

Reputation for playing mixed actions: A characterization theorem

Harry Pei ¹

Department of Economics, Northwestern University, United States of America

Received 28 June 2020; final version received 14 February 2022; accepted 15 February 2022

Available online 22 February 2022

Abstract

A patient player privately observes a persistent state that directly affects his opponents' payoffs, and can be one of the several commitment types that plays the same mixed action in every period. I characterize the set of environments under which the patient player obtains at least his commitment payoff in all equilibria regardless of his stage-game payoff function. Due to the presence of interdependent values, the patient player cannot guarantee his mixed commitment payoff by imitating the mixed commitment type, and small perturbations to a pure commitment action can significantly reduce the patient player's lowest equilibrium payoff.

© 2022 Elsevier Inc. All rights reserved.

JEL classification: C73; D82; D83

Keywords: Reputation; Interdependent values; Supermartingales; Doob's Upcrossing Inequality

E-mail address: harrydp@northwestern.edu.

¹ I am indebted to Daron Acemoglu, Drew Fudenberg, Juuso Toikka, and Alex Wolitzky for guidance and support. I thank Heski Bar-Isaac, Daniel Clark, Martin Cripps, Joyee Deb, Mehmet Ekmekci, Jack Fanning, Yuhta Ishii, Elliot Lipnowski, Qingmin Liu, Shuo Liu, Lucas Maestri, Marcin Peński, Bruno Strulovici, Can Urgan, Nicolas Vieille, Geyu Yang, an associate editor, and two referees for helpful comments. I thank Jin Yang, Maren Vairo, and Tomer Yehoshua-Sandak for excellent research assistance and the NSF Grant SES-1947021 for financial support.

<https://doi.org/10.1016/j.jet.2022.105438>

0022-0531/© 2022 Elsevier Inc. All rights reserved.

1. Introduction

I examine patient players' returns from building reputations for playing mixed actions. To fix ideas, consider a profit-maximizing firm that needs to decide whether to imitate the behavior of an ethical firm that intrinsically cares about its worker and customers. Suppose the ethical firm commits to provide good customer service unless its worker is sick, and consumers can only observe the quality of service but not whether the worker is sick or healthy, then the ethical firm behaves *as if* it is mixing between providing good service and bad service.

The reputation results in Fudenberg and Levine (1989, 1992) imply that when consumers' payoffs depend only on their actions and the firm's action, there is no qualitative difference between establishing reputations for playing pure actions and establishing reputations for playing mixed actions. They show that in every equilibrium, a patient firm receives at least its commitment payoff if it plays its commitment action in every period, no matter whether this commitment action is pure or mixed.

This paper shows that in interdependent value environments, whether the commitment action is pure or mixed has significant effects on a patient player's payoff. I study a repeated game between a patient player 1 (e.g., firm) and a sequence of short-lived player 2s (e.g., consumers). Player 1 privately observes the realization of a state (e.g., product safety or durability) that is constant over time and affects both players' stage-game payoffs, in addition to knowing whether he is *strategic* or *committed*. The strategic player 1 maximizes his discounted average payoff. The committed player 1 mechanically plays the same *commitment action* in every period, which can be pure or mixed and can depend on the state. This differs from Pei (2020) which assumes that all commitment types play pure strategies. Player 2s can observe all the actions taken in the past, but they cannot directly observe the state or the mixing probabilities.

My main result characterizes the set of interdependent value environments under which the patient player receives at least his commitment payoff in all equilibria regardless of his stage-game payoff function. My characterization implies that securing commitment payoffs from mixed actions requires more demanding conditions than securing commitment payoffs from pure actions. I also show by example that small perturbations to a pure commitment action can significantly reduce the patient player's lowest equilibrium payoff.

Intuitively, when a commitment action is mixed, some pure actions in the support of this mixed commitment action can be played with strictly higher probability by some strategic types than by the mixed-strategy commitment type. If this is the case, then playing these pure actions *increases* the likelihood ratios between these strategic types and the mixed-strategy commitment type. In contrast, when the patient player plays a pure commitment action in every period, the likelihood ratio between every strategic type and that pure-strategy commitment type must be non-increasing. Therefore, more posterior beliefs are plausible in equilibrium when the commitment action is mixed, making it harder for player 1 to secure his commitment payoff in *all* equilibria.

My analysis unveils another difference between private and interdependent value environments, that when player 2's best reply to the mixed commitment action depends on the state, player 1 *cannot* secure his mixed commitment payoff by imitating the mixed-strategy commitment type. Again, this is because playing some actions in the support of a mixed commitment action can increase the likelihood ratio between some strategic types and this mixed-strategy commitment type. Since the state is persistent and affects player 2's best reply to the commitment action, player 2's belief about the state in any given period can have a long-lasting effect on her future actions, which in turn affects the patient player's continuation value. This suggests the

need for the patient player to *take actions selectively* in the support of the mixed commitment action. To the best of my knowledge, this observation is novel in the reputation literature, since the existing reputation results are shown by bounding a patient player's discounted average payoff when he imitates one of the commitment types.

However, taking actions selectively in the support of the mixed commitment action raises two new concerns. First, player 1 may play some low-payoff actions too frequently, in which case his expected payoff may fall below his commitment payoff. Second, given that player 1 may not play the mixed commitment action in every period, he may fail to convince his opponents that the mixed commitment action will be played in the future.

I establish a learning result that addresses both concerns. It shows that for every strategy profile, player 1 can find a deviation under which (1) player 2 has an incentive to play the desirable best reply to the commitment action in every period under her posterior belief; (2) with probability close to one, the discounted frequency of player 1's action is close to the mixed commitment action; (3) in expectation, player 2 believes that player 1's action is close to the mixed commitment action in all except for a bounded number of periods. My proof uses the upcrossing inequality, the central limit theorem, and the entropy techniques in Gossner (2011).

This paper contributes to the reputation literature by examining players' guaranteed returns from building reputations when values are interdependent and their opponents *cannot* perfectly monitor whether they have honored their commitment. It highlights the differences between building reputations for playing mixed actions and that for playing pure actions, as well as the role of interdependent values in driving these differences.

My analysis unveils the challenges to build reputations when learning is *confounded*. Even though the informed player can convince his opponents about his future actions, he may not teach them how to best reply when their payoff functions depend on a persistent state. This is related to Yang (2019) and Deb and Ishii (2021), in which confounded learning is caused by uncertainty in the monitoring structure. Yang (2019) focuses on private value environments and provides sufficient conditions under which the patient player can secure his commitment payoff. Deb and Ishii (2021) allow for interdependent values and uncertainty in the monitoring structure. They assume that for every pair of states, there exists an action of the long-run player such that the distribution over public signals induced by this action in the first state is different from that induced by any action in the second state. Their identification condition is violated in my model where the uninformed players learn about the informed player's type *only* through the latter's actions.

Ekmekci and Maestri (2019) and Ekmekci, Gorno, Maestri, Sun and Wei (2021) characterize an informed player's payoffs and behaviors when monitoring is imperfect and a long-lived uninformed player decides whether to irreversibly stop interacting with the informed player. By contrast, my result highlights the challenges for a patient player to build reputations when his opponents can freely choose their actions.

2. Model

Time is discrete, indexed by $t = 0, 1, 2, \dots$. A long-lived player 1 (he, e.g., a seller) with discount factor $\delta \in (0, 1)$ interacts with an infinite sequence of short-lived player 2s (she, e.g., consumer), arriving one in each period and each plays the game only once. In period t , players simultaneously choose their actions $(a_{1,t}, a_{2,t}) \in A_1 \times A_2$.

Player 1 has private information about the state $\theta \in \Theta$, and whether he is *strategic* or *committed*. Both are drawn and fixed before period 0. If player 1 is strategic, then he can flexibly

choose his actions in order to maximize his discounted average payoff. If player 1 is committed, then he mechanically follows one of the several *commitment plans*. A typical commitment plan is denoted by $\gamma : \Theta \rightarrow \Delta(A_1)$, according to which the committed player plays $\gamma(\theta) \in \Delta(A_1)$ in every period when the realized state is θ . Let Γ be an exogenous set of *feasible commitment plans* that the committed player 1 can follow. Let

$$\mathcal{A}_1^* \equiv \{\alpha_1^* \in \Delta(A_1) \mid \text{there exist } \gamma \in \Gamma \text{ and } \theta \in \Theta \text{ such that } \gamma(\theta) = \alpha_1^* \in \Delta(A_1)\}, \quad (2.1)$$

be the set of *commitment actions*. Intuitively, α_1^* belongs to $\mathcal{A}_1^*$ if and only if α_1^* is played in at least one state under at least one commitment plan in Γ . Let γ^s stand for player 1 being strategic. Let

$$\mu \in \Delta\left(\Theta \times \underbrace{(\{\gamma^s\} \cup \Gamma)}_{\text{player 1's characteristics}}\right), \quad (2.2)$$

be player 2's prior belief, which is a joint distribution of the state θ and player 1's *characteristics*, namely, whether he is strategic or committed, and if he is committed, which commitment plan in Γ he is following.

For every $\theta \in \Theta$, I say that player 1 is (*strategic*) *type* θ if he is strategic and knows that the state is θ . Let $\mu(\theta)$ be the prior probability of type θ . For every $\alpha_1^* \in \mathcal{A}_1^*$, I say that player 1 is (*commitment*) *type* α_1^* if he is committed and plays α_1^* in every period. Let $\mu(\alpha_1^*)$ be the prior probability of type α_1^* .

Assumption 1. Sets Θ , Γ , A_1 , and A_2 are finite, $|A_1| \geq 2$, $|A_2| \geq 2$, and μ has full support.²

Let $h^t \equiv \{a_{1,s}, a_{2,s}\}_{s=0}^{t-1} \in \mathcal{H}^t$ be a public history. Let $\mathcal{H} \equiv \bigcup_{t=0}^{+\infty} \mathcal{H}^t$ be the set of public histories. Player 2's history coincides with the public history. Player 1's history consists of the public history and his type. Let $\sigma_1 \equiv (\sigma_\theta)_{\theta \in \Theta}$ be player 1's strategy, with $\sigma_\theta : \mathcal{H} \rightarrow \Delta(A_1)$ the strategy for type θ . Let $\sigma_2 : \mathcal{H} \rightarrow \Delta(A_2)$ be player 2's strategy. Let $\sigma \equiv (\sigma_1, \sigma_2)$ be a typical strategy profile. Let Σ be the set of strategy profiles.

For $i \in \{1, 2\}$, player i 's stage-game payoff in period t is $u_i(\theta, a_{1,t}, a_{2,t})$, which is naturally extended to mixed actions and distributions over states. This formulation allows for interdependent values since u_2 can depend on player 1's private information θ . The solution concept is Bayes Nash equilibrium, or *equilibrium* for short. Let $\text{NE}(\delta)$ be the set of equilibria. For every $\theta \in \Theta$, type θ 's *lowest equilibrium payoff* is denoted by:

$$\underline{v}_\theta(\delta) \equiv \inf_{((\sigma_\theta)_{\theta \in \Theta}, \sigma_2) \in \text{NE}(\delta)} \mathbb{E}^{(\sigma_\theta, \sigma_2)} \left[\sum_{t=0}^{+\infty} (1 - \delta)^t u_1(\theta, a_{1,t}, a_{2,t}) \right], \quad (2.3)$$

where $\mathbb{E}^{(\sigma_\theta, \sigma_2)}[\cdot]$ is the expectation induced by type θ 's strategy σ_θ and player 2's strategy σ_2 .

For every $\phi \in \Delta(\Theta)$ and $\alpha_1 \in \Delta(A_1)$, let $\text{BR}_2(\phi, \alpha_1) \subset A_2$ be player 2's pure best replies to α_1 when the state is distributed according to ϕ . Abusing notation, let $\text{BR}_2(\theta, \alpha_1) \subset A_2$ be player 2's pure best replies to α_1 when the state is θ . I make the following assumption that is satisfied for generic u_2 :

Assumption 2. For every $\alpha_1^* \in \mathcal{A}_1^*$ and $\theta \in \Theta$, $\text{BR}_2(\theta, \alpha_1^*)$ is a singleton.

² Cases in which $|A_1| = 1$ or $|A_2| = 1$ are trivial since either player 1 or player 2 has no choice to make.

