Theoretical Economics 16 (2021), 449–475 1555-7561/20210449

Trust and betrayals: Reputational payoffs and behaviors without commitment

Harry Pei
Department of Economics, Northwestern University

I study a repeated game in which a patient player wants to win the trust of some myopic opponents, but can strictly benefit from betraying them. His benefit from betrayal is strictly positive and is his persistent private information. I characterize every type of patient player's highest equilibrium payoff and construct equilibria that attain this payoff. Since the patient player's Stackelberg action is mixed and motivating the lowest-benefit type to play mixed actions is costly, every type's highest equilibrium payoff is strictly lower than his Stackelberg payoff. In every equilibrium where the patient player approximately attains his highest equilibrium payoff, no type of the patient player plays stationary strategies or completely mixed strategies.

Keywords. Reputation, no commitment type, equilibrium payoff, equilibrium behavior.

JEL classification. C73, D82, D83.

1. Introduction

I examine patient players' returns from good reputations when it is common knowledge that they have strict incentives to betray their opponents' trust. My model features a sequential-move trust game (Figure 1) played repeatedly between a patient seller (player 1) and an infinite sequence of myopic buyers (players 0), arriving one in each period and each plays the game only once. In every period, the seller wants to win a buyer's trust, but has a strict incentive to exert low effort once trust is granted. Her cost of effort is perfectly persistent and is her private information, which I call her type. Each buyer observes all the past outcomes and prefers to trust the seller only when the probability of high effort exceeds some cutoff.[1]

My contribution is to show that when all types of the seller can strictly benefit from betraying buyers' trust and their optimal commitment action is nontrivially mixed, the cost of providing the lowest-cost type incentives to play mixed actions leads to a clean characterization of every type's highest equilibrium payoff. In addition, the lowest-cost

Harry Pei: harrydp@northwestern.edu
I thank Daron Acemoglu, Ricardo Alonso, Guillermo Caruana, Isa Chaves, Mehmet Ekmekci, Jeff Ely, Drew Fudenberg, Olivier Gossner, Bruno Jullien, Aditya Kuvalekar, Elliot Lipnowski, Debraj Ray, Joel Sobel, Bruno Strulovici, Saturo Takahashi, Jean Tirole, Juuso Toikka, Tristan Tomala, Alex Wolitzky, Muhamet Yildiz, and two anonymous referees for helpful comments.

[1]Aside from business transactions, the trust game is also applicable to the study of capital taxation (Phelan 2006), monetary policy (Barro 1986), political economy (Tirole 1996), and so on.

 https://doi.org/10.3982/TE4182

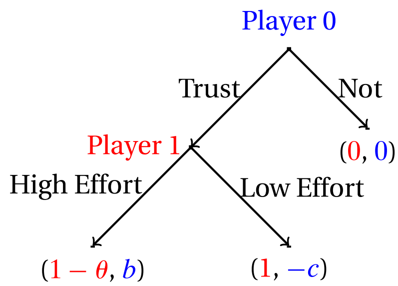

FIGURE 1. The stage game, where $\theta \in (0, 1)$, $b > 0$, $c > 0$.

type's behavior in every seller-optimal equilibrium differs from that of the Stackelberg commitment types in the existing reputation literature who mechanically play the same mixed action in every period.

Theorem 1 characterizes every type of patient seller's highest equilibrium payoff, which equals the product of her Stackelberg payoff and a term called the *purchasing probability*. The latter is the ratio between the lowest-cost type's payoff from exerting high effort and the lowest-cost type's (mixed) Stackelberg payoff.

To understand this characterization, notice that first, the lowest-cost type has no good candidate to imitate in the repeated incomplete information game. As a result, her payoff is no more than her highest equilibrium payoff in a repeated complete information game where her cost is common knowledge. The latter equals her payoff from high effort according to Fudenberg et al. (1990). Intuitively, she needs to exert effort with positive probability so as to receive the buyer's trust and, therefore, exerting effort whenever she is trusted is one of her best replies, from which her payoff is no more than that from high effort.

For the purchasing probability term, consider a high-cost type's payoff from imitating the lowest-cost type in a static setting. Keeping the lowest-cost type's payoff fixed, reducing the probability of high effort increases the high-cost type's payoff. The buyer's incentive to trust yields a lower bound on the probability of high effort, which is binding under the seller's Stackelberg action. The purchasing probability term is the maximal probability that the buyer plays trust when the lowest-cost type plays the Stackelberg action and receives payoff no more than her payoff from high effort.

My result implies that the highest equilibrium payoff for every high-cost type is strictly greater than her payoff under complete information, but is strictly lower than her Stackelberg payoff. The intuition for the first part is similar to models with mixed-strategy commitment types surveyed by Mailath and Samuelson (2006), wherein every high-cost type can free-ride on the effort of the lowest-cost type, and when the lowest-cost type mixes, the high-cost type can benefit from her private information in multiple periods. The second part is novel and is driven by the fact that the lowest-cost type must be incentivized to play mixed actions. Since the seller's Stackelberg action is mixed and buyers cannot observe the seller's mixed actions, the cost of providing mixing incentives prevents every type of seller from attaining her Stackelberg payoff. As the lowest cost in the support of the buyers' prior belief goes to zero, the cost of providing incentives vanishes, and every type's highest equilibrium payoff converges to her Stackelberg payoff.

Theorem 2 shows that if the seller is patient and has at least two types, then in every equilibrium that approximately attains her highest equilibrium payoff, no type mixes between high and low effort at every on-path history. This conclusion extends to a type whose cost of high effort is zero. It implies that every type has a strict incentive at some histories even when she is indifferent between high and low effort in the one-shot game.

For some intuition, suppose toward a contradiction that the lowest-cost type is indifferent in every period. Then exerting low effort in every period is the lowest-cost type's best reply, which implies that every high-cost type exerts low effort for sure in every period. The reputation result in Fudenberg and Levine (1989) implies that when the buyers are facing the high-cost type, they will learn her equilibrium behavior in a bounded number of periods, after which they will stop trusting, leaving the seller with a continuation value of zero. This contradicts the presumption that every high-cost type approximately attains her highest equilibrium payoff.

Related literature The standard reputation models such as Fudenberg and Levine (1989, 1992), Gossner (2011), and Mathevet et al. (2019) assume that with positive probability, the patient player is committed and mechanically plays the same action.[2] They derive results that rule out equilibria in which the patient player receives a low payoff. By contrast, I study a model in which all types are rational. Although my model cannot rule out low-payoff equilibria, the rational types' incentive constraints lead to a sharp characterization of the patient player's highest equilibrium payoff when the Stackelberg action is mixed.[3] Another strand of work (e.g., Board and Meyer-ter-Vehn 2013, Liu and Skrzypacz 2014) characterizes the patient player's behavior in Markov equilibria. By contrast, my Theorem 2 studies the common properties of all seller-optimal equilibria.

In terms of rationalizing commitment types, Weinstein and Yildiz (2016) provide a strategic foundation for nonstationary pure-strategy commitment types in finitely repeated games by constructing a strategic type whose behavior coincides with that of the nonstationary commitment type. My results are complementary to theirs by examining the strategic foundations for mixed-strategy commitment types in infinitely repeated games.

[2] These commitment types are also present in models of credibility such as Sobel (1985) and Benabou and Laroque (1992).

[3] Such a characterization is not obtained in commitment type models for sequential-move stage games. This is because when the buyer chooses N in period t, future buyers cannot observe the seller's action in period t, i.e., the seller's actions are not statistically identified. Theorem 3.1 in Fudenberg and Levine (1992) implies that in commitment-type reputation models with mixed-strategy commitment types, the patient seller's payoff lower bound is zero and her payoff upper bound is her Stackelberg payoff. However, their results do not imply that the payoff upper bound is tight and, therefore, cannot characterize what the patient seller's highest equilibrium payoff is.

My paper also contributes to the rational-type repeated game literature. Aumann and Maschler (1995) study repeated zero-sum games with one-sided private information. Hart (1985) and Shalev (1994) characterize the set of equilibrium payoffs in repeated games with one-sided private information and without discounting. In discounted repeated games, Hörner et al. (2011) and Pęski (2014) characterize the set of equilibrium payoffs when all players' discount factors are arbitrarily close to 1. Cripps and Thomas (2003) show that when the informed player is arbitrarily patient and the uninformed player's discount factor is bounded away from 1, Shalev's characterization remains a necessary condition for being an equilibrium payoff, but it is not sufficient in general. Focusing on games in which the informed player is patient and the uninformed players are impatient, I identify conditions that are both necessary and sufficient for being an equilibrium payoff.

2. Model

Time is discrete, indexed by $t = 0, 1, 2, \dots$. A long-lived seller (she, player 1) with discount factor $\delta \in (0, 1)$ plays the stage-game in Figure 1 against an infinite sequence of myopic buyers (he, player 0), arriving one in each period and each playing the game only once. The buyer moves first, deciding whether to trust the seller (action T) or not (action N). If he chooses N, then both players' payoffs are normalized to 0. If he chooses T, then the seller chooses between high effort (action H) and low effort (action L). If the seller chooses L, then her payoff is 1 and the buyer's payoff is $-c$. If the seller chooses H, then her payoff is $1 - \theta$ and the buyer's payoff is b. In period t, let $y_t \in \{N, H, L\}$ be the stage-game outcome and let $u_1(\theta, y_t)$ be the seller's stage-game payoff.

Both b and c are strictly positive and are common knowledge among players, while $\theta \in \Theta \equiv (0, 1)$ is constant over time and is the seller's private information (or her type). To simplify the exposition, I assume that $\Theta \equiv \{\theta_1, \theta_2\}$ with $0 < \theta_1 < \theta_2 < 1$. All my results generalize to any finite number of types. I briefly discuss this extension in Section 5, and the details can be found in the working paper version (Pei 2020b). The buyers have a full support prior belief $\pi \in \Delta(\Theta)$.

The game's past outcomes can be perfectly monitored. Let $h^t = \{y_s\}_{s=0}^{t-1} \in \mathcal{H}^t$ be a public history with $\mathcal{H} \equiv \bigcup_{t=0}^{+\infty} \mathcal{H}^t$. Let $\sigma_0 : \mathcal{H} \to \Delta\{T, N\}$ be the buyers' strategy. Let $\sigma_1 \equiv (\sigma_\theta)_{\theta \in \Theta}$ be the seller's strategy with the set of public histories $\sigma_\theta : \mathcal{H} \to \Delta\{H, L\}$. Type θ chooses σ_θ to maximize her discounted average payoff $\mathbb{E}^{(\sigma_0, \sigma_\theta)}[\sum_{t=0}^{\infty} (1 - \delta)\delta^t u_1(\theta, y_t)]$, where $\mathbb{E}^{(\sigma_0, \sigma_\theta)}[\cdot]$ is the expectation operator under the probability measure induced by $(\sigma_0, \sigma_\theta)$.

