Reputation Effects under Interdependent Values

Harry PEI*

May 16, 2020

Abstract: A patient player privately observes a persistent state, and interacts with an infinite sequence of myopic uninformed players. The patient player is either a strategic type who maximizes his payoff, or one of several commitment types that mechanically plays the same action in every period. I focus on situations in which the uninformed player's best reply to a commitment action depends on the state, and the total probability of commitment types is sufficiently small. I show that the patient player's equilibrium payoff is bounded below his commitment payoff in some equilibria under some of his payoff functions. This is because he faces a trade-off between building his reputation for commitment and signaling favorable information about the state. When players' stage-game payoff functions are *monotone-supermodular*, the patient player receives high payoffs in all states and in all equilibria. Under an additional condition on the state distribution, my reputation model yields a unique prediction on the patient player's equilibrium payoff and on-path behavior.

Keywords: reputation, interdependent values, commitment payoff, robust behavioral prediction

1 Introduction

Economists have long recognized that reputation lends credibility to agents' threats and promises. This intuition has been formalized in a series of works starting with Kreps and Wilson (1982), Milgrom and Roberts (1982), Fudenberg and Levine (1989,1992) and others, who show that having the option to build a reputation dramatically affects a patient individual's gains in long-term relationships. These reputation results are *robust* as they apply across all equilibria,[1] which enable researchers to make sharp predictions in many decentralized markets where there is no mediator helping participants to coordinate on a particular equilibrium.

However, previous works on robust reputation effects exclude settings in which reputation-building agents' private information directly affects their opponents' payoffs. For example, consumers' willingness to pay for a product depends not only on its seller's effort but also on its *quality*. The latter is persistent over time and is the seller's private information. In this environment, the seller's incentive to sustain his reputation for exerting high effort interacts dynamically with his motive to signal high quality. Such interaction introduces new economic

*Northwestern University. I am indebted to Daron Acemoglu, Drew Fudenberg, Juuso Toikka, and Alex Wolitzky for guidance and support, and Jie Bai, Heski Bar-Isaac, Martin Cripps, Dan Clark, Joyee Deb, Miaomiao Dong, Esther Duflo, Mehmet Ekmekci, Yingni Guo, Yuhta Ishii, Qingmin Liu, Shuo Liu, Alessandro Pavan, Bruno Strulovici, Nicolas Vieille, and four referees for helpful comments.

[1] Throughout this paper, a property is *robust* if and only if it is true for all equilibria of a given model.

forces that are omitted by existing reputation models. In addition, existing reputation results deliver robust predictions on players' equilibrium payoffs, but not on their equilibrium *behaviors*.

This paper studies the effects of *interdependent values* on reputation-building players' payoffs and behaviors. In my model, a patient player 1 (e.g., firm) interacts with an infinite sequence of myopic player 2s (e.g., consumers), arriving one in each period and each plays the game only once. Different from existing reputation models, player 1 privately observes the realization of a payoff-relevant state (e.g., quality) that is constant over time and affects both players' stage-game payoffs, in addition to knowing whether he is strategic or committed. If player 1 is strategic, then he maximizes his discounted average payoff. If player 1 is committed, then he mechanically plays the same action (i.e., a *commitment action*) in every period, which can depend on the state. Player 2s observe all the actions taken in the past, but not their predecessors' payoffs.

My results focus on properties that apply to *all* equilibria. I make three conceptual contributions. First, I identify conflicts between reputation building and signaling under interdependent values. Second, I provide sufficient conditions under which the patient player guarantees high returns from building reputations despite facing such conflicts. Third, I show that interdependent values lead to a *disciplinary effect*, which motivates the patient player to sustain his reputation, and results in a unique prediction on his equilibrium behavior.

Theorem 1 shows that for every state, commitment action, and payoff function of player 2's, if player 2's best reply against this commitment action depends on the state, and the probability of commitment types is sufficiently small, then there *exists* a payoff function for player 1 and an equilibrium, in which his equilibrium payoff in the chosen state is strictly bounded below his complete information commitment payoff. This contrasts to Fudenberg and Levine (1989)'s result for private-value models, which says that player 1 receives at least his commitment payoff, regardless of his own payoff function and the probability of commitment types.

Intuitively, in order to motivate player 2s to play their best reply against the commitment action in the chosen state, player 1 needs to both convince them that this commitment action will be played in the future, and to signal them the correct information about the state. These two objectives are in conflict when player 2s believe that player 1 is more likely to play the commitment action in another state, under which her best reply against the commitment action is different. When facing this conflict, player 1 either abandons his reputation, after which he loses the credibility for playing his commitment action in the future; or he signals negative information about the state, after which player 2s do not have incentives to choose player 1's desired best reply even when they are convinced that player 1 will play his commitment action.

Next, I restrict attention to games with *monotone-supermodular* (or *MS*) payoffs. My MS condition requires that states and actions be ranked such that (1) player 1's payoff is strictly decreasing in his action and is strictly increasing in player 2's action; (2) action profiles and states are complements in both players' payoff functions. For example, the following product choice game between a firm and its customers satisfies MS when states and actions are ranked according to high quality $\succ$ low quality, high effort $\succ$ low effort, trust $\succ$ not trust:

high quality	trust	not trust
high effort	$\mathbf{1}, 1$	$-1, 0$
low effort	$2, -1$	$0, 0$

low quality	trust	not trust
high effort	$\mathbf{2/3}, -1$	$-4/3, 0$
low effort	$2, -2$	$0, 0$

I establish robust predictions on player 1's payoff and behavior when he can build a reputation for playing his *highest action*. Theorems 2 and 3 consider two cases separately, depending on player 2's *prior belief about the state*. To yield clear comparisons with Theorem 1, these results allow commitment types to be arbitrarily rare and player 2's best reply against player 1's highest action to depend on the state.

When player 2's prior belief about the state is such that her highest action *best replies against* player 1's highest action,[2] Theorem 2 shows that in every state and in every equilibrium, a patient player 1 receives at least his payoff from the *highest action profile*, which is no less than his commitment payoff from playing his highest action. In the example, it implies that when the probability of high-quality state exceeds $1/2$, a high-quality firm secures payoff 1 and a low-quality firm secures payoff $2/3$ by exerting high effort in every period.

The difference between Theorems 1 and 2 is driven by the MS condition on stage-game payoffs, which is assumed in the latter but not in the former. Intuitively, MS implies that player 1 has a stronger preference towards higher action profiles in higher states. When this stage-game is played *repeatedly*, despite MS *not ruling out* the aforementioned conflict between reputation-building and signaling,[3] it rules out situations in which player 1 *plays his highest action in every period* in a lower state, but not in a higher state. This leads to a *uniform lower bound* on player 2's posterior belief about the state, which applies to all histories such that player 1 has played his highest action in all previous periods. In the example, it implies that player 2's posterior attaches probability more than $1/2$ to the high-quality state if player 1 has exerted high effort in all previous periods. This belief lower bound implies patient player 1's payoff lower bound, since player 2 has a strict incentive to choose player 1's desired best reply once she is convinced that player 1 will play his highest action.

When player 2's belief about the state is such that her highest action *does not best reply against* player 1's highest action, Theorem 3 shows that player 1 has a *unique equilibrium payoff* and a *unique on-path behavior*. The latter also pins down player 2s' beliefs on the equilibrium path. Player 1's unique payoff is strictly lower than his payoff from the highest action profile, but is greater than his minmax payoff. Player 1's unique on-path behavior is characterized by a *cutoff state* (in the example, low-quality state), such that he plays his highest action in every period when the state is above this cutoff, plays his lowest action in every period when the state is below this cutoff, and randomizes between playing his highest action in every period and playing his lowest

[2]The current paper focuses on MS games in which player 2's action choice is binary. Generalizations to games in which player 2 has three or more actions can be found in the appendix of the working paper version (Pei 2018).

[3]This is because player 1's action not only signals the persistent state, but also affects the continuation equilibrium being played through repeated game effects. Online Appendix D.1 constructs an equilibrium in which exerting high effort signals low quality.

action in every period at the cutoff state. His mixing probability is such that player 2 is indifferent between her highest action and her lowest action after observing player 1's highest action in the first period.

The unique behavioral prediction in Theorem 3 contrasts to the private-value reputation game in Fudenberg and Levine (1989) and the case studied by Theorem 2, in which there are multiple equilibria with different on-path behaviors. My behavioral uniqueness result is driven by a novel *disciplinary effect*, that the strategic long-run player is guaranteed to receive a high payoff when pooling with commitment type, and is guaranteed to receive a low payoff after separating from commitment type.

I illustrate this effect using the product choice game. First, the disciplinary effect is absent in the private-value case of Fudenberg and Levine (1989), i.e., when the high-quality state occurs for sure. This is because after separating from the commitment type that exerts high effort, a patient firm's continuation payoff can be anything between its minmax payoff 0 and its commitment payoff 1. This multiplicity in continuation values leads to multiple on-path behaviors. For example, if at a given history, the strategic firm can attain its commitment payoff after separating from the commitment type, then it has a strict incentive to exert low effort at that history; if the firm can only receive its minmax payoff after separation, then it has a strict incentive to exert high effort. A similar intuition applies under the conditions of Theorem 2 (i.e., the probability of high-quality state is above $1/2$), that after separating from the commitment type, there exist equilibria in which the firm attains commitment payoffs in all states, and equilibria in which it receives minmax payoffs in all states.

Under the conditions required by Theorem 3 (i.e., the probability of high-quality state is below $1/2$), exerting low effort signals low firm-quality *in all equilibria*, after which the firm's continuation payoff equals its minmax payoff. To see this, suppose toward a contradiction that exerting low effort signals high firm-quality. Since belief is a martingale, exerting high effort signals low firm-quality, after which consumers' posterior belief about firm-quality is more pessimistic compared to their prior. As a result, there exists at least one state in the support of this posterior in which the strategic firm's continuation payoff equals its minmax payoff. Since exerting high effort is costly, the firm has a strict incentive to deviate to low effort in this state, which strictly increases its stage-game payoff while not lowering its continuation payoff. This leads to a contradiction.

