

Online Appendix: Reputation Effects with Endogenous Records

Harry PEI

March 5, 2026

A Proof of Lemma 1

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'' are all best replies to $\lambda^* a_1^* + (1 - \lambda^*)\underline{a}_1$. Then Assumption 3 implies that there exist $\lambda, \lambda', \lambda'' \in [0, 1] \setminus \{\lambda^*\}$ 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 , that is,

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

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 $0 \equiv \lambda_0 < \lambda_1 < \dots < \lambda_n \equiv 1$ such that once we restrict attention to $\lambda \in (0, 1)$, for every $a_2 \in A_2^*$, there exists $j \in \{1, 2, \dots, n\}$ 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 \underline{a}_2 is a strict best

reply to \underline{a}_1 . This implies that \underline{a}_2 also best replies to $\lambda a_1^* + (1 - \lambda)\underline{a}_1$ for all λ close to 0. Let λ_1 denote 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. 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 $\lambda_1, \lambda_2, \dots$, until λ reaches 1.

The above construction implies that for every $j \in \{1, 2, \dots, n - 1\}$, 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 $u_1(a_1, \beta) = v$.

B High Cost of Erasing Actions

I maintain all the assumptions in the baseline model except for Assumption 4. Instead, I assume that the opportunistic type's cost c of erasing each action is large enough such that

$$c > \tilde{c} \equiv \max_{a_2 \in A_2} \left\{ u_1(\underline{a}_1, a_2) - u_1(a_1^*, a_2) \right\}. \quad (\text{B.1})$$

In the following product choice game, we have $\tilde{c} = g$ and condition (B.1) translates into $c > g$.

seller \ consumer	Large Quantity	Small Quantity
Good Products	1, 1	$-g, x$
Bad Products	$1 + g, -x$	0, 0

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

I show that as long as player 1's *effective discount factor* is greater than some cutoff, his payoff in every equilibrium is bounded below by something that is close to his commitment payoff $u_1(a_1^*, a_2^*)$.

Similar to the payoff lower bound in Fudenberg and Levine (1989), this payoff lower bound applies to all values of $\pi \in (0, 1)$ and furthermore, it depends on $\hat{\delta}$ and $\bar{\delta}$ only through the product of the two δ .

Theorem. *Suppose (u_1, u_2) satisfies Assumptions 1, 2, and 3. For every $c > \tilde{c}$, $\pi \in (0, 1)$, and $\varepsilon > 0$,*

there exists $\delta^* \in (0, 1)$ such that in every equilibrium where $\delta > \delta^*$, the opportunistic type's discounted average payoff is at least $u_1(a_1^*, a_2^*) - \varepsilon$ and he never erases any action on the equilibrium path.

This result is not straightforward since even when playing \underline{a}_1 and then erasing it leads to a strictly lower stage-game payoff compared to playing a_1^* , the opportunistic type may have an incentive to do so if it increases his continuation value. My proof below rules out the above concern.

Proof. Fix any equilibrium. Player 1 will not take any action a_1 that is strictly greater than \underline{a}_1 and then erase it. This is because doing so is strictly dominated by taking action \underline{a}_1 and then erasing it since the two strategies lead to the same history for player 2's in the next period (and hence, the same continuation value for player 1) but the latter results in a strictly higher stage-game payoff. Let h_*^k denote the history where player 1 has k unerased actions, all of which are a_1^* . Let β_k denote player 2's action at h_*^k . Let V_k denote the opportunistic type's continuation value at history h_*^k . The rest of the proof proceeds in three steps.

Step 1: Let μ_k denote the probability that the history is h_*^k conditional on player 1 being the opportunistic type. Let q_k denote the probability that the opportunistic type takes action \underline{a}_1 at h_*^k and then erase it. Let p_k denote the probability that the opportunistic type takes action other than a_1^* at h_*^k and then not erase it. By definition, the opportunistic type takes action a_1^* at h_*^k with probability $1 - p_k - q_k$. Since player 2's prior belief assigns probability $(1 - \bar{\delta})\bar{\delta}^k$ to player 1's age in the game being k , the state distribution lemma presented in Appendix A of the main text implies that

$$\mu_0 = (1 - \bar{\delta}) + \bar{\delta}q_0\mu_0$$

$$\text{and } \mu_k = \bar{\delta} \left\{ (1 - p_{k-1} - q_{k-1})\mu_{k-1} + q_k\mu_k \right\} \text{ for every } k \geq 1,$$

which is equivalent to

$$\mu_0 = \frac{1 - \bar{\delta}}{1 - \bar{\delta}q_0} \tag{B.2}$$

$$\text{and } \frac{\mu_k}{\mu_{k-1}} = \frac{\bar{\delta}(1 - p_{k-1} - q_{k-1})}{1 - \bar{\delta}q_k} \text{ for every } k \geq 1. \tag{B.3}$$

Equations (B.2) and (B.3) together imply that for every $k \geq 1$, we have:

$$\frac{\mu_k}{(1 - \bar{\delta})\bar{\delta}^k} = \frac{1}{1 - \bar{\delta}q_k} \cdot \prod_{i=0}^{k-1} \frac{1 - p_i - q_i}{1 - \bar{\delta}q_i}. \tag{B.4}$$

I show that for every $x \in (0, 1)$ and $y > 0$, there exist at most a finite number of k such that $\frac{\mu_k}{(1-\delta)\delta^k} \geq y$ and $p_k + q_k \geq 1 - x$. This is because for every $k \in \mathbb{N}$, we have

$$\frac{1 - p_k - q_k}{1 - \bar{\delta}q_k} \leq 1,$$

and for every k where $p_k + q_k \geq 1 - x$, we have

$$\frac{1 - p_k - q_k}{1 - \bar{\delta}q_k} \leq \frac{x}{1 - \bar{\delta} + \bar{\delta}x} < 1. \quad (\text{B.5})$$

Moreover,

$$\frac{1}{1 - \bar{\delta}q_k} \leq \frac{1}{1 - \bar{\delta}}. \quad (\text{B.6})$$

Suppose by way of contradiction that there are infinitely many k such that $\frac{\mu_k}{(1-\delta)\delta^k} \geq y$ and $p_k + q_k \geq 1 - x$. Combining (B.4), (B.5) and (B.6), we know that there exists $t \in \mathbb{N}$ such that as long as there are t such periods before history h_*^s ,

$$\frac{\mu_s}{(1-\delta)\delta^s} \leq \left(\frac{x}{1 - \bar{\delta} + \bar{\delta}x} \right)^t \cdot \frac{1}{1 - \bar{\delta}} < y. \quad (\text{B.7})$$

This contradiction implies that there exist at most a finite number of such k .

