Reputation Effects with Endogenous Records

Harry PEI*

First Draft: August 31, 2023. This Draft: April 1, 2024

Abstract: A patient player interacts with a sequence of short-run players. The patient player is either an *honest type* who always takes a commitment action and never erases any record, or an *opportunistic type* who decides which action to take and whether to erase his action at a low cost. I show that in every equilibrium, the patient player will take the commitment action with positive probability until he has accumulated a long enough good record, at which point he can secure a payoff that is strictly greater than his commitment payoff. However, as long as the patient player is sufficiently long-lived, his equilibrium payoff must be close to his minmax value. Although a tiny probability of opportunistic type can wipe out all of the patient player's returns from building reputations, it only has a negligible effect on the short-run players' welfare.

Keywords: record length, erasing records, reputation effects, reputation failure.

1 Introduction

In most of the existing works on repeated games and reputations, the lengths of players' records are *exogenous*. These include the models of Fudenberg and Maskin (1986) and Fudenberg and Levine (1989) where players' records contain the full history of play and the models of Liu and Skrzypacz (2014), Bhaskar and Thomas (2019), Levine (2021), and Pei (2024a) where players' record length is some bounded number.

However, in many situations, the lengths of players' records are *endogenous* and are affected by their strategic behaviors. To fix ideas, sellers in online platforms may bribe consumers for deleting negative reviews, and may even threaten to sue them for defamation if the negative reviews are not removed.[1] Through these bribes and threats, the seller could affect the *number* of reviews future consumers observe.

This paper takes a first step to study reputation formation when players' record lengths are *endogenous*. I analyze a novel reputation model in which a long-lived player might be able to erase his past actions

*Department of Economics, Northwestern University. Email: harrydp@northwestern.edu. I thank Jeff Ely, Nina Fluegel, Drew Fudenberg, Marina Halac, David Levine, Teddy Mekonnen, Ayça Kaya, Alessandro Pavan, Larry Samuelson, Ali Shourideh, Andrzej Skrzypacz, Egor Starkov, Philipp Strack, Bruno Strulovici, Yiman Sun, Alex Wolitzky, and Yutong Zhang for helpful comments. I thank the NSF Grants SES-1947021 and SES-2337566 as well as the Cowles Foundation for financial support.

[1]According to Section 5 in Tadelis (2016), the lack of negative reviews due to seller reciprocity, retaliation, and harassment has caused significant biases in online reviews. A 2019 report in CNBC documents that many consumers who left negative reviews on Yelp were sued by firms in SLAPP lawsuits. Reports from CNET and the Guardian document that in the US and the UK, many Amazon sellers bribe consumers for deleting negative reviews. Bolton, Greiner, and Ockenfels (2013) provide empirical evidence for seller reciprocity and retaliation, making it more costly for buyers to post negative reviews than posting positive ones. Nosko and Tadelis (2015) document that only 0.07% of the reviews on eBay are negative despite a much larger fraction of the consumers complained to consumer service. Cai, et al (2014) and Tadelis (2016) report similar findings on EachNet and Airbnb.

from his records. The main takeaway is that when a player is patient, the possibility that he can erase records *cannot* eliminate his incentives to build reputations. However, it can wipe out all of his returns from building reputations even when his opponents only assign a low probability to types who can erase records.

I study a repeated game between a long-run player (e.g., a seller) and a sequence of short-run players (e.g., consumers). The long-run player discounts future payoffs and exits the game with some exogenous probability after each period. Players' stage-game payoffs are monotone-supermodular. The product choice game in Mailath and Samuelson (2001) satisfies my assumption, which I use to illustrate my results:

seller \ consumer	**L**arge Quantity	**S**mall Quantity
Good Products	$\mathbf{1}, 1$	$-g, x$
Bad Products	$1+g, -x$	$0, 0$

with $g > 0$ and $x \in (0, 1)$.

By the end of each period, the seller can erase his action in that period at a positive cost c.[2] I focus on the case where the cost of erasing an action c is strictly lower than the cost of supplying good products g.[3]

The seller has persistent private information about his type: He is either an *honest type* who always supplies good products and never erases any action, or an *opportunistic type* who decides which products to supply and whether to erase his actions in order to maximize his payoff. Each consumer *can* observe the seller's unerased actions (i.e., the seller's *record*) but *cannot* observe how many actions were erased. Consistent with the literature on reputation effects with limited memories such as Liu and Skrzypacz (2014), the consumers *cannot* directly observe the seller's age in the game, or equivalently calendar time.[4] They have a prior belief about the seller's age, which is determined by the seller's exit rate. After observing the seller's record, the consumers update their beliefs about the seller's age and type via Bayes rule.[5]

When there is no honest type, the opportunistic seller will always have a *strict* incentive to supply bad products and will receive his minmax value 0 no matter how patient he is. However, as long as the honest type occurs with positive probability, Theorem 1 shows that a patient opportunistic seller will supply good products with positive probability for a long time until he has a sufficiently long good record at which point

[2]The seller in my model can only erase reviews but cannot modify the content of reviews. Arguably, it is harder to persuade dissatisfied consumers to write positive reviews than to ask them to stay silent. My main result shows that reputation effects will fail when sellers can manipulate their records, which is *stronger* when they can only erase reviews but cannot modify their content.

[3]I study the case where $c > g$ in Online Appendix B. My assumption that $c < g$ seems reasonable since the consumers' losses from their bad experiences are sunk, so they might be willing to remove their negative reviews in exchange for a small bribe or to avoid a defamation lawsuit. The firms' marginal costs of issuing giftcards and making legal threats seem to be reasonably low.

[4]Information about the seller's age on the market is not available or cannot be easily obtained by consumers in online platforms such as Yelp, Amazon, and TMall since they only display in a salient place the number of reviews each seller received, the number times that he received each rating, together with some comments. Online Appendix E extends my results to settings where either the consumers arrive *stochastically* or they *post reviews with probability less than* 1, in which cases the seller's age on the market, the number of consumers that he has interacted with, and the number of reviews that he has received may not be the same.

[5]The canonical reputation model in Fudenberg and Levine (1989) is consistent with my formulation. This is because when the long-run player *cannot* erase any action, his age in the game (i.e., calendar time) equals the length of his record. In this case, the short-run players will have a *degenerate posterior belief* about calendar time *after* they observe the long-run player's record.

his continuation value will be strictly greater than his commitment payoff 1. The intuition is that the seller can signal his honesty via the *length* of his good record, and in every equilibrium, both the seller's reputation and the probability with which the consumers play L are *increasing* in the length of the seller's good record.

Despite Theorem 1 shows that the seller can secure a high payoff *after* accumulating a long good record, it does not imply that he will receive a high payoff *in equilibrium*. My main result, Theorem 2, shows that as long as a patient seller is *sufficiently long-lived*, (i) his equilibrium payoff must be close to its minmax value 0 and (ii) the consumers' equilibrium welfare, measured by the expected sum of their payoffs, must be close to that in an auxiliary complete information game where the seller's type is common knowledge.

My result implies that the presence of *a small fraction of opportunistic types* who may supply bad products and may erase records can wipe out *all* of the seller's returns from building reputations. However, it has little effect on consumer welfare. It also implies that when a seller *promises* to the consumers that he will supply good products and will never erase any action, his benefit from such a promise will be seriously compromised as long as the consumers entertain a grain of doubt about his willingness to honor his promise.

Intuitively, the seller's returns from building reputations is wiped out by the conflict between two effects, both of which are driven by his ability to erase records. First, the opportunistic seller can sustain his current continuation value by supplying bad products and then erasing his action. Therefore, he has an incentive to supply good products only when his continuation value increases fast enough with the length of his good record. This leads to an *upper bound* on the maximal length of good record that the opportunistic type may have, or equivalently, the minimal length of good record required for the seller to have a perfect reputation.

Second, the opportunistic type's ability to erase records enables him to *pool with a younger honest type*. According to Bayes rule, this will lower the seller's reputation when he has a short good record, which leads to a *lower bound* on the length of good record required for the seller to have a perfect reputation. This lower bound *increases* in the seller's expected lifespan since a decrease in the exit rate increases the likelihood of the *old opportunistic type* relative to the *young honest type*. Once this lower bound on the length of good record exceeds the upper bound implied by the need to provide the opportunistic type incentives, the opportunistic type has to separate from the honest type in the beginning of the game in order to boost his initial reputation. If this is the case, then the opportunistic type's equilibrium payoff must be close to his minmax value since his continuation value after separating from the honest type equals his minmax value.

In terms of consumer welfare, although the opportunistic type will play G with probability bounded above 0 at some histories, I show that when the seller is sufficiently long-lived, such histories must occur with probability close to 0. An implication of this observation is that the *average* probability with which the opportunistic-type seller plays G vanishes to zero as his expected lifespan goes to infinity. Hence, in every equilibrium of the incomplete information game, consumers' welfare must be close to that in an auxiliary complete information game where the seller's type is known and the opportunistic type never plays G.

This paper contributes to the reputation literature by taking a first step to analyze reputation formation when players' record lengths are endogenous and are determined by their strategic behaviors. It stands in contrast to Ekmekci (2011), Vong (2022), Kovbasyuk and Spagnolo (2023), and Wong (2023) which study games where the record systems are designed by social planners who can commit, Liu (2011) which studies a reputation model where the uninformed players decide how much information to acquire about the game's history, as well as Ekmekci, Gorno, Maestri, Sun and Wei (2022) and Saeedi and Shourideh (2023) in which a long-run player can manipulate the *content* of the public signals rather than the *length* of his records.

My reputation model has several merits relative to some of the existing ones. First, it leads to sharp predictions not only on the long-run player's equilibrium payoff, but also on players' behaviors and the short-run players' welfare. In contrast, the results in Fudenberg and Levine (1989) and many follow-up results focus exclusively on the long-run player's equilibrium payoff but cannot deliver sharp predictions on players' behaviors and the short-run players' welfare. The details are explained in Li and Pei (2021).

Second, the predictions of my reputation model fit some of the empirical findings in online marketplaces. For example, Livingston (2005) finds that sellers' sales on eBay depend mostly on the lengths of their good records. Nosko and Tadelis (2015) document that 99.3% of the reviews on eBay are positive despite a much larger fraction of consumers are dissatisfied and complained to customer service. These findings match the predictions of my model that when the seller's cost of erasing actions is low, (i) consumers trust the seller with higher probability when the latter has a longer good record and (ii) the number of negative reviews is much smaller than the number of times that the consumers are dissatisfied (i.e., the seller supplied low quality). Meanwhile, in Fudenberg and Levine (1989), the relationship between the consumers' actions and the length of the seller's record is not necessarily monotone. In Liu and Skrzypacz (2014), the consumers' actions depend only on the timing of the latest bad review rather than on the number of good reviews.

Third, my results suggest that when a long-run player's record length is endogenous, the equilibrium outcomes depend not only on his *effective discount factor*, but also on whether it results from his *time preference* or his *survival probability*. This novel feature stands in contrast to Fudenberg and Levine (1989) and existing reputation models where calendar time (i.e., the long-run player's age) is not observed such as Liu (2011), Liu and Skrzypacz (2014), Levine (2021), and Pei (2024a) in which players' time preferences and their survival probabilities play the same role and there is no need to distinguish between the two.

My reputation failure result is related to Ely and Välimäki (2003) and Ely, Fudenberg and Levine (2008). They focus on *participation games* in which the short-run players can shut down learning and show that the patient player will receive a low payoff when there are *bad commitment types*. In my model, the short-run players *cannot* shut down learning and reputation failure is caused by the low rate of learning.[6]

My model can be interpreted as a continuum of sellers and consumers being randomly matched in each period. Each consumer observes the *record* of the seller she is matched with, but the sellers cannot observe the consumers' records. This is related to community enforcement models with a continuum of players, such as Takahashi (2010), Heller and Mohlin (2018), Bhaskar and Thomas (2019), and Clark, Fudenberg and Wolitzky (2021).[7] My contribution is to introduce *endogenous record length* to this literature.

In a follow-up work (Pei 2024b), I study a community enforcement model with *complete information*, where *all* players are long-lived and each player can erase signals from his records. In contrast, the current model has only *one long-run player* but allows for *incomplete information* about his type. The results in the current paper shed light on the *dynamics* of players' behaviors, which is not the case in Pei (2024b).

My work is also related to several recent papers on dynamic information censoring such as Smirnov and Starkov (2022), Hauser (2023), and Sun (2024). Compared to their works, I highlight the roles of *record length* and *the informed player's expected lifespan* on the value of reputations. Unlike those papers in which the uninformed player's payoff depends only on the informed player's *type*, the uninformed player's payoff in my model depends only on the informed player's *action*. My formulation is standard in models of repeated games and reputations, which fits markets where quality provision is subject to moral hazard.

The long-run player in my model decides whether to disclose each of his past actions but cannot fabricate information. This is related to the literature on disclosing hard information pioneered by Grossman (1981), Milgrom (1981), and Dye (1985). More closely related is Dziuda (2011) in which a sender has some pieces of good evidence and bad evidence and chooses a subset of them to disclose. In contrast to her static model in which the distribution of the evidence available to the sender is exogenous, my model is dynamic in which the numbers of good and bad actions are endogenously determined by the long-run player's behaviors.

2 The Baseline Model

Time is indexed by $k = 0, 1, 2, ...$ A long-lived player 1 (e.g., seller) interacts with an infinite sequence of short-lived player 2s (e.g., consumers), arriving one in each period and each plays the game only in the period she arrives. In each period $k \in \mathbb{N}$, players choose their actions $a_{1,k} \in A_1$ and $a_{2,k} \in A_2$ simultaneously

[6] The conclusion that delays are necessary for the patient player to signal his type also appears in repeated signaling games with *interdependent values* where the receiver's payoff depends directly on the sender's type, such as Kaya (2009). In contrast, my model has *private values* and costly delays are caused by the patient player's ability to erase actions from his records.

[7] Sugaya and Wolitzky (2020) show that players' equilibrium payoffs are arbitrarily close to their minmax values when there are bad commitment types who always defect. However, their model has a finite number of players and focuses on *symmetric* stage games with a pairwise dominant action (e.g., the prisoner's dilemma). Both features stand in contrast to my model.

from finite sets A_1 and A_2. Their stage-game payoffs are $u_1(a_{1,k}, a_{2,k})$ and $u_2(a_{1,k}, a_{2,k})$.

Unlike in existing reputation models where player 1's time preference and survival probability play the same role, their roles are different when record length is endogenous. In order to capture this new feature, I distinguish between the *two reasons* for why player 1 discounts future payoffs. First, by the end of each period, he exits the game for exogenous reasons with probability $1 - \overline{\delta}$ where $\overline{\delta} \in (0, 1)$, after which the game ends and both players receive zero payoffs. Second, conditional on surviving in period $k + 1$, he is indifferent between $\widehat{\delta} \in (0, 1)$ unit of utility in period k and 1 unit of utility in period $k + 1$. Therefore, player 1 discounts future payoffs by $\delta \equiv \overline{\delta} \cdot \widehat{\delta}$, which I call his *effective discount factor*. Player 1's *expected lifespan* is $(1 - \overline{\delta})^{-1}$, which depends only on his survival probability $\overline{\delta}$ but not on his time preference $\widehat{\delta}$.

Player 1's type is fixed over time and is denoted by ω. With probability $\pi \in (0, 1)$, he is an *honest type* ω_h, who takes a fixed *commitment action* $a_1^* \in A_1$ in every period and never erases any action. With probability $1 - \pi$, he is an *opportunistic type* ω_o, who can choose any action and by the end of each period k (but before period $k + 1$) can decide whether to erase his period-k action $a_{1,k}$ at cost $c > 0$.[8] I denote the decision of whether to erase $a_{1,k}$ by $c_k \in \{0, c\}$, where $c_k = c$ stands for $a_{1,k}$ being erased and vice versa.

Before choosing $a_{1,k}$, player 1 observes his type $\omega \in \{\omega_h, \omega_o\}$ and the *full history* of the game $\{a_{1,s}, a_{2,s}, c_s\}_{s=0}^{k-1}$. Player 1 observes ω, $\{a_{1,s}, a_{2,s}, c_s\}_{s=0}^{k-1}$, and $(a_{1,k}, a_{2,k})$ before choosing c_k.

Each player 2 observes player 1's *unerased actions* but *cannot directly observe* player 1's type, the number of actions player 1 erased, and the time at which each unerased action was taken. Formally, player 2's period k history is $\{a_{1,\tau_0}, ..., a_{1,\tau_{m(k)}}\}$ where $0 \leq \tau_0 < ... < \tau_{m(k)} \leq k - 1$ such that for every $s \in \{0, 1, ..., k - 1\}$, there exists $i \in \{0, 1, ..., m(k)\}$ such that $s = \tau_i$ if and only if $c_s = 0$. Let $\mathcal{H}$ denote the set of player 2's histories, or equivalently, *player 1's records*. Let h denote a typical element of $\mathcal{H}$. My results extend to the case in which player 2 only observes *the number of times* that each $a_1 \in A_1$ occurred in the sequence $\{a_{1,\tau_0}, ..., a_{1,\tau_{m(k)}}\}$, that is, the *summary statistics* of player 1's unerased actions.