I use $a_2^*(\theta, \alpha_1^*)$ to denote the unique element of $\text{BR}_2(\theta, \alpha_1^*)$. Type θ 's commitment payoff from α_1^* is³

$$v_\theta(\alpha_1^*) \equiv u_1\left(\theta, \alpha_1^*, a_2^*(\theta, \alpha_1^*)\right). \quad (2.4)$$

For given $\theta^* \in \Theta$ and $\alpha_1^* \in \mathcal{A}_1^*$, my main result provides conditions on μ and u_2 under which

$$\underbrace{\liminf_{\delta \rightarrow 1} \underline{v}_{\theta^*}(\delta)}_{\text{type } \theta^*'s \text{ lowest equilibrium payoff}} \geq \underbrace{v_{\theta^*}(\alpha_1^*)}_{\text{type } \theta^*'s \text{ commitment payoff from } \alpha_1^*} \quad \text{for all } u_1. \quad (2.5)$$

Benchmark: When u_2 does not depend on θ , or more generally, player 2's best reply to α_1^* does not depend on θ , inequality (2.5) is implied by the results in Fudenberg and Levine (1989, 1992) and player 1 can guarantee his commitment payoff by playing α_1^* in every period. The intuition is that after observing player 1's action frequency matches α_1^* for a long time, player 2s will be convinced that player 1's action is close to α_1^* in all future periods, in which case they will play their myopic best reply to α_1^* . Hence, by playing α_1^* in every period, type θ^* receives payoff $v_{\theta^*}(\alpha_1^*)$ in all except for a bounded number of periods.

3. Result

3.1. An example

I use an example to explain why establishing reputations for playing pure actions and that for playing mixed actions are different under interdependent values. Suppose $\Theta \equiv \{\theta^*, \theta_1, \theta_2\}$ and players' payoffs are:

θ^*	G	M_1	M_2	θ_1	G	M_1	M_2	θ_2	G	M_1	M_2
H	1, 3	$-\frac{1}{2}, 0$	$-\frac{1}{2}, 0$	H	$2, \frac{1}{2}$	$2, \frac{3}{2}$	2, 0	H	$2, \frac{1}{2}$	2, 0	$2, \frac{3}{2}$
I	2, -1	$0, -\frac{1}{2}$	$0, -\frac{1}{2}$	I	2, 0	2, 1	$2, -\frac{1}{2}$	I	2, 0	$2, -\frac{1}{2}$	2, 1
L	3, -2	$\frac{1}{2}, -1$	$\frac{1}{2}, -1$	L	2, -2	2, -1	2, -1	L	2, -2	2, -1	2, -1

Intuitively, player 1 is a firm that chooses between high (H), intermediate (I), and low (L) effort, and each player 2 is a consumer who chooses between buying a good product (G), a mediocre product with the first characteristic (M_1), and a mediocre product with the second characteristic (M_2). If the state is θ^* , then exerting effort is costly for the firm and purchasing the good product is strictly optimal for consumers when the firm exerts high effort. If the state is $\theta_i \in \{\theta_1, \theta_2\}$, then effort is not costly and consumers prefer the mediocre product with characteristic i as long as the firm exerts intermediate or high effort.

Let $\alpha_1^* \equiv (1 - \varepsilon)H + \varepsilon I$ and let $\mathcal{A}_1^* \equiv \{\alpha_1^*\}$. The parameter of interest is $\varepsilon \in [0, \frac{1}{8}]$. Player 2's best reply to α_1^* in state θ^* is G . Suppose state θ^* occurs with probability 1 conditional on player 1 being commitment type α_1^* , player 2 has a strict incentive to play G if her belief satisfies the following three assumptions:

³ Notice that $\text{BR}_2(\cdot)$ and $a_2^*(\cdot)$ also depend on u_2 , $v_\theta(\cdot)$ also depends on u_1 and u_2 , and $\text{NE}(\cdot)$ and $\underline{v}_\theta(\cdot)$ also depend on μ , u_1 , and u_2 . I omit μ , u_1 , and u_2 in those functions in order to avoid cumbersome notation.

1. Both the likelihood ratio between strategic type θ_1 and commitment type α_1^* , and the likelihood ratio between strategic type θ_2 and commitment type α_1^* are no more than $\frac{5}{2}$.
2. Player 1 is either strategic type θ_1 or strategic type θ_2 or commitment type α_1^* .
3. Both strategic type θ_1 and strategic type θ_2 play α_1^* .

In what follows, I assume that player 2's *prior belief* satisfies⁴:

$$\frac{\mu(\theta_1)}{\mu(\alpha_1^*)} = \frac{\mu(\theta_2)}{\mu(\alpha_1^*)} = \frac{5}{2}.$$

I make no restriction on the probability of type θ^* . When I vary ε later on in order to discuss the differences between pure and mixed commitment actions, I keep the probability of type θ^* fixed throughout.

In the benchmark where $\varepsilon = 0$, i.e., $\alpha_1^* = H$. All the commitment actions are pure. Theorem 1' in Pei (2020) implies that as $\delta \rightarrow 1$, type θ^* receives at least his commitment payoff from α_1^* in all equilibria. This is because $\frac{\mu(\theta_1)}{\mu(\alpha_1^*)} = \frac{\mu(\theta_2)}{\mu(\alpha_1^*)} = \frac{5}{2}$ under player 2's prior belief, and the likelihood ratio between each strategic type and commitment type α_1^* cannot increase when player 1 plays the commitment action $\alpha_1^* = H$ in every period.

However, for every $\varepsilon \in (0, \frac{1}{8}]$, there are equilibria where type θ^* 's payoff is no more than $\frac{3}{4}$ when he is arbitrarily patient, and this payoff is lower than his commitment payoff $v_{\theta^*}(\alpha_1^*) = 1 + \varepsilon$. First, I construct such a low-payoff equilibrium when $\varepsilon = \frac{1}{8}$, which is relatively simple and can highlight the differences between pure and mixed commitment actions. After that, I construct low-payoff equilibria for any $\varepsilon \in (0, \frac{1}{8}]$.

- Type θ^* plays L in every period. In period 0, type θ_1 plays H and type θ_2 plays I . In every period after period 1, both type θ_1 and type θ_2 play α_1^* .
 - Player 2's beliefs at on-path histories are pinned down by her prior belief and player 1's behavior.
 - At histories where player 1 played H in period 0 and L has not occurred before, player 2's posterior belief assigns probability $\frac{5}{7-2\varepsilon}$ to type θ_1 and probability $\frac{2-2\varepsilon}{7-2\varepsilon}$ to type α_1^* .
 - At histories where player 1 played I in period 0 and L has not occurred before, player 2's posterior belief assigns probability $\frac{5}{5+2\varepsilon}$ to type θ_2 and probability $\frac{2\varepsilon}{5+2\varepsilon}$ to type α_1^* .
 - At histories where player 1 played L in all previous periods, player 2's posterior belief assigns probability $\frac{\mu(L)}{\mu(L)+\mu(\theta^*)}$ to type L and probability $\frac{\mu(\theta^*)}{\mu(L)+\mu(\theta^*)}$ to type θ^* .
- Next, I specify player 2's posterior beliefs at off-path histories. If L was played in period 0 but H or I has been played after period 0, then player 2's belief assigns probability 1 to type θ^* . If H was played in period 0 but L has been played after period 0, or if I was played in period 0 but L has been played after period 0, then player 2's belief assigns probability $\frac{1}{2}$ to type θ_1 and assigns probability $\frac{1}{2}$ to type θ_2 .
- Player 2's action in period 0 is irrelevant for player 1's incentives and payoffs when δ is close to 1. I left it unspecified in order to avoid detours. Starting from period 1,

⁴ I set $(\frac{\mu(\theta_1)}{\mu(\alpha_1^*)}, \frac{\mu(\theta_2)}{\mu(\alpha_1^*)}) = (\frac{5}{2}, \frac{5}{2})$ since when player 2's prior belief satisfies this condition, player 1 can guarantee his commitment payoff from α_1^* if α_1^* is pure, but cannot do so when α_1^* is mixed. The same conclusion holds as long as $(\frac{\mu(\theta_1)}{\mu(\alpha_1^*)}, \frac{\mu(\theta_2)}{\mu(\alpha_1^*)})$ belongs to an open neighborhood of $(\frac{5}{2}, \frac{5}{2})$.

- Player 2 plays M_1 if player 1 played H in period 0 and L has not occurred before. I verify her incentive to do that. First, conditional on facing type θ_1 , her payoff gain from playing M_1 instead of G is 1, and type θ_1 occurs with probability $\frac{5}{7-2\varepsilon}$ under her posterior. Second, since $\varepsilon = \frac{1}{8}$, conditional on facing type α_1^* , her payoff gain from playing G instead of M_1 is $\frac{41}{16}$ and type α_1^* occurs with probability $\frac{2-2\varepsilon}{7-2\varepsilon}$ under her posterior. Her best reply is M_1 since $\frac{5}{7-2\varepsilon} > \frac{41}{16} \cdot \frac{2-2\varepsilon}{7-2\varepsilon}$.
- Player 2 plays M_2 if player 1 played I in period 0 and L has not occurred before. I verify her incentive to do that. First, conditional on facing type θ_2 , her payoff gain from playing M_2 instead of G is 1, and type θ_2 occurs with probability $\frac{5}{5+2\varepsilon}$ under her posterior. Second, conditional on facing type α_1^* , her payoff gain from playing G instead of M_2 is $\frac{41}{16}$, and type α_1^* occurs with probability $\frac{2\varepsilon}{5+2\varepsilon}$ under her posterior. Her best reply is M_2 since $\frac{5}{5+2\varepsilon} > \frac{2\varepsilon}{5+2\varepsilon} \cdot \frac{41}{16}$ when $\varepsilon = \frac{1}{8}$.
- Player 2 plays $\frac{1}{2}M_1 + \frac{1}{2}M_2$ at other histories. This is incentive compatible since at those histories, either she believes that L will be played for sure (by type θ^* and type L), or she believes that both type θ_1 and type θ_2 occur with probability $\frac{1}{2}$ and that none of these two types will play L .
- Type θ^* 's payoff from playing L in every period is no more than $3(1-\delta) + \frac{\delta}{2}$, which is close to $\frac{1}{2}$ as $\delta \rightarrow 1$. He has no incentive to play H or I since doing so reduces his stage-game payoff but cannot increase his continuation value. The latter is because player 2 will never play G starting from period 1, so type θ^* 's continuation value after period 1 is *at most* $\frac{1}{2}$ regardless of his actions.

Intuitively, when α_1^* is mixed, playing some actions in the support of α_1^* can *increase* the likelihood ratio between some strategic type and mixed commitment type α_1^* . This cannot happen when α_1^* is pure, in which case playing α_1^* cannot increase the likelihood ratio between any strategic type and commitment type α_1^* . The above equilibrium highlights this difference between pure and mixed commitment actions. In particular, (1) if player 1 plays H in period 0, then the likelihood ratio between type θ_1 and type α_1^* increases, after which player 2 prefers M_1 to G ; (2) if player 1 plays I in period 0, then the likelihood ratio between type θ_2 and type α_1^* increases, after which player 2 prefers M_2 to G ; (3) if player 1 plays L in period 0, then he separates from commitment type α_1^* , after which player 2 believes that player 1 will play L instead of α_1^* .⁵

In order to highlight the discontinuity in type θ^* 's equilibrium payoff at $\varepsilon = 0$, for any $\varepsilon \in (0, \frac{1}{8}]$, I construct an equilibrium in which type θ^* 's payoff is approximately $\frac{3}{4}$ when δ is close to 1. Let T be the unique integer that satisfies $\frac{5}{2(1-\varepsilon)^T} \in [5(1-\varepsilon)^2, 5(1-\varepsilon))$. By definition, T depends on ε but does not depend on δ . Since $\varepsilon \leq \frac{1}{8}$, we have $T \geq 1$. Type θ_1 plays H from period 0 to T , and plays α_1^* starting from period $T+1$. Type θ^* plays L from period 0 to T , plays I in period $T+1$, and starting from period $T+2$, he plays H in odd periods and L in even periods.⁶ For the behavior of type θ_2 ,

⁵ Under other strategies of player 1, there can exist an action such that taking that action increases the likelihood ratio between type θ_1 and type α_1^* and the likelihood ratio between type θ_2 and type α_1^* . For example, suppose both type θ_1 and type θ_2 play H with probability 1 in period 0, and $\alpha_1^* = (1-\varepsilon)H + \varepsilon I$ with $\varepsilon \in (0, 1)$. Then after observing H in period 0, both the likelihood ratio between type θ_1 and type α_1^* and the likelihood ratio between type θ_2 and type α_1^* increase.

⁶ When ε is large enough such that $T = 1$, I can let type θ^* playing L in every period starting from period 1. When $T > 1$, player 2 may have an incentive to play G in period $t \in \{1, 2, \dots, T-1\}$ after observing H in period 0. I let type

- In period $t \in \{0, \dots, T-1\}$, if I has never been played before, then type θ_2 plays H with probability

$$P(\varepsilon, t) \equiv \frac{2(1-\varepsilon)^{t+1} - 1}{2(1-\varepsilon)^t - 1}, \quad (3.1)$$

and plays I with probability $1 - P(\varepsilon, t)$. Since $\frac{5}{2(1-\varepsilon)^T} < 5(1-\varepsilon)$, $P(\varepsilon, t) \in (0, 1)$ for every $t \leq T-1$.