Step 2: At every history h_*^k where with positive probability, the opportunistic type takes action \underline{a}_1 and then erases it, his incentive constraint at that history implies that

$$u_1(\underline{a}_1, \beta_k) - c \geq (1 - \delta)u_1(a_1^*, \beta_k) + \delta V_{k+1}. \quad (\text{B.8})$$

Since $c > \tilde{c}$, we know that $u_1(\underline{a}_1, \beta_k) - c < u_1(a_1^*, \beta_k)$. Therefore, inequality (B.8) implies that

$$V_{k+1} \leq \frac{u_1(\underline{a}_1, \beta_k) - c - (1 - \delta)u_1(a_1^*, \beta_k)}{\delta} < \frac{u_1(a_1^*, \beta_k) - (1 - \delta)u_1(a_1^*, \beta_k)}{\delta} = u_1(a_1^*, \beta_k).$$

This implies that there exists $s > k$ such that player 2 takes action a_2^* with probability strictly less than 1 at history h_*^s . This is because otherwise, the opportunistic type can secure payoff $u_1(a_1^*, a_2^*)$ at h_*^{k+1} by playing a_1^* in every subsequent period. This secured payoff is weakly greater than $u_1(a_1^*, \beta_k)$. This will contradict our earlier conclusion that $V_{k+1} < u_1(a_1^*, \beta_k)$.

Since player 2 has an incentive to take actions strictly lower than a_2^* at h_*^k , there exists $x \in (0, 1)$ such that (i) player 1's reputation at h_*^k , measured by the probability that he is the honest type, is no more than x , and (ii) the opportunistic type takes action a_1^* at h_*^k with probability less than x , or equivalently,

$p_k + q_k \geq 1 - x$. Since the honest type induces history h_*^k with probability $(1 - \bar{\delta})\bar{\delta}^k$, player 1's reputation at h_*^k , denoted by π_k , satisfies:

$$\frac{1 - \pi_k}{\pi_k} = \frac{1 - \pi}{\pi} \cdot \frac{\mu_k}{(1 - \bar{\delta})\bar{\delta}^k}. \quad (\text{B.9})$$

Combining Step 1 and Step 2, we know that in every equilibrium, there exist at most a finite number of k such that the opportunistic type erases action \underline{a}_1 with positive probability at h_*^k .

Step 3: If there exists no $k \in \mathbb{N}$ such that the opportunistic type erases action \underline{a}_1 with positive probability at h_*^k , then the desired conclusion follows from the argument in Fudenberg and Levine (1989).

Suppose by way of contradiction that there exists at least one $k \in \mathbb{N}$ such that the opportunistic type erases action \underline{a}_1 with strictly positive probability at h_*^k . Let t denote the largest of such integer k . First, we show that $\pi_{t+1} \geq \pi$. By definition, $q_{t+1} = 0$, and according to (B.4),

$$\frac{\mu_{t+1}}{(1 - \bar{\delta})\bar{\delta}^{t+1}} = \frac{1}{1 - \bar{\delta}q_{t+1}} \prod_{i=0}^t \frac{1 - p_i - q_i}{1 - \bar{\delta}q_i} = \prod_{i=0}^t \frac{1 - p_i - q_i}{1 - \bar{\delta}q_i} \leq 1.$$

Therefore, $\pi_{t+1} \geq \pi$ follows from (B.9). The result in Fudenberg and Levine (1989) implies that for every $\pi \in (0, 1)$ and $\varepsilon > 0$, there exists $\delta^* \in (0, 1)$ such that in every equilibrium where $\delta > \delta^*$, the opportunistic type's continuation value at h_*^{t+1} is at least $u_1(a_1^*, a_2^*) - \varepsilon$. Let us consider the opportunistic type's incentive at h_*^t . By definition, at history h_*^t , he prefers to take action \underline{a}_1 and then erases it, to taking action a_1^* . This leads to the following incentive constraint:

$$u_1(\underline{a}_1, \beta_t) - c \geq (1 - \delta)u_1(a_1^*, \beta_t) + \delta V_{t+1} \geq (1 - \delta)u_1(a_1^*, \beta_t) + \delta(u_1(a_1^*, a_2^*) - \varepsilon). \quad (\text{B.10})$$

However, since $c > \tilde{c}$, inequality (B.10) cannot be true for any ε that satisfies $u_1(a_1^*, a_2^*) - \varepsilon > u_1(\underline{a}_1, \beta_t) - c$. This contradiction implies that when δ is close enough to 1, there exists no equilibrium in which the opportunistic type erases action with positive probability at any history where he has a strictly positive reputation. It also implies that for every $\varepsilon > 0$, there exists $\delta^* \in (0, 1)$, such that in every equilibrium, the opportunistic type's payoff is at least $u_1(a_1^*, a_2^*) - \varepsilon$ as long as $\delta > \delta^*$. \square

C Type-Dependent Survival Probability

My baseline model assumes that both types survive with the same probability $\bar{\delta}$. Suppose now that the honest type's survival probability, denoted by $\xi \in (0, 1)$, can be different from that of the opportunistic type's,

denoted by $\bar{\delta} \in (0, 1)$. My main result, Theorem 1, extends to this general setting: Suppose (u_1, u_2, c) satisfies Assumptions 1, 2., 3 and 4. For every $\pi \in (0, 1)$, $\hat{\delta} \in (0, 1)$, and $\varepsilon > 0$, there exists $\delta^* \in (0, 1)$ such that for every $\bar{\delta} \in (\delta^*, 1)$ and $\xi \in (0, 1)$, in every equilibrium under $(u_1, u_2, c, \pi, \hat{\delta}, \bar{\delta}, \xi)$:

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

To see why this is true, let us revisit the proof of Theorem 1 in Section 3.3 of the main text. The derivation of the upper bound on t , which is inequality (3.2) in the main text:

$$t \leq \bar{T} \equiv (1 - \delta)^{-1} \frac{u_1(a_1, a_2^*)}{\bar{c} - c}. \quad (\text{C.1})$$

only uses the opportunistic type's incentive constraint. Therefore, this upper bound remains proportional to $(1 - \delta)^{-1}$ in this general setting. As for the derivation of the lower bound on t , equations (3.3) and (3.4) in the main text remain unchanged:

$$\mu_0^* = (1 - \bar{\delta}) + \bar{\delta}\mu_0^*q_0^* \quad \text{and} \quad \mu_k^* = \bar{\delta}(\mu_{k-1}^*p_{k-1}^* + \mu_k^*q_k^*) \quad \text{for every } k \in \{1, \dots, t\}. \quad (\text{C.2})$$

$$\mu_0^* = \frac{1 - \bar{\delta}}{1 - \bar{\delta}(1 - p_0^*)} \quad \text{and} \quad \frac{\mu_k^*}{\mu_{k-1}^*} = \frac{\bar{\delta}p_{k-1}^*}{1 - \bar{\delta}(1 - p_k^*)} \quad \text{for every } k \in \{1, \dots, t\}. \quad (\text{C.3})$$