Benchmarks Since every type of seller has a strict incentive to choose L after the buyer plays T, the unique stage-game equilibrium outcome is N from which the seller's payoff is 0. If the seller can optimally commit to an action $\alpha_1 \in \Delta\{H, L\}$ before the buyer moves, then every type's optimal commitment (i.e., Stackelberg action) is to play H with probability $\gamma^* \equiv \frac{c}{b+c}$ and play L with complementary probability. For every $j \in \{1, 2\}$, type θ_j's Stackelberg payoff is $1 - \gamma^* \theta_j$.

In a repeated complete information game in which the buyers know θ, the seller's equilibrium payoff cannot exceed $1-\theta$ according to the folk theorem result in Fudenberg et al. (1990), which is strictly lower than her Stackelberg payoff $1-\gamma^*\theta$. This is because each buyer has an incentive to play T only when he expects the seller to choose H with strictly positive probability. Therefore, playing H at every history in which the buyer plays T is the seller's best reply, from which her discounted average payoff is no more than $1-\theta$.

3. Results

The seller's payoff is $v \equiv (v_1, v_2) \in \mathbb{R}^2$, in which v_j represents type θ_j's discounted average payoff. Let $v^* \equiv (v_1^*, v_2^*)$ with

$$v_j^* \equiv \underbrace{\big(1-\gamma^*\theta_j\big)}_{\text{type } \theta_j\text{'s Stackelberg payoff}} \underbrace{\frac{1-\theta_1}{1-\gamma^*\theta_1}}_{\text{purchasing probability}} \quad \text{for } j \in \{1,2\}. \tag{3.1}$$

Theorem 1 shows that v_j^* is type θ_j patient seller's highest equilibrium payoff and that the highest equilibrium payoffs for all types can be approximately attained in the same equilibrium.

Theorem 1. *For every $\varepsilon > 0$, there exists $\underline{\delta} \in (0,1)$ such that for every $\delta \in (\underline{\delta}, 1)$, the following statements hold:*

(i) There exists no Bayes Nash equilibrium (BNE) in which type θ_1's payoff is strictly more than v_1^*. There exists no BNE in which type θ_2's payoff is more than $v_2^* + \varepsilon$.

(ii) There exists a sequential equilibrium in which the seller's payoff is within an ε neighborhood of (v_1^*, v_2^*).[4]

My formula for the patient seller's highest equilibrium payoff admits an intuitive interpretation. First, $v_1^* = 1-\theta_1$. The intuition is that type θ_1 has the lowest cost in the support of the buyers' prior belief and, therefore, cannot benefit from imitating other types. For the intuition behind v_2^*, consider type θ_2's payoff from imitating type θ_1 in a static model. Let $x \in [0,1]$ be the probability that the buyer chooses T and let $\gamma \in [0,1]$ be the probability that the type θ_1 seller chooses H. The buyer has an incentive to trust only when $\gamma \geq \gamma^*$, and type θ_1's payoff being no more than $1-\theta_1$ implies that

$$x(1-\gamma\theta_1) \leq 1-\theta_1. \tag{3.2}$$

Type θ_2's payoff from imitating type θ_1 is $(1-\gamma\theta_2)x$. Under constraint (3.2) and $\gamma \geq \gamma^*$, we have

$$(1-\gamma\theta_2)x \leq (1-\gamma\theta_2)\frac{1-\theta_1}{1-\gamma\theta_1} \leq \big(1-\gamma^*\theta_2\big)\frac{1-\theta_1}{1-\gamma^*\theta_1} = v_2^*,$$

[4]The notion of sequential equilibrium in infinitely repeated games is introduced by Pęski (2014). I use different solution concepts in the two statements of Theorem 1 to strengthen my result: the necessary conditions for being an equilibrium payoff applies under a weak solution concept (BNE) and the constructed equilibria that attain v^* can survive demanding refinements (sequential equilibrium).

in which the purchasing probability term in (3.1) is the value of x when $\gamma = \gamma^*$ and type θ_1's payoff is $1 - \theta_1$. The substantial part is to show that the payoff upper bound (v_1^*, v_2^*) is attainable when δ is large enough. I provide a constructive proof in Section 4.1 and show that the seller's payoff cannot exceed (v_1^*, v_2^*) in Appendix B.

Theorem 1 suggests that when the Stackelberg action is mixed and all types of the patient player have strict incentives to betray, the incentive problem of the lowest-cost type leads to a tight characterization of all types of the patient player's highest equilibrium payoff. So as to obtain payoff strictly greater than $1 - \theta_2$, type θ_2 needs to free-ride on the effort of type θ_1, which means that at some histories, she plays L for sure while type θ_1 mixes between H and L. Different from models with commitment types, the low-cost type in my model also faces incentive problems and needs to be motivated to play mixed actions. The buyers' inability to observe the seller's mixed actions introduces a cost of providing incentives, which increases in θ_1 and explains why type θ_2's highest equilibrium payoff is bounded below her Stackelberg payoff. As $\theta_1 \to 0$, the cost of providing incentives vanishes, and according to (3.1), v_2^* converges to type θ_2's Stackelberg payoff.

Theorem 2 examines the common properties of the patient seller's behavior in equilibria that approximately attain payoff v^*. For every $\sigma \equiv (\sigma_0, \{\sigma_\theta\}_{\theta \in \Theta})$ and $\theta \in \Theta$, let $\mathcal{H}^{(\sigma_0, \sigma_\theta)}$ be the set of histories that (i) occur with positive probability under $(\sigma_0, \sigma_\theta)$ and (ii) the buyer plays T with positive probability. I say that σ_θ is stationary (with respect to σ) if it prescribes the same action at every history that belongs to $\mathcal{H}^{(\sigma_0, \sigma_\theta)}$, and σ_θ is completely mixed if it prescribes a nontrivially mixed action at every history that belongs to $\mathcal{H}^{(\sigma_0, \sigma_\theta)}$.

Theorem 2. *For every small enough $\varepsilon > 0$, there exists $\underline{\delta} \in (0, 1)$ such that when $\delta > \underline{\delta}$, in every BNE that attains payoff within ε of v^*, no type of seller uses stationary strategies or has a completely mixed best reply.*

The proof is provided in Appendix C. For an intuitive explanation, suppose by way of contradiction that type θ_1 mixes between high and low effort at every history. Then exerting low effort at every history is her best reply. A single-crossing argument in repeated signaling games implies that type θ_2 exerts low effort for sure in every period and behaves as a commitment type who mechanically plays L. The learning argument in Fudenberg and Levine (1989) implies that after observing low effort for a bounded number of periods, the buyers believe that low effort occurs with probability close to 1 in all future periods, after which they will stop trusting the seller, leaving the latter with a payoff close to 0. This contradicts the presumption that the high-cost type's equilibrium payoff is approximately v_2^*.

Suppose toward a contradiction that type θ_2 mixes between high and low effort at every history. Then exerting high effort at every history is her best reply, and according to the same single-crossing argument, type θ_1 exerts high effort for sure in every period. This implies that type θ_2 separates from type θ_1 as soon as she exerts low effort, after which her continuation value is no more than $1 - \theta_2$. This suggests that the high-cost type's equilibrium payoff is no more than $(1 - \delta) + \delta(1 - \theta_2)$, which is strictly lower than v_2^* when δ is close to 1.

The above argument implies that Theorem 2 extends to the case in which $\theta_1 = 0$. My result suggests that in every seller-optimal equilibrium, each type of the seller must have a strict incentive at some on-path histories and her strategy must exhibit nontrivial history dependence, even when she is indifferent between high and low effort in the one-shot game. This stands in contrast to the Stackelberg commitment types in Fudenberg and Levine (1992) and Gossner (2011), who mechanically play the same mixed action in every period.

4. Proof of Theorem 1: Attaining payoff v^*

For every $\gamma \in [\gamma^*, 1]$, let $x(\gamma)$ be the largest $x \in \mathbb{R}_+$ that satisfies inequality (3.2). For every $j \in \{1, 2\}$, let $u_j(\gamma) \equiv x(\gamma)(1-\gamma\theta_j)$ and let $u(\gamma) \equiv (u_1(\gamma), u_2(\gamma))$. I depict $u(\gamma)$ in Figure 2. Since $u(\gamma^*) = v^*$ and $u(\cdot)$ is continuous at γ^*, the second statement of Theorem 1 is implied by Proposition 1.

Proposition 1. *For every $\eta \in (0, 1)$ and $\gamma \in (\gamma^*, 1)$, there exists $\underline{\delta} \in (0, 1)$ such that when $\delta > \underline{\delta}$ and type θ_1 occurs with probability more than η, there exists a sequential equilibrium that attains $u(\gamma)$.*

I summarize the challenges and ideas before presenting my constructive proof. For type θ_2 to obtain payoff close to v_2^*, she needs to have a strict incentive to play L at some histories where the buyer plays T, and the buyer's incentive to play T implies that type θ_1 plays H with probability at least γ^*. At those histories, the seller's action reveals information about her type. When δ is close to 1, the payoff consequence of any given period becomes negligible and leads to the following tension: On one hand, type θ_2 needs to reveal her private information in an unbounded number of periods to attain payoff close to v_2^*; on the other hand, she cannot benefit from her private information in the future if she has already revealed a lot of information.

To overcome this challenge, I construct an equilibrium that keeps track of the seller's continuation value and her *reputation*, defined as the probability that the buyers' belief attaches to θ_1, which is denoted by $\eta(h^t)$.

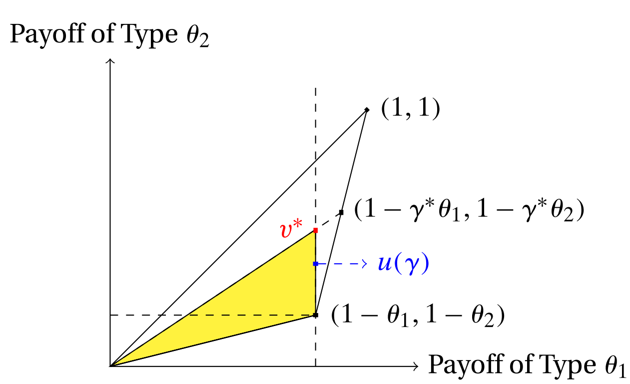

Figure 2. Player 1's highest equilibrium payoff v^*, her equilibrium payoff set (shaded area), and payoff $u(\gamma)$ for some $\gamma \in (\gamma^*, 1)$.

Play starts from an active learning phase in which the buyers play T and both types of seller mix between H and L, with type θ_1 playing H with strictly higher probability. An exception is when $\eta(h^t)$ is close to 1, in which case type θ_1 mixes between H and L while type θ_2 plays L for sure. As a result, the seller's reputation improves after playing H and deteriorates after playing L. This arrangement enables type θ_2 to rebuild her reputation every time she milks it, and reduces her reputation loss when she shirks for sure and free rides on type θ_1.