- If I has never been played before period T , then type θ_2 plays I for sure in period T .
- For every $t \geq 1$, if I has been played before period t , then type θ_2 plays α_1^* in period t .

At every off-path history, player 2's posterior belief assigns zero probability to both commitment types, and assigns positive probability only to strategic types. She believes that all strategic types play L at those off-path histories, in which case player 2 plays $\frac{1}{2}M_1 + \frac{1}{2}M_2$. Next, I specify player 2's beliefs and behaviors at *on-path histories* starting from period 1.⁷

- If she observes L in period 0, then her posterior belief assigns positive probability only to strategic type θ^* and commitment type L . She plays $\frac{1}{2}M_1 + \frac{1}{2}M_2$ from period 1 to $T+1$. Starting from period $T+2$, if period $T+1$ action was I , and H has been played in every odd period starting from period $T+2$, she plays G in odd periods and plays $\frac{1}{2}M_1 + \frac{1}{2}M_2$ in even periods. Otherwise, she plays $\frac{1}{2}M_1 + \frac{1}{2}M_2$.
- After observing H from period 0 to T , player 2's belief assigns positive probability only to type θ_1 and type α_1^* . Since $\frac{\mu(\theta_1)}{\mu(\alpha_1^*)} = \frac{5}{2}$, the posterior likelihood ratio between type θ_1 and type α_1^* equals $\frac{5}{2(1-\varepsilon)^{T+1}}$. This is at least $5(1-\varepsilon) \geq \frac{35}{8}$ since $\frac{5}{2(1-\varepsilon)^T} \in [5(1-\varepsilon)^2, 5(1-\varepsilon))$ and $\varepsilon \leq \frac{1}{8}$. She strictly prefers M_1 to G starting from period $T+1$, since her payoff gain from playing M_1 instead of G is 1 when facing type θ_1 and her payoff gain from playing G instead of M_1 is at most 3 when facing type α_1^* .
- At every on-path history where I has been played in at least one period from 0 to T , player 2's belief assigns positive probability only to type θ_2 and type α_1^* . This is because type θ_1 plays H from period 0 to T and type θ^* plays L from period 0 to T . I verify in Online Appendix B that the likelihood ratio between type θ_2 and type α_1^* is at least 5 at such histories. This uses the assumption that player 2's prior belief satisfies $\frac{\mu(\theta_2)}{\mu(\alpha_1^*)} = \frac{5}{2}$. Under such a posterior belief, player 2 strictly prefers M_2 to G starting from period $T+1$, since her payoff gain from playing M_2 instead of G is 1 when facing type θ_2 , and her payoff gain from playing G instead of M_2 is $\frac{41}{16}$ when facing type α_1^* .

Type θ^* has a strict incentive to follow his equilibrium strategy when δ is close to 1 regardless of player 2's actions from period 0 to T . This is because his equilibrium payoff is at least $\delta^{T+1} \cdot \frac{3}{4}$, while his continuation value after any deviation is at most $3(1-\delta^{T+1}) + \delta^{T+1} \frac{1}{2}$. When δ is close

⁷ θ^* playing H with positive frequency in order to make sure that his payoff from following his equilibrium strategy is no less than his payoff from any deviation.

⁷ I leave some of player 2's actions from period 0 to period T unspecified. This is because type θ_1 and type θ_2 's incentive constraints are trivial, and in my equilibrium construction, type θ^* 's continuation value after playing L for T periods is approximately $\frac{3}{4}$, and his continuation value when he does not play L from period 0 to T is $\frac{1}{2}$. Hence, player 2's behaviors from period 0 to T do not affect type θ^* 's incentives when δ is close enough to 1.

to 1, his equilibrium payoff is approximately $\frac{3}{4}$, which is bounded below $v_{\theta^*}(\alpha_1^*)$ but is strictly greater than his payoff after any deviation.

Remark 1. The above constructions do not work when $\varepsilon = 0$. This is because the key to these constructions is that in period 0, playing H increases the likelihood ratio between strategic type θ_1 and commitment type α_1^* , and playing I increases the likelihood ratio between strategic type θ_2 and commitment type α_1^* . When $\varepsilon = 0$, the posterior likelihood ratios between type θ_1 and type α_1^* and that between type θ_2 and type α_1^* *cannot* increase when player 1 plays α_1^* in every period. Given that $\frac{\mu(\theta_1)}{\mu(\alpha_1^*)} = \frac{\mu(\theta_2)}{\mu(\alpha_1^*)} = \frac{5}{2}$, player 2 has a strict incentive to play G against α_1^* under her posterior belief. This implies that type θ^* 's payoff is no less than 1 as $\delta \rightarrow 1$.

Remark 2. In this example, when δ is close to 1, type θ^* 's payoff is at least 1 in every equilibrium when $\varepsilon = 0$, but is no more than $\frac{3}{4}$ in some equilibria when $\varepsilon \in (0, \frac{1}{8}]$. Type θ^* 's commitment payoff from α_1^* equals $1 + \varepsilon$, which is increasing in ε . However, due to the presence of interdependent values, his lowest equilibrium payoff decreases when ε increases from 0 to something strictly positive. That being said, a small perturbation to a pure commitment action can significantly reduce the patient player's lowest equilibrium payoff.

3.2. Statement of result

Let $m \equiv |\Theta|$. For every $\theta \in \Theta$ and $\alpha_1^* \in \mathcal{A}_1^*$, let $\lambda_\theta(\alpha_1^*) \equiv \frac{\mu(\theta)}{\mu(\alpha_1^*)}$ be the prior likelihood ratio between strategic type θ and commitment type α_1^* . Let $\lambda(\alpha_1^*) \equiv \{\lambda_\theta(\alpha_1^*)\}_{\theta \in \Theta} \in \mathbb{R}_+^m$ be the *prior likelihood ratio vector* with respect to α_1^* . Let $\phi_{\alpha_1^*} \in \Delta(\Theta)$ be the state distribution *conditional on* player 1 being commitment type α_1^* . Assumption 1 requires μ to have full support, under which both $\lambda(\alpha_1^*)$ and $\phi_{\alpha_1^*}$ are well-defined and can be computed from μ . Let $\Lambda(\theta^*, \alpha_1^*) \subset \mathbb{R}_+^m$ be the set of $\{\lambda_\theta\}_{\theta \in \Theta}$ such that $a_2^*(\theta^*, \alpha_1^*)$ is the unique element of

$$\arg \max_{a_2 \in A_2} \{u_2(\phi_{\alpha_1^*}, \alpha_1^*, a_2) + \sum_{\theta \in \Theta} \tilde{\lambda}_\theta u_2(\theta, \alpha_1^*, a_2)\}$$

for every $\{\tilde{\lambda}_\theta\}_{\theta \in \Theta}$ that satisfies $0 \leq \tilde{\lambda}_\theta \leq \lambda_\theta$ for all $\theta \in \Theta$.

For the example of Subsection 3.1, I depict $\Lambda(\theta^*, \alpha_1^*)$ in the left panel of Fig. 1. When μ is such that $\frac{\mu(\theta_1)}{\mu(\alpha_1^*)} = \frac{\mu(\theta_2)}{\mu(\alpha_1^*)} = \frac{5}{2}$, we have $\lambda_{\theta_1}(\alpha_1^*) = \lambda_{\theta_2}(\alpha_1^*) = \frac{5}{2}$, so that $\lambda(\alpha_1^*) \in \Lambda(\theta^*, \alpha_1^*)$ for every $\varepsilon \in [0, \frac{1}{8}]$.

Theorem 1' in Pei (2020) shows that when *all* commitment actions are pure, inequality (2.5) is true if $\lambda(\alpha_1^*) \in \Lambda(\theta^*, \alpha_1^*)$. Intuitively, $\lambda(\alpha_1^*) \in \Lambda(\theta^*, \alpha_1^*)$ implies that (1) $a_2^*(\theta^*, \alpha_1^*)$ is player 2's best reply to α_1^* conditional on the event that player 1 is either strategic or is the commitment type who plays α_1^* in every period, and (2) $a_2^*(\theta^*, \alpha_1^*)$ remains player 2's best reply to α_1^* when the likelihood ratio between every strategic type and commitment type α_1^* weakly decreases (relative to the prior likelihood ratio). The patient player receives at least his commitment payoff in all equilibria since when α_1^* is pure, no entry of the likelihood ratio vector can increase when he plays α_1^* in every period. The argument in Fudenberg and Levine (1989) implies that when player 1 plays α_1^* in every period, player 2 plays $a_2^*(\theta^*, \alpha_1^*)$ in all except for a bounded number of periods, so type θ^* 's payoff is at least his commitment payoff from α_1^* when he is sufficiently patient. This result straightforwardly extends when α_1^* is pure while other commitment actions in $\mathcal{A}_1^*$ may be mixed, since the probabilities of other commitment types vanish exponentially when player 1 plays α_1^* in every period.




Fig. 1. The relevant sets of likelihood ratio vectors in the example of Subsection 3.1 when $\varepsilon = \frac{1}{8}$ and $\frac{\mu(\theta_1)}{\mu(\alpha_1^*)} = \frac{\mu(\theta_2)}{\mu(\alpha_1^*)} = \frac{5}{2}$, with $\Lambda(\theta^*, \alpha_1^*)$ in the left panel, and $\underline{\Lambda}(\theta^*, \alpha_1^*)$ in the right panel.

By contrast, inequality (2.5) fails when $\lambda(\alpha_1^*)$ does not belong to the closure of $\Lambda(\theta^*, \alpha_1^*)$. The proof is substantially different from that in Pei (2020) due to the presence of other commitment types that can play mixed strategies. I state this observation as Statement 1 in Theorem 1.

When α_1^* is nontrivially mixed, $\lambda \in \Lambda(\theta^*, \alpha_1^*)$ is no longer sufficient for (2.5) since playing some actions in the support of α_1^* may *increase* some entries of the likelihood ratio vector. Such a possibility is highlighted in the example of Subsection 3.1 when the commitment action is $\alpha_1^* \equiv (1 - \varepsilon)H + \varepsilon I$ with $\varepsilon \in (0, \frac{1}{8}]$.

I propose a sufficient condition under which the aforementioned problem disappears. Let

$$\Theta_{(\theta^*, \alpha_1^*)}^b \equiv \{\theta \in \Theta \mid a_2^*(\theta^*, \alpha_1^*) \neq a_2^*(\theta, \alpha_1^*)\} \quad (3.2)$$

be the set of states under which player 2's best reply to α_1^* differs from that in state θ^* . In the example of Subsection 3.1, $\Theta_{(\theta^*, \alpha_1^*)}^b = \{\theta_1, \theta_2\}$. The result and proof in Fudenberg and Levine (1992) extend to the case where $\Theta_{(\theta^*, \alpha_1^*)}^b = \emptyset$. If $\Theta_{(\theta^*, \alpha_1^*)}^b \neq \emptyset$, then for every $\theta \in \Theta_{(\theta^*, \alpha_1^*)}^b$, let ψ_θ^* be the largest $\psi \in \mathbb{R}_+$ such that:

$$a_2^*(\theta^*, \alpha_1^*) \in \arg \max_{a_2 \in A_2} \{u_2(\phi_{\alpha_1^*}, \alpha_1^*, a_2) + \psi u_2(\theta, \alpha_1^*, a_2)\}. \quad (3.3)$$

Intuitively, ψ_θ^* is the intercept of $\Lambda(\theta^*, \alpha_1^*)$ on the axis for λ_θ , which is depicted in Fig. 1. Let

$$\underline{\Lambda}(\theta^*, \alpha_1^*) = \left\{ (\lambda_\theta)_{\theta \in \Theta} \in \mathbb{R}_+^m \mid \sum_{\theta \in \Theta_{(\theta^*, \alpha_1^*)}^b} \lambda_\theta / \psi_\theta^* < 1 \right\}. \quad (3.4)$$

Since $\underline{\Lambda}(\theta^*, \alpha_1^*)$ is characterized by a linear inequality, both $\underline{\Lambda}(\theta^*, \alpha_1^*)$ and $\mathbb{R}_+^m \setminus \underline{\Lambda}(\theta^*, \alpha_1^*)$ are convex sets. For the example of Subsection 3.1, I depict $\underline{\Lambda}(\theta^*, \alpha_1^*)$ in the right panel of Fig. 1 (which is the gray triangle). When the prior belief μ is such that $\frac{\mu(\theta_1)}{\mu(\alpha_1^*)} = \frac{\mu(\theta_2)}{\mu(\alpha_1^*)} = \frac{5}{2}$, we have $\lambda(\alpha_1^*) \notin \underline{\Lambda}(\theta^*, \alpha_1^*)$.