This is because these equations only examine the probabilities conditional on player 1 being the opportunistic type, which are independent of the honest type's survival probability ξ . Equation (3.5) in the main text, which is

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

should be modified as:

$$\frac{\pi_k}{1 - \pi_k} = \frac{\pi}{1 - \pi} \cdot \frac{(1 - \xi)\xi^k}{\mu_k^*} \quad \Rightarrow \quad \mu_k^*/\mu_{k-1}^* = \xi \frac{1 - \pi_k}{\pi_k} \cdot \frac{\pi_{k-1}}{1 - \pi_{k-1}} \leq \xi. \quad (\text{C.4})$$

Inequality (C.4) together with (C.3) implies that

$$p_{k-1}^* \leq \frac{\xi}{\bar{\delta}} \left((1 - \bar{\delta}) + \bar{\delta}p_k^* \right) \leq \frac{1}{\bar{\delta}} \left((1 - \bar{\delta}) + \bar{\delta}p_k^* \right). \quad (\text{C.5})$$

Inequality (C.5) implies that $p_{k-1}^* - p_k^* \leq \frac{1 - \bar{\delta}}{\bar{\delta}}$, so the lower bound on t remains proportional to $(1 - \bar{\delta})^{-1}$. The argument which shows that p_0^* being uniformly bounded above 0 remains unchanged. The argument

which shows that the short-run players' welfare is close to that in the repeated complete information game is based on the probabilities conditional on player 1 being the opportunistic type, which also extends to this more general setting.

D Stochastic Arrivals and Stochastic Reviews

I discuss two extensions: one in which the consumers arrive with probability less than one in each period and one in which the consumers post reviews with probability less than 1 after interacting with the seller. In these extensions, the seller's age in the game may not coincide with the length of his record, even when he is honest and has not erased any action. I argue that my main result, Theorem 1, continues to hold in these environments and the qualitative features of the equilibria are similar to those in the baseline model.

Stochastic Arrivals: I start from an extension in which the short-run players arrive *stochastically* over time. Suppose in each period, a short-run player arrives with some exogenous probability $p \in (0, 1)$. If a short-run player arrives, then players play the stage game, and by the end of that period, the long-run player decides whether to erase his action. If no short-run player arrives in a given period, then the long-run player's record remains the same regardless of his behavior and his stage-game payoff is normalized to 0.

In this setting, the opportunistic type maximizes $p \sum_{k=0}^{+\infty} (1 - \delta) \delta^k (u_1(a_{1,k}, a_{2,k}) - c_k)$. The short-run players' prior belief assigns probability $(1 - \bar{\delta}) \bar{\delta}^k$ to the long-run player's *age in the game* being k and assigns probability $(1 - \tilde{\delta}) \tilde{\delta}^k$ to the honest type having interacted with k short-run players, where

$$\tilde{\delta} \equiv 1 - \frac{1 - \bar{\delta}}{1 - \bar{\delta}(1 - p)}. \quad (\text{D.1})$$

Using the same method as in the proof of Theorem 1, one can show that for every $c < \bar{c}$, $p \in (0, 1)$, $\pi \in (0, 1)$, $\hat{\delta} \in (0, 1)$, and $\varepsilon > 0$, there exists $\delta^* \in (0, 1)$ such that when $\bar{\delta} > \delta^*$, the opportunistic type's payoff is no more than

$$u(p) \equiv \frac{(1 - \delta)(1 - \delta + \delta p)}{\delta} c \quad (\text{D.2})$$

in every equilibrium. This payoff converges to his minmax value 0 as $\delta \rightarrow 1$. Under an additional assumption that $u_2(a_1, a_2)$ is weakly increasing in a_1 , the short-run players' welfare, measured by the expected sum of their payoffs, is ε -close to $p\pi u_2(a_1^*, a_2^*) + p(1 - \pi)u_2(\underline{a}_1, \underline{a}_2)$ in every equilibrium.

I explain how to extend the proof of Theorem 1 to the case with stochastic arrivals. I omit the details in order to avoid repetition. Suppose by way of contradiction that there exists an equilibrium in which the

opportunistic type of player 1 strictly prefers to erase action \underline{a}_1 after taking it at the null history h_*^0 . Since his continuation value at h_*^k is strictly increasing in k for every $k \leq t$, he also has a strict incentive to erase \underline{a}_1 at every h_*^k . On the one hand, player 1 has an incentive to take action a_1^* at history h_*^k only if

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

A necessary condition for the above inequality is that player 1's continuation value increases by something linear in $1 - \delta$ when his record length increases from k to $k + 1$. This leads to an upper bound on the number of periods with which the opportunistic type can take action a_1^* , which is an affine function of $(1 - \delta)^{-1}$.

On the other hand, the argument in the proof of Theorem 1 implies that the rate with which player 1's reputation increases is bounded above by something proportional to $1 - \tilde{\delta}$. This implies that the number of periods with which the opportunistic type needs to take action a_1^* is at least a linear function of $(1 - \tilde{\delta})^{-1}$. As $\bar{\delta} \rightarrow 1$, expression (D.1) implies that $\tilde{\delta}$ also goes to 1, and the lower bound on t will exceed the upper bound. Therefore, in every equilibrium, the opportunistic type has an incentive not to erase \underline{a}_1 at h_*^0 .

I bound player 1's payoff in equilibria where he has an incentive *not to erase* \underline{a}_1 at history h_*^0 . Suppose a short-run player arrives at history h_*^0 and that the opportunistic type took action \underline{a}_1 at that history, the opportunistic type prefers not to erase \underline{a}_1 if

$$(1 - \delta)u_1(\underline{a}_1, \beta_0) \geq (1 - \delta)(u_1(\underline{a}_1, \beta_0) - c) + \delta p(u_1(\underline{a}_1, \beta_0) - c),$$

or equivalently,

$$u_1(\underline{a}_1, \beta_0) \leq \frac{1 - \delta + \delta p}{\delta p} c = u(p). \quad (\text{D.3})$$

In order to bound the short-run players' payoffs, I establish a generalized version of Lemma 3 in the main text, which says that the probability of event \mathcal{E}_k is no more than $1 - \tilde{\delta}$.