To provide incentives for type θ_1 to play mixed actions, I introduce an absorbing phase in which learning stops and the outcome in the continuation game is either H or N. This absorbing phase is reached either when the seller's reputation reaches 1 or when she has played L too frequently in the past. Although the seller can flexibly choose her actions in the active learning phase, her action choices affect the calendar time at which play enters the absorbing phase and her continuation value after reaching the absorbing phase.

Intuitively, if the seller plays L with higher frequency in the beginning, then the absorbing phase is reached sooner and her continuation value after reaching the absorbing phase is lower. A more subtle situation arises when the high-cost type plays H too frequently: if her current-period continuation value is strictly greater than $1-\theta_2$ and she plays H with positive probability in equilibrium, then her continuation value increases after she plays H. This raises the concern that type θ_2's continuation value will be too high at some on-path histories so that it cannot be delivered in the continuation game. To address this issue, I show that when H is played with frequency above γ, the seller's reputation reaches 1 in a bounded number of periods, after which her continuation value is a convex combination of $(0,0)$ and $(1-\theta_1, 1-\theta_2)$. The convex weights are chosen such that (i) type θ_1 is indifferent between building a perfect reputation and milking her reputation right before it reaches 1, and (ii) type θ_2 strict prefers to milk her reputation right before it reaches 1 and is indifferent otherwise. This bounds type θ_2's continuation value from above, which guarantees its attainability in the continuation game.

4.1 *Constructive Proof of Proposition 1*

I specify players' actions and the seller's continuation values at on-path histories. At every off-path history, the buyer plays N and all types of the seller play L. I keep track of two state variables:

(i) The probability of type θ_1, denoted by $\eta(h^t) \in [0,1]$, with $\eta(h^0)$ the prior probability of type θ_1.

(ii) The seller's continuation value $v(h^t) \in \mathbb{R}^2$ with $v(h^0) \equiv u(\gamma)$. Lemma A.1 shows by induction that for every $h^t \in \mathcal{H}$, $v(h^t)$ can be written as a convex combination of $v^N \equiv (0,0)$, $v^H \equiv (1-\theta_1, 1-\theta_2)$, and $v^L \equiv (1,1)$, i.e., $v(h^t) = p^N(h^t)v^N + p^H(h^t)v^H + p^L(h^t)v^L$. This implies that keeping track of $v(h^t)$ is equivalent to keeping track of the convex weights $p^N(h^t)$, $p^H(h^t)$, and $p^L(h^t)$.

I partition the set of on-path histories into three classes, depending on the value of $p^L(h^t)$.

- Class 1 Histories: $p^L(h^t) \geq 1-\delta$.
- Class 2 Histories: $p^L(h^t) \in (0, 1-\delta)$.
- Class 3 Histories: $p^L(h^t) = 0$.

Learning about the seller's type happens at Class 1 and 2 histories (i.e., the *active learning phase*) and stops at Class 3 histories (i.e., *the absorbing phase*). Play starts from the active learning phase and reaches the absorbing phase in finite time, after which play remains in that phase forever.

Class 1 histories Play starts from a Class 1 history where $p^L(h^t) \geq 1-\delta$. If h^t belongs to Class 1, then the following statements hold:

- The buyer plays T for sure.
- The buyer's posterior beliefs, $\eta(h^t, H)$ and $\eta(h^t, L)$, are functions of $\eta(h^t)$, which are given by

$$\eta(h^t, H) = \eta^* + \min\{1-\eta^*, (1+\lambda(1-\gamma^*))(\eta(h^t)-\eta^*)\} \tag{4.1}$$

$$\eta(h^t, L) = \eta^* + (1-\lambda\gamma^*)(\eta(h^t)-\eta^*), \tag{4.2}$$

  where $\lambda > 0$ measures the speed of buyers' learning with the requirement specified in (A.2), and $\eta^* \in (\gamma^*\eta(h^0), \eta(h^0))$ is a lower bound on the seller's reputation. Given that $\eta(h^0) > \eta^*$, one can verify by induction that $\eta(h^t) > \eta^*$ for every history h^t that belongs to Class 1 (and also Class 2).

Equations (4.1) and (4.2) pin down both types of seller's actions at h^t. According to Bayes' rule, the probability that type θ_1 plays H at h^t is

$$\frac{\eta(h^t) - \eta(h^t, L)}{\eta(h^t, H) - \eta(h^t, L)} \cdot \frac{\eta(h^t, H)}{\eta(h^t)}, \tag{4.3}$$

and the probability that type θ_2 plays H at h^t is

$$\frac{\eta(h^t) - \eta(h^t, L)}{\eta(h^t, H) - \eta(h^t, L)} \cdot \frac{1-\eta(h^t, H)}{1-\eta(h^t)}. \tag{4.4}$$

Plugging (4.1) and (4.2) into (4.3) and (4.4), every type's action at h^t can be written as a function of $\eta(h^t)$.

Let $p_H(h^t)$ be the unconditional probability of H at h^t. Since the buyers' belief is a martingale, $p_H(h^t)\eta(h^t, H) + (1-p_H(h^t))\eta(h^t, L) = \eta(h^t)$. Equation (4.2) and $\eta(h^t) > \eta^*$ imply that $\eta(h^t, L) \neq \eta(h^t)$ and, therefore,

$$\frac{1-p_H(h^t)}{p_H(h^t)} = \frac{\eta(h^t, H) - \eta(h^t)}{\eta(h^t) - \eta(h^t, L)}.$$

Plugging (4.1) and (4.2) into the above equation, we have

$$\frac{1-p_H(h^t)}{p_H(h^t)} = \frac{\eta(h^t,H)-\eta(h^t)}{\eta(h^t)-\eta(h^t,L)} \leq \frac{1-\gamma^*}{\gamma^*}.$$

This implies that $p_H(h^t) \geq \gamma^*$, so the buyer has an incentive to play T at h^t. The seller's continuation value after playing L at h^t is

$$v(h^t,L) = \frac{p^N(h^t)}{\delta}v^N + \frac{p^L(h^t)-(1-\delta)}{\delta}v^L + \frac{p^H(h^t)}{\delta}v^H. \tag{4.5}$$

If h^t is such that $\eta(h^t,H) < 1$, then the seller's continuation value after playing H at h^t is

$$v(h^t,H) = \frac{p^N(h^t)}{\delta}v^N + \frac{p^L(h^t)}{\delta}v^L + \frac{p^H(h^t)-(1-\delta)}{\delta}v^H. \tag{4.6}$$

If h^t is such that $\eta(h^t,H) = 1$, then the seller's continuation value after playing H at h^t is

$$v(h^t,H) = \frac{v_1(h^t,H)}{1-\theta_1}v^H + \left(1-\frac{v_1(h^t,H)}{1-\theta_1}\right)v^N, \quad \text{with} \tag{4.7}$$

$$v_1(h^t,H) \equiv \frac{v_1(h^t)-(1-\delta)(1-\theta_1)}{\delta} \quad \text{and} \quad v_1(h^t) \text{ is the first entry of } v(h^t). \tag{4.8}$$

If h^t is such that $\eta(h^t,H) < 1$, then (4.5) and (4.6) imply that both types of seller are indifferent between H and L. If h^t is such that $\eta(h^t,H) = 1$, then (4.5), (4.7), and (4.8) imply that type θ_1 is indifferent while type θ_2 strictly prefers L. This verifies both types of the seller's incentive constraints at every Class 1 history h^t.

Class 2 histories For every history h^t that belongs to Class 2, i.e., $p^L(h^t) \in (0, 1-\delta)$, the followingnstatements hold:

- The buyer plays T for sure.
- Type θ_1 plays H for sure. Type θ_2 plays L with probability $\min\{1, \frac{1-\gamma^*}{1-\eta(h^t)}\}$. The probability of L at h^t is $(1-\eta(h^t))\min\{1, \frac{1-\gamma^*}{1-\eta(h^t)}\}$, which is less than $1-\gamma^*$. This implies the buyer's incentive to play T at h^t.

The seller's mixing probabilities imply that $\eta(h^t,L) = 0$ and $\eta(h^t,H) = \min\{1, \frac{\eta(h^t)}{\gamma^*}\}$. The seller's continuation value after playing L at h^t is

$$v(h^t,L) \equiv \frac{Q(h^t)}{\delta}v^H + \frac{\delta-Q(h^t)}{\delta}v^N, \tag{4.9}$$

where

$$Q(h^t) \equiv p^H(h^t) - \frac{1-\delta-p^L(h^t)}{1-\theta_2}.$$

If $\eta(h^t,H) < 1$, then the seller's continuation value after playing H at h^t is given by (4.6). If $\eta(h^t,H) = 1$, then the seller's continuation value after playing H at h^t is given by (4.7).

If $\eta(h^t, H) < 1$, then (4.6) and (4.9) imply that type θ_2 is indifferent at h^t and type θ_1 strictly prefers to play H at h^t. If $\eta(h^t, H) = 1$, then (4.7) and (4.9) imply that type θ_2 strictly prefers to play L at h^t and type θ_1 strictly prefers to play H at h^t. The seller's incentive constraints at h^t are satisfied since type θ_1 is required to play H, while type θ_2 is required to mix only if $\eta(h^t, H) < 1$ and is required to play L if $\eta(h^t, H) = 1$.

Class 3 histories If h^t is such that $p^L(h^t) = 0$, then all types of the seller play the same action at h^t and learning about the seller's type stops. The seller's continuation value at every subsequent history is a convex combination of v^H and v^N. The construction of equilibrium play after reaching any Class 3 history uses Lemma 3.7.2 in Mailath and Samuelson (2006, p. 99; MS).

Lemma 3.7.2 in MS. *For every $\varepsilon > 0$, there exists $\underline{\delta} \in (0, 1)$ such that for every $\delta \in (\underline{\delta}, 1)$ and every $v \in \mathbb{R}^2$ that is a convex combination of $v(1), v(2), \dots, v(k) \in \mathbb{R}^2$, there exists $\{v^s\}_{s=0}^{\infty}$ with $v^s \in \{v(1), \dots, v(k)\}$ such that (i) $v = \sum_{s=0}^{\infty}(1-\delta)\delta^s v^s$ and (ii) $\sum_{s=t}^{\infty}(1-\delta)\delta^{s-l} v^s$ is within ε of v for every $t \in \mathbb{N}$.*

In words, this lemma implies that for any v that is a convex combination of v^H and v^N, and for any large enough discount factor, one can construct a deterministic sequence of v^H and v^N such that the discounted average value of this sequence equals v, and the discounted average value starting from any calendar time is close to v.

For every Class 3 history h^t, let $\{v^s\}_{t=0}^{\infty}$ be a sequence of v^N and v^H that has a discounted average value of $v(h^t)$. The continuation play after h^t is as follows:

- For every $s \in \mathbb{N}$ such that $v^s = v^H$, the buyer plays T and all types of seller play H in period $t + s$.
- For every $s \in \mathbb{N}$ such that $v^s = v^N$, the buyer plays N and all types of seller play L in period $t + s$.