In order to understand intuitively why the patient player can secure his commitment payoff from mixed action α_1^* when $\lambda(\alpha_1^*) \in \underline{\Lambda}(\theta^*, \alpha_1^*)$, notice that every entry of the likelihood ratio

vector is a non-negative supermartingale under the probability measure induced by type α_1^* .⁸ Since $\mathbb{R}_+^m \setminus \underline{\Delta}(\theta^*, \alpha_1^*)$ is convex, if the prior likelihood ratio vector belongs to $\underline{\Delta}(\theta^*, \alpha_1^*)$, then there exists at least one pure action a_1 in the support of α_1^* such that the posterior likelihood ratio vector belongs to $\underline{\Delta}(\theta^*, \alpha_1^*)$ after observing a_1 . Hence, as long as player 1 plays a_1 in that period, player 2 will have an incentive to play $a_2^*(\theta^*, \alpha_1^*)$ in the next period when she believes that α_1^* will be played with high enough probability.

By contrast, $\mathbb{R}_+^m \setminus \Lambda(\theta^*, \alpha_1^*)$ is not necessarily convex. Hence, it is possible that $\lambda(\alpha_1^*)$ belongs to $\Lambda(\theta^*, \alpha_1^*)$ but the posterior likelihood ratio vector does not belong to $\Lambda(\theta^*, \alpha_1^*)$ no matter which action player 1 plays in the current period. As a result, player 2 may have an incentive not to play $a_2^*(\theta^*, \alpha_1^*)$ even when they are convinced that player 1 will play commitment action α_1^* . In the example of Subsection 3.1, $\lambda(\alpha_1^*)$ belongs to $\Lambda(\theta^*, \alpha_1^*) \setminus \underline{\Delta}(\theta^*, \alpha_1^*)$. In the equilibrium I constructed for $\varepsilon = \frac{1}{8}$, no matter which action player 1 plays in period 0, the posterior likelihood ratio vector in period 1 does not belong to $\Lambda(\theta^*, \alpha_1^*)$ and player 2 has no incentive to play G even when she is convinced that player 1 will play α_1^* . Similarly, in the equilibrium I constructed for any $\varepsilon \in (0, \frac{1}{8}]$, the posterior likelihood ratio vector after period T does not belong to $\Lambda(\theta^*, \alpha_1^*)$ no matter which actions player 1 played from period 0 to period T .

Theorem 1. For every pure commitment action $\alpha_1^* \in \mathcal{A}_1^*$ and every $\theta^* \in \Theta$,

1. If $\lambda(\alpha_1^*)$ does not belong to the closure of $\Lambda(\theta^*, \alpha_1^*)$ and $BR_2(\phi_{\alpha_1^*}, \alpha_1^*)$ is a singleton, then there exists u_1 such that $\limsup_{\delta \rightarrow 1} \underline{v}_{\theta^*}(\delta) < v_{\theta^*}(\alpha_1^*)$.

For every nontrivially mixed commitment action $\alpha_1^* \in \mathcal{A}_1^*$ and every $\theta^* \in \Theta$,

2. If $\lambda(\alpha_1^*)$ does not belong to the closure of $\underline{\Delta}(\theta^*, \alpha_1^*)$, $BR_2(\phi_{\alpha_1^*}, \alpha_1^*)$ is a singleton, and α_1^* does not belong to the convex hull of $\mathcal{A}_1^* \setminus \{\alpha_1^*\}$, then there exists u_1 such that $\limsup_{\delta \rightarrow 1} \underline{v}_{\theta^*}(\delta) < v_{\theta^*}(\alpha_1^*)$.
3. If $\lambda(\alpha_1^*) \in \underline{\Delta}(\theta^*, \alpha_1^*)$, then $\liminf_{\delta \rightarrow 1} \underline{v}_{\theta^*}(\delta) \geq v_{\theta^*}(\alpha_1^*)$ for every u_1 .

Theorem 1 points out the failure of reputation effects in repeated incomplete information games with interdependent values. According to this interpretation, my model is obtained by perturbing a repeated incomplete information game with a small probability of commitment types. When every commitment type is *arbitrarily unlikely* relative to every strategic type and player 2's best reply to α_1^* depends on the state, the prior likelihood ratio vector $\lambda(\alpha_1^*)$ does not belong to the closures of $\Lambda(\theta, \alpha_1^*)$ and $\underline{\Delta}(\theta, \alpha_1^*)$ for any $\theta \in \Theta$.

Theorem 1 also evaluates the robustness of reputation effects in private value reputation games against *interdependent value perturbations*. Under this interpretation, a private value reputation game is perturbed with a small probability of *other strategic types*. Such a perturbation captures situations such as buyers facing uncertainty about the safety or durability of the seller's products,

⁸ The likelihood ratio may not be a martingale since α_1^* may not have full support. For example, suppose the commitment type plays $\alpha_1^* \equiv a_1^* \in A_1$, and type θ plays a_1^* with probability 1/2 and plays another action a_1 with probability 1/2. Conditional on the probability measure induced by α_1^* , action a_1^* occurs with probability 1. After observing a_1^* , the likelihood ratio between type θ and type α_1^* is multiplied by 1/2. Hence, the expected posterior likelihood ratio is strictly lower than the expected prior likelihood ratio.

which the seller knows more about. My sufficient conditions are satisfied when the short-run players' doubt on their own payoffs is sufficiently small.

In the rest of this subsection, I explain the ideas behind the proof of Theorem 1. The full proofs can be found in Subsection 3.3 (statement 1), Online Appendix A (statement 2), and Subsection 3.4 (statement 3).

Idea behind the proof of Statement 1: I replace α_1^* with a_1^* given that α_1^* is pure. I replace $a_2^*(\theta^*, a_1^*)$ with a_2^* in order to simplify notation. Similar to Pei (2020) which studies the case where all commitment actions are pure, I take player 1's stage-game payoff function to be $u_1(\theta, a_1, a_2) \equiv \mathbf{1}\{\theta = \theta^*, a_1 = a_1^*, a_2 = a_2^*\}$.

The key challenge is to handle commitment types other than a_1^* who play mixed strategies. To illustrate, suppose for example, $\Theta = \{\theta^*, \tilde{\theta}\}$, $A_1 = \{a_1^*, a_1'\}$, $\tilde{\theta} \in \Theta_{(\theta^*, a_1^*)}^b$, $\mathcal{A}_1^* = \{a_1^*, \alpha_1'\}$, α_1' is mixed, with $\text{BR}_2(\phi_{a_1^*}, a_1^*) = \text{BR}_2(\phi_{\alpha_1'}, \alpha_1') = a_2^*$. Given the assumption that $\text{BR}_2(\phi_{\alpha_1'}, \alpha_1') = a_2^*$ and the fact that my result makes no restriction on the probability of type α_1' , the presence of commitment type α_1' encourages player 2 to play a_2^* , which can help type θ^* to obtain a high payoff. Unlike commitment types who play pure strategies, commitment type α_1' plays a completely mixed strategy, so it occurs with positive probability at every history. Hence, type α_1' needs to be taken into account when verifying player 2's incentive constraints at every history.

My proof constructs an equilibrium in which the strategic types in $\Theta_{(\theta^*, a_1^*)}^b$ play *non-stationary mixed strategies*. In the above example, type θ^* plays a pure strategy on the equilibrium path. In all periods except for a few periods in the beginning, he plays either a_1^* or a_1' . With probability close to 1, type $\tilde{\theta}$ plays a_1^* in every period, which makes sure that player 2 has no incentive to play a_2^* after observing a_1^* in period 0. With complementary probability, type $\tilde{\theta}$ uses strategy $\sigma_{\alpha_1'}$ described as follows:

- He plays α_1' at histories that occur with positive probability under type θ^* 's equilibrium strategy.
- At every history that occurs with zero probability under type θ^* 's equilibrium strategy, he plays a completely mixed action $\hat{\alpha}_1'$ that attaches a higher probability to a_1^* compared to α_1' .

The above construction ensures that type θ^* 's continuation value is bounded away from 1 after he deviates from his equilibrium strategy, no matter how and when he deviates. This is because first, the likelihood ratio between type $\tilde{\theta}$ who plays $\sigma_{\alpha_1'}$ and type α_1' remains unchanged when type θ^* follows his equilibrium strategy. Hence, no matter when and how type θ^* deviates from his equilibrium strategy, the above likelihood ratio is the same after his first deviation. Second, after type θ^* deviates in period $t \geq 1$, player 2's posterior belief attaches positive probability only to type $\tilde{\theta}$ and type α_1' .⁹ Since $\hat{\alpha}_1'$ attaches higher probability to a_1^* compared to α_1' , player 2's posterior belief about type $\tilde{\theta}$ increases every time she observes a_1^* . When the probability of type $\tilde{\theta}$ is large enough, player 2 has no incentive to play a_2^* . This leads to an upper bound on the frequency with which player 1 can play a_1^* while inducing player 2 to play a_2^* , which also bounds type θ^* 's payoff from above.

⁹ Another possibility is that type θ^* deviates in period 0 by playing a_1^* in every period, in which case player 2's posterior belief attaches positive probability to type α_1' , type a_1^* , and type $\tilde{\theta}$. Since type $\tilde{\theta}$ plays a_1^* in every period with probability close to 1 and $\lambda(a_1^*) \notin \text{cl}(\Lambda(\theta^*, a_1^*))$, player 2 has no incentive to play a_2^* after observing player 1 has played a_1^* in every period.

One can then pick the frequency with which type θ^* plays a_1^* on the equilibrium path to be such that his continuation value at every on-path history is bounded below his commitment payoff 1 but is strictly greater than the upper bound on his continuation value after any of his deviation. This verifies type θ^* 's incentive to play his equilibrium strategy and finishes the construction of a low-payoff equilibrium.

Idea behind the proof of Statement 2: The proof of Statement 2 requires four additional steps, which are summarized below with details in Online Appendix A. First, player 1's stage-game payoff is 0 if $\theta \neq \theta^*$ or if $a_2 \neq a_2^*(\theta^*, \alpha_1^*)$. Player 1's stage-game payoff when $\theta = \theta^*$ and $a_2 = a_2^*(\theta^*, \alpha_1^*)$ satisfies:

$$u_1(\theta^*, \alpha_1^*, a_2^*(\theta^*, \alpha_1^*)) > \max_{\alpha_1 \in \mathcal{A}_1^* \setminus \{\alpha_1^*\}} u_1(\theta^*, \alpha_1, a_2^*(\theta^*, \alpha_1^*)). \quad (3.5)$$

According to the separating hyperplane theorem, such u_1 exists given that $\alpha_1^* \notin \text{co}(\mathcal{A}_1^* \setminus \{\alpha_1^*\})$.

Expression (3.5) implies that type θ^* 's commitment payoff from every commitment action in $\mathcal{A}_1^* \setminus \{\alpha_1^*\}$ is strictly lower than his commitment payoff from α_1^* . This is required since my result imposes no restriction on the probabilities of commitment types other than α_1^* . When other commitment types occur with high probability, type θ^* may secure his commitment payoff from actions other than α_1^* . This is why I require his commitment payoff from any action in $\mathcal{A}_1^* \setminus \{\alpha_1^*\}$ to be strictly lower than his commitment payoff from α_1^* .

In Online Appendix A.1, I construct strategies for strategic types other than θ^* under which there exists $T \in \mathbb{N}$ such that player 2's posterior in period T is bounded away from $\Lambda(\theta^*, \alpha_1^*)$ no matter which actions player 1 played from period 0 to T . This is feasible as long as the prior likelihood ratio vector $\lambda(\alpha_1^*)$ does not belong to the closure of $\underline{\Lambda}(\theta^*, \alpha_1^*)$. An example of such a T -period strategy is displayed in Subsection 3.1 when constructing low-payoff equilibria for ε arbitrarily close to 0. In general, when there are multiple strategic types in $\Theta_{(\theta^*, \alpha_1^*)}^b$ and α_1^* is nontrivially mixed,¹⁰ one can construct strategies for strategic types that belong to $\Theta_{(\theta^*, \alpha_1^*)}^b$ such that no matter which pure action player 1 plays in the support of α_1^* , the likelihood ratio between some strategic type in $\Theta_{(\theta^*, \alpha_1^*)}^b$ and commitment type α_1^* increases, and moreover, the posterior likelihood ratio vector does not belong to $\Lambda(\theta^*, \alpha_1^*)$ after T periods no matter which actions player 1 played in T periods.