The short-run players' equilibrium welfare is no less than their payoff when they take action a_2^* if and only if the length of player 1's good record exceeds t , and takes action \underline{a}_2 otherwise. This lower bound converges to $p\pi u_2(a_1^*, a_2^*) + p(1 - \pi)u_2(\underline{a}_1, \underline{a}_2)$ as $\tilde{\delta} \rightarrow 1$, since the *expected* number of periods for the honest type to obtain a record length t is a linear function of $(1 - \delta)^{-1}$, which is lower than the decay rate of their belief $\tilde{\delta}$. Their equilibrium welfare is no more than their payoff when they can observe the realized pure action of player 1. This upper bound converges to $p\pi u_2(a_1^*, a_2^*) + p(1 - \pi)u_2(\underline{a}_1, \underline{a}_2)$ as $\tilde{\delta} \rightarrow 1$, since the average probability with which the opportunistic type takes action a_1^* vanishes as $\tilde{\delta} \rightarrow 1$. The lower and the upper bounds coincide, which pins down the short-run players' equilibrium welfare.

Stochastic Online Reviews: I discuss an extension in which the short-run players post reviews with probability less than 1. Suppose after interacting with the long-run player, the short-run player does not leave any review with probability $p \in [0, 1)$, in which case the long-run player's record does not change regardless of his action. Conditional on the short-run player posts a review, the long-run player decides whether to erase it at cost c . My baseline model assumes that $p = 0$.

I explain how to extend Theorem 1 to any arbitrary $p \in [0, 1)$. The opportunistic type's continuation value equals his minmax value 0 after separating from the honest type. After taking action \underline{a}_1 at history h , he prefers to erase it if and only if his continuation value at history h , denoted by $V(h)$, satisfies $V(h) \geq \frac{(1-\delta)c}{\delta}$. Hence, in any equilibrium where the opportunistic type has an incentive not to erase \underline{a}_1 after taking it in period 0, player 1's equilibrium payoff is no more than $\frac{(1-\delta)c}{\delta}$.

Suppose by way of contradiction that for every $\delta^* \in (0, 1)$, there exist $\bar{\delta} > \delta^*$ and an equilibrium under which the opportunistic type strictly prefers to erase \underline{a}_1 after taking it in period 0. Let $t \in \mathbb{N}$ be such that the opportunistic type plays a_1^* with positive probability at h_*^k if and only if $k < t - 1$. Let V_k denote player 1's continuation value at h_*^k and let β_k denote player 2's action at h_*^k . Since at every h_*^k with $k < t - 1$, the opportunistic type is indifferent between playing a_1^* and playing \underline{a}_1 and then erasing it, we have:

$$(1 - \delta)u_1(a_1^*, \beta_k) + \delta \left\{ pV_k + (1 - p)V_{k+1} \right\} = (1 - \delta) \left\{ u_1(\underline{a}_1, \beta_k) - (1 - p)c \right\} + \delta V_k,$$

which implies that

$$V_{k+1} - V_k = \frac{1 - \delta}{\delta(1 - p)} \left\{ u_1(\underline{a}_1, \beta_k) - u_1(a_1^*, \beta_k) - c(1 - p) \right\}. \quad (\text{D.4})$$

Since $u_1(\underline{a}_1, \beta_k) - u_1(a_1^*, \beta_k) - c(1 - p) > 0$, there exists a constant $\lambda > 0$ such that $t \leq \frac{\lambda}{1 - \bar{\delta}}$.

Recall the definitions of μ_k^* and p_k^* in the proof of Theorem 1. When the short-run players do not leave reviews with probability p , we have

$$\mu_0^* = (1 - \bar{\delta}) + \bar{\delta} \mu_0^* \left\{ (1 - p_0^*) + p_0^* p \right\}$$

and

$$\mu_k^* = \bar{\delta} \left\{ \mu_{k-1}^* p_{k-1}^* (1 - p) + \mu_k^* (1 - p_k^* + p_k^* p) \right\}$$

or equivalently,

$$\mu_0^* = \frac{1 - \bar{\delta}}{1 - \bar{\delta}(1 - p_0^* + p_0^* p)} \quad (\text{D.5})$$

and

$$\frac{\mu_k^*}{\mu_{k-1}^*} = \frac{\bar{\delta} p_{k-1}^* (1-p)}{1 - \bar{\delta} (1 - p_k^* + p_k^* p)}. \quad (\text{D.6})$$

Let x_k denote player 2's belief about the probability of a_1^* at history h_*^k . Since the probability with which player 2 does not leave reviews is the same regardless of player 1's type and action, as in the baseline model, we have the following two equations:

$$\frac{\pi}{1-\pi} (1-\bar{\delta}) \bar{\delta}^k = \mu_k^* \left\{ \frac{x_k}{1-x_k} (1-p_k^*) - p_k^* \right\} = \mu_k^* \frac{x_k - p_k^*}{1-x_k} \quad (\text{D.7})$$

and

$$\frac{\mu_k^*}{\mu_{k-1}^*} = \bar{\delta} \frac{x_{k-1} - p_{k-1}^*}{x_k - p_k^*} \cdot \frac{1-x_k}{1-x_{k-1}} \leq \bar{\delta} \frac{x_k - p_{k-1}^*}{x_k - p_k^*}. \quad (\text{D.8})$$

Plugging $k = 0$ into (D.7) and applying equation (D.5), we know that p_0^* is bounded above 0 as $\bar{\delta} \rightarrow 1$. Equations (D.6) and (D.8) together imply that

$$\frac{x_k - p_{k-1}^*}{x_k - p_k^*} \geq \frac{p_{k-1}^* (1-p)}{1 - \bar{\delta} (1 - p_k^* + p_k^* p)}, \quad (\text{D.9})$$

As in the baseline model, the above inequality leads to an upper bound on t , which is still proportional to $(1-\bar{\delta})^{-1}$.

When $\bar{\delta} \rightarrow 1$, the lower bound on t exceeds the upper bound on t driven by the opportunistic type's incentives to take action a_1^* , which rules out equilibria in which the opportunistic type has a strict incentive to erase \underline{a}_1 in period 0 and implies that player 1's payoff in every equilibrium is no more than $\frac{(1-\delta)c}{\delta}$.

The short-run players' welfare is arbitrarily close to $\pi u_2(a_1^*, a_2^*) + (1-\pi)u_2(\underline{a}_1, \underline{a}_2)$ in every equilibrium. This is because conditional on player 1 being the opportunistic type, the probability of event \mathcal{E}^k , defined before the statement of Lemma 3 in the main text, is $\mu_k^* p_k^*$ and satisfies

$$\mu_k^* p_k^* \leq (1-\bar{\delta}) \bar{\delta}^k \frac{1}{1-\bar{\delta} p}. \quad (\text{D.10})$$

This can be shown using the same induction argument as in the proof of Lemma 3 of the main text. Since t is bounded above by a linear function of $(1-\delta)^{-1}$, the average probability with which the opportunistic type takes action a_1^* is close to 0 when $\bar{\delta}$ is close to 1 in every equilibrium.

