Limited Records and Reputation*

Qingmin Liu† Andrzej Skrzypacz‡

October 27, 2009

Abstract

We study the impact of limited records on reputation dynamics, that is, how the set of equilibria and equilibrium payoffs changes in a model in which one long-lived player faces a sequence of short-lived players who observe only limited information about past play (the last K periods of the long-lived player's actions). We show that limited records dramatically change the equilibrium behavior. Moreover, with limited records, equilibria in games with complete and incomplete information are strikingly different (in contrast to games with complete records). We also obtain a lower bound for equilibrium payoffs at any moment of the game, not only at the beginning, thus providing a stronger long-run prediction.

1 Introduction

In this paper we study the economics of reputation in dynamic games with limited record-keeping, i.e., games in which new players observe only the last few periods of play, instead of the full history of the game. We characterize a class of games in which limited records play a major role. First, we demonstrate that without incomplete information there is no possibility of sustaining reputation and providing dynamic incentives, which is in great contrast to what can be achieved in equilibrium if the complete history of the game is observed. Second, with incomplete information, we show that equilibrium dynamics are very different under limited records than under complete records. Finally, we show that despite the difference in equilibrium behavior in games with incomplete information, similar payoff bounds can be established under limited record-keeping (if players are sufficiently patient and if the records are long enough). However,

---

*We thank Aaron Bodoh-Creed, Alex Frankel, Philippe Jehiel, George Mailath, Robert Wilson and seminar participants at Stanford, University of Maryland, University of Western Ontario, and University of Toronto for comments and feedback on this project.

†University of Pennsylvania, qingmin@econ.upenn.edu

‡Stanford GSB, andy@gsb.stanford.edu

unlike in games with complete records, our payoff bound applies not only at the beginning of the game, but at any moment of the game.

Our results reveal that unlike in the more traditional games with complete records, with finite records the two common ways of modeling reputation, namely as repeated-game effects or as beliefs about commitment types lead to very different sets of equilibria. Typically, if "reputation behavior/incentives" can be sustained in an equilibrium of a model with commitment types, it can be also sustained in an equilibrium of a model without types and vice-versa (although, see below for discussion of differences in finite horizon games). In our games with limited records, the set of equilibria is very different – there is a complete collapse of reputation effects in the game without commitment types and there are rich reputation dynamics with much higher payoffs in the game with commitment types. (The only previous result of that kind to our knowledge is in Faingold and Sannikov (2007) – they study reputation games in continuous time with imperfect monitoring via Brownian motions and show that rich dynamics can be achieved with commitment types but not without.)

We analyze dynamic games in which a long-lived player faces a sequence of short-lived players. In each period the players act simultaneously: the short-lived player decides how much to trust the long-lived player who in turn decides how much to exploit the trust. We characterize the perfect Bayesian equilibria (PBE) of the game under the assumption that a short-lived player entering the game sees only the last K periods of the long-lived player's actions (in the games with complete information we assume that the short-lived players see the calendar time, while for tractability in the games with incomplete information we assume that a short-lived player does not know how long the game has been played so far, but has a prior over the times at which he enters the game). We make several assumptions about the stage game payoffs. The crucial ones are Assumption 1, that the long-lived player always has a static incentive to fully exploit the short-lived player's trust, and Assumption 2, that the incentives to exploit grows as the short-lived player becomes more trusting (submodularity of the long-lived player's payoffs). Indeed, our results stem from the interaction of the limited records and these payoff assumptions – it is easy to demonstrate counterexamples in the case the two payoff assumptions are not satisfied.

More precisely, our results are summarized as follows. In the class of games with limited records (short-lived players observing the last K periods of the long-lived player) under our assumptions on stage-game payoffs (the main being submodularity of the long-lived player payoffs):

1. In an infinitely repeated game without types (i.e. with complete information, Model 1), there is a unique perfect Bayesian equilibrium – the repetition of the static Nash play. In other words, the lack of complete records leads to a collapse of reputation.

2. In a game with an Honorable Type, who always plays some non-opportunistic action c, all stationary PBE depend on a simple statistic of the observed history. In any equilibrium after

a history containing at least one exploiting action, the long-lived player mixes between exploiting a short-lived player's trust and mimicking the Honorable Type. When the history is "clean" (i.e. consistent with the Honorable Type), the rational long-lived player exploits his opponent's trust for sure. The short-lived players choose a positive amount of trust in every period and the trust grows in the number of periods since the last exploitation.

3. For games with an Honorable Type, we prove a payoff bound for the game with limited records if K is large enough. As the long-lived player becomes infinitely patient, this bound converges to the same value as the bound established by Fudenberg and Levine (1989) for games with complete records (so that the payoff bounds are similar despite very different equilibrium behavior). In contrast to their result, our bound applies at any time of the equilibrium play, not just at the beginning of the game, providing a stronger long-run prediction.[1,2]

These results establish the difference in equilibrium behavior between the reputation effects in games with and without types (sometimes referred to as the bootstrap and Bayesian reputation mechanisms). The main difference established before in the literature is that reputation can be sustained in a model with types even if the game has finite horizon, see for example Kreps and Wilson (1982), Milgrom and Roberts (1982) or Kreps, Milgrom, Roberts and Wilson (1982). In contrast, if the stage game has a unique Nash Equilibrium, in a finite game with complete information no dynamic incentives can be provided. We stress that our result is not a mere consequence of the horizon being in some sense "finite" due to the limited record-keeping. In fact, if we changed the assumptions on stage game payoffs, even if the stage game still had a unique Nash equilibrium, the negative result in point 1 would not necessarily hold.[3]

Our analysis may also be relevant in the context of informal/relational contracts. In markets where precise and impartial record keeping is hard or even impossible, as is the case for example in developing economies or in grey markets, transactions are often conducted with limited records. Since lack of good record keeping is often correlated with a weak legal enforcement of contracts (or the lack of it), reputation effects are believed to be an important substitute for the formal legal system.

Interestingly, similar concerns have motivated previous research on relationship building in such environments, e.g., Ghosh and Ray (1996, 2001), Watson (1999, 2002) and Watson and

---

[1] In the game with complete records there exist equilibria with high payoffs at all dates, but there also exist equilibria with low payoffs after long enough histories. We show that the reputation effect with limited records is permanent in all equilibria.

[2] A related question is about persistence of reputation in games with imperfect monitoring - as shown in Cripps, Mailath and Samuelson (2004), under complete records reputation is necessarily temporary. Ekmekci (2006) showed if histories are restricted to a rating system, there exists an equilibrium with high continuation payoffs after all histories. We conjecture that in our games with limited records for sufficiently precise monitoring all equilibria would have high payoffs after all histories.

[3] This distinction is analogous to characterizing equilibria in finitely repeated games versus infinitely repeated games but restricting equilibrium strategies to condition only on a finite history of play.

Rauch (2003). Interestingly, all equilibria in the model with an Honorable Type feature a graduate increase of trust, echoing the “starting small” feature in this literature.

As we point out below, our paper is also related to models with switching types, for example, Mailath and Samuelson (2001) and Phelan (2006). Phelan (2006) studies a game in which a government with private type decides on tax rates for a population of agents that make investment decisions. One of the types never taxes the investments and another type is opportunistic. His key assumption is that the “long-lived” player is stochastically replaced by a new player with a different type. In his model it is important to assume that when type changes, the player is replaced by a new player (so that the opportunistic player maximizes his payoff until the replacement) since otherwise equilibria become non-tractable.

In that game, the equilibrium behavior of rational type, who will be replaced by a different type eventually, resembles the equilibrium behavior of our game – the opportunistic government mixes between taxing and not taxing current investment and the more periods since the last tax, the more trusting is the population. That continues until K^* periods elapse since the last taxation (K^* being determined endogenously in equilibrium), when the government is maximally trusted and the opportunistic type taxes for sure.

There are several features that make the two models quite different. First, in Phelan (2006) since types change, it is possible to see a past taxation followed by an arbitrary number of periods of no taxation, which is never the case in our model. In fact, if types changed only once the dynamics would be very different in the two models. Second, adding commitment types in our model is a small perturbation of the complete information model in the usual sense that players believe that the complete information game happens with a high probability. In contrast, introducing changing types does not seem to be such a simple perturbation. Third, with a player dying upon a change in types, it is difficult to interpret player’s dynamic payoffs – our model with long-lived player makes interpretation of the payoff bounds much more natural.

The rest of the paper is organized as follows. We introduce the model without types in Section 2 and show the complete collapse of reputation in Section 3. Section 4 introduces the model with types and Section 5 characterizes the reputation dynamics. Section 6 studies payoff bounds and Section 7 concludes.

1.1 Why Study Limited Records?

In our opinion, there are several reasons why the study of games with limited records is relevant and interesting. First, in many situations records are limited by institutional design. For example, the Better Business Bureau reports customer complaints over the last 36 months, the credit history for US customers contains the last 7 years of their activities[4], points for driving violations are in many states removed from driving records after time of safe driving.[5] Since in some situations the amount of record keeping is a part of the market design, we think it is important to understand the incentives players face in markets with limited records. For example, the current design by eBay is such that any feedback received from buyers becomes a permanent part of seller's record.[6] Yet, eBay may consider changing this policy and reporting only more recent feedback,[7] and given the way eBay currently displays summary statistics on the full history, for many buyers the relevant information may be similar to the limited-records in our model.[8]

Second, in many practical problems of repeated interactions, the history is conveyed by word of mouth or the short-lived players represent overlapping generations. For example, college students transmit information from seniors to juniors about professors' behavior. In many such environments, players would naturally observe only limited past data. Our modeling approach may then be of interest for analyzing the university's decision of how many past years of student evaluations to reveal about current professors.

Third, collecting information about past behavior of new trading partners (e.g. background checks) can be costly and in equilibrium short-lived players would choose to obtain limited information.[9]

Fourth, records limited to the last K periods can be treated as a stylized simplification of information depreciation that allows us to sharply characterize equilibria. Based on our analysis, our intuition is that what is crucial for the results is that information depreciates over time. In our case after information moves beyond the K horizon, it is lost. A different modeling approach has been used in Mailath and Samuelson (2001), Phelan (2006) and Wiseman (2008) – these papers study games in which player 1 is stochastically replaced by a new player 1 with

---

[4] Outright bankruptcy may be difficult to erase from the record, but many other events that reduce credit score are erased after seven years.

[5] See http://www.carinsurance.com/Articles/content246.aspx: In Michigan, "points are only displayed on a driving record for two years from the date of conviction". Pennsylvania "allows points to be removed from your driving record for safe driving. You can get 3 points removed from your driving history for every 12 consecutive months (from the date of the last violation) you go without a violation, which results in points, license suspension, or revocation. Once a driving record is reduced to zero and remains at zero points for 12 consecutive months, any further accumulation of points is treated as the first accumulation of points."

[6] As of June 2009, see http://pages.ebay.com/services/forum/feedback.html.

[7] That may be beneficial especially since eBay right now shows most prominently the percentage of positive feedback. After long enough histories this statistic may contain little information about recent conduct and could undermine the system. For a recent paper on reputation effects when only some statistics of the history are displayed to short-lived players, see Ekmekci (2006).

[8] Interestingly, on page "detailed feedback" eBay reports a summary of feedback received by a seller. It is grouped by Last Month, Last 6 Months and Last 12 Months. Moreover, eBay computes Average Detailed Ratings on a rolling 12-month basis and shows summary of detailed feedback over the last 12 months, see http://pages.ebay.com/help/feedback/contextual/detailed-feedback-ratings.html (as of June 2009).

[9] Here we take the limits on monitoring as exogenous. Liu (2009) investigates endogenous costly monitoring for a subclass of our games.

different type.[10] As a result, as time passes, observing a non-commitment action at some past date conveys less and less information about the current type, leading to similar dynamic forces. Tadelis (1999), as in Mailath and Samuelson (2001), identifies a different form of impermanent reputations: the possibility that a firm secretly changes its name or buys an existing name. We note that though Tadelis (1999) has an OLG structure with firms living only for two periods, it is clear that even in a more general model his assumptions would deliver at least one feature consistent with our model: "clean names," i.e. histories with no bad outcomes, are not totally trusted in equilibrium. In his model the reason is that old information depreciates because there is a chance that the name has been sold.

Fifth, it is known from the psychological literature that people do not act on the whole history they observed, but display several biases, one of which is to pay more attention to recent history. This so-called "recency effect" has been documented and studied extensively; see, e.g., Murdock (1962), Bjork and Whitten (1974), Broadbent and Broadbent (1981). That has motivated a new string of literature on the economics with finite memory.[11]

Sixth, it has been recently shown that whether equilibria in repeated games are robust to perturbing them to private monitoring depends crucially on whether the strategies depend on finite or infinite histories (see Mailath and Morris (2002, 2006)). These results motivate a wider interest in what can be achieved with strategies that depend only on a limited number of periods in the history (see Jehiel (1995), Bhaskar (1998) and Cole, Kocherlakota (2005) for early negative examples and Hörner and Olszewski (2008), Mailath and Olszewski (2008) for a Folk Theorem). Our paper contributes to this literature by looking at the class of games with one long-lived and a sequence of short-lived players.

