of probability distributions as long as the Euclidean distance between the original action and the set of probability distributions is sufficiently small.

Let $BR_2(\Delta(X)) := \{\alpha_2 \in \Delta(J) : \alpha_2 \in BR_2(\alpha_1) \text{ for some } \alpha_1 \in \Delta(X)\}$ for any $X \subseteq I$.

Lemma 1 *There exists an $\bar{\varepsilon} > 0$ for all non-empty $X \subsetneq I$ such that $\min_{\sigma_X \in \Delta(X)} ||\alpha_1 - \sigma_X|| \in (0, \bar{\varepsilon})$, then $BR_2(\alpha_1) \subseteq BR_2(\Delta(X))$.*

It is essential to note about Lemma 1 that the $\bar{\varepsilon}$ is uniform over $X \subsetneq I$, i.e., a fixed $\bar{\varepsilon}$ works for all $X \subsetneq I$.

Next, we use Lemma 1 to show that for any $\varepsilon < \bar{\varepsilon}$ operators, T and T_ε coincide. Note that this means $\bar{\varepsilon}$ is uniform over all $W \subset \mathbb{R}$.

Lemma 2 *For any $0 < \varepsilon < \bar{\varepsilon}$ and $W \subset \mathbb{R}$, $T_\varepsilon(W) = T(W)$.*

Since $\bar{\varepsilon}$ is uniform over $W \subset \mathbb{R}$, any arbitrary number of iterations of T_ε and T will coincide for every $\varepsilon < \bar{\varepsilon}$. Our next lemma formalizes this fact for W_0. Recall that $W_0 = [-M, M]$, where $M = \max_{i,j} |u_1(i,j)|$. Let $T^t_\varepsilon(W)$ be defined recursively as $T^1_\varepsilon(W) = T_\varepsilon(W)$; $T^k_\varepsilon(W) = T_\varepsilon(T^{k-1}_\varepsilon(W))$ for all $k \in \mathbb{N}$.

Lemma 3 *For any $0 < \varepsilon < \bar{\varepsilon}$, $T^t_\varepsilon(W_0) = T^t(W_0)$.*

4.4 Proof of the Main Result

Now, we are ready to give the proof of our main result:

Proof of Theorem 1.

Define $d_w = \min_{v \in V(0,\delta)} |v - w|$.[21] Given any $\zeta > 0$, the fact $T^\infty(W_0) = V(0,\delta)$ implies that there must exist a t^* such that for $t > t^*$, $\max_{w \in T^t(W_0)} d_w < \zeta$.[22]

Let $\eta := \frac{\varepsilon}{\sqrt{2}\kappa^{t^*}}$ for some $\varepsilon < \bar{\varepsilon}$ of Lemma 1, and let

$$\kappa := \sup_{\alpha_1 \in \Delta(I), j \in J} \max_y \{ \frac{\sum_i \rho^y_{i,j} s_1(i)}{\sum_i \rho^y_{i,j} \alpha_1(i)} \}$$

[21] d_w is well defined since $V(0,\delta)$ is compact as shown by Theorem 4 of Abreu, Pearce, and Stacchetti (1990)] and the Euclidean distance $|\cdot|$ is continuous.

[22] Note here that $\max_{w \in T^t(W_0)} d_w$ is well defined as well since by Lemma 1 of Abreu, Pearce, and Stacchetti (1990) the operator T is monotone and preserves compactness and $W_0 = [-M, M]$ is compact.

Note first that by the full-support assumption (Assumption 1), $\kappa < \infty$.[23]

In the incomplete information game with commitment prior $p_0 \in (0, \eta)$, at any period $t \leq t^*$, the probability with which Player 1 is a commitment type is not more than $\frac{\varepsilon}{\sqrt{2}}$. To see why, observe that the posterior belief of Player 2 about the commitment type in any period can be at most κ times his prior from the preceding period and hence $p_t < p_0 \kappa^t < \frac{\varepsilon}{\sqrt{2}}$ for all $t \leq t^*$.