The buyers' incentive constraints at Class 3 histories are trivially satisfied. To verify the seller's incentive constraints, I show in Lemma A.4 of Appendix A that $p^H(h^t)$ is bounded away from 0 for every on-path history h^t. Picking ε in the above lemma to be small enough, one can ensure that the seller's continuation value at every on-path history is strictly bounded away from 0. This implies that when δ is large enough, all types of the seller have strict incentives to comply no matter whether she is asked to play H or L. This is because her continuation value equals 0 if she does not comply and is bounded away from 0 if she complies.

Promise-keeping constraints In the last step, I verify that the continuation play at every on-path history delivers every type of seller her continuation value. I show that under the above strategy profile, play reaches a Class 3 history in finite time, which is implied by Lemmas A.2 and A.3 in Appendix A. Moreover, the seller's continuation value at Class 3 histories can be delivered via a sequence of payoffs consisting of v^H and v^N.

5. Concluding remarks

I conclude by discussing extensions and robustness of my results. Further generalizations can be found in the working paper version (Pei 2020b).

More than two types In the working paper version (Pei 2020b), I extend Theorems 1 and 2 to any finite number of types. Suppose $\Theta \equiv \{\theta_1, \dots, \theta_m\}$ with $0 < \theta_1 < \cdots < \theta_m < 1$ and $m \geq 2$. Let $\pi \in \Delta(\Theta)$ be the buyers' full support prior belief. For every $j \in \{1, 2, \dots, m\}$, let

$$v_j^* \equiv \big(1 - \gamma^* \theta_j\big) \frac{1 - \theta_1}{1 - \gamma^* \theta_1}.$$

When the seller is sufficiently patient, one can show that first, for any $j \in \{1, 2, \dots, m\}$, type θ_j's payoff cannot exceed v_j^* and, moreover, there exists an equilibrium that attains payoff within an ε neighborhood of $(v_1^*, \dots, v_m^*)$. Second, in any equilibrium that approximately attains payoff $(v_1^*, \dots, v_m^*)$, no type of the seller uses stationary strategies or completely mixed strategies.

The proofs are similar to those in the two-type case except for the construction of equilibria that approximately attain $(v_1^*, \dots, v_m^*)$. In particular, one needs to keep track of an additional state variable, which is the highest-cost type in the support of the buyers' posterior belief. This new state variable affects the seller's actions at Class 2 histories (defined in Section 4), at which all types play the same action except for the highest-cost type in the support of the buyers' posterior.[5]

Forward-looking buyer My results are robust when the seller faces a single buyer whose discount factor is strictly positive but close to 0. Let δ_1 be the seller's discount factor and let δ_0 be the buyer's discount factor.

First, the constructed equilibrium that approximately attains v^* remains an equilibrium under any $\delta_0 \in (0, 1)$. This is because at every off-path history, the buyer plays N and all types of the seller play L, in which case the buyer receives his minmax payoff. Hence, the buyer's strategy in the constructed equilibrium maximizes both his stage-game payoff and his continuation value.

Second, I show by the end of Appendix B that the patient seller's payoff cannot significantly exceed v^* when δ_0 is low enough. Intuitively, this is because, first, the seller still needs to play H with positive probability so as to induce the buyer to play T, which implies that the lowest-cost type's payoff cannot exceed $1 - \theta_1$. Moreover, the buyers can learn the seller's behavior in a bounded number of periods, which means that the frequency with which each type of seller plays H cannot be significantly lower than γ^*.

Rationalizing mixed commitment types Theorem 2 implies that mixed-strategy commitment types (such as those that play the Stackelberg action in every period) cannot be rationalized by a payoff type that has a cost of effort that is low or zero. In Pei (2020b), I allow the seller to have arbitrary preferences and show that if a type plays a mixed action at every on-path history, then she must be indifferent between all outcomes in the one-shot game, which means that she faces no cost to supply high quality and receives no benefit from buyers' purchases.

[5]Similar to the construction in the two-type case, all types play the same action at Class 3 histories, and all types except for type θ_1 play the same action at Class 1 histories.

Simultaneous-move stage game My results extend to simultaneous-move stage games. For example, suppose players' payoffs are

$1\backslash 0$	T	N
H	$1-\theta, b$	$-d(\theta), 0$
L	$1, -c$	$0, 0$

,

where $b, c > 0$, $\theta \in \Theta \equiv \{\theta_1, \theta_2\} \subset (0, 1)$ is player 1's persistent private information, and $d(\theta) \geq 0$ is player 1's loss from exerting high effort when player 0 does not trust. In the repeated version of this game, players' past action choices are perfectly monitored and the public history $h^t \equiv \{a_{0,s}, a_{1,s}\}_{s=0}^{t-1}$ consists of both players' past action choices. Other features of the game remain the same as in the baseline model.

Recall the definition of v^* in (3.1). A construction similar to that in Section 4 implies that v^* is approximately attainable when δ is close to 1. Under a supermodularity condition $0 \leq d(\theta_1) - d(\theta_2) \leq \theta_1 - \theta_2$, type θ_j's highest equilibrium payoff is v_j^*, and the conclusion of Theorem 2 also extends, namely, no type of seller uses a stationary strategy or a completely mixed strategy in any equilibrium that approximately attains v^*.

Appendix A: Proof of Proposition 1: Technical details

Recall the definitions of v^H, v^N, and v^L. Let $\gamma \in (\gamma^*, 1)$ and let $u(\gamma)$ be the equilibrium payoff of the seller. Since $\gamma > \gamma^*$, there exists a rational number $\widehat{n}/\widehat{k} \in (\gamma^*, \gamma)$ with $\widehat{n}, \widehat{k} \in \mathbb{N}$. This also suggests the existence of an integer $j \in \mathbb{N}$ such that

$$\frac{\widehat{n}}{\widehat{k}} = \frac{\widehat{n}j}{\widehat{k}j} < \frac{\widehat{n}j}{\widehat{k}j - 1} < \gamma.$$

Let $n \equiv \widehat{n}j$ and $k \equiv \widehat{k}j$. Let $\delta \in (0, 1)$ be large enough such that

$$\frac{\delta + \delta^2 + \cdots + \delta^n}{\delta + \delta^2 + \cdots + \delta^k} < \widetilde{\gamma} < \frac{\delta^{k-n-1}\big(\delta + \delta^2 + \cdots + \delta^n\big)}{\delta + \delta^2 + \cdots + \delta^{k-1}}. \tag{A.1}$$

I introduce two additional requirements on δ later on, namely (A.4) and (A.13). These requirements are compatible with (A.1) since all of them are satisfied when δ is above some cutoff. Let

$$\widetilde{\gamma} \equiv \frac{1}{2}\left(\frac{n}{k} + \frac{n}{k-1}\right) \quad \text{and} \quad \widehat{\gamma} \equiv \frac{1}{2}\left(\frac{n}{k} + \gamma^*\right).$$

By construction, $\gamma^* < \widehat{\gamma} < \frac{n}{k} < \widetilde{\gamma} < \frac{n}{k-1} < \gamma$. Let $\lambda > 0$ be small enough such that

$$\lambda < \frac{1 - \sqrt{\gamma^*}}{\gamma^*} \quad \text{and} \quad \big(1 - \lambda\gamma^*\big)^{1-\widehat{\gamma}}\big(1 + \lambda\big(1 - \gamma^*\big)\big)^{\widehat{\gamma}} > 1. \tag{A.2}$$

The promise-keeping constraint is established via a sequence of lemmas, with proofs provided in Appendices A.1–A.4.

Lemma A.1. *The seller's continuation value at every history is a convex combination of v^N, v^H, and v^L.*

Lemma A.1 implies that one can keep track of the convex weights instead of the continuation payoff. The next lemma establishes a lower bound on the seller's reputation if her frequency of playing H is more than $\widetilde{\gamma}$.

Lemma A.2. *For every $\underline{\eta} \in (\eta^*, 1)$, there exist $T \in \mathbb{N}$ and $\underline{\delta} \in (0,1)$, such that when $\eta(h^r) \geq \underline{\eta}$ and $\delta > \underline{\delta}$, if $h^t \equiv (y_0, \ldots, y_{t-1}) \succ h^r$ and all histories between h^r and h^t belong to Class 1, then*

$$\underbrace{(1-\delta)\sum_{s=r}^{t-1} \delta^{s-r} \mathbf{1}\{y_s = H\}}_{\text{weight of outcome } H \text{ from } r \text{ to } t} \leq \underbrace{\left(1-\delta^T\right)}_{\text{weight of the initial } T \text{ periods}} + \underbrace{(1-\delta)\sum_{s=r}^{t-1} \delta^{s-r} \mathbf{1}\{y_s = L\}}_{\text{weight of outcome } L \text{ from } r \text{ to } t} \cdot \frac{\widetilde{\gamma}}{1-\widetilde{\gamma}}.$$

I apply Lemma A.2 by setting $h^r = h^0$ and $\underline{\eta} = \eta(h^0)$. If h^t and all its predecessors belong to Class 1,

$$(1-\delta)\sum_{s=0}^{t-1} \delta^s \mathbf{1}\{y_s = L\} \leq p^L\left(h^0\right) = \frac{(1-\theta_1)(1-\gamma)}{1-\gamma\theta_1}.$$

Lemma A.2 leads to an upper bound on $(1-\delta)\sum_{s=0}^{t-1} \delta^s \mathbf{1}\{y_s = H\}$, which implies that if δ is large enough, then

$$p^H\left(h^t\right) \geq Y \equiv \frac{1}{2} \underbrace{\left(\gamma - (1-\gamma)\frac{\widetilde{\gamma}}{1-\widetilde{\gamma}}\right)}_{>0} \frac{1-\theta_1}{1-\gamma\theta_1} \tag{A.3}$$

for every h^t such that h^t and all its predecessors belong to Class 1. Lemma A.3 establishes an upper bound on the number of Class 2 histories along every on-path play.

Lemma A.3. *There exist $\underline{\delta} \in (0,1)$ and $M \in \mathbb{N}$, such that when $\delta > \underline{\delta}$, the number of Class 2 histories along every path of equilibrium play is at most M.*

Lemma A.4 establishes a uniform lower bound on $p^H(h^t)$ for every h^t that belongs to Class 1 or Class 2.

Lemma A.4. *There exist $\underline{\delta} \in (0,1)$ and $\underline{Q} > 0$, such that when $\delta > \underline{\delta}$ and for every h^t that belongs to Class 1 and Class 2, $p^H(h^t) \geq \underline{Q}$.*

Lemma A.4 also implies a lower bound on $p^H(h^t)$ if h^t is the first history that belongs to Class 3, i.e., h^t is such that $p^L(h^t) = 0$ and $p^L(h^s) > 0$ for all $h^s \prec h^t$. These lemmas imply that, first, play reaches the absorbing phase with probability 1 on the equilibrium path and, second, the seller's continuation value when play first reaches the absorbing phase is bounded away from 0. When δ is large enough such that

$$(1-\delta) \leq \delta\big(\underline{Q}/2-(1-\delta)\big), \tag{A.4}$$

Lemma 3.7.2 in Mailath and Samuelson (2006) implies that the seller's continuation value can be delivered via a deterministic sequence consisting of v^N and v^H, which implies the promise-keeping constraint.