E Multiple Honest and Opportunistic Types

This appendix extends my results to settings with multiple honest and opportunistic types. Formally, the set of player 1's types is denoted by $\Omega \equiv \Theta \cup \tilde{A}_1$, where each $\theta \in \Theta$ stands for an opportunistic type who is characterized by a stage-game payoff function $u_1(\theta, a_1, a_2)$ and a cost of erasing actions $c(\theta)$, and each $a_1 \in \tilde{A}_1 \subset A_1$ stands for an honest type who takes action a_1 in every period and never erases his action.

I assume that the type distribution $\pi \in \Delta(\Omega)$ has full support, player 2's payoff $u_2(a_1, a_2)$ does not depend on player 1's type, and that for every $\theta \in \Theta$ and for every action a_1^* taken by some honest type, $u_1(\theta, a_1, a_2)$ and $u_2(a_1, a_2)$ satisfy Assumptions 1, 2 and 3. I also adopt the normalization that $u_1(\theta, \underline{a}_1, \underline{a}_2) = 0$ for every $\theta \in \Theta$. As in the baseline model, I restrict attention to public equilibria which are Nash equilibria where each type of player 1's strategy depends only on player 2's history.

I start from the benchmark without any honest type. Recall that a_1' is the lowest action in A_1 under which \underline{a}_2 is *not* a best reply. Let

$$\bar{c}(\theta) \equiv \min_{\beta \in \Delta(A_2)} \left\{ u_1(\theta, \underline{a}_1, \beta) - u_1(\theta, a_1', \beta) \right\}. \quad (\text{E.1})$$

Similar to the case with only one opportunistic type, when the cost of erasing actions $c(\theta)$ is lower than $\bar{c}(\theta)$ for every $\theta \in \Theta$, I show that player 1 plays \underline{a}_1 and player 2 plays \underline{a}_2 at every on-path history of every Nash equilibrium, regardless of $\hat{\delta}$ and $\bar{\delta}$. For each strategy profile (σ_1, σ_2) , let $\mathcal{H}(\sigma_1, \sigma_2 | \theta)$ denote the set of player 2's histories that occur with positive probability under $(\theta, \sigma_1, \sigma_2)$.

Proposition. *Suppose (u_1, u_2) satisfies Assumptions 1, 2, and 3 for every $\theta \in \Theta$, $c(\theta) < \bar{c}(\theta)$ for every $\theta \in \Theta$, and the probability of honest types is 0. If (σ_1, σ_2) is a Nash equilibrium, then for every $\theta \in \Theta$, type θ of player 1 plays \underline{a}_1 and player 2 plays \underline{a}_2 at every history that belongs to $\mathcal{H}(\sigma_1, \sigma_2 | \theta)$.*

This proposition implies that the presence of incomplete information by itself *cannot* alleviate the inefficiencies caused by the long-run player's ability to erase actions at a low cost. Even when there are multiple opportunistic types with potentially different stage-game payoffs and different costs of erasing actions, each opportunistic type will take action \underline{a}_1 at every on-path history and there is no cooperation in any equilibrium.

Proof. I use h to denote a generic history of player 2's. Let $\beta(h) \in \Delta(A_2)$ denote player 2's action at history h . Let $V_\theta(h)$ denote type θ 's continuation value at h .

Suppose by way of contradiction that there exists an equilibrium (σ_1, σ_2) in which some opportunistic

type $\theta_1 \in \Theta$ of player 1 plays $a_1'' \neq \underline{a}_1$ with positive probability at some history $h \in \mathcal{H}(\sigma_1, \sigma_2 | \theta_1)$. Let

$$\bar{V}_{\theta_1} \equiv \sup_{h \in \mathcal{H}(\sigma_1, \sigma_2 | \theta_1)} V_{\theta_1}(h), \quad (\text{E.2})$$

which is type θ_1 's highest continuation value in this equilibrium. Suppose by way of contradiction that $\bar{V}_{\theta_1} = 0$, then $V_{\theta_1}(h) = 0$ for every $h \in \mathcal{H}(\sigma_1, \sigma_2 | \theta_1)$ since type θ_1 can secure payoff 0 by playing \underline{a}_1 in every period. It is never optimal for player 1 to play a_1'' and then erase the record since it is strictly dominated by playing \underline{a}_1 and erasing the record. This implies that at any history h where type θ_1 has an incentive to play a_1'' , type θ_1 's continuation value $V_{\theta_1}(h)$ satisfies

$$V_{\theta_1}(h) = (1 - \delta)u_1(\theta_1, a_1'', \beta(h)) + \delta V_{\theta_1}(h, a_1'').$$

Therefore, $u_1(\theta_1, a_1'', \beta(h)) = 0$. Since u_1 is strictly increasing in a_2 and is strictly decreasing in a_1 , $a_1'' \succ \underline{a}_1$ implies that $\beta(h)$ FOSDs \underline{a}_2 . This implies that type θ_1 can secure payoff $(1 - \delta)u_1(\theta_1, \underline{a}_1, \beta(h))$ by playing \underline{a}_1 in every period, which is strictly positive. This contradicts the hypothesis that $\bar{V}_{\theta_1} = 0$.

Hence, it must be the case that $\bar{V}_{\theta_1} > 0$. For every ε that satisfies:

$$0 < \varepsilon < \min \left\{ \frac{\bar{V}_{\theta_1}}{2}, \frac{(1 - \delta)(\bar{c}(\theta_1) - c(\theta_1))}{\delta} \right\}, \quad (\text{E.3})$$

there exists a history $h(1) \in \mathcal{H}(\sigma_1, \sigma_2 | \theta_1)$ such that $V_{\theta_1}(h(1)) > \bar{V}_{\theta_1} - \varepsilon$. I consider type θ_1 's incentive at $h(1)$. His continuation value for playing \underline{a}_1 and then erasing it is at least

$$(1 - \delta) \left(u_1(\theta_1, \underline{a}_1, \beta(h(1))) - c(\theta_1) \right) + \delta(\bar{V}_{\theta_1} - \varepsilon).$$

His continuation value for playing any action $a_1 \succ a_1'$ is at most $(1 - \delta)u_1(\theta_1, a_1, \beta(h(1))) + \delta\bar{V}_{\theta_1}$. This upper bound is strictly less than $(1 - \delta)\{u_1(\theta_1, \underline{a}_1, \beta(h(1))) - c(\theta_1)\} + \delta(\bar{V}_{\theta_1} - \varepsilon)$. This implies that type θ_1 has no incentive to play any action weakly greater than a_1' at $h(1)$. Player 2 cannot have a strict incentive to play \underline{a}_2 at $h(1)$. This is because otherwise,

$$\bar{V}_{\theta_1} - \varepsilon < V_{\theta_1}(h(1)) \leq (1 - \delta)u_1(\theta_1, \underline{a}_1, \underline{a}_2) + \delta\bar{V}_{\theta_1}$$

which implies that $\bar{V}_{\theta_1} - \frac{\varepsilon}{1 - \delta} < u_1(\theta_1, \underline{a}_1, \underline{a}_2) = 0$ for every ε that satisfies (E.3). Therefore, $\bar{V}_{\theta_1} \leq 0$, which contradicts our earlier conclusion that $\bar{V}_{\theta_1} > 0$.