Therefore, the set of continuation payoffs at period t^* is a subset of $T_\varepsilon(W_0)$. This is because in any equilibrium of the incomplete information game when the normal type plays according to $\tilde{\sigma}_1$ with $\tilde{\sigma}_1(h_{1t^*}) = \alpha_1$, then at t^* Player 2 is best-responding to $\alpha_1' = p_{t^*} s_1 + (1 - p_{t^*})\alpha_1$, and since $p_{t^*} < \frac{\varepsilon}{\sqrt{2}}$, we have $||\alpha_1' - \alpha_1|| < \varepsilon$.[24] Similarly, at $t^* - 1$ the set of continuation payoffs is a subset of $T_\varepsilon^2(W_0)$. Thus, iterating backwards, the set of equilibrium payoffs at period 0, $V(p_0, \delta)$, is a subset of $T_\varepsilon^{t^*+1}(W_0)$.

By Lemma 3, $T_\varepsilon^{t^*+1}(W_0) = T^{t^*+1}(W_0)$. Therefore $\max_{w \in T_\varepsilon^{t^*+1}(W_0)} d_w < \zeta$. Hence, we have the incomplete information equilibrium payoff set of Player 1 (the long-run player) $V(p_0, \delta) \subset T_\varepsilon^{t^*+1}(W_0)$, and it follows that $V(p_0, \delta)$ is in the ζ neighborhood of $V(0, \delta)$. ■

4.5 Order of Limits

To clarify how our main result identifies the role of the order of limits in reputations, let $V : [0, 1)^2 \to 2^{\mathbb{R}}$ be the function which gives the equilibrium payoff set of long-run player 1 for any commitment prior, discount factor pair, (p_0, δ), where $2^{\mathbb{R}}$ denotes the power set of $\mathbb{R}$. That is, as before, $V(p_0, \delta)$ is the equilibrium payoff set of Player 1 when the commitment prior is p_0 and the discount factor of player 1 is δ.

Unfortunately, we cannot provide the following inequality:

$$\lim_{\delta \to 1} \lim_{p_0 \to 0} V(p_0, \delta) \neq \lim_{p_0 \to 0} \lim_{\delta \to 1} V(p_0, \delta) \tag{7}$$

The technical reason why one *cannot* provide inequality (7) is because there is no standard

[23] Note that this is where Assumption 1 (full-support monitoring) bites. Assumption 1 is crucial for our result not only because under Assumption 1 Nash equilibrium payoffs are the same as perfect public equilibrium payoffs, but also, without full-support monitoring we cannot bound κ. This is why our proof fails for the case of perfect monitoring as well.

[24] $||\alpha_1' - \alpha_1|| \leq \varepsilon$ since $||s_1 - \alpha_1|| \leq \sqrt{2}$.

topology (or metric) defined on the power set of $\mathbb{R}$ where the limits in inequality (7) are well-defined.

A commonly used metric for defining limits of sequences of sets is the Hausdorff metric, which is defined as follows

$$d(V, W) = \max\{\sup_{v \in V} \inf_{w \in W} d(v, w), \sup_{w \in W} \inf_{v \in V} d(v, w)\}$$

But, the Hausdorff metric is a metric *only for the compact subsets* of $\mathbb{R}$. Yet, we do not know whether $V(p_0, \delta)$ is compact for when p_0 is positive.[25]

However, whenever the Stackelberg payoff is not attainable in the stage game, it is possible to obtain the following order of limits result in terms of upper and lower bounds of the equilibrium payoff sets which will imply that the sets in question in inequality (7) are indeed not close according to the intuition behind the Hausdorff metric.[26]

Corollary 2 *If the Stackelberg payoff S is not a Nash equilibrium payoff of the stage game and the commitment type of Player 1 is associated with the Stackelberg action with corresponding Stackelberg payoff S then*

$$\lim_{\delta \to 1} \limsup_{p_0 \to 0} \overline{v}(p_0, \delta) \leq S \leq \liminf_{p_0 \to 0} \lim_{\delta \to 1} \underline{v}(p_0, \delta) \tag{8}$$

Proof. Proposition 3 of Fudenberg, Kreps and Maskin (1990) implies that any complete information payoff $v(0, \delta) < S$. Therefore by Theorem 1, we have $v(p_0, \delta) < S + \zeta$ whenever $p_0 < \eta$. Since this is true for any $\zeta > 0$ and for any arbitrarily small p_0 we obtain $\limsup_{p_0 \to 0} \overline{v}(p_0, \delta) \leq S$