In Online Appendix A.2, I construct a strategy for strategic types other than θ^* at histories where player 2 has ruled out type θ^* . This is to make sure that type θ^* 's continuation value is bounded below $v_{\theta^*}(\alpha_1^*)$ after any of his deviations. Similar to the proof of Statement 1, I construct strategies for types in $\Theta_{(\theta^*, \alpha_1^*)}^b$ such that playing actions where type θ^* receives a high payoff increases the probability of types in $\Theta_{(\theta^*, \alpha_1^*)}^b$.¹¹ Hence, if player 1 plays high-payoff actions too frequently, player 2 will have no incentive to play $a_2^*(\theta^*, \alpha_1^*)$ until she observes some low-payoff actions. This bounds the frequency with which type θ^* can play high-payoff actions while inducing player 2 to play $a_2^*(\theta^*, \alpha_1^*)$. This in turn bounds type θ^* 's continuation value from above.

¹⁰ If $\Theta_{(\theta^*, \alpha_1^*)}^b$ has only one element, then by definition, $\Lambda(\theta^*, \alpha_1^*) = \underline{\Lambda}(\theta^*, \alpha_1^*)$. If this is the case, then $\lambda(\alpha_1^*) \notin \underline{\Lambda}(\theta^*, \alpha_1^*)$ implies that $\lambda(\alpha_1^*) \notin \Lambda(\theta^*, \alpha_1^*)$, which means that the step in Online Appendix A.1 is redundant.

¹¹ Type θ^* 's payoff is 0 if $a_2 \neq a_2^*(\theta^*, \alpha_1^*)$, and his payoff if $a_2 = a_2^*(\theta^*, \alpha_1^*)$ is constructed via the separating hyperplane theorem.

In Online Appendix A.3, I construct players' strategies at histories where player 2's belief attaches positive probability to type θ^* . The key step is to ensure that type θ^* 's continuation value is bounded below his commitment payoff $v_{\theta^*}(\alpha_1^*)$, but is greater than the upper bound on his continuation value after reaching any history where player 2's posterior belief attaches zero probability to type θ^* . There exist such strategies since the upper bound on type θ^* 's continuation value after he deviates is bounded below $v_{\theta^*}(\alpha_1^*)$. This accomplishes the construction of an equilibrium where type θ^* 's payoff is bounded below $v_{\theta^*}(\alpha_1^*)$.

Idea behind the proof of Statement 3: Since playing actions in the support of α_1^* can *increase* the likelihood ratios between some strategic types in $\Theta_{(\theta^*, \alpha_1^*)}^b$ and commitment type α_1^* , player 1 may not secure his commitment payoff by playing α_1^* in every period. This is because when some of these likelihood ratios are too large so that the posterior likelihood ratio vector does not belong to $\underline{\Delta}(\theta^*, \alpha_1^*)$, player 2 may not have incentive to play $a_2^*(\theta^*, \alpha_1^*)$ even if she is convinced that player 1 will play α_1^* .

My proof constructs a strategy for type θ^* under which he can guarantee his commitment payoff from mixed action α_1^* . Under such a strategy, type θ^* takes actions selectively from the support of α_1^* . The key step of the proof is to show that as long as the prior likelihood ratio vector satisfies $\lambda(\alpha_1^*) \in \underline{\Delta}(\theta^*, \alpha_1^*)$, one can construct a deviation for type θ^* that achieves the following objectives (see Proposition 1 in Subsection 3.4):

1. The posterior likelihood ratio vector belongs to $\underline{\Delta}(\theta^*, \alpha_1^*)$ in every period.
2. For every $a_1 \in \text{supp}(\alpha_1^*)$, the frequency with which type θ^* plays a_1 is approximately $\alpha_1^*(a_1)$.
3. Player 2s believe that actions close to α_1^* will be played in all except for a bounded number of periods.

3.3. Proof of Statement 1: constructing low payoff equilibria

In this subsection, I replace α_1^* with a_1^* , $a_2^*(\theta^*, \alpha_1^*)$ with a_2^* , and $\lambda(a_1^*)$ with λ .

Step 1: I show that when the prior likelihood ratio vector with respect to a_1^* , denoted by λ , does not belong to the closure of $\Lambda(\theta^*, a_1^*)$, there exist $a_2' \neq a_2^*$ and $\lambda' \equiv \{\lambda'_\theta\}_{\theta \in \Theta}$ that satisfy the following four conditions:

1. $\lambda'_{\theta^*} = 0$,
2. $0 \leq \lambda'_\theta \leq \lambda_\theta$ for every $\theta \in \Theta$,
- 3.

$$\sum_{\theta \in \Theta} \lambda'_\theta \left(u_2(\theta, a_1^*, a_2') - u_2(\theta, a_1^*, a_2^*) \right) > 0, \quad (3.6)$$

- 4.

$$u_2(\phi_{a_1^*}, a_1^*, a_2') - u_2(\phi_{a_1^*}, a_1^*, a_2^*) + \sum_{\theta \in \Theta} \lambda'_\theta \left(u_2(\theta, a_1^*, a_2') - u_2(\theta, a_1^*, a_2^*) \right) > 0. \quad (3.7)$$

Recall Statement 1 requires that $\text{BR}_2(\phi_{a_1^*}, a_1^*)$ is a singleton. I consider two cases. First, suppose the unique element in $\text{BR}_2(\phi_{a_1^*}, a_1^*)$ is a_2^* . According to the definition of $\Lambda(\theta^*, a_1^*)$, there exists $\lambda'' \equiv \{\lambda''_\theta\}_{\theta \in \Theta}$ such that $0 \leq \lambda''_\theta \leq \lambda_\theta$ for every $\theta \in \Theta$ and

$$a_2^* \notin \arg \max_{a_2 \in A_2} \left\{ u_2(\phi_{a_1^*}, a_1^*, a_2) + \sum_{\theta \in \Theta} \lambda''_{\theta} u_2(\theta, a_1^*, a_2) \right\}.$$

Let $\lambda' \in \mathbb{R}_+^m$ be defined as $\lambda'_{\theta^*} \equiv 0$ and $\lambda'_{\theta} \equiv \lambda''_{\theta}$ for all $\theta \neq \theta^*$. By definition, there exists $a'_2 \neq a_2^*$ such that

$$u_2(\phi_{a_1^*}, a_1^*, a'_2) + \sum_{\theta \in \Theta} \lambda'_{\theta} u_2(\theta, a_1^*, a'_2) > u_2(\phi_{a_1^*}, a_1^*, a_2^*) + \sum_{\theta \in \Theta} \lambda'_{\theta} u_2(\theta, a_1^*, a_2^*).$$

Inequalities (3.6) and (3.7) hold under likelihood ratio vector λ' and action a'_2 .

Second, suppose the unique element in $\text{BR}_2(\phi_{a_1^*}, a_1^*)$ is some a'_2 that is not a_2^* . Then there exists $\theta' \in \Theta$ such that $u_2(\theta', a_1^*, a'_2) > u_2(\theta', a_1^*, a_2^*)$. Let $\lambda' \equiv (\lambda'_{\theta})_{\theta \in \Theta} \in \mathbb{R}_+^m$ be defined as $\lambda'_{\theta'} \equiv \lambda_{\theta'}$ and $\lambda'_{\theta} \equiv 0$ for all $\theta \neq \theta'$. Then (3.6) and (3.7) are true for likelihood ratio vector λ' and action a'_2 .

Step 2: Let

$$u_1(\theta, a_1, a_2) \equiv \mathbf{1}\{\theta = \theta^*, a_1 = a_1^*, a_2 = a_2^*\}. \quad (3.8)$$

By definition, type θ^* 's commitment payoff from a_1^* is 1. On the equilibrium path, type θ^* plays a different pure action in each period from period 0 to $|A_1| - 1$, in order to separate from all pure-strategy commitment types. Starting from period $|A_1|$, there exists an integer k^* specified by the end of Step 3 such that type θ^* 's action rotates every $k^* + 1$ periods: he plays a_1^* for k^* periods and then plays some action $a_1 \neq a_1^*$ for one period.

I construct equilibrium strategies for strategic types other than θ^* . Find $\lambda' \in \mathbb{R}_+^m$ and $a'_2 \neq a_2^*$ according to Step 1. Since inequality (3.7) is strict, there exists $\epsilon > 0$ such that:

$$u_2(\phi_{a_1^*}, a_1^*, a'_2) - u_2(\phi_{a_1^*}, a_1^*, a_2^*) + (1 - \epsilon) \sum_{\theta \in \Theta} \lambda'_{\theta} (u_2(\theta, a_1^*, a'_2) - u_2(\theta, a_1^*, a_2^*)) > 0. \quad (3.9)$$

For every $\tilde{\theta} \neq \theta^*$, with probability $(\lambda_{\tilde{\theta}} - \lambda'_{\tilde{\theta}}) / \lambda_{\tilde{\theta}}$, strategic type $\tilde{\theta}$ plays $a'_1 \neq a_1^*$ in every period; with probability $(1 - \epsilon) \lambda'_{\tilde{\theta}} / \lambda_{\tilde{\theta}}$, strategic type $\tilde{\theta}$ plays a_1^* in every period. For every $\alpha_1 \in \mathcal{A}_1^*$ that is nontrivially mixed, strategic type $\tilde{\theta}$ plays strategy σ_{α_1} with probability $\frac{\epsilon \lambda'_{\tilde{\theta}}}{k} / \lambda_{\tilde{\theta}}$, where $k \in \mathbb{N}$ is the number of nontrivially mixed actions in $\mathcal{A}_1^*$. I will specify σ_{α_1} in the next paragraph. If $k = 0$, then one can set $\epsilon = 0$.

Now I describe strategy σ_{α_1} . Since $\mathcal{A}_1^*$ is a finite set, there exists $\eta > 0$ such that $\max_{\alpha_1 \in \mathcal{A}_1^* \setminus \{a_1^*\}} \alpha_1[a_1^*] < 1 - \eta$, where $\alpha_1[a_1^*]$ is the probability α_1 attaches to a_1^* , and furthermore, given that inequality (3.6) is strict, one can find $\eta > 0$ small enough such that

$$\sum_{\theta \in \Theta} \lambda'_{\theta} u_2(\theta, \alpha'_1, a'_2) > \sum_{\theta \in \Theta} \lambda'_{\theta} u_2(\theta, \alpha'_1, a_2^*), \quad (3.10)$$

for every $\alpha'_1 \in \Delta(A_1)$ satisfying $\alpha'_1[a_1^*] \geq 1 - \eta$. At every h^t that occurs with positive probability under type θ^* 's equilibrium strategy, $\sigma_{\alpha_1}(h^t) = \alpha_1$. At every h^t that occurs with zero probability under type θ^* 's equilibrium strategy, $\sigma_{\alpha_1}(h^t) = \hat{\alpha}_1(\alpha_1) \in \Delta(A_1)$ where

$$\hat{\alpha}_1(\alpha_1) \equiv (1 - \frac{\eta}{2}) \alpha_1^* + \frac{\eta}{2} \tilde{\alpha}_1(\alpha_1) \quad (3.11)$$

and $\tilde{\alpha}_1(\alpha_1) \in \Delta(A_1)$ attaches probability 0 to a_1^* and probability $\frac{\alpha_1[a_1]}{1 - \alpha_1[a_1]}$ to every $a_1 \neq a_1^*$.

Step 3: I verify type θ^* 's incentive constraints by deriving a *uniform upper bound* on his continuation payoff *after his first deviation*, i.e., at a history h^t such that h^t occurs with zero probability under type θ^* 's equilibrium strategy but its immediate predecessor h^{t-1} occurs with positive probability under type θ^* 's equilibrium strategy. For every $\alpha_1 \in \mathcal{A}_1^* \setminus \{a_1^*\}$, let $\mu_t(\theta(\alpha_1))$ be the probability that player 1 is strategic and follows strategy σ_{α_1} . Let $\beta_t(\alpha_1) \equiv \mu_t(\theta(\alpha_1)) / \mu_t(\alpha_1)$. The value of $\beta_t(\alpha_1)$ equals $\beta_0(\alpha_1)$ at period- t histories that occur with positive probability under type θ^* 's equilibrium strategy since σ_{α_1} asks player 1 to play α_1 at those histories.