In summary, we think there are many economic problems in which limited records are relevant, either directly, or as a design option or as a parable for other frictions. In this paper we focus on models of reputation and show how limited records change the equilibria of those games.

As we show, the extent to which reputation can be sustained is not only a question of how many periods back are kept in the records, but also whose actions are kept. For most of the paper we assume that only the long-lived player's actions are kept in the records, and a short-lived player's actions are not observable to future short-lived players. We show that this does matter – there are examples in which players could sustain much higher payoffs using reputation effects/intertemporal incentives if the short-lived players' recent actions could be observed.

---

[10] Phelan (2006) is closer to our model since he and we assume that actions are observed perfectly, while Mailath and Samuelson (2001) assume imperfect monitoring. Wiseman (2008) analyzes the chain-store paradox game with Markovian types.

[11] Even though our paper can be interpreted as a game with the short-lived players not paying attention to or forgetting the past, we stress that our model has perfect recall, so that we can still apply standard equilibrium notions.

There are at least two reasons we think studying the case where short-lived players' actions are not observable is relevant. First, the designer of the record-keeping system may not want to reveal the data for privacy reasons (for example, a university may not want to reveal the grades of the students that wrote negative feedback for a course, eBay may not want to reveal the size of the transactions of buyers providing feedback, etc.). Second, in some situations the long-lived player may observe the short-lived players' actions privately and is not able to reveal them credibly. Beyond the economic relevance of unobservability of the short-lived players' actions, we stress that the failure of the Folk Theorem of Fudenberg, Kreps and Maskin (1990) is due to the limited record keeping but not the unobservability of the short-lived players' actions.

2 Model 1 (without Types)

The repeated games we study are defined as follows. There is one long-lived player (player 1) who over time $t \in \{0, 1, ...\}$ plays with an infinite sequence of short-lived players. We refer to a generic short-lived player as player 2. A short-lived player who arrives at time t plays one stage game (the trust game) with the long-lived player and exits the game.

In the stage *trust game* the players move simultaneously. Player 2 chooses the amount of trust, $y \in [0, 1]$ (denote the action space by Y). Player 1 decides how much to honor player 2's trust by choosing $x \in [0, 1]$ (denote the action space by X), with $1 - x$ being a measure of how much player 1 "abuses" player 2's trust.[12] This stage game generalizes the stylized product-choice game described in Mailath and Samuelson (2001, 2006).

The main novel assumption in our analysis is about the information of the short-lived players. Unlike most of the reputation literature, we assume that the short-lived players observe only a finite and partial history of past play.[13]

In particular, we assume the game has K-recall and that only the actions of player 1 are public. That is, a short-lived player, upon entering the game, observes only the actions chosen by player 1 in the previous K periods. We refer to this assumption by saying that only finite records are kept in the game or that the game has finite monitoring, as opposed to traditional models where short-lived players see the full history, which we refer to as complete records or full monitoring. For most of the paper we assume short-lived players do not observe the actions

---

[12] Of course, other interpretations of x and y are possible, see examples below.

[13] As we discussed in the Introduction, this assumption captures the fact often records are kept for a limited amount of time or that players decide not to condition on information far past in time. Liu (2009) studies a different reputation model (the product-choice game) in the case that any short-lived player has potentially access to all past information, but has to pay a cost to learn about every additional past period. Hence, he shows a way to endogenize the finiteness of records that we assume exogenously in this paper. However, under costly information acquisition in Liu (2009), the equilibrium sizes of records are necessarily random.

of past short-lived players, but in Section 3.1 we show how reputation could be enhanced if the records also included the actions of short-lived players.

Formally, denote the set of period $t > 0$ histories of player 1's play by $H^t = X^t$, and the initial history by the null set $H^0 = X^0 = \{\varnothing\}$. The set of all histories is $H = \cup_{t=0}^{\infty} H^t$. A history $h = (x^0, x^1, ..., x^{t-1}) \in H^t$ specifies player 1's play in period 0 through period $t-1$. For any $h \in H^t$ and $k > 0$, let

$$\tau^k(h) = \begin{cases} (x^0, x^1, ..., x^{t-1}), & \text{if } t < k \\ (x^{t-k}, x^{t-k+1}, ..., x^{t-1}), & \text{if } t \geq k \end{cases}$$

be the k-period truncation of history h. Consistent with the definition, write $\tau^0(h) = \varnothing$. The strategy of player 2 who enters the game at period t is $\sigma^t : H^t \rightarrow \Delta(Y)$, where $\Delta(Y)$ is the set of Borel measures on Y. Since player 2's information is summarized by the truncation function τ^K, σ^t is τ^K-measurable. Player 1's strategy is $\pi : H \rightarrow \Delta(X)$.[14]

The stage game payoffs are $g_1(x, y), g_2(x, y)$ for the two players respectively and they are both continuous. Player 1 discounts future payoffs by $\delta \in (0, 1)$ and maximizes the expected sum of discounted payoffs. Each short-lived player maximizes his stage game payoffs.

We make the following assumptions about the stage game payoffs and provide several examples in the next section to show that all can be generically satisfied:

Assumption 1 (myopic best response of player 1) $g_1(x, y)$ is weakly decreasing in x for each y, and strictly decreasing when $y > 0$. As a consequence, $x = 0$ is a dominant strategy for player 1 in a single repetition of the game.

Assumption 2 (submodularity of player 1) $g_1(x, y) - g_1(x', y)$ is strictly increasing in y for any $x < x'$. In words, player 1 has stronger static incentives to abuse player 2 for a higher (trusting) action of player 2.

Assumption 3 (unique best response of player 2) For any player 1's (mixed) action, player 2 has a unique best response, denoted by $y^*(x)$ for pure actions and $y^*(\pi)$ for mixed actions.

Assumption 4 (monotone best response of player 2) Player 2's unique best response to player 1's mixed action (an element of $\Delta(X)$) is strictly increasing when $\Delta(X)$ is partially ordered by first order stochastic dominance. We normalize the payoffs and remove strictly dominated strategies for player 2, so that player 2's best response to player 1's action 0 is $y^*(0) = 0$. That makes $(0, 0)$ the unique Nash equilibrium of the stage game.

An easy-to-check sufficient condition for Assumption 3 to hold is strict quasi-concavity of $g_2(x, y)$ in y for each x, because it implies the strict quasi-concavity of $g_2(\pi, y)$ in y for each π. Pure

---

[14] Although we allow mixed strategies for both players, Assumption 3 below guarantees that every player 2 in equilibrium plays a pure strategy.

strategy is usually considered in games with continuum action sets (see Mailath and Samuelson (2006)). Nevertheless, mixed strategies are important for a player with private information. The merit of Assumption 3 is that it guarantees the set of player 2's best responses is simply ordered. This fact is used in conjunction with Assumption 4 in our proofs.[15] An easy-to-check sufficient condition for Assumption 4 to hold is:

Lemma 1 *Assumption 4 is implied if:*
a) $\frac{\partial}{\partial y} g_2(x,y)$ *is strictly increasing in* x *(differentiable version of supermodularity)*[16] *and*
b) player 2*'s best response to any* π *that is not degenerate at* 0 *is interior.*

Our equilibrium concept is the perfect Bayesian equilibrium of the repeated game (PBE). Note that despite player 2 observing only a finite history, the game has perfect recall. Therefore we can use this standard equilibrium notion in our analysis.

Note that by Assumption 3, player 2 plays a pure strategy in any equilibrium (but player 1 can potentially mix). Therefore in what follows we abuse notation and use $\sigma^t(h)$ to denote the *pure* equilibrium strategy of player 2 in time t.

Given an equilibrium $(\pi, \{\sigma^t\}_{t=0}^{\infty})$ we denote by $V(h)$ the expected equilibrium payoff of player 1 (rescaled to average per-period payoffs) given a history h. We write (h, x^t) to represent a history h followed by player 1 choosing x^t.

2.1 Examples

Before we analyze the repeated trust games, we provide a few examples that motivate this class of games.

Consider the following transactions.

Example 1 *Player 1 is a software developer offering a freeware program for download. He decides on the quantity of adware/spyware to include in the program (with* x *decreasing in the amount and aggressiveness of these add-ons). The short-lived players represent customers that may use the program, with* y *being the number of people that download the program. Users leave feedback about the amount and aggressiveness of the adware on a download site (for example,*

---

[15] If player 2's action set is binary, say $Y = \{0,1\}$, then Assumption 3 can be relaxed since the set of mixed strategies with binary choices is simply ordered according to first order stochastic dominance. In this case, the folloiwng modified version of Assumption 4 suffices for our proofs: player 2's set of pure strategy best responses is nondecreasing according to the strong set order (c.f. Topkis (1998)) in player 1's mixed strategies (partially ordered by first order stochastic dominance), and there exists a cutoff $d \in (0,1]$ on player 1's action set $X = [0,1]$ such that player 2's best response to player 1's pure action $x \in [d,1]$ is $y = 1$ and his best response to $x \in [0,d)$ is $y = 0$. We demonstrate how the proof of Theorem 1 works with these alternative assumptions in Appendix.

[16] It turns out that if we only assume that g_2 is supermodular (instead of the differentiable version), we can only establish that the best response is weakly increasing in π – see the discussion in the end of proof.

download.com) and new users read entries from the last K periods. Player 1 benefits from installing adware proportionally to the number of people that download the freeware. Users prefer programs with less adware.

Example 2 *An agent, player* 1, *receives investment* $y \in [0,1]$ *from the principal, player* 2, *and produces output* $f(y)$, *where* $f(0) = 0$, $f'(y) > 0$, $f''(y) < 0$, *and* $f(y) > y$. *The agent then decides the fraction of output to conceal for his own consumption. Let* x *be the fraction that is returned to the investor. Therefore* $g_1(x,y) = (1-x)f(y)$ *and* $g_2(x,y) = xf(y) - y$. *Phelan (2006) is a special case in which the government (player 1) chooses the tax rate* $1 - x$ *and the households (player 2) choose the investment* y. *The central trade-off is that player 1 always wants to conceal more output for his own consumption, but player 2 invests only when he anticipates a positive return.*

Example 3 *A firm, player* 2, *decides the amount of time* $y \in [0,1]$ *to spend with a consultant (or lawyer), player* 1, *who charges a rate* p *per unit of time. The consultant decides his effort per unit of time* $x \in [0,1]$, *yielding* xy *units of effective effort. Player* 2*'s payoff is* $g_2(x,y) = v(xy) - py$ *and player 1's payoff is* $g_1(x,y) = py - xy$.

3 Collapse of Reputation with Finite Records

Our first main result is that every repeated trust game (with finite records and under the assumptions listed above) has a unique PBE outcome: in every period the players play the static Nash equilibrium $(0,0)$. That is true for every finite K, no matter how patient is player 1 :

Theorem 1 *For any* $K < \infty$ *and* $\delta < 1$, *the infinite repetition of* $(0,0)$ *is the unique PBE of the repeated trust game .*

We present a detailed proof in the Appendix and provide the main economic intuition here. The collapse of reputation sustained by dynamic incentives is mostly caused by the finite records, but the assumptions about the stage-game payoffs also play an important role in the result.

Suppose that the best response to action $x = 1$ is $y = 1$ and suppose the players try to support as equilibrium full trust and no abuse $(1,1)$. For that player 1 needs to be punished by some future short-lived players if he deviates from $x = 1$ today. Suppose such a deviation happened exactly K periods ago and therefore today the players see a history consisting of honorable play $(x = 1)$ in the last $K - 1$ periods and some deviation K periods ago.

Can it be that player 1 is punished in the current period? We claim not. The reason is that the current player 2 may not have incentives to punish player 1 today if he expects player 1 to behave honorably, which would allow player 1 to “clean up” his history (entering the next period the history observed by player 2 would be straight “1”). The reasoning is as follows. Note that punishment requires that in the current period player 2 plays $y < 1$. Now, if an infinite play of $(1, 1)$ is an equilibrium it must be that player 1 weakly prefers to play $x = 1$ when player 2 plays $y = 1$. But in the current period $y < 1$, and hence player 1’s current-period temptation to abuse player 2’s trust is *strictly less* than on the equilibrium path. Therefore even if a deviation has happened K periods ago, player 1 would play $x = 1$ for sure in the current period, because future player 2s cannot distinguish whether the current period’s deviation is the first or second deviation in the game. That leads to a contradiction since if player 1 plays $x = 1$ today, player 2 does not want to follow up on the punishment (and since he exits the game immediately, there is nobody to punish him!).

Hence, we see that player 1 cannot be punished after histories that have a publicly recorded deviation only K periods ago. By induction, repeating the reasoning, player 2 will not punish player 1 after histories that possibly have deviations only K and $K - 1$ periods ago and so on... As a result, player 1 gets the same continuation payoff, no matter what is the history of play. That means that he plays a static best response in any equilibrium – and that is unique, by Assumption 1. We generalize this argument to any candidate equilibrium and that implies that the unique equilibrium is the infinite repetition of the static Nash equilibrium.