My model unifies the commitment type approach in Fudenberg and Levine (1989, 1992) and Gossner (2011) with the dynamic signaling approach in Bar-Isaac (2003), Kaya (2009), and Lee and Liu (2013). My results deliver novel insights on seller reputations. Theorem 1 implies that conflicts between building reputations for exerting high effort and signaling high quality can have persistent negative effects on firms' profits. It provides an explanation for why patient firms refrain from building reputations in markets with adverse selection. Theorem 2 implies that in markets with MS payoffs, firms can secure high profits in the long run by establishing reputations for exerting high effort, despite doing so may occasionally trigger negative inferences about their qualities. Theorem 3 implies that when consumers are pessimistic about a firm's quality, the firm has a strong incentive to sustain its reputation for exerting high effort, leading to consistent firm behavior in equilibrium.

Theorems 2 and 3 can also be viewed as equilibrium refinements for repeated incomplete information games with interdependent values. Theorem 2 suggests that introducing reputational types can rule out equilibria in which the patient player receives low payoffs. Theorem 3 advances this line of research one-step further by delivering a unique prediction on a patient player's on-path behavior. This contrasts to dynamic signaling games without commitment types and private-value reputation games, both of which have multiple predictions on the patient player's on-path behavior. Since agents' behaviors are more likely to be observed relative to their payoffs, my robust behavioral predictions bring us closer to empirically testing reputation models.

2 Baseline Model

Time is discrete, indexed by $t = 0, 1, 2....$ A long-lived player 1 (he) with discount factor $\delta \in (0, 1)$ interacts with an infinite sequence of short-lived player 2s (she), 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 perfectly persistent private information about a payoff-relevant state $\theta \in \Theta$, and whether he is *strategic* or *committed*. If player 1 is strategic, then he can flexibly choose his actions. If player 1 is committed, then he mechanically follows one of the several *commitment plans*. A typical commitment plan is denoted by $\gamma : \Theta \to A_1$, according to which the committed long-run player plays $\gamma(\theta)$ in every period when the state is θ. Let Γ be the set of possible commitment plans, and let γ^* stand for player 1 being strategic. Let

$$\mu \in \Delta\Big(\Theta \times \underbrace{(\{\gamma^*\} \cup \Gamma)}_{\text{player 1's characteristics}}\Big) \tag{2.1}$$

be player 2's prior belief about player 1's private information, which is a joint distribution of θ and player 1's *characteristics*, namely, whether he is strategic or committed, and if he is committed, which plan in Γ is he following. Let $\phi \in \Delta(\Theta)$ be the marginal distribution of μ on Θ, namely, the *prior distribution of states*. Let

$$A_1^* \equiv \{a_1 \in A_1 | \text{ there exist } \gamma \in \Gamma \text{ and } \theta \in \Theta \text{ such that } \gamma(\theta) = a_1\}, \tag{2.2}$$

be the set of *commitment actions*. Intuitively, an action a_1^* belongs to A_1^* if and only if it is played in some state under some commitment plan. For every $\theta \in \Theta$, player 1 is *strategic type* θ if he is strategic and knows the state is θ. For every $a_1^* \in A_1^*$, player 1 is *commitment type* a_1^* if he is committed and plays a_1^* in every period.

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 1's private history consists of the public history and his type. Player 2's private history coincides with the public history, which means that she cannot observe her predecessors' payoffs.[4] For every $\theta \in \Theta$, let $\sigma_\theta : \mathcal{H} \to \Delta(A_1)$ be strategic type θ's strategy. Let $\sigma_2 : \mathcal{H} \to \Delta(A_2)$ be player 2's strategy.

For $i \in \{1, 2\}$, player i's stage-game payoff in period t is $u_i(\theta, a_{1,t}, a_{2,t})$. This formulation allows for interdependent values since u_2 depends on θ. The solution concept is Bayesian Nash equilibrium (or *equilibrium* for short), which is a strategy profile $\sigma \equiv \big((\sigma_\theta)_{\theta\in\Theta}, \sigma_2\big)$ such that for every $\theta \in \Theta$, σ_θ maximizes the expected value of $\sum_{t=0}^{\infty}(1-\delta)\delta^t u_1(\theta, a_{1,t}, a_{2,t})$; and σ_2 maximizes player 2's stage-game payoff.

I assume that Θ, Γ, A_1, and A_2 are finite sets, $|A_1|, |A_2| \geq 2$, and μ has full support. This implies that an equilibrium exists and every type occurs with strictly positive probability. To simplify the exposition, I assume that the distribution of θ and the distribution of player 1's characteristics are *independent* according to μ.

2.1 Example: Product Choice Game

I introduce a product choice game that fits into markets with two characteristics. First, consumers' willingness to pay depends not only on a firm's *effort* on designing its products, but also on the firm's *quality*, such as the quality of its upstream suppliers. Arguably, the firm has private information about its quality relative to the market. Second, the firm's quality is persistent, and informative signals about quality, other than the firm's observable efforts, are rarely available to consumers or are unlikely to arrive for a long time.

A firm (player 1) privately observes its intrinsic quality $\theta \in \{\theta_h, \theta_l\}$, and in every period, chooses between high effort H and low effort L. Each consumer (player 2) chooses between a *trusting action* T (e.g., purchase) and a *non-trusting action* N (e.g., do not purchase). Players' stage-game payoffs are:

$\theta = \theta_h$	T	N
H	$1, 1$	$-1, 0$
L	$2, -1$	$0, 0$

$\theta = \theta_l$	T	N
H	$1-\eta, -1$	$-1-\eta, 0$
L	$2, -2$	$0, 0$

where $\eta \in (-1, 1)$ is a parameter. When the firm's quality is high, its cost of playing H is 1, and a consumer has an incentive to trust when H is played with probability more than $1/2$. When the firm's quality is low, its cost of playing H is $1+\eta$, and a consumer has no incentive to trust regardless of the firm's action.

The firm is either *strategic* or *committed*. Suppose there is only one commitment plan γ, according to which the firm plays its pure Stackelberg action in every state:

$$\gamma(\theta) \equiv \begin{cases} H & \text{if } \theta = \theta_h \\ L & \text{if } \theta = \theta_l. \end{cases} \tag{2.3}$$

[4] This feature of my model makes reputation building challenging since player 2s can only learn the persistent state from player 1's actions. This assumption is important for Theorem 1. It also highlights the novelty of Theorems 2 and 3.

The set of commitment actions is $A_1^* = \{H, L\}$. The firm has four types: two strategic types, θ_h and θ_l; and two commitment types, H and L. The consumers' prior belief μ is a joint distribution of the state and whether the firm is committed or strategic, from which one can derive the prior state distribution $\phi \in \Delta\{\theta_h, \theta_l\}$.

Remarks: My Theorems 2 and 3 continue to hold under the following variations of my model. First, a sequence of short-lived consumers can be replaced by *a continuum of long-lived consumers*, given that the firm and future consumers can observe the *aggregate distribution* of the consumers' actions in every period, but not the action of each individual consumer. This is because long-lived consumers play their myopic best replies when their individual actions have negligible impact on the aggregate distribution.[5] Second, one can perturb the model by introducing a small fraction of non-strategic buyers who mechanically choose to buy (take action T) in every period. This addresses the practical concern that future buyers cannot observe the seller's actions when no buyer in the current period chooses to buy.

3 Results

I examine properties of a patient player 1's payoff and behavior that apply to *all* equilibria. For every $a_1 \in A_1$, $\theta \in \Theta$, and u_2, let

$$\mathrm{BR}_2(\theta, a_1|u_2) \equiv \arg\max_{a_2 \in A_2} u_2(\theta, a_1, a_2). \tag{3.1}$$

Given u_2 and $a_1^* \in A_1^*$, interdependent values are *nontrivial* under (u_2, a_1^*) if there exist $\theta', \theta'' \in \Theta$ such that:

$$\mathrm{BR}_2(\theta', a_1^*|u_2) \bigcap \mathrm{BR}_2(\theta'', a_1^*|u_2) = \{\varnothing\}. \tag{3.2}$$

Given state $\theta \in \Theta$ and commitment action $a_1^* \in A_1^*$, type θ's *commitment payoff* from a_1^* is given by:

$$v_\theta(a_1^*, u_1, u_2) \equiv \min_{a_2 \in \mathrm{BR}_2(\theta, a_1^*|u_2)} u_1(\theta, a_1^*, a_2). \tag{3.3}$$

Let $\underline{v}_\theta(\delta, \mu, u_1, u_2)$ be type θ's *lowest equilibrium payoff* under parameters (δ, μ, u_1, u_2). I make the following assumption which is satisfied for generic u_2, including player 2's payoff function in the product choice game:

Assumption 1. *For every $\theta \in \Theta$ and $a_1 \in A_1$, $BR_2(\theta, a_1|u_2)$ is a singleton.*

[5] The equivalence between a continuum of long-lived players and a sequence of short-lived players is well-known, see for example, Fudenberg and Levine (2009) and Faingold and Sannikov (2011). The details can be found in a supplementary appendix on the author's webpage https://sites.northwestern.edu/harrypei/research/.

3.1 Reputation Failure in Games with Unrestricted Payoffs

Theorem 1 shows that if commitment types are rare and player 2's best reply against a commitment action depends on the state, then there *exists* u_1 under which patient player 1's lowest equilibrium payoff is strictly bounded below his complete information commitment payoff.

Theorem 1. *For every u_2 that satisfies Assumption 1, $a_1^* \in A_1^*$, $\theta \in \Theta$, and full support $\phi \in \Delta(\Theta)$. If interdependent values are nontrivial under (u_2, a_1^*), then there exist u_1 and $\overline{\varepsilon} > 0$, such that for every prior belief μ which has state distribution ϕ and attaches probability less than $\overline{\varepsilon}$ to all commitment types:*

$$\limsup_{\delta \to 1} \underline{v}_\theta(\delta, \mu, u_1, u_2) < v_\theta(a_1^*, u_1, u_2). \tag{3.4}$$

The proof is in Appendix C. Theorem 1 contrasts to the private-value reputation result in Fudenberg and Levine (1989), which implies that if $a_1^* \in A_1^*$ and player 2's best reply against a_1^* does not depend on the state, then $\liminf_{\delta \to 1} \underline{v}_\theta(\delta, \mu, u_1, u_2) \geq v_\theta(a_1^*, u_1, u_2)$ for every u_1, θ, and full support μ.