The fact that $\liminf_{p_0 \to 0} \lim_{\delta \to 1} \underline{v}(p_0, \delta) \geq S$ follows from Corollary 3.2 of Fudenberg and Levine (1992) since it implies $\lim_{\delta \to 1} \underline{v}(p_0, \delta) \geq S$ for any $p_0 \in (0, 1)$. ∎

Corollary 2 implies that, according to the intuition behind the Hausdorff metric – two sets are close if every point of either set is close to some point of the other set – the corresponding limit equilibrium payoff sets in question above in inequality (7) are not close. An upper bound for the limit set on the left hand side of (7) is less than a lower bound of the limit set on the right hand side of (7) This means if the equilibrium payoff sets were all compact for any commitment prior and

[25] We know when $p_0 = 0$, $V(0, \delta)$ is compact by Theorem 4 of Abreu, Pearce, and Stachetti (1990). But, we do not have a similar result for the case of repeated games with incomplete information.

[26] Recall that Stackelberg payoff in the stage game is the highest payoff that Player 1 can get by publicly committing to a (possibly mixed) action. Formally, $S = \max_{\alpha_1 \in \Delta(I), \alpha_2 \in BR_2(\alpha_1)} u_1(\alpha_1, \alpha_2)$.

any discount factor then limits in inequality (7) will be well defined with respect to the Hausdorff metric and hence inequality (7) will hold true.

As discussed earlier, Fudenberg and Levine (1992)'s reputation result show that when the long-run player (Player 1) becomes arbitrarily patient ($\delta \to 1$) she guarantees herself a payoff close to her Stackelberg payoff as long as ex-ante probability of the Stackelberg (commitment) type is positive – no matter how small it is. The intuition behind their result is that by mimicking the Stackelberg type often enough, the long-run player can convince short-run players that she is the Stackelberg (commitment) type with sufficiently high probability. Hence, the short-run players will best respond to the Stackelberg action except for a finite number of periods. When $\delta \to 1$ this finite number of periods will not matter.

On the other hand, our main result implies that when the commitment prior is arbitrarily small ($p_0 \to 0$) any incomplete information payoff will be close to a complete information payoff – no matter how large the discount factor is. The intuiton behind our result is simple: When the discount factor δ is fixed, there will be a time period t^* such that the effect of periods after t^* on the average discounted sum of payoffs are negligible. Hence, for arbitrarily small commitment priors, it will take longer than t^* for the long-run player to convince short-run players that he is the Stackelberg type with sufficiently high probability – to induce them to best respond to the Stackelberg action. Therefore, the effect of introducing arbitrarily small incomplete information on the equilibrium payoffs will be negligible as well.

5 Conclusion

The main result of this paper is essentially an upper-hemi continuity result concerning the equilibrium payoffs in reputation games where a long-run player faces an infinite sequence of short-run players. Technically, we showed that in these games the Nash equilibrium correspondence is, for a fixed discount factor, upper-hemi continuous in the prior probability that the long-run player is a commitment type at zero when there is full-support imperfect public monitoring.

To the best of our knowledge, this is the first result that explicitly provides a proof for this particular upper-hemi continuity property, which highlights the order of limits issue in the reputations literature: Our result highlights that – as anticipated – the order of limits is important in some of the standard reputation results. If the discount factor of the long-run player tends to 1 before the commitment type's ex-ante probability tends to 0, then the aforementioned reputation results à la Fudenberg and Levine (1992) are expected to hold true; however, if the commitment type's ex-ante probability tends to 0 before the discount factor of the long-run player tends to 1, then the incomplete information equilibrium payoffs cannot be far from the complete information equilibrium payoff set.

Unfortunately, even though an affirmative necessary condition was conjectured by Cripps, Mailath, and Samuelson (2004), the corresponding lower-hemi continuity result remains to be an open problem for further research.

Our result may not hold when the full-support monitoring assumption is missing. For example, in the extreme case of perfect monitoring, it can be shown that the incomplete information equilibrium payoff of Player 1 can only differ from a complete information equilibrium payoff by one period payoff. But since we fix $\delta < 1$, this fact does not provide a counterpoint to the continuity result in this paper.