A.1 *Proof of Lemma A.1*

By definition, $v(h^0)=u(\gamma)$ is a convex combination of v^N, v^H, and v^L. I establish the following claim.

- Suppose h^t is an on-path history and $v(h^t)$ is a convex combination of v^N, v^H, and v^L. Then for every outcome $y_t \in \{N, H, L\}$ that occurs with positive probability at h^t, the seller's continuation value after y_t, given by $v(h^t, y_t)$, is also a convex combination of v^N, v^H, and v^L.

First, consider the case in which h^t belongs to Class 3. Given that $p^L(h^t)=0$ or, equivalently, $v(h^t)$ is a convex combination of v^N and v^H, the only on-path outcomes at h^t are N and H. As a result, the seller's continuation values $v(h^t, N)$ and $v(h^t, H)$ are both convex combinations of v^N and v^H.

Second, consider the case in which h^t belongs to Class 1. The possible outcomes at h^t are H and L. If h^t is such that $\eta(h^t, H)\neq 1$, then according to (4.5) and (4.6), the seller's continuation value remains a convex combination of v^N, v^H, and v^L. If h^t is such that $\eta(h^t, H)=1$, then according to (4.5) and (4.7), the seller's continuation value remains a convex combination of v^N, v^H, and v^L.

Third, consider the case in which h^t belongs to Class 2. The possible outcomes at h^t are H and L. If the seller plays L, then her continuation value is (4.9), which is a convex combination of v^N and v^H. If she plays H, then her continuation value is (4.6) when $\eta(h^t, H)\neq 1$ and is (4.7) when $\eta(h^t, H)=1$. In both cases, $v(h^t, H)$ is a convex combination of v^N, v^L, and v^H. □

A.2 *Proof of Lemma A.2*

Let $X \in \mathbb{N}$ be the smallest integer such that

$$\big(1+\lambda(1-\gamma^*)\big)^X \geq \frac{1-\eta^*}{\eta(h^0)-\eta^*}.$$

Intuitively, the seller's reputation reaches 1 after playing H for X consecutive periods. For every h^t, let $\Delta(h^t)\equiv\eta(h^t)-\eta^*$. For every $t\in\mathbb{N}$, let $N_{L,t}$ and $N_{H,t}$ be the number of periods in which outcomes L and H occur from period 0 to $t-1$, respectively. Choose T such that $T \geq 2k+X$. The proof is done by induction on $N_{L,t}$.

When $N_{L,t}\leq 2(k-n)$, the conclusion holds since $N_{H,t}\geq 2n+X$. Moreover, $\Delta(h^T)$ reaches $1-\eta^*$ before period T, after which play reaches a Class 3 history.

Suppose the conclusion holds for $N_{L,t} \le N$ with $N \ge 2(k-n)$, and suppose toward a contradiction that there exists h^T with $T \ge k+X$ and $N_{L,T} = N+1$, such that every $h^t \preceq h^T$ belongs to Class 1, but

$$(1-\delta)\sum_{t=0}^{T-1}\delta^t \mathbf{1}\{y_t = H\} - \bigl(1-\delta^X\bigr) > (1-\delta)\sum_{t=0}^{T-1}\delta^t \mathbf{1}\{y_t = L\} \cdot \frac{\widetilde{\gamma}}{1-\widetilde{\gamma}}. \tag{A.5}$$

I obtain a contradiction in three steps.

Step 1 I show that for every $s < T$,

$$(1-\delta)\sum_{t=s}^{T-1}\delta^t \mathbf{1}\{y_t = H\} \ge (1-\delta)\sum_{t=s}^{T-1}\delta^t \mathbf{1}\{y_t = L\}\frac{\widetilde{\gamma}}{1-\widetilde{\gamma}}. \tag{A.6}$$

Suppose toward a contradiction that (A.6) fails. This together with (A.5) implies that

$$(1-\delta)\sum_{t=0}^{s-1}\delta^t \mathbf{1}\{y_t = H\} - \bigl(1-\delta^X\bigr) > (1-\delta)\sum_{t=0}^{s-1}\delta^t \mathbf{1}\{y_t = L\}\frac{\widetilde{\gamma}}{1-\widetilde{\gamma}} \tag{A.7}$$

and

$$(1-\delta)\sum_{t=s}^{T-1}\delta^t \mathbf{1}\{y_t = L\} > 0. \tag{A.8}$$

According to (A.8), $N_{L,s} < N_{L,T}$. Since $N_{L,T} = N+1$, we have $N_{L,s} \le N$. Applying the induction hypothesis and (A.7), we know that play reaches a Class 3 history before h^s. This leads to a contradiction.

Step 2 I show that for every k consecutive periods, the number of outcome H in sequence $\{y_r, y_{r+1}, \dots, y_{r+k-1}\}$ is at least $n+1$. According to (A.6) shown in the previous step, outcome H occurs at least $n+1$ times in the last k periods, namely, in the set $\{y_{T-k+1}, \dots, y_T\}$.

Suppose toward a contradiction that there exists k consecutive periods in which outcome H occurs no more than n times. Then the preceding conclusion that outcome (T, H) occurs at least $n+1$ times in the last k periods implies that there exists k consecutive periods $\{y_r, \dots, y_{r+k-1}\}$ in which outcome H occurs exactly n times and outcome L occurs exactly $k-n$ times. According to (A.2), we have

$$(1-\delta)\sum_{t=r}^{r+k-1}\delta^t \mathbf{1}\{y_t = H\} < (1-\delta)\sum_{t=r}^{r+k-1}\delta^t \mathbf{1}\{y_t = L\}\frac{\widetilde{\gamma}}{1-\widetilde{\gamma}}, \tag{A.9}$$

but then

$$\Delta\bigl(h^{r+k}\bigr) > \Delta\bigl(h^{r+1}\bigr). \tag{A.10}$$

Next, let us consider the new sequence of outcomes with length $T-k$,

$$\widetilde{h}^{T-k} \equiv \{\widetilde{y}_0, \widetilde{y}_1, \dots, \widetilde{y}_{T-k-1}\} \equiv \{y_0, y_1, \dots, y_{r-1}, y_{r+k}, \dots, y_{T-1}\},$$

which is obtained by removing $\{y_r, \dots, y_{r+k-1}\}$ from the original sequence and front-loading the subsequent play $\{y_{r+k}, \dots, y_{T-1}\}$. The number of outcome L in this new sequence is at most $N+1-(n-k)$, which is no more than N. According to the conclusion in Step 1,

$$(1-\delta)\sum_{t=r+k}^{T-1}\delta^t\mathbf{1}\{y_t=H\} > (1-\delta)\sum_{t=r+k}^{T-1}\delta^t\mathbf{1}\{y_t=L\}\frac{\widetilde{\gamma}}{1-\widetilde{\gamma}}.$$

This together with (A.9) and (A.5) implies that

$$(1-\delta)\sum_{t=0}^{T-k-1}\delta^t\mathbf{1}\{\widetilde{y}_t=H\} - \big(1-\delta^X\big) > (1-\delta)\sum_{t=0}^{T-k-1}\delta^t\mathbf{1}\{\widetilde{y}_t=L\}\frac{\widetilde{\gamma}}{1-\widetilde{\gamma}}.$$

According to the induction hypothesis, play reaches a Class 3 history before period $T-k$ if the seller plays according to $\{\widetilde{y}_0, \widetilde{y}_1, \dots, \widetilde{y}_{T-k-1}\}$.

(i) Suppose $\widetilde{h}^{T-k}$ reaches a Class 3 history before period r. Then play reaches a Class 3 history before period r according to the original sequence.

(ii) Suppose $\widetilde{h}^{T-k}$ reaches a Class 3 history in period s, with $s>t$. Then $\Delta(\widetilde{h}^s) \leq \Delta(h^{s+k})$ according to (A.10). This implies that play reaches a Class 3 history in period $s+k$ under the original sequence.

This contradicts the hypothesis that play has never reached a Class 3 history before h^T.

Step 3 For every history $h^T \equiv \{y_0, y_1, \dots, y_{T-1}\} \in \{H, L\}^T$ and $t \in \{1, \dots, T-1\}$, define the operator $\Omega_t : \{H, L\}^T \to \{H, L\}^T$ as

$$\Omega_t\big(h^T\big) = (y_0, \dots, y_{t-2}, y_t, y_{t-1}, y_{t+1}, \dots, y_{T-1});$$

in other words, swap the order between y_{t-1} and y_t. Recall the belief updating formula in Class 1 histories and let

$$\mathcal{H}^{T,*} \equiv \big\{h^T \big| \Delta\big(h^t\big) < 1-\eta^* \text{ for all } h^t \prec h^T\big\}.$$

If $h^T \in \mathcal{H}^{T,*}$, then $\Omega_t(h^T) \in \mathcal{H}^{T,*}$ unless

- $y_{t-1}=L$, $y_t=H$, and $(1+\lambda(1-\gamma^*))\Delta(h^{t-1}) \geq 1-\eta^*$.

Next, I show that the above situation cannot occur except in the last k periods. Suppose toward a contradiction that there exists $t \leq T-k$ such that $h^T \in \mathcal{H}^{T,*}$ but $\Omega_t(h^T) \notin \mathcal{H}^{T,*}$. Based on the conclusion in Step 2, outcome H occurs at least $n+1$ times in the sequence $\{y_t, \dots, y_{t+k-1}\}$. Consider another sequence $\{y_{t-1}, \dots, y_{t+k-1}\}$, in which outcome H occurs at least $n+1$ times and outcome L occurs at most $k-n$ times. This implies that

$$\begin{aligned}\Delta\big(h^{t+k}\big) &\geq \Delta\big(h^{t-1}\big)\big(1+\lambda\big(1-\gamma^*\big)\big)^{n+1}\big(1-\lambda\gamma^*\big)^{k-n}\\ &= \Delta\big(h^{t-1}\big)\underbrace{\big(1+\lambda\big(1-\gamma^*\big)\big)^{n}\big(1-\lambda\gamma^*\big)^{k-n}}_{\geq 1}\big(1+\lambda\big(1-\gamma^*\big)\big)\end{aligned}$$

$$\geq \Delta(h^{t-1})\big(1+\lambda(1-\gamma^*)\big)$$
$$\geq 1-\eta^*, \tag{A.11}$$

where the second inequality follows from $n/k > \widehat{\gamma}$ and the third inequality follows from the hypothesis that $\Omega_t(h^T) \notin \mathcal{H}^{T,*}$. Inequality (A.11) implies that play reaches the high phase before period $t+k \leq T$. This contradicts the hypothesis that $h^T \in \mathcal{H}^{T,*}$.