In order for player 2 to play actions other than \underline{a}_2 at $h(1)$, there must exist a type of player 1, denote it

by θ_2 , that occurs with positive probability at $h(1)$ and plays some action $a_1'' \succcurlyeq a_1'$ with positive probability at $h(1)$. As I shown before, type θ_2 has no incentive to erase a_1'' after playing it at history $h(1)$.

Consider the continuation game at history $(h(1), a_1'')$. Type θ_1 occurs with zero probability at that history since he never plays any action that is weakly greater than a_1' at $h(1)$. Let

$$\bar{V}_{\theta_2} \equiv \sup_{h \in \mathcal{H}(\sigma_1, \sigma_2 | \theta_2), h \succeq (h(1), a_1'')} V_{\theta_2}(h),$$

I show that $\bar{V}_{\theta_2} > 0$. This is because otherwise, $\bar{V}_{\theta_2} = 0 = V_{\theta_2}(h(1), a_1'')$, in which case type θ_2 has a strict incentive to deviate to \underline{a}_1 at $h(1)$. Applying the same argument as before, one can obtain that there exists $h(2) \succeq (h(1), a_1'')$ such that type θ_2 has no incentive to play any action weakly greater than a_1' and there exists another type θ_3 that occurs with positive probability at $h(2)$ and plays some action $a_1''' \succcurlyeq a_1'$ with positive probability. Since Θ is finite, one can obtain a contradiction after a finite number of iterations. \square

Next, I consider games with honest types. The next theorem generalizes Theorem 1 in the main text by showing that for any level of patience $\hat{\delta}$ and any full support distribution π , every opportunistic type's equilibrium payoff cannot be significantly greater than his minmax value 0 as his expected lifespan diverges.

Theorem. *For every $\pi \in \Delta(\Omega)$, $c(\theta) < \bar{c}(\theta)$ for every $\theta \in \Theta$, and $\hat{\delta} \in (0, 1)$, there exists $\delta^* \in (0, 1)$ such that for every $\bar{\delta} > \delta^*$ and $\theta \in \Theta$, type θ 's payoff in every equilibrium is no more than $\frac{(1-\bar{\delta})c(\theta)}{\bar{\delta}}$.*

The intuition is similar to that behind Theorem 1 in the main text. As long as there exists one opportunistic type $\theta \in \Theta$ whose payoff is strictly greater than $\frac{(1-\bar{\delta})c(\theta)}{\bar{\delta}}$, then the probability of the event that *player 1 is the opportunistic type and will erase \underline{a}_1 at every history* is bounded below by $\pi(\theta)$. As in the proof of Theorem 1, one can obtain that the number of periods it takes for player 1 to have a perfect reputation must be bounded from below by something proportional to $(1 - \bar{\delta})^{-1}$. However, since each opportunistic type can erase actions at a low cost, his incentive to take actions greater than a_1' implies that his continuation value needs to increase by something proportional to $1 - \delta$ as the length of his good record increases. This implies that the number of periods with which each opportunistic type may take actions other than \underline{a}_1 must be bounded above by something proportional to $(1 - \delta)^{-1}$. As $\bar{\delta} \rightarrow 1$, the lower bound will exceed the upper bound and this contradiction rules out equilibria in which the opportunistic type obtains high payoffs.

Proof. At any history where player 1's record length is no less than 1, player 2's belief assigns strictly positive probability to *at most one honest type*. I establish this theorem in a model with only one honest type a_1^* . This proof easily extends to models with multiple honest types playing stationary pure strategies.

Let a_2^* be player 2's strict best reply to a_1^* . Player 1 plays only a_1^* and \underline{a}_1 with positive probability, and according to the proposition earlier in this appendix, he plays a_1^* with positive probability only at histories that belong to \mathcal{H}_* .

Recall the definitions of h_*^k , β_k , and p_k in the main text. Let x_k denote the probability player 2's belief assigns to a_1^* at history h_*^k . Let $V_\theta(k)$ be type θ 's continuation value at history h_*^k . Recall the definition of $\mathcal{B} \subset \Delta(A_2)$ in the main text. Since u_2 does not depend on θ and satisfies Assumptions 1 and 3, every pair of elements in \mathcal{B} can be ranked according to FOSD. The case where $u_1(\theta, a_1^*, a_2^*) \leq 0$ for every $\theta \in \Theta$ is trivial. In what follows, I focus on the interesting case where $u_1(\theta, a_1^*, a_2^*) > 0$ for some $\theta \in \Theta$.

Type θ of player 1 prefers not to erase \underline{a}_1 at h_*^k if and only if $(1 - \delta)u_1(\theta, \underline{a}_1, \beta_k) \geq u_1(\theta, \underline{a}_1, \beta_k) - c(\theta)$, or equivalently,

$$u_1(\theta, \underline{a}_1, \beta_k) \leq \frac{c(\theta)}{\delta}. \quad (\text{E.4})$$

I only need to show that in every equilibrium, every opportunistic type $\theta \in \Theta$ has an incentive to play \underline{a}_1 and then not erase it at h_*^0 , since (E.4) will then imply that every type θ 's payoff is no more than $\frac{(1-\delta)c(\theta)}{\delta}$.

Suppose by way of contradiction that there exists an equilibrium such that there exists a type θ who has no incentive to *play \underline{a}_1 and then not erase it at h_*^0* . The rest of the proof consists of five steps.

Step 1: I show that in every equilibrium, there exists $t \in \mathbb{N}$ such that player 2 assigns probability 1 to the honest type starting from history h_*^t . Let \bar{V}_θ be type θ 's highest continuation value. Fix any ε that satisfies:

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

there exists $t_\theta \in \mathbb{N}$ such that $V_\theta(t_\theta) > \bar{V}_\theta - \varepsilon$. According to the proof of the proposition earlier in this appendix, type θ has no incentive to play a_1^* at $h_*^{t_\theta}$, which implies that player 2's belief assigns zero probability to type θ at history $h_*^{t_\theta+1}$. Suppose player 2's belief assigns positive probability to some opportunistic type at $h_*^{t_\theta+1}$, pick any type θ_* that it assigns positive probability to. Let \bar{V}_{θ_*} be type θ_* 's highest continuation value at histories that succeed $h_*^{t_\theta+1}$. Type θ_* has no incentive to play a_1^* when his continuation value is sufficiently close to \bar{V}_{θ_*} . Iterate this process finitely many times, we can find a history h_*^t at which player 2's belief assigns zero probability to all opportunistic types. In what follows, I use t to denote the smallest integer such that player 2's belief assigns zero probability to all opportunistic types at h_*^t .