To finish this section, we contrast our impossibility result with the Folk-theorem result under full monitoring. If player 1 is sufficiently patient, under the same assumptions on stage-game payoffs it is easy to sustain other equilibrium outcomes if the players observe the full history of play. In fact, it is sufficient that the short-lived players observe the full history of player 1 past actions to sustain other outcomes.

To see this, consider a simple grim-trigger equilibrium: on the equilibrium path players choose $(x, y^* (x))$ for some $x \in X$. If player 1 ever deviates, the players revert to playing $(0, 0)$ forever. Standard arguments show that if $g_1 (x, y^* (x)) > g_1 (0, 0)$, then for high enough δ the grim trigger strategy profile indeed forms an equilibrium (which payoff-dominates the unique outcome that can be sustained if records, K, are long but finite).

There is a subtle reason for the stark difference in the sustainability of high payoffs with dynamic incentives when finite versus full records are kept. It may seem that the key is the infinite-punishment of the grim trigger strategies – no matter how long ago player 1’s non-honorable behavior was observed, the strategy calls for continued punishment. Yet it is not, since for sufficiently patient players we can construct equilibria with finite length of punishment. For example, a non-honorable behavior of player 1 (deviation from x) is followed by m periods of punishment play $(0, 0)$, followed by a return to the original equilibrium (which allows players to forget about past deviations).

The subtle reason this cannot work with finite records, even if $m < K$, is that it is difficult to fully punish player 1 for not following up on such punishment. And once player 2 does not expect player 1 to follow up on the punishment in the current period, he has no incentives to punish him. It would be better collectively for the short-lived players if each of them followed up on punishments, but each of them separately is going to behave opportunistically.

More precisely, in the equilibrium with full records, if either of the players deviates from the punishment, we would in general need to restart the punishment (or start some other punishment for deviating from punishment) – if such multiple deviations are not kept in the history, player 1 has a way to "clean up" his history and the short-lived players have individual incentives to help him!

3.1 Records of Short-lived Player Actions and Comparison with Existing Results

We are assuming throughout the paper that the actions of player 2 are not recorded, either because of privacy reasons (players that exit the game may not want the information about their past actions to be available to others) or because player 1 observed these actions privately and cannot credibly reveal them.

We now show that, at least in some games, it is possible to sustain better outcomes even if records are finite in terms of their length, but they contain both player 1 and player 2 past actions. To see this possibility, consider the following equilibrium with $K = 1$ and both x and y from last period being observed. In period 0 player 1 plays $c > 0$ and player 2 plays $y^* (c)$. In any future period if $(c, y^* (c))$ is observed, both players continue with $(c, y^* (c))$. If any other history is observed, then players play $(0, 0)$. If $g_1 (c, y^* (c)) > g_1 (0, 0)$, then there clearly exists δ high enough so that the described profile forms an equilibrium. The difference comes from the impossibility to "clean up the history" - suppose that player 1 deviates to 0 in period 0 and then continues playing c afterwards. Unlike our original model, even though past deviation of player 1 is no longer visible in period $t = 2$, the lack of trust of the short-lived player in period 1 is enough of a coordination signal for the short-lived player in period 2 that he knows that his best response is to play 0. Technically, the past play of player 2 allows the short-lived players to keep track of past deviation of player 1 in an incentive-compatible way, a form of contagion. In everyday language, player 1 is not trusted by the short-lived players not because they know details of his past non-honorable behavior, but because they know player 1 was not trusted by the recent short-lived players, which is bad news about player 1.

The proof of Theorem 1 uses a similar logic to Cole and Kocherlakota (2005), but our result is different for at least two reasons. First, we cover all equilibria in our class of games (while they consider a subset of equilibria in prisoners' dilemma games). Second, as discussion in this subsection illustrates, our result depends on the limited records not containing the short-lived player actions, a condition that plays no role in Cole and Kocherlakota (2005).

Jehiel (1995) and Bhaskar (1998) obtain degeneracy of the set of pure strategy equilibria under limited memory in sequential move games.[17] In both papers only one player chooses an action in a given period and this structure turns out to be important for their argument: when a player moves in period t, he knows that a player at $t+1$ will not know the play in period $t-K$. Since the pure strategy played at $t+1$ is a function of actions in periods $t-K+1, ..., t$, the player choosing action at t knows that the action taken at $t-K$ has no impact on his payoffs and hence his best response does not depend on it. Our logic is quite different. In the game with simultaneous moves, even though players know that what will happen tomorrow will not depend on what players saw $t-K$ periods ago, they can still coordinate today based on what they saw $t-K$ periods ago. Therefore, the induction/unraveling step requires assumptions on the stage-game payoffs and does not follow simply from the information structure of the game. Finally, notice that Theorem 1 covers both pure and mixed strategy equilibria, while the degeneracy results in Jehiel (1995) and Bhaskar (1998) cover only pure strategy equilibria.

4 Model 2 (with an Honorable Type)

The second model we study introduces a behavioral type for player 1, as in the large literature on reputation with types. We explore two questions:

1) What is the impact of finite records on the possible equilibrium behavior in a game with an Honorable Type?

2) Given that the repeated-game effects and behavioral types are the two main ways we model reputation in economics and these two models under full records typically can derive the same predictions/rationalizations,[18] are the predictions different under finite records?

Consider the following modification of Model 1. Player 1 has two types. With probability $\mu^* \in (0,1)$ player 1 is a commitment type (the Honorable Type) who always plays $c \in (0,1]$ in the repeated game. With probability $1-\mu^*$ player 1 is rational with the strategies and payoffs specified above. The type space is $\Theta = \{r, c\}$, where r stands for "rational type" and with some abuse of notation, c for "honorable type" who always plays action c. We emphasize that c is not

---

[17] Jehiel's (1995) model assumes infinitely lived players while Bhaskar (1998) considers two-period lived players in an OLS model.

[18] See Mailath and Samuelson (2006 Ch.15) for detailed discussion about the two approaches, sometimes referred to as the bootstrap and Bayesian mechanisms. Models of reputation via the first approach have for example been used to investigate the time consistency of government policies (see Barro and Gordon (1983) and Stokey (1989)). See also Ljungqvist and Sargent (2004) for a discussion of that reputation literature.

necessarily the Stackelberg type.

For instance, in example 2 above type c agent always returns fraction c of the output and in example 3 the type c consultant always puts effort c.

We add two additional assumptions about the stage game payoffs to capture the trade-off of reputations. These assumptions are satisfied by all our examples above.

Assumption 5 (monotonicity) $g_1(x, y)$ is weakly increasing in y for each x and strictly increasing in y when $x = 0$; namely, player 1 prefers a higher action of player 2.

Assumption 6 (value of reputation) $g_1(c, y^*(c)) > g_1(0, 0)$ (recall $y^*(c)$ is player 2's best response to the commitment action c). In words, player 1 prefers player 2 to believe that he is a commitment type and to act accordingly than to play the static Nash equilibrium of the game without an Honorable Type.

We will also focus on the case of high-enough discount factor. Let

$$\bar{\delta} = \frac{g_1(0, y^*(c)) - g_1(c, y^*(c))}{g_1(0, y^*(c)) - g_1(0, 0)}$$

Assumptions 1 and 6 guarantee $\bar{\delta} \in (0, 1)$. From now on we assume $\delta > \bar{\delta}$.

4.1 Stationary Strategies

For tractability, in this model we introduce another restriction about the information sets of the short-lived players. In particular, assume that the short-lived players do not observe the calendar time, and before entering the game, they share a common prior belief P over the time in which they are likely to enter the game - i.e., they believe that they will enter in period $t \geq 0$ with probability $P(t) \geq 0$.[19]

As a result, the collection of information sets in which player 2 can be is $\cup_{t=0}^{K} H^t$ - the set of histories with length at most K. We denote that state space H_+^K. We denote by H^K the subset of the state space that corresponds to histories of length at least K (H_+^K is the set of histories of length K, H^K, *plus* all histories of shorter length).

Given this assumption, we shall write player 2's strategy as $\sigma : H_+^K \to Y$. Finally, we restrict analysis to (stationary) PBE in which player 1 also plays a strategy which depends only on the information set of player 2. That is, we look at PBE in which $\pi : H_+^K \to \Delta(X)$.[20] From now on we abuse the notation by writing $h \in H_+^K$ when we mean $\tau^K(h) \in H_+^K$.

---

[19] Note that despite player 2's not observing the calendar time, the game still has perfect recall and hence we can apply standard equilibrium notion.

[20] By a standard argument, if player 2's strategy is measurable with respect to H_+^K then player 1 has a best response which is also H_+^K-measurable. Hence, for every equilibrium in which player 1 strategy is a more complicated function of the history, there exists a stationary equilibrium with the same payoff to player 1.

Standard Bayesian updating implies that π, P and μ^* induce a posterior belief for player 2 over player 1's type space Θ, $\mu(\theta|h)$, for each $h \in H_+^K$ that is reached in equilibrium with positive probability. For any off the equilibrium histories in which player 1 played at least one action different from c in the last K periods, we assume $\mu(c|h) = 0$.[21] Of course, for all histories reached on the equilibrium path that contain at least one action different from c in the last K periods, the Bayes rule implies $\mu(c|h) = 0$ as well). Following the literature, we call player 2's posterior belief $\mu(c|h)$ player 1's reputation (as a committed player).

Our goal is to characterize all stationary PBE:

Definition 1 (π, σ, μ) *is a stationary PBE if* π *and* σ *(which depend only on* H_+^K*) are best responses to each other given* μ *and on the equilibrium path* μ *is consistent with the Bayes rule given* π *and priors* P*,* μ^* *(while off-the path, if player* 2 *observes any action different from* c*, then* $\mu(c|h) = 0$*).*

Note that in Model 1 (which is Model 2 with $\mu^* = 0$ and all players observing the calendar time), the stationary PBE is also well-defined. Since stationary PBE are a subset of all PBE, Theorem 1 holds also for this equilibrium notion.

For parts of the analysis it will be easiest to describe the equilibria in case P is the uniform improper prior, that is, if we take the limit of the Bayesian updating to compute $\mu(\theta|h)$ as $P(t)/P(t+1) \to 1$ uniformly for all t. For $h \in H^K$ the resulting beliefs are equal to the fraction of time type θ reaches a given state on the equilibrium path (the off-path beliefs are still $\mu(c|h) = 0$).

5 Equilibrium with an Honorable Type

In this section we characterize strategies in all stationary PBE.

5.1 Sufficient Statistics of History

The set H_+^K is a continuum, making it potentially difficult to describe the equilibrium strategies. We decompose the set into three subsets. The first subset consists of histories in which a non-commitment action has been played, and we denote this set by H_{++}^K. The second subset is a singleton - history $(\underbrace{c, ..., c}_{K})$ i.e. a history in which player 1 played c in the previous K periods.

---

[21] The only other off-equilibrium histories are such that $P(t) = 0$. For $t < K$, we assume that in that case player 2 updates using π and σ. For $t > K$, we assume that $P(t)$ assigns positive probability to at least one such time.

We call it the "clean history". The residual is a finite set $\{\varnothing, C^1, ..., C^{K-1}\}$ where C^k is a history of length k and the commitment action has been played in all periods so far.

We can partition H^K_{++} further into a finite collection by counting how close the player is to a clean history. Formally, $I(h)$ for a history $h = (x^0, x^1, ..., x^t) \in H^K_{++}$, $0 < t < K$ is

$$I(h) = K - 1 - \max\{k : x^k \neq c\}$$

Therefore, $I(h)$ measures the number of commitment actions in h since the most recent non-commitment action. For convenience, let us define $I(h) = K$ if $h = (\underbrace{c, ..., c}_{K})$ and call $I(h)$ the commitment index of a history. A commitment action played on a history h with $I(h) = k$ will increase this index by 1 (or the index remains at K if $k = K$). A non-commitment action will reduce the index down to 0.

For a fixed K, we denote the set of indices as $\mathbf{I} = \{0, 1, 2, ..., K\}$. We claim that $I(h)$ contains all the strategically relevant information in any stationary PBE upon history $h \in H^K_{++}$:[22]

Proposition 1 *In any stationary PBE, equilibrium strategies for histories in H^K_{++}, on and off the equilibrium path, depend only on $I(h)$. That is, if $I(h) = I(h')$, then $\sigma(h) = \sigma(h')$ and $\pi(h) = \pi(h')$.*

This result greatly simplifies our analysis by allowing us to focus on a finite state space $\mathbf{I} \cup \{\varnothing, C^1, ..., C^{K-1}\}$. We call $\mathbf{I}$ the set of *regular indices for regular histories* and $\{\varnothing, C^1, ..., C^{K-1}\}$ the *special indices for special histories.* From now on write strategies $\pi : \mathbf{I} \cup \{\varnothing, C^1, ..., C^{K-1}\} \rightarrow \Delta(X)$ and $\sigma : \mathbf{I} \cup \{\varnothing, C^1, ..., C^{K-1}\} \rightarrow Y$. We write $\mu_k = \mu(c|h)$ for a regular history h with $I(h) = k$. By definition, $\mu_0 = \mu_1 = \cdots = \mu_{K-1} = 0$ and $\mu_K > 0$.