Intuitively, player 1's action not only shows his propensity to play a_1^* in the future, but also signals the persistent state. When interdependent values are nontrivial, player 2's belief about state affects her best reply against a_1^*. In order to motivate player 2s to play the action in $\text{BR}_2(\theta, a_1^*|u_2)$, player 1 needs to convince them that a_1^* will be played with high probability, and the state *does not* belong to the following subset:

$$\Big\{\theta' \in \Theta \Big| \text{BR}_2(\theta', a_1^*|u_2) \neq \text{BR}_2(\theta, a_1^*|u_2)\Big\}. \tag{3.5}$$

This is the set of states such that player 2's best reply against a_1^* differs from that under state θ.

A conflict between these two objectives arises when player 2s believe that strategic types outside (3.5) separate from commitment type a_1^*, and those in (3.5) play a_1^* in every period. Under this belief, player 1 cannot pool with commitment type a_1^* while separating away from strategic types in (3.5). Theorem 1 confirms that under some u_1, such a belief arises in equilibrium and negatively affects a patient player's payoff.

The above argument also explains why Theorem 1 applies to *every* full support state distribution ϕ, regardless of the probability it attaches to states in (3.5). This is because under the aforementioned self-fulfilling belief, player 2 has no incentive to play the action in $\text{BR}_2(\theta, a_1^*|u_2)$ as long as the probability of commitment type a_1^* is small relative to that of strategic types who knew that the state belongs to (3.5).

Theorem 1 is applicable to the product choice game for $a_1^* = H$ and $\theta = \theta_h$. This is because u_2 satisfies Assumption 1 and interdependent values are nontrivial under (u_2, H). Let u_1 be the one in the matrices with $\eta \in (-1, 0]$. Fix any full support state distribution ϕ, when the probability of commitment type H is lower than the probability of strategic type θ_l, there exists equilibrium in which strategic type θ_l plays H in every period, and player 2's posterior belief attaches probability more than $1/2$ to state θ_l upon observing action H. As a result, player 2s prefer action N despite being convinced that H will be played with high probability. Strategic type θ_h's payoff in such an equilibrium is 0. The details of the equilibrium construction are in Appendix C.

3.2 Reputation Results in Monotone-Supermodular Games

Motivated by Theorem 1, Theorems 2 and 3 derive robust predictions on player 1's payoff and behavior when u_1 and u_2 satisfy a *monotone-supermodularity* condition (or MS).

Assumption 2 (MS)**.** *There exist a ranking on* Θ*, a ranking on* A_1*, and a ranking on* A_2*, under which:*

1. $u_1(\theta, a_1, a_2)$ *is strictly decreasing in* a_1*, and is strictly increasing in* a_2*.*
2. $u_1(\theta, a_1, a_2)$ *has strictly increasing differences in* θ *and* (a_1, a_2)*.*
3. $u_2(\theta, a_1, a_2)$ *has strictly increasing differences in* a_2 *and* (θ, a_1)*.*[6]

In the product choice game with rankings $\theta_h \succ \theta_l$, $H \succ L$, and $T \succ N$, Assumption 2 is satisfied when $\eta > 0$. I examine a patient player 1's payoff and behavior when he can build a reputation for playing *his highest action*. Let $\overline{a}_i \equiv \max A_i$ and $\underline{a}_i \equiv \min A_i$ be player $i \in \{1, 2\}$'s highest action and lowest action, respectively. According to Assumption 2, strategic type θ's minmax payoff is $u_1(\theta, \underline{a}_1, \underline{a}_2)$. Let

$$\Theta^* \equiv \left\{\theta \in \Theta \middle| u_1(\theta, \overline{a}_1, \overline{a}_2) > u_1(\theta, \underline{a}_1, \underline{a}_2)\right\} \tag{3.6}$$

be the set of states under which $\overline{a}_1$ is *individually rational* for player 1. I focus on games in which Θ^* is nonempty, and player 2's action choice is binary:[7]

Assumption 3. Θ^* *is non-empty, and* $|A_2| = 2$*.*

A reputation for playing $\overline{a}_1$ is *potentially valuable* only when $\overline{a}_1$ is one of the commitment actions, and player 2 has an incentive to choose player 1's desired action $\overline{a}_2$ when she knew that player 1 is committed and plays $\overline{a}_1$ in every period. Formally, for every $a_1^* \in A_1^*$, let $\phi_{a_1^*} \in \Delta(\Theta)$ be the distribution of states *conditional*

[6] Assumption 2 rules out zero-sum games, common interest games, and coordination games more generally. Reputation for commitment is not valuable in zero-sum games. In Online Appendix D.2, I provide an example of common interest game with nontrivial interdependent values, under which player 1's guaranteed payoff is arbitrarily low compared to his pure Stackelberg payoff. My MS condition differs from the *monotone-submodular* condition in Liu (2011) and Liu and Skrzypacz (2014) for two reasons. First, they require the long-run player to have stronger incentives to play low actions when the short-run player's action is higher, which is not required by my condition. Second, the payoff-relevant state is not present in their models, but is crucial for mine.

[7] Binary action games have been a primary focus of the reputation literature, examples of which include Mailath and Samuelson (2001), Ekmekci (2011), and Liu (2011). Extensions of Theorems 2 and 3 to games with $|A_2| \geq 3$ can be found in Pei (2018).

on player 1 being commitment type a_1^*. This can be derived from player 2's prior belief μ. Let

$$\mathrm{BR}_2(\phi_{a_1^*}, a_1^*|u_2) \equiv \arg\max_{a_2 \in A_2} \Big\{ \sum_{\theta\in\Theta} \phi_{a_1^*}(\theta) u_2(\theta, a_1^*, a_2) \Big\}$$

be the set of player 2's pure best replies against commitment type a_1^*.

Assumption 4. $\overline{a}_1 \in A_1^*$ *and* $BR_2(\phi_{\overline{a}_1}, \overline{a}_1|u_2) = \{\overline{a}_2\}$.

Assumptions 3 and 4 are satisfied in the product choice game under the aforementioned rankings over states and actions. This is because $\theta_h \in \Theta^*$, the set of commitment actions is $A_1^* = \{H, L\}$, which includes the highest action H. According to (2.3), ϕ_H attaches probability 1 to state θ_h, which implies that $\mathrm{BR}_2(\phi_H, H|u_2) = \{T\}$.

Theorems 2 and 3 focus on opposite conditions on the prior state distribution $\phi \in \Delta(\Theta)$. In particular, ϕ is *optimistic* if:

$$\overline{a}_2 \in \arg\max_{a_2\in A_2} \Big\{ \sum_{\theta\in\Theta^*} \phi(\theta) u_2(\theta, \overline{a}_1, a_2) \Big\}, \tag{3.7}$$

and ϕ is *pessimistic* otherwise. In the product choice game, $\Theta^* = \{\theta_h, \theta_l\}$, and ϕ is optimistic if it attaches probability more than $1/2$ to state θ_h. In general, ϕ is optimistic if it attaches high enough probability to high states, such that $\overline{a}_2$ best replies against $\overline{a}_1$ when $\overline{a}_1$ is played in all states under which it is individually rational. When the probability of commitment types is sufficiently small, ϕ being optimistic is equivalent to the *existence of equilibrium* under which every strategic type $\theta \in \Theta^*$ receives payoff at least $u_1(\theta, \overline{a}_1, \overline{a}_2)$. Theorem 2 shows that a patient player 1 receives at least this payoff in all states and in *all equilibria*:

Theorem 2. *If ϕ is optimistic, and the game satisfies Assumptions 1, 2, 3, and 4, then for every $\theta \in \Theta$:*

$$\liminf_{\delta\to 1} \underline{v}_\theta(\delta, \mu, u_1, u_2) \geq \max\{u_1(\theta, \overline{a}_1, \overline{a}_2), u_1(\theta, \underline{a}_1, \underline{a}_2)\}.^8 \tag{3.8}$$

The proofs of Theorem 2 and the next theorem are in Appendix D and Online Appendices A and B. Since the RHS of (3.8) is weakly greater than type θ's commitment payoff from $\overline{a}_1$, player 1 is guaranteed to receive *at least* his commitment payoff from $\overline{a}_1$. This difference between Theorems 1 and 2 is driven by the MS condition on stage-game payoffs, under which player 1 has a stronger preference toward higher action profiles in higher states. It implies that for each of player 2s' equilibrium strategy σ_2, suppose there exists θ such that playing $\overline{a}_1$ in every period best replies against σ_2 in state θ, then in every state higher than θ, $\overline{a}_1$ is chosen for sure in every period under every best reply against σ_2. For every equilibrium in which such θ exists, player 2's posterior

[8] In Online Appendix C, I show that this lower bound is tight in the sense that no strategic type can guarantee a strictly higher payoff when (1) commitment types are rare, and (2) player 2's best reply against player 1's highest action depends on the state.

belief about the state cannot become more pessimistic upon observing $\overline{a}_1$.

However, MS *cannot* rule out conflicts between building reputation for playing $\overline{a}_1$ and signaling high θ. This is because under some equilibrium strategies of player 2s', playing $\overline{a}_1$ in every period is *not* a best reply against it in any state (call them *irregular equilibria*).[9] To circumvent this complication, I establish the following *belief lower bound*, that in every irregular equilibrium, and at every history where $\overline{a}_1$ has always been played in the past, player 2's posterior belief about the state must be optimistic.[10] This belief lower bound implies that player 2 has a strict incentive to play $\overline{a}_2$ as long as she is convinced that $\overline{a}_1$ will be played. It implies that a patient player 1 can secure his payoff from the highest action profile by playing $\overline{a}_1$ in every period.