I also assume that the short-run players *cannot* directly observe the long-run player's *age* in the game, or equivalently, calendar time. This assumption is common in reputation models with limited memories such as Liu (2011), Liu and Skrzypacz (2014), Levine (2021), and Pei (2024a). It is consistent with the context that motivates my study, namely, the consumers may not know the extent to which the seller has erased his records. It seems reasonable on platforms such as Yelp, Amazon, and TMall, which show the number of each rating the seller has received but do *not* display in a salient place the seller's time on the market.

As in Liu and Skrzypacz (2014), the short-run players have a *prior belief* about player 1's *age* and *type*, which is their belief *after* knowing that the game has not ended but *before* they observe their history. After

[8] All my results except for Proposition 6 extend to the case where the cost of erasing actions is 0. In my baseline model, the opportunistic type *cannot* erase actions taken in previous periods. I will discuss this modeling assumption in Section 4.

they observe their history, they update their belief about player 1's age and type using Bayes rule.[9]

Recall that the short-run player's prior belief assigns probability π to the honest type, which might be different from her posterior belief after she observes that the long-run player has no action in his record. The probability that her posterior belief at $h \in \mathcal{H}$ assigns to the honest type is called *player* 1*'s reputation* at h.

In terms of the short-run player's prior belief about the long-run player's age, recall that the long-run player exits the game with probability $1-\overline{\delta}$ after each period. Therefore, the probability that the short-run player's prior belief assigns to the long-run player's age being $k+1$ should equal $\overline{\delta}$ times the probability that her prior belief assigns to the long-run player's age being k. The unique prior that satisfies this condition for every $k \in \mathbb{N}$ is the one that assigns probability $(1-\overline{\delta})\overline{\delta}^k$ to the long-run player's age being k.[10]

The opportunistic type's strategy σ_1 and the short-run player's strategy σ_2 are mappings from their histories to their mixed actions, and for the opportunistic type, it also includes the probability with which he erases each of his actions. Let $\mathcal{H}(\sigma_1,\sigma_2) \subset \mathcal{H}$ denote the set of player 2's histories that occur with positive probability when (i) the honest type always plays a_1^* and never erases any action, (ii) the opportunistic type uses strategy σ_1, and (iii) player 2 uses strategy σ_2. I call $\mathcal{H}(\sigma_1,\sigma_2)$ the set of *on-path histories*.

A *Nash equilibrium* is a strategy profile (σ_1,σ_2) such that (i) σ_1 maximizes the opportunistic type's *discounted average payoff* $\sum_{k=0}^{+\infty}(1-\delta)\delta^k\{u_1(a_{1,k},a_{2,k})-c_k\}$ against σ_2 and (ii) at every on-path history, σ_2 maximizes player 2's stage-game payoff against σ_1. To simplify the description of players' equilibrium behaviors, I examine the common properties of all *public equilibria* (hereafter, *equilibria*), which are Nash equilibria such that *the opportunistic type's strategy* is measurable with respect to player 2's history, or equivalently, the history that both players observe. Focusing on public equilibria is common in reputation models with limited memories or with incomplete records, see for example, Liu (2011) and Liu and Skrzypacz (2014). In fact, I show that restricting attention to public equilibria is without loss of generality:

Lemma 1. *For any Nash equilibrium (σ_1,σ_2), there exists a public equilibrium (σ_1^*,σ_2) such that (i) $\mathcal{H}(\sigma_1,\sigma_2)=\mathcal{H}(\sigma_1^*,\sigma_2)$, (ii) at every $h \in \mathcal{H}(\sigma_1,\sigma_2)$, player 2's beliefs about player 1's behavior and type at h are the same under (σ_1,σ_2) and under (σ_1^*,σ_2), and (iii) the expected values of $\sum_{k=0}^{+\infty}(1-\delta)\delta^k\{u_1(a_{1,k},a_{2,k})-c_k\}$ and $\sum_{k=0}^{+\infty}\overline{\delta}^k u_2(a_{1,k},a_{2,k})$ are the same under (σ_1,σ_2) and under (σ_1^*,σ_2).*

The proof is in Online Appendix A. In that appendix, I also establish a *state distribution lemma* (Page 2), which applies to *all* Nash equilibria and will be used not only in the proof of Lemma 1 but also in the proofs of other results in this paper. Lemma 1 implies that focusing on public equilibria is without loss of

[9]The short-run players' *posterior belief* about player 1's age and type will depend on players' equilibrium strategies. Fudenberg and Levine (1989)'s model is also consistent with my formulation, since after observing the long-run player's record, which consists of the entire sequence of his actions, the short-run players' posterior belief will assign probability 1 to the true calendar time.

[10]Hu (2020) provides a foundation for this exponential prior belief about the long-run player's age by examining a game in which the short-run players face uncertainty about the entry process and establishing the equivalence between these two models.

generality in the sense that for any Nash equilibrium, there exists an equivalent public equilibrium such that the opportunistic type's discounted average payoff, the short-run players' strategy, their beliefs about player 1's behavior and type at every on-path history, and the sum of their expected payoffs (taking into account the fact that the game ends with probability $1-\overline{\delta}$ after each period) remain the same. Therefore, as long as some properties on player 1's equilibrium payoff, player 1's reputation, the sum of player 2's payoffs, and player 2's behavior apply to *all* public equilibria, they must also apply to *all* Nash equilibria.

Next, I introduce three assumptions on (u_1, u_2). I start from a standard monotone-supermodularity assumption, which is also satisfied in Mailath and Samuelson (2001), Ekmekci (2011), and Liu (2011).

Assumption 1. *There exist a complete order $\succsim_1$ on A_1 and a complete order $\succsim_2$ on A_2 such that $u_1(a_1, a_2)$ is strictly decreasing in a_1 and is strictly increasing in a_2, and u_2 satisfies strictly increasing differences, i.e., $u_2(a_1, a_2) - u_2(\widetilde{a}_1, a_2) > u_2(a_1, \widetilde{a}_2) - u_2(\widetilde{a}_1, \widetilde{a}_2)$ for every $a_1 \succ_1 \widetilde{a}_1$ and $a_2 \succ_2 \widetilde{a}_2$.*[11]

The product choice game in the introduction satisfies Assumption 1 once the row player's actions are ranked according to $G \succ_1 B$ and the column player's actions are ranked according to $L \succ_2 S$.

Let $\underline{a}_1$ denote the lowest action in A_1, which by Assumption 1, is strictly dominant in the stage game. Let $\underline{a}_2$ denote player 2's *lowest* best reply to $\underline{a}_1$. Under Assumption 1, $u_1(\underline{a}_1, \underline{a}_2)$ is player 1's *minmax value* in the sense of Fudenberg, et al (1990), which requires player 2 to play a best reply to some $\alpha_1 \in \Delta(A_1)$. I normalize players' stage-game payoffs so that $u_1(\underline{a}_1, \underline{a}_2) = u_2(\underline{a}_1, \underline{a}_2) \equiv 0$. Let a_2^* denote player 2's *highest* best reply to a_1^*. I focus on the interesting case in which player 1's *commitment payoff* $u_1(a_1^*, a_2^*)$ is *strictly* greater than his minmax value 0. Otherwise, player 1 will have no incentive to build reputations.

Assumption 2. $u_1(a_1^*, a_2^*) > u_1(\underline{a}_1, \underline{a}_2) \equiv 0$.

In order to simply the analysis, I also make an assumption that is *generically* satisfied.

Assumption 3. *For every $i \in \{1, 2\}$ and $a_{-i} \in A_{-i}$, player i has a strict best reply to a_{-i}. For every $a_2 \in A_2$ and $\lambda \in [0, 1]$, if a_2 best replies to player 1's mixed action $\lambda a_1^* + (1-\lambda)\underline{a}_1$, then there exists $\widetilde{\lambda} \in [0, 1] \backslash \{\lambda\}$ such that a_2 also best replies to mixed action $\widetilde{\lambda} a_1^* + (1-\widetilde{\lambda})\underline{a}_1$.*

The first part of Assumption 3 requires each player to have a *strict* best reply to each of their opponent's pure actions. This assumption is generically satisfied since A_1 and A_2 are finite sets. Under this assumption, $\underline{a}_2$ is the *unique* best reply to $\underline{a}_1$ and a_2^* is the *unique* best reply to a_1^*. This part of Assumption 3 together with Assumptions 1 and 2 implies that $a_1^* \succ_1 \underline{a}_1$ and $a_2^* \succ_2 \underline{a}_2$. The second part of Assumption 3 rules

[11] Strict orders on A_1 and A_2, $\succ_1$ and $\succ_2$, are defined naturally from $\succsim_1$ and $\succsim_2$. Abusing notation, for every $\beta, \beta' \in \Delta(A_2)$, I use $\beta \succsim_2 \beta'$ to denote β first-order-stochastically dominates β' under the order induced by $\succsim_2$ and similarly for $\beta \succ_2 \beta'$.

out player 2's action that *only* best replies to one mixture between a_1^* and $\underline{a}_1$. It allows some of player 2's actions to be strictly dominated and also allows some actions to best reply to an open set of mixtures. Let

$$\mathcal{B} \equiv \Big\{ \beta \in \Delta(A_2) \Big| \beta \text{ best replies to } \lambda a_1^* + (1-\lambda)\underline{a}_1 \text{ for some } \lambda \in [0,1] \Big\}, \tag{2.1}$$

which is the set of player 2's mixed actions that best reply to some mixtures between a_1^* and $\underline{a}_1$. Lemma 2 shows that under my assumptions, each pair of player 2's mixed actions in $\mathcal{B}$ can be ranked according to FOSD and that one can generate a rich set of payoffs for player 1 by varying player 2's actions in $\mathcal{B}$.

Lemma 2. *Suppose (u_1, u_2) satisfies Assumptions 1, 2, and 3. The set of elements in $\mathcal{B}$ can be completely ranked according to FOSD, with $\underline{a}_2$ the lowest element and a_2^* the highest element. For every $a_1 \in A_1$ and $v \in [u_1(a_1, \underline{a}_2), u_1(a_1, a_2^*)]$, there exists a unique $\beta \in \mathcal{B}$ that satisfies $u_1(a_1, \beta) = v$.*

Appendix A shows Lemma 2 and explains why the second part of Assumption 3 is generically satisfied. Later, we will see that Lemma 2 implies that player 2's mixed actions on the equilibrium path can be ranked.

Remark: An alternative interpretation of my model is that before the game starts, player 1 *promises* to all the short-run players that he will always take action a_1^* and that he will never erase any action. Player 2 believes that player 1 will keep his promise with probability π. With probability $1-\pi$, player 1 can *renege* in the sense that he may take actions other than a_1^* and may also erase his actions. Whether player 1 can renege is determined by a random draw in period 0 *after* he makes his promise and remains fixed throughout the infinite horizon game.[12] I will use this interpretation to discuss the implications of my theorems.

2.1 Discussions on the Modeling Assumptions

My baseline model assumes that a consumer will arrive in each period, in which case the *honest type*'s age will coincide with the number of actions in his record. In Online Appendix E, I discuss extensions to settings where the consumers either *arrive stochastically* or *do not post reviews with positive probability*. In these cases, even the honest type's age may not coincide with the number of actions in his record.

In my model, player 1 can erase actions from his records but *cannot* modify the content of his records. This is motivated by an observation in Tadelis (2016) that most of the consumers post reviews because they are intrinsically motivated to share their opinions, to reward sellers' good behaviors and to punish bad ones, or to provide future consumers useful information. If this is the case, then it seems more challenging for

[12] This form of imperfect promises is studied by Lipnowski, Ravid and Shishkin (2022) in a one-shot communication game. This formulation has also been used in dynamic settings, such as in the reputational bargaining models of Kambe (1999) and Wolitzky (2012), where a player first announces a bargaining posture and then becomes committed to it with positive probability.

sellers to convince consumers to explicitly lie about their experiences than to ask them to stay silent. My main result shows that sellers will receive low payoffs when they cannot commit not to manipulate their records, which is stronger when they can only erase reviews but cannot modify the content of reviews.[13]

My baseline model focuses on the case in which (i) the short-run players' best reply does not depend on whether the long-run player will erase his action and (ii) the cost of erasing actions does not depend on the action profile being played. In practice, if we interpret erasing actions as offering a partial refund or a giftcard in exchange for deleting a review, then the consumers' demands for refunds or giftcards may depend on the quantity they purchased and on the seller's action (e.g., the quality he supplies). My results can be extended to cases where (i) each consumer's best reply also depends on the probability with which the seller erasing his action, and (ii) the seller's cost of erasing actions depends on the actions being played.

My baseline model focuses on the case in which there is only one honest type and one opportunistic type and assumes that the honest type *cannot* erase actions. Section 4 studies an extension where the honest type mechanically takes action a_1^* in every period but can *strategically* decide whether to erase his action at a strictly positive cost $c > 0$ in order to maximize his discounted average payoff. In Online Appendix D, I study another extension where (i) there are multiple honest types taking different pure actions and (ii) there are multiple opportunistic types with different stage-game payoffs and different costs of erasing actions.

My baseline model assumes that the short-run players *cannot* observe previous short-run players' actions. This assumption is standard in reputation models with limited memories, which is also made in Liu (2011), Liu and Skrzypacz (2014), and Pei (2024a). When a short-run player who arrives in period $s > k$ observes $a_{1,k}$ and $a_{2,k}$, the realized $a_{2,k}$ is an informative signal about the time at which $a_{1,k}$ was taken, leading to an intractable statistical inference problem. I leave the analysis of this case to future work.

My baseline model also assumes that the short-run players can perfectly observe the long-run player's *unerased actions*. This rules out situations in which they can only observe noisy signals about those actions. Analyzing repeated games with incomplete information, imperfect monitoring, and limited observations is challenging. This explains why most of the existing analysis on reputation games with finite record lengths such as Liu (2011), Liu and Skrzypacz (2014), and Heller and Mohlin (2018) all focus on perfect monitoring.[14] The case with endogenous record length and imperfect monitoring is left for future work.

[13] If the opportunistic type can modify the content of his records at a low cost, e.g., he can change his record to a_1^* when his action was $\underline{a}_1$, then he will never play a_1^* since doing so is strictly dominated by playing $\underline{a}_1$ and then modifying his record to a_1^*. Compared to the predictions of my model that the seller will build a reputation, the prediction that the seller never supplies high quality does not seem to match the stylized fact in Tadelis (2016) that online reviews motivate many sellers to supply high quality.

[14] Bhaskar and Thomas (2019) allow for imperfect monitoring, but they focus on a *complete information game* in which the long-run player has only one type. Levine (2021) allows for imperfect monitoring but assumes that players' record length is 1.

3 Analysis

Section 3.1 shows that the patient player can secure his commitment payoff in all equilibria when he *cannot* erase actions. This is also the case when his cost of erasing actions is high enough. My subsequent analysis focuses on the case where the cost of erasing actions is low. Section 3.3 characterizes the dynamics of players' behaviors. Section 3.4 states a reputation failure result, which implies that when the patient player is *sufficiently long-lived*, his payoff in every equilibrium must be close to his minmax value 0 even when the honest type occurs with high probability. My analysis also pins down the short-run players' welfare.

3.1 Benchmark: The Opportunistic Type Cannot Erase Actions

I start from a benchmark scenario in which player 1 *cannot* erase any action. Since A_1 and A_2 are finite sets and a_2^* is a strict best reply to a_1^*, there exists a constant $\gamma \in (0,1)$ that depends only on (u_1, u_2) such that that a_2^* is a *strict* best reply to any mixed action of player 1's that assigns probability at least γ to a_1^*.

Proposition 1. *If player 1 cannot erase actions and (u_1, u_2) satisfies Assumptions 1, 2, and 3, then player 1's payoff in every equilibrium is at least* $u_1(a_1^*, a_2^*) - (1 - \delta^{\frac{\log \pi}{\log \gamma}})\big(u_1(a_1^*, a_2^*) - u_1(a_1^*, \underline{a}_2)\big)$.

The proof of Proposition 1 is in Online Appendix B, which follows from Fudenberg and Levine (1989). Intuitively, when the opportunistic type deviates to playing a_1^* in every period, either player 2 has a strict incentive to play a_2^*, or she believes that a_1^* will be played with probability less than γ in which case after observing a_1^*, her posterior belief about the honest type will be multiplied by at least $1/\gamma$. Therefore, in periods where player 2 does not play a_2^*, player 1's reputation grows at an *exponential* rate, so there can be at most $\log \pi / \log \gamma$ such periods. This leads to a lower bound on player 1's discounted average payoff under such a deviation and his equilibrium payoff must be weakly greater than his payoff under any deviation.

Proposition 1 implies that when player 1 cannot erase actions, as long as the effective discount factor δ exceeds some cutoff δ^*, his equilibrium payoff is bounded below by something ε-close to his commitment payoff $u_1(a_1^*, a_2^*)$. In Online Appendix B, I extend this conclusion to settings where the cost of erasing actions c is greater than some cutoff, in which I also show that player 1 will *never* erase any action on the equilibrium path. In the product choice game, the cutoff cost equals the cost g of supplying good products. Importantly, the payoff lower bounds in these results depend on $\widehat{\delta}$ and $\overline{\delta}$ *only* through their product $\delta \equiv \overline{\delta} \cdot \widehat{\delta}$.

3.2 The Threshold Cost of Erasing Actions

Motivated by the benchmark result in Section 3.1, my subsequent analysis focuses on settings where the cost of erasing actions c is *lower* than the cost of supplying good products g. This restriction seems reasonable since consumers' losses from their bad experiences are sunk,[15] so they might be willing to remove their negative reviews in exchange for a small bribe, or to avoid a defamation lawsuit. From the firm's perspective, the marginal costs of paying bribes (e.g., issuing giftcards) and making legal threats are usually not that high.