To simplify notation, we write $y_k \equiv \sigma(k)$ for regular indices. Since as we established in Proposition 1 the continuation payoff of player 1 depends only on the commitment index of history, if he chooses to play an action different from c, his best response is to play 0 (his myopic best response; see Lemma 6 in the Appendix for a formal proof). In other words, $\pi(k)$ assigns positive probability to at most two actions: $\{0, c\}$. Abusing notation a bit, we write player 1's strategy hence as β_k where β_k is the probability player 1 assigns to action c (which means he assigns probability $1 - \beta_k$ to action 0) and player 2's strategy as y_k. We also write $y^*(\beta_k)$ as player 2's best response when he believes that action c is played with probability β_k.

---

[22] Even though we have assumed $\delta > \overline{\delta}$, this proposition holds for all δ.

5.2 Equilibrium Characterization

We now focus on the strategies in states in H^K, i.e. after histories h^t for $t \geq K$. We return to the initial histories in Section 5.3. We claim the following characterization of stationary PBE. In every equilibrium, when the index of the history is $k < K$, player 1 mixes between 0 and c. When he achieves the "clean history", index K, he plays 0 for sure, exploiting his reputation. In every period player 2 trusts player 1 to some extent – and the equilibrium degree of trust is increasing in how close the history is to the clean history – and is maximized at the clean history. Formally:

Theorem 2 *For $\delta > \bar{\delta}$, any stationary PBE takes the following form:*

1. There exists a strictly increasing sequence $\{\beta_k\}_{k=0}^{K-1} \subset (0,1)$, such that if the index of the history is $k < K$, player 1 plays c with probability β_k and 0 with probability $1-\beta_k$; player 2 plays $y_k = y^*(\beta_k)$

2. If the index is K, player 1 plays action 0 (with prob.1); player 2 plays $y^*(\mu_K)$, where $\mu_K > \beta_{K-1}$.

Since player 1 is indifferent in states $k < K$, the IC constraints require for $k \in \{0, ..., K-1\}$:

$$g_1(0, y_k) - g_1(c, y_k) = \delta \left[g_1(0, y_{k+1})) - g_1(0, y_0) \right], \tag{1}$$

where player 2's best response requires $y_k = y^*(\beta_k)$. These equations almost pin down the equilibrium. The only remaining equation is for state K to pin down μ_K and y_K. That in general depends on $P(t)$, μ^* and the equilibrium strategies after the special histories (in the initial $K-1$ periods with no deviations from c).

In the special case of improper uniform prior we can characterize the equilibrium even further:

Proposition 2 *Assume P is the "improper uniform prior". The equilibrium strategies $\beta_0, \beta_1, ..., \beta_{K-1}$ and the posterior belief μ_K are completely characterized by (1) and:*

$$\mu_K = \frac{\mu^*(1 + \beta_0 + \beta_0\beta_1 + \cdots + \beta_0\beta_1...\beta_{K-1})}{\mu^*(1 + \beta_0 + \beta_0\beta_1 + \cdots + \beta_0\beta_1...\beta_{K-2}) + \beta_0\beta_1...\beta_{K-1}} \tag{2}$$

$$y_K = y^*(\mu_K). \tag{3}$$

The proof of Theorem 2 is developed through a series of lemmas that show the following:

1. When the history is clean (with index K), then type r of player 1 plays 0 for sure. If not, he would play c after every history because submodularity of g_1 implies that if he is not tempted to play 0 when player 2 plays y, he is not tempted for all $y' < y$. But if he plays c after every history, then for all histories player 2 would play $y^*(c)$ – there would be no intertemporal

incentives. As a result, it would be a strict best response for player 1 to play 0, a contradiction. Similarly, for any regular history it cannot be the case that player 1 plays c with probability 1 - that would be the best period for him to deviate to 0!

2. After a history with index k, player 1 either plays 0 with probability 1 or he plays c with probability strictly higher than in state $k-1$. The reasoning is as follows: suppose player 1 plays c with positive probability in states $k-1$ and k. The continuation payoff from playing 0 is the same in both cases, but playing 0 in state $k-1$ delivers the extra myopic payoff sooner. For player 1 to be willing to play c in state $k-1$ it must be that he is rewarded next period with an even higher reward for playing 0, to compensate for discounting (this is captured by equation (1)).

3. If $\delta > \bar{\delta}$ then playing 0 in every period is not an equilibrium. In fact, playing 0 for sure in any of the regular histories is not an equilibrium. The reason is that otherwise in state K player 2 would assign a very high probability to type c and play close to $y^*(c)$. If δ is high enough then type r of player 1 prefers to mimic type c instead of deviating once and getting $g_1(0,0)$ afterwards. Combining 1 and 3, $\beta_k \in (0,1)$ for all $k \in \{0,...,K-1\}$ and (1) requires that $\{\beta_k\}_{k=0}^{K-1}$ is strictly increasing and $\mu_K > \beta_{K-1}$.

To summarize, we have established in this section that the equilibrium behavior in the model with an Honorable Type and finite records is quite different from the equilibrium behavior in Model 1. We can also argue that it is different from possible equilibrium behavior if we allowed the short-lived players to always see complete history of player 1 actions. For example, for sufficiently high δ and full records it is easy to construct grim-trigger equilibria in which player 1 always plays c. In contrast, the recurring building and exploiting of reputation that is a feature of any equilibrium with finite records cannot be supported as an equilibrium behavior in the game with complete records. To see this, note that for our equilibrium in which player 1 mixes, it is necessary that y_k is strictly increasing in k. But y_k is bounded by $y^*(1)$. So in some period player 1 would strictly prefer to play 0. In our case it is when the index reaches K. But in a game with complete records if player 1 ever played an action different than c in the past, player 2 would know that he faces the r type and would expect $x = 0$ is coming. As a result, player 2 would not trusting player 1 and avoid being exploited (play $y = 0$). But that would mean that player 1 should deviate to 0 one period before and the whole equilibrium would unravel. So if player 2 sees the full history of player 1's play, then there is no room for player 1's "manipulation" – player 2 can be "surprised" on the equilibrium path only once.

5.3 Initial K periods

In the initial K periods of the game, player 2 knows the calendar time simply from the length of the history. As a result, he (and hence player 1) can play differently in these initial periods than later in the game. We have already established in Proposition 1 that if any action $x \neq c$ is observed in the initial periods, the equilibrium depends only on the commitment index of the history.

So to finish characterization of the equilibrium we need to pin down $\sigma\left(C^k\right)$ and $\beta\left(C^k\right)$ for $k \in \{0, ...K-1\}$ (with $C^0 \equiv \varnothing$ being the empty history). First of all, a stationary PBE always exists: it is described by a finite collection of $\beta_0, ..., \beta_K, \beta\left(C^0\right), ..., \beta\left(C^{K-1}\right), \mu_K, \left\{\mu\left(c|C^k\right)\right\}_{k=1}^{K-1}$ (and the corresponding unique player 2 best responses) – since the payoffs are continuous and the action space is compact, standard arguments guarantee existence. For a general $P(t)$, $\sigma\left(C^k\right)$ and $\beta\left(C^k\right)$ can be computed as a fixed point of best-response conditions and Bayes rule equations, noting that they depend on μ_K via $V(K)$ and in turn, if $P(t) > 0$ for $t \in \{K, 2K-1\}$, μ_K depends on $\beta\left(C^k\right)$.

Finally, in the end of proof of Theorem 2 we have also established that for all $k \in \{0, K-1\}$, if in equilibrium $\beta\left(C^k\right) > 0$:

$$\begin{aligned} y_k &\geq \sigma(C^k) \\ V(k) &\geq V\left(C^{k-1}\right) \end{aligned}$$

(and if $\beta\left(C^k\right) > 0$ then $\sigma(C^{k+1}) > \sigma(C^k)$). The inequality $y_k \geq \sigma(C^k)$ additionally implies that

$$\beta_k > \beta\left(C^K\right)$$

since at history C^k player 2 believes that player 1 will play c with probability at least $\mu^* + (1-\mu^*)\beta\left(C^K\right)$ (and his response is weakly lower than that to β_k). Note that this last bound holds even if $\beta\left(C^k\right) = 0$ (which can be the case for some μ^*).

6 Payoff Bounds

We now turn to the equilibrium payoffs. Using the technique of Fudenberg and Levine (1989) it is easy to obtain the following bound:

Lemma 2 *For any $\varepsilon > 0$ and $\mu^* \in (0,1)$, there exists an integer $K(\varepsilon, \mu^*) > 0$ (independent of δ) such that if $K > K(\varepsilon, \mu^*)$ then in any stationary PBE player 1's payoff computed at period 0 is at least $B(\delta, K) = (1 - \delta^{K(\varepsilon,\mu^*)})g_1(c, 0) + (\delta^{K(\varepsilon,\mu^*)} - \delta^K)g_1(c, y^*(c)) + \delta^K g_1(0, 0) - \varepsilon$, which converges to $(1 - \delta^{K(\varepsilon,\mu^*)})g_1(c, 0) + \delta^{K(\varepsilon,\mu^*)}g_1(c, y^*(c)) - \varepsilon$ as $K \to \infty$.*

The proof follows almost exactly the steps in Fudenberg and Levine (1989) (see also Mailath and Samuelson (2001, Ch.15) for a detailed exposition). For brevity we omit the proof, but provide the general idea. Player 2's best response is continuous (the singleton valued u.s.c. best response correspondence is a continuous function). Therefore, for any $\varepsilon > 0$, there exists $\eta(\varepsilon) > 0$ such that if player 2 believes that player 1 plays c with a probability more than $1-\eta(\varepsilon)$, then player 2's best response will be very close to $y^*(c)$. If that happens, player 1's payoff from playing c will be ε-close to $g_1(c, y^*(c))$. The bound is then constructed by considering a deviation by player 1 to playing $c > 0$ for the first K periods and then playing 0. Fudenberg and Levine (1989) show that if player 1 plays c constantly, there exists $K(\varepsilon, \mu^*)$ (it can be taken as $\frac{\ln \mu^*}{\ln \eta(\varepsilon)}$) such that there are no more than $K(\varepsilon, \mu^*)$ periods in which player 2 expects player 1 to play c with a probability less than $1-\eta(\varepsilon)$.[23] In the worst case scenario player 1 receives $g_1(c, 0)$ in each of these periods and these periods happen early. After that, for the next $K - K(\varepsilon, \mu^*)$ periods, player 1 is guaranteed to receive a payoff of at least $g(c, y^*(c)) - \varepsilon$. After the K periods (when player 2 can no longer observe the calendar time), player 1's payoff is at least $g_1(0, 0)$. That yields the bound.

Remark 1 *Lemma 2 is not satisfactory for two reasons.*

First, the bound comes strictly from the part of the game in which the information sets are exactly as in a standard game with full records. Hence, it does not tell us much about the impact of the finite records. To learn about that one may be more interested in bounding the equilibrium payoffs at $t \geq K$.

Second, the limit of this payoff bound for patient players depends crucially on the order of taking limits: $\lim_{\delta\to1}\lim_{K\to\infty} B(\delta, K) = g_1(c, y^*(c)) - \varepsilon$, *but* $\lim_{\delta\to1} B(\delta, K) = g_1(0, 0) - \varepsilon$ *for any* $K > K(\varepsilon, \mu^*)$. *Note that only the first limit matters in games with complete records where* $K = \infty$ *by assumption, but with finite records the second limit seems to be of interest as well.*

To address these two issues, we provide a tighter bound. We focus on PBE instead of Nash — this is necessary if one wants to bound any continuation payoffs, because there are Nash equilibria in which continuation payoffs are low off the equilibrium path.

To provide a tighter bound for the payoffs when $\delta \to 1$ for any fixed K, it is not sufficient to apply the methods from Fudenberg and Levine (1989). The main trick used to prove Lemma 2 is that along the path of observing a straight sequence of c, the number of periods in which player 2 anticipates an action far away from c is bounded above. However, with finite monitoring, when player 2 is in a state with index $I(h) = K$, he might believe that he is at period $t = K$ and

---

[23] The intuition is that for player 2 not to play close to $y^*(c)$, he must expect player 1 not to play c with a high enough probability. But that implies that after observing a sequence of c the belief player 2 assigns to the c type has to increase close to 1, making the best response close to $y^*(c)$.

$\mu_K < 1$ and player 1 will play 0 with a high probability. That would make player 2 unwilling to play an action close to $y^*(c)$.

Our new approach is to show that the posterior belief on the commitment type, μ_K, is *bounded from below* and the lower bound is close to 1 when K is large enough:

Lemma 3 *For any* $\eta \in (0,1)$, *whenever* $K > \frac{\ln \frac{\eta\mu^*}{1-\mu^*}}{\ln(1-\eta)} - 1$, *in any stationary PBE we have* $\mu_K > 1-\eta$.

The general idea behind the proof is as follows. Suppose to the contrary that $\mu_K \leq 1-\eta$. Then, as we have shown in Theorem 2, $\beta_k < \mu_K \leq (1-\eta)$, so from the history with index 0, the rational type of player 1 reaches the clean history (with index K) with a probability at most $\prod_{k=0}^{K-1} \beta_k < (1-\eta)^K$. But then, for high enough K seeing the clean history player 2 would assign a very high probability to type c, a contradiction.