When ϕ is *pessimistic* and commitment types are rare, Theorem 3 uniquely pins down player 1's equilibrium payoff and on-path behavior. Let

$$A_1^g \equiv \Big\{ a_1^* \in A_1^* \Big| \text{BR}_2(\phi_{a_1^*}, a_1^* | u_2) = \{\overline{a}_2\} \Big\}. \tag{3.9}$$

Assumption 4 implies that $\overline{a}_1 \in A_1^g$. For every pessimistic ϕ, let $\theta^*(\phi)$ be the *largest* $\theta \in \Theta^*$ such that:

$$\{\underline{a}_2\} = \arg\max_{a_2 \in A_2} \Big\{ \sum_{\theta' \succeq \theta} \phi(\theta') u_2(\theta', \overline{a}_1, a_2) \Big\}. \tag{3.10}$$

Let

$$r(\phi) \equiv \frac{u_1(\theta^*(\phi), \underline{a}_1, \underline{a}_2) - u_1(\theta^*(\phi), \overline{a}_1, \underline{a}_2)}{u_1(\theta^*(\phi), \overline{a}_1, \overline{a}_2) - u_1(\theta^*(\phi), \overline{a}_1, \underline{a}_2)}, \tag{3.11}$$

which is strictly between 0 and 1 since $\theta^*(\phi) \in \Theta^*$, and $u_1(\theta, \overline{a}_1, \overline{a}_2) > u_1(\theta, \underline{a}_1, \underline{a}_2) > u_1(\theta, \overline{a}_1, \underline{a}_2)$ for every $\theta \in \Theta^*$. Let

$$w_\theta(\phi) \equiv \begin{cases} u_1(\theta, \underline{a}_1, \underline{a}_2) & \text{if } \theta \preceq \theta^*(\phi) \\ r(\phi) u_1(\theta, \overline{a}_1, \overline{a}_2) + \big(1 - r(\phi)\big) u_1(\theta, \overline{a}_1, \underline{a}_2) & \text{if } \theta \succ \theta^*(\phi). \end{cases} \tag{3.12}$$

Theorem 3. *Under Assumptions 1, 2, and 3. For every pessimistic ϕ, there exist $\overline{\varepsilon} \in (0,1)$ and $\underline{\delta} \in (0,1)$, such that if $\delta > \underline{\delta}$, μ has state distribution ϕ, attaches probability less than $\overline{\varepsilon}$ to all commitment types, and satisfies Assumption 4, then in every equilibrium, strategic type θ's payoff is $w_\theta(\phi)$ for every $\theta \in \Theta$, and*

1. *For every $\theta \succ \theta^*(\phi)$, strategic type θ plays $\overline{a}_1$ at each of his on-path history.*[11]

[9] I construct an irregular equilibrium in Online Appendix D.1, in which playing $\overline{a}_1$ signals low θ at some on-path histories.

[10] Suppose toward a contradiction that in an irregular equilibrium, player 2's posterior belief about the state is pessimistic at some history where player 1 has played $\overline{a}_1$ in all previous periods. On one hand, the definition of irregular equilibrium implies that all strategic types eventually separate from commitment type $\overline{a}_1$, after which player 2's belief about state is optimistic given Assumption 4. On the other hand, some strategic types are supposed to separate from commitment type $\overline{a}_1$ at the *last history* where posterior belief is pessimistic, after which at least one of these types receives his minmax payoff. However, if this type deviates at this last history by pooling with commitment type $\overline{a}_1$, then his continuation payoff is no less than $u_1(\theta, \overline{a}_1, \overline{a}_2)$, which is strictly greater than his minmax payoff. This contradicts his incentive to separate from commitment type $\overline{a}_1$ at that last history.

[11] For any given equilibrium $\big((\sigma_\theta)_{\theta \in \Theta}, \sigma_2\big)$ and state $\theta \in \Theta$, a history h^t is an *on-path history* for strategic type θ if h^t occurs with positive probability under the probability measure induced by $(\sigma_\theta, \sigma_2)$.

2. *For every $\theta \prec \theta^*(\phi)$, strategic type θ plays $\underline{a}_1$ at each of his on-path history.*

3. *In period 0, type $\theta^*(\phi)$ plays a mixed action supported in $A_1^g \cup \{\underline{a}_1\}$. His mixing probabilities are chosen such that for every $a_1^* \in A_1^g \backslash \{\underline{a}_1\}$, after observing a_1^* in period 0, player 2 is indifferent between $\overline{a}_2$ and $\underline{a}_2$ against a_1^* according to her posterior belief about the state. Starting from period 1, type $\theta^*(\phi)$ repeats the same action that he has played in period 0 on the equilibrium path.*

I explain the intuition behind $\theta^*(\phi)$, $r(\phi)$, and $w_\theta(\phi)$ in the context of Theorem 3. When payoffs are MS and commitment types occur with *small but positive probability*, Assumption 4 and the definition of $\theta^*(\phi)$ in (3.10) imply the existence of $q \in (0,1)$ such that:[12]

- if all strategic types above $\theta^*(\phi)$ play $\overline{a}_1$ with probability 1, all strategic types below $\theta^*(\phi)$ play $\underline{a}_1$ with probability 1, and strategic type $\theta^*(\phi)$ plays $\overline{a}_1$ with probability q, then after observing $\overline{a}_1$ in period 0, player 2 is indifferent between $\overline{a}_2$ and $\underline{a}_2$ against $\overline{a}_1$ under her posterior belief about the state.

Since $\{\underline{a}_2\} = \mathrm{BR}_2(\theta^*(\phi), \overline{a}_1 | u_2)$ and $\underline{a}_1$ is played only when $\theta \preceq \theta^*(\phi)$, player 2 plays $\underline{a}_2$ in every period if player 1 plays $\underline{a}_1$ in every period. For every $a_1^* \in A_1^g \backslash \{\underline{a}_1\}$, type $\theta^*(\phi)$'s indifference condition in period 0 uniquely pins down the discounted average probability with which $\overline{a}_2$ is played, conditional on a_1^* being played in every period. For commitment action $\overline{a}_1$, this discounted average probability equals $r(\phi)$.

As a result, every strategic type below $\theta^*(\phi)$ receives his minmax payoff, and every strategic type $\theta \succeq \theta^*(\phi)$ receives payoff $w_\theta(\phi) \equiv r(\phi) u_1(\theta, \overline{a}_1, \overline{a}_2) + \big(1 - r(\phi)\big) u_1(\theta, \overline{a}_1, \underline{a}_2)$. The latter is strictly lower than his guaranteed payoff under an optimistic ϕ.[13] This is because when ϕ is pessimistic and commitment types are rare, player 2 has no incentive to play $\overline{a}_2$ until some strategic types in Θ^* separate from commitment type $\overline{a}_1$, after which at least one of these types receives his minmax payoff. The definition of Θ^* implies that this type strictly prefers $(\overline{a}_1, \overline{a}_2)$ to $(\underline{a}_1, \underline{a}_2)$. In order to prevent this type from imitating other strategic types, the equilibrium payoff of every type in Θ^* must be strictly lower than his payoff from $(\overline{a}_1, \overline{a}_2)$.

Theorem 3 also offers a unique prediction on player 1's on-path behavior, according to which he repeats the same action over time and sustains his reputation.[14] This further implies the uniqueness of player 2's on-path beliefs, according to which observing $\overline{a}_1$ in the first period is treated as a positive signal about the state, observing $\underline{a}_1$ is treated as a negative signal about the state, and player 2s' learning stops after the first period.

These sharp predictions on behavior and learning contrast to private-value reputation games, the optimistic prior case studied by Theorem 2, and repeated signaling games without commitment types. In those models,

[12] If Θ is an interval, ϕ has no atom, and u_2 is continuous in θ, then q is not needed to describe player 1's unique on-path behavior.

[13] Proposition 1.2 in Pei (2018) shows that for every $\theta \in \Theta$, $w_\theta(\phi)$ is a patient player 1's *highest equilibrium payoff* in a repeated incomplete information game with state distribution ϕ but without commitment types.

[14] Player 2's on-path behavior is not unique. This is because the cutoff type's indifference condition only pins down the discounted average frequency of $\overline{a}_2$, but does not pin down how the play of $\overline{a}_2$ is allocated over time.

the informed player has multiple on-path behaviors, switching actions over time is strictly optimal for him in many equilibria, and uninformed players' posterior beliefs vary across equilibria.

My behavioral uniqueness result is driven by a novel *disciplinary effect*, implied by the joint forces of interdependent values and commitment types. In particular, player 1 can guarantee payoff strictly greater than his minmax payoff by imitating commitment type $\overline{a}_1$ (i.e., *a guaranteed reward*), and is guaranteed to receive his minmax payoff after separating from commitment types (i.e., *a guaranteed punishment*).

The guaranteed reward part is driven by Assumption 4, which says that building a reputation for playing $\overline{a}_1$ is feasible and player 2 has a strict incentive to play $\overline{a}_2$ once she is convinced that player 1 is commitment type $\overline{a}_1$. This effect also occurs in private-value reputation games and interdependent-value reputation games studied by Theorem 2. However, it is missing in repeated signaling games without commitment types.

The guaranteed punishment part is driven by interdependent values and the high likelihood of low states, which is absent in existing reputation models. The key observation is that in all equilibria, separating from commitment type $\overline{a}_1$ triggers negative inference about the state, after which player 1's continuation payoff equals his minmax payoff. To see this, take a simplified setting where $A_1 \equiv \{\overline{a}_1, \underline{a}_1\}$ and suppose toward a contradiction that playing $\underline{a}_1$ makes player 2's belief about the state more optimistic. Since player 2's prior belief is pessimistic and belief is a martingale, observing $\overline{a}_1$ leads to a more pessimistic belief about the state. This implies that there exists at least one strategic type who (1) plays $\overline{a}_1$ with positive probability, and (2) receives his minmax payoff after playing $\overline{a}_1$. Since playing $\overline{a}_1$ is strictly costly for player 1, the above strategic type has a strict incentive to deviate by playing $\underline{a}_1$. This leads to a contradiction.

In contrast, this guaranteed punishment is missing in private-value reputation games. According to the folk theorem in Fudenberg, Kreps and Maskin (1990), strategic player 1's continuation payoff after separating from commitment type $\overline{a}_1$ can be anything between his minmax payoff and his payoff from $(\overline{a}_1, \overline{a}_2)$. This multiplicity in continuation values leads to multiple on-path behaviors. This is because at any given history, whether player 1 has an incentive to pool with or separate from commitment type $\overline{a}_1$ depends on his continuation value after separation. In particular, he strictly prefers to pool with commitment type $\overline{a}_1$ if he can only receive his minmax payoff, and strictly prefers to separate from commitment type $\overline{a}_1$ if he can still receive payoff from $(\overline{a}_1, \overline{a}_2)$.

A similar intuition applies to the case studied by Theorem 2. Each strategic type's continuation payoff after separating from commitment type $\overline{a}_1$ can be anything between his minmax payoff and his payoff from $(\overline{a}_1, \overline{a}_2)$. This leads to multiple on-path behaviors for patient player 1. In addition, observing $\overline{a}_1$ is interpreted as a positive signal about θ in some equilibria, and is treated as a negative signal in other equilibria.

