Online Appendix: Reputation Effects with Endogenous Records

Harry PEI

April 1, 2024

A Proof of Lemma 1

Overview of the Proof: Throughout this proof, I will *fix* the parameters of the game $(u_1, u_2, c, \pi, \widehat{\delta}, \overline{\delta})$ as well as a Nash equilibrium (σ_1, σ_2) under that parameter configuration, I construct a strategy σ_1^* for the opportunistic type and verify that (σ_1^*, σ_2) is a public equilibrium under $(u_1, u_2, c, \pi, \widehat{\delta}, \overline{\delta})$ *using the hypothesis* that (σ_1, σ_2) is a Nash equilibrium under $(u_1, u_2, c, \pi, \widehat{\delta}, \overline{\delta})$. I start from introducing some notation.

Notation: Recall that $\mathcal{H}$ is the set of player 2's histories. Let $\overline{\mathcal{H}}$ denote the set of player 1's histories, which consists of his type, the sequence of actions players took in the past, as well as whether he has erased each of his past action. The opportunistic type's strategy is $\sigma_1 : \overline{\mathcal{H}} \rightarrow \Delta\big(A_1 \times \{0, c\}\big)$, which is a mapping from his histories to a distribution over his actions and his erasing decisions.[1] Player 2's strategy is $\sigma_2 : \mathcal{H} \rightarrow \Delta(A_2)$.

For any strategy profile (σ_1, σ_2), let $\mathcal{H}(\sigma_1, \sigma_2) \subset \mathcal{H}$ denote the set of player 2's histories that occur with positive probability under (σ_1, σ_2). For any $h \in \mathcal{H}(\sigma_1, \sigma_2)$, let $\alpha_1(h, \sigma_1, \sigma_2) \in \Delta(A_1)$ denote player 2's expectation of the *opportunistic type*'s action conditional on her history being h. For every a_1 that belongs to the support of $\alpha_1(h, \sigma_1, \sigma_2)$, let $c(a_1|h, \sigma_1, \sigma_2) \in [0, 1]$ denote the expected probability with which the *opportunistic type* erases a_1 after taking it conditional on player 2's history being h. Since A_1 and A_2 are finite sets, the definition of $\mathcal{H}(\sigma_1, \sigma_2)$ implies that these conditional probabilities are well-defined.

For every $h \in \mathcal{H}(\sigma_1, \sigma_2)$, let $\phi(h)$ denote the probability that player 2's history is h *conditional on* player 1 being the honest type. Since the honest type always plays a_1^* and never erases any action and player 2 *cannot* observe previous short-run players' actions, $\phi(h)$ is the same for all strategy profiles (σ_1, σ_2). Let $p_0(h)$ denote the probability that player 2's history is h conditional on calendar time is 0, which by definition, is the Dirac measure on h_*^0 and is independent of player 1's type. Let $\mu(h|\sigma_1, \sigma_2)$ denote the probability that player 2's history is h *conditional on* (i) player 1 being the opportunistic type and (ii) players behaving according to (σ_1, σ_2). Let $Q(h' \rightarrow h|\sigma_1, \sigma_2)$ denote the probability that player 2's history in the next period

[1] For every $a_1 \in A_1$, (a_1, c) stands for taking action a_1 and then erasing it and $(a_1, 0)$ stands for taking action a_1 and then not erasing it. For every $\overline{h} \in \overline{\mathcal{H}}$, $\sigma_1(\overline{h})$ is a distribution over the product set $A_1 \times \{0, c\}$.

is h *conditional on* (i) player 2's current-period history being h', (ii) player 1 being the opportunistic type, and (iii) players behaving according to (σ_1, σ_2). I establish the following lemma:

State Distribution Lemma. *For any strategy profile* (σ_1, σ_2) *and any* $h \in \mathcal{H}(\sigma_1, \sigma_2)$*, we have*

$$\mu(h|\sigma_1, \sigma_2) = (1 - \overline{\delta})p_0(h) + \overline{\delta} \sum_{h' \in \mathcal{H}} \mu(h'|\sigma_1, \sigma_2)Q(h' \to h|\sigma_1, \sigma_2). \tag{A.1}$$

Proof of Lemma: I will omit (σ_1, σ_2) in this proof in order to avoid cumbersome notation. For every $t \in \mathbb{N}$, let $p_t(h)$ denote the probability that player 2's history being h conditional on player 1 being the opportunistic type and the current calendar time being t, and let $q_t(h' \to h)$ denote the probability that player 2's current-period history being h conditional on her history in the period before being h', player 1 being the opportunistic type, and the calendar time in the period before being t. The law of total probability implies that $p_{t+1}(h) = \sum_{h' \in \mathcal{H}} p_t(h')q_t(h' \to h)$. Since player 2's prior belief assigns probability $(1 - \overline{\delta})\overline{\delta}^t$ to calendar time being t, by the law of total probability and Bayes rule, we have

$$\mu(h') = \sum_{t=0}^{+\infty} (1 - \overline{\delta})\overline{\delta}^t p_t(h') \quad \text{and} \quad Q(h' \to h) = \frac{\sum_{t=0}^{+\infty} (1 - \overline{\delta})\overline{\delta}^t p_t(h')q_t(h' \to h)}{\sum_{t=0}^{+\infty} (1 - \overline{\delta})\overline{\delta}^t p_t(h')}.$$

These equations imply the following two equations:

$$\sum_{h' \in \mathcal{H}} \mu(h')Q(h' \to h) = \sum_{h' \in \mathcal{H}} \sum_{t=0}^{+\infty} (1-\overline{\delta})\overline{\delta}^t p_t(h')q_t(h' \to h) = \sum_{t=0}^{+\infty} (1-\overline{\delta})\overline{\delta}^t \sum_{h' \in \mathcal{H}} p_t(h')q_t(h' \to h) = \sum_{t=0}^{+\infty} (1-\overline{\delta})\overline{\delta}^t p_{t+1}(h)$$

and

$$\overline{\delta}^{-1}\Big\{\mu(h) - (1 - \overline{\delta})p_0(h)\Big\} = \overline{\delta}^{-1}\Big\{\sum_{t=0}^{+\infty} (1 - \overline{\delta})\overline{\delta}^t p_t(h) - (1 - \overline{\delta})p_0(h)\Big\} = \sum_{t=0}^{+\infty} (1 - \overline{\delta})\overline{\delta}^t p_{t+1}(h).$$

These two equations together imply the following equation that is equivalent to (A.1):

$$\sum_{h' \in \mathcal{H}} \mu(h')Q(h' \to h) = \overline{\delta}^{-1}\Big\{\mu(h) - (1 - \overline{\delta})p_0(h)\Big\}.$$

□

This lemma implies that the distribution over player 2's histories $\{\mu(h|\sigma_1, \sigma_2)\}_{h \in \mathcal{H}}$ conditional on player 1 being the opportunistic type depends on the strategy profile (σ_1, σ_2) *only through* $\{Q(h' \to h|\sigma_1, \sigma_2)\}_{h,h' \in \mathcal{H}}$. The latter depends on (σ_1, σ_2) *only through* the distribution over the opportunistic type's behavior (i.e., action and erasing decision) *conditional on* player 2's history.