Now, if $P(t)$ assigns positive probability to periods $\{K, ...2K-1\}$, the belief depends also on the strategies after the special histories (with indices $\varnothing, C^1, ...C^{K-1}$), but we have argued that $\beta_k > \beta\left(C^k\right)$ and $\beta_0 > \beta(\varnothing)$, so the same bound applies.

That lemma leads to our new bound:

Theorem 3 *For any* $\varepsilon > 0$ *and* $\mu^* \in (0,1)$, *there exists an integer* $K(\varepsilon,\mu^*)$ *independent of the equilibrium and the discount factor such that player* 1*'s payoff with finite records of length* $K > K(\varepsilon,\mu^*)$ *at* any time *is bounded from below by:*

$$(1-\delta^K)g_1(c,0) + \delta^K g_1(c,y^*(c)) - \varepsilon$$

which converges to $g_1(c,y^*(c)) - \varepsilon$ *as* δ *goes to* 1.

Proof. From the continuity of player 2's best response and the continuity of player 1's payoff function, for any $\varepsilon > 0$ there exists a $\eta(\varepsilon) > 0$ such that if player 2 believes that player 1 plays c with probability at least $1-\eta(\varepsilon)$, then player 1's payoff is ε-close to $g_1(c,y^*(c))$ this period.

Consider player 1's deviation to playing c constantly. If $\mu_K > 1-\eta(\varepsilon)$, the result is immediate: in the worst case, player 1 gets $g_1(c,0)$ in the first K periods, and stays in state K with a payoff of at least $g_1(c,y^*(c)) - \varepsilon$. If $\mu_K \leq 1-\eta(\varepsilon)$, then take $K(\varepsilon,\mu^*) = \ln \frac{\eta(\varepsilon)\mu^*}{1-\mu^*} / \ln(1-\eta(\varepsilon)) - 1$. The rest follows from Lemma 3. ■

It is worth noting that the record length K plays a dual role *on* the equilibrium path. It makes a clean history a convincing signal for the commitment type, but it also makes it harder to clean a history.

Summing up this section, we have shown that even though with finite records the equilibrium behavior is quite different from the complete records, if the records are long enough and player 1 is patient enough, he can still achieve payoffs close to $g_1(c, y^*(c))$. The difference from the game with complete records is that we are able to establish our payoff bound for any time in the game. In contrast, no such uniform-over-time bound has been proven for games in which short-lived players observe complete histories of play.

7 Conclusion

In this paper we have studied the impact of limited records on reputation dynamics. We have shown that limited records dramatically change the equilibrium behavior. Moreover, we have shown that under limited records, equilibria in games with complete and incomplete information are also quite different.

Our characterization stems from the interaction of submodularity of stage game payoffs and limited records. Investigating the implications of submodularity in more general dynamic setups can be an interesting direction for future research. Submodularity has found important applications in static games with sharp predictions. Athey and Schmutzler (2001) and Jun and Vives (2004) also explored its applications in the dynamic oligopoly context. Our work demonstrates the consequences of submodularity in repeated games.

Another interesting direction for future research is to consider games with imperfect monitoring: in our model the actions of player 1 are observed without noise, yet in some applications it is natural to ask what happens if the monitoring is imperfect. Some preliminary work we have done shows that in the model with an Honorable Type the strategies will depend in a more complex way on the history – it will no longer be sufficient to keep track only of the commitment index, since even a history containing some non-commitment actions is consistent with player 1 being the c type. Nevertheless, we conjecture that for sufficiently low noise in monitoring there would exist equilibria similar to the ones described in this paper (we have constructed examples for $K = 1$).[24] More importantly, we do not expect the result that reputations are short-lived (see Cripps, Mailath and Samuelson (2004)) to hold in our games with limited records (again, we have constructed counterexamples). The intuition is based on the results of this paper: even if the beliefs assigned to the c type get very low, player 1 will have opportunity to "clean up his history" by playing c for K periods (which with high probability should increase the beliefs). Under our assumptions on stage game payoffs, player 1 would have incentives to do so, if the probability of succeeding in cleaning up the reputation is sufficiently high.

Yet, it remains to be discovered what other new dynamic effects would limited records introduce in games with imperfect monitoring. One might also take a mechanism design approach to study the limited records. These richer models, though beyond the scope of the current

---

[24] Of course, it may depend on the way one models noisy monitoring.

work, might help to understand the reputation phenomenon observed in online auctions; see e.g., Cabral and Hortaçsu (2004) and Kato and Jin (2006).

8 Appendix A (Proofs)

Proof of Lemma 1. Let $g_2(\pi, y)$ be the expected payoff of player 2 when he plays y and expects player 1 to follow strategy π. Integrability implies that

$$\frac{\partial}{\partial y} g_2(\pi, y) = \frac{\partial}{\partial y} \int g_2(x, y)\pi(dx) = \int \frac{\partial}{\partial y} g_2(x, y)\pi(dx).$$

If π first order dominates π', then $\int \frac{\partial}{\partial y} g_2(x, y)\pi(dx) > \int \frac{\partial}{\partial y} g_2(x, y)\pi'(dx)$ because $\frac{\partial}{\partial y} g_2(x, y)$ is strictly increasing in x. That is $\frac{\partial}{\partial y} g_2(\pi, y) > \frac{\partial}{\partial y} g_2(\pi', y)$. Therefore, a result in Milgrom and Shannon (1994) implies that $y^*(\pi) \geq y^*(\pi')$. Since $y^*(\pi)$ and $y^*(\pi')$ are interior if π and π' are not degenerate at 0, we then have

$$\frac{\partial}{\partial y} g_2(\pi, y^*(\pi)) = 0 = \frac{\partial}{\partial y} g_2(\pi', y^*(\pi')).$$

If $y^*(\pi) = y^*(\pi')$, then we have a contradiction with the monotonicity of $\frac{\partial}{\partial y} g_2(\pi, y)$ in π. Therefore, $y^*(\pi) > y^*(\pi')$.[25] It remains to show that $y^*(\pi) > y^*(0) = 0$ for any π that does not put probability 1 on 0. Assume to the contrary that we can find π so that $y^*(\pi) = 0$. Then consider a compound lottery π' that is obtained by assigning probability $\frac{1}{2}$ on π and probability $\frac{1}{2}$ on 0. By definition, π first order dominates π'. But then we will have $0 = y^*(\pi) > y^*(\pi') \geq 0$. A contradiction.

As we remarked in the footnote in this lemma, interestingly, it is not sufficient to assume that g_2 is supermodular (instead of the differentiable version in the lemma). In that case we could only establish that the best response is weakly increasing in π. In particular, assume π first order dominates π' and $y > y'$. If $g_2(x, y)$ is supermodular, then $g_2(x, y) - g_2(x, y')$ is strictly increasing in x. Therefore, by the definition of first order domination,

$$\begin{aligned} & [g_2(\pi, y) - g_2(\pi, y')] - [g_2(\pi', y) - g_2(\pi', y')] \\ = & \int (g_2(x, y) - g_2(x, y'))\pi(dx) - \int (g_2(x, y) - g_2(x, y'))\pi'(dx) \\ > & \ 0. \end{aligned}$$

Therefore, $g_2(\pi, y)$ is supermodular in (π, y) and the result from Milgrom and Shannon (1994)

---

[25] This line follows from the proof of Edlin and Shannon (1998 Theorem 1).

implies that player 2's best response is non-decreasing in π. However, strict comparative statics are not guaranteed; see Athey, Milgrom and Roberts (1996) for a counterexample. ■

Proof of Theorem 1. We prove by induction that the continuation payoff of player 1 is independent of the history of the game. As a result, he plays a static best response in every stage, yielding the result (recall that by **Assumption 1** player 1 has a dominant strategy 0 in the stage game and by **Assumption 4** player 2's unique best response to $x = 0$ is $y = 0$).

The induction argument is based on a partition of the set histories. First, consider any two histories that differ only in what happened more than K periods ago – since the strategies of the short-lived players are τ^K-measurable, the payoff to player 1 has to be the same after any such two histories.

Second, consider any period $t \geq K$. Truncation functions τ^k partition H^t (for $t \geq K \geq k$) into sets of histories that differ only with respect to what happened k or more periods ago. For any $h, h' \in H^t$, we must have $\tau^k(h) = \tau^k(h')$ for some $k \in \{0, \cdots, K\}$ (note that $\tau^0(h) = \tau^0(h')$ for any h and h'). Therefore, our strategy is to show that if $V(h) = V(h')$ for any two histories that differ only with respect to what happened k or more periods ago, then it is also true for any two histories that differ only with respect to what happened $k-1$ or more periods ago. By induction that implies $V(h) = V(h')$ for any two histories.

Finally, it is straightforward to extend the claim for all periods $t < K$ using backward induction. We present the induction argument for $t \geq K > k$. The claim for $k = K$ is immediate since strategies of player 2 are τ^K-measurable.

CLAIM : *If for a given $k \in \{1, ...K\}$ it is true that for any two histories h, h' such that $\tau^k(h) = \tau^k(h')$ it holds that $V(h) = V(h')$ and $\sigma^t(h) = \sigma^t(h')$, then the same is true for $k-1$.*

Proof of Claim: Suppose there exist two histories h and h' such that $\tau^{k-1}(h) = \tau^{k-1}(h')$ and yet $\sigma^t(h) \neq \sigma^t(h')$. Suppose WLOG that $\sigma^t(h) > \sigma^t(h')$ (recall that by **Assumption 3** player 2 is playing a pure strategy in any equilibrium hence the inequality is well-defined; for the opposite inequality just reorder h and h').

We claim that this implies that there exists $d > 0$ such that $\pi(h)$ puts a positive probability on the interval $[d, 1]$. To see $\pi(h)([d, 1]) > 0$, note that if for any $d > 0$, $\pi(h)([d, 1]) = 0$, then we would have $\pi(h)(0) = 1$ and $\sigma^t(h) = 0$, contradicting $\sigma^t(h) > \sigma^t(h')$.[26] Let D be the set of such d.

We also claim that $\pi(h')([0, d)) > 0$ for some $d \in D$. Suppose to the contrary, $\pi(h')([d, 1]) = 1$ for any $d \in D$, then $\pi(h')$ first order stochastically dominates $\pi(h)$. By **Assumption 4,** we shall have $\sigma^t(h) \leq \sigma^t(h')$, contradicting the assumption we start with.

---

[26] In our proof we allow π to depend only on $\tau^K(h)$ and calendar time. In general, since player 1 observes the whole history of the game, π can depend on other aspects of the history. But for any equilibrium in which π depends on something more than $\tau^K(h)$ and calendar time, there exists another outcome-equivalent equilibrium in which π is measurable only with respect the information set of player 2. Hence our restriction is WLOG.

Combining these two observations, there exists $d > 0$ such that

$$\begin{aligned} \pi(h)([d,1]) &> 0 \\ \pi(h')([0,d)) &< 1. \end{aligned}$$

Take $d^* \in [d,1]$ from the support of $\pi(h)$ and $d_* \in [0,d)$ from the support of $\pi(h')$. The necessary incentive compatibility constraints for player 1 after history h are:

$$\begin{aligned} V(h) &= (1-\delta)g_1(d^*, \sigma^t(h)) + \delta V((h,d^*)) \\ &\geq (1-\delta)g_1(d_*, \sigma^t(h)) + \delta V((h,d_*)) \end{aligned}$$

(recall (h, d^*) is the history obtained by appending d^* to h). Hence

$$(1-\delta)(g_1(d^*, \sigma^t(h)) - g_1(d_*, \sigma^t(h))) \geq \delta(V((h,d_*)) - V((h,d^*))). \tag{4}$$

On the other hand, the IC constraint after history h' is:

$$\begin{aligned} V(h') &= (1-\delta)g_1(d_*, \sigma^t(h')) + \delta V((h',d_*)) \\ &\geq (1-\delta)g_1(d^*, \sigma^t(h')) + \delta V((h',d^*)). \end{aligned}$$

Therefore,

$$(1-\delta)(g_1(d^*, \sigma^t(h')) - g_1(d_*, \sigma^t(h'))) \leq \delta(V((h',d_*)) - V((h',d^*))). \tag{5}$$

Now, since $\tau^{k-1}(h) = \tau^{k-1}(h')$ we have $\tau^k((h',d_*)) = \tau^k((h,d_*))$ and $\tau^k((h',d^*)) = \tau^k((h,d^*))$. Therefore the RHS in (4) and (5) are equal. Combining these two inequalities we obtain:

$$g_1(d^*, \sigma^t(h')) - g_1(d_*, \sigma^t(h')) \leq g_1(d^*, \sigma^t(h)) - g_1(d_*, \sigma^t(h))$$

This contradicts **Assumption 2**.

Therefore $\sigma^t(h) = \sigma^t(h')$. That immediately implies $V(h) = V(h')$, finishing the proof of the claim.