Next, consider histories that occur with zero probability under type θ^* 's equilibrium strategy. Recall that $\max_{\alpha_1 \in \mathcal{A}_1^* \setminus \{a_1^*\}} \alpha_1[a_1^*] < 1 - \eta$, so when a_1^* is observed in period t , $\beta_{t+1}(\alpha_1) \geq \frac{1-\eta/2}{1-\eta} \beta_t(\alpha_1)$ for every $\alpha_1 \in \mathcal{A}_1^* \setminus \{a_1^*\}$. Let $\kappa \equiv 1 - \min_{\alpha_1 \in \mathcal{A}_1^* \setminus \{a_1^*\}} \alpha_1[a_1^*]$. If $a_1 \neq a_1^*$ is observed in period t , then the definition of $\tilde{\alpha}_1(\alpha_1)$ implies that $\beta_{t+1}(\alpha_1) \geq \frac{\eta}{2\kappa} \beta_t(\alpha_1)$. Let $\bar{k} \equiv \left\lceil \log \frac{2\kappa}{\eta} / \log \frac{1-\eta/2}{1-\eta} \right\rceil$. For every $\alpha_1 \in \mathcal{A}_1^*$, let $\bar{\beta}(\alpha_1)$ be the smallest $\beta \in \mathbb{R}_+$ such that:

$$u_2(\phi_{\alpha_1}, \alpha_1, a'_2) + \beta \sum_{\theta \in \Theta} \lambda'_\theta u_2(\theta, \hat{\alpha}_1(\alpha_1), a'_2) \geq u_2(\phi_{\alpha_1}, \alpha_1, a_2^*) + \beta \sum_{\theta \in \Theta} \lambda'_\theta u_2(\theta, \hat{\alpha}_1(\alpha_1), a_2^*) \quad (3.12)$$

Let $\bar{\beta} \equiv 2 \max_{\alpha_1 \in \mathcal{A}_1^* \setminus \{a_1^*\}} \bar{\beta}(\alpha_1)$ and $\underline{\beta} \equiv \min_{\alpha_1 \in \mathcal{A}_1^* \setminus \{a_1^*\}} \frac{\mu(\theta(\alpha_1))}{\mu(\alpha_1)}$. Let $T_1 \equiv \left\lceil \log \frac{\bar{\beta}}{\underline{\beta}} / \log \frac{1-\eta/2}{1-\eta} \right\rceil$. At any history right after type θ^* 's first deviation, $\beta_t(\alpha_1) \geq \underline{\beta}$ for all $\alpha_1 \in \mathcal{A}_1^* \setminus \{a_1^*\}$. After player 2 observes a_1^* for T_1 consecutive periods, a_2^* is strictly dominated by a'_2 until some $a'_1 \neq a_1^*$ is observed. Moreover, every time player 1 plays some $a'_1 \neq a_1^*$, he can induce outcome (a_1^*, a_2^*) for at most $\bar{k}$ consecutive periods before a_2^* is strictly dominated by a'_2 again. Therefore, type θ^* 's continuation payoff after his first deviation is at most $(1 - \delta^{T_1}) + \delta^{T_1} \left\{ (1 - \delta^{\bar{k}-1}) + \delta^{\bar{k}} (1 - \delta^{\bar{k}-1}) + \delta^{2\bar{k}} (1 - \delta^{\bar{k}-1}) + \dots \right\}$, which converges to $\frac{\bar{k}}{1+\bar{k}}$ as $\delta \rightarrow 1$.

Let $k^* \equiv 2\bar{k}$. When $\delta \rightarrow 1$, type θ^* 's payoff at any on-path history converges to $\frac{2\bar{k}}{2\bar{k}+1}$, which is strictly greater than $\frac{\bar{k}}{1+\bar{k}}$. This verifies type θ^* 's incentive to play his equilibrium strategy.

3.4. Proof of Statement 3: bounding equilibrium payoffs

For every $\psi \equiv (\psi_\theta)_{\theta \in \Theta} \in \mathbb{R}_+^m$ and $\chi > 0$, let

$$\underline{\Delta}(\psi, \chi) \equiv \left\{ (\tilde{\lambda}_\theta)_{\theta \in \Theta} \in \mathbb{R}_+^m \mid \sum_{\theta \in \Theta} \tilde{\lambda}_\theta / \psi_\theta < \chi \right\}. \quad (3.13)$$

Abusing notation, let $\mu(h^t)$ and $\lambda(h^t)$ be player 2's posterior belief and the likelihood ratio vector with respect to α_1^* at h^t . Let $\mathcal{A}_1^* \equiv \text{supp}(\alpha_1^*)$. For every $\sigma_\theta : \mathcal{H} \rightarrow \Delta(A_1)$ and $\sigma_2 : \mathcal{H} \rightarrow \Delta(A_2)$, let $\mathcal{P}^{(\sigma_\theta, \sigma_2)}$ be the probability measure over $\mathcal{H}$ induced by $(\sigma_\theta, \sigma_2)$, let $\mathcal{H}^{(\sigma_\theta, \sigma_2)}$ be the set of histories that occur with positive probability under $\mathcal{P}^{(\sigma_\theta, \sigma_2)}$, and let $\mathbb{E}^{(\sigma_\theta, \sigma_2)}$ be the expectation induced by probability measure $\mathcal{P}^{(\sigma_\theta, \sigma_2)}$.

Proposition 1. Suppose $\lambda \in \underline{\Delta}(\psi, \chi)$. For every $\epsilon > 0$, there exist $T \in \mathbb{N}$ and $\underline{\delta} \in (0, 1)$ such that for every $\delta > \underline{\delta}$ and every equilibrium $\sigma \equiv ((\sigma_\theta)_{\theta \in \Theta}, \sigma_2)$, we can find a deviation $\tilde{\sigma}_\theta : \mathcal{H} \rightarrow \Delta(A_1^*)$ and a continuous function $\beta(\delta)$ satisfying $\lim_{\delta \rightarrow 1} \beta(\delta) = 0$ such that:

$$\lambda(h^t) \in \underline{\Delta}(\psi, \chi + \epsilon) \quad \text{for every } h^t \in \mathcal{H}^{(\tilde{\sigma}_\theta, \sigma_2)}, \quad (3.14)$$

$$\mathcal{P}^{(\tilde{\sigma}_\theta, \sigma_2)} \left(\left| \sum_{t=0}^{\infty} (1 - \delta) \delta^t \mathbf{1}\{h_t^\infty = a_1\} - \alpha_1^*(a_1) \right| < \epsilon \text{ for every } a_1 \in A_1 \right) > 1 - \beta(\delta), \quad (3.15)$$

$$\mathbb{E}^{(\tilde{\sigma}_\theta, \sigma_2)} \left[\# \left\{ t \in \mathbb{N} \mid \|\alpha_1^* - \alpha_1(\cdot|h^t)\| > \epsilon \right\} \right] < T. \quad (3.16)$$

Proof. I show Proposition 1 in three steps.

Step 1: Let $\mathcal{P}^{(\alpha_1^*, \sigma_2)}$ be the probability measure over $\mathcal{H}$ when player 1 plays α_1^* in every period and player 2 plays according to σ_2 . Let $\chi(h^t) \equiv \sum_{i=1}^m \frac{\lambda_i(h^t)}{\psi_i}$. By definition, $\lambda \in \underline{\Delta}(\psi, \chi)$ if and only if $\chi(h^0) < \chi$. Let $\{\mathcal{F}^t\}_{t \in \mathbb{N}}$ be the filtration induced by the public history. Since $\{\lambda_i(h^t), \mathcal{P}^{(\alpha_1^*, \sigma_2)}, \mathcal{F}^t\}_{t \in \mathbb{N}}$ is a non-negative supermartingale for every $i \in \{1, 2, \dots, m\}$, $\{\chi_t, \mathcal{P}^{(\alpha_1^*, \sigma_2)}, \mathcal{F}^t\}_{t \in \mathbb{N}}$ is also a non-negative supermartingale. For every $a < b$, let $U(a, b)$ be the number of upcrossings from a to b .¹² The Doob's Upcrossing Inequality implies:

$$\mathcal{P}^{(\alpha_1^*, \sigma_2)} \left\{ U(\chi, \chi + \epsilon) = 0 \right\} \geq \frac{\epsilon}{\chi + \epsilon}. \quad (3.17)$$

Let $\tilde{\mathcal{H}}^\infty$ be the set of histories such that $\chi_t < \chi + \epsilon$ for every $t \in \mathbb{N}$. According to (3.17), $\tilde{\mathcal{H}}^\infty$ occurs with probability at least $\frac{\epsilon}{\chi + \epsilon}$ under probability measure $\mathcal{P}^{(\alpha_1^*, \sigma_2)}$.

Let $\tilde{\mathcal{P}}$ be a probability measure defined as $\tilde{\mathcal{P}}(E) \equiv \frac{\mathcal{P}^{(\alpha_1^*, \sigma_2)}(E \cap \tilde{\mathcal{H}}^\infty)}{\mathcal{P}^{(\alpha_1^*, \sigma_2)}(\tilde{\mathcal{H}}^\infty)}$. I construct a strategy $\tilde{\sigma}_\theta$ such that when player 1 uses $\tilde{\sigma}_\theta$ and player 2s use their equilibrium strategy, the induced probability measure over histories is $\tilde{\mathcal{P}}$. For every h^t such that $\chi(h^t) < \chi + \epsilon$, let $A_1(h^t) \subset \text{supp}(\alpha_1^*)$ be such that $a_1 \in A_1(h^t)$ if and only if $\chi(h^t, a_1) < \chi + \epsilon$. The set $A_1(h^t)$ is not empty since $\chi(h^t) < \chi + \epsilon$ and $\{\chi_t, \mathcal{P}^{(\alpha_1^*, \sigma_2)}, \mathcal{F}^t\}_{t \in \mathbb{N}}$ is a supermartingale, and moreover, $h^t \in \tilde{\mathcal{H}}^\infty$ if and only if for every $s < t$, $h^s \in \tilde{\mathcal{H}}^\infty$ and player 1's action in period s belongs to $A_1(h^s)$. Let $\tilde{\mathcal{P}}(\cdot|h^t)$ be the probability measure induced by $\tilde{\mathcal{P}}$ conditional on the history being h^t , which is well-defined for every $h^t \in \tilde{\mathcal{H}}^\infty$. Suppose $\tilde{\sigma}_\theta$ is such that at every h^t satisfying $\chi(h^t) < \chi + \epsilon$, player 1 plays a_1 with zero probability if $a_1 \notin A_1(h^t)$, and plays a_1 with probability $\tilde{\mathcal{P}}(a_{1,t} = a_1|h^t)$ if $a_1 \in A_1(h^t)$; at every h^t such that $\chi(h^t) \geq \chi + \epsilon$, $\tilde{\sigma}_\theta$ can be arbitrary. By construction, $\tilde{\sigma}_\theta$ induces probability measure $\tilde{\mathcal{P}}$.

Step 2: I show that when δ is close enough to 1, there exists a subset of $\mathcal{H}^\infty$ that occurs with probability close to 1 under probability measure $\mathcal{P}^{(\alpha_1^*, \sigma_2)}$, such that the discounted frequency of every $a_1 \in A_1$ is close to $\alpha_1^*(a_1)$. For every $a_1 \in A_1$, let $\{X_t\}$ be a sequence of i.i.d. random variables such that:

$$X_t = \begin{cases} 1 & \text{when } a_{1,t} = a_1 \\ 0 & \text{otherwise.} \end{cases}$$

Under probability measure $\mathcal{P}^{(\alpha_1^*, \sigma_2)}$, $X_t = 1$ occurs with probability $\alpha_1^*(a_1)$. Let $n \equiv |A_1|$.

¹² For real numbers $a < b$, the number of upcrossings of $\{X_t\}_{t \in \mathbb{N}}$ from a to b is the maximum nonnegative integer n such that there exist $s_k, t_k \in \mathbb{N}$ such that $s_1 < t_1 < s_2 < \dots < s_n < t_n$ and $X_{s_k} < a < b < X_{t_k}$ for every $k \in \{1, 2, \dots, n\}$.

Lemma 3.1. *For every $\eta > 0$, there exists $\delta^* \in (0, 1)$, such that for all $\delta \in (\delta^*, 1)$,*

$$\limsup_{\delta \rightarrow 1} \mathcal{P}^{(\alpha_1^*, \sigma_2)} \left(\left| \sum_{t=0}^{+\infty} (1-\delta)^t X_t - \alpha_1^*(a_1) \right| \geq \eta \right) \leq \frac{\eta}{n}. \quad (3.18)$$

The proof of this lemma is standard, which is relegated to Online Appendix C. According to Lemma 3.1, for every $a_1 \in A_1$ and $\eta > 0$, there exists $\delta^* \in (0, 1)$, such that for every $\delta > \delta^*$, there exists $\mathcal{H}_{\eta, a_1}^\infty(\delta) \subset \mathcal{H}^\infty$, such that the discounted frequency of a_1 is η -close to $\alpha_1^*(a_1)$ for every $h^\infty \in \mathcal{H}_{\eta, a_1}^\infty(\delta)$, and $\mathcal{P}^{(\alpha_1^*, \sigma_2)}(\mathcal{H}_{\eta, a_1}^\infty(\delta)) \geq 1 - \frac{\eta}{n}$. Let $\mathcal{H}_\eta^\infty(\delta) \equiv \bigcap_{a_1 \in A_1} \mathcal{H}_{\eta, a_1}^\infty(\delta)$, we have $\mathcal{P}^{(\alpha_1^*, \sigma_2)}(\mathcal{H}_\eta^\infty(\delta)) \geq 1 - \eta$.