Finally, we believe that the method of proof we provide is a novel application of Abreu, Pearce, and Stacchetti's (1990) techniques in repeated games with incomplete information. We hope that our work will inspire other researchers to use these techniques to tackle similar problems in the literature.

6 Appendix

6.1 Proofs of Lemmata

Lemma 1 *There exists an $\bar{\varepsilon} > 0$ for all non-empty $X \subsetneq I$ such that $\min_{\sigma_X \in \Delta(X)} ||\alpha_1 - \sigma_X|| \in (0, \bar{\varepsilon})$, then $BR_2(\alpha_1) \subseteq BR_2(\Delta(X))$.*

Proof. Suppose this is not the case. Since $\Delta(I)$ is compact, it is possible to construct convergent sequences $\alpha_{1,n} \in (\Delta(I) \backslash \Delta(X))$ and $\alpha_{2,n} \in \Delta(J)$ such that $\lim_{n\to\infty} \min_{\sigma_X \in \Delta(X)} ||\alpha_{1,n} - \sigma_X|| = 0$ for some $X \subset I$, with $\alpha_{2,n} \in BR_2(\alpha_{1,n})$ and $\alpha_{2,n} \notin BR_2(\Delta(X))$ for all n. Let $\alpha_{1,n} \to \bar{\alpha}_1$. Observe that we must have $\bar{\alpha}_1 \in \Delta(X)$.

Let $2^J = \{Y_1, Y_2, \ldots Y_m\}$. We know that $BR_2(\Delta(X)) = \bigcup_{j\in H} \Delta(Y_j)$ for some $H \subset \{1, \ldots, m\}$. We must have for all n, $\alpha_{2,n} \notin \Delta(Y_j)$ for all $j \in H$. Therefore, $supp(\alpha_{2,n}) = Y_n$, where $Y_n \neq Y_j$ for any $j \in H$. But since the number of subsets of J is finite, there exists a subsequence Y_k of Y_n such that $Y_k = \bar{Y}$ for some $\bar{Y} \neq Y_j$ for any $j \in H$. This implies $supp(\alpha_{2,k}) = \bar{Y}$ for all k.

Consider the stage-game action α_2^* of Player 2 that gives equal probability to all the actions in $\bar{Y}$. It follows that α_2^* is a best-response to $\alpha_{1,k}$ for all k. But since the best-response correspondence has a closed graph, this implies that α_2^* is a best-response to $\bar{\alpha}_1 \in \Delta(X)$. This contradicts the fact that $\bar{Y} \neq Y_j$ for any $j \in H$. ■

Lemma 2 *For any* $0 < \varepsilon < \bar{\varepsilon}$ *and* $W \subset \mathbb{R}$, $T_\varepsilon(W) = T(W)$.

Proof. $T(W) \subseteq T_\varepsilon(W)$ is true for any $\varepsilon > 0$, and since $\varepsilon_1 > \varepsilon_2$ implies $T_{\varepsilon_2}(W) \subseteq T_{\varepsilon_1}(W)$, it is enough to show that $T_{\bar{\varepsilon}}(W) \subseteq T(W)$.

Suppose $v \in T_{\bar{\varepsilon}}(W)$. Therefore, there exists $I_v \subseteq I$, $\alpha_2 \in \Delta(J)$ and $w = (w_1, w_2, \ldots, w_Y) \in W^{|Y|}$ such that (i), (ii), and $(iii_{\bar{\varepsilon}})$ hold.

We need to identify an $\hat{I}_v \subseteq I$, $\hat{\alpha}_2 \in \Delta(J)$ and $\hat{w} = (\hat{w}_1, \hat{w}_2, \ldots, \hat{w}_Y) \in W^{|Y|}$ such that (i), (ii), and (iii) hold.

If $\alpha_1' \in \Delta(I_v)$, we are trivially done. (Simply take $\hat{I}_v = I_v$, $\hat{\alpha}_2 = \alpha_2$, $\hat{w} = w$ and $\alpha_1 = \alpha_1'$ in (iii).)