In order to define the threshold cost of erasing actions more generally, let a_1' denote the lowest action in A_1 such that $\underline{a}_2$ does *not* best reply to a_1'. Since $\underline{a}_2$ best replies to $\underline{a}_1$, Assumption 1 implies that $a_1' \succ_1 \underline{a}_1$. Assumption 2 requires that $u_1(a_1^*, a_2^*) > u_1(\underline{a}_1, \underline{a}_2)$, which implies that $a_2^* \succ_2 \underline{a}_2$ and $a_1^* \succeq_1 a_1'$.

Assumption 4. *The cost of erasing actions c satisfies $c < \overline{c} \equiv \min_{a_2 \in A_2} \big\{ u_1(\underline{a}_1, a_2) - u_1(a_1', a_2) \big\}$.*[16]

Assumption 4 requires c to be less than $\overline{c}$, which is defined as the lowest cost that player 1 needs to incur in order to increase his action from $\underline{a}_1$ to a_1'. In the product choice game, $\overline{c} = g$. Since $a_1^* \succeq_1 a_1'$ and $u_1(a_1, a_2)$ is strictly decreasing in a_1, Assumption 4 implies that c is lower than the cost of playing a_1^*.

3.3 The Dynamics of Players' Behaviors

In this section, I characterize players' equilibrium behaviors when the opportunistic type's cost of erasing actions is strictly less than $\overline{c}$. I start from a result which implies that players will have *strict* incentives to play the minmax action profile $(\underline{a}_1, \underline{a}_2)$ *after* player 1 is revealed to be the opportunistic type.

Proposition 2. *Suppose $\pi = 0$ and (u_1, u_2, c) satisfies Assumptions 1, 2, 3, and 4. In every equilibrium (σ_1, σ_2) and at every on-path history $h \in \mathcal{H}(\sigma_1, \sigma_2)$, players have strict incentives to play $\underline{a}_1$ and $\underline{a}_2$.*

Intuitively, even after his opponents rule out the honest type, it is not obvious why player 1 has no incentive to take actions above $\underline{a}_1$ since he could be rewarded or punished based on *the number of high actions he took*, which *cannot* be fabricated. I show that this logic breaks down in absence of the honest type since the opportunistic type has no incentive to take high actions at histories where his continuation value is close to its maximum. Anticipating this, the short-run players strictly prefer $\underline{a}_2$ at such histories, which implies that they will have no incentive to reward the opportunistic type when the latter is supposed to receive a high continuation value. The incentive to take high actions disappears due to the lack of rewards.

Proof of Proposition 2: Fix any equilibrium (σ_1, σ_2). Since the opportunistic type's continuation value depends only on player 2's history $h \in \mathcal{H}$, I use $V(h)$ to denote player 1's continuation value at h. Let

[15] An interpretation is that the seller decides whether to supply high quality experienced goods and the consumers decide how much to purchase *without* knowing product quality. The consumers can observe the seller's action, i.e., quality, *after* purchase. Then they post a review that honestly reflects the seller's action. The seller then decides whether to erase that review at cost c.

[16] Even when player 1 has *a continuum of actions*, the threshold cost of erasing actions $\overline{c}$ is also strictly bounded above 0. As a concrete example, suppose a seller chooses his effort from the unit interval $[0, 1]$ and the consumers have an incentive trust him only when his expected effort is more than $x \in (0, 1)$, then $\overline{c}$ equals the seller's minimal cost of exerting effort x.

$\overline{V} \equiv \sup_{h \in \mathcal{H}(\sigma_1,\sigma_2)} V(h)$. Suppose by way of contradiction that $\overline{V} > 0$. The definition of $\overline{V}$ implies that for every ε that satisfies

$$0 < \varepsilon < \min\Big\{\frac{\overline{V}}{2}, \frac{(1-\delta)(\overline{c}-c)}{\delta}\Big\}, \tag{3.1}$$

there exists $h \in \mathcal{H}(\sigma_1,\sigma_2)$ such that $V(h) > \overline{V} - \varepsilon$. I examine the opportunistic type's incentive at such a history h. His payoff from playing $\underline{a}_1$ and then erasing it is at least $(1-\delta)(u_1(\underline{a}_1, \beta(h)) - c) + \delta V(h)$, where $\beta(h) \in \Delta(A_2)$ is player 2's equilibrium action at h. His payoff from playing any $a_1 \succsim_1 a_1'$ is at most $(1-\delta)u_1(a_1,\beta(h)) + \delta\overline{V}$, which is strictly less than $(1-\delta)(u_1(\underline{a}_1,\beta(h)) - c) + \delta V(h)$ given that $c < \overline{c}$ and ε satisfies (3.1). This comparison implies that player 1 has no incentive to play any $a_1 \succsim_1 a_1'$ at h. Since $\pi = 0$, the definition of a_1' implies that player 2 has a strict incentive to play $\underline{a}_2$ at h. Therefore,

$$\overline{V} - \varepsilon < V(h) \leq (1-\delta)u_1(\underline{a}_1, \underline{a}_2) + \delta\overline{V} = \delta\overline{V},$$

which implies that $\overline{V} - \frac{\varepsilon}{1-\delta} < 0$ for every ε that satisfies (3.1). Hence, $\overline{V} \leq 0$. Since the opportunistic type's minmax value is $u_1(\underline{a}_1, \underline{a}_2) \equiv 0$, his continuation value is always 0. Since $u_1(a_1, a_2)$ is strictly decreasing in a_1, player 1 strictly prefers to play $\underline{a}_1$ and in response, player 2 strictly prefers to play $\underline{a}_2$. □

Next, I examine the case where the honest type occurs with positive probability. A useful observation is that player 1 will never take any action other than a_1^* and $\underline{a}_1$ with positive probability. This is because if in equilibrium, he takes any action $a_1 \notin \{a_1^*, \underline{a}_1\}$ with positive probability and does not erase it afterwards, then he will be separated from the honest type and Proposition 2 implies that his continuation value will be 0. In order to obtain a strictly positive continuation value after taking action a_1, he needs to erase a_1. But *taking action* a_1 *and then erasing it* is strictly dominated by *taking action* $\underline{a}_1$ *and then erasing it*, since they lead to the same history for player 2 in the next period but the latter results in a higher stage-game payoff.

Since player 1 only takes a_1^* and $\underline{a}_1$ with positive probability, player 2's action at every $h_2 \in \mathcal{H}(\sigma_1,\sigma_2)$ belongs to $\mathcal{B}$ defined in (2.1). Let h_*^k denote player 2's history where she observes k actions, all of which are a_1^*. Let $\mathcal{H}_* \equiv \{h_*^k \big| k \geq 0\}$ denote the set of h_*^k. Let $\beta_k \in \Delta(A_2)$ denote player 2's action at h_*^k. Let $\pi_k \in [0,1]$ denote player 1's reputation at h_*^k. Theorem 1 characterizes some common properties of players' behaviors and reputations, which apply to *all* equilibria under *all* $\pi \in (0,1)$, $\widehat{\delta} \in (0,1)$, and $\overline{\delta} \in (0,1)$.

Theorem 1. *Fix any* $\pi > 0$ *and* (u_1, u_2, c) *that satisfies Assumptions 1, 2, 3, and 4. In every equilibrium, there exist two cutoffs for player* 1*'s record length* $t_0 \in \mathbb{N}$ *and* $t \in \mathbb{N}$ *with* $0 \leq t_0 \leq t$ *such that:*

1. *The opportunistic type mixes between* a_1^* *and* $\underline{a}_1$ *at* h *if and only if* $h = h_*^k$ *for some* $k < t$*, and at other on-path histories, he plays* $\underline{a}_1$ *for sure. For every* $k \leq t$*,* $\pi_k < 1$*. For every* $k \geq t+1$*,* $\beta_k = a_2^*$ *and*

$\pi_k = 1$. *The opportunistic type's continuation value at* h_*^{t+1} *is* $\max\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta)u_1(\underline{a}_1, a_2^*)\}$.

2. *When* $k \leq t$, π_k *is strictly increasing in* k *and* β_k *is strictly increasing in* k *in the sense of FOSD.*

3. *At history* h_*^k, *the opportunistic type never erases* $\underline{a}_1$ *if* $k < t_0$ *and erases* $\underline{a}_1$ *for sure if* $t_0 < k \leq t$. *He never erases* a_1^* *at any history and never erases any action at histories that do not belong to* $\mathcal{H}_*$.

According to Theorem 1, the opportunistic type's equilibrium behavior is characterized by two cutoffs t_0 and t such that (i) he mixes between a_1^* and $\underline{a}_1$ *if and only if* his record length is *less* than t and does not contain any action other than a_1^*, and plays $\underline{a}_1$ for sure otherwise, (ii) he erases $\underline{a}_1$ *if and only if* his record length is *more* than t_0 and does not contain any action other than a_1^*. Hence, the maximal length of good record that the opportunistic type will have is t. By Bayes rule, player 1's reputation equals 1 when the length of his good record reaches $t+1$. If the opportunistic type has a good record of length at least $t+1$, then his continuation value is $\max\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta)u_1(\underline{a}_1, a_2^*)\}$, which under Assumption 4, is *strictly greater* than his commitment payoff $u_1(a_1^*, a_2^*)$. In terms of the dynamics, as the length of player 1's good record increases, his reputation increases and player 2's action increases in the sense of FOSD.

The behaviors described in Theorem 1 stand in contrast to the conclusion in Proposition 2, according to which players have *strict* incentives to play $\underline{a}_1$ and $\underline{a}_2$ at every on-path history of every equilibrium. Proposition 3 highlights this comparison even further by providing a *uniform lower bound* on the equilibrium value of t, which is the maximal length of good record that the opportunistic type will have.

Proposition 3. *Fix any* $\pi > 0$ *and* (u_1, u_2, c) *that satisfies Assumptions 1, 2, 3, and 4. There exists a constant* $\lambda > 0$ *which is independent of* $\widehat{\delta}$ *and* $\overline{\delta}$ *such that* $t \geq \lambda(1-\delta)^{-1} - 1$ *in every equilibrium.*

Proposition 3 suggests that in any equilibrium, the maximal length of good record that the opportunistic type will have must be *uniformly* bounded below by an affine function of $(1-\delta)^{-1}$. This lower bound diverges to infinity as $\delta \to 1$. That is to say, as long as the opportunistic type has a high effective discount factor δ, he has an incentive to build a reputation for a long time and will do so with positive probability.

The proofs are in Appendices B and C. The rest of this section offers an intuitive explanation for Theorem 1. As in the benchmark *without* honest type, the opportunistic type has no incentive to play a_1^* at h when his continuation value at h, denoted by $V(h)$, is close to his highest continuation value $\overline{V}$. However, player 2 might still be willing to take actions strictly greater than $\underline{a}_2$ even when they know that the opportunistic type will play $\underline{a}_1$ for sure, since their posterior belief may assign *a high probability to the honest type*. Their willingness to reward player 1 with high reputations provides the opportunistic type an incentive to build a reputation via accumulating a long record of a_1^*, until his continuation value becomes close to $\overline{V}$.

Since the opportunistic type has the option to play $\underline{a}_1$ and then erase it at a low cost, he has an incentive to take the commitment action a_1^* only if a *longer good record* results in a *higher continuation value*. Since $u_1(a_1, a_2)$ increases in a_2, a higher continuation value for the long-run player translates into a *higher action for the short-run players*. When opportunistic type decides whether to erase action $\underline{a}_1$, he trades off the benefit from doing so (which is to sustain his current continuation value) and its cost. As a result, he has a stronger incentive to erase actions when his continuation value is higher,[17] or equivalently, when he has a longer good record. The long-run player's reputation is strictly increasing in the length of his good record since the honest type takes the commitment action with probability 1 while the opportunistic type does so with probability less than 1. That is to say, a longer good record signals the long-run player's honesty.

3.4 The Long-Lived Player's Equilibrium Payoff & The Short-Run Players' Welfare

Although Theorem 1 implies that the opportunistic type of the patient player can secure a payoff that is strictly greater than his commitment payoff *after* he accumulates a long enough good record, it remains silent about the patient player's *equilibrium payoff* and the short-run players' welfare. This is because it does not characterize the *time* it takes for the opportunistic type to secure the high continuation value.

Theorem 2 characterizes the opportunistic type's equilibrium payoff and the short-run players' welfare when the patient player is *sufficiently long-lived*. Recall that $\widehat{\delta}$ stands for player 1's time preference, $\overline{\delta}$ stands for his survival probability, and that stage-game payoffs are normalized so that $u_1(\underline{a}_1, \underline{a}_2) = u_2(\underline{a}_1, \underline{a}_2) = 0$.

Theorem 2. *Suppose (u_1, u_2, c) satisfies Assumptions 1, 2, 3, and 4. For every $\pi \in (0,1)$, $\widehat{\delta} \in (0,1)$, and $\varepsilon > 0$, there exists $\delta^* \in (0,1)$ such that when $\overline{\delta} \in (\delta^*, 1)$, in every equilibrium:*

1. *The opportunistic type's discounted average payoff is no more than $(1-\delta)c/\delta$.*
2. *Under an additional assumption that $u_2(a_1, a_2)$ is weakly increasing in a_1, player 2's welfare, measured by $U_2 \equiv \sum_{k=0}^{+\infty}(1-\overline{\delta})\overline{\delta}^k u_2(a_{1,k}, a_{2,k})$, belongs to an ε-neighborhood of $\pi u_2(a_1^*, a_2^*)$.*

Theorem 2 implies that even when player 1 has a high effective discount factor δ and is likely to be the honest type, as long as he is *sufficiently long-lived*, his payoff in every equilibrium must be close to his minmax value 0. When the two types share the same stage-game payoff function, the honest type's payoff is also no more than $(1-\delta)c/\delta$ since the opportunistic type can imitate the honest type. Under the interpretation that a sufficiently long-lived player *promises* to his opponents that he will always play a_1^* and will never erase any action, although he can obtain his commitment payoff $u_1(a_1^*, a_2^*)$ when his opponents

[17] A similar intuition appears in the bad news model of Board and Meyer-ter-Vehn (2013), in which an agent has stronger incentives to exert costly effort when his continuation value is higher.

fully believe in his promise (i.e., $\pi = 1$), his benefit from such a promise is wiped out almost entirely when his opponents entertain a grain of doubt about his willingness to honor his promise (i.e., $\pi < 1$).

My conclusion stands in contrast to the result in Fudenberg and Levine (1989, or Proposition 1 in the current paper) which shows that when player 1 *cannot* erase actions and his effective discount factor δ is *above some cutoff*, his payoff in every equilibrium must be bounded below by something ε-close to his commitment payoff $u_1(a_1^*, a_2^*)$. However, when player 1 *can* erase actions, his equilibrium payoff is close to his minmax value 0 even if his effective discount factor $\delta \equiv \widehat{\delta} \cdot \overline{\delta}$ is arbitrarily close to 1, which is the case when his time preference $\widehat{\delta}$ is fixed to be something close to 1 and his survival probability $\overline{\delta}$ converges to 1. This comparison highlights how endogenous record length affects the long-run player's equilibrium payoff.

In Section 4, instead of assuming that the survival probability $\overline{\delta}$ goes to 1 at a faster rate than the time preference $\widehat{\delta}$, I consider the opposite case where $\widehat{\delta}$ goes to 1 at a faster rate than $\overline{\delta}$, or even $\widehat{\delta} = 1$. Proposition 4 shows that the long-run player's payoff is no more than $\frac{(1-\delta)c}{\delta}$ when the probability of the honest type π is below some cutoff, which implies that Theorem 2 *partially* extends to the case where $\widehat{\delta}$ is much closer to 1 relative to $\overline{\delta}$. Proposition 5 shows that for any payoff $v < u_1(\underline{a}_1, a_2^*) - c$, there exist equilibria in which the opportunistic type's discounted average payoff is more than v as long as π is above some cutoff. The comparison between this result and Theorem 2 highlights the distinction between players' time preference and survival probability in determining the game's equilibrium outcomes: Fix a large enough π, player 1's equilibrium payoffs when $\overline{\delta}$ goes to 1 faster than $\widehat{\delta}$ are different from those when $\widehat{\delta}$ goes to 1 faster than $\overline{\delta}$. This stands in contrast to the canonical reputation models such as the one in Fudenberg and Levine (1989) where the time preference and the survival probability play the same role. It also justifies my modeling choice of distinguishing between the two reasons for why the long-run player discounts future payoffs.

If in addition that $u_2(a_1, a_2)$ is weakly increasing in a_1, which translates into the consumers' payoff increases in the seller's effort, then Theorem 2 also leads to a sharp prediction on the short-run players' welfare. In particular, the sum of their equilibrium payoffs in the incomplete information game must be arbitrarily close to that in an *auxiliary complete information game* where they can observe player 1's type.[18] This implies that although a tiny probability of opportunistic type can wipe out all of the long-lived player's returns from building reputations, it only has a negligible effect on the short-run players' welfare.