Construction of Public Strategy σ_1^*: I construct a strategy for the opportunistic type $\sigma_1^* : \mathcal{H} \to \Delta\big(A_1 \times \{0, c\}\big)$ from (σ_1, σ_2), which is *measurable with respect to player 2's history* and is defined as follows:

- At every $h \in \mathcal{H}(\sigma_1, \sigma_2)$, strategy σ_1^* asks the opportunistic type to play $\alpha_1(h, \sigma_1, \sigma_2)$, and to erase his action with probability $c(a_1|h, \sigma_1, \sigma_2)$ if his realized pure action is $a_1 \in \text{supp}(\alpha_1(h, \sigma_1, \sigma_2))$. This construction implies that $\alpha_1(h, \sigma_1, \sigma_2) = \alpha_1(h, \sigma_1^*, \sigma_2)$ and $c(a_1|h, \sigma_1, \sigma_2) = c(a_1|h, \sigma_1^*, \sigma_2)$.

I will leave the behaviors at $h \notin \mathcal{H}(\sigma_1, \sigma_2)$ unspecified since they are not relevant for verifying the conditions of public equilibria. The rest of my proof proceeds in five steps.

Step 1: I show that $\mathcal{H}(\sigma_1, \sigma_2) = \mathcal{H}(\sigma_1^*, \sigma_2)$ and $\mu(h|\sigma_1, \sigma_2) = \mu(h|\sigma_1^*, \sigma_2)$ for every $h \in \mathcal{H}(\sigma_1, \sigma_2)$, that is, the distribution over on-path histories are the same under (σ_1, σ_2) and under (σ_1^*, σ_2).

Suppose by way of contradiction that there exists $h \in \mathcal{H}(\sigma_1, \sigma_2) \backslash \mathcal{H}(\sigma_1^*, \sigma_2)$. Since the empty history $\emptyset$ belongs to $\mathcal{H}(\sigma_1, \sigma_2) \cap \mathcal{H}(\sigma_1^*, \sigma_2)$, there exists $h^* \in \mathcal{H}(\sigma_1, \sigma_2) \backslash \mathcal{H}(\sigma_1^*, \sigma_2)$ such that the immediate predecessor of h^*, denoted by h^{**}, belongs to $\mathcal{H}(\sigma_1, \sigma_2) \bigcap \mathcal{H}(\sigma_1^*, \sigma_2)$. Therefore, *either* there exists an action $a_1 \in A_1$ such that $\alpha_1(h^{**}, \sigma_1, \sigma_2)$ assigns positive probability to a_1 but $\alpha_1(h^{**}, \sigma_1^*, \sigma_2)$ does not, *or* there exists a_1 that occurs with positive probability under both $\alpha_1(h^{**}, \sigma_1, \sigma_2)$ and $\alpha_1(h^{**}, \sigma_1, \sigma_2)$ but the opportunistic type erases a_1 with different expected probabilities under the two strategies. Both of these cases are at odds with the construction of σ_1^* under which $\alpha_1(h^{**}, \sigma_1, \sigma_2) = \alpha_1(h^{**}, \sigma_1^*, \sigma_2)$ and the expected probabilities with which they erase each action are the same. Similarly, one can show that there exists no h that belongs to $\mathcal{H}(\sigma_1^*, \sigma_2) \backslash \mathcal{H}(\sigma_1, \sigma_2)$. The two parts together imply that $\mathcal{H}(\sigma_1, \sigma_2) = \mathcal{H}(\sigma_1^*, \sigma_2)$.

My construction of σ_1^* implies that for every $h \in \mathcal{H}(\sigma_1, \sigma_2)$, (i) player 2's expectations of player 1's action at h are the same under (σ_1, σ_2) and under (σ_1^*, σ_2), that is, $\alpha_1(h, \sigma_1, \sigma_2) = \alpha_1(h, \sigma_1^*, \sigma_2)$, and (ii) her expected probabilities with which player 1 will erase each action $a_1 \in \text{supp}(\alpha_1(h, \sigma_1, \sigma_2))$ at h are the same under (σ_1, σ_2) and under (σ_1^*, σ_2), that is, $c(a_1|h, \sigma_1, \sigma_2) = c(a_1|h, \sigma_1^*, \sigma_2)$. The lemma I established earlier then implies that $\mu(h|\sigma_1, \sigma_2) = \mu(h|\sigma_1^*, \sigma_2)$ for every $h \in \mathcal{H}(\sigma_1, \sigma_2) = \mathcal{H}(\sigma_1^*, \sigma_2)$.

Step 2: I show that (σ_1^*, σ_2) is a public equilibrium as long as (σ_1, σ_2) is a Nash equilibrium. For this purpose, I only need to show that (σ_1^*, σ_2) is a Nash equilibrium as long as (σ_1, σ_2) is a Nash equilibrium. This is because by construction, σ_1^* depends only on player 2's history.

First, I show that σ_2 best replies to σ_1^* as long as σ_2 best replies to σ_1. If σ_2 is player 2's best reply to σ_1, then for every $h \in \mathcal{H}(\sigma_1, \sigma_2)$, player 2's action at h, denoted by $\sigma_2(h)$, best replies to player 2's *expectation* of player 1's action at h under (σ_1, σ_2), which according to Bayes rule, equals

$$\underbrace{\frac{(1-\pi)\mu(h|\sigma_1,\sigma_2)}{(1-\pi)\mu(h|\sigma_1,\sigma_2)+\pi\phi(h)}}_{\text{prob of opportunistic type conditional on } h} \alpha_1(h,\sigma_1,\sigma_2) + \underbrace{\frac{\pi\phi(h)}{(1-\pi)\mu(h|\sigma_1,\sigma_2)+\pi\phi(h)}}_{\text{prob of honest type conditional on } h} a_1^*. \tag{A.2}$$

Alternatively, under (σ_1^*, σ_2), player 2's expectation of player 1's action at h is:

$$\frac{(1-\pi)\mu(h|\sigma_1^*,\sigma_2)}{(1-\pi)\mu(h|\sigma_1^*,\sigma_2)+\pi\phi(h)}\alpha_1(h,\sigma_1^*,\sigma_2) + \frac{\pi\phi(h)}{(1-\pi)\mu(h|\sigma_1^*,\sigma_2)+\pi\phi(h)}a_1^*. \tag{A.3}$$

Since $\alpha_1(h, \sigma_1, \sigma_2) = \alpha_1(h, \sigma_1^*, \sigma_2)$, $\mu(h|\sigma_1, \sigma_2) = \mu(h^*|\sigma_1, \sigma_2)$ for every $h \in \mathcal{H}(\sigma_1, \sigma_2)$, and the distribution over player 2's histories *conditional on player* 1 *being the honest type* $\phi(\cdot)$ does not depend on (σ_1, σ_2), player 2's expectations of player 1's action at h are the same under (σ_1, σ_2) and under (σ_1^*, σ_2). Therefore, $\sigma_2(h)$ also best replies to player 2's expectation of player 1's action at h under (σ_1^*, σ_2). Since $\mathcal{H}(\sigma_1, \sigma_2) = \mathcal{H}(\sigma_1^*, \sigma_2)$, we know that player 2's strategy σ_2 best replies to player 2's expectation of player 1's action at h under (σ_1^*, σ_2) for every $h \in \mathcal{H}(\sigma_1^*, \sigma_2)$. This implies that σ_2 best replies to σ_1^*.