Inductively, that implies that even if $\tau^0(h) = \tau^0(h')$ (which is true for all histories) in any equilibrium $\sigma^t(h') = \sigma^t(h)$ and $V(h) = V(h')$. Hence in any equilibrium player 1 plays a static best response. As we argued above, that implies a unique PBE which is the repetition of the unique static Nash equilibrium.

Remark 2 *The proof still goes through with the assumptions in footnote 15 if player 2's action* *set is binary ($Y = \{0, 1\}$) and mixed strategies are allowed. With σ^t a mixed strategy, we need reinterpret "$\sigma^t(h) > \sigma^t(h')$" according to first order stochastic dominance (i.e. $\sigma^t(h)$ puts higher weight on $y = 1$.) From the modification of Assumption 4, there exist $d^* \in [d, 1]$ from the support of $\pi(h)$ and $d_* \in [0, d)$ from the support of $\pi(h')$, where d is defined in footnote 15. The logic from inequality (4) onwards go through.*

■

8.1 Proof of Proposition 1

In this subsection we prove Proposition 1.

In any stationary PBE, equilibrium strategies for states in H^K_{++}, on and off the equilibrium path, depend only on $I(h)$. That is, if $I(h) = I(h')$, then $\sigma(h) = \sigma(h')$ and $\pi(h) = \pi(h')$.

Our proof utilizes the following auxiliary relation between histories, which we termed "n-similarity". In words, if h and h' are n-similar, then h and h' share the same most recent $n-1$ elements and their nth recent entries are different from the commitment action c. Let $L(h)$ be the length of the history h. Formally:

Definition 2 *For any $h, h' \in H^K_{++}$, we say that h and h' are n-similar, for $n = 1, ..., \min\{L(h), L(h')\}$, denoted by $h \sim^n h'$, if*
1) $\tau^{n-1}(h) = \tau^{n-1}(h')$ and
2) $\tau^n(h) \neq (\tau^{n-1}(h), c)$ and $\tau^n(h') \neq (\tau^{n-1}(h'), c)$.

For example, take $K = 5$:

1) $h = (0, 0, c, c, c)$ and $h' = (0, c, c, c, c)$ are not n-similar for any n (even though last 3 entries are the same, since the 4th most recent entry in h' is c they are not 4-similar).

2) $h = (0, 0, c, c, c)$ and $h' = (c, 0, c, c, c)$ are 4-similar

3) $h = (0, 0, c, 0, 0)$ and $h' = (0, 0, 0, 0, 0)$ are 1-similar and 2-similar

4) $h = (0, c)$ and $h' = (c, c, 0, 0, c)$ are 1-similar

We use the following relationship between the commitment index of the histories and their n-similarity:

Lemma 4 *For any $h, h' \in H^K_{++}$, if $I(h) = I(h')$, then either $h = h'$ or $h \sim^n h'$ for some $n \in \{1, 2, ..., K\}$.*

Proof. Obviously, if $h = h'$ then $I(h) = I(h')$. Suppose $I(h) = I(h') = i$ and $h \neq h'$. That implies $i < \min\{L(h), L(h')\} \leq K$ and that there exist some $x, x' \neq c$ such that $\tau^{i+1}(h) = (x, \underbrace{c, c, ..., c}_{i})$ and $\tau^{i+1}(h') = (x', \underbrace{c, c, ..., c}_{i})$. Therefore, by definition, $h \sim^{i+1} h'$. ■

Our next (main) step is to show that if two regular histories have the same index, then player 2 plays the same strategy after seeing these two strategies and player 1 payoffs in these two states are the same. The proof is similar to that of Theorem 1 and proceeds via induction on the n-similarity relation.

Lemma 5 *If $I(h) = I(h')$, then $\sigma(h) = \sigma(h')$ and $V(h) = V(h')$.*

Proof. By Lemma 4, we only need to show that if $h \sim^n h'$ for some $n \in \{1, 2, ..., K\}$, then $\sigma(h) = \sigma(h')$ and $V(h) = V(h')$. We do so by induction. Consider any stationary PBE, (π, σ, μ).

Step 1: If $h \sim^K h'$ then $\sigma(h) = \sigma(h')$ and $V(h) = V(h')$

Suppose to the contrary that $\sigma(h) \neq \sigma(h')$ and take WLOG $\sigma(h) > \sigma(h') \geq 0$.

Since player 1's type is known to player 2 after each of the two histories, (if two histories are n-similar, they both have to contain at least one action different from c) there exists $d > 0$ such that $\pi(h)([d, 1]) > 0$ (otherwise, we would have player 1 play 0 for sure and $\sigma(h) = 0$). Let D be the set of such d.

Next, it must be that $\pi(h')([0, d)) > 0$ for some $d \in D$. Suppose to the contrary, $\pi(h')([d, 1]) = 1$ for any $d \in D$, then $\pi(h')$ first order dominates $\pi(h)$. That would violate Assumption 4 and the assumption $\sigma(h) > \sigma(h')$.

Therefore, there exists $d > 0$ such that

$$\begin{aligned} \pi(h)([d,1]) &> 0 \\ \pi(h')([0,d)) &< 1. \end{aligned}$$

Take $d^* \in [d, 1]$ from the support of $\pi(h)$ and $d_* \in [0, d)$ from the support of $\pi(h')$. The IC constraint for player 1 in state h implies

$$(1-\delta)(g_1(d^*, \sigma(h)) - g_1(d_*, \sigma(h))) \geq \delta(V((h, d_*)) - V((h, d^*))), \tag{6}$$

where h, d^* is the history obtained by appending d^* to h and then dropping the oldest entry of h.

Similarly, the IC constraint in state h' implies:

$$(1-\delta)(g_1(d^*, \sigma(h')) - g_1(d_*, \sigma(h'))) \leq \delta(V((h', d_*)) - V((h', d^*))). \tag{7}$$

Note that $(h', d_*) = (h, d_*)$ and $(h', d^*) = (h, d^*)$ since $h \sim^K h'$ (since the two histories are K-similar, after appending the same action today and dropping the latest action, they become identical). Combining (6) and (7), we get:

$$g_1(d^*, \sigma(h')) - g_1(d_*, \sigma(h')) \leq g_1(d^*, \sigma(h)) - g_1(d_*, \sigma(h))$$

which contradicts Assumption 2.

Therefore $\sigma(h) = \sigma(h')$ if $h \sim^K h'$. It follows immediately that $V(h) = V(h')$.

Step 2: Assume for any $k \geq 2$, that if $h \sim^k h'$, then $V(h) = V(h')$ and $\sigma(h) = \sigma(h')$. We claim that implies that the same is true for $k-1$.

Suppose $h \sim^{k-1} h'$, but $\sigma(h) \neq \sigma(h')$. Then either $\sigma(h) > \sigma(h')$ or $\sigma(h') > \sigma(h)$. Assume the former WLOG. Following the argument in Step 1, we obtain (6) and (7). But since $(h', d_*) \sim^k (h, d_*)$ and $(h', d^*) \sim^k (h, d^*)$, we again derive a contradiction with Assumption 2 using the induction assumption. ■

As we argued in the text, in any equilibrium player 1 plays only either c or 0 after any history (not only after regular histories)

Lemma 6 *In any equilibrium after any history, player* 1 *plays with positive probability only* 0 *or* c*.*

Proof. Assume that player 1 plays a strategy $\pi(h)$ that puts positive probability on $(0, c) \cup (c, 1]$. Then, player 2's best response is $\sigma(h) > 0$ by Assumption 4. Now consider player 1's intertemporal incentive at h. For any $x \in (0, c) \cup (c, 1]$ in the support of $\pi(h)$,

$$\begin{aligned} V(h) &= (1-\delta) g_1(x, \sigma(h)) + \delta V((h, x)) \\ &< (1-\delta) g_1(0, \sigma(h)) + \delta V((h, x)) \\ &= (1-\delta) g_1(0, \sigma(h)) + \delta V((h, 0)) \end{aligned}$$

where the last equality follows from Lemma 5 because $I((h, x)) = I((h, 0))$. Therefore, x is not a best response. ■

Finally, we can pin down the strategy of player 1 :

Lemma 7 $\pi(h) = \pi(h')$ *if* $I(h) = I(h')$*.*

Proof. From Lemma 5, if $I(h) = I(h')$ then $\sigma(h) = \sigma(h')$. From Lemma 6, player 1 only plays 0 and c. If $\pi(h) \neq \pi(h')$ then $\pi(h)$ and $\pi(h')$ can be ranked according to FOSD. Therefore, when $I(h) < K$, $\sigma(h) \neq \sigma(h')$, a contradiction. ■

This finishes the proof of Proposition 1.

8.2 Proof of Theorem 2

In this subsection we prove the general characterization Theorem 2 (the structure of all stationary PBE with sufficiently patient players).

Since Proposition 1 established that equilibrium strategies depend only on the index of the history, we now write the continuation payoffs $V(i)$ as a function of the index alone. Recall that we now denote by $y^*(\beta)$ player 2 best response if he expects player 1 will play c with probability β and 0 with probability $(1-\beta)$ and we denote by y_k the equilibrium strategy of player 2 as a function of the history index k.

Our first lemma and corollary establish that in state K (when player 2 observes a clean history consisting of K observations of c), the rational type of player 1 plays 0 for sure:

Lemma 8 *If player* 1 *weakly prefers action* c *in state* K, *then in each state* $i = 0, 1, ..., K-1$, *we have* (1) *player* 1 *weakly prefers action* c, (2) $y_i = y_K$, (3) $V(i) = V(K)$.

Proof. Since player 1 weakly prefers action c in state K, $V(K) = g_1(c, y_K)$. Assume by induction that for $i = k+1, ..., K$, the three properties hold. Consider $i = k$.

Step 1: We first show that $y_k \leq y_{k+1} = y_K$.

Suppose to the contrary that $y_k > y_{k+1}$. Let $r_k(x)$ be player 1's payoff from playing x in state k once and then returning to the equilibrium strategy. These payoffs are:

$$\begin{aligned} r_k(0) &= (1-\delta)g_1(0, y_k) + \delta V(0) \\ r_k(c) &= (1-\delta)g_1(c, y_k) + \delta V(K) \end{aligned}$$

and analogously

$$\begin{aligned} r_{k+1}(0) &= (1-\delta)g_1(0, y_{k+1}) + \delta V(0) \\ r_{k+1}(c) &= (1-\delta)g_1(c, y_{k+1}) + \delta V(K) \end{aligned}$$

Note that $y_k > 0$ by the assumption $y_k > y_{k+1}$. Thus $r_k(0) - r_k(c) \leq 0$. Therefore

$$\begin{aligned} & [r_{k+1}(0) - r_{k+1}(c)] - \underbrace{[(r_k(0) - r_k(c)]}_{\leq 0} \\ = \; & (1-\delta)\{[g_1(0, y_{k+1}) - g_1(c, y_{k+1})] - [g_1(0, y_k) - g_1(c, y_k)]\} \\ < \; & 0 \end{aligned}$$

where the inequality follows from Assumption 2 and $y_k > y_{k+1}$. That requires $r_{k+1}(0) - r_{k+1}(c) < 0$.

But then $\beta_{k+1} = 1$ which contradicts $y_k > y_{k+1}$ (because y_k is never higher than $y^*(c) = y^*(1) = y_{k+1}$).

Step 2: We finish the induction.

We either have $\beta_k = 1$ or $\beta_k < 1$. If $\beta_k = 1$, then $y_k = y_{k+1}$ because $\beta_k = 1$ implies $y_k = y^*(c)$, Step 1 implies $y_{k+1} \geq y_k$ and $y_j \leq y^*(c)$ for any j. Properties (1) and (3) hold in this case as well.

Now, suppose $\beta_k < 1$ and $y_k < y_{k+1}$, then from calculations in Step 1,

$$[r_{k+1}(0) - r_{k+1}(c)] - [r_k(0) - r_k(c)] > 0. \tag{8}$$

Now, $\beta_k < 1$ requires $r_k(0) \geq r_k(c)$ (0 is a best response in state k). To satisfy inequality (8) we require $r_{k+1}(0) > r_{k+1}(c)$, i.e., player 1 strictly prefers action 0 in state $k+1$, contradicting the induction hypothesis.

Finishing up, we have shown $y_k = y_{k+1}$. Using calculations in step 1 it yields $r_k(0) - r_k(c) = r_{k+1}(0) - r_{k+1}(c)$. In turn, that implies that player 1 weakly prefers action c in state k and obviously property (3) holds as well. ■

Corollary 1 *Player* 1 *strictly prefers action* 0 *in state* K *and* $\beta_K = 0$.

Proof. If player 1 weakly prefers action c, then player 2's strategies do not depend on the history of the game, according to the previous lemma (and since $\mu_K \geq \mu^*$, $y>0$ in every period). As a result, player 1 would strictly prefer action 0, a contradiction. ■

The next step is to show that y_k is 0 or $y_k - y_{k-1} > 0$ and analogously the payoffs are weakly increasing in k.

Lemma 9 *Player 2's strategy* $\{y_i\}_{i=0}^K$ *and player 1's payoff* $\{V(i)\}_{i=0}^K$ *are weakly increasing in* i. *Moreover, if for some* $j < K$, $y_{j+1} > 0$, *then they are strictly increasing for all* $i \geq j$. *As a result, for any* $i < K$, $y_i < y^*(c)$ *and* $\beta_i < 1$.