Nevertheless, my result evaluates the short-run players' welfare via the *expected sum* of their payoffs. Under this criteria, *conditional on* the game will continue to the next period, the planner who evaluates welfare does *not* value the current player's payoff more than the next one's. However, consumer welfare

[18]The additional assumption is not needed to show that player 2's welfare is at most $\pi u_2(a_1^*, a_2^*) + \varepsilon$ but is needed to show that player 2's welfare is at least $\pi u_2(a_1^*, a_2^*) - \varepsilon$. This is because Assumption 1 only requires u_2 to be supermodular but does not say anything about player 2's preference over player 1's actions, so there is no guarantee that $u_2(\underline{a}_1, \underline{a}_2)$ is less than $u_2(a_1^*, \underline{a}_2)$.

can be low once it is evaluated by a planner who weights the current consumer's payoff much more than the ones in the future. To see this, take the product choice game example and recall the definition of t in the statement of Theorem 1. In equilibrium, the first t consumers find it optimal to play S, from which their payoff is no more than $x \in (0, 1)$. Since Proposition 3 implies that t is bounded below by an affine function of $(1-\delta)^{-1}$, we know that even when the seller is the honest type with probability close to 1, there will be a large number of consumers whose payoffs are bounded below their first-best level 1 (in fact, less than x).

The proof of Theorem 2 is in Sections 3.5 and 3.6. Intuitively, the opportunistic type receives a low payoff due to the conflict between two forces that are driven by his ability to erase records: (i) the need to motivate him to take the commitment action and (ii) the time it takes for him to establish a reputation.

First, the opportunistic type has an incentive to take the commitment action a_1^* only when a longer good record increases his continuation value by at least something proportional to $1-\delta$. This is because he has the option to play $\underline{a}_1$ and then erase it, by which he can sustain his current continuation value. The need to motivate him to accumulate longer good records leads to an *upper bound* on the maximal length of good record that he may have in any equilibrium and this upper bound is proportional to $(1-\delta)^{-1}$.

Second, the opportunistic type's ability to erase records enables him to *pool with a younger honest type*. According to Bayes rule, this will lower player 1's reputation when he has a short good record, which will in turn increase the length of good record required for him to have a perfect reputation. An increase in player 1's survival probability will further increase the length of good record required for a perfect reputation since it increases the likelihood of the *old opportunistic type* relative to the *young honest type*.

The key step of my proof is to show that as the length of player 1's good record increases by 1, his reputation increases by at most $1-\overline{\delta}$. This finding is consistent with my intuition that it takes longer for player 1 to establish a reputation when his expected lifespan increases. The speed of reputation building in my model stands in contrast to models where player 1 cannot erase records (e.g., Fudenberg and Levine 1989), in which case player 1's reputation grows exponentially in periods where his opponents refuse to play a_2^*. Since player 1's reputation reaches 1 once the length of his good record reaches $t+1$, the upper bound on the speed of reputation building leads to a lower bound on t, which is an affine function of $(1-\overline{\delta})^{-1}$.

Once we fix player 1's time preference $\widehat{\delta}$ and let his survival probability $\overline{\delta}$ go to 1, the ratio between $1-\overline{\delta}$ and $1-\delta$ goes to 0. The opportunistic type must separate from the honest type with positive probability at time 0 in order to boost his initial reputation (i.e., his reputation at the null history h_*^0). This is because otherwise, the lower bound on t implied by the speed of reputation building will exceed the upper bound on t implied by the need to provide the opportunistic type incentives. The opportunistic type's equilibrium payoff must be close to his minmax value 0 since separating from the honest type at time 0 is optimal for him and by Proposition 2, his continuation value after separation equals his minmax value 0.

As for the second part, it is not obvious why consumer welfare is arbitrarily close to their welfare under complete information. This is because first, by Proposition 3, the opportunistic type will play a_1^* with probability bounded above 0 in many periods and second, histories $h_*^0, ..., h_*^{t-1}$ may occur with very high probability when player 1 is the opportunistic type due to his ability to erase past actions.

The key observation is that for any $k \in \mathbb{N}$, it can never be the case that history h_*^k occurs with a high probability *and* the opportunistic type plays a_1^* at h_*^k with a high probability. In fact, I show that *either* the probability of history h_*^k *or* the probability that the opportunistic type plays a_1^* at h_*^k is at most proportional to $1-\overline{\delta}$. I formally state this observation as Lemma 3, with proof in Section 3.5. Intuitively, when the opportunistic type *cannot* erase actions, h_*^k occurs with probability $(1-\overline{\delta})\overline{\delta}^k$, which vanishes to 0 as player 1 becomes infinitely long-lived. Similarly, h_*^k also occurs with a low probability when the opportunistic type erases actions at h_*^k with a low probability, which is the case when he takes action a_1^* with a high probability.

The above observation together with my earlier conclusion that t is bounded above by a linear function of $(1-\delta)^{-1}$ implies that when player 1's time preference is fixed but his expected lifespan diverges to infinity, the *average probability* with which the opportunistic type plays a_1^* vanishes to 0. This upper bound on the average probability of playing a_1^* leads to the conclusion in the second part of Theorem 2.

3.5 Proof of Theorem 2: Part 1

Let $\beta_k \in \Delta(A_2)$ denote player 2's action at h_*^k. The opportunistic type prefers to erase $\underline{a}_1$ at h_*^k if and only if:

$$\underbrace{u_1(\underline{a}_1, \beta_k) - c}_{\text{player 1's payoff from playing } \underline{a}_1 \text{ in every period and then erasing it}} \geq \underbrace{(1-\delta)u_1(\underline{a}_1, \beta_k)}_{\text{player 1's payoff from playing } \underline{a}_1 \text{ and not erasing it}},$$

or equivalently,

$$u_1(\underline{a}_1, \beta_k) \geq c/\delta. \tag{3.2}$$

Fix any equilibrium (σ_1, σ_2) in which the opportunistic type does *not* have a strict incentive to erase $\underline{a}_1$ after taking it at h_*^0. I show that the opportunistic type's payoff under (σ_1, σ_2) is no more than $(1-\delta)c/\delta$. This is because if the opportunistic type finds it weakly optimal *not* to erase $\underline{a}_1$ at h_*^0, then (3.2) implies that $u_1(\underline{a}_1, \beta_0) \leq c/\delta$. According to Proposition 2, his continuation value after he plays $\underline{a}_1$ (and does not erase it) is 0. As a result, his equilibrium payoff equals $(1-\delta)u_1(\underline{a}_1, \beta_0)$, which is no more than $(1-\delta)c/\delta$.

Next, suppose by way of contradiction that for every $\delta^* \in (0, 1)$, there exist $\overline{\delta} \in (\delta^*, 1)$ and an equilibrium (σ_1, σ_2) under $(u_1, u_2, c, \pi, \widehat{\delta}, \overline{\delta})$ such that the opportunistic type has a *strict* incentive to erase $\underline{a}_1$ after taking it at h_*^0. Part 1 of Theorem 2 is established once I find a contradiction and rule out such equilibria.

Step 1: I use the opportunistic type's incentive to take action a_1^* to derive an *upper bound* on the maximal length of good record that the opportunistic type may have in any equilibrium, which is denoted by t.

The definition of t implies that for every $k < t$, the opportunistic type plays a_1^* with strictly positive probability at h_*^k. Let V_k denote the opportunistic type's continuation value at h_*^k. If the opportunistic type strictly prefers to erase $\underline{a}_1$ after taking it at h_*^0, then by Theorem 1, he will erase $\underline{a}_1$ with probability 1 at h_*^k for every $k < t$. This implies that at every h_*^k with $k < t$, the opportunistic type is indifferent between *playing* a_1^* and *playing* $\underline{a}_1$ *and then erasing it*. This leads to the following expression for V_k:

$$V_k = u_1(\underline{a}_1, \beta_k) - c = (1-\delta)u_1(a_1^*, \beta_k) + \delta V_{k+1} \text{ for every } k < t. \tag{3.3}$$

Plugging $V_{k+1} = u_1(\underline{a}_1, \beta_{k+1}) - c$ into (3.3), we obtain that

$$u_1(\underline{a}_1, \beta_{k+1}) - u_1(\underline{a}_1, \beta_k) = (1-\delta)\big(u_1(\underline{a}_1, \beta_{k+1}) - c - u_1(a_1^*, \beta_k)\big)$$

$$> (1-\delta)\big(u_1(\underline{a}_1, \beta_k) - c - u_1(a_1^*, \beta_k)\big) \geq (1-\delta)(\overline{c} - c), \tag{3.4}$$

where the first inequality uses Assumption 1 and the conclusion in Theorem 1 that β_{k+1} FOSDs β_k and the second inequality uses Assumption 4. Since the opportunistic type's continuation value is at most $u_1(\underline{a}_1, a_2^*)$ and is at least 0, we have

$$t \leq \overline{T} \equiv (1-\delta)^{-1} \frac{u_1(\underline{a}_1, a_2^*)}{(\overline{c} - c)}. \tag{3.5}$$

Step 2: I use player 2's incentives to derive a *lower bound* on t. For every $k \leq t-1$, let μ_k^* denote the probability of history h_*^k *conditional on* player 1 being the opportunistic type. Let q_k^* denote the probability that the opportunistic type plays $\underline{a}_1$ at h_*^k and then erases it. Let p_k^* denote the probability with which the opportunistic type plays a_1^* at h_*^k. Since the game ends with probability $1-\overline{\delta}$ after each period, one can apply the state distribution lemma on page 2 of Online Appendix A and obtain that

$$\mu_0^* = (1-\overline{\delta}) + \overline{\delta}\mu_0^* q_0^* \quad \text{and} \quad \mu_k^* = \overline{\delta}(\mu_{k-1}^* p_{k-1}^* + \mu_k^* q_k^*) \text{ for every } k \in \{1, ..., t\}. \tag{3.6}$$

My hypothesis that the opportunistic type strictly prefers to erase action $\underline{a}_1$ after taking it at h_*^0 implies that $q_k^* = 1 - p_k^*$ for every $k \in \{0, ..., t\}$. Therefore, (3.6) implies that

$$\mu_0^* = \frac{1-\overline{\delta}}{1-\overline{\delta}(1-p_0^*)} \quad \text{and} \quad \frac{\mu_k^*}{\mu_{k-1}^*} = \frac{\overline{\delta} p_{k-1}^*}{1-\overline{\delta}(1-p_k^*)} \text{ for every } k \in \{1, ..., t\}. \tag{3.7}$$

Applying Bayes rule to player 1's probability of being the honest type, we obtain that

$$\frac{\pi_k}{1-\pi_k} = \frac{\pi}{1-\pi} \cdot \frac{(1-\overline{\delta})\overline{\delta}^k}{\mu_k^*} \quad \Rightarrow \quad \frac{\mu_k^*}{\mu_{k-1}^*} = \overline{\delta} \cdot \frac{1-\pi_k}{\pi_k} \cdot \frac{\pi_{k-1}}{1-\pi_{k-1}}. \tag{3.8}$$

According to Theorem 1, $\pi_{k-1} < \pi_k$ for every $k \leq t$. This together with (3.8) implies that $\mu_k^*/\mu_{k-1}^* \leq \overline{\delta}$. Combining this upper bound on μ_k^*/μ_{k-1}^* with its expression in (3.7), we obtain that

$$p_{k-1}^* \leq (1-\overline{\delta}) + \overline{\delta} p_k^* \quad \Rightarrow \quad p_{k-1}^* - p_k^* \leq (1-\overline{\delta})(1-p_k^*) \leq 1-\overline{\delta}. \tag{3.9}$$

Since the opportunistic type plays a_1^* with zero probability at history h_*^t, we have $p_t^* = 0$. This implies that

$$t \geq p_0^*(1-\overline{\delta})^{-1}. \tag{3.10}$$

I show that there exists $p^* \in (0,1)$ such that $p_0^* > p^*$ for *all* $\overline{\delta} \in (0,1)$ that is greater than some cutoff. Let x_k denote the probability with which player 2's belief at h_*^k assigns to player 1's current-period action being a_1^*. According to Bayes rule, we have:

$$\frac{x_k}{1-x_k} = \frac{\pi(1-\overline{\delta})\overline{\delta}^k + (1-\pi)\mu_k^* p_k^*}{(1-\pi)\mu_k^*(1-p_k^*)},$$

or equivalently,

$$\frac{\pi}{1-\pi}(1-\overline{\delta})\overline{\delta}^k = \mu_k^* \Big\{ \frac{x_k}{1-x_k}(1-p_k^*) - p_k^* \Big\} = \mu_k^* \frac{x_k - p_k^*}{1-x_k}. \tag{3.11}$$

Take $k=0$ in (3.11), replace μ_0^* with its expression in (3.7), and divide both sides by $1-\overline{\delta}$, we have

$$\frac{\pi}{1-\pi} = \frac{1}{1-\overline{\delta}(1-p_0^*)} \cdot \frac{x_0 - p_0^*}{1-x_0}. \tag{3.12}$$

Recall the normalization that $u_1(\underline{a}_1, \underline{a}_2) = 0$. The hypothesis that the opportunistic type strictly prefers to erase $\underline{a}_1$ at h_*^0 implies that $u_1(\underline{a}_1, \beta_0) > c/\delta \geq 0$. Under Assumption 1, this is the case *only if* player 2's action β_0 assigns positive probability to some $a_2' \in A_2$ that is strictly greater than $\underline{a}_2$. Since $\underline{a}_2$ is a strict best reply to $\underline{a}_1$, there exists $x^* > 0$ such that player 2 has a strict incentive to play $\underline{a}_2$ as long as player 1 plays a_1^* with probability less than x^*. Player 2's incentive to play a_2' implies that $x_0 \geq x^*$. Since the RHS of (3.12) diverges to $+\infty$ as $p_0^* \to 0$ and $\overline{\delta} \to 1$, for every $\pi \in (0,1)$, there exist $\delta^*, p^* \in (0,1)$ such that

$$\frac{\pi}{1-\pi} < \frac{1}{1-\delta^*(1-p^*)} \cdot \frac{x^* - p^*}{1-x^*}. \tag{3.13}$$

Since the RHS of (3.12) is strictly decreasing in p_0^*, is strictly increasing in x_0, and is strictly increasing in $\overline{\delta}$, inequality (3.13) implies that as long as $\overline{\delta} \geq \delta^*$ and $x_0 \geq x^*$, (3.12) does not hold when $p_0^* \leq p^*$. Hence, it must be the case that $p_0^* > p^*$ for every $\overline{\delta} > \delta^*$, that is to say, p_0^* is uniformly bounded from below by p^*.

For any $\pi \in (0,1)$ and $\widehat{\delta} \in (0,1)$, there exists $\delta^* \in (0,1)$ such that when $\overline{\delta} > \delta^*$, the upper bound on t obtained in (3.5) is strictly less than the lower bound on t obtained in (3.10). This contradiction rules out equilibria in which the opportunistic type has a strict incentive to erase $\underline{a}_1$ at h_*^0. As a result, the opportunistic type's discounted average payoff in every equilibrium is no more than $(1-\delta)c/\delta$.

3.6 Proof of Theorem 2: Part 2

Step 1: Recall the definitions of t and t_0 in the statement of Theorem 1. I show that there exist $m \in \mathbb{N}$ that depends only on (π, u_1, u_2) and $\lambda > 0$ that depends only on (u_1, u_2) such that $t \leq m + \lambda(1-\delta)^{-1}$.

According to (3.4), the opportunistic type's incentive to play a_1^* from $h_*^{t_0}$ to h_*^{t-1} implies that $t - t_0$ is bounded above by an affine function of $(1-\delta)^{-1}$, with the coefficients of this function depending only on (u_1, u_2). At every h_*^k with $k < t_0$, player 1 has no incentive to erase any action at h_*^k and player 2 takes actions other than a_2^* with strictly positive probability. Let $\gamma \in (0,1)$ be such that a_2^* is a strictly best reply to all mixed actions that assign probability more than γ to a_1^*. Hence, the unconditional probability that player 1 plays a_1^* at h_*^k is no more than γ for every $k < t_0$. According to Bayes rule, player 1's reputations at h_*^k and h_*^{k+1}, denoted by π_k and π_{k+1}, satisfy $\pi_{k+1} \geq \pi_k/\gamma$ for every $k < t_0$. Since player 1's reputation at $h_*^{t_0}$ is no more than 1, we know that $t_0 \leq m \equiv \frac{\log \pi}{\log \gamma}$. The two parts together imply that $t \leq m + \lambda(1-\delta)^{-1}$.