Next, I show that σ_1^* best replies to σ_2 as long as σ_1 best replies to σ_2. This part uses the observation that player 1 can observe player 2's history, which implies that player 2's strategy is also measurable with respect to player 1's history. For any player 1's history $\overline{h} \in \overline{\mathcal{H}}$ and player 2's history $h \in \mathcal{H}$, I say that $\overline{h}$ is *consistent with* h if player 1's history being $\overline{h}$ implies that player 2's history being h.

For every $h \in \mathcal{H}(\sigma_1^*, \sigma_2) = \mathcal{H}(\sigma_1, \sigma_2)$ and every $(a_1, \widetilde{c}) \in A_1 \times \{0, c\}$ that the opportunistic type of player 1 chooses with positive probability under σ_1^* at history h, the construction of σ_1^* implies that there exists a player 1's history $\overline{h}$ that is (i) consistent with h and (ii) occurs with positive probability under (σ_1, σ_2), such that player 1 plays $(a_1, \widetilde{c})$ with positive probability at $\overline{h}$ under (σ_1, σ_2). Since player 2's strategy σ_2 is measurable with respect to player 2's history and the opportunistic type's payoff depends only on his action and player 2's action, the opportunistic type's best reply problem is also measurable with respect to player 2's history. Under the hypothesis that (σ_1, σ_2) is a Nash equilibrium, we know that $(a_1, \widetilde{c})$ is optimal for the opportunistic type at $\overline{h}$ when player 2's strategy is σ_2. As a result, $(a_1, \widetilde{c})$ must also be optimal for the opportunistic type at any player 1's history $\overline{h}'$ that is consistent with h when player 2's strategy is σ_2. Hence, every pair $(a_1, \widetilde{c})$ that σ_1^* chooses with positive probability at any $h \in \mathcal{H}(\sigma_1^*, \sigma_2)$ is optimal. This implies that σ_1^* is the opportunistic type's best reply against σ_2.

Step 3: I show that at every $h \in \mathcal{H}(\sigma_1, \sigma_2)$, player 2's belief about player 1's type and behavior (action and erasing probability) are the same under (σ_1, σ_2) and under (σ_1^*, σ_2).

According to Bayes rule, player 1's reputation at h under (σ_1, σ_2) is $\frac{\pi\phi(h)}{(1-\pi)\mu(h|\sigma_1,\sigma_2)+\pi\phi(h)}$ and his reputation at h under (σ_1^*, σ_2) is $\frac{\pi\phi(h)}{(1-\pi)\mu(h|\sigma_1^*,\sigma_2)+\pi\phi(h)}$, which are the same since $\mu(h|\sigma_1, \sigma_2) = \mu(h|\sigma_1^*, \sigma_2)$.

According to (A.2) and (A.3), player 2's beliefs about player 1's actions are the same conditional on every $h \in \mathcal{H}(\sigma_1, \sigma_2)$. Conditional on every $h \in \mathcal{H}(\sigma_1, \sigma_2)$ and every $a_1 \in \text{supp}\big(\alpha_1(h, \sigma_1, \sigma_2)\big)$, when the opportunistic type uses strategy σ_1, the expected probability with which player 1 erases a_1 equals $c(a_1|h, \sigma_1, \sigma_2)$ for every $a_1 \neq a_1^*$, and when $a_1 = a_1^*$, it equals

$$\underbrace{\frac{(1-\pi)\mu(h|\sigma_1,\sigma_2)p(h|\sigma_1)}{(1-\pi)\mu(h|\sigma_1,\sigma_2)p(h|\sigma_1)+\pi\phi(h)}}_{\text{prob of opportunistic type conditional on taking action } a_1^* \text{ at history } h} \qquad c(a_1^*|h,\sigma_1,\sigma_2) \tag{A.4}$$

where $p(h|\sigma_1)$ is the *expected* probability that the opportunistic type takes action a_1^* conditional on player 2's history being h when the opportunistic type plays according to strategy σ_1. When the opportunistic type uses strategy σ_1^*, the expected probability with which player 1 erases a_1 equals $c(a_1|h, \sigma_1^*, \sigma_2)$ for every $a_1 \neq a_1^*$, and when $a_1 = a_1^*$, it equals

$$\frac{(1-\pi)\mu(h|\sigma_1^*,\sigma_2)p(h|\sigma_1^*)}{(1-\pi)\mu(h|\sigma_1^*,\sigma_2)p(h|\sigma_1^*)+\pi\phi(h)}c(a_1^*|h,\sigma_1^*,\sigma_2) \tag{A.5}$$

where $p(h|\sigma_1^*)$ is the *expected* probability that the opportunistic type takes action a_1^* conditional on player 2's history being h when the opportunistic type plays according to strategy σ_1^*. The values of (A.4) and (A.5) are the same given my earlier conclusion that $\mu(h|\sigma_1, \sigma_2) = \mu(h|\sigma_1^*, \sigma_2)$ and $p(h|\sigma_1) = p(h|\sigma_1^*)$.

Step 4: I show that the opportunistic type receives the same discounted average payoff under (σ_1^*, σ_2) and under (σ_1, σ_2). Let Σ_1 denote the set of the opportunistic type's strategies in the repeated game. Since (σ_1, σ_2) is a Nash equilibrium, we have $U_1(\sigma_1, \sigma_2) \geq U_1(\sigma_1', \sigma_2)$ for every $\sigma_1' \in \Sigma_1$, where $U_1(\cdot, \cdot)$ denotes the opportunistic type's discounted average payoff as a function of his own strategy σ_1 and player 2's strategy σ_2. Since (σ_1^*, σ_2) is also a Nash equilibrium, we have $U_1(\sigma_1^*, \sigma_2) \geq U_1(\sigma_1', \sigma_2)$ for every $\sigma_1' \in \Sigma_1$. Since both σ_1 and σ_1^* are best replies to σ_2, they must yield the same discounted average payoff. Therefore, the opportunistic type receives the same discounted average payoff under (σ_1^*, σ_2) and under (σ_1, σ_2).

Step 5: I show that the sum of the short-run players' payoffs are the same under (σ_1^*, σ_2) and under (σ_1, σ_2). I use two observations. First, for every $h \in \mathcal{H}(\sigma_1, \sigma_2)$, conditional on observing history h, player 2's payoffs are the same under (σ_1, σ_2) and (σ_1^*, σ_2), which I denote by $v(h)$. This is because her expectations of player 1's actions are the same and σ_2 is player 2's best reply against both σ_1^* and σ_1.