Player 1 also has multiple on-path behaviors in private-value reputation games with a persistent state, for example, when he has persistent private information about his discount factor or his cost of taking a high action. This is because a strategic type with high cost or low discount factor *either* separates from the commitment type, after which the disciplinary effect disappears; *or* he pools with the commitment type, in which case he is equivalent to the commitment type from player 2's perspective thanks to the private-value assumption.

3.3 Insights on the Product Choice Game

I map the product choice game into the US toy industry. I use the variation with a continuum of long-lived consumers, with a small fraction of them being mechanical and choose to buy in every period. I explain the fitness of my modeling assumptions and deliver novel testable predictions based on my analysis in section 3.2. When citing empirical evidence, I associate a positive fraction of firms having a certain characteristic (or taking a particular action) to a positive probability of the single firm having that characteristic (or taking that action).

Fitness of Assumptions: Player 1 is the *headquarter of a US toy company*. Player 2s are a continuum of consumers (e.g., parents who want to purchase a particular type of toys for their kids), who are potentially long-lived but find it optimal to play their myopic best replies.

In every period, each consumer decides whether to buy a toy from the company (action T) or not (action N), and the company chooses its effort, either high (H) or low (L), that affects the *design of its toys*. This includes for example, whether the toys are fun to play with and whether their designs have flaws.

I model the design of toys as the company's action, which it chooses in every period. This is because according to Beamish and Bapuji (2008), the headquarters of US toy companies are mainly in charge of designing toys, while production is done mostly in developing countries. According to Freedman, Kearney and Lederman (2012, page 503), toy companies need to frequently redesign its toys in order to introduce new features and to meet customers' changing demands.

The toy company has persistent private information about the quality of its upstream supplier,[15] which is either high ($\theta = \theta_h$) or low ($\theta = \theta_l$). Consumers' payoffs depend not only on the design of the toy (i.e., the company's effort), but also on the company's supplier quality. This is because the latter has a significant impact on the safety of toys being manufactured, such as the lead-content in the paint of the toys. Consumers can observe the toy company's effort choices in the past, but *cannot* directly observe the quality of their suppliers or the safety of their toys. I explain the fitness of these modeling assumptions in four steps:

1. There is heterogeneity in supplier quality among US toy companies, and this is correlated with the heterogeneity in product safety. According to Beamish and Bapuji (2008), the production of toys is typically done in developing countries, either through FDI or through outsourcing to independent local manufac-

[15] I do not model the toy company's supplier as a player in the game. Instead, each company is associated with its own supplier(s) and the company privately knows the quality of its supplier(s), which is interpreted as the persistent state in my model.

turers.[16] The New York Times published an article titled "*Toymaking in China, Mattel's Way*," which suggests that firms that conduct FDI typically produce higher quality toys and are less likely to suffer from product safety problems compared to those whose production is outsourced. These claims are substantiated by Hansman, Hjort, León and Teachout (2018) who show empirically that firms that conduct FDI are significantly better in terms of their product safety.

2. Supplier quality (i.e., whether the toy company conducts FDI or not) is *persistent* over time, and cannot be easily changed by the company's headquarter. This is supported empirically from the trade literature, which has documented that firms face very large fixed costs to conduct FDI (Helpman 2006).

   Evidence for persistence is also found in the empirical work of Freedman, Kearney and Lederman (2012), who show that consumers significantly update their beliefs about the safety of a toy that is *currently active on the market*, after a toy that was *no longer actively selling* was recalled due to safety reasons. Under the presumption that consumers are Bayesian, this suggests that product safety is persistent over time. This is because otherwise, consumers' belief about the safety of a currently active product should not change after receiving negative information about the safety of an old product that was no longer active.

3. Consumers can observe the company's effort since it is reflected in the design of the toys. According to Ni, Flynn, and Jacobs (2016), attributes related to the product's design are typical examples of *experience quality* (Nelson 1970), which can be observed by consumers *after purchase*. This information can be transmitted to future consumers through product reviews and word-of-mouth communication.[17]

4. Consumers cannot easily observe the safety of products (affected by supplier quality) even after purchase.[18] This is because the long-term health impact of a product is a classic example of *credence quality* (Darby and Karni 1973). In the toy industry, a significant fraction of safety issues in recent years are

[16] FDI (or foreign direct investment) is an investment in the form of a controlling ownership in a business in one country by an entity based in another country. In this application, FDI refers to a US toy company directly owning manufacturing facilities in China.

[17] This uses the extension that a small fraction of buyers are mechanical and automatically buy in every period. It addresses the concern that future consumers cannot learn about the seller's action in a given period if no buyer buys in that period.

[18] My model hinges on two assumptions. First, consumers *face uncertainty* about the company's supplier quality after observing the publicly available information, which includes the company's name and advertisements, the country in which the toy was manufactured, and so on. Second, after observing these public signals, consumers *do not receive additional information about supplier quality over time* except for the company's effort. Despite a company's name is informative about its supplier quality, it has already been incorporated into consumers' prior belief about the state, i.e., it does not give consumers *new information* about the state in every period. Consumers who want to purchase *a particular type of toy* from the company can still face uncertainty about supplier quality after observing the company's name. This is because even within the same company, different categories of toys are produced in different factories, some are directly owned by the toy company and some are owned by local independent manufacturers. According to New York Times article Toymaking in China, Mattel's Way, "*about 50% of Mattel's toy revenue comes from products made in company-run plants*". Another way of saying this is, about 50% of Mattel's products are manufactured by *independent local suppliers*. Similar to the company's name, its advertisements also cannot give consumers new information about supplier quality in every period, i.e., by repeatedly watching the same advertisement by the same company every day, a consumer does not learn new things about the supplier quality every day. In addition, a toy company might be reluctant to disclose which of its products are produced in-house in their advertisements. This is because doing do can hurt the sales of the company's other products that are produced by independent manufacturers.

related to lead-based paint (Freedman, et al. 2012). Lead-based paint reduces consumers' willingness to pay for the toy since lead can be absorbed by the human body and cause long-lasting health conditions. However, most children who are exposed to lead have no immediate symptoms, and therefore, parents cannot observe them even after purchasing the toy.[19]

In the product choice game of section 2.1, my monotone-supermodularity assumption translates into $\eta > 0$, i.e., a toy company with higher quality supplier faces lower effort cost to introduce a good toy design. Empirical evidence that supports this assumption is documented in Helpman (2006), who shows that firms that conduct FDI (i.e., those with higher quality suppliers) are usually more efficient and are more productive compared to those that outsource their production.

Testable Predictions: My analysis in section 3.2 leads to the following empirical predictions:

1. When consumers entertain an optimistic belief about the firm's quality (i.e., their prior belief attaches probability more than $1/2$ to state θ_h), the firm has a strict incentive to *behave inconsistently* in many equilibria, i.e., exerting high effort for some periods and exerting low effort in other periods.

2. When consumers entertain a pessimistic belief about the firm's quality (i.e., their prior belief attaches probability less than $1/2$ to state θ_h), the firm's effort exhibits perfect serial correlation in *all equilibria*: as long as the firm exerts high effort in the first period, it exerts high effort in all subsequent periods; as long as the firm exerts low effort in the first period, it exerts low effort in all subsequent periods.

These predictions can be tested in the US toy industry by comparing the *serial correlation* in a toy company's effort *before* and *after* one of its products was recalled due to safety reasons. My theory predicts that the *intertemporal correlation* of a firm's observable effort increases after the recall. To the best of my knowledge, this prediction that links *consumers' beliefs about a persistent state* to the *intertemporal correlation of a firm's observable effort* is novel and has not been explored in the existing reputation literature.

The assumption behind this empirical exercise is that a product safety recall causes a *significant drop in consumers' belief* about the toy company's quality.[20] Such an assumption has been tested both in the toy industry and in other industries by estimating the stock market's reactions to product recalls. For example,

[19] Despite consumers can receive informative signals about a product's safety after it is recalled, Freedman, Kearney and Lederman (2012) document that (1) the number of recalls is small relative to the varieties of toys on the market (i.e., such signals rarely arrive), and (2) most recalls happen to toys that have not been actively selling on the market for a long time (i.e., there is a significant delay in the arrival of such signals). That being said, my assumption that consumers cannot observe additional signals about the state other than the company's observable effort is a good approximation of what happens in the US toy industry.

[20] In an interdependent value reputation game, a discontinuous change in consumers' beliefs about the persistent state can also lead to a discontinuous change in the company's Stackelberg payoff. Take the product choice game example, when consumers' belief attaches probability more than $1/2$ to state θ_h, the company's Stackelberg payoff is 1 in state θ_h and $1-\eta$ in state l; when consumers' belief attaches probability less than $1/2$ to state θ_h, the company's Stackelberg payoff is 0 in both states.

Jarrell and Peltzman (1985) study the automobile and pharmaceutical industries, and find a significant drop in stock prices after a recall announcement. Hoffer, Pruitt and Reilly (1987) study the stock prices of Chrysler, Ford, and General Motors around the time of announcement of severe automotive recalls. They also find a significant decline in stock prices after the news of recall becomes public. In the toy industry, similar patterns are documented by Ni, Flynn, and Jacobs (2016), who show that the announcement of a toy recall is associated with a 0.5% wealth loss by the source firm. These findings suggest that product recalls have significant negative effects on the market's belief about firm quality.

In terms of the reasons behind such significant belief changes, Jarrell and Peltzman (1985) propose a model where the extent to which the financial market reacts to a product recall depends on *how well the recall is anticipated by the market ex ante*: a product recall has a more significant impact on stock prices when similar recalls rarely happen, or in another word, the market believes that such recalls are unlikely.[21] In the toy industry, product recalls rarely happen to toys actively selling on the market. As documented by Freedman, Kearney and Lederman (2012), only 82 toys were recalled in 2007 and between 30-38 toys were recalled each year from 2004 to 2006. Among those recalled toys, 78% of them were no longer actively selling on the market at the time of their recall announcements. Jarrell and Peltzman (1985)'s model implies that once such a recall happens, it leads to a significant change in consumers' beliefs about the firm's quality.