Step 2: I bound the short-run players' welfare from above by deriving an upper bound on the opportunistic type's *average probability* of playing a_1^*. Fix any equilibrium. For every $k \in \mathbb{N}$, I define event $\mathcal{E}^k$ as

$$\mathcal{E}^k \equiv \{\text{the current history being } h_*^k \text{ and player 1 playing } a_1^* \text{ in the current period}\}.$$

Lemma 3. *In every equilibrium and for every $k \in \mathbb{N}$, the probability of event $\mathcal{E}^k$ conditional on player 1 being the opportunistic type is no more than $(1-\overline{\delta})\overline{\delta}^k$.*

Proof. By definition, the probability of event $\mathcal{E}^k$ conditional on player 1 being the opportunistic type is $\mu_k^* p_k^*$. I show by induction that $\mu_k^* p_k^* \leq (1-\overline{\delta})\overline{\delta}^k$ for every $k \in \mathbb{N}$. The first part of (3.7) implies that:

$$\mu_0^* p_0^* = \frac{(1-\overline{\delta})p_0^*}{1-\overline{\delta}q_0^*} \leq \frac{(1-\overline{\delta})p_0^*}{1-\overline{\delta}(1-p_0^*)} \leq 1-\overline{\delta} \text{ for every } p_0^* \in [0,1].$$

Next, suppose $\mu^*_{j-1}p^*_{j-1} \leq (1-\overline{\delta})\overline{\delta}^{j-1}$ for some $j \geq 1$, then the second part of (3.7) implies that

$$\mu^*_j p^*_j = \frac{\overline{\delta} p^*_j p^*_{j-1}\mu^*_{j-1}}{1-\overline{\delta} q^*_j} \leq \frac{\overline{\delta} p^*_j p^*_{j-1}\mu^*_{j-1}}{1-\overline{\delta}(1-p^*_j)} \leq \overline{\delta}^j \frac{(1-\overline{\delta})p^*_j}{1-\overline{\delta}(1-p^*_j)} \leq (1-\overline{\delta})\overline{\delta}^j \text{ for every } p^*_j \in [0,1].$$

□

Conditional on player 1 being the opportunistic type, the probability of event $\mathcal{E}^k$ is 0 for every $k > t$ since the opportunistic type never reaches h^{t+1}_*. The *ex ante* probability that the opportunistic type takes action a^*_1 is $\sum_{k=0}^t \Pr(\mathcal{E}^k)$, which by Lemma 3, is no more than $1-\overline{\delta}^t$ in any equilibrium. Fix any equilibrium (σ_1, σ_2) as well as the resulting distribution over player 1's actions. Player 2's payoff is no more than her payoff when she can observe player 1's *realized pure action* before choosing her action, which is at most

$$(1-\pi)\Big\{1-(1-\overline{\delta}^t)\Big\}u_2(\underline{a}_1,\underline{a}_2)+\Big\{\pi+(1-\overline{\delta}^t)(1-\pi)\Big\}u_2(a^*_1,a^*_2) = \Big\{\pi+(1-\overline{\delta}^t)(1-\pi)\Big\}u_2(a^*_1,a^*_2). \quad (3.14)$$

Since t is bounded above by an affine function of $(1-\delta)^{-1}$, for every $\widehat{\delta} \in (0,1)$ and $\varepsilon > 0$, there exists $\delta^* \in (0,1)$ such that when $\overline{\delta} > \delta^*$, the value of (3.14) is no more than $\pi u_2(a^*_1, a^*_2) + \varepsilon$.

Step 3: I bound the short-run players' welfare from below. Fix any equilibrium. Suppose player 2 deviates by playing a^*_2 at h^k_* for every $k \geq t+1$ and playing $\underline{a}_2$ at any other history. Since player 1's strategy does not depend on player 2's actions, the expected value of $\sum_{k=0}^{+\infty}(1-\overline{\delta})\overline{\delta}^k u_2(a_{1,k}, a_{2,k})$ under this deviation, denoted by $\underline{U}_2$, is weakly lower than its expected value in equilibrium. Since the honest type reaches record h^{t+1}_* in period $t+1$ and the opportunistic type never reaches h^{t+1}_*, we obtain that

$$\underline{U}_2 \geq \Big\{(1-\overline{\delta}^{t+1})\pi + (1-\pi)\Big\}u_2(\underline{a}_1,\underline{a}_2) + \overline{\delta}^{t+1}\pi u_2(a^*_1,a^*_2) = \overline{\delta}^{t+1}\pi u_2(a^*_1,a^*_2). \quad (3.15)$$

Inequality (3.15) relies on $u_2(a_1,a_2)$ being weakly increasing in a_1 since it ensures that $u_2(a^*_1,\underline{a}_2) \geq u_2(\underline{a}_1,\underline{a}_2) \equiv 0$. Recall from the first step of this proof that $t \leq m + \lambda(1-\delta)^{-1}$. For every $\widehat{\delta} \in (0,1)$ and $\varepsilon > 0$, there exists $\delta^* \in (0,1)$ such that when $\overline{\delta} > \delta^*$, the RHS of (3.15) is at least $\pi u_2(a^*_1,a^*_2) - \varepsilon$.

4 Concluding Remarks

I take a first step to analyze reputation formation when a patient player's record length is determined endogenously by his own strategic behaviors. Although the patient player has a strict incentive to take his strictly dominant action when he cannot build reputations, he will have an incentive to take the commitment action for a long time when he can build a reputation. This is because a longer good record signals his honesty and leads to a higher continuation value. In fact, the patient player can secure a continuation value that is strictly greater than his commitment payoff once the length of his good record exceeds some cutoff.

However, as long as the patient player is sufficiently long-lived, his equilibrium payoff must be close to his minmax value. This is because the opportunistic type's ability to erase records enables him to pool with a younger honest type, which lowers his reputation when he has a short good record. A longer expected lifespan (or equivalently, a higher survival probability) further slows down the reputation building process since it increases the likelihood ratio between the old opportunistic type and the young honest type.

My results imply that (i) the possibility of erasing actions *cannot* eliminate patient players' incentives to build reputations, (ii) a small probability of opportunistic type can entirely wipe out the patient player's returns from building reputations, although it has a negligible effect on the short-run players' welfare. I conclude with a discussion of my modeling assumptions and results beyond those discussed in Section 2.1.

Time Preference vs Survival Probability: Theorem 2 shows that regardless of the probability of honest type, as long as player 1's survival probability $\overline{\delta}$ goes to 1 much faster his time preference $\widehat{\delta}$ goes to 1, his equilibrium payoff must be close to his minmax value 0. In what follows, I consider the opposite case where $\widehat{\delta}=1$ but $\overline{\delta}$ is less than 1. The results that I derived here also apply to the case where $\widehat{\delta}$ goes to 1 much faster than $\overline{\delta}$ goes to 1. I start from a result which shows that my reputation failure result (Theorem 2) extends to the case where $\widehat{\delta}=1$ as long as the probability of the honest type is lower than some cutoff.

Proposition 4. *Suppose* $\widehat{\delta}=1$ *and* (u_1,u_2,c) *satisfies Assumptions 1, 2, 3, and 4. There exists* $\overline{\pi}\in(0,1)$ *such that for every* $\pi\in(0,\overline{\pi})$ *and* $\overline{\delta}\in(0,1)$*, the opportunistic type's discounted average payoff in every equilibrium is no more than* $(1-\delta)c/\delta$*.*

The proofs of this result and the next one are in Appendices D and E. According to Proposition 4, as long as the probability of the honest type is lower than some cutoff $\overline{\pi}$, player 1's equilibrium payoff is arbitrarily close to his minmax value 0 even when his effective discount factor $\delta\equiv\widehat{\delta}\cdot\overline{\delta}$ is arbitrarily close to 1. This stands in contrast to the result in Fudenberg and Levine (1989, also see Proposition 1 of the current paper), that fix any probability of the honest type $\pi>0$ and any $\varepsilon>0$, player 1 can secure a payoff that is ε-close to his commitment payoff $u_1(a_1^*,a_2^*)$ when his effective discount factor δ is above some cutoff.

Next, I state a result which shows that player 1 can receive a high payoff in *some* equilibria when the probability of honest type π is high enough. This conclusion stands in contrast to the one in Theorem 2 that for every $\widehat{\delta}\in(0,1)$ and $\pi\in(0,1)$, the long-run player's payoff in every equilibrium is no more than $(1-\delta)c/\delta$ when $\overline{\delta}$ is arbitrarily close to 1. This comparison highlights the different roles of player 1's time preference $\widehat{\delta}$ and his survival probability $\overline{\delta}$ in determining the game's equilibrium outcomes, which is a novel feature compared to existing reputation models such as the one in Fudenberg and Levine (1989). This also justifies my choice of modeling players' time preference and survival probability separately, since they are not only conceptually different but also have different implications on the equilibrium outcomes.

Proposition 5. *Suppose* $\widehat{\delta} = 1$ *and* (u_1, u_2, c) *satisfies Assumptions 1, 2, 3, and 4. For every* $v < u_1(\underline{a}_1, a_2^*) - c$*, there exist* $\pi^* \in (0,1)$ *and* $\delta^* \in (0,1)$ *such that for every* $\pi > \pi^*$ *and* $\overline{\delta} > \delta^*$*, there exists an equilibrium in which the opportunistic type's payoff is at least* v*.*

Stochastic Arrivals & Stochastic Reviews: My baseline model assumes that a consumer arrives in each period, in which case the *honest type*'s age coincides with the number of actions in his record. In Online Appendix E, I discuss extensions when the consumers either *arrive stochastically* or *do not post reviews with positive probability*. In these cases, the honest type's age may not coincide with the number of actions in his record. Nevertheless, the qualitative features of the equilibria and the results remain unchanged.

The Honest Type Erasing Actions: In my baseline model, the honest type cannot erase any action. My results can be extended to the case where the honest type mechanically takes action a_1^* in every period, as in canonical reputation models, but he can strategically decide whether to erase his action in order to maximize his discounted average payoff, defined as

$$\sum_{k=0}^{+\infty}(1-\delta)\delta^k\Big\{u_1(a_1^*, a_{2,k}) - c_k\Big\}.$$

Proposition 6. *Suppose* (u_1, u_2, c) *satisfies Assumptions 1, 2, 3, and 4 and* $c > 0$*, then the honest type will never erase his action at any on-path history of any equilibrium even when he has the option to do so.*

The proof is in Appendix F. This result does not rely on both types sharing the same stage-game payoff function. It remains valid when the honest type's stage-game payoff $\widetilde{u}_1(a_1, a_2)$ is different from that of the opportunistic type's $u_1(a_1, a_2)$, as long as $\widetilde{u}_1(a_1, a_2)$ is also strictly increasing in a_2.

Erasing Past Actions: My baseline model assumes that the patient player can only erase $a_{1,k}$ by the end of period k but cannot do so after period k. My result for the complete information game, Proposition 2, continues to hold when the patient player also has the ability to erase past actions.

However, my characterization results for the incomplete information game, Theorems 1 and 2, rely on this simplifying assumption, under which the opportunistic type's continuation value equals his minmax value 0 after he separates from the honest type. Intuitively, when player 1 can erase past actions, his continuation value at histories that contain actions other than a_1^* might be strictly positive since he can erase these actions in the future, after which the short-run players will still assign strictly positive probability to the honest type. Such a possibility makes the model intractable.

The Uniqueness & Multiplicity of Equilibrium: Theorem 1 characterizes the common properties of all equilibria and Theorem 2 provides a sharp characterization of players' equilibrium payoffs. One may wonder whether there is a unique equilibrium, and if not, what are the sources for multiplicity.

Recall the definition of $\mathcal{B}$ in (2.1) and that by Lemma 2, every pair of elements in $\mathcal{B}$ can be ranked according to FOSD. Recall that h_*^k denotes player 2's history in which there are k actions, all of which are a_1^*, μ_k^* denotes the probability of h_*^k conditional on player 1 being the opportunistic type, π_k denotes player 1's reputation at h_*^k, β_k denotes player 2's action at h_*^k, and p_*^k denotes the probability with which the opportunistic type plays a_1^* at h_*^k. Fix any $\beta_0 \in \mathcal{B}$, player 2's actions when player 1 has a positive reputation, $\beta_1, \beta_2, ..., \beta_t$, are pinned down by player 1's indifference condition at h_*^k for $k \in \{0, 1, 2, ..., t-1\}$:

$$\underbrace{\max\Big\{u_1(\underline{a}_1, \beta_k) - c, (1-\delta)u_1(\underline{a}_1, \beta_k)\Big\}}_{\text{player 1's continuation value at } h_*^k} = (1-\delta)u_1(a_1^*, \beta_k) + \delta \underbrace{\max\Big\{u_1(\underline{a}_1, \beta_{k+1}) - c, (1-\delta)u_1(\underline{a}_1, \beta_{k+1})\Big\}}_{\text{player 1's continuation value at } h_*^{k+1}}.$$

This recursive process also pins down the value of t since β_t must be weakly lower than a_2^* but must be high enough so that player 1 does not have an incentive to play a_1^* at h_*^t.

Let $\beta^\dagger \in \mathcal{B}$ be such that $u_1(\underline{a}_1, \beta^\dagger) = c/\delta$. When $c < \overline{c}$, such an action exists when $\delta > \frac{c}{u_1(\underline{a}_1, a_2^*)}$ and is unique by Lemma 2. One can also show that $\beta^\dagger$ is nontrivially mixed when δ is large enough. To see this, I consider two cases. If there exists a pure action $\beta \in \mathcal{B}$ such that $u_1(\underline{a}_1, \beta) = c$, then $\beta^\dagger$ must be nontrivially mixed when δ is close to 1. If the unique $\beta \in \mathcal{B}$ that satisfies $u_1(\underline{a}_1, \beta) = c$ is nontrivially mixed, then a continuity argument implies that $\beta^\dagger$ is also nontrivially mixed for every δ close enough to 1.

When $\overline{\delta}$ and $\widehat{\delta}$ are bounded below 1, player 2's action in period 0 can be bounded below $\beta^\dagger$. If her action in period 0 is a pure action, then there are multiple probabilities with which player 1 can play $\underline{a}_1$ in period 0, leading to a multiplicity of equilibrium outcomes.

Fix any π. When δ is close enough to 1, it must be the case that $\beta_0 = \beta^\dagger$ or β_0 is close to $\beta^\dagger$. This is because the speed with which β increases in t is proportional to $1-\delta$ and similar to Fudenberg and Levine (1989), the speed with which player 1's reputation increases when $\beta < \beta^\dagger$ is bounded above 0. If it takes too many periods for β to reach $\beta^\dagger$, then player 1's reputation will exceed 1 before β reaches $\beta^\dagger$, which will lead to a contradiction. If π is small enough such that β_0 is strictly lower than $\beta^\dagger$, then even when both $\beta^\dagger$

and β_0 are nontrivially mixed, there may exist multiple values of β_0 in equilibrium, which is another source of multiplicity. However, as long as δ is close to 1, β_0 must be close to $\beta^\dagger$, and the values of $\beta_1^*, ..., \beta_t^*$, $p_0^*, ..., p_t^*$, and $\mu_0^*, ..., \mu_t^*$ are also close across different equilibria.

When π is above some cutoff and δ is close to 1, I can show that $\beta_0 = \beta^\dagger$ in all equilibria, from which I can pin down the values of t as well as $\beta_1, \beta_2, ..., \beta_{t-1}, \beta_t$. If all of $\beta_0, ..., \beta_{t-1}, \beta_t$ are nontrivially mixed, which happens under generic δ, then the conclusion that $p_t^* = 0$ as well as player 2's indifference conditions pin down the values of $p_0^*, ..., p_t^*$ and $\mu_0^*, ..., \mu_t^*$. When some actions in $\{\beta_0, \beta_1, ..., \beta_{t-1}, \beta_t\}$ are pure actions, there are multiple actions of player 1's under which player 2 has an incentive to play that pure action. This implies that there are multiple values of $p_0^*, ..., p_t^*$ and $\mu_0^*, ..., \mu_t^*$ that can satisfy player 2's incentive constraints, leading to multiple equilibrium outcomes. However, even at these degenerate values of δ where multiple equilibrium outcomes occur, the equilibrium values of $p_0^*, ..., p_t^*$ and $\mu_0^*, ..., \mu_t^*$ are pinned down except for periods in which player 2 takes a pure action in equilibrium.

The Imperfect Promise Interpretation & Information Disclosure Policies: Recall the *imperfect promise* interpretation in Section 2 that before the game starts, player 1 promises to all short-run players that he will always take action a_1^* and that he will disclose all his past actions. Player 2 believes that player 1 will honor his promise with probability π. With complementary probability, player 1 can *renege* in the sense that he may take actions other than a_1^* and may also erase his actions. Whether player 1 can renege is determined by a random draw after he makes his promise and is fixed over time. Under this interpretation, Theorem 2 implies that when player 1 is sufficiently long-lived, his benefit from making such a promise is entirely wiped out even when his opponents only entertain a grain of doubt about his willingness to honor his promise.

In Online Appendix C, I examine a natural follow-up question, namely, can player 1 obtain higher payoffs from *alternative promises* when his opponents believe that he will renege with positive probability?

In the spirit of commitment types in the reputation literature, I still assume that player 1 commits to play a_1^* in every period. However, instead of committing to fully disclose all of his past actions, he can commit to alternative *information disclosure policies*, such as only disclosing his last K actions, disclosing his last K actions with probability $1/2$ and disclosing nothing otherwise, and so on. My analysis and conclusion in that appendix also apply to several alternative scenarios, such as when an online platform commits to reveal at most $K \in \mathbb{N}$ unerased actions from each seller (i.e., K reviews each seller received) to the consumers.

When the patient player is sufficiently long-lived, I show that under a large class of disclosure policies, the opportunistic type's payoff is no more than his payoff in the game where the patient player commits to reveal *no action* to his opponents. My finding implies that a long-lived player can benefit from committing to alternative information disclosure policies (relative to fully disclosing all past actions) *only if* his opponents believe that he will honor his promise with probability above some cutoff. It also implies that as long as a sufficiently long-lived player can selectively erase his records with positive probability, he *cannot* benefit from being monitored in the sense that his payoff under any disclosure policy is no more than his payoff under no disclosure. This stands in contrast to the usual lessons from the theories of repeated games, that a sufficiently long-lived and patient player can attain strictly higher payoffs (at least in some equilibria) when his actions are being monitored compared to the case in which his actions cannot be monitored at all.