Second, since I have shown before that $\mu(h|\sigma_1, \sigma_2) = \mu(h|\sigma_1^*, \sigma_2)$ for every $h \in \mathcal{H}(\sigma_1, \sigma_2)$ and that the distribution of h conditional on player 1 being the honest type does not depend on players' strategy profile, the unconditional distributions over player 2's histories are the same under (σ_1, σ_2) and (σ_1^*, σ_2), which I denote by $\widetilde{\mu} \in \Delta(\mathcal{H})$. I use 2_k to denote the player 2 who arrives in period k. Let $\widetilde{\mu}_k(h)$ denote the probability that player 2_k observes history h under (σ_1, σ_2). Let $\widetilde{\mu}_k^*(h)$ denote the probability that player 2_k observes history h under (σ_1^*, σ_2). Since player 2's prior belief assigns probability $(1-\overline{\delta})\overline{\delta}^k$ to the calendar time being k, we have

$$\widetilde{\mu}(h) = \sum_{k=0}^{+\infty} (1-\overline{\delta})\overline{\delta}^k \widetilde{\mu}_k(h) = \sum_{k=0}^{+\infty} (1-\overline{\delta})\overline{\delta}^k \widetilde{\mu}_k^*(h).$$

This implies that

$$\mathbb{E}^{(\sigma_1,\sigma_2)}\Big[\sum_{k=0}^{+\infty}(1-\overline{\delta})\overline{\delta}^k u_2(a_{1,k}, a_{2,k})\Big] = \sum_{k=0}^{+\infty}\Big\{(1-\overline{\delta})\overline{\delta}^k \sum_{h\in\mathcal{H}(\sigma_1,\sigma_2)} v(h)\widetilde{\mu}_k(h)\Big\} = \sum_{h\in\mathcal{H}(\sigma_1,\sigma_2)} v(h)\Big(\sum_{k=0}^{+\infty}(1-\overline{\delta})\overline{\delta}^k \widetilde{\mu}_k(h)\Big)$$

$$= \sum_{h\in\mathcal{H}(\sigma_1,\sigma_2)} v(h)\widetilde{\mu}(h) = \sum_{h\in\mathcal{H}(\sigma_1,\sigma_2)} v(h)\Big(\sum_{k=0}^{+\infty}(1-\overline{\delta})\overline{\delta}^k \widetilde{\mu}_k^*(h)\Big) = \mathbb{E}^{(\sigma_1^*,\sigma_2)}\Big[\sum_{k=0}^{+\infty}(1-\overline{\delta})\overline{\delta}^k u_2(a_{1,k}, a_{2,k})\Big].$$

This implies that the expected sum of player 2's payoffs are the same under (σ_1, σ_2) and under (σ_1^*, σ_2).

B Inability to Erase Actions or High Cost of Erasing Actions

Proof of Proposition 1: Recall from the definition of γ that player 2 has a strict incentive to play a_2^* when player 1's mixed action assigns probability at least $\gamma \in (0,1)$ to a_1^*. Fix the parameter values (u_1, u_2, δ, π) as well as an arbitrary Nash equilibrium (σ_1, σ_2) under these parameters. I bound the opportunistic type's payoff from below when he deviates from σ_1 and instead, plays a_1^* in every period, which I denote by U_1^*. Since (σ_1, σ_2) is a Nash equilibrium, the opportunistic type's equilibrium payoff is weakly greater than U_1^*.

Suppose player 1 deviates to playing a_1^* in every period. In every period, (i) either player 2 plays a_2^* with probability 1, or (ii) she has an incentive to take actions other than a_2^*. In the second case, her belief assigns probability less than γ to a_1^*. According to Bayes rule, in every such period, the posterior belief she assigns to the honest type is multiplied by at least $1/\gamma$. Since her prior belief assigns probability π to the honest type and the probability of the honest type is non-decreasing over time when player 1 plays a_1^* in every period, there can be at most $\frac{\log \pi}{\log \gamma}$ periods in which player 2 has an incentive to take actions other than a_2^*. Since $u_1(a_1, a_2)$ is strictly increasing in a_2, player 1's stage-game payoff is at least $u_1(a_1^*, \underline{a}_2)$ when he plays a_1^*. As a result,

$$U_1^* \geq (1-\delta^{\frac{\log \pi}{\log \gamma}}) u_1(a_1^*, \underline{a}_2) + \delta^{\frac{\log \pi}{\log \gamma}} u_1(a_1^*, a_2^*). \tag{B.1}$$

The payoff lower bound in Proposition 1 is obtained once we rearrange the terms in (B.1).

High Cost of Erasing Actions: I maintain all the assumptions in the baseline model except for Assumption 4. Instead, I assume that player 1's cost of erasing actions c is large enough in the sense that

$$c > \widetilde{c} \equiv \max_{a_2 \in A_2} \Big\{ u_1(\underline{a}_1, a_2) - u_1(a_1^*, a_2) \Big\}. \tag{B.2}$$

In the following product choice game, we have $\widetilde{c} = g$ and condition (B.2) translates into $c > g$.

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 $x \in (0,1)$.

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 $\widehat{\delta}$ and $\overline{\delta}$ only through the product of the two δ.

Theorem. *Suppose (u_1, u_2) satisfies Assumptions 1, 2, and 3. For every $c > \widetilde{c}$, $\pi \in (0,1)$, 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 $(1-\overline{\delta})\overline{\delta}^k$ to player 1's age in the game being k, the state distribution lemma in Online Appendix A implies that

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

$$\text{and} \quad \mu_k = \overline{\delta}\Big\{(1-p_{k-1}-q_{k-1})\mu_{k-1} + q_k \mu_k\Big\} \text{ for every } k \geq 1,$$

which is equivalent to

$$\mu_0 = \frac{1-\overline{\delta}}{1-\overline{\delta} q_0} \tag{B.3}$$

$$\text{and} \quad \frac{\mu_k}{\mu_{k-1}} = \frac{\overline{\delta}(1-p_{k-1}-q_{k-1})}{1-\overline{\delta} q_k} \text{ for every } k \geq 1. \tag{B.4}$$

Equations (B.3) and (B.4) together imply that for every $k \geq 1$, we have:

$$\frac{\mu_k}{(1-\overline{\delta})\overline{\delta}^k} = \frac{1}{1-\overline{\delta} q_k} \cdot \prod_{i=0}^{k-1} \frac{1-p_i-q_i}{1-\overline{\delta} q_i}. \tag{B.5}$$

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

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

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

$$\frac{1-p_k-q_k}{1-\overline{\delta} q_k} \leq \frac{x}{1-\overline{\delta}+\overline{\delta} x} < 1. \tag{B.6}$$

Moreover,

$$\frac{1}{1-\overline{\delta} q_k} \leq \frac{1}{1-\overline{\delta}}. \tag{B.7}$$