A practical challenge to implement this empirical analysis is that researchers cannot directly observe a toy company's effort in designing its toys. A proxy for effort is the consumers' product reviews. In the toy industry example, consumers' reviews reflect mostly their assessments of the product's design and other attributes they can *directly observe*, instead of product safety that they cannot observe. A better but more costly way to obtain data related to the company's effort is by conducting surveys among consumers, i.e., sending consumers questionnaires and explicitly asking them about their satisfaction with the *design* of the toy they purchased.

4 Concluding Remarks

My analysis unveils challenges to reputation building when uninformed players' learning is confounded. This is related to Deb and Ishii (2019) who study reputation building when uninformed players do not know the monitoring structure. Their model assumes that there exists a public signal which can statistically identify the state.[22] They construct a commitment type that plays a *nonstationary* strategy under which the patient informed player can secure his complete information commitment payoff.

[21]Large discontinuous drop in belief also occurs in Poisson bad news learning models such as Keller and Rady (2010). In their model, bad news rarely arrives, but once it arrives, it leads to a large discontinuous change in players' beliefs about the persistent state.

[22]Their Assumption 2.3 requires for every $\theta, \theta' \in \Theta$, there exists $\alpha_1 \in \Delta(A_1)$ such that the signal distribution under (θ, α_1) *cannot* be induced by *any action distribution* in state θ'. This is violated in my model and the repeated incomplete information games of Aumann and Maschler (1995), Hart (1985), Hörner and Lovo (2009), and Pęski (2014), and repeated signaling games in Kaya (2009).

In contrast, I study a model in which the state affects the uninformed players' best replies, but can only be learnt via the informed player's actions. The lack of exogenous signals that can statistically identify the state introduces new challenges for uninformed players to learn the correct best reply against the commitment action. I also derive a unique prediction on the informed player's on-path behavior in addition to obtaining lower bounds on his equilibrium payoff.

Theorem 1 is related to Ely and Välimäki (2003) and Ely, Fudenberg and Levine (2008), who show that reputation fails in a class of *private-value* reputation games called *participation games*. They show that an informed player's equilibrium payoff is low when commitment types that discourage the uninformed players from participating occur with high enough probability relative to Stackelberg commitment types. Their results rely on uninformed players' ability to choose a *non-participating action*, under which the public signal becomes uninformative about the informed player's action. The informed player receives low payoff when his opponents do not participate, but once they participate, he has a strong incentive to build his reputation.

In my model, the uninformed players *cannot* stop the informed player from signaling his type. However the informational content of the informed player's actions is sensitive to equilibrium selection. In particular, there exist equilibria in which uninformed players believe that the strategic informed player is more likely to choose the commitment action in some alternative state, under which they have no incentive to play the informed player's desired best reply. This discourages the informed player from building his reputation even when he has an opportunity to do so, which differs from the predictions in bad reputation models.

My Theorem 3 suggests that interdependent values can contribute to reputation sustainability. Following Cripps, Mailath and Samuelson (2004), a reputation for playing $\overline{a}_1$ is sustained in an equilibrium if *conditional on* player 1 being strategic, there exists a positive probability event under which player 1's reputation (i.e., probability with which player 2's posterior assigns to commitment type $\overline{a}_1$) does not vanish as $t \to \infty$. According to this definition, reputation for playing $\overline{a}_1$ is sustained in all equilibria in games studied by Theorem 3. This contrasts to Fudenberg and Levine (1989) in which player 1 loses his reputation in some equilibria, and Cripps, Mailath and Samuelson (2004) in which player 1 loses his reputation in all equilibria.

Player 1's unique on-path behavior in Theorem 3 is close to one of his equilibrium behaviors in a benchmark repeated game with the same state distribution but *without* commitment types, which is also the case in Cripps, Mailath and Samuelson (2004). The difference is, playing a costly action (such as $\overline{a}_1$) in every period is *suboptimal* in any equilibrium of a repeated game with private values, full support monitoring, but without commitment types. This explains why reputation vanishes in all equilibria of their model. In contrast, playing $\overline{a}_1$ in every period is optimal in *some equilibria* of the repeated game with interdependent values, perfect monitoring, but without commitment types. Theorem 3 shows that introducing commitment type selects equilibria with this particular on-path behavior, which indicates the sustainability of reputation in my model.

A General Characterization Theorem

I generalize Theorem 1 by allowing for arbitrary correlation between θ and player 1's characteristics. I provide a *sufficient and almost necessary condition* on the joint distribution μ under which the commitment payoff theorem applies to every u_1.

Let $m \equiv |\Theta|$. For every $\theta \in \Theta$, let $\mu(\theta)$ be the probability of *strategic type* θ. For every $a_1^* \in A_1^*$, let $\mu(a_1^*)$ be the probability of *commitment type* a_1^*. Let $\phi_{a_1^*} \in \Delta(\Theta)$ be the state distribution conditional on player 1 being commitment type a_1^*. Let $\lambda(\mu, a_1^*) \equiv \big\{\lambda_\theta(\mu, a_1^*)\big\}_{\theta \in \Theta} \in \mathbb{R}_+^m$ be the *likelihood ratio vector* with respect to a_1^*, with $\lambda_\theta(\mu, a_1^*) \equiv \mu(\theta)/\mu(a_1^*)$. Let $a_2^*(\theta^*, a_1^*|u_2)$ be the unique element in $\mathrm{BR}_2(\theta^*, a_1^*|u_2)$. Let $\Lambda(\theta^*, a_1^*, u_2)$ be the subset of $\mathbb{R}_+^m$ such that $\lambda \in \Lambda(\theta^*, a_1^*, u_2)$ *if and only if*:

$$\big\{a_2^*(\theta^*, a_1^*|u_2)\big\} = \arg\max_{a_2 \in A_2} \big\{u_2(\phi_{a_1^*}, a_1^*, a_2) + \sum_{\theta \in \Theta} \lambda'_\theta u_2(\theta, a_1^*, a_2)\big\} \text{ for all } \lambda' \equiv \{\lambda'_\theta\}_{\theta \in \Theta} \text{ with } 0 \leq \lambda' \leq \lambda. \tag{A.1}$$

According to (A.1), whether $\lambda(\mu, a_1^*)$ belongs to $\Lambda(\theta^*, a_1^*, u_2)$ does not depend on the probability of commitment types other than a_1^*, nor does it depend on the probability of strategic types under which player 2's best reply against a_1^* is the same as that under state θ^*. This is because first, player 2s rule out other pure strategy commitment types after observing a_1^*. Second, in the worst equilibrium, those good strategic types separate from commitment type a_1^* and the bad strategic types pool with commitment type a_1^*.

Theorem 1'. *Suppose Assumption 1 holds and μ has full support. For every $(\theta^*, a_1^*) \in \Theta \times A_1^*$,*

1. *If $\lambda(\mu, a_1^*) \in \Lambda(\theta^*, a_1^*, u_2)$, then $\liminf_{\delta \to 1} \underline{v}_{\theta^*}(\delta, \mu, u_1, u_2) \geq v_{\theta^*}(a_1^*, u_1, u_2)$ for every u_1.*

2. *If $\lambda(\mu, a_1^*)$ does not belong to the closure of $\Lambda(\theta^*, a_1^*, u_2)$, then there exists u_1 such that $\limsup_{\delta \to 1} \underline{v}_{\theta^*}(\delta, \mu, u_1, u_2) < v_{\theta^*}(a_1^*, u_1, u_2)$.*

The proof is in Appendices B and C, and a generalization that incorporates mixed-strategy commitment types is stated as Theorem 1.1 in Pei (2018). To avoid cumbersome notation, I replace $a_2^*(\theta^*, a_1^*|u_2)$, $\lambda(\mu, a_1^*)$, $\Lambda(\theta^*, a_1^*, u_2)$, and $\mathrm{BR}_2(\theta, a_1|u_2)$ with a_2^*, λ, Λ, and $\mathrm{BR}_2(\theta, a_1)$, respectively. Theorem 1 is implied by statement 2 of Theorem 1' since all entries of λ go to infinity when the probability of commitment types vanishes to 0, and when interdependent values are nontrivial, there exists $\theta' \neq \theta$ such that $a_2^* \notin \mathrm{BR}_2(\theta', a_1^*)$. Therefore, if $\lambda_{\theta'}$ is large enough, then the likelihood ratio vector λ does not belong to the closure of Λ.

I explain the intuition behind the set Λ. First, for type θ^* to secure payoff $v_{\theta^*}(a_1^*, u_1, u_2)$ under every u_1, it is necessary that a_2^* is player 2's strict best reply against a_1^* under her prior belief about the state. However, this is *not sufficient* since player 2's belief is updated over time. As a result, player 1 needs to find a strategy under which he can pool with commitment type a_1^*, while making sure that player 2 has an incentive to play a_2^* under her posterior belief about the state. Given that a_1^* is a pure action, each entry of λ is nonincreasing as long as player 1 plays a_1^* in every period. As a result, Λ requires a_2^* to be a strict best reply against a_1^* after *any* fraction of strategic types separate from commitment type a_1^*.