A Proof of Lemma 2

First, I show that for every $\lambda \in [0,1]$, player 2 has at most two pure-strategy best replies to player 1's mixed action $\lambda a_1^* + (1-\lambda)\underline{a}_1$. Suppose by way of contradiction that there exist a_2, a_2', a_2'' with $a_2 \succ a_2' \succ a_2''$ and $\lambda^* \in [0,1]$ such that a_2, a_2', a_2'' best reply to $\lambda^* a_1^* + (1-\lambda^*)\underline{a}_1$ Then the last part of Assumption 3 implies that there exist $\lambda, \lambda', \lambda'' \in [0,1]\backslash\{\lambda^*\}$ where $\lambda, \lambda', \lambda''$ *can be the same* such that a_2 best replies to $\lambda a_1^* + (1-\lambda)\underline{a}_1$, a_2' best replies to $\lambda' a_1^* + (1-\lambda')\underline{a}_1$, and a_2'' best replies to $\lambda'' a_1^* + (1-\lambda'')\underline{a}_1$.

Since $\lambda \neq \lambda^*$, $\lambda' \neq \lambda^*$, and $\lambda'' \neq \lambda^*$, we know that *either* at least two of λ, λ' and λ'' are strictly more than λ^*, *or* at least two of λ, λ' and λ'' are strictly less than λ^*. In the first case, a_2 best replies to $\lambda^* a_1^* + (1-\lambda^*)\underline{a}_1$ and at least one of a_2' and a_2'' best replies to an action that FOSDs $\lambda^* a_1^* + (1-\lambda^*)\underline{a}_1$. This is because $\lambda' a_1^* + (1-\lambda')\underline{a}_1$ FOSDs $\lambda^* a_1^* + (1-\lambda^*)\underline{a}_1$ when $\lambda' > \lambda^*$ and $\lambda'' a_1^* + (1-\lambda'')\underline{a}_1$ FOSDs $\lambda^* a_1^* + (1-\lambda^*)\underline{a}_1$ when $\lambda'' > \lambda^*$. Since a_2' and a_2'' are strictly lower than a_2, this contradicts Assumption 1 that $u_2(a_1, a_2)$ has strictly increasing differences. Using a symmetric argument, one can also derive a contradiction in the second case where at least two of λ, λ' and λ'' are strictly less than λ^*. This implies that for every $\lambda \in [0,1]$, player 2 has at most two pure-strategy best replies to $\lambda a_1^* + (1-\lambda)\underline{a}_1$.

Next, let A_2^* denote the set of player 2's *pure* best replies against mixtures between a_1^* and $\underline{a}_1$:

$$A_2^* \equiv \{a_2 \in A_2 | \text{there exists } \lambda \in [0,1] \text{ s.t. } a_2 \text{ best replies to } \lambda a_1^* + (1-\lambda)\underline{a}_1\}. \tag{A.1}$$

Recall that $u_2(a_1, a_2)$ has strictly increasing differences. Since a_2^* best replies to a_1^*, $\underline{a}_2$ best replies to $\underline{a}_1$, and $a_1^* \succ \underline{a}_1$, we know that a_2^* is the highest action in A_2^* and $\underline{a}_2$ is the lowest action in A_2^*. Assumption 3 then implies that there exist $0 \equiv \lambda_0 < \lambda_1 < ... < \lambda_n \equiv 1$ such that for every $a_2 \in A_2^*$, there exists $j \in \{1, 2, ..., n\}$ such that a_2 is a strict best reply to $\lambda a_1^* + (1-\lambda)\underline{a}_1$ if and only if $\lambda \in (\lambda_{j-1}, \lambda_j)$.

To see why this is true as well as how the cutoffs for λ are constructed, note that the first part of Assumption 3 implies that player 2 has a strict best reply $\underline{a}_2$ to $\underline{a}_1$, which implies that $\underline{a}_2$ also best replies to

$\lambda a_1^* + (1-\lambda)\underline{a}_1$ when λ is close to 0. Let λ_1 be the largest λ such that $\underline{a}_2$ best replies to $\lambda a_1^* + (1-\lambda)\underline{a}_1$. Since $a_1^* \succ_1 \underline{a}_1$ and $a_2^* \succ_2 \underline{a}_2$, we know that $\lambda_1 < 1$. By definition, $\underline{a}_2$ does not best reply to $\lambda a_1^* + (1-\lambda)\underline{a}_1$ for any $\lambda = \lambda_1 + \varepsilon$ with $\varepsilon > 0$. The upper-hemi-continuity of best reply correspondences implies that there exists $a_2' \succ_2 \underline{a}_2$ that best replies to $\lambda_1 a_1^* + (1-\lambda_1)\underline{a}_1$. Since player 2 has at most two pure best replies to every $\lambda a_1^* + (1-\lambda)\underline{a}_1$, such $a_2' \succ_2 \underline{a}_2$ is unique. The second part of Assumption 3 then implies that a_2' must also best reply to some $\lambda a_1^* + (1-\lambda)\underline{a}_1$ with $\lambda \neq \lambda_1$. Since $u_2(a_1, a_2)$ has strictly increasing differences, such λ must be strictly greater than λ_1. Let λ_2 denote the largest λ such that a_2' best replies to $\lambda a_1^* + (1-\lambda)\underline{a}_1$... Iterate this process, we can obtain the cutoffs, denoted by $\lambda_1, \lambda_2, ...$, until λ reaches 1.

The above construction implies that for every $j \in \{1, 2, ..., n-1\}$, player 2 has two pure-strategy best replies to $\lambda_j a_1^* + (1-\lambda_j)\underline{a}_1$, which are her strict best replies when $\lambda \in (\lambda_{j-1}, \lambda_j)$ and when $\lambda \in (\lambda_j, \lambda_{j+1})$, respectively. Therefore, $\mathcal{B}$ consists of all actions in A_2^* and all mixtures between pairs of *adjacent elements* in A_2^*. Hence, every pair of elements in set $\mathcal{B}$ can be ranked according to FOSD. Since $u_1(a_1, a_2)$ is strictly increasing in a_2, for every $a_1 \in A_1$ and $v \in [u_1(a_1, \underline{a}_2), u_1(a_1, a_2^*)]$, there exists a unique element $\beta \in \mathcal{B}$ that satisfies $u_1(a_1, \beta) = v$.

Remark: Let us focus on the set of (u_1, u_2) that satisfies Assumptions 1, 2, and the first part of Assumption 3. I explain why the second part of Assumption 3 is satisfied for *generic* (u_1, u_2) that belongs to this set.

Suppose there exists $a_2 \in A_2$ such that there exists a unique $\lambda \in [0, 1]$ such that a_2 best replies to $\lambda a_1^* + (1-\lambda)\underline{a}_1$. I consider three cases. First, suppose $\lambda = 0$. Then $\lambda a_1^* + (1-\lambda)\underline{a}_1$ is a pure action $\underline{a}_1$ and there exists a sequence $\{\widetilde{\lambda}_n\}_{n\in\mathbb{N}}$ such that $\widetilde{\lambda}_n > 0$ and $\lim_{n\to+\infty} \widetilde{\lambda}_n = 0$. The upper-hemi-continuity of best reply correspondences implies that any limit point of the best replies to $\{\widetilde{\lambda}_n a_1^* + (1-\widetilde{\lambda}_n)\underline{a}_1\}_{n\in\mathbb{N}}$, denoted by $\widetilde{a}_2$, is a best reply to $\lambda a_1^* + (1-\lambda)\underline{a}_1$. Since a_2 only best replies to $\underline{a}_1$ and $\widetilde{a}_2$ best replies to some $\lambda a_1^* + (1-\lambda)\underline{a}_1$ with $\lambda > 0$, we know that $a_2 \neq \widetilde{a}_2$. This contradicts the first part of Assumption 3 that player 2 has a unique best reply against each pure action of player 1's. Similarly, $\lambda \neq 1$.

Next, consider the case in which $\lambda \in (0, 1)$. Let $\{\widehat{\lambda}_n\}_{n\in\mathbb{N}}$ and $\{\widetilde{\lambda}_n\}_{n\in\mathbb{N}}$ denote two sequences such that $\widehat{\lambda}_n < \lambda$, $\lim_{n\to+\infty} \widehat{\lambda}_n = \lambda$, $\widetilde{\lambda}_n > \lambda$, and $\lim_{n\to+\infty} \widetilde{\lambda}_n = \lambda$. The upper-hemi-continuity of best reply correspondences implies that any limit of the best replies to $\widehat{\lambda}_n a_1^* + (1-\widehat{\lambda}_n)\underline{a}_1$, denoted by $\widehat{a}_2$, and any limit of the best replies to $\widetilde{\lambda}_n a_1^* + (1-\widetilde{\lambda}_n)\underline{a}_1$, denoted by $\widetilde{a}_2$, are best replies to $\lambda a_1^* + (1-\lambda)\underline{a}_1$. Since a_2 only best replies to $\lambda a_1^* + (1-\lambda)\underline{a}_1$, we know that $a_2 \neq \widetilde{a}_2$ and $a_2 \neq \widehat{a}_2$. Since $u_2(a_1, a_2)$ has strictly increasing differences, we know that $\widetilde{a}_2 \succ a_2 \succ \widehat{a}_2$. Therefore, player 2 has three-strategy pure best replies to mixed action $\lambda a_1^* + (1-\lambda)\underline{a}_1$. Depict player 2's payoff from each of her pure actions as a function of the probability that player 1 plays a_1^* (as opposed to $\underline{a}_1$), having three pure best replies to $\lambda a_1^* + (1-\lambda)\underline{a}_1$

implies that three of these affine functions intersect at λ. This can only happen at degenerate (u_1, u_2).

B Proof of Theorem 1

Fix any equilibrium (σ_1, σ_2). Recall that at any $h \in \mathcal{H}(\sigma_1, \sigma_2)$, player 1 never takes actions other than a_1^* and $\underline{a}_1$. Hence, player 2's action at every $h \in \mathcal{H}(\sigma_1, \sigma_2)$ belongs to $\mathcal{B}$. By Lemma 2, any pair of elements in $\mathcal{B}$ can be ranked according to FOSD. Proposition 2 implies that at every $h \in \mathcal{H}(\sigma_1, \sigma_2) \backslash \mathcal{H}_*$, players take actions $(\underline{a}_1, \underline{a}_2)$, player 1's continuation value equals his minmax value $u_1(\underline{a}_1, \underline{a}_2) = 0$, and therefore, he has no incentive to pay cost $c > 0$ to erase any action at any such h. Let V_k denote player 1's continuation value at h_*^k and let $\overline{V} \equiv \sup_{k \in \mathbb{N}} V_k$. I show that (σ_1, σ_2) satisfies all the properties listed in Theorem 1.

Step 1: I show that $\overline{V} > 0$. Suppose by way of contradiction that $\overline{V} = 0$. Since the opportunistic type can obtain a strictly positive payoff at any history h_*^k where player 2's action is strictly greater than $\underline{a}_2$, by playing $\underline{a}_1$ and not erasing it, $\overline{V} = 0$ only if player 2 plays $\underline{a}_2$ at every $h \in \mathcal{H}_*$. Since player 2 will also play $\underline{a}_2$ at every $h \in \mathcal{H}(\sigma_1, \sigma_2) \backslash \mathcal{H}_*$, the opportunistic type will have no intertemporal incentive and will play $\underline{a}_1$ at h_*^0. By Bayes rule, player 2 will assign probability 1 to the honest type at history h_*^1, which implies that she will have a strict incentive to play $a_2^* \neq \underline{a}_2$ at h_*^1. This leads to a contradiction and implies that $\overline{V} > 0$.

Step 2: I show that there exists $t \in \mathbb{N}$ such that (i) the opportunistic type plays a_1^* with positive probability at h_*^k if and only if $k < t$, (ii) $\pi_k = 1$ if and only if $k > t$, and (iii) $\beta_{t+1} = a_2^*$ and $V_{t+1} = \max\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta) u_1(\underline{a}_1, a_2^*)\}$. Since $\overline{V} > 0$, fix any ε that satisfies

$$0 < \varepsilon < \min\left\{\frac{\overline{V}}{2}, \frac{(1-\delta)(\overline{c} - c)}{\delta}\right\},$$

and any $h \in \mathcal{H}(\sigma_1, \sigma_2)$ that satisfies $V(h) > \overline{V} - \varepsilon$. Using the same argument as in the proof of Proposition 2, we know that the opportunistic type plays a_1^* with zero probability at h, since it is strictly dominated by playing $\underline{a}_1$ and then erasing it. The requirement on ε implies that $V(h) > 0$, and therefore, there exists $k \in \mathbb{N}$ such that $h = h_*^k$. Let $t \in \mathbb{N}$ be the *smallest* integer $k \in \mathbb{N}$ such that the opportunistic type plays a_1^* with zero probability at h_*^k. Such an integer t exists and by construction, satisfies the first requirement. According to Bayes rule, t also satisfies the second requirement. When $k = t+1$, $\beta_k = a_2^*$ since player 2's posterior belief assigns probability 1 to the honest type at h_*^{t+1}. The opportunistic type's continuation value at h_*^{t+1} is $\max\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta) u_1(\underline{a}_1, a_2^*)\}$. This is because playing a_1^* is not optimal for him at h_*^{t+1} and that his payoff is $\max\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta) u_1(\underline{a}_1, a_2^*)\}$ by playing $\underline{a}_1$.

Step 3: I show that for every $k < t$, the opportunistic type plays $\underline{a}_1$ with strictly positive probability at h_*^k. Suppose by way of contradiction that the opportunistic type plays a_1^* with probability 1 at h_*^k for some $k < t$. Then player 2 strictly prefers to play a_2^* at h_*^k since both the honest type and the opportunistic type play a_1^* at h_*^k. The opportunistic type's incentive to play a_1^* instead of $\underline{a}_1$ at h_*^k implies that

$$V_k = (1-\delta)u_1(a_1^*, a_2^*) + \delta V_{k+1} \geq \max\Big\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta)u_1(\underline{a}_1, a_2^*)\Big\}, \tag{B.1}$$

where $u_1(\underline{a}_1, a_2^*) - c$ is his continuation value when he plays $\underline{a}_1$ and then erases it (and does this in every subsequent period), and $(1-\delta)u_1(\underline{a}_1, a_2^*)$ is his continuation value when he plays $\underline{a}_1$ and does not erase it.

Since $c < \overline{c} \leq u_1(\underline{a}_1, a_2^*) - u_1(a_1^*, a_2^*)$, we have $u_1(a_1^*, a_2^*) < u_1(\underline{a}_1, a_2^*) - c \leq \max\big\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta)u_1(\underline{a}_1, a_2^*)\big\}$. This together with (B.1) implies that

$$V_{k+1} > V_k \geq \max\big\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta)u_1(\underline{a}_1, a_2^*)\big\}. \tag{B.2}$$

I show that

$$V_t > ... > V_{k+1} > V_k \geq \max\big\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta)u_1(\underline{a}_1, a_2^*)\big\}. \tag{B.3}$$

For every s that satisfies $t > s \geq k+1$, the construction of t in Step 2 implies that taking action a_1^* is optimal at h_*^s, which implies that $V_s = (1-\delta)u_1(a_1^*, \beta_s) + \delta V_{s+1}$. The rest of the proof is done by induction on s. First, inequality (B.2) implies that $V_{k+1} > V_k \geq \max\big\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta)u_1(\underline{a}_1, a_2^*)\big\}$. Suppose we know that $V_s > ... > V_{k+1} > \max\big\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta)u_1(\underline{a}_1, a_2^*)\big\}$ for some $s \leq t-1$, let us compare V_s and V_{s+1}. Since $V_s = (1-\delta)u_1(a_1^*, \beta_s) + \delta V_{s+1}$ and $V_s > \max\big\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta)u_1(\underline{a}_1, a_2^*)\big\} \geq u_1(\underline{a}_1, a_2^*) - c > u_1(a_1^*, a_2^*) \geq u_1(a_1^*, \beta_s)$, we know that $V_{s+1} > V_s > u_1(a_1^*, \beta_s)$. This establishes (B.3).

The construction of t in Step 2 also implies that playing a_1^* is *not* optimal for the opportunistic type at h_*^t, which implies that

$$V_t = \max\{u_1(\underline{a}_1, \beta_t) - c, (1-\delta)u_1(\underline{a}_1, \beta_t)\} \leq \max\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta)u_1(\underline{a}_1, a_2^*)\}, \tag{B.4}$$

where the inequality comes from Assumption 1 that $u_1(a_1, a_2)$ is strictly increasing in a_2 and a_2^* is greater than β_t. Inequalities (B.1) and (B.4) together imply that $V_t \leq V_k$, which contradicts (B.3) that $V_t > V_k$. This implies that the opportunistic type plays a_1^* with probability strictly less than 1 at every history.

Step 4: I show that for every $k \leq t$, β_k is strictly increasing in k in the sense of FOSD. Step 3 implies that at every h_*^k with $k \leq t$, *either* it is optimal for player 1 to play $\underline{a}_1$ and then erase it, *or* it is optimal for him to play $\underline{a}_1$ and then not erase it. In the first case, player 1's continuation value is $u_1(\underline{a}_1, \beta_k) - c$. In the second case, player 1's continuation value is $(1-\delta)u_1(\underline{a}_1, \beta_k)$. This leads to the following formula for V_k:

$$V_k = \max\left\{u_1(\underline{a}_1, \beta_k) - c, (1-\delta)u_1(\underline{a}_1, \beta_k)\right\} \text{ for every } k \leq t. \tag{B.5}$$

Since $\beta_k, \beta_{k-1} \in \mathcal{B}$, Lemma 2 implies that β_k and β_{k-1} can be ranked according to FOSD. That is to say, either $\beta_k \succ_2 \beta_{k-1}$ or $\beta_{k-1} \succsim_2 \beta_k$ where $\succ_2$ and $\succsim_2$ are the orders on A_2 defined in Assumption 1.