Suppose by way of contradiction that there are infinitely many k such that $\frac{\mu_k}{(1-\overline{\delta})\overline{\delta}^k} \geq y$ and $p_k + q_k \geq 1 - x$. Combining (B.5), (B.6) and (B.7), we know that there exists $t \in \mathbb{N}$ such that as long as there are t such periods before history h_*^s,

$$\frac{\mu_s}{(1-\overline{\delta})\overline{\delta}^s} \leq \Big(\frac{x}{1-\overline{\delta}+\overline{\delta} x}\Big)^t \cdot \frac{1}{1-\overline{\delta}} < y. \tag{B.8}$$

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

$$u_1(\underline{a}_1, \beta_k) - c \geq (1-\delta) u_1(a_1^*, \beta_k) + \delta V_{k+1}. \tag{B.9}$$

Since $c > \widetilde{c}$, we know that $u_1(\underline{a}_1, \beta_k) - c < u_1(a_1^*, \beta_k)$. Therefore, inequality (B.9) 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 $u_1(a_1^*, a_2^*)$, which is weakly greater than $u_1(a_1^*, \beta_k)$ by taking action a_1^* in every period starting from history h_*^{k+1}. 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 $(1-\overline{\delta})\overline{\delta}^k$, player 1's reputation at h_*^k, denoted by π_k, satisfies:

$$\frac{1-\pi_k}{\pi_k} = \frac{1-\pi}{\pi} \cdot \frac{\mu_k}{(1-\overline{\delta})\overline{\delta}^k}. \tag{B.10}$$

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.5),

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

Therefore, $\pi_{t+1} \geq \pi$ follows from (B.10). 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_*^k, he prefers to take action $\underline{a}_2$ and then erases it, to taking action a_1^*. This leads to the following incentive constraint:

$$u_1(\underline{a}_1, \beta_t) - c \geq (1-\delta)u_1(a_1^*, \beta_t) + \delta V_{t+1} \geq (1-\delta)u_1(a_1^*, \beta_t) + \delta\big(u_1(a_1^*, a_2^*) - \varepsilon\big). \tag{B.11}$$

However, since $c > \widetilde{c}$, inequality (B.11) 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 $\delta > \delta^*$. □

C Alternative Promises & Information Disclosure Policies

I adopt the "imperfect promise" interpretation of my model. The long-run player commits to play a_1^* in every period. However, instead of committing to fully reveal all past actions, he commits to a general *disclosure policy* according to which he discloses his past actions to the short-run players. The short-run players believe that the long-run player will honor his promise with probability π and with complementary probability, he may take other actions and may also erase actions from his records at cost c. Whether the long-run player can directly observe the signals received by the short-run players does not affect my results.

I restrict attention to *disclosure policies* that take the form of a mapping $\mathbf{q} \equiv \{q_m(n)\}_{n,m\in\mathbb{N}}$ such that if the long-run player's record length is $m \in \mathbb{N}$, then he commits to reveal a record with length $n \in \mathbb{N}$ consisting only of a_1^* with probability $q_m(n)$.[2] This class of disclosure policies includes disclosing the last K actions, randomizing between disclosing all past actions and disclosing no past action, and so on.[3] It rules out disclosure policies that requires the long-run player to fabricate records, namely, revealing an action that he has never taken to the short-run players.

I derive an *upper bound* on the long-run player's discounted average payoff *when he can renege*. This is because his payoff when he honors his promise is weakly lower than his payoff when he can renege, which

[2]Since the long-run player commits to play a_1^* in every period, all his past actions are a_1^* conditional on the event that he honors his promise, in which case his history can be summarized by the number of a_1^* that he has taken.

[3]There are other reasonable restrictions on the set of disclosure policies that the long-run player can commit to, such as the length of record he reveals cannot exceed the number of actions that he has taken. I do not include any additional restriction since my result applies to *all* disclosure policies that belong to the class I just specified, which makes the result stronger.

implies that player 1's expected payoff is also lower than the upper bound I derived.

I start from a benchmark in which the long-run player commits to reveal the null history to the short-run players regardless of his past actions, that is, $q_m(0) = 1$ for every $m \in \mathbb{N}$. For every $\pi \in (0,1)$, let $\overline{a}_2^{\pi}$ denote player 2's *highest* best reply to player 1's mixed action $\pi a_1^* + (1-\pi)\underline{a}_1$ and let $\underline{a}_2^{\pi}$ denote player 2's *lowest* best reply to $\pi a_1^* + (1-\pi)\underline{a}_1$. The following lemma derives lower and upper bounds on player 1's equilibrium payoff conditional on the event that he can renege on his promise:

Lemma. *If $q_m(0) = 1$ for every $m \in \mathbb{N}$, then conditional on player 1 can renege, his payoff in every equilibrium is at least* $\max\{(1-\delta)u_1(\underline{a}_1, \underline{a}_2^{\pi}), u_1(\underline{a}_1, \underline{a}_2^{\pi}) - c\}$ *and is at most* $\max\{\frac{1-\delta}{\delta}c, u_1(\underline{a}_1, \overline{a}_2^{\pi}) - c\}$.

For generic probability with which the long-run player will honor his promise π, player 2 has a unique best reply to $\pi a_1^* + (1-\pi)\underline{a}_1$, in which case $\underline{a}_2^{\pi} = \overline{a}_2^{\pi}$ and the payoff lower and upper bounds derived in the lemma will converge to the same value as $\delta \to 1$.

Proof. Since the long-run player commits to reveal the null record, when he can renege, it is optimal for him to take action $\underline{a}_1$ in every period. This is because $\underline{a}_1$ is player 1's strictly dominant action in the stage game and after player 1 shows any action to the short-run players, he reveals that he will not honor his promise and will therefore, receive a continuation value of 0. Therefore, at history h_*^0, the short-run player's action is at least $\underline{a}_2^{\pi}$ and is at most $\overline{a}_2^{\pi}$, and at any other on-path history, the short-run player's action is $\underline{a}_2$.

Let us bound the long-run player's payoff conditional on the event that he can renege on his promise. His payoff is at least $(1-\delta)u_1(\underline{a}_1, \underline{a}_2^{\pi})$ if he plays $\underline{a}_1$ in period 0 and does not erase it; his payoff is at least $u_1(\underline{a}_1, \underline{a}_2^{\pi}) - c$ if he plays $\underline{a}_1$ in every period and then erases it. Hence, his equilibrium payoff is at least $\max\{(1-\delta)u_1(\underline{a}_1, \underline{a}_2^{\pi}), u_1(\underline{a}_1, \underline{a}_2^{\pi}) - c\}$. His payoff is at most $u_1(\underline{a}_1, \overline{a}_2^{\pi}) - c$ if he plays $\underline{a}_1$ and erases it after each period and is at most $\frac{1-\delta}{\delta}c$ if he plays $\underline{a}_1$ in period 0 and then does not erase it. Since playing $\underline{a}_1$ is optimal, his equilibrium payoff is at most $\max\{\frac{1-\delta}{\delta}c, u_1(\underline{a}_1, \overline{a}_2^{\pi}) - c\}$. □

I state a theorem which shows that as long as player 1 is sufficiently long-lived, regardless of the disclosure policy that he commits to and regardless of the equilibrium being played, his payoff when he can renege cannot be significantly greater than his highest equilibrium payoff when he promises to show the null record no matter what. I also show that the short-run players' welfare cannot be significantly greater than that when the long-run player commits to reveal all his actions, i.e., $q_n(m) = 1$ if and only if $n = m$.