If not, $\alpha_1' \in \Delta(I) \backslash \Delta(I_v)$ with $\min_{\sigma\in\Delta(I_v)} ||\alpha_1' - \sigma|| < ||\alpha_1' - \tilde{\alpha}_1|| < \bar{\varepsilon}$ since $\tilde{\alpha}_1 \in \Delta(I_v)$. Hence, by Lemma 1, we must have $BR_2(\alpha_1') \subseteq BR_2(\Delta(I_v))$. So there exists α_1'' in $\Delta(I_v)$ with $\alpha_2 \in BR_2(\alpha_1'')$. Let $\bar{X}$ be the support of α_1''. Clearly, $\bar{X} \subseteq I_v$.

Now take $\hat{I}_v = \bar{X}$, $\hat{\alpha}_2 = \alpha_2$, and $\hat{w} = w$. Observe that (i) and (ii) hold since $\bar{X} \subseteq I_v$. The fact that (iii) holds is clear; simply take $\alpha_1 = \alpha_1''$. ■

Lemma 3 *For any* $0 < \varepsilon < \bar{\varepsilon}$, $T_\varepsilon^t(W_0) = T^t(W_0)$.

Proof. By Lemma 2, for any $0 < \varepsilon < \bar{\varepsilon}$, $T_\varepsilon(W_0) = T(W_0)$. Since the ε in Lemma 2 is independent of W, with the same ε, $T_\varepsilon^2(W_0) = T^2(W_0)$. Iterating gives $T_\varepsilon^t(W_0) \subseteq T^t(W_0)$. ■

References

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

Atakan, A., and M. Ekmekci (2012): "Reputation in Long-Run Relationships," *Review of Economic Studies*, 79(2), 451–480.

——— (2013): "A two-sided reputation result with long-run players," *Journal of Economic Theory*, 148(1), 376–392.

——— (2014): "Reputation in the Long-Run with Imperfect Monitoring," working paper.

Cripps, M., G. Mailath, and L. Samuelson (2004): "Imperfect Monitoring and Impermanent Reputations," *Econometrica*, 72(2), 407–432.

Cripps, M., and J. Thomas (2003): "Some Asymptotic Results in Discounted Repeated Games of One-Sided Incomplete Information," *Mathematics of Operations Research*, 28(3), 433–462.

Ekmekci, M. (2011): "Sustainable Reputations with Rating Systems," *Journal of Economic Theory*, 146(2), 479–503.

Ekmekci, M., O. Gossner, and A. Wilson (2012): "Impermanent Types and Permanent Reputation," *Journal of Economic Theory*, 147(1), 162–178.

Fudenberg, D., D. M. Kreps, and E. Maskin (1990): "Repeated Games with Long-Run and Short-Run Players," *Review of Economic Studies*, 57, 555–573.

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

——— (1992): "Maintaining a Reputation when Strategies are Imperfectly Observed," *Review of Economic Studies*, 59(3), 561–579.

——— (1994): "Efficiency and Observability with Long-Run and Short-Run Players," *Journal of Economic Theory*, 62(1), 103–135.

Fudenberg, D., D. Levine, and E. Maskin (1994): "The Folk Theorem with Imperfect Public Information," *Econometrica*, 62(5), 997–1039.

Fudenberg, D., and E. Maskin (1986): "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54(3), 533–554.

Fudenberg, D., and J. Tirole (1991): *Game Theory*. The MIT Press, Cambridge, Massachusetts.

Kandori, M., and H. Matsushima (1998): "Private Observation, Communication and Collusion," *Econometrica*, 66(3), 627–652.

Kreps, D. M., P. Milgrom, D. Roberts, and R. Wilson (1982): "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 27(2), 245–252.

Kreps, D. M., and R. Wilson (1982): "Reputation and Imperfect Information," *Journal of Economic Theory*, 27(2), 253–279.

Liu, Q. (2011): "Information Acquisition and Reputation Dynamics," *Review of Economic Studies*, 78(4), 1400–1425.

Liu, Q., and A. Skrzypacz (2014): "Limited Records and Reputation Bubbles," *Journal of Economic Theory*, 151, 2–29.

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

Milgrom, P., and D. Roberts (1982): "Predation, Reputation and Entry Deterrence," *Journal of Economic Theory*, 27(2), 280–312.

Peski, M. (2008): "Repeated Games with Incomplete Information on One Side," *Theoretical Economics*, 3(1), 29–84.

Sekiguchi, T. (1997): "Efficiency in Repeated Prisoners' Dilemma with Private Monitoring," *Journal of Economic Theory*, 76(2), 345–361.