Proof. We first show that $y_{K-1} < y_K$.

There are two possibilities: $y_{K-1} = 0$ and $y_{K-1} > 0$. In the first case, immediately $y_{K-1} < y_K$ since $y_K \geq y^*(\mu^*)$. Also, immediately $V(K) > V(K-1)$ (since $y_{K-1} = 0 \Longleftrightarrow \beta_{K-1} = 0$ and we have $y_K > 0$ and $\beta_K = 0$).

Now consider $y_{K-1} > 0 \Longleftrightarrow \beta_{K-1} > 0$

Since player 1 strictly prefers action 0 in state K and weakly prefers c in state $K-1$, we have

$$\begin{aligned}
(1-\delta)g_1(0, y_K) + \delta V(0) &> (1-\delta)g_1(c, y_K) + \delta V(K) \\
(1-\delta)g_1(c, y_{K-1}) + \delta V(K) &\geq (1-\delta)g_1(0, y_{K-1}) + \delta V(0)
\end{aligned}$$

Summing up the two inequalities, we have $g_1(0, y_K) - g_1(c, y_K) > g_1(0, y_{K-1}) - g_1(c, y_{K-1})$. By Assumption 2, we have $y_K > y_{K-1}$, as claimed. Moreover,

$$\begin{aligned} V(K) &> (1-\delta)g_1(c, y_K) + \delta V(K) \\ &\geq (1-\delta)g_1(c, y_{K-1}) + \delta V(K) \\ &= V(K-1). \end{aligned}$$

which establishes the claim for K and $K-1$.

Now consider any i and suppose $y_i = 0$ for some i. Let $i^* < K$ be the highest such i. We claim that for all $i < i^*$, y_i is zero as well. Suppose not. That implies there are two states i and $i+1$ such that $\beta_i > 0$ and $\beta_{i+1} = 0$ and $i+1 < K$. That means that $V(i+1) = V(0)$ and the incentive compatibility constraint is:

$$(1-\delta)g_1(c, y_i) + \delta V(i+1) \geq (1-\delta)g_1(0, y_i) + \delta V(0)$$

which cannot be satisfied (if player 1 plays 0 for sure next period and this period $y > 0$, he prefers to speed up playing 0).

Now suppose that $y_i > 0$ and $y_i - y_{i-1}$ is not strictly positive. Let i^* be the largest state in which $y_{i^*-1} \geq y_{i^*} > 0$. By the result for K and $K-1$, we know that $1 \leq i^* < K$. It must also be the case that $\beta_{i^*} < 1$. Otherwise, $y_{i^*} \geq y_{i^*+1}$, contradicting the definition of i^*. Since $\beta_{i^*} > 0$ and $\beta_{i^*-1} > 0$,

$$\begin{aligned} V(i^*-1) &= (1-\delta)[g_1(c, y_{i^*-1}) + \delta g_1(c, y_{i^*}) + \cdots + \delta^{K-i^*-1} g_1(c, y_{K-2})] + \delta^{K-i^*} V(K-1) \\ &\geq (1-\delta)g_1(0, y_{i^*-1}) + \delta V(0) \end{aligned}$$

and

$$\begin{aligned} V(i^*) &= (1-\delta)[g_1(c, y_{i^*}) + \delta g_1(c, y_{i^*+1}) + \cdots + \delta^{K-i^*-1} g_1(c, y_{K-1})] + \delta^{K-i^*} V(K) \\ &= (1-\delta)g_1(0, y_{i^*}) + \delta V(0) \end{aligned}$$

The last equality follows because $0 < \beta_{i^*} < 1$.

Subtracting the second expression from the first, and re-arranging terms, we have,

$$\begin{aligned} &\sum\nolimits_{k=i^*+1}^{K-1} \delta^{k-i^*}[g_1(c, y_{k-1}) - g_1(c, y_k)] + \delta^{K-i^*}[V(K-1) - V(K)] \\ \geq\ &[g_1(0, y_{i^*-1}) - g_1(0, y_{i^*})] - [g_1(c, y_{i^*-1}) - g_1(c, y_{i^*})] \end{aligned}$$

For each $k > i^*$, we have $y_{k-1} \leq y_k$, and hence $g_1(c, y_{k-1}) - g_1(c, y_k) \leq 0$. Furthermore, $V(K-1) <$

$V(K)$ as we have shown above. Therefore, from the above inequality,

$$g_1(0, y_{i^*-1}) - g_1(c, y_{i^*-1}) < g_1(0, y_{i^*}) - g_1(c, y_{i^*}).$$

By Assumption 2, $y_{i^*-1} < y_{i^*}$, contradicting the definition of i^*. So it must be that $y_i > y_{i-1}$ whenever $y_i > 0$.

When y_i is strictly increasing, $V(i) = (1-\delta)g_1(0, y_i) + \delta V(0)$ is also strictly increasing.

Since $y_K \leq y^*(c)$, it must be that for all k, $y_k < y^*(c)$ and hence $\beta_k < 1$. ■

So we have established that in any stationary PBE, in state K player 1 plays 0 and in all other states he plays 0 with positive probability. The next lemma shows that he must also play c with positive probability in all states $i < K$:

Lemma 10 *If $\delta > \bar{\delta}$, then $\beta_i > 0$ for each $i < K$. In words, player 1 plays the commitment action c with positive probability if $i < K$.*

Proof. Assume to the contrary that the statement is not true, i.e., player 1 plays 0 with probability 1 in some states lower than K. Let i^* be the smallest such state at which player 1 plays 0 with probability 1.

We first show $i^* = 0$. Suppose instead $i^* > 0$. By definition, $i^* < K$. Then player 2 plays 0 at i^* because his belief is $\mu_{i^*} = 0$. Player 1 payoffs at states $i = 0, 1, ..., i^* - 1$ and i^* are

$$\begin{aligned} V(i) &= (1-\delta)g_1(c, y_i) + \delta V(i+1), i = 0, 1, ..., i^* - 1. \\ V(i^*) &= (1-\delta)g_1(0, 0) + \delta V(0). \end{aligned}$$

Note that

$$\begin{aligned} V(i^*-1) &= (1-\delta)g_1(c, y_{i^*-1})) + \delta V(i^*) \\ &< (1-\delta)g_1(0, y_{i^*-1})) + (1-\delta)\delta g_1(0,0) + \delta^2 V(0) \\ &< (1-\delta)g_1(0, y_{i^*-1})) + (1-\delta)\delta g_1(0, y_0)) + \delta^2 V(0) \\ &\leq (1-\delta)g_1(0, y_{i^*-1})) + \delta V(0) \end{aligned}$$

where the first and second inequalities follow from the monotonicity of g_1 given $y_i > 0$; the third inequality follows from the incentive constraints in state 0. Therefore, we conclude that player 1 should play action 0 instead of action c in state $i^* - 1$, contradicting $i^* > 0$. Therefore, $i^* = 0$.

Second, suppose that $i^* = 0$, which means $\beta_0 = 0$ and implies $y_0 = 0$. Therefore, $V(0) = g_1(0, 0)$.

CASE 1: Suppose $y_K = y^*(c)$.

Note that if $\delta > \frac{g_1(0,y^*(c)) - g_1(c,y^*(c))}{g_1(0,y^*(c)) - g_1(0,0)} = \bar{\delta}$ then

$$(1-\delta)g_1(0, y^*(c)) + \delta g_1(0,0) < g_1(c, y^*(c)).$$

The LHS is player 1's payoff if he plays 0 in state K and the right-hand side is his payoff if he plays c forever. Therefore, player 1 has a unique best response at state K - to play c.

CASE 2: Suppose $y_K < y^*(c)$.

That implies that there exists t such that $P(t) > 0$ and with positive probability player 1 plays c in periods $(t-1, ...t-K)$ and at t, if the index is K, player 1 plays 0 with positive probability. If $\beta_0 = 0$, as we assumed, that can happen only if player 1 plays c with positive probabilities in periods $0, ...K-1$ and then plays 0 with positive probability.

Consider history $h_c^n = (c, c, ..., c) \in H^n$ for $n < K$ which has index C^n. That is a history in period $t = n < K$ in which player 1 has played c in all the periods so far. We claim that if $\beta_0 = 0$ then in period 0 player 1 plays 0 for sure (and hence we have a contradiction: a history with index K is never reached by type r of player 1).

We show this claim by induction. First, consider history with index C^{K-1}. Since player 1 plays c with positive probability after that history, his IC constraint is:

$$(1-\delta)g_1(c, \sigma(C^{K-1})) + \delta V(K) \geq (1-\delta)g_1(0, \sigma(C^{K-1})) + \delta V(0)$$

We combine it with the IC constraint in state $K-1$ (which is a history for $t > K$). Since, as we have shown above, in state $k < K$ player 1 plays 0 with positive probability:

$$(1-\delta)g_1(0, y_{K-1})) + \delta V(0) \geq (1-\delta)g_1(c, y_{K-1})) + \delta V(K)$$

Adding these two IC constraints together, and re-arranging terms, we obtain

$$g_1(0, y_{K-1}) - g_1(c, y_{K-1}) \geq g_1(0, \sigma(C^{K-1})) - g_1(c, \sigma(C^{K-1}))$$

Using Assumption 2 (submodularity), this implies:

$$\begin{aligned} y_{K-1} &\geq \sigma(C^{K-1}) \\ V(K-1) &\geq V\left(C^{K-1}\right) \end{aligned}$$

We follow by induction. Suppose that for all $\{n, n+1, ...K-1\}$ we have

$$\begin{aligned} y_n &\geq \sigma(C^n) \\ V(n) &\geq V(C^n) \end{aligned}$$

then the IC constraints in states C^{n-1} and $k = n-1$ are (recall that $\beta_n < 1$ and we supposed $\beta_{C^n} > 0$):

$$(1-\delta)g_1(c, \sigma(C^{n-1})) + \delta V(C^n) \geq (1-\delta)g_1(0, \sigma(C^{n-1})) + \delta V(0)$$

$$(1-\delta)g_1(0, y_{n-1})) + \delta V(0) \geq (1-\delta)g_1(c, y_{n-1})) + \delta V(n)$$

where we used the result that even for $t \in \{1, K-1\}$, if player 1 plays any action other than c, the continuation equilibrium depends only on the commitment index of the history, which after the first non-c action is 0.

Adding these two IC constraints together, and re-arranging terms, we obtain:

$$g_1(0, y_{n-1}) - g_1(c, y_{n-1}) \geq g_1(0, \sigma(C^{n-1})) - g_1(c, \sigma(C^{n-1}))$$

and again by Assumption 2 this implies

$$\begin{aligned} y_{n-1} &\geq \sigma(C^{n-1}) \\ V(n-1) &\geq V\left(C^{n-1}\right) \end{aligned}$$

Going all the way to $n = 1$, the next iteration establishes a bound on the strategy of player 1 at the beginning of the game ($t = 0$): he plays c with probability at most β_0. But that leads to a contradiction: we cannot have $\beta_0 = 0$ and $y_K < y^*(c)$, since then the only type ever reaching history with index K would be the commitment type, which would imply $\mu_K = 1$ and $y_K = y^*(c)$. ■

Corollary 2 *$\{y_i\}_{i=0}^{K}$ and $\{\beta_i\}_{i=0}^{K-1}$ are strictly increasing and $0 < \beta_i < 1$, $i = 0, 1, ..., K-1$.*

Proof. We have shown that $\beta_i > 0$ for all $i \in \{0, ...K-1\}$. Hence $y_i > 0$ and by Lemma 9 $\{y_i\}_{i=0}^{K}$ is a strictly increasing sequence. Since $\{y_i\}_{i=0}^{K}$ is strictly increasing, $\{\beta_i\}_{i=0}^{K-1}$ must be strictly increasing. Since $y_{K-1} < y_K$, we have $\beta_{K-1} < 1$. $\beta_i > 0$ is shown in a previous lemma. ■

Corollary 3 *$\{V(i)\}_{i=0}^{K}$ is strictly increasing.*

Proof. We have already shown that $V(K-1) < V(K)$ in Lemma 9. For each $i = 0, 1, ..., K-1$, we have $0 < \beta_i < 1$ and hence

$$V(i) = (1-\delta)g_1(0, y_i) + \delta V(0)$$

Since y_i is strictly increasing, $V(i)$ is strictly increasing. ■

Proof of Proposition 2. This result follows from Theorem 2. In state $k \leq K-1$, player 1 is indifferent between (1) playing 0 immediately, returning to state 0, and staying there by playing 0 repeatedly, and (2) playing 1 once, moving to state $k+1$, playing 0 there, and returning to state 0. The indifference condition is

$$(1-\delta)g_1(0, y_k)) + \delta g_1(0, y_0)) = (1-\delta)g_1(c, y_k)) + \delta(1-\delta)g_1(0, y_{k+1})) + \delta^2 g_1(0, y_0))$$

which simplifies to Eq. (1). Equation (2) is immediate from the Bayes' rule. ■

8.3 Proof of Payoff Bound

Proof of Lemma 3. Assume instead $\mu_K \leq 1-\eta$. In Theorem 2 we have established that in any stationary PBE, $\beta_k < \mu_K$ and $\beta\left(C^k\right) < \beta_k$ for the initial histories (see Section 5.3).