Theorem. *Fix any $\pi > 0$, $\widehat{\delta} \in (0,1)$, and (u_1, u_2, c) that satisfies Assumptions 1,2,3 and 4. For every $\varepsilon > 0$, there exists $\delta^* \in (0,1)$ such that for every $\overline{\delta} > \delta^*$, in every equilibrium under every disclosure policy,* *player* 1*'s equilibrium payoff when he can renege is at most*

$$\max\Big\{\frac{1-\delta}{\delta}c,\ u_1(\underline{a}_1,\overline{a}_2^{\pi})-c\Big\}. \tag{C.1}$$

If in addition that $u_2(a_1,a_2)$ *is weakly increasing in* a_1*, then player* 2*'s welfare* U_2 *in any equilibrium under any disclosure policy is no more than* $\pi u_2(a_1^*,a_2^*)+\varepsilon$.

This theorem implies that as long as the short-run players suspect that (i) the long-run player may renege with positive probability and (ii) the reneging long-run player has the option to erase past actions at a low cost, the long-run player can no longer benefit from having his actions monitored relative to the benchmark scenario in which his actions cannot be monitored at all.[4] This conclusion stands in contrast to usual lessons from the theories of repeated games, that a patient player can obtain strictly higher payoffs when his opponents can monitor his actions. Compared to the baseline model where player 1 commits to disclose all past actions to the short-run players, this theorem implies that committing to alternative disclosure policies cannot benefit the short-run players but can benefit the long-run player when his opponents believe that he will honor his promise with probability above some cutoff.

Proof. Using the same argument as in the proof of Theorem 1, I can show that in every equilibrium (i) there exists $t\in\mathbb{N}$ such that when the long-run player can renege, (ii) he takes action a_1^* with positive probability at every history h_*^k with $k<t$, and that he takes action $\underline{a}_1$ with positive probability at every history, and (iii) player 2's action at h_*^k, denoted by β_k, is strictly increasing in k in the sense of FOSD.

Therefore, even conditional on the event that player 1 can renege, his equilibrium payoff is bounded above by $\max\{\frac{(1-\delta)c}{\delta},u_1(\underline{a}_1,\beta_0)-c\}$. This is because $\frac{(1-\delta)c}{\delta}$ is an upper bound on player 1's equilibrium payoff when he does not erase $\underline{a}_1$ with positive probability and $u_1(\underline{a}_1,\beta_0)-c$ is his equilibrium payoff when he takes action $\underline{a}_1$ and then erases it in every period. This payoff upper bound is no more than (C.1) unless (i) $u_1(\underline{a}_1,\beta_0)>c/\delta$ and (ii) there exists an action strictly greater than $\overline{a}_2^{\pi}$ that belongs to the support of β_0.

The first requirement implies that at every h_*^k, when player 1 can renege, he strictly prefers to erase $\underline{a}_1$ after taking it at h_*^k. This implies that the reneging player 1 never reveals to the short-run players that he can renege on his promise. Recall the definition of $\overline{a}_2^{\pi}$. Since β_k is non-decreasing in k, the second requirement implies that for every β_k with $k\geq 0$, there exists an action that is strictly greater than $\overline{a}_2^{\pi}$ that belongs to the support of β_k. This implies that there exists $p>0$ depends only on (u_1,u_2) such that *conditional on* player

[4] The current theorem neither implies nor is implied by Theorem 2 in the main text. This is because Theorem 2 derives a stronger payoff upper bound for the long-run player and a sharper characterization of the short-run players' welfare under full disclosure while the current result derives weaker conclusions but allows for a larger class of disclosure policies.

1 can renege, the expected probability with which he takes action a_1^* is at least p. This is because otherwise, player 2 will have no incentive to take actions strictly greater than $\overline{a}_2^{\pi}$ at h_*^k for every $k \in \mathbb{N}$.

Fix any $\widehat{\delta} \in (0,1)$. Suppose by way of contradiction that for every $\delta^* \in (0,1)$, there exists a survival probability $\overline{\delta} > \delta^*$, a disclosure policy $\mathbf{q}$, and an equilibrium under $(\mathbf{q}, \overline{\delta}, \widehat{\delta})$ such that player 1's payoff when he can renege is strictly greater than (C.1). Let $\widetilde{q} \in \Delta(\mathbb{N})$ denote the distribution over player 1's revealed record conditional on he honors his promise, which is pinned down the disclosure policy $\mathbf{q}$. Since at every history h_*^k with $k < t$, player 1 is indifferent between playing a_1^* and playing $\underline{a}_1$ and then erasing it conditional on he can renege, we know that

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

This implies that t is bounded from above by some linear function of $(1-\delta)^{-1}$.

Since player 2 has an incentive to take some action that is strictly greater than $\overline{a}_2^{\pi}$ at h_*^0 and player 2's action increases in the length of player 1's good record in the sense of FOSD, there exists $x > \pi$ such that player 2's belief assigns probability at least x to a_1^* at every h_*^k with $k \leq t-1$. Importantly, this x depends only on (u_1, u_2) and does not depend on $\widehat{\delta}$, $\overline{\delta}$, and the promised disclosure policy. Recall that μ_k^* is the probability that the history is h_*^k conditional on player 1 being the opportunistic type and that p_k^* is the probability with which the opportunistic type plays a_1^* at h_*^k. Player 2's incentive constraint at history h_*^k implies that

$$\frac{\pi \widetilde{q}(k) + (1-\pi)\mu_k^* p_k^*}{(1-\pi)\mu_k^*(1-p_k^*)} \geq \frac{x}{1-x} \text{ for every } k \in \{1, ..., t-1\}. \tag{C.3}$$

Since $\pi < x$, (C.3) is true for every $k \leq t-1$ only when $\sum_{j=0}^{t-1} \mu_j^* p_j^*$ is uniformly bounded above 0 as $\overline{\delta} \to 1$. Notice that Lemma 3 in the main text applies to all disclosure policies. It implies that $\mu_j^* p_j^* \leq 1 - \overline{\delta}$ for every $j \leq t-1$. Therefore, $\sum_{j=0}^{t-1} \mu_j^* p_j^*$ is bounded above 0 if and only if t is bounded below by a linear function of $(1-\overline{\delta})^{-1}$. For any fixed $\widehat{\delta} \in (0,1)$, there exists $\overline{\delta}$ close to 1 such that the lower bound on t is strictly greater than the upper bound on t implied by the opportunistic type's incentive constraints. This rules out equilibria in which player 1's payoff being strictly greater than (C.1). Similarly, fix any equilibrium and the resulting distribution over player 1's actions, player 2's payoff cannot be greater than their payoff when they can observe player 1's realized pure action before choosing their action. This upper bound cannot be greater than $\pi u_2(a_1^*, a_2^*) + (1-\pi)u_2(\underline{a}_1, \underline{a}_2)$ as $\overline{\delta} \to 1$ since $\Pr(\mathcal{E}_k) \leq 1 - \overline{\delta}$ for every $k \leq t$ and t is at most proportional to $(1-\delta)^{-1}$. $\square$

D 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 $\Omega \equiv \Theta \bigcup \widetilde{A}_1$, 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 $a_1 \in \widetilde{A}_1 \subset A_1$ 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 $\theta \in \Theta$, $u_1(\theta, a_1, a_2)$ 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 $\theta \in \Theta$.