To summarize, for every $t \leq T-k$, if $h^T \in \mathcal{H}^{T,*}$, then $\Omega_t(h^T) \in \mathcal{H}^{T,*}$. For every $t > T-k$, if $h^T \in \mathcal{H}^{T,*}$, then $\Omega_t(h^T) \in \mathcal{H}^{T,*}$ unless $y_{t-1} = L$ and $y_t = H$. Therefore, one can freely front-load outcome H from period 0 to $T-k-1$ and obtain the revised sequence

$$\{H, H, \dots, H, L, \dots, L, y_{T-k}, \dots, y_{T-1}\}, \tag{A.12}$$

which meets the following two requirements: First, the revised sequence (A.12) still belongs to set $\mathcal{H}^{T,*}$; second, the sequence in (A.12) satisfies (A.5).

According to the conclusion in Step 2, the number of outcome L from period 0 to $T-k-1$ cannot exceed $k-n-1$, and the number of outcome L from period $T-k$ to $T-1$ cannot exceed $k-n-1$. This is because otherwise, there exists a sequence of length k that has at most n periods of outcome H, contradicting the conditions that the sequence of outcomes in (A.12) must satisfy. Therefore, the number of L in this sequence is at most $2(k-n-1)$. This contradicts the induction hypothesis that the number of L exceeds $2(k-n)$. □

A.3 *Proof of Lemma A.3*

Let $N \equiv \lceil \frac{1}{1-\gamma} \rceil$ and recall the integer T in the statement of Lemma A.2. In addition to the requirements on δ mentioned earlier, I also require δ to satisfy

$$\delta^{T+1}\big(1+\delta+\dots+\delta^N\big) > N \quad \text{and} \quad 2\delta^{T+N+2} > 1. \tag{A.13}$$

These are compatible given that all of them require δ to be sufficiently large. The rest of the proof consists of two steps.

In the first step, I show that after the seller plays H at h^t, it takes at most $T+N$ periods for play to reach a history that belongs to either Class 2 or Class 3. According to the continuation value at (h^t, H), we have

$$p^L\big(h^t, H\big) = \frac{p^L(h^t)}{\delta} < \frac{1-\delta}{\delta}. \tag{A.14}$$

The last inequality comes from h^t belonging to Class 2, so that $p^L(h^t) < 1-\delta$ by definition. According to Lemma A.2, for every Class 1 history h^s such that $h^s \succ (h^t, H)$ and all histories between (h^t, H) and h^s belong to Class 1,

$$(1-\delta)\sum_{r=t+1}^{s} \delta^{r-(t+1)}\mathbf{1}\{y_r = H\} \leq \big(1-\delta^T\big) + (1-\delta)\sum_{r=t+1}^{s} \delta^{r-(t+1)}\mathbf{1}\{y_r = L\}\frac{\widetilde{\gamma}}{1-\widetilde{\gamma}}.$$

Moreover, (A.14) and the requirement that all histories between (h^t, H) and h^s belong to Class 1 imply that

$$(1-\delta)\sum_{r=t+1}^{s}\delta^{r-(t+1)}\mathbf{1}\{y_r=L\}<\frac{1-\delta}{\delta}. \tag{A.15}$$

Given that only outcomes L and H occur with positive probability at Class 1 and Class 2 histories,

$$\begin{aligned}1-\delta^{s-(t+1)} &= (1-\delta)\sum_{r=t+1}^{s}\delta^{r-(t+1)}\mathbf{1}\{y_r=L\}+(1-\delta)\sum_{r=t+1}^{s}\delta^{r-(t+1)}\mathbf{1}\{y_r=H\}\\ &\leq \left(1-\delta^T\right)+\frac{1-\delta}{\delta}+\frac{1-\delta}{\delta}\frac{\widetilde{\gamma}}{1-\widetilde{\gamma}}\leq \left(1-\delta^T\right)+\frac{1-\delta}{\delta}\frac{1}{1-\widetilde{\gamma}}\\ &\leq \left(1-\delta^T\right)+\frac{1-\delta}{\delta}\frac{1}{1-\gamma}.\end{aligned}$$

To show that $s-(t+1)\leq T+N$, suppose toward a contradiction that $s-(t+1)\geq T+N+1$. Then

$$\left(1-\delta^T\right)+\frac{1-\delta}{\delta}N\geq \left(1-\delta^T\right)+\frac{1-\delta}{\delta}\frac{1}{1-\gamma}\geq 1-\delta^{s-(t+1)}\geq 1-\delta^{T+N+1},$$

which yields

$$\frac{1-\delta}{\delta}N\geq \delta^T\left(1-\delta^{N+1}\right).$$

Dividing both sides by $\frac{1-\delta}{\delta}$, we have

$$N\geq \delta^{T+1}\left(1+\delta+\cdots+\delta^N\right),$$

which contradicts the first inequality of (A.13). The above contradiction implies that $s-(t+1)\leq T+N$.

In the second step, I focus on every history h^s that has the following two features:

(i) h^s belongs to Class 2

(ii) $h^s \succeq (h^t, H)$ and all histories between (h^t, H) and h^s, excluding h^s, belong to Class 1.

I show that there exists at most one period from (h^t, H) to h^s such that the stage-game outcome is (T, L). Suppose toward a contradiction that there exist two or more such periods. Then

$$(1-\delta)\sum_{r=t+1}^{s}\delta^{r-(t+1)}\mathbf{1}\{y_r=L\}\geq 2(1-\delta)\delta^{T+N+1}.$$

The last inequality comes from the previous conclusion that $s-(t+1)\leq T+N$. This is because h^s belongs to Class 2 and h^{s-1} belongs to Class 1, and, therefore, $(s-1)-(t+$

$1) \le T+N$ or, equivalently, $s-(t+1) \le T+N+1$. According to (A.15),

$$2(1-\delta)\delta^{T+N+1} < (1-\delta)\sum_{r=t+1}^{s}\delta^{r-(t+1)}\mathbf{1}\{y_r = L\} < \frac{1-\delta}{\delta}.$$

The above inequality contradicts the second inequality of (A.13) that $2\delta^{T+N+2} > 1$.

Letting h^t be the first time play reaches a history that belongs to Class 2 with $\overline{\theta}(h^t) = \theta_2$, we have $\eta(h^t, H) \ge \frac{\eta^*}{\gamma^*} \ge \eta(h^0)$. Let h^s be the next history that belongs to Class 2 with $h^s \succ (h^t, H)$. Since we have shown that outcome L occurs at most once between (h^t, H) and h^s, we know that

$$\eta(h^s, H) = \min\left\{1, \frac{\eta(h^s)}{\gamma^*}\right\} \ge \min\left\{1, \frac{\eta(h^t, H)}{\gamma^*}(1-\lambda\gamma^*)\right\}.$$

Therefore, conditional on (h^s, H) not being a Class 3 history, the buyer's belief at (h^s, H) attaches probability at least

$$\eta(h^s, H) \ge \eta(h^t, H)\frac{1-\lambda\gamma^*}{\gamma^*} \ge \eta(h^t, H)\sqrt{\frac{1}{\gamma^*}}$$

to type θ_1, where the last inequality comes from $\lambda \in (0, \frac{1-\sqrt{\gamma^*}}{\gamma^*})$. Let

$$M \equiv \frac{\log(1/\pi_1)}{\log\sqrt{\frac{1}{\gamma^*}}} + 1.$$

Since $\eta(h^t, H) \ge \pi_1$ for the first history that belongs to Class 2, there can be at most M Class 2 histories with θ_2 being the highest-cost type along every path of play. This is because otherwise the buyer's posterior belief attaches probability greater than

$$\pi_1\left(\frac{1}{\sqrt{\gamma^*}}\right)^M > 1$$

at the $M+1$th Class 2 history. This leads to a contradiction. □

A.4 *Proof of Lemma A.4*

Let h^t be a Class 2 history such that no predecessor of h^t belongs to Class 2. According to (A.3), $p^H(h^{t-1}) \ge Y$, which implies that $p^H(h^t) \ge Y-(1-\delta)$. As a result

$$Q(h^t) = p^H(h^t) - \frac{1-\delta-p^L(h^t)}{1-\theta_2} \ge Y-(1-\delta)\left(1+\frac{1}{1-\theta_2}\right) > 0.$$

If play remains at a Class 1 or Class 2 history after h^t, then the seller must be playing H at h^t, after which

$$p^H(h^t, H) \ge p^H(h^t) - (1-\delta) \ge Y-2(1-\delta) \quad \text{and} \quad p^L(h^t, H) \le \frac{1-\delta}{\delta}.$$

Since $\eta(h^t, H) \geq \eta(h^0)$, one can then apply Lemma A.2 again, which implies that at every Class 1 history h^s such that only one predecessor of h^s belongs to Class 2, we have

$$p^H(h^s) \geq Z \equiv Y - 2(1-\delta) - \frac{1-\delta}{\delta}\frac{\widetilde{\gamma}}{1-\widetilde{\gamma}} - (1-\delta^T).$$

When δ is large enough, $Z \geq Y/2$. One can then show that for every Class 2 history h^s such that there is only one strict predecessor history that belongs to Class 2,

$$Q(h^s) = p^H(h^s) - \frac{1-\delta-p^L(h^s)}{1-\overline{\theta}(h^s)} \geq Z - (1-\delta)\left(1+\frac{1}{1-\theta_m}\right) > 0.$$

Iteratively apply this process. Since

(i) the number of Class 2 histories along every path of play is at most M (Lemma A.3)

(ii) for every Class 2 history h^t, $p^L(h^t, H) = \frac{1-\delta}{\delta}$ and $\eta(h^t, H) \geq \eta(h^0)$,

there exist $\underline{\delta} \in (0,1)$ and $\underline{Q} > 0$ such that when $\delta > \underline{\delta}$, $p^H(h^t) \geq \underline{Q}$ for every Class 1 or Class 2 history h^t. □

Appendix B: Proof of Theorem 1: Patient player's payoff cannot exceed v^*

My proof consists of three lemmas. Lemma B.1 relates v_j^* to the value of a constrained optimization problem.