Suppose by way of contradiction that $\beta_{k-1} \succsim_2 \beta_k$ for some $1 \leq k \leq t$. Since $u_1(a_1, a_2)$ is strictly increasing in a_2, equation (B.5) implies that V_k is a strictly increasing function of β_k. The hypothesis that $\beta_{k-1} \succsim_2 \beta_k$ implies that $V_{k-1} \geq V_k$. Since $k \leq t$, the definition of t implies that the opportunistic type reaches h_*^k with positive probability. Hence, it is optimal for the opportunistic type to play a_1^* at h_{k-1}^*, which implies that $V_{k-1} = (1-\delta)u_1(a_1^*, \beta_{k-1}) + \delta V_k$. This together with $V_{k-1} \geq V_k$ implies that

$$V_{k-1} = (1-\delta)u_1(a_1^*, \beta_{k-1}) + \delta V_k \leq (1-\delta)u_1(a_1^*, \beta_{k-1}) + \delta V_{k-1},$$

or equivalently, $V_{k-1} \leq u_1(a_1^*, \beta_{k-1})$. Since $c < \overline{c} \leq u_1(\underline{a}_1, \beta_{k-1}) - u_1(a_1^*, \beta_{k-1})$, we know that

$$V_{k-1} \leq u_1(a_1^*, \beta_{k-1}) < u_1(\underline{a}_1, \beta_{k-1}) - c. \tag{B.6}$$

This contradicts $V_{k-1} = \max\left\{u_1(\underline{a}_1, \beta_{k-1}) - c, (1-\delta)u_1(\underline{a}_1, \beta_{k-1})\right\}$ and implies that $\beta_k \succ_2 \beta_{k-1}$.

Step 5: I examine the opportunistic type's incentive to erase $\underline{a}_1$. According to (B.5), not erasing $\underline{a}_1$ is preferred to erasing $\underline{a}_1$ if and only if $(1-\delta)u_1(\underline{a}_1, \beta_k) \geq u_1(\underline{a}_1, \beta_k) - c$, or equivalently, $u_1(\underline{a}_1, \beta_k) \leq c/\delta$. Since player 2's action at every on-path history belongs to $\mathcal{B}$, which by Lemma 2 can be completely ranked according to FOSD, we know that $u_1(\underline{a}_1, \beta_k) \leq c/\delta$ is equivalent to β_k being lower than some cutoff in the sense of FOSD. The conclusion in Step 4 that β_k is strictly increasing in k implies the existence of a cutoff $0 \leq t_0 \leq t$ such that at h_*^k, the opportunistic type has a strict incentive to erase $\underline{a}_1$ if $k > t_0$, has a strict incentive not to erase $\underline{a}_1$ if $k < t_0$, and might be indifferent between erasing and not erasing if $k = t_0$.

Step 6: I show that π_k is strictly increasing in k for every $k \leq t$. Suppose by way of contradiction that $\pi_{k-1} \geq \pi_k$ for some $k \leq t$. Let μ_k^* denote the probability that player 2's history is h_*^k conditional on player 1 being the opportunistic type. Applying Bayes Rule to player 1's reputations at h_*^k and h_*^{k-1}, we obtain that

$$\frac{\pi_k}{1-\pi_k} = \frac{\pi(1-\overline{\delta})\overline{\delta}^k}{(1-\pi)\mu_k^*} \leq \frac{\pi_{k-1}}{1-\pi_{k-1}} = \frac{\pi(1-\overline{\delta})\overline{\delta}^{k-1}}{(1-\pi)\mu_{k-1}^*},$$

which implies that

$$\mu_k^*/\mu_{k-1}^* \geq \overline{\delta}. \tag{B.7}$$

Let p_k^* denote the probability that the opportunistic type plays a_1^* at h_*^k. Let q_k^* denote the probability that the opportunistic type plays $\underline{a}_1$ and then erases it at h_*^k. By definition, $q_k^* \leq 1 - p_k^*$. Applying the state distribution lemma on page 2 of Online Appendix A, we obtain that

$$\mu_k^* = \overline{\delta}\Big(\mu_{k-1}^* p_{k-1}^* + \mu_k^* q_k^*\Big) \leq \overline{\delta}\Big(\mu_{k-1}^* p_{k-1}^* + \mu_k^*(1-p_k^*)\Big),$$

or equivalently,

$$\frac{\mu_k^*}{\mu_{k-1}^*} \leq \frac{\overline{\delta} p_{k-1}^*}{1-\overline{\delta}(1-p_k^*)}. \tag{B.8}$$

Inequalities (B.7) and (B.8) together imply that

$$p_{k-1}^* \geq (1-\overline{\delta}) + \overline{\delta} p_k^*. \tag{B.9}$$

Due to the conclusion in Step 3 that $p_k^* < 1$, the above inequality implies that $p_{k-1}^* > p_k^*$. Let x_k denote the probability that player 2's belief at h_*^k assigns to player 1's current-period action being a_1^*. By the law of total probabilities, we have

$$x_k = \pi_k + p_k^*(1-\pi_k). \tag{B.10}$$

Since $p_{k-1}^* > p_k^*$, my hypothesis that $\pi_{k-1} \geq \pi_k$ implies that $x_{k-1} > x_k$. However, Assumption 1 requires that $u_2(a_1, a_2)$ has strictly increasing differences. Recall my conclusion in Step 4 that player 2's action β_k is strictly increasing in k in the sense of FOSD, we know that x_k is weakly increasing in k, that is, $x_{k-1} \leq x_k$. This contradicts my earlier conclusion that $x_{k-1} > x_k$ and implies that $\pi_{k-1} < \pi_k$ for every $k \leq t$.

C Proof of Proposition 3

I start from defining some useful constants. Recall the definition of A_2^* in (A.1) and that a_2^* is the highest action that belongs to A_2^*. Let a_2' denote the *second highest action* in A_2^*. There exists $x^* \in (0,1)$ such that player 2 is indifferent between a_2' and a_2^* when player 1 plays the mixed action $x^* a_1^* + (1-x^*)\underline{a}_1$. Let us consider equation

$$\frac{\pi}{1-\pi} = \frac{1}{(1-\delta)+\delta p} \cdot \frac{x^*-p}{1-x^*}. \tag{C.1}$$

The RHS of (C.1) is strictly decreasing in p, which equals 0 when $p = x^*$. Hence, there exists $p \in (0, x^*]$ that solves (C.1) if and only if the RHS is strictly greater than $\frac{\pi}{1-\pi}$ when $p = 0$. Let

$$\delta_1 \equiv 1 - \frac{1-\pi}{2\pi} \cdot \frac{x^*}{1-x^*}.$$

By definition, when $\delta = \delta_1$, the value of $p \in (0, x^*]$ that solves (C.1) is strictly positive, which I denote by p^*. Since the RHS of (C.1) is strictly increasing in δ, the value of $p \in [0, x^*]$ that solves (C.1) is strictly greater than p^* for every $\delta > \delta_1$.

Assumptions 2 and 4 together imply that $u_1(\underline{a}_1, a_2^*) - c > u_1(a_1^*, a_2^*) > u_1(\underline{a}_1, \underline{a}_2) \equiv 0$. Let $\delta_2 \in (0,1)$ be defined via

$$\frac{1-\delta_2}{\delta_2} c = \frac{1}{2}\Big(u_1(\underline{a}_1, a_2^*) - c\Big).$$

By definition, $u_1(\underline{a}_1, a_2^*) - c > (1-\delta) u_1(\underline{a}_1, a_2^*)$ for every $\delta > \delta_2$. Let $\delta^* \equiv \max\{\delta_1, \delta_2\}$ and let $\lambda^* \equiv 1 - \delta^*$. By definition, if $\delta \leq \delta^*$, then $\lambda^*(1-\delta)^{-1} - 1 \leq 0$, in which case $t \geq \lambda^*(1-\delta)^{-1} - 1$ is trivially satisfied for every $t \in \mathbb{N}$. Note that $\delta_1, \delta_2, \delta^*, \lambda^*, p^*$ depend only on (u_1, u_2, c, π) but not on $\delta, \overline{\delta}, \widehat{\delta}$.

Recall the definitions of t_0 and t in Theorem 1. For every $t_0 < k < t$, the opportunistic type has an incentive to play $\underline{a}_1$ and then erase it at h_*^k, which implies that $V_k = u_1(\underline{a}_1, \beta_k) - c$. He also has an incentive to play a_1^* at h_*^k, which implies that $V_k = (1-\delta) u_1(a_1^*, \beta_k) + \delta V_{k+1}$. Therefore,

$$u_1(\underline{a}_1, \beta_k) - c = (1-\delta) u_1(a_1^*, \beta_k) + \delta(u_1(\underline{a}_1, \beta_{k+1}) - c). \tag{C.2}$$

Since player 2's mixed action at every history belongs to set $\mathcal{B}$ and $\underline{a}_2$ is the lowest element in $\mathcal{B}$, equation (C.2) implies that for every $t_0 < k < t$, we have

$$V_{k+1} - V_k = u_1(\underline{a}_1, \beta_{k+1}) - u_1(\underline{a}_1, \beta_k) \leq (1-\delta)\Big(u_1(\underline{a}_1, \beta_{k+1}) - c - u_1(a_1^*, \beta_k)\Big) \leq (1-\delta)\Big(u_1(\underline{a}_1, a_2^*) - c - u_1(a_1^*, \underline{a}_2)\Big),$$

where the first inequality uses Assumption 4 and the second inequality uses Assumption 1. Let $\Delta \equiv u_1(\underline{a}_1, a_2^*) - c - u_1(a_1^*, \underline{a}_2)$. Assumptions 1 and 4 imply that $\Delta > 0$. I classify equilibria into three classes and derive a lower bound on t for each class. Then I construct a lower bound that apply to all equilibria.

Class 1: Consider any equilibrium where the opportunistic type weakly prefers *not* to erase $\underline{a}_1$ at h_*^0. Step 5 in the proof of Theorem 1 implies that $u_1(\underline{a}_1, \beta_0) \leq c/\delta$, so the opportunistic type's continuation value at h_*^0 is at most $\frac{(1-\delta)c}{\delta}$. From Theorem 1, his continuation value at h_*^{t+1} is $\max\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta) u_1(\underline{a}_1, a_2^*)\}$, which equals $u_1(\underline{a}_1, a_2^*) - c$ when $\delta > \delta_2$. This together with $V_{k+1} - V_k \leq \Delta(1-\delta)$ implies that

$$t+1 \geq \frac{u_1(\underline{a}_1, a_2^*) - c - \frac{(1-\delta)c}{\delta}}{\Delta(1-\delta)} \geq \underbrace{\frac{u_1(\underline{a}_1, a_2^*) - c}{2\Delta}}_{\equiv \lambda_1}(1-\delta)^{-1} = \lambda_1(1-\delta)^{-1} \text{ for every } \delta > \delta^*.$$

Class 2: Consider any equilibrium in which at h_*^0, the opportunistic type has a strict incentive to erase $\underline{a}_1$ and player 2 takes action a_2^* with probability 0. Since a_2' is the second highest action in A_2^* and $u_1(a_1, a_2)$ is strictly increasing in a_2, player 1's continuation value at h_*^0 is at most $u_1(\underline{a}_1, a_2') - c$. Since the opportunistic type's continuation value at h_*^{t+1} is at least $u_1(\underline{a}_1, a_2^*) - c$, we know that

$$t+1 \geq \underbrace{\frac{u_1(\underline{a}_1, a_2^*) - u_1(\underline{a}_1, a_2')}{\Delta}}_{\equiv \lambda_2}(1-\delta)^{-1} = \lambda_2(1-\delta)^{-1}.$$

Class 3: Consider any equilibrium in which at h_0^*, the opportunistic type has a strict incentive to erase $\underline{a}_1$ and player 2 takes action a_2^* with strictly positive probability. Let p_k^* denote the probability that the opportunistic type plays a_1^* at h_*^k. The definition of t implies that $p_t^* = 0$. Inequality (3.9) implies that $p_{k-1}^* - p_k^* \leq 1 - \overline{\delta}$. Player 2's incentive to take action a_2^* at h_*^0 implies that she expects player 1 to take action a_1^* at h_*^0 with probability at least x^*. This together with Bayes rule implies that there exists $x_0 \geq x^*$ such that

$$\frac{\pi}{1-\pi} = \frac{1}{(1-\overline{\delta}) + \overline{\delta} p_0^*} \cdot \frac{x_0 - p_0^*}{1 - x_0}, \tag{C.3}$$

If $\delta > \delta_1$, then $\overline{\delta} > \delta_1$ and the value of p_0^* that solves (C.3) is at least p^*, as defined earlier in this proof. This together with $p_{k-1}^* - p_k^* \leq 1 - \overline{\delta}$ and $p_t^* = 0$ implies that $t \geq p^*(1-\overline{\delta})^{-1} \geq p^*(1-\delta)^{-1}$.

Let $\lambda \equiv \min\{\lambda_1, \lambda_2, p^*\}$, which is independent of $\delta, \overline{\delta}, \widehat{\delta}$ since $\lambda_1, \lambda_2, p^*$ are all independent of $\delta, \overline{\delta}, \widehat{\delta}$. Hence, for every $\delta \leq \delta^*$, we have $\lambda(1-\delta)^{-1} - 1 \leq 0$, so $t \geq \lambda(1-\delta)^{-1} - 1$ for every $t \in \mathbb{N}$. For every $\delta > \delta^*$, as I have shown earlier, $t+1$ is at least $\lambda(1-\delta)^{-1}$ for all three classes of equilibria.

D Proof of Proposition 4

Since $\widehat{\delta} = 1$, I replace $\overline{\delta}$ with δ. Suppose by way of contradiction that for every $\pi > 0$, there exists $\delta \in (0,1)$ and an equilibrium under (π, δ) such that player 1's payoff is *strictly* more than $(1-\delta)c/\delta$. Then, it must be the case that the opportunistic type has a *strict* incentive to erase $\underline{a}_1$ after taking it at h_*^0.

Recall that t denotes the maximal length of good record that the opportunistic type will have. From Step 1 in the proof of Theorem 2, we know that there exists a constant $\phi > 0$ such that $t \leq \phi(1-\delta)^{-1}$ in every equilibrium. Recall that π_0 is player 1's reputation at history h_*^0. Since the honest type reaches history h_*^0 with probability $1-\delta$ and the opportunistic type reaches history h_*^0 with probability μ_0^*, the following equation is obtained from Bayes Rule:

$$\frac{\pi_0}{1-\pi_0} = \frac{\pi}{1-\pi}\frac{1-\delta}{\mu_0^*}. \tag{D.1}$$

The expression for μ_0^* in (3.7) as well as the conclusion in Theorem 1 that $p_0^* < 1$ together imply that $\mu_0^* > 1-\delta$, and therefore, $\pi_0 < \pi$.