I start from the benchmark without any honest type. Similar to the case with only one opportunistic type, when the cost of erasing actions $c(\theta)$ is lower than some cutoff $\overline{c}(\theta)$ for every $\theta \in \Theta$, player 1 plays $\underline{a}_1$ and player 2 plays $\underline{a}_2$ at every on-path history of every Nash equilibrium, regardless of $\widehat{\delta}$ and $\overline{\delta}$. Formally, for each strategy profile (σ_1, σ_2), let $\mathcal{H}(\sigma_1, \sigma_2|\theta)$ be the set of player 1 histories that occur with positive probability under $(\theta, \sigma_1, \sigma_2)$. Recall that a_1' is the lowest action in A_1 under which $\underline{a}_2$ is *not* a best reply. Let

$$\overline{c}(\theta) \equiv \min_{\beta \in \Delta(A_2)} \Big\{ u_1(\theta, \underline{a}_1, \beta) - u_1(\theta, a_1', \beta) \Big\}. \tag{D.1}$$

Proposition. *Suppose (u_1, u_2) satisfies Assumptions 1, 2, and 3 for every $\theta \in \Theta$, $c(\theta) < \overline{c}(\theta)$ 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

$$\overline{V}_{\theta_1} \equiv \sup_{h \in \mathcal{H}(\sigma_1, \sigma_2|\theta_1)} V_{\theta_1}(h), \tag{D.2}$$

which is type θ_1's highest continuation value in this equilibrium. Suppose by way of contradiction that

$\overline{V}_{\theta_1} = 0$, 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 $\overline{V}_{\theta_1} = 0$.

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

$$0 < \varepsilon < \min\Big\{\frac{\overline{V}_{\theta_1}}{2}, \frac{(1-\delta)(\overline{c}(\theta_1) - c(\theta_1))}{\delta}\Big\}, \tag{D.3}$$

there exists a history $h(1) \in \mathcal{H}(\sigma_1, \sigma_2|\theta_1)$ such that $V_{\theta_1}(h(1)) > \overline{V}_{\theta_1} - \varepsilon$. 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

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

His continuation value for playing any action $a_1 \neq \underline{a}_1$ is at most $(1-\delta)u_1(\theta_1, a_1, \beta(h(1))) + \delta\overline{V}_{\theta_1}$. This upper bound is strictly less than $(1-\delta)\{u_1(\theta_1, \underline{a}_1, \beta(h(1))) - c(\theta_1)\} + \delta(\overline{V}_{\theta_1} - \varepsilon)$. 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,

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

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

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 $a_1'' \succsim a_1'$ 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

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

I show that $\overline{V}_{\theta_2} > 0$. This is because otherwise, $\overline{V}_{\theta_2} = 0 = V_{\theta_2}(h(1), a_1'')$, 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 $a_1''' \succsim a_1'$ 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 2 in the main text by showing that for any level of patience $\widehat{\delta}$ 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 $\pi \in \Delta(\Omega)$, $c(\theta) < \overline{c}(\theta)$ for every $\theta \in \Theta$, and $\widehat{\delta} \in (0,1)$, there exists $\delta^* \in (0,1)$ such that for every $\overline{\delta} > \delta^*$ and $\theta \in \Theta$, type θ's payoff in every equilibrium is no more than $\frac{(1-\delta)c(\theta)}{\delta}$.*

The intuition is similar to that behind Theorem 2 in the main text. As long as there exists one opportunistic type $\theta \in \Theta$ whose payoff is strictly greater than $\frac{(1-\delta)c(\theta)}{\delta}$, 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 2, 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 $(1-\overline{\delta})^{-1}$. 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 $\overline{\delta} \to 1$, 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,

$$u_1(\theta, \underline{a}_1, \beta_k) \leq \frac{c(\theta)}{\delta}. \tag{D.4}$$

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 (D.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 $\overline{V}_\theta$ be type θ's highest continuation value. Fix any ε that satisfies:

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

there exists $t_\theta \in \mathbb{N}$ such that $V_\theta(t_\theta) > \overline{V}_\theta - \varepsilon$. 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 $\overline{V}_{\theta_*}$ 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 $\overline{V}_{\theta_*}$. 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 plays a_1^* with positive probability. Let this type be θ. His continuation value at h_*^k satisfies:

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

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

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

Since $c(\theta)<\overline{c}(\theta)$, 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,

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

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

$$t\leq\sum_{\theta\in\Theta}\frac{\delta\cdot\overline{u}_1(\theta)}{(1-\delta)\Delta(\theta)}. \tag{D.7}$$

Step 3: I show that for every type $\theta\in\Theta$ and integer $k\leq t-1$, 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 $\beta_k\succeq\beta_{k+1}$. 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

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

and

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

Since $c(\theta)<\overline{c}(\theta)$, we have $u_1(\theta,\underline{a}_1,\beta_k)-c(\theta)>u_1(\theta,a_1^*,\beta_k)$. Therefore, (D.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 (D.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 2 in the main text), for every $\theta\in\Theta$, there exists at most one

$\beta^*(\theta) \in \mathcal{B}$ such that inequality (D.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 Online Appendix A implies that:

$$\mu_0^* = (1-\overline{\delta}) + \overline{\delta}\mu_0^*(1-p_0^*), \tag{D.10}$$

$$\mu_k^* = \overline{\delta}\mu_{k-1}^* p_{k-1}^* + \overline{\delta}\mu_k^*(1-p_k^*) \text{ for every } k \in \{1,2,...,t-1\}, \tag{D.11}$$

and

$$\widehat{\mu}_k = (1-\overline{\delta})\overline{\delta}^k \Pi_{j=0}^{k-1} \widehat{p}_j. \tag{D.12}$$

Equation (D.11) implies that

$$\frac{\mu_k^*}{\mu_{k-1}^*} = \frac{\overline{\delta} p_{k-1}^*}{1-\overline{\delta}(1-p_k^*)} \text{ for every } k \in \{1,2,...,t-1\}. \tag{D.13}$$