Lemma B.1. *For every $j \in \{1,2\}$, the value of the constrained optimization problem*

$$\max_{\alpha\in\Delta\{N,H,L\}} \Big\{(1-\theta_j)\underbrace{\alpha(H)}_{\text{probability of outcome } H} + \underbrace{\alpha(L)}_{\text{probability of outcome } L}\Big\} \tag{B.1}$$

is v_j^* subject to

$$(1-\theta_1)\alpha(H) + \alpha(L) \leq 1-\theta_1 \tag{B.2}$$

and

$$\alpha(H) \geq \frac{\gamma^*}{1-\gamma^*}\alpha(L). \tag{B.3}$$

Proof. Constraint (B.3) implies that

$$(1-\theta_1)\alpha(H) + \alpha(L) \leq (1-\theta_1)\alpha(H) + \frac{1-\gamma^*}{\gamma^*}\alpha(H) = \left(\frac{1}{\gamma^*} - \theta_1\right)\alpha(H)$$

or, equivalently,

$$\alpha(H) \geq \frac{\gamma^*}{1-\gamma^*\theta_1}\big((1-\theta_1)\alpha(H) + \alpha(L)\big).$$

Rewrite the objective function (B.1) as

$$\begin{aligned}(1-\theta_j)\alpha(H)+\alpha(L) &= (1-\theta_1)\alpha(H)+\alpha(L)-(\theta_j-\theta_1)\alpha(H)\\ &\le \left(1-(\theta_j-\theta_1)\frac{\gamma^*}{1-\gamma^*\theta_1}\right)\big((1-\theta_1)\alpha(H)+\alpha(L)\big)\\ &= \frac{1-\gamma^*\theta_j}{1-\gamma^*\theta_1}\big((1-\theta_1)\alpha(H)+\alpha(L)\big).\end{aligned}$$

Constraint (B.2) implies that

$$(1-\theta_j)\alpha(H)+\alpha(L)\le\frac{1-\gamma^*\theta_j}{1-\gamma^*\theta_1}\big((1-\theta_1)\alpha(H)+\alpha(L)\big)\le\big(1-\gamma^*\theta_j\big)\frac{1-\theta_1}{1-\gamma^*\theta_1}=v_j^*.$$

The above upper bound is attained by the following distribution over action profiles that satisfies (B.2) and (B.3):

$$\alpha(H)=\frac{(1-\theta_1)\gamma^*}{1-\gamma^*\theta_1},\qquad \alpha(L)=\frac{(1-\theta_1)\big(1-\gamma^*\big)}{1-\gamma^*\theta_1},\quad\text{and}\quad \alpha(N)=\frac{\theta_1\big(1-\gamma^*\big)}{1-\gamma^*\theta_1}.$$ □

Next, I map the choice variable in the optimization problem α into the repeated incomplete information game. For given Bayes Nash equilibrium $\sigma\equiv(\sigma_0,(\sigma_\theta)_{\theta\in\Theta})$, type θ_j's equilibrium payoff in the repeated game equals her expected payoff in a one-shot game under outcome distribution $\alpha^j\in\Delta\{N,H,L\}$, where

$$\alpha^j(y)\equiv\mathbb{E}^{(\sigma_0,\sigma_{\theta_j})}\left[\sum_{t=0}^{\infty}(1-\delta)\delta^t\mathbf{1}\{y_t=y\}\right]\quad\text{for every } y\in\{N,H,L\}.$$

Replacing α with α^j, (B.1) is type θ_j's equilibrium payoff in the repeated game. According to Lemma B.1, the necessity of (B.2) and (B.3) implies that type θ_j's equilibrium payoff cannot exceed v_j^* for every $j\in\{1,2\}$.

Lemma B.2 and Lemma B.3 establish the necessity of (B.2) and (B.3), respectively. For every strategy profile σ, let $\mathcal{H}^\sigma$ be the set of histories that occur with positive probability under σ. For every $h^t\in\mathcal{H}^\sigma$, let $\Theta^\sigma(h^t)\subset\Theta$ be the support of buyers' belief at h^t.

Lemma B.2. *For every prior belief π, including those that do not have full support, if type θ_i is the lowest-cost type in the support of this prior belief, then type θ_i's equilibrium payoff is no more than $1-\theta_i$ in every BNE.*

Proof. Rank the seller's actions according to $H\succ L$. Given strategy profile σ and history $h^t\in\mathcal{H}^\sigma$, let

$$\overline{a}_1^\sigma\big(h^t\big)\equiv\max\left\{\bigcup_{\theta\in\Theta^\sigma(h^t)}\operatorname{supp}\big(\sigma_\theta\big(h^t\big)\big)\right\}$$

be the highest action played by the seller with positive probability at h^t. By definition, for every BNE σ and $h^t\in\mathcal{H}^\sigma$, if $\sigma_0(h^t)$ assigns positive probability to T, then $\overline{a}_1^\sigma(h^t)=H$.

Step 1 I show that when $|\Theta^\sigma(h^0)|=1$, the only type in the support of the buyer's prior belief, denoted by θ_i, receives payoff no more than $1-\theta_i$. This also implies that for every equilibrium σ and for every $h^t \in \mathcal{H}^\sigma$, if $\Theta^\sigma(h^t) = \{\theta_i\}$ for some $\theta_i \in \Theta$, then type θ_i's continuation payoff at h^t cannot exceed $1-\theta_i$. This conclusion does not follow from the result in Fudenberg et al. (1990) since the solution concept is BNE, and a type that occurs with zero probability is not equivalent to a type that is excluded from the type space.

This is because $\Theta^\sigma(h^0) = \{\theta_i\}$ implies that $\Theta^\sigma(h^t) = \{\theta_i\}$ for every $h^t \in \mathcal{H}^\sigma$. Therefore, $\overline{a}_1^\sigma(h^t)$ is played by type θ_i with positive probability at every $h^t \in \mathcal{H}^\sigma$. Given type θ_i's equilibrium strategy σ_{θ_i}, the strategy $\widetilde{\sigma}_{\theta_i} : \mathcal{H} \to \Delta(A_1)$, defined as

$$\widetilde{\sigma}_{\theta_i}(h^t) \equiv \begin{cases} \overline{a}_1^\sigma(h^t) & \text{if } h^t \in \mathcal{H}^\sigma \\ \sigma_{\theta_i}(h^t) & \text{otherwise,} \end{cases}$$

also best replies against the buyers' equilibrium strategy σ_0, from which type θ_i receives her equilibrium payoff. If type θ_i plays according to $\widetilde{\sigma}_{\theta_i}$ and the buyers play according to σ_0, then the outcome at every history in $\mathcal{H}^\sigma$ is either H or N. Therefore, type θ_i's stage-game payoff at every history in $\mathcal{H}^\sigma$ cannot exceed $1-\theta_i$, and her discounted average payoff cannot exceed $1-\theta_i$.

Step 2 I show that type θ_1's payoff is no more than $1-\theta_1$ when there are two types in the support of buyers' prior belief. I define $\overline{\mathcal{H}}_t^\sigma$ for every $t \in \mathbb{N}$ recursively. Let $\overline{\mathcal{H}}_0^\sigma \equiv \{h^0\}$. Given the definition of $\overline{\mathcal{H}}_t^\sigma$, let

$$\overline{\mathcal{H}}_{t+1}^\sigma \equiv \{h^{t+1} \in \mathcal{H}^\sigma \mid \exists h^t \in \overline{\mathcal{H}}_t^\sigma \text{ s.t. } h^{t+1} \succ h^t \text{ and either}$$
$$h^{t+1} = (h^t, N) \text{ or } h^{t+1} = (h^t, (T, \overline{a}_1^\sigma(h^t)))\}.$$

Intuitively, $\overline{\mathcal{H}}_{t+1}^\sigma$ is the set of period $t+1$ on-path histories such that the seller has played her highest on-path action from period 0 to period t. Let $\overline{\mathcal{H}}^\sigma \equiv \cup_{t=0}^\infty \overline{\mathcal{H}}_t^\sigma$. Given type θ_1's equilibrium strategy σ_{θ_1}, let $\widehat{\sigma}_{\theta_1} : \mathcal{H} \to \Delta\{H, L\}$ be defined as

$$\widehat{\sigma}_{\theta_1}(h^t) \equiv \begin{cases} \overline{a}_1^\sigma(h^t) & \text{if } h^t \in \mathcal{H}^\sigma \text{ and } \overline{a}_1^\sigma(h^t) \in \operatorname{supp}(\sigma_{\theta_1}(h^t)) \\ \sigma_{\theta_1}(h^t) & \text{otherwise.} \end{cases}$$

By construction, $\widehat{\sigma}_{\theta_1}$ is type θ_1's best reply against σ_0. Let $\mathcal{H}^{(\sigma_0, \widehat{\sigma}_{\theta_1})}$ be the set of histories that occur with positive probability under $(\sigma_0, \widehat{\sigma}_{\theta_1})$. Let

$$\overline{\mathcal{H}}^{\sigma,\theta_1} \equiv \{h^t \in \overline{\mathcal{H}}^\sigma \mid \theta_1 \in \Theta^\sigma(h^t) \text{ and } \overline{a}_1^\sigma(h^t) \notin \operatorname{supp}(\sigma_{\theta_1}(h^t))\}.$$

Consider type θ_1's payoff if he plays $\widehat{\sigma}_{\theta_1}$ and the buyers play σ_0. For any given $h^t \in \mathcal{H}^{(\sigma_0, \widehat{\sigma}_{\theta_1})}$, the following situations occur:

(i) If no $h^s \preceq h^t$ such that $h^s \in \overline{\mathcal{H}}^{\sigma,\theta_i}$ exists, then type θ_1's stage-game payoff at h^t and at all histories preceding h^t is no more than $1-\theta_1$.

(ii) If $h^s \preceq h^t$ such that $h^s \in \overline{\mathcal{H}}^{\sigma,\theta_1}$ exists, then I show below that type θ_1's continuation payoff at h^s is no more than $1-\theta_1$.

First, since $h^s \in \overline{\mathcal{H}}^{\sigma,\theta_1}$, after the buyer observes $\overline{a}_1^{\sigma}(h^s)$ at h^s, type θ_1 is no longer in the support of her posterior belief and, according to Step 1, type θ_2's continuation value at h^s is no more than $1-\theta_2$. As a result, type θ_2's continuation payoff by deviating to $\widehat{\sigma}_{\theta_1}$ starting from h^s is no more than $1-\theta_2$. Since $\theta_1 < \theta_2$, and the maximal difference between type θ_1 and θ_2's stage-game payoff is $\theta_2 - \theta_1$, we know that type θ_1's continuation value at h^s by playing $\widehat{\sigma}_{\theta_1}$ is no more than $1-\theta_1$.