B Proof of Theorem 1' Statement 1

If $\lambda \in \Lambda$, then for every $\lambda' \equiv \{\lambda'_\theta\}_{\theta\in\Theta}$ with $0 \leq \lambda' \leq \lambda$, we have:

$$\{a_2^*\} = \arg\max_{a_2\in A_2} \Big\{u_2(\phi_{a_1^*}, a_1^*, a_2) + \sum_{\theta\in\Theta} \lambda'_\theta u_2(\theta, a_1^*, a_2)\Big\}. \tag{B.1}$$

Let $\overline{h}^t$ be a history such that a_1^* has been played in every period. For every $\theta \in \Theta$, let $q_t(\theta)$ be the ex ante probability of the event that $h^t = \overline{h}^t$ and player 1 is strategic type θ. Player 2's maximization problem at $\overline{h}^t$ is:

$$\max_{a_2\in A_2} \Big\{\mu(a_1^*)u_2(\phi_{a_1^*}, a_1^*, a_2) + \sum_{\theta\in\Theta} \Big[q_{t+1}(\theta)u_2(\theta, a_1^*, a_2) + (q_t(\theta) - q_{t+1}(\theta))u_2(\theta, \alpha_{1,t}(\theta), a_2)\Big]\Big\} \tag{B.2}$$

in which $\alpha_{1,t}(\theta) \in \Delta(A_1\backslash\{a_1^*\})$ is the distribution of strategic type θ's action at $\overline{h}^t$ conditional on $a_{1,t} \neq a_1^*$. If type θ plays a_1^* at $\overline{h}^t$ with probability 1, then let $\alpha_{1,t}(\theta)$ be any arbitrary distribution over A_1. Since $\frac{q_t(\theta)}{\mu(a_1^*)} \leq \frac{\mu(\theta)}{\mu(a_1^*)} = \lambda_\theta$, equation (B.1) implies that:

$$\{a_2^*\} = \arg\max_{a_2\in A_2} \Big\{\mu(a_1^*)u_2(\phi_{a_1^*}, a_1^*, a_2) + \sum_{\theta\in\Theta} q_t(\theta)u_2(\theta, a_1^*, a_2)\Big\}.$$

This together with (B.2) implies that there exists $\rho > 0$, such that player 2 has a strict incentive to play a_2^* at $\overline{h}^t$ as long as

$$\sum_{\theta\in\Theta} q_{t+1}(\theta) > \sum_{\theta\in\Theta} q_t(\theta) - \rho. \tag{B.3}$$

If player 1 plays a_1^* in every period, then there exist at most $\overline{T} \equiv \lceil 1/\rho \rceil$ periods in which player 2 does not have a strict incentive to play a_2^*. For every $\theta \in \Theta$, type θ's payoff by playing a_1^* in every period is at least:

$$(1 - \delta^{\overline{T}}) \min_{(a_1,a_2)\in A_1\times A_2} u_1(\theta, a_1, a_2) + \delta^{\overline{T}} v_\theta(a_1^*, u_1, u_2). \tag{B.4}$$

Given that $\overline{T}$ is independent of δ, (B.4) converges to $v_\theta(a_1^*, u_1, u_2)$ as $\delta \to 1$. This implies statement 1 of Theorem 1' since type θ's equilibrium payoff is no less than (B.4).

C Proof of Theorem 1 and Theorem 1' Statement 2

Throughout the proof, let player 1's stage-game payoff function be:

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

First, consider the case in which $\{a_2^*\} = \mathrm{BR}_2(\phi_{a_1^*}, a_1^*)$. Since λ does not belong to the closure of Λ, there exists $\lambda' \equiv \{\lambda'_\theta\}_{\theta\in\Theta}$ such that $0 \leq \lambda'_\theta < \lambda_\theta$ for every $\theta \in \Theta$, and

$$\{a_2^*\} \neq \arg\max_{a_2\in A_2} \Big\{u_2(\phi_{a_1^*}, a_1^*, a_2) + \sum_{\theta\in\Theta} \lambda'_\theta u_2(\theta, a_1^*, a_2)\Big\}. \tag{C.2}$$

The presumption $\{a_2^*\} = \mathrm{BR}_2(\phi_{a_1^*}, a_1^*)$ implies that $\mathbf{0} \in \Lambda$, and therefore, there exists $\lambda'' \equiv \{\lambda''_\theta\}_{\theta\in\Theta}$ with $\lambda'_\theta \leq \lambda''_\theta < \lambda_\theta$ for every $\theta \in \Theta$, and

$$a_2^* \notin \arg\max_{a_2\in A_2} \Big\{u_2(\phi_{a_1^*}, a_1^*, a_2) + \sum_{\theta\in\Theta} \lambda''_\theta u_2(\theta, a_1^*, a_2)\Big\}. \tag{C.3}$$

Let $\lambda^\dagger \in \mathbb{R}_+^m$ be such that $\lambda^\dagger_{\theta^*} \equiv 0$, and $\lambda^\dagger_\theta \equiv \lambda''_\theta$ for all $\theta \neq \theta^*$. Since $\{a_2^*\} = \mathrm{BR}_2(\theta^*, a_1^*)$, we have

$$\arg\max_{a_2\in A_2} \Big\{u_2(\phi_{a_1^*}, a_1^*, a_2) + \sum_{\theta\in\Theta} \lambda^\dagger_\theta u_2(\theta, a_1^*, a_2)\Big\},$$

does not contain a_2^*, but contains some $a_2' \neq a_2^*$. Since $\{a_2^*\} = \mathrm{BR}_2(\phi_{a_1^*}, a_1^*)$, there exists $a_2'' \neq a_2^*$, which can be the same as a_2', such that:

$$a_2'' \in \arg\max_{a_2\in A_2} \Big\{\sum_{\theta\in\Theta} \lambda^\dagger_\theta u_2(\theta, a_1^*, a_2)\Big\}. \tag{C.4}$$

Let $\varepsilon > 0$ be small enough such that first, $(1+\varepsilon)\lambda^\dagger_\theta < \lambda_\theta$ for all $\theta \in \Theta$; and second, a_2^* is player 2's strict best reply against a_1^* when player 2's belief attaches probability more than $\frac{\lambda_{\theta^*}}{\lambda_{\theta^*}+\varepsilon\sum_{\theta\neq\theta^*}\lambda_\theta}$ to state θ^*.

Consider the following strategy profile. On the equilibrium path, type θ^* plays $a_1' \neq a_1^*$ in every even period and a_1^* in every odd period. For every $\theta \neq \theta^*$, type θ plays a_1^* in every period with probability $\lambda^\dagger_\theta/\lambda_\theta$; he plays a_1' in every even period and a_1^* in every odd period with probability $\varepsilon\lambda^\dagger_\theta/\lambda_\theta$; he plays a_1' in every period with probability $1-(1+\varepsilon)\lambda^\dagger_\theta/\lambda_\theta$. Starting from period 1, player 2 plays a_2' at on-path histories where a_1^* was played in period 0. She plays a_2'' at off-path histories. In odd periods starting from period 3, she plays a_2^* *if and only if* a_1^* has been played in every odd period before, and a_1' has been played in every even period before. At other on-path histories, she plays any of her myopic best replies according to her belief about θ and player 1's action.

This is an equilibrium since type θ^*'s continuation payoff at each of his on-path histories is approximately

$1/2$ when δ is close enough to 1, and his continuation payoff after he deviates is at most $1-\delta$, no matter when and how he deviates. Player 2's incentive constraints at on-path histories are satisfied, and other types of player 1's incentive constraints are trivially satisfied. Type θ^*'s equilibrium payoff is approximately $1/2$ as $\delta \to 1$, which is strictly bounded below his commitment payoff from a_1^*, equal to 1.

Next, consider the case in which $\{a_2^*\} \neq \text{BR}_2(\phi_{a_1^*}, a_1^*)$, i.e., Λ is empty. Let $a_2' \neq a_2^*$ be such that $a_2' \in \text{BR}_2(\phi_{a_1^*}, a_1^*)$. Since $\{a_2^*\} = \text{BR}_2(\theta^*, a_1^*)$, there exist $\theta' \neq \theta^*$ and $a_2'' \neq a_2^*$ such that $a_2'' \in \text{BR}_2(\theta', a_1^*)$.

On the equilibrium path, type $\theta'' \notin \{\theta^*, \theta'\}$ plays $a_1' \neq a_1^*$ in every period. Type θ^* plays a_1' in every even period and a_1^* in every odd period. With probability $1-\varepsilon$, type θ' plays a_1' in every period; with probability ε, he plays a_1' in every even period and a_1^* in every odd period, with ε being small enough such that a_2^* is player 2's strict best reply against a_1^* when her belief attaches probability $\frac{\mu(\theta^*)}{\mu(\theta^*)+\varepsilon\mu(\theta')}$ to state θ^* and complementary probability to state θ'. Starting from period 1, player 2 plays a_2' at on-path histories where a_1^* was played in period 0. She plays a_2'' at off-path histories. In odd periods starting from period 3, she plays a_2^* if and only if a_1^* has been played in every odd period before, and a_1' has been played in every even period before. At other on-path histories, she plays any of her myopic best replies according to her belief about θ and player 1's action.

This is an equilibrium since type θ^*'s continuation payoff at each of his on-path histories is approximately $1/2$ when δ is close enough to 1, and his continuation payoff after he deviates is at most $1-\delta$, no matter when and how he deviates. Player 2's incentive constraints at on-path histories are satisfied, and other types of player 1's incentive constraints are trivially satisfied. Type θ^*'s equilibrium payoff is approximately $1/2$ as $\delta \to 1$, which is strictly bounded below his commitment payoff from a_1^*, equal to 1.

Product Choice Game: I apply Theorem 1 to commitment action H and state θ_h in the product choice game. Take a full support state distribution $\phi(\theta_h) = 0.99$ and $\phi(\theta_l) = 0.01$, and let the firm's payoff function be the one in the matrices with $-1 < \eta \leq 0$. When the probability of commitment types is less than 0.01 and δ is close to 1, the following strategy profile is an equilibrium:

1. Strategic type θ_l plays H if L has never been played before, and plays L otherwise.

2. In period 0, strategic type θ_h plays H with probability β, and plays L with probability $1-\beta$, in which $\beta \in (0,1)$ is such that when H is observed in period 0, player 2's posterior attaches probability $1/2$ to state θ_h. Such β exists when the probability of commitment type H is less than that of strategic type θ_l.

   In period $t \geq 1$, strategic type θ_h plays H if L has never been played before, and plays L otherwise.

3. Consumer plays N in period 0. Starting from period 1, she plays T with probability $1/(2\delta)$ if L has never been played before, and plays N otherwise.

In the above equilibrium, the strategic high-quality firm's payoff is 0, which is strictly lower than his complete information commitment payoff from H, equals to 1.

D Proofs of Theorems 2 and 3

I show Theorems 2 and 3 under a simplifying assumption that player 2 can only observe player 1's past actions, which conveys the key ideas and intuition. The full proofs are in Online Appendices A and B.