Step 3: Recall that $\tilde{\mathcal{H}}^\infty$ is the set of histories such that $\chi_t < \chi + \epsilon$ for every $t \in \mathbb{N}$, which occurs with probability at least $\frac{\epsilon}{\chi + \epsilon}$. Therefore, the probability of $\tilde{\mathcal{H}}^\infty \equiv \tilde{\mathcal{H}}^\infty \cap \mathcal{H}_\eta^\infty(\delta)$ conditional on $\tilde{\mathcal{H}}^\infty$ is at least $1 - \frac{\eta(\chi + \epsilon)}{\epsilon}$. Intuitively, $\tilde{\mathcal{H}}^\infty$ is the event in which $\chi_t < \chi + \epsilon$ for every $t \in \mathbb{N}$ and the discounted frequency of every player 1's pure action is η -close to its probability in α_1^* . Since η is arbitrarily close to 0 as $\delta \rightarrow 1$, $1 - \frac{\eta(\chi + \epsilon)}{\epsilon}$ can be arbitrarily close to 1, which means that the probability that the discounted frequency of every action being close to its probability in the mixed commitment action is arbitrarily close to 1 conditional on $\tilde{\mathcal{H}}^\infty$.

Let $d(\cdot \| \cdot)$ denote the Kullback-Leibler divergence between two distributions. Gossner (2011)'s result implies that:

$$\mathbb{E}^{(\alpha_1^*, \sigma_2)} \left[\sum_{\tau=0}^{+\infty} d(\alpha_1^* \| \alpha_1(\cdot | h^\tau)) \right] \leq -\log \mu(\alpha_1^*). \quad (3.19)$$

Since the Kullback-Leibler divergence must be non-negative, Markov Inequality implies that:

$$\mathbb{E}^{(\alpha_1^*, \sigma_2)} \left[\sum_{\tau=0}^{+\infty} d(\alpha_1^* \| \alpha(\cdot | h^\tau)) \middle| \tilde{\mathcal{H}}^\infty \right] \leq -\frac{(\chi + \epsilon) \log \mu(\alpha_1^*)}{\epsilon}. \quad (3.20)$$

Recall that $\tilde{\sigma}_\theta$ is strategic-type player 1's strategy that induces probability measure $\tilde{\mathcal{P}}$, i.e., the probability measure such that $\tilde{\mathcal{P}}(E) \equiv \frac{\mathcal{P}^{(\alpha_1^*, \sigma_2)}(E \cap \tilde{\mathcal{H}}^\infty)}{\mathcal{P}^{(\alpha_1^*, \sigma_2)}(\tilde{\mathcal{H}}^\infty)}$. If player 1 deviates to strategy $\tilde{\sigma}_\theta$, then the expected number of periods in which $d(\alpha_1^* \| \alpha(\cdot | h^t)) > \epsilon^2/2$ is at most $T \equiv \left\lceil -\frac{2(\chi + \epsilon) \log \mu(\alpha_1^*)}{\epsilon^3} \right\rceil$. The Pinsker's inequality implies that the expected number of periods where $\|\alpha_1^* - \alpha(\cdot | h^t)\| > \epsilon$ is at most T . $\square$

Proposition 1 implies the following corollary:

Corollary 1. *If $\lambda(\alpha_1^*) \in \underline{\Delta}(\theta^*, \alpha_1^*)$ and δ is large, then for every equilibrium σ , there exists a deviation for strategic type θ^* , denoted by $\tilde{\sigma}_{\theta^*} : \mathcal{H} \rightarrow \Delta(A_1^*)$ such that when player 1 uses $\tilde{\sigma}_{\theta^*}$ and player 2s use their equilibrium strategy,*

1. *With probability 1, player 2's posterior likelihood ratio vector in every period belongs to $\Delta(\psi, 1 - \epsilon)$.*
2. *With probability close to 1, the discounted frequency of every $a_1 \in A_1$ is approximately $\alpha_1^*(a_1)$.*

3. In all but a bounded number of periods, player 2's prediction about player 1's action is close to α_1^*

This is because $\underline{\Delta}(\theta^*, \alpha_1^*)$ coincides with $\underline{\Delta}(\psi, \chi)$ when $\chi = 1$, $\psi_\theta \equiv \psi_\theta^*$ for every $\theta \in \Theta_{(\theta^*, \alpha_1^*)}^b$, and $\psi_\theta \equiv +\infty$ for every $\theta \notin \Theta_{(\theta^*, \alpha_1^*)}^b$. Let $\epsilon \equiv \frac{1}{2} \left(1 - \sum_{\theta \in \Theta} \frac{\lambda_\theta(\alpha_1^*)}{\psi_\theta} \right)$. Since $\lambda(\alpha_1^*) \in \underline{\Delta}(\theta^*, \alpha_1^*)$, we have $\epsilon > 0$.

Corollary 1 does not directly imply that type θ^* can guarantee payoff $v_{\theta^*}(\alpha_1^*)$ for every u_1 in every equilibrium. This is because due to the potential correlation between player 1's action and the state θ , player 2s may not have incentives to play a_2^* despite $\lambda(\alpha_1^*)$ belongs to the interior of $\underline{\Delta}(\theta^*, \alpha_1^*)$ and player 1's average action is close to α_1^* . The remaining proof proceeds in two steps, which I summarize below before presenting the details.

1. Suppose all entries of $\lambda(\alpha_1^*)$ except for at most one are sufficiently close to 0, then player 2 has a strict incentive to play a_2^* when player 1's average action is close to α_1^* . Let Λ^0 be the set of type distributions with this feature. One can then directly apply Corollary 1 to establish inequality (2.5).
2. If player 1's average action is close to α_1^* but player 2 does not have a strict incentive to play a_2^* , then different types of player 1's actions at that history must be significantly different. This implies that player 1's action at that history must be informative about his type, in which case he can pick a particular action that induces player 2 to learn. I show that for every $\lambda(\alpha_1^*) \in \underline{\Delta}(\theta^*, \alpha_1^*)$, there exists an integer $K(\lambda)$ and a strategy for type θ^* such that if type θ^* follows this strategy, then after at most $K(\lambda)$ such periods, player 2's belief about his type belongs to Λ^0 , which concludes the proof.

Step 1: When at most one entry of $\lambda(\alpha_1^*)$ is large In order to simplify notation, I index the set of states by $\{1, 2, \dots, m\}$ instead of $\theta \in \Theta$ when doing summation. For every $\xi > 0$, a likelihood ratio vector λ is of 'size ξ ' if there exists $\tilde{\psi} \equiv (\tilde{\psi}_1, \dots, \tilde{\psi}_m) \in \mathbb{R}_+^m$ such that $\tilde{\psi}_i \in (0, \psi_i)$ for all i and moreover,

$$\lambda \in \left\{ \tilde{\lambda} \in \mathbb{R}_+^m \mid \sum_{i=1}^m \tilde{\lambda}_i / \tilde{\psi}_i < 1 \right\} \subset \left\{ \tilde{\lambda} \in \mathbb{R}_+^m \mid \#\{i \mid \tilde{\lambda}_i \leq \xi\} \geq m - 1 \right\}. \quad (3.21)$$

Intuitively, λ is of size ξ if there exists a downward sloping hyperplane such that every non-negative likelihood ratio vector below this hyperplane has at least $m - 1$ entries no larger than ξ . By definition, for every $\xi' \in (0, \xi)$, if λ is of size ξ' , then it is also of size ξ . Proposition 2 establishes (2.5) when λ is of size ξ for ξ small enough.

Proposition 2. *There exists $\xi > 0$, s.t. $\liminf_{\delta \rightarrow 1} v_{\theta^*}(\delta) \geq v_{\theta^*}(\alpha_1^*)$ for every λ of size ξ .*

Proof. Let $\Omega \equiv \Theta \cup \mathcal{A}_1^*$ stand for the set of types, where every element in Θ stands for a strategic type and every element in $\mathcal{A}_1^*$ stands for a commitment type. I use $\omega \in \Omega$ to denote a typical element of Ω . Let $\alpha_1(\cdot | h^t, \omega_i) \in \Delta(A_1)$ be the equilibrium action of type ω_i at history h^t . Let

$$B_{i, a_1}(h^t) \equiv \lambda_i(h^t) \left(\alpha_1^*(a_1) - \alpha_1(a_1 | h^t, \omega_i) \right). \quad (3.22)$$

Let $\alpha_1(\cdot | h^t)$ be the average action expected by player 2 at h^t . For every $\lambda \in \underline{\Delta}(\theta^*, \alpha_1^*)$ and $\epsilon > 0$, there exists $\varepsilon > 0$ such that for every likelihood ratio vector $\tilde{\lambda} \equiv (\tilde{\lambda}_i)_{i=1}^m$ satisfying:

$$\sum_{i=1}^m \tilde{\lambda}_i / \psi_i < \frac{1}{2} \left(1 + \sum_{i=1}^m \lambda_i / \psi_i \right), \quad (3.23)$$

$a_2^*(\theta^*, \alpha_1^*)$ is player 2's strict best reply to every $\{\alpha_1(\cdot|h^t, \omega_i)\}_{i=1}^m$ satisfying the following two conditions

1. $|B_{i,a_1}(h^t)| < \varepsilon$ for all i and a_1 .
2. $\|\alpha_1^* - \alpha_1(\cdot|h^t)\| \leq \epsilon$.

This is because when the prior likelihood ratio vector satisfies (3.23), $a_2^*(\theta^*, \alpha_1^*)$ is player 2's strict best reply when all types of player 1 play α_1^* . When ϵ and ε are both small enough, this strictness cannot be overturned.

According to the Pinsker's Inequality, we know that for every $\epsilon > 0$, $\|\alpha_1^* - \alpha_1(\cdot|h^t)\| \leq \epsilon$ is implied by $d(\alpha_1^* || \alpha_1(\cdot|h^t)) \leq \epsilon^2/2$. Let $\bar{\psi} \equiv \max\{\tilde{\psi}_1, \dots, \tilde{\psi}_m\}$, where $\tilde{\psi}_i$ is given by (3.21). Pick $\epsilon > 0$ and $\xi > 0$ to be small enough such that:

$$\epsilon < \frac{\varepsilon}{2(1 + \bar{\psi})} \quad \text{and} \quad \xi < \frac{\varepsilon}{(m-1)(1 + \varepsilon)}. \quad (3.24)$$

Suppose without loss of generality that $\lambda_i(h^t) \leq \xi$ for all $i \geq 2$, since $\|\alpha_1^* - \alpha_1(\cdot|h^t)\| \leq \epsilon$, we have:

$$\frac{\|(\alpha_1^* - \alpha_1(a_1|h^t, \omega_1))\lambda_1(h^t) + \sum_{i=2}^m (\alpha_1^* - \alpha_1(a_1|h^t, \omega_i))\lambda_i(h^t)\|}{1 + \lambda_1(h^t) + (m-1)\xi} \leq \epsilon.$$

The triangular inequality implies that:

$$\begin{aligned} \|(\alpha_1^* - \alpha_1(a_1|h^t, \omega_1))\lambda_1(h^t)\| &\leq \sum_{i=2}^m \|(\alpha_1^* - \alpha_1(a_1|h^t, \omega_i))\lambda_i(h^t)\| \\ &\quad + \epsilon(1 + \lambda_1(h^t) + (m-1)\xi) \\ &\leq (m-1)\xi + \epsilon(1 + \bar{\psi} + (m-1)\xi) \leq \varepsilon, \end{aligned} \quad (3.25)$$

where the last inequality uses (3.24). Inequality (3.25) implies that $\|B_{1,a_1}(h^t)\| \leq \varepsilon$. As a result, for every λ of size ξ , $a_2^*(\theta^*, \alpha_1^*)$ is player 2's strict best reply at every history h^t satisfying $d(\alpha_1^* || \alpha_1(\cdot|h^t)) \leq \epsilon^2/2$. $\square$