Hence, the probability of observing a history with index K in period $t \geq K$, conditional on player 1 being the r type is bounded above by:

$$\prod_{k=0}^{K-1} \beta_k < \mu_K^K \leq (1-\eta)^K$$

We now use the Bayes rule to show that this contradicts $\mu_K \leq 1-\eta$ for high enough K. Write C^K as the history $(c, c, ..., c) \in H^K$ (the history with index K). Using the notation in the derivation of $\mu(\theta|h)$ in Section 9 below, $\Pi^t((\text{Proj})^{-1}(C^K)) < (1-\eta)^K$ for $t \geq K$ by our analysis above. Therefore,

$$\begin{aligned} Q(\{c\} \times X^{t-K} \times (\text{Proj})^{-1}(C^K)) &= \mu^*, \\ Q(\{r\} \times X^{t-K} \times (\text{Proj})^{-1}(C^K)) &< (1-\mu^*)(1-\eta)^K. \end{aligned}$$

The joint probability of player 1 being the c type and player 2 observing the C^K history, denoted $\mu(c, C^K)$, is:

$$\mu(c, C^K) = \mu^* \sum\nolimits_{t=K}^{\infty} P(t),$$

We also get:

$$\begin{aligned}\mu(r|C^K) &= \sum\nolimits_{t=K}^{\infty} P(t) Q(\{r\} \times X^{t-K} \times (\text{Proj})^{-1}(C^K)) \\ &\leq (1-\mu^*) \sum\nolimits_{t=K}^{\infty} P(t)(1-\eta)^K\end{aligned}$$

Therefore

$$\begin{aligned}\mu_K &= \mu(c|C^K) \geq \frac{\mu^* \sum_{t=K}^{\infty} P(t)}{\mu^* \sum_{t=K}^{\infty} P(t) + (1-\mu^*) \sum_{t=K}^{\infty} P(t)(1-\eta)^K} \\ &= \frac{\mu^*}{\mu^* + (1-\mu^*)(1-\eta)^K}\end{aligned}$$

Therefore, $\mu_K > 1-\eta$ whenever $K > \frac{\ln \frac{\eta\mu^*}{1-\mu^*}}{\ln(1-\eta)} - 1$, a contradiction. ■

9 Appendix B (Formal Derivation of $\mu(\theta|h)$)

Let $\Omega = \Theta \times X^{\infty}$ be the set of all possible outcomes. For example, a state $(\theta, x^0, x^1, ...) \in \Omega$ describes a situation where player 1's type is θ, and player 1's period-$t \geq 0$ play is x^t. We consider the σ-algebra on Ω generated by $(t+1)$-dimensional calendar sets $\{\theta\} \times B_0 \times B_1 \times \cdots \times B_{t-1}$, where $\theta \in \Theta$ and $\{B_t\}$ is a sequence of Borel measurable subsets of X. For each history $h = (x^0, x^1, ..., x^{t-1}) \in H^t$, we identify h with the subset $\{(\hat{x}^0, \hat{x}^1, ...) \in X^{\infty} : \hat{x}^0 = x^0, \hat{x}^1 = x_1, ..., \hat{x}^{t-1} = x^{t-1}\}$. Denote this set by $(\text{Proj})^{-1}(h)$.

Step 1. We first define the probability measure Q on Ω induced by $\pi : H \to \Delta X$ and μ^*.

The rational player 1's strategy $\pi : H \to \Delta(X)$ induces a probability measure Π^t over measurable rectangles $B_0 \times B_1 \times \cdots \times B_t$ of the $(t+1)$-dimensional space X^{t+1}, $t \geq 0$, inductively as follows.

$$\begin{aligned}\Pi^0(B_0) &= \pi(\varnothing)(B_0) \\ \Pi^t(B_0 \times B_1 \times \cdots \times B_t) &= \int_{B_0 \times B_2 \times \cdots \times B_{t-1}} \pi(x^0, x^1, ..., x^{t-1})(B_t) \; d\Pi^{t-1}\end{aligned}$$

By Kolmogorov Extension Theorem (See Billingsley (1995)), the sequence of measures $\{\Pi^t\}$ extends uniquely to a probability measure Π over X^{∞}. Then Π and μ^* induce $Q \in \Delta(\Omega)$ as follows. For any measurable set $A \subset X^{\infty}$.

$$\begin{aligned}Q(\{c\} \times A) &= \begin{cases} \mu^*, & \text{if } (c, c, ...) \in A \\ 0, & \text{otherwise} \end{cases} \\ Q(\{r\} \times A) &= (1-\mu^*)\Pi(A)\end{aligned}$$

In words, Q is an outside observer's belief about the outcomes if player 1 plays π.

Step 2. We derive player 2's ex ante belief μ over $\Theta \times \tau^K(H)$ before entering the game.

Formally, μ is induced by Q and P as follows. For each $\theta \in \Theta$, if $A \subset \cup_{t=0}^{K-1} H^t \subset \tau^K(H)$, then

$$\mu(\{\theta\} \times A) = P(t) Q(\{\theta\} \times (\mathrm{Proj})^{-1}(A)),$$

and if $A \subset H^K \subset \tau^K(H)$, then

$$\mu(\{\theta\} \times A) = \lim_{n \to \infty} \sum_{t=K}^{K+n} P(t) Q(\{\theta\} \times X^{t-K} \times (\mathrm{Proj})^{-1}(\tau^K(A)))$$

Step 3. After player 2 enters the game and observes a truncated history $h \in \tau^K(H) = \cup_{t=0}^{K} H^t$, he updates his belief on the type of player 1. Let's denote this posterior belief by $\mu(\theta|h)$. We take $\mu(\theta|h)$ as a version of the conditional probability $\mu(\{\theta\} \times \tau^K(H))|\Theta \times \{h\})$ such that $\mu(c|h) = 0$ if $\mu((c, h)) = 0$.

Remark 3 *In parts of the paper we use the special case of P being the improper uniform prior. In computing beliefs in this case, only Step 2 needs to be refined. For each $\theta \in \Theta$, if $A \subset \cup_{t=0}^{K-1} H^t$, then*

$$\mu(\{\theta\} \times A) = \frac{1}{K} Q(\{\theta\} \times (\mathrm{Proj})^{-1}(A)),$$

For $A \subset H^K \subset \tau^K(H)$, define

$$\mu(\{\theta\} \times A) = \lim_{n \to \infty} \frac{1}{n} \sum_{t=K}^{K+n} Q(\{\theta\} \times X^{t-K} \times (\mathrm{Proj})^{-1}(\tau^K(A)))$$

if the limit exists (and μ is arbitrary otherwise).[27]

References

[1] Athey, S., P. Milgrom, P., and J. Roberts (1996), *Robust Comparative Statics,* Manuscript, Harvard University.

[2] Athey, S., and A. Schmutzler (2001), *Investment and market dominance,* RAND Journal of Economics, 32(1): 1–26.

[3] Barro, R., and D. B. Gordon (1983), *Rules, Discretion and Reputation in a Model of Monetary Policy,* Journal of Monetary Economics, 12(1): 101-121.

---

[27] The existence is not guaranteed for an arbitrary strategy π. Alternatively, we can define μ from the ergodic limit of the Markov process induced by Π on H^K with the initial measure Π^{K-1}. The ergodic limit need not exist because H^K is a continuum. (See, Stokey and Lucas, (1989), Theorem 11.9). However, we show that for equilibrium π the limit exists and coincides with the unique invariant measure.

[4] Bhaskar, V., (1998), *Informational Constraints and the Overlapping Generations Model: Folk and Anti-Folk Theorems*, Review of Economic Studies, 65(1):135-49l.

[5] Billingsley, P., (1995), *Probability and Measure,* Wiley, New York .

[6] Bjork, R.A, and W. B. Whitten (1974), *Recency-Sensitive Retrieval Process in Long-Term free Recall*, Cognitive Psychology 6:173-189.

[7] Broadbent, D. E., and M. H. P Broadbent (1981), *Recency effects in visual memory,* The Quarterly Journal of Experimental Psychology, 33A, 1-15.

[8] Cabral, L. M. B., and A. Hortaçsu (2004) *The Dynamics of Seller Reputation: Theory and Evidence from eBay,* NBER Working Paper No. W10363.

[9] Cole, H., and N, Kocherlakota (2005), *Finite Memory and Imperfect Monitoring*, Games and Economic Behavior, 53(1): 59-72.

[10] Cripps, M., G. Mailath and L. Samuelson (2004), *Imperfect Monitoring and Impermanent Reputations,* Econometrica, 72: 407-432.

[11] Ekmekci, M. (2006) *Sustainable Reputations with Rating Systems*, mimeo.

[12] Faingold, E. and Y. Sannikov (2007), *Reputation Effects and Equilibrium Degeneracy in Continuous-Time Games*, Cowles Foundation Discussion Paper No. 1624.

[13] Fudenberg, D., D. M. Kreps and E. S. Maskin (1990), *Repeated Games with Long-run and Short-run Players*, Review of Economic Studies, 57(4): 555-73.

[14] Fudenberg, D., and D. Levine (1989), *Reputation and Equilibrium Selection in Games with a Single Long-Run Player*, Econometrica, 57:759-778.

[15] Ghosh, P., and D. Ray (1996), *Cooperation in Community Interaction without Information Flows,* Review of Economic Studies, 63(3): 491-519.

[16] Ghosh, P., and D. Ray (2001), *Information and Enforcement in Informal Credit Markets*, Working paper, New York University.

[17] Hörner, J. and W. Olszewski (2008), *How Robust is the Folk Theorem with Imperfect Public Monitoring?* Quarterly Journal of Economics, forthcoming.

[18] Jehiel, P. (1995), *Limited Horizon Forecast in Repeated Alternate Games*, Journal of Economic Theory, 67(2): 497-519.

[19] Jin, G. Z., and A. Kato (2006), *Price, Quality and Reputation: Evidence From An Online Field Experiment*, RAND Journal of Economics, 37(4):983-1005.

[20] Jun, B., and X. Vives (2004), *Strategic incentives in dynamic duopoly,* Journal of Economic Theory, 116(2): 249-281.

[21] Kreps, D. M, P. Milgrom, J. Roberts, and R. Wilson (1982), *Rational Cooperation in the Finitely Repeated Prisoners' Dilemma*, Journal of Economic Theory, 27: 245-52.

[22] Kreps, D. M., and R. Wilson (1982), *Reputation and Imperfect Information*, Journal of Economic Theory, 27: 253-79.

[23] Liu, Q. (2009), *Information Acquisition and Reputation Dynamics*, Working paper, University of Pennsylvania.

[24] Ljungqvist, L., and T. J. Sargent (2004), *Recursive Macroeconomic Theory*, 2nd Edition. MIT Press, Cambridge, MA.

[25] Mailath, G. J., and S. Morris (2002), *Repeated Games with Almost-Public Monitoring*, Journal of Economic Theory, 102(1): 189-228.

[26] Mailath, G. J., and S. Morris (2006), *Coordination failure in repeated games with almost-public monitoring*, Theoretical Economics, 1(3): 311-340.

[27] Mailath, G. J., and L. Samuelson (2001), *Who Wants a Good Reputation?* Review of Economic Studies 68: 415-441.

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

[29] Mailath, G. J., and W. Olszewski (2008), *Folk Theorems with Bounded Recall under (Almost) Perfect Monitoring*, PIER Working Paper Archive 08-027, University of Pennsylvania.

[30] Milgrom, P., and J. Roberts (1982), *Predation, reputation, and entry deterrence*, Journal of Economic Theory 27: 280–312.

[31] Murdock, B.B., Jr. (1962) *The Serial Position Effect of Free Recall*, Journal of Experimental Psychology, 64, 482-488.

[32] Phelan, C. (2006), *Public Trust and Government Betrayal*, Journal of Economic Theory 130:27–43.

[33] Rauch, J., and J. Waston (2003), *Starting Small in an Unfamiliar Environment*, International Journal of Industrial Organization, 21: 1021-1042.

[34] Stokey, N. (1989), *Reputation and Time Consistency*, The American Economic Review, 79(2), 134-39.

[35] Stokey, N., and R. Lucas (1989), *Recursive Methods in Economic Dynamics*, Havard University Press.

[36] Tadelis S. (1999), *What's in a Name? Reputation as a Tradeable Asset*, American Economic Review, 89(3):548-563.

[37] Topkis, D. M. (1998), *Supermodularity and Complementarity*, Princeton: Princeton University Press.

[38] Waston, J. (1999), *Starting Small and Renegotiation*, Journal of Economic Theory, 8: 52-90.

[39] Waston, J. (2002), *Starting Small and Commitment*, Games and Economic Behavior, 38: 176-199.

[40] Wiseman, T. (2008) *Reputation and impermanent types,* Games and Economic Behavior, 62(1), 190-210.