Step 2: I derive an upper bound on t . At every history h_*^k with $k < t - 1$, there exists at least one opportunistic type that occurs with positive probability at h_*^k who plays a_1^* with positive probability. If

a type does not play a_1^* with positive probability at h_*^k , then it occurs with zero probability at h_*^s for every $s > k$. Hence, there exists $\theta \in \Theta$ such that type θ plays a_1^* with positive probability at h_*^k for every $k < t-1$. Type θ 's continuation value at h_*^k satisfies:

$$V_\theta(k) = (1 - \delta)u_1(\theta, a_1^*, \beta_k) + \delta V_\theta(k+1) \geq u_1(\theta, \underline{a}_1, \beta_k) - c(\theta), \quad (\text{E.5})$$

where the RHS is type θ 's payoff if he plays \underline{a}_1 and erases it in every subsequent period. Therefore,

$$V_\theta(k+1) - V_\theta(k) = \frac{1 - \delta}{\delta} \left\{ V_\theta(k) - u_1(\theta, a_1^*, \beta_k) \right\} \geq \frac{1 - \delta}{\delta} \left\{ u_1(\theta, \underline{a}_1, \beta_k) - c(\theta) - u_1(\theta, a_1^*, \beta_k) \right\}$$

Since $c(\theta) < \bar{c}(\theta)$, there exists $\Delta(\theta) > 0$ such that $u_1(\theta, \underline{a}_1, b) - c(\theta) - u_1(\theta, a_1^*, b) \geq \Delta(\theta)$ for every $b \in B$. Therefore,

$$V_\theta(k+1) - V_\theta(k) \geq \frac{1 - \delta}{\delta} \Delta(\theta) > 0. \quad (\text{E.6})$$

Since type θ 's continuation value is at least 0 and is at most $\bar{u}_1(\theta) \equiv \max_{a_1, a_2} u_1(\theta, a_1, a_2)$, an upper bound on t is given by

$$t \leq \sum_{\theta \in \Theta} \frac{\delta \cdot \bar{u}_1(\theta)}{(1 - \delta)\Delta(\theta)}. \quad (\text{E.7})$$

Step 3: I show that for every $k \leq t-1$ and every type $\theta \in \Theta$ that occurs with positive probability at both h_*^k and h_*^{k+1} , if type θ has no incentive to play \underline{a}_1 and not erase it at h_*^k , then he has no incentive to play \underline{a}_1 and not erase it at h_*^{k+1} . Suppose by way of contradiction that there exist such k and θ , then it must be the case that $\beta_k \succeq \beta_{k+1}$. Since type θ occurs with positive probability at h_*^{k+1} , my hypothesis implies that type θ weakly prefers playing a_1^* at h_*^k and then playing \underline{a}_1 and not erasing at h_*^{k+1} to the following two strategies (i) playing \underline{a}_1 and erasing in every subsequent period after reaching h_*^k as well as (ii) playing \underline{a}_1 and not erasing in every subsequent period after reaching h_*^k . These two incentive constraints imply that

$$(1 - \delta)u_1(\theta, a_1^*, \beta_k) + \delta(1 - \delta)u_1(\theta, \underline{a}_1, \beta_{k+1}) \geq u_1(\theta, \underline{a}_1, \beta_k) - c(\theta) \quad (\text{E.8})$$

and

$$(1 - \delta)u_1(\theta, a_1^*, \beta_k) + \delta(1 - \delta)u_1(\theta, \underline{a}_1, \beta_{k+1}) \geq (1 - \delta)u_1(\theta, \underline{a}_1, \beta_k) \quad (\text{E.9})$$

Since $c(\theta) < \bar{c}(\theta)$, we have $u_1(\theta, \underline{a}_1, \beta_k) - c(\theta) > u_1(\theta, a_1^*, \beta_k)$. Therefore, (E.8) together with $\beta_k \succeq \beta_{k+1}$ implies that $(1 - \delta)u_1(\theta, \underline{a}_1, \beta_{k+1}) > u_1(\theta, a_1^*, \beta_k) \geq u_1(\theta, a_1^*, \beta_{k+1})$. Inequality (E.9) implies that $u_1(\theta, a_1^*, \beta_k) \geq u_1(\theta, \underline{a}_1, \beta_k) - \delta u_1(\theta, \underline{a}_1, \beta_{k+1}) \geq (1 - \delta)u_1(\theta, \underline{a}_1, \beta_k)$. This leads to a contradiction.

Step 4: For any equilibrium (σ_1, σ_2) in which players' strategies depend only on player 2's history, we modify player 1's strategy to σ_1^* such that (σ_1^*, σ_2) remains a Nash equilibrium and that player 2's expectation about player 1's action at every on-path history is the same under (σ_1, σ_2) and under (σ_1^*, σ_2) . That is to say, the two equilibria are equivalent.

Since player 2's action belongs to \mathcal{B} at every on-path history and every pair of elements in \mathcal{B} can be ranked according to FOSD (see Lemma 1 in the main text), for every $\theta \in \Theta$, there exists at most one $\beta^*(\theta) \in \mathcal{B}$ such that inequality (E.4) holds with equality. The conclusion in Step 3 implies that for every $\theta \in \Theta$, there exists *at most one* period $k(\theta) \leq t - 1$ such that $\beta_{k(\theta)} = \beta^*(\theta)$.

I describe every opportunistic type's strategy under σ_1^* . For every $\theta \in \Theta$ such that $k(\theta)$ does not exist, type θ 's strategies under σ_1 and under σ_1^* are the same. For every $\theta \in \Theta$ such that $k(\theta)$ exists, type θ 's actions under σ_1 and under σ_1^* are the same at every history except for $h_*^{k(\theta)}$. At history $h_*^{k(\theta)}$, type θ erases \underline{a}_1 with probability 1 if calendar time is strictly above $k(\theta)$, and erases \underline{a}_1 with probability $p(\theta) \in [0, 1]$ if calendar time equals $k(\theta)$. There exists $p(\theta)$ such that player 2's belief about type θ 's action at $h_*^{k(\theta)}$ remains the same. This is because (i) when $p(\theta) = 1$, player 2 believes that type θ plays \underline{a}_1 and then does not erase it, (ii) when $p(\theta) = 0$, player 2 believes that type θ either plays a_1^* or plays \underline{a}_1 and then erases it, and (iii) player 2's belief changes continuously with $p(\theta)$.