Step 2: When multiple entries of $\lambda(\alpha_1^*)$ are large I apply the conclusion of Proposition 2 in order to establish inequality (2.5) for every $\lambda(\alpha_1^*) \in \underline{\Delta}(\theta^*, \alpha_1^*)$. Recall the definition of $B_{i,a_1}(h^t)$ in (3.22) and recall that A_1^* is defined as the support of mixed commitment action α_1^* . Let $\lambda(h^t, a_1)$ be the posterior likelihood ratio vector after player 2 observes a_1 at h^t , and let $\lambda_i(h^t, a_1)$ be the i th entry of that vector. According to Bayes rule, if player 1 plays $a_1 \in A_1^*$ at h^t , then

$$\lambda_i(h^t) - \lambda_i(h^t, a_1) = \frac{B_{i,a_1}(h^t)}{\alpha_1^*(a_1)} \quad \text{and} \quad \sum_{a_1 \in A_1^*} \alpha_1^*(a_1) (\lambda_i(h^t) - \lambda_i(h^t, a_1)) \geq 0.$$

Let $D(h^t, a_1) \equiv \left(\lambda_i(h^t) - \lambda_i(h^t, a_1) \right)_{i=1}^m \in \mathbb{R}^m$. Suppose $B_{i,a_1}(h^t) \geq \varepsilon$ for some i and $a_1 \in A_1^*$, we have $\|D(h^t, a_1)\| \geq \varepsilon$. Pick $\xi > 0$ small enough in order to meet the requirement in Proposition 2. I define two sequences of subsets of $\underline{\Delta}(\theta^*, \alpha_1^*)$, $\{\Lambda^k\}_{k=0}^\infty$ and $\{\widehat{\Lambda}^k\}_{k=1}^\infty$, as follows:

- Let Λ^0 be the set of likelihood ratio vectors that are of size ξ ,
- For every $k \geq 1$, let $\widehat{\Lambda}^k$ be the set of likelihood ratio vectors that belong to $\underline{\Delta}(\theta^*, \alpha_1^*)$ such that if $\lambda(h^t) \in \widehat{\Lambda}^k$, then either $\lambda(h^t) \in \Lambda^{k-1}$ or for every $\{\alpha_1(\cdot|h^t, \omega_i)\}_{i=1}^m$ such that $\|D(h^t, a_1)\| \geq \varepsilon$ for some $a_1 \in A_1^*$ (i.e., non-trivial learning takes place at h^t), there exists $a'_1 \in A_1^*$ such that the posterior likelihood ratio vector belongs to Λ^{k-1} after player 2 observes a'_1 at h^t .
- Let Λ^k be the set of likelihood ratio vectors in $\underline{\Delta}(\theta^*, \alpha_1^*)$ such that for every $\lambda \in \Lambda^k$, there exists $\tilde{\psi} \equiv (\tilde{\psi}_1, \dots, \tilde{\psi}_m) \in \mathbb{R}_+^m$ such that $\tilde{\psi}_i \in (0, \psi_i)$ for all i and

$$\lambda \in \left\{ \tilde{\lambda} \in \mathbb{R}_+^m \mid \sum_{i=1}^m \tilde{\lambda}_i / \tilde{\psi}_i < 1 \right\} \subset \left(\bigcup_{j=0}^{k-1} \Lambda^j \right) \cup \widehat{\Lambda}^k. \quad (3.26)$$

By construction,

$$\left\{ \tilde{\lambda} \in \mathbb{R}_+^m \mid \sum_{i=1}^m \tilde{\lambda}_i / \tilde{\psi}_i < 1 \right\} \subset \bigcup_{j=0}^k \Lambda^j = \Lambda^k. \quad (3.27)$$

Since $(0, \dots, \psi_i - \nu, \dots, 0) \in \Lambda^0$ for any $i \in \{1, 2, \dots, m\}$ and $\nu \in (0, \psi_i)$, we have $\text{co}(\Lambda^0) = \underline{\Delta}(\theta^*, \alpha_1^*)$. By definition, $\{\Lambda^k\}_{k \in \mathbb{N}}$ is an increasing sequence of sets (in the set inclusion sense) with $\Lambda^k \subset \underline{\Delta}(\theta^*, \alpha_1^*) = \text{co}(\Lambda^k)$ for any $k \in \mathbb{N}$. Hence, it is bounded from above by a compact set. Therefore $\Lambda^\infty \equiv \lim_{k \rightarrow \infty} \bigcup_{j=0}^k \Lambda^j$ is well-defined and Λ^∞ is a subset of $\text{cl}(\underline{\Delta}(\theta^*, \alpha_1^*))$.

The next lemma shows that $\text{cl}(\Lambda^\infty)$ coincides with $\text{cl}(\underline{\Delta}(\theta^*, \alpha_1^*))$.

Lemma 3.2. *We have $\text{cl}(\Lambda^\infty) = \text{cl}(\underline{\Delta}(\theta^*, \alpha_1^*))$.*

Proof. Since $\Lambda^k \subset \underline{\Delta}(\theta^*, \alpha_1^*)$ for every $k \in \mathbb{N}$, we have $\text{cl}(\Lambda^\infty) \subset \text{cl}(\underline{\Delta}(\theta^*, \alpha_1^*))$. The rest of the proof shows the other direction. Suppose by way of contradiction that $\text{cl}(\Lambda^\infty) \subsetneq \text{cl}(\underline{\Delta}(\theta^*, \alpha_1^*))$.

1. Let $\widehat{\Lambda} \subset \underline{\Delta}(\theta^*, \alpha_1^*)$ be such that if $\lambda(h^t) \in \widehat{\Lambda}$, then either $\lambda(h^t) \in \Lambda^\infty$, or for every $\{\alpha_1(\cdot|h^t, \omega_i)\}_{i=1}^m$ such that $\|D(h^t, a_1)\| \geq \varepsilon$ for some $a_1 \in A_1^*$, there exists $a'_1 \in A_1^*$ such that $\lambda(h^t, a'_1) \in \Lambda^\infty$, where $\lambda(h^t, a'_1)$ denotes the posterior likelihood ratio after a'_1 is observed at h^t .
2. Let $\check{\Lambda}$ be the set of likelihood ratio vectors in $\underline{\Delta}(\theta^*, \alpha_1^*)$ such that for every $\lambda \in \check{\Lambda}$, there exists $\tilde{\psi} \equiv (\tilde{\psi}_1, \dots, \tilde{\psi}_m) \in \mathbb{R}_+^m$ such that:

$$\tilde{\psi}_i \in (0, \psi_i^*) \text{ for all } i \text{ and } \lambda \in \left\{ \tilde{\lambda} \in \mathbb{R}_+^m \mid \sum_{i=1}^m \tilde{\lambda}_i / \tilde{\psi}_i < 1 \right\} \subset (\Lambda^\infty \cup \widehat{\Lambda}). \quad (3.28)$$

Since Λ^∞ is defined as the limit of the above operator, so $\text{cl}(\Lambda^\infty) \subsetneq \text{cl}(\underline{\Lambda}(\theta^*, \alpha_1^*))$ implies that either $\check{\Lambda} = \Lambda^\infty$, or $\Xi \cap \check{\Lambda} = \{\emptyset\}$ where

$$\Xi \equiv \text{cl}(\underline{\Lambda}(\theta^*, \alpha_1^*)) \setminus \text{cl}(\Lambda^\infty). \quad (3.29)$$

Set Ξ is convex and has non-empty interior. For every $\varrho > 0$, there exist $x \in \Xi$, $\theta \in (0, \pi/2)$, and a halfspace $H(\chi) \equiv \left\{ \tilde{\lambda} \mid \sum_{i=1}^m \tilde{\lambda}_i / \chi_i < \chi \right\}$ with $\phi > 0$ that satisfy:

1. $\sum_{i=1}^m x_i / \psi_i^* = \chi$.¹³
2. $\partial B(x, r) \cap H(\chi) \cap \underline{\Lambda}(\theta^*, \alpha_1^*) \subset \Lambda^\infty$ for every $r \geq \varrho$.
3. For every $r \geq \rho$ and $y \in \partial B(x, r) \cap \underline{\Lambda}(\theta^*, \alpha_1^*)$, either $y \in \Lambda^\infty$ or $d(y, H(\chi)) > r$.

The second and third property use the presumption that $\text{cl}(\Lambda^\infty)$ is not convex. Suppose $\lambda(h^t) = x$ for some h^t and there exists $a_1 \in A_1^*$ such that $\|D(h^t, a_1)\| \geq \varepsilon$,

- Either $\lambda(h^t, a_1) \in \Lambda^\infty$, in which case $x \in \check{\Lambda}$ but $x \in \Xi$, leading to a contradiction.
- Or $\lambda(h^t, a_1) \notin \Lambda^\infty$. Requirement 3 implies that $d(\lambda(h^t, a_1), H(\chi)) > \varepsilon$. On the other hand,

$$\sum_{a'_1 \in A_1^*} \alpha_1^*(a'_1) \lambda_i(h^t, a'_1) \leq \lambda_i(h^t) \text{ for every } i. \quad (3.30)$$

Requirement 1 then implies that $\sum_{a'_1 \in A_1^*} \alpha_1^*(a'_1) \lambda_i(h^t, a'_1) \in H(\chi)$, which implies that

$$\sum_{a'_1 \in A_1^*} \alpha_1^*(a'_1) \sum_{i=1}^m \lambda_i(h^t, a'_1) / \psi_i^* \leq \chi. \quad (3.31)$$

According to Requirement 2, $\lambda(h^t, a_1) \notin H(\chi)$. In another word, $\sum_{i=1}^m \lambda_i(h^t, a_1) / \psi_i^* > \chi + \varepsilon\kappa$ for some $\kappa > 0$. Let $\rho \equiv \frac{\varepsilon\kappa}{2} \min_{a_1 \in A_1^*} \{\alpha_1^*(a_1)\}$. Inequality (3.30) implies the existence of $a'_1 \in A_1^* \setminus \{a_1\}$ such that the likelihood ratio vector after observing a'_1 at h^t belongs to $H(\chi) \cap B(x, \rho)$. Requirement 2 then implies that $x = \lambda(h^t) \in \check{\Lambda}$. Since $x \in \Xi$, this leads to a contradiction and validates Lemma 3.2. $\square$

Lemma 3.2 implies that for every prior likelihood ratio vector $\lambda(\alpha_1^*) \in \underline{\Lambda}(\theta^*, \alpha_1^*)$, there exists $K \in \mathbb{N}$ such that $\lambda(\alpha_1^*) \in \Lambda^K$. Statement 3 of Theorem 1 can then be shown by induction on K , which is finite for every $\lambda(\alpha_1^*) \in \underline{\Lambda}(\theta^*, \alpha_1^*)$. The case where $K = 0$ is implied by Proposition 2. Suppose the statement applies to every $K \leq K^* - 1$, let us consider the case where $\lambda(\alpha_1^*) \in \Lambda^{K^*}$. According to the definition of Λ^{K^*} , there exists a strategy for type θ^* such that in every period where (1) $d(\alpha_1^* \|\alpha_1(\cdot|h^t)) < \varepsilon^2/2$ and (2) player 2 does not have a strict incentive to play $a_2^*(\theta^*, \alpha_1^*)$, there exists $a'_1 \in A_1^*$ such that player 2's posterior likelihood ratio vector belongs to Λ^{K^*-1} after observing a'_1 in that period. After the posterior likelihood ratio vector belongs to Λ^{K^*-1} , the induction hypothesis implies that type θ^* can secure $v_{\theta^*}(\alpha_1^*)$ when δ is close to 1.

¹³ Since we index the states via $i \in \{1, 2, \dots, m\}$ in this proof, I replace ψ_θ^* defined in (3.4) by ψ_i^* .

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jet.2022.105438>.

References

- Deb, Joyee, Ishii, Yuhta, 2021. Reputation Building under Uncertain Monitoring. Working Paper.
- Ekmekci, Mehmet, Maestri, Lucas, 2019. Reputation and Screening in a Noisy Environment with Irreversible Actions. Working Paper.
- Ekmekci, Mehmet, Gorno, Leandro, Maestri, Lucas, Sun, Jian, Wei, Dong, 2021. Learning from Manipulable Signals. Working Paper.
- Fudenberg, Drew, Levine, David, 1989. Reputation and equilibrium selection in games with a patient player. *Econometrica* 57 (4), 759–778.
- Fudenberg, Drew, Levine, David, 1992. Maintaining a reputation when strategies are imperfectly observed. *Rev. Econ. Stud.* 59 (3), 561–579.
- Gossner, Olivier, 2011. Simple bounds on the value of a reputation. *Econometrica* 79 (5), 1627–1641.
- Pei, Harry, 2020. Reputation effects under interdependent values. *Econometrica* 88 (5), 2175–2202.
- Yang, Geyu, 2019. Robustness of Reputation Effects under Uncertain Monitoring. Working Paper.