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

$$\frac{x_k}{1-x_k} = \frac{\pi(1-\overline{\delta})\overline{\delta}^k + \pi^* \mu_k^* p_k^* + \widehat{\pi}\widehat{\mu}_k \widehat{p}_k}{\pi^* \mu_k^*(1-p_k^*) + \widehat{\pi}\widehat{\mu}_k(1-\widehat{p}_k)} = \frac{\pi^*\mu_k^* p_k^* + (1-\overline{\delta})\overline{\delta}^k I_k}{\pi^*\mu_k^*(1-p_k^*) + (1-\overline{\delta})\overline{\delta}^k J_k}, \tag{D.14}$$

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

$$\pi^*\mu_k^*(x_k - p_k^*) = (1-\overline{\delta})\overline{\delta}^k \Big\{ I_k - x_k(I_k + J_k) \Big\} \leq (1-\overline{\delta})\overline{\delta}^k(\pi + \widehat{\pi}). \tag{D.15}$$

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

$$\overline{\mu}_k(x_k - \overline{p}_k) = (1-\overline{\delta})\overline{\delta}^k \frac{\pi + \widehat{\pi}}{\pi^*}, \tag{D.16}$$

$$\overline{\mu}_0 = \frac{1-\overline{\delta}}{1-\overline{\delta}(1-\overline{p}_0)} \text{ and } \overline{\mu}_k = \frac{\overline{\delta}\overline{p}_{k-1}\overline{\mu}_{k-1}}{1-\overline{\delta}(1-\overline{p}_k)}. \tag{D.17}$$

Similar to the proof of Theorem 2 in the main text, one can show that $\overline{p}_0$ is bounded above 0 and $\overline{p}_{k-1} - \overline{p}_k < 1-\overline{\delta}$, which implies that $\overline{p}_t = 0$ if and only if t is bounded below by something proportional to $(1-\overline{\delta})^{-1}$.

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

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

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

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

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 $\{p_0^*, ..., p_k^*\}$ all reach their minimal values $\{\overline{p}_0, ..., \overline{p}_k\}$. The definitions of $\{\overline{p}_k\}_{k=0}^t$ and $\{\overline{\mu}_k\}_{k=0}^t$ then imply that $p_{k+1}^* \geq \overline{p}_{k+1}$. This completes the proof that t is bounded below by something proportional to $(1-\overline{\delta})^{-1}$. As $\overline{\delta} \to 1$, 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 $\frac{(1-\delta)c(\theta)}{\delta}$. $\square$

E Extensions: 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 2, 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 $p\sum_{k=0}^{+\infty}(1-\delta)\delta^k\Big(u_1(a_{1,k},a_{2,k})-c_k\Big)$. The short-run players' prior belief assigns probability $(1-\overline{\delta})\overline{\delta}^k$ to the long-run player's *age in the game* being k and assigns probability $(1-\widetilde{\delta})\widetilde{\delta}^k$ to the honest type having interacted with k short-run players, where

$$\widetilde{\delta} \equiv 1 - \frac{1-\overline{\delta}}{1-\overline{\delta}(1-p)}. \tag{E.1}$$

Using the same method as in the proof of Theorem 2, one can show that for every $c < \overline{c}$, $p \in (0,1)$, $\pi \in (0,1)$, $\widehat{\delta} \in (0,1)$, and $\varepsilon > 0$, there exists $\delta^* \in (0,1)$ such that when $\overline{\delta} > \delta^*$, the opportunistic type's payoff is no more than

$$u(p) \equiv \frac{(1-\delta)(1-\delta+\delta p)}{\delta}c \tag{E.2}$$

in every equilibrium. This payoff converges to his minmax value 0 as $\delta \to 1$. 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 2 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 2 implies that the rate with which player 1's reputation increases is bounded above by something proportional to $1-\widetilde{\delta}$. 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 $(1-\widetilde{\delta})^{-1}$. As $\overline{\delta} \to 1$, expression (E.1) implies that $\widetilde{\delta}$ 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,

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

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

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 $\widetilde{\delta} \to 1$, 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 $\widetilde{\delta}$. 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 $\widetilde{\delta} \to 1$, since the average probability with which the opportunistic type takes action a_1^* vanishes as $\widetilde{\delta} \to 1$. 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 2 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 $\overline{\delta} > \delta^*$ 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:

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

which implies that

$$V_{k+1} - V_k = \frac{1-\delta}{\delta(1-p)}\Big\{u_1(\underline{a}_1, \beta_k) - u_1(a^*, \beta_k) - c(1-p)\Big\}. \tag{E.4}$$

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

Recall the definitions of μ_k^* and p_k^* in the proof of Theorem 2. When the short-run players do not leave reviews with probability p, we have

$$\mu_0^* = (1-\overline{\delta}) + \overline{\delta}\Big\{(1-p_0^*) + p_0^* p\Big\}$$

and

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

or equivalently,

$$\mu_0^* = \frac{1-\overline{\delta}}{1-\overline{\delta}(1-p_0^* + p_0^* p)} \tag{E.5}$$

and

$$\frac{\mu_k^*}{\mu_{k-1}^*} = \frac{\overline{\delta} p_{k-1}^*(1-p)}{1-\overline{\delta}(1-p_k^*+p_k^* p)}. \tag{E.6}$$

Let x_k denote player 2's belief about the probability of a_1^* at history h_*^k. As in the baseline model, we have

$$\frac{\pi}{1-\pi}(1-\overline{\delta})\overline{\delta}^k = \mu_k^* \Big\{ \frac{x_k}{1-x_k}(1-p_k^*) - p_k^* \Big\} = \mu_k^* \frac{x_k - p_k^*}{1-x_k} \tag{E.7}$$

and

$$\frac{\mu_k^*}{\mu_{k-1}^*} = \overline{\delta} \frac{x_{k-1} - p_{k-1}^*}{x_k - p_k^*} \cdot \frac{1-x_k}{1-x_{k-1}} \leq \overline{\delta} \frac{x_k - p_{k-1}^*}{x_k - p_k^*}. \tag{E.8}$$

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

$$\frac{x_k - p_{k-1}^*}{x_k - p_k^*} \geq \frac{\overline{\delta} p_{k-1}^*(1-p)}{1-\overline{\delta}(1-p_k^*+p_k^* p)}, \tag{E.9}$$

which is equivalent to

$$p_{k-1}^* - p_k^* \leq (1-\overline{\delta}) \frac{x - p_{k-1}^*}{x(1-p)} (1 - p_k^*(1-p)). \tag{E.10}$$

This leads to an upper bound on t, which is proportional to $(1-\overline{\delta})^{-1}$.

When $\overline{\delta} \to 1$, 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 the probability of event $\mathcal{E}^k$, defined before the statement of Lemma 3, is $\mu_k^* p_k^*$ and satisfies

$$\mu_k^* p_k^* \leq (1-\overline{\delta})\overline{\delta}^k \frac{1}{1-\overline{\delta} p}. \tag{E.11}$$

This bound coincides with the one in Lemma 3 of the main text when $p = 0$. 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 $\overline{\delta}$ is close to 1 in every equilibrium.