Recall that $x_k \in [0,1]$ denotes the probability player 2's posterior assigns to player 1's current-period action being a_1^* after observing history h_*^k. Due to the conclusion in Theorem 1 that $\beta_k \succ_2 \beta_{k-1}$ for every $k \leq t$, we know from Assumption 1 that $x_k \geq x_{k-1}$. Applying equation (3.11) to both k and $k-1$, we obtain that

$$\frac{\mu_k^*}{\mu_{k-1}^*} = \delta\frac{x_{k-1}-p_{k-1}^*}{x_k - p_k^*}\cdot\frac{1-x_k}{1-x_{k-1}} \leq \delta\frac{x_{k-1}-p_{k-1}^*}{x_k-p_k^*}. \tag{D.2}$$

This together with (3.8) implies that

$$\frac{\delta p_{k-1}^*}{1-\delta(1-p_k^*)} = \frac{\mu_k^*}{\mu_{k-1}^*} \leq \delta\frac{x_{k-1}-p_{k-1}^*}{x_k-p_k^*}. \tag{D.3}$$

Taking the inverse on both sides of (D.3) and using the conclusion that $x_k \geq x_{k-1}$, we obtain that

$$\frac{p_k^* + (1-\delta)(1-p_k^*)}{p_{k-1}^*} \geq \frac{x_k - p_k^*}{x_{k-1}-p_{k-1}^*} \geq \frac{x_{k-1}-p_k^*}{x_{k-1}-p_{k-1}^*} = 1 + \frac{p_{k-1}^*-p_k^*}{x_{k-1}-p_{k-1}^*}.$$

This further implies that

$$\frac{(1-\delta)(1-p_k^*)}{p_{k-1}^*} \geq \frac{p_{k-1}^*-p_k^*}{p_{k-1}^*} + \frac{p_{k-1}^*-p_k^*}{x_k-p_{k-1}^*},$$

from which we obtain that

$$p_{k-1}^* - p_k^* \leq (1-\delta)\frac{x_k - p_{k-1}^*}{x_k}. \tag{D.4}$$

Since $u_1(\underline{a}_1, \beta_0) > c/\delta > 0$ and $u_1(\underline{a}_1, \underline{a}_2) \equiv 0$, the probability that β_0 assigns to actions strictly greater than $\underline{a}_2$ is bounded above 0. Since $\underline{a}_2$ is a strict best reply to $\underline{a}_1$, there exists $x^* > 0$ such that player 2 has a strict incentive to play $\underline{a}_2$ against player 1's mixed action $xa_1^* + (1-x)\underline{a}_1$ as long as $x < x^*$. This implies that $x_0 \geq x^*$. Let

$$K \equiv \Big\{k\Big|k \in \{0,1,...,t\} \text{ and } x_k > x_{k-1}\Big\}.$$

Since $u_2(a_1,a_2)$ has strictly increasing differences and β_k is strictly increasing in k (Theorem 1), we know that $x_k \geq x_{k-1}$ for every $k \leq t$. The definition of K then implies that $x_k = x_{k-1}$ for every $k \notin K$. Recall from the proof of Lemma 2 that under the complete order $\succsim_2$, β_k are mixtures between *adjacent* elements in A_2^*, defined in (A.1). Therefore, player 1's expected actions at h_*^k and h_*^{k-1} must be the same (i.e., $x_k = x_{k-1}$) as long as β_k and β_{k-1} are non-trivial mixtures between the same pair of elements in A_2^*. Since β_k is *strictly* increasing in k in FOSD, we know that $|K| \leq |A_2^*| \leq |A_2|$. Recall from (B.10) that $x_t = \pi_t + (1-\pi_t)p_t^*$, we know that (i) p_t^* converges to x_t as π_t converges to 0 and (ii) there exists $p^* \in (0,1)$ such that $p_k^* < p^*$ for every $k \notin K$, and (iii) applying (B.10) to both k and $k-1$, we obtain

$$\pi_k - \pi_{k-1} = x_k - x_{k-1} - p_k^*(1-\pi_k) + p_{k-1}^*(1-\pi_{k-1}) = (p_{k-1}^* - p_k^*)(1-\pi_{k-1}) - p_k^*(\pi_{k-1}-\pi_k). \quad \text{(D.5)}$$

This further implies that

$$\pi_k - \pi_{k-1} = \frac{1-\pi_{k-1}}{1-p_k^*} \cdot (p_{k-1}^* - p_k^*) \leq \frac{p_{k-1}^* - p_k^*}{1-p_k^*} \leq (p_{k-1}^* - p_k^*)(1-p^*)^{-1}. \quad \text{(D.6)}$$

The last inequality in (D.6) relies on $p_{k-1}^* - p_k^* > 0$, which is obtained from (i) $\pi_k > \pi_{k-1}$ (Theorem 1) and (ii) the first part of (D.6) that $\pi_k - \pi_{k-1} = \frac{1-\pi_{k-1}}{1-p_k^*} \cdot (p_{k-1}^* - p_k^*)$ which is derived solely from (D.5).

Since $x_0 \geq x^*$, we know that for every $\varepsilon > 0$, there exists $\overline{\pi} > 0$ such that $p_0^* \geq x^* - \varepsilon$ whenever $\pi < \overline{\pi}$. The definition of t implies that $p_t^* = 0$. The gap between p_0^* and p_t^* together with (D.4) and (D.6) leads to a lower bound on t, which diverges to infinity as $\pi \to 0$. Hence, for every $\phi > 0$, there exists $\overline{\pi} > 0$ such that for every $\pi < \overline{\pi}$, the lower bound on t implied by the speed with which p_k^* decreases in k will be strictly greater than $2\phi(1-\delta)^{-1}$. Under such a π, the existence of an equilibrium where player 1's payoff being strictly greater than $(1-\delta)c/\delta$ contradicts my earlier conclusion that $t < \phi(1-\delta)^{-1}$.

E Proof of Proposition 5

I replace $\overline{\delta}$ with δ. Let $x^{**} \in (0,1)$ denote the smallest $x \in (0,1)$ such that a_2^* best replies to $xa_1^* + (1-x)\underline{a}_1$. Throughout the proof, I also assume that δ is large enough in the sense that it satisfies $u_1(\underline{a}_1, a_2^*) - c > (1-\delta)u_1(\underline{a}_1, a_2^*)$, in which case $V_{t+1} = \max\{u_1(\underline{a}_1, a_2^*) - c, (1-\delta)u_1(\underline{a}_1, a_2^*)\} = u_1(\underline{a}_1, a_2^*) - c$.

I provide a constructive proof to the following claim, which implies Proposition 5: For every $\eta > 0$, there exists $\overline{\pi} > 0$ such that for every $\pi > \overline{\pi}$, there exists an equilibrium in which (i) the opportunistic type has a strict incentive to erase $\underline{a}_1$ at h_*^0, (ii) at every h_*^k with $k < t$, player 2 believes that player 1 will play a_1^* with probability x^{**}, and (iii) player 2's action at h_*^0 assigns probability more than $1-\eta$ to a_2^*.

The law of total probabilities implies that $x_k = \pi_k + (1-\pi_k)p_k^*$. Recall from (D.1) and (3.7) that

$$\frac{\pi_0}{1-\pi_0} = \frac{\pi}{1-\pi} \cdot (1-\delta(1-p_0^*)). \tag{E.1}$$

Equation (E.1) implies that fixing the value of π, π_0 is strictly increasing in p_0^*. It then implies that for every $\varepsilon > 0$, there exists $\overline{\pi} > 0$ such that when $\pi > \overline{\pi}$ and $p_0^* = \varepsilon$, we have $\pi_0 + (1-\pi_0)p_0^* \geq x^{**}$. When $x_0 = x_1 = ... = x_{t-1} = x^{**}$, the conclusion in Theorem 1 that π_k being strictly increasing in k implies that p_k is strictly decreasing in k for every $k < t$. Moreover, we know from (D.2) and (D.3) that when $x_{k-1} = x_k = x^{**}$,

$$\delta \frac{x_{k-1} - p_{k-1}^*}{x_k - p_k^*} = \delta \frac{x^{**} - p_{k-1}^*}{x^{**} - p_k^*} = \frac{\mu_k^*}{\mu_{k-1}^*} = \frac{\delta p_{k-1}^*}{1-\delta(1-p_k^*)},$$

or equivalently,

$$\frac{x^{**} - p_{k-1}^*}{x^{**} - p_k^*} = \frac{p_{k-1}^*}{1-\delta(1-p_k^*)}.$$

After doing some algebra, we can obtain the following expression for $p_{k-1}^* - p_k^*$, which is linear in $1-\delta$:

$$\Big(1 + \delta \frac{x^{**} - p_k^*}{1-\delta(1-p_k^*)}\Big)(p_{k-1}^* - p_k^*) = (1-\delta)\frac{(x^{**} - p_k^*)(1-p_{k-1}^*)}{1-\delta(1-p_k^*)}. \tag{E.2}$$

Since p_k^* is strictly decreasing in k for every $k \leq t$, equation (E.2) implies that there exists $\psi > 0$ such that for every $\varepsilon > 0$ small enough, $p_{k-1}^* - p_k^* \geq \psi(1-\delta)$ for every $k \leq t$ when $p_0^* = \varepsilon$. Hence, when π is large enough such that $\pi_0 + (1-\pi_0)\varepsilon \geq x^{**}$, there exists an equilibrium in which $t \leq \frac{\varepsilon}{\psi(1-\delta)}$. Since the opportunistic type is indifferent between playing a_1^* and playing $\underline{a}_1$ and then erasing it at h_*^k, we have

$$u_1(\underline{a}_1, \beta_{k+1}) - u_1(\underline{a}_1, \beta_k) = (1-\delta)\big(u_1(\underline{a}_1, \beta_{k+1}) - c - u_1(a_1^*, \beta_k)\big) \leq (1-\delta)(\overline{u} - \underline{u} - c).$$

Therefore, player 1's equilibrium payoff $u_1(\underline{a}_1, \beta_0) - c$ is at least $V_{t+1} - (1-\delta)(t+1)(\overline{u} - \underline{u} - c)$ where $\overline{u}$ and $\underline{u}$ are the highest value and the lowest value of player 1's stage-game payoff $u_1(a_1, a_2)$. Since $V_{t+1} = u_1(\underline{a}_1, a_2^*) - c$ and $t \leq \frac{\varepsilon}{\psi(1-\delta)}$, player 1's equilibrium payoff is at least $u_1(\underline{a}_1, a_2^*) - c - \varepsilon \frac{\overline{u} - \underline{u} - c}{\psi}$. The claim in the beginning of this appendix is then established once we set $\eta \equiv \varepsilon \frac{\overline{u} - \underline{u} - c}{\psi}$.

F Proof of Proposition 6

Suppose by way of contradiction that there exists an equilibrium in which the honest type erases a_1^* with strictly positive probability at some history. Since the honest type plays a_1^* in every period, there exists $t \in \mathbb{N}$, which is the smallest integer k such that the honest type erases a_1^* with positive probability at h_*^k.

Hence, it is optimal for the honest type *not* to erase any action until he reaches h_*^t in period t, after which he erases a_1^* in every subsequent period. Under such a strategy, player 2's history will remain at h_*^t after reaching it. I denote this strategy by σ_1^*. Since β_{k+1} and β_t belong to $\mathcal{B}$, Lemma 2 implies that they can be ranked according to FOSD. I consider two cases separately: $\beta_{k+1} \succsim_2 \beta_k$ and $\beta_k \succ_2 \beta_{k+1}$.

First, consider the case in which β_{t+1} weakly FOSDs β_t. Since the honest type chooses a_1^* in every period, his continuation value at h_*^t is $u_1(a_1^*, \beta_t) - c$ if he follows strategy σ_1^* and is $(1-\delta)u_1(a_1^*, \beta_t) + \delta(u_1(a_1^*, \beta_{t+1}) - c)$ if he does not erase a_1^* at h_*^t and erases a_1^* at h_*^{t+1}. The latter is strictly greater than the former when $\beta_{t+1} \succsim_2 \beta_t$ and $c > 0$. This contradicts the hypothesis that erasing a_1^* at h_*^t is optimal.

Next, consider the case in which β_t strictly FOSDs β_{t+1}. Since a_2^* is the highest action in $\mathcal{B}$ and $\beta_t, \beta_{t+1} \in \mathcal{B}$, player 2 *cannot* play a_2^* for sure at h_*^{t+1}. This implies that the opportunistic type must reach h_*^{t+1} with positive probability (which implies that he needs to play a_1^* at h_*^t with positive probability) and moreover, he must play $\underline{a}_1$ with positive probability at h_*^{t+1}. Consider the opportunistic type's continuation value at h_*^t. His continuation value from playing a_1^* at h_*^t and playing $\underline{a}_1$ at h_*^{t+1} is

$$(1-\delta)u_1(a_1^*, \beta_t) + \delta \max\{u_1(\underline{a}_1, \beta_{t+1}) - c, (1-\delta)u_1(\underline{a}_1, \beta_{t+1})\}, \tag{F.1}$$

whereas his continuation value from playing $\underline{a}_1$ at h_*^t is at least

$$\max\{u_1(\underline{a}_1, \beta_t) - c, (1-\delta)u_1(\underline{a}_1, \beta_t)\}. \tag{F.2}$$

Under the hypothesis that $\beta_t \succ_2 \beta_{t+1}$ as well as Assumptions 1 and 4, (F.2) is strictly greater than (F.1). This contradicts my earlier conclusion that the opportunistic type needs to play a_1^* at h_*^t with positive probability.

References

[1] Bhaskar, V and Caroline Thomas (2019) "Community Enforcement of Trust with Bounded Memory," *Review of Economic Studies*, 86(3), 1010-1032.

[2] Board, Simon and Moritz Meyer-ter-Vehn (2013) "Reputation for Quality," *Econometrica*, 81(6), 2381-2462.

[3] Bolton, Gary, Ben Greiner and Axel Ockenfels (2013) "Engineering Trust: Reciprocity in the Production of Reputation Information," *Management Science*, 59(2), 265-285.

[4] Cai, Hongbin, Ginger Zhe Jin, Chong Liu and Li-an Zhou (2014) "Seller Reputation: From Word-of-Mouth to Centralized Feedback," *International Journal of Industrial Organization*, 34, 51-65.

[5] Clark, Daniel, Drew Fudenberg and Alexander Wolitzky (2021) "Record-Keeping and Cooperation in Large Societies," *Review of Economic Studies*, 88, 2179-2209.

[6] Dziuda, Wioletta (2011) "Strategic Argumentation," *Journal of Economic Theory*, 146(4), 1362-1397.

[7] Dye, Ronald (1985) "Disclosure of Nonproprietary Information," *Journal of Accounting Research*, 23(1), 123-145.

[8] Ekmekci, Mehmet (2011) "Sustainable Reputations with Rating Systems," *Journal of Economic Theory*, 146, 479-503.

[9] Ekmekci, Mehmet, Leandro Gorno, Lucas Maestri, Jian Sun and Dong Wei (2022) "Learning from Manipulable Signals," *American Economic Review*, 112(12), 3995-4040.

[10] Ely, Jeffrey and Juuso Välimäki (2003) "Bad Reputation," *Quarterly Journal of Economics*, 118(3), 785-814.

[11] Ely, Jeffrey, Drew Fudenberg and David Levine (2008) "When is Reputation Bad?" *Games and Economic Behavior*, 63(2), 498-526.

[12] Fudenberg, Drew, David Kreps and Eric Maskin (1990) "Repeated Games with Long-run and Short-run Players," *Review of Economic Studies*, 57(4), 555-573.

[13] Fudenberg, Drew and Eric Maskin (1986) "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54(3), 533-554.

[14] Fudenberg, Drew and David Levine (1989) "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, 57(4), 759-778.

[15] Grossman, Sanford (1981) "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law & Economics*, 24(3), 461-483.

[16] Hauser, Daniel (2023) "Censorship and Reputation," *American Economic Journal: Microeconomics*, 15(1), 497-528.

[17] Heller, Yuval and Erik Mohlin (2018) "Observations on Cooperation," *Review of Economic Studies*, 88, 1892-1935.

[18] Hu, Ju (2020) "On the Existence of the Ex Post Symmetric Random Entry Model," *Journal of Mathematical Economics*, 90, 42-47.

[19] Kambe, Shinsuke (1999) "Bargaining with Imperfect Commitment," *Games and Economic Behavior*, 28(2), 217-237.

[20] Kaya, Ayca (2009) "Repeated Signaling Games," *Games and Economic Behavior*, 66(2), 841-854.

[21] Kovbasyuk, Sergey and Giancarlo Spagnolo (2023) "Memory and Markets," *Review of Economic Studies*, forthcoming.

[22] Levine, David (2021) "The Reputation Trap," *Econometrica*, 89(6), 2659-2678.

[23] Li, Yingkai and Harry Pei (2021) "Equilibrium Behaviors in Repeated Games," *Journal of Economic Theory*, 193, 105222.

[24] Lipnowski, Elliot, Doron Ravid, and Denis Shishkin (2022) "Persuasion via Weak Institutions," *Journal of Political Economy*, 130(10), 2705-2730.

[25] Livingston, Jeffrey (2005) "How Valuable Is a Good Reputation? A Sample Selection Model of Internet Auctions," *Review of Economics and Statistics*, 87(3),453-465.

[26] Liu, Qingmin (2011) "Information Acquisition and Reputation Dynamics," *Review of Economic Studies*, 78(4), 1400-1425.

[27] Liu, Qingmin and Andrzej Skrzypacz (2014) "Limited Records and Reputation Bubbles," *Journal of Economic Theory* 151, 2-29.

[28] Mailath, George and Larry Samuelson (2001) "Who Wants a Good Reputation?" *Review of Economic Studies*, 68(2), 415-441.

[29] Milgrom, Paul (1981) "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12(2), 380-391.

[30] Nosko, Chris and Steven Tadelis (2015) "The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment," NBER Working Paper, 20830.

[31] Pei, Harry (2024a) "Reputation Effects under Short Memories," *Journal of Political Economy*, forthcoming.

[32] Pei, Harry (2024b) "Community Enforcement with Endogenous Records," arXiv:2401.00839.

[33] Phelan, Christopher (2006) "Public Trust and Governmental Betrayals," *Journal of Economic Theory*, 130(1), 27-43.

[34] Saeedi, Maryam and Ali Shourideh (2023) "Optimal Rating Design under Moral Hazard," arXiv:2008.09529.

[35] Smirnov, Aleksei, and Egor Starkov (2022) "Bad News Turned Good: Reversal Under Censorship," *American Economic Journal: Microeconomics*, 14(2), 506-560.

[36] Sugaya, Takuo and Alexander Wolitzky (2020) "Do a Few Bad Apples Spoil the Barrel?: An Anti-Folk Theorem for Anonymous Repeated Games with Incomplete Information," *American Economic Review*, 110, 3817-3835.

[37] Sun, Yiman (2024) "A Dynamic Model of Censorship," *Theoretical Economics*, 19(1), 29-60.

[38] Tadelis, Steven (2016) "Reputation and Feedback Systems in Online Platform Markets," *Annual Review of Economics*, 8, 321-340.

[39] Takahashi, Satoru (2010) "Community Enforcement When Players Observe Partners' Past Play," *Journal of Economic Theory*, 145, 42-62.

[40] Vong, Allen (2022) "Certification for Consistent Quality Provision," Working Paper.

[41] Wolitzky, Alexander (2012) "Reputational Bargaining with Minimal Knowledge of Rationality," *Econometrica*, 80(5), 2047-2087.

[42] Wong, Yu Fu (2023) "Dynamic Monitoring Design," Working Paper.