Step 5: I derive a lower bound on t based on the equilibrium (σ_1^*, σ_2) we constructed in Step 4. This lower bound also applies to (σ_1, σ_2) since the two are equivalent. Let \mathcal{E}^* denote the event that *player 1 is opportunistic and erases \underline{a}_1 whenever he plays it*. Let $\widehat{\mathcal{E}}$ denote the event that *player 1 is opportunistic and does not erase \underline{a}_1 after he plays it*. The conclusions in Step 3 and Step 4 imply that when player 1 is one of the opportunistic types, either event \mathcal{E}^* or event $\widehat{\mathcal{E}}$ will happen under the probability measure induced by (σ_1^*, σ_2) . Let π^* denote the probability of event \mathcal{E}^* . Let $\widehat{\pi}$ denote the probability of event $\widehat{\mathcal{E}}$. Let π denote the probability that player 1 is the honest type. Let θ denote the type such that playing \underline{a}_1 and then erasing it at h_*^0 is strictly suboptimal. The conclusion in Step 3 implies that $\pi^* \geq \pi(\theta) > 0$.

Let μ_k^* denote the probability of history h_*^k conditional on event \mathcal{E}^* . Let p_k^* denote the probability that player 1 plays a_1^* conditional on event \mathcal{E}^* and the history in the current period being h_*^k . Let $\widehat{\mu}_k$ denote the probability of history h_*^k conditional on event $\widehat{\mathcal{E}}$. Let \widehat{p}_k denote the probability that player 1 plays a_1^* conditional on event $\widehat{\mathcal{E}}$ and the history in the current period being h_*^k . The state distribution lemma in Appendix A of the main text implies that:

$$\mu_0^* = (1 - \bar{\delta}) + \bar{\delta}\mu_0^*(1 - p_0^*), \quad (\text{E.10})$$

$$\mu_k^* = \bar{\delta}\mu_{k-1}^*p_{k-1}^* + \bar{\delta}\mu_k^*(1 - p_k^*) \text{ for every } k \in \{1, 2, \dots, t-1\}, \quad (\text{E.11})$$

and

$$\hat{\mu}_k = (1 - \bar{\delta})\bar{\delta}^k \prod_{j=0}^{k-1} \hat{p}_j. \quad (\text{E.12})$$

Equation (E.11) implies that

$$\frac{\mu_k^*}{\mu_{k-1}^*} = \frac{\bar{\delta}p_{k-1}^*}{1 - \bar{\delta}(1 - p_k^*)} \text{ for every } k \in \{1, 2, \dots, t-1\}. \quad (\text{E.13})$$

Let x_k denote the probability that player 2's belief assigns to a_1^* at h_*^k . According to Bayes rule,

$$\frac{x_k}{1 - x_k} = \frac{\pi(1 - \bar{\delta})\bar{\delta}^k + \pi^*\mu_k^*p_k^* + \hat{\pi}\hat{\mu}_k\hat{p}_k}{\pi^*\mu_k^*(1 - p_k^*) + \hat{\pi}\hat{\mu}_k(1 - \hat{p}_k)} = \frac{\pi^*\mu_k^*p_k^* + (1 - \bar{\delta})\bar{\delta}^k I_k}{\pi^*\mu_k^*(1 - p_k^*) + (1 - \bar{\delta})\bar{\delta}^k J_k}, \quad (\text{E.14})$$

where $I_k \equiv \pi + \hat{\pi}\prod_{j=0}^k \hat{p}_j$ and $J_k \equiv \hat{\pi}(1 - \hat{p}_k)\prod_{j=0}^{k-1} \hat{p}_j$. Equation (E.14) implies that

$$\pi^*\mu_k^*(x_k - p_k^*) = (1 - \bar{\delta})\bar{\delta}^k \left\{ I_k - x_k(I_k + J_k) \right\} \leq (1 - \bar{\delta})\bar{\delta}^k (\pi + \hat{\pi}). \quad (\text{E.15})$$

I define two new sequences $\{\bar{p}_k\}_{k=0}^t$ and $\{\bar{\mu}_k\}_{k=0}^t$ according to

$$\bar{\mu}_k(x_k - \bar{p}_k) = (1 - \bar{\delta})\bar{\delta}^k \frac{\pi + \hat{\pi}}{\pi^*}, \quad (\text{E.16})$$

$$\bar{\mu}_0 = \frac{1 - \bar{\delta}}{1 - \bar{\delta}(1 - \bar{p}_0)} \text{ and } \bar{\mu}_k = \frac{\bar{\delta}\bar{p}_{k-1}\bar{\mu}_{k-1}}{1 - \bar{\delta}(1 - \bar{p}_k)}. \quad (\text{E.17})$$

Similar to the proof of Theorem 1 in the main text, one can show that \bar{p}_0 is bounded above 0 and $\bar{p}_{k-1} - \bar{p}_k < 1 - \bar{\delta}$, which implies that $\bar{p}_t = 0$ if and only if t is bounded below by something proportional to $(1 - \bar{\delta})^{-1}$.

In order to complete the proof, I only need to show by induction that $p_k \geq \bar{p}_k$ for every $k \leq t$, regardless of $\{I_k, J_k\}_{k=0}^t$. When $k = 0$, we have

$$\mu_0^* = \frac{1 - \bar{\delta}}{1 - \bar{\delta}(1 - p_0^*)} \text{ and } \mu_0^*(x_0 - p_0^*) \leq (1 - \bar{\delta}) \frac{\pi + \hat{\pi}}{\pi^*}.$$

Since μ_0^* is strictly decreasing in p_0^* , we obtain that p_0^* is bounded above 0, and moreover, $p_0^* \geq \bar{p}_0$. If $p_j^* \geq \bar{p}_j$ for every $j \leq k$, then when $j = k + 1$, we have

$$\mu_{k+1}^* = \frac{\bar{\delta}\mu_k^*p_k^*}{1 - \bar{\delta}(1 - p_{k+1}^*)} \text{ and } \mu_{k+1}^*(x_{k+1} - p_{k+1}^*) \leq (1 - \bar{\delta}) \frac{\pi + \hat{\pi}}{\pi^*}.$$

Since μ_j^* is strictly decreasing in p_j^* for every $j \leq k$, we know that the value of p_{k+1}^* is minimized when $\{p_0^*, \dots, p_k^*\}$ all reach their minimal values $\{\bar{p}_0, \dots, \bar{p}_k\}$. The definitions of $\{\bar{p}_k\}_{k=0}^t$ and $\{\bar{\mu}_k\}_{k=0}^t$ then imply that $p_{k+1}^* \geq \bar{p}_{k+1}$. This completes the proof that t is bounded below by something proportional to $(1 - \bar{\delta})^{-1}$. As $\bar{\delta} \rightarrow 1$, this lower bound exceeds the upper bound I derived in Step 2, which rules out equilibria in which some opportunistic type θ receives a payoff strictly greater than $\frac{(1-\delta)c(\theta)}{\delta}$. \square