The two parts together imply that type θ_1's discounted average payoff in period 0 is no more than $1-\theta_1$. □

Lemma B.3. *For every $\varepsilon > 0$, there exists $\underline{\delta} \in (0,1)$, such that in every BNE where $\delta > \underline{\delta}$,*

$$\frac{\alpha^j(H)}{\alpha^j(L)} = \frac{\mathbb{E}^{(\sigma_0,\sigma_{\theta_j})}\left[\sum_{t=0}^{\infty}(1-\delta)\delta^t \mathbf{1}\{y_t = H\}\right]}{\mathbb{E}^{(\sigma_0,\sigma_{\theta_j})}\left[\sum_{t=0}^{\infty}(1-\delta)\delta^t \mathbf{1}\{y_t = L\}\right]} \geq \frac{\gamma^* - \varepsilon}{1-\gamma^*+\varepsilon} \quad \textit{for every } j \in \{1,2\}.$$

Proof. Under the probability measure over $\mathcal{H}$ induced by $(\sigma_0, \sigma_{\theta_j})$, let $X^{(\sigma_0,\sigma_{\theta_j})}$ be the occupation measure of outcome (T, H) and let $Y^{(\sigma_0,\sigma_{\theta_j})}$ be the occupation measure of outcome (T, L). Suppose toward a contradiction that

$$\frac{X^{(\sigma_0,\sigma_{\theta_j})}}{Y^{(\sigma_0,\sigma_{\theta_j})}} < \frac{\gamma^*}{1-\gamma^*}.$$

Then there exists $\gamma \in [0, \gamma^*)$ such that the value of the left-hand side is $\frac{\gamma}{1-\gamma}$. For every $h^\tau \in \mathcal{H}$, let $\sigma_{\theta_j}(h^\tau) \in \Delta(A_1)$ be the (mixed) action prescribed by σ_{θ_j} at h^τ and let $\alpha_1(\cdot|h^\tau)$ be the buyer's expectation of the seller's action at h^τ. Letting $d(\cdot\|\cdot)$ denote the Kullback–Leibler divergence, we have

$$\mathbb{E}^{(\sigma_0,\sigma_{\theta_j})}\left[\sum_{\tau=0}^{+\infty} d\big(\sigma_{\theta_j}(h^\tau) \,\|\, \alpha_1(\cdot|h^\tau)\big)\right] \leq -\log \pi_0(\theta_j). \tag{B.4}$$

This implies that for every $\epsilon > 0$, the expected number of periods such that $d(\sigma_{\theta_j}(h^\tau) \,\|\, \alpha_1(\cdot|h^\tau)) > \epsilon$ is no more than

$$T(\epsilon) \equiv \left\lceil \frac{-\log \pi_0(\theta_j)}{\epsilon} \right\rceil. \tag{B.5}$$

Let

$$\epsilon \equiv d\left(\frac{\gamma+2\gamma^*}{3}H + \left(1 - \frac{\gamma+2\gamma^*}{3}\right)L \,\|\, \gamma^* H + (1-\gamma^*)L\right) \tag{B.6}$$

and let δ be large enough such that

$$\frac{X^{(\sigma_0,\sigma_{\theta_j})}}{Y^{(\sigma_0,\sigma_{\theta_j})} - \left(1-\delta^{T(\epsilon)}\right)} < \frac{2\gamma+\gamma^*}{3-2\gamma-\gamma^*}.$$

According to (B.4) and (B.5), if type θ_j plays according to her equilibrium strategy, then there exist at most $T(\epsilon)$ periods in which the buyers' expectation over the seller's action differs from σ_{θ_j} by more than ϵ. According to (B.6), aside from $T(\epsilon)$ periods, the buyers will trust the seller at h^t only when $\sigma_{\theta_j}(h^t)$ assigns probability at least $\frac{\gamma+2\gamma^*}{3}$ to H. Therefore, under the probability measure induced by $(\sigma_0, \sigma_{\theta_j})$, the occupation measure with which the buyer trusts the seller is at most

$$\big(1-\delta^{T(\epsilon)}\big) + \big(X^{(\sigma_0,\sigma_{\theta_j})} + Y^{(\sigma_0,\sigma_{\theta_j})} - \big(1-\delta^{T(\epsilon)}\big)\big)\frac{2\gamma+\gamma^*}{\gamma+2\gamma^*},$$

which is strictly less than $X^{(\sigma_0,\sigma_{\theta_j})} + Y^{(\sigma_0,\sigma_{\theta_j})}$ when δ is close enough to 1. This leads to a contradiction. □

Forward-looking buyers I show that the patient seller's payoff cannot significantly exceed v^* when the buyer is forward-looking and has a discount factor δ_0 close to 0.

The necessity of constraint (B.2) relies on the observation that at every on-path history, the buyer has no incentive to play T unless he expects H to be played with positive probability. This remains valid when $\delta_0 < \gamma^*$. Suppose toward a contradiction that at some on-path history h^t, all types of seller play L for sure, but the buyer plays T with strictly positive probability. The buyer's discounted average payoff by playing T at h^t is at most

$$\underbrace{(1-\delta_0)(-c)}_{\text{buyer's stage-game payoff if he plays } T \text{ while seller plays } L \text{ for sure}} + \underbrace{\delta_0 b}_{\text{buyer's maximal continuation payoff after playing } T}.$$

Since $\delta_0 < \gamma^* \equiv \frac{c}{b+c}$, the above expression is strictly less than 0. This contradicts the buyer's incentive to play T at h^t since he can secure payoff 0 by playing N in every subsequent period.

In addition, when δ_0 is close to 0, the buyer has no incentive to play T at h^t unless he expects H to be played with probability more than $\gamma^* - \varepsilon$, with ε vanishing to 0 as $\delta_0 \to 0$. This implies an approximate version of constraint (B.3) when the seller's discount factor δ_1 is close enough to 1:

$$\alpha^j(H) \geq \frac{\gamma^* - \varepsilon}{1-\gamma^*+\varepsilon}\alpha^j(L), \tag{B.7}$$

with

$$\alpha^j(y) \equiv \mathbb{E}^{(\sigma_0,\sigma_{\theta_j})}\left[\sum_{t=0}^{\infty}(1-\delta_1)\delta_1^t \mathbf{1}\{y_t = y\}\right] \quad \text{for every } y \in \{N, H, L\}.$$

Replacing (B.3) with constraint (B.7), the value of the constrained optimization problem is close to v_j^*, which converges to v_j^* as $\delta_2 \to 0$. This implies the robustness of Theorem 1 to perturbations of δ_0. Given that the proof of Theorem 2 does not use the buyers' incentive constraints aside from the conclusion that v^* is a patient seller's highest equilibrium payoff, those results are also robust to small perturbations of δ_0.

Appendix C: Proof of Theorem 2

Suppose toward a contradiction that there exists $j \in \{1, 2\}$ such that type θ_j's best reply is to mix at every on-path history. Then both playing L at every on-path history and playing H at every on-path history are her best replies against σ_0. If we order the states and actions according to $T \succ N$, $H \succ L$, and $\theta_1 \succ \theta_2$, then players' stage-game payoffs satisfy the monotone-supermodularity condition in Liu and Pei (2020).

If $j = 1$, then Lemma D.2 in Pei (2020a) implies that type θ_2 plays L with probability 1 at every history in $\mathcal{H}^{(\sigma_0, \sigma_{\theta_2})}$. Hence, there exists $T \in \mathbb{N}$ such that under the probability measure induced by $(\sigma_0, \sigma_{\theta_2})$, there are at most T periods in which buyers believe that H is played with probability at least γ^*. Type θ_2's payoff is no more than $(1 - \delta^T)$, which converges to 0 as $\delta \to 1$. This contradicts the presumption that her payoff is ε close to v_2^*.

If $j = 2$, then Lemma D.2 in Pei (2020a) implies that type θ_1 plays H with probability 1 at every history in $\mathcal{H}^{(\sigma_0, \sigma_{\theta_1})}$. After type θ_2 plays L, she separates from the type θ_1 and, according to Lemma B.2, her continuation value is no more than $1 - \theta_2$. Therefore, type θ_2's discounted average payoff in period 0 is at most $(1 - \delta) + \delta(1 - \theta_2)$. Since $v_2^* > 1 - \theta_2$, $(1 - \delta) + \delta(1 - \theta_2) < v_2^*$ when δ is close to 1.

References

Aumann, Robert J. and Michael B. Maschler (1995), *Repeated Games With Incomplete Information*. MIT Press, Cambridge. [452]

Barro, Robert (1986), "Reputation in a model of monetary policy with incomplete information." *Journal of Monetary Economics*, 17, 3–20. [449]

Benabou, Roland and Guy Laroque (1992), "Using privileged information to manipulate markets: Insiders, gurus, and credibility." *Quarterly Journal of Economics*, 107, 921–958. [451]

Board, Simon and Moritz Meyer-ter-Vehn (2013), "Reputation for quality." *Econometrica*, 81, 2381–2462. [451]

Cripps, Martin W. and Jonathan P. Thomas (2003), "Some asymptotic results in discounted repeated games of one-sided incomplete information." *Mathematics of Operations Research*, 28, 433–462. [452]

Fudenberg, Drew, David M. Kreps, and Eric S. Maskin (1990), "Repeated games with long-run and short-run players." *The Review of Economic Studies*, 57, 555–573. [450, 453, 471]

Fudenberg, Drew and David Levine (1992), "Maintaining a reputation when strategies are imperfectly observed." *Review of Economic Studies*, 59, 561–579. [451, 455]

Fudenberg, Drew and David K. Levine (1989), "Reputation and equilibrium selection in games with a patient player." *Econometrica*, 57, 759–778. [451, 454]

Gossner, Olivier (2011), "Simple bounds on the value of a reputation." *Econometrica*, 79, 1627–1641. [451, 455]

Hart, Sergiu (1985), "Nonzero-sum two-person repeated games with incomplete information." *Mathematics of Operations Research*, 10, 117–153. [452]

Hörner, Johannes, Stefano Lovo, and Tristan Tomala (2011), "Belief-free equilibria in games with incomplete information: Characterization and existence." *Journal of Economic Theory*, 146, 1770–1795. [452]

Liu, Qingmin and Andrzej Skrzypacz (2014), "Limited records and reputation bubbles." *Journal of Economic Theory*, 151, 2–29. [451]

Liu, Shuo and Harry Pei (2020), "Monotone equilibria in signalling games." *European Economic Review*, 124. [474]

Mailath, George J. and Larry Samuelson (2006), *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press, New York, New York. [450, 459, 463]

Mathevet, Laurent, David Pearce, and Ennio Stacchetti (2019), "Reputation and information design." Unpublished paper, New York University. [451]

Pei, Harry (2020a), "Reputation effects under interdependent values." *Econometrica*, 88, 2175–2202. [474]

Pei, Harry (2020b), "Trust and betrayals: Reputational payoffs and behaviors without commitment." Unpublished paper, arXiv:2006.08071. [452, 460]

Pęski, Marcin (2014), "Repeated games with incomplete information and discounting." *Theoretical Economics*, 9, 651–694. [452, 453]

Phelan, Christopher (2006), "Public trust and government betrayal." *Journal of Economic Theory*, 130, 27–43. [449]

Shalev, Jonathan (1994), "Nonzero-sum two-person repeated games with incomplete information and known-own payoffs." *Games and Economic Behavior*, 7, 246–259. [452]

Sobel, Joel (1985), "A theory of credibility." *Review of Economic Studies*, 52, 557–573. [451]

Tirole, Jean (1996), "A theory of collective reputations (with applications to the persistence of corruption and to firm quality)." *Review of Economic Studies*, 63, 1–22. [449]

Weinstein, Jonathan and Muhamet Yildiz (2016), "Reputation without commitment in finitely repeated games." *Theoretical Economics*, 11, 157–185. [451]

Co-editor Simon Board handled this manuscript.

Manuscript received 20 March, 2020; final version accepted 14 August, 2020; available online 18 August, 2020.