D.1 Partition of Θ

For every $\phi \in \Delta(\Theta)$ and $\alpha_1 \in \Delta(A_1)$, let $\mathcal{D}(\phi, \alpha_1) \equiv u_2(\phi, \alpha_1, \overline{a}_2) - u_2(\phi, \alpha_1, \underline{a}_2)$. Let

$$\Theta_g \equiv \left\{\theta \middle| \mathcal{D}(\theta, \overline{a}_1) \geq 0 \text{ and } \theta \in \Theta^*\right\}, \quad \Theta_p \equiv \left\{\theta \middle| \mathcal{D}(\theta, \overline{a}_1) < 0 \text{ and } \theta \in \Theta^*\right\} \tag{D.1}$$

and $\Theta_n \equiv \Theta \backslash \Theta^*$. One can verify that first, $\{\Theta_g, \Theta_p, \Theta_n\}$ is a partition of Θ, and second, $\Theta^* = \Theta_g \cup \Theta_p$. I focus on the nontrivial case in which neither Θ_g nor $\Theta_p \cup \Theta_n$ is empty. I show the following lemma:

Lemma D.1. *When u_1 and u_2 satisfy Assumption 2:*

1. *If $\theta_g \in \Theta_g$, $\theta_p \in \Theta_p$, and $\theta_n \in \Theta_n$, then $\theta_g \succ \theta_p$, $\theta_p \succ \theta_n$ and $\theta_g \succ \theta_n$.*
2. *If both Θ_p and Θ_n are nonempty, then $\mathcal{D}(\theta_n, \overline{a}_1) < 0$ for every $\theta_n \in \Theta_n$.*

Proof. To show statement 1, first consider the case in which $\Theta_p \neq \{\varnothing\}$. Since $\mathcal{D}(\theta_g, \overline{a}_1) \geq 0$, $\mathcal{D}(\theta_p, \overline{a}_1) < 0$, and u_2 has SID in θ and a_2, we have $\theta_g \succ \theta_p$. Since $u_1(\theta_p, \overline{a}_1, \overline{a}_2) > u_1(\theta_p, \underline{a}_1, \underline{a}_2)$, $u_1(\theta_n, \overline{a}_1, \overline{a}_2) \leq u_1(\theta_n, \underline{a}_1, \underline{a}_2)$, and u_1 has SID in θ and (a_1, a_2), we have $\theta_p \succ \theta_n$. Next, consider the case in which $\Theta_p = \{\varnothing\}$. Since $u_1(\theta_g, \overline{a}_1, \overline{a}_2) > u_1(\theta_g, \underline{a}_1, \underline{a}_2)$ and $u_1(\theta_n, \overline{a}_1, \overline{a}_2) \leq u_1(\theta_n, \underline{a}_1, \underline{a}_2)$, we have $\theta_g \succ \theta_n$. To show statement 2, given $\Theta_p, \Theta_n \neq \{\varnothing\}$, u_2 has SID in θ and a_2 implies that $\mathcal{D}(\theta_n, \overline{a}_1) < \mathcal{D}(\theta_p, \overline{a}_1) < 0$. $\square$

D.2 Implication of Stage-Game MS & Two Classes of Equilibria

I derive an implication of *MS stage-game payoffs* on a *repeated MS game*. For any strategy profile $\sigma \equiv \big((\sigma_\theta)_{\theta \in \Theta}, \sigma_2\big)$ and $\theta \in \Theta$, let $\mathcal{P}^{(\sigma_\theta, \sigma_2)}$ be the probability measure over public histories induced by $(\sigma_\theta, \sigma_2)$.

Lemma D.2. *Under Assumptions 2 and 3, for every $\widehat{\theta} \succ \widetilde{\theta}$ and in every equilibrium $\sigma \equiv \big((\sigma_\theta)_{\theta \in \Theta}, \sigma_2\big)$,*

1. *if playing $\overline{a}_1$ in every period is type $\widetilde{\theta}$'s best reply against σ_2, then according to $\sigma_{\widehat{\theta}}$, strategic type $\widehat{\theta}$ plays $\overline{a}_1$ with probability 1 at every history that occurs with positive probability under $\mathcal{P}^{(\sigma_{\widehat{\theta}}, \sigma_2)}$.*

2. *if playing $\underline{a}_1$ in every period is type $\widehat{\theta}$'s best reply against σ_2, then according to $\sigma_{\widetilde{\theta}}$, strategic type $\widetilde{\theta}$ plays $\underline{a}_1$ with probability* 1 *at every history that occurs with positive probability under* $\mathcal{P}^{(\sigma_{\widetilde{\theta}},\sigma_2)}$.

This lemma is implied by Theorem 1 in Liu and Pei (2020), which says that in a *1-shot signaling game* where the sender's payoff u_1 and the receiver's payoff u_2 satisfy Assumption 2, and the receiver's action choice is binary, then in every Nash equilibrium, the sender's action is nondecreasing in θ.

In the *repeated MS game* of this paper, each type of player 1 chooses $\sigma_\theta : \mathcal{H} \rightarrow \Delta(A_1)$ and induces a discounted average distribution over player 2s' actions. Liu and Pei (2020)'s theorem implies that if a lower type player 1 finds it optimal to *play* $\overline{a}_1$ *in every period* of the repeated game, then playing actions other than $\overline{a}_1$ on the equilibrium path must be suboptimal for any higher type. A similar argument applies to $\underline{a}_1$. However, it *does not* imply that at any *given history*, a higher strategic type is more likely to play $\overline{a}_1$ than a lower strategic type. This is because player 1's action affects the equilibrium being played in the continuation game.

I categorize the set of equilibria into two classes. An equilibrium $\sigma \equiv \big((\sigma_\theta)_{\theta\in\Theta}, \sigma_2\big)$ is *regular* if there exists $\theta \in \Theta_p \cup \Theta_n$ such that playing $\overline{a}_1$ in every period is type θ's best reply against σ_2. Otherwise, σ is *irregular*. Let $\overline{q}$ be the probability player 2's prior belief μ attaches to commitment type $\overline{a}_1$. Let $\overline{h}^t \equiv (\overline{a}_1, \overline{a}_1, ..., \overline{a}_1)$. Given σ, let $q_t^\sigma(\theta)$ be the probability of the event that the state is θ, player 1 is strategic, and the history is $\overline{h}^t$.

D.3 Analysis of Regular Equilibria

I show that first, if ϕ is optimistic, then there exists a constant $C \in \mathbb{R}_+$ such that for every $\theta \in \Theta^*$, type θ's payoff in any regular equilibrium is at least $u_1(\theta, \overline{a}_1, \overline{a}_2) - (1-\delta)C$. Second, if ϕ is pessimistic and the probability of commitment types is small enough, then player 1's payoff and on-path behavior are the same in all regular equilibria, and are given by the ones characterized in Theorem 3.

For any given regular equilibrium σ, Lemmas D.1 and D.2 imply that for every $\theta_g \in \Theta_g$ and $t \in \mathbb{N}$, type θ_g plays $\overline{a}_1$ with probability 1 at $\overline{h}^t$. Let θ^* be the *lowest* $\theta \in \Theta_p \cup \Theta_n$ such that playing $\overline{a}_1$ in every period is type θ's best reply against σ_2. In what follows, I show that $\theta^* \in \Theta_p$.

Suppose towards a contradiction that $\theta^* \in \Theta_n$. In order for type θ^*'s equilibrium payoff to be no less than his minmax payoff $u_1(\theta, \underline{a}_1, \underline{a}_2)$, player 2 needs to play $\overline{a}_2$ with probability 1 at $\overline{h}^t$ for every $t \in \mathbb{N}$. But then, type θ^*'s payoff from playing $\underline{a}_1$ in every period is at least $(1-\delta)u_1(\theta^*, \underline{a}_1, \overline{a}_2) + \delta u_1(\theta^*, \underline{a}_1, \underline{a}_2)$, which is strictly greater than $u_1(\theta^*, \overline{a}_1, \overline{a}_2)$. The latter is type θ^*'s highest possible payoff from playing $\overline{a}_1$ in every period. This leads to a contradiction which implies that $\theta^* \in \Theta_p$.

Let t^* be the smallest $t \in \mathbb{N}$ such that $q_t(\theta) = 0$ for all $\theta \in \Theta_n$. If $t^* \geq 1$, then there exists $\theta_n \in \Theta_n$ such that one of type θ_n's best reply against σ_2 is to play $\overline{a}_1$ from period 0 to t^*-1, from which his payoff is at most:

$$\sum_{t=0}^{t^*-1}(1-\delta)\delta^t u_1(\theta_n,\overline{a}_1,\alpha_{2,t})+(1-\delta)\delta^{t^*}u_1(\theta_n,\underline{a}_1,\alpha_{2,t^*})+\delta^{t^*+1}u_1(\theta_n,\underline{a}_1,\underline{a}_2), \tag{D.2}$$

in which $\alpha_{2,t}\in\Delta(A_2)$ is player 2's action at $\overline{h}^t$. The above payoff must be higher than type θ_n's minmax payoff $u_1(\theta_n,\underline{a}_1,\underline{a}_2)$. Since $u_1(\theta_n,\underline{a}_1,\underline{a}_2)\geq u_1(\theta_n,\overline{a}_1,\overline{a}_2)$, we have:

$$\sum_{t=0}^{t^*-1}(1-\delta)\delta^t\Big(u_1(\theta_n,\overline{a}_1,\overline{a}_2)-u_1(\theta_n,\overline{a}_1,\alpha_{2,t})\Big)\leq(1-\delta)\delta^{t^*}\Big(u_1(\theta_n,\underline{a}_1,\alpha_{2,t^*})-u_1(\theta_n,\overline{a}_1,\overline{a}_2)\Big). \tag{D.3}$$

Since Θ is finite, and u_1 is strictly increasing in a_2, the value of the following expression is positive and finite:

$$\max_{\theta',\theta''\in\Theta}\Big\{\frac{u_1(\theta',\overline{a}_1,\overline{a}_2)-u_1(\theta',\overline{a}_1,\underline{a}_2)}{u_1(\theta'',\overline{a}_1,\overline{a}_2)-u_1(\theta'',\overline{a}_1,\underline{a}_2)}\Big\}.$$

Therefore, (D.3) implies the existence of a constant $C_0>0$ such that for every $\theta\in\Theta^*$,

$$\sum_{t=0}^{t^*-1}(1-\delta)\delta^t\Big(u_1(\theta,\overline{a}_1,\overline{a}_2)-u_1(\theta,\overline{a}_1,\alpha_{2,t})\Big)\leq(1-\delta)C_0. \tag{D.4}$$

For periods after t^*, I examine optimistic and pessimistic ϕ separately.

Case 1: ϕ is Optimistic For every $t\geq t^*$, player 2 does not have a strict incentive to play $\overline{a}_2$ at $\overline{h}^t$ *only if*:

$$\overline{q}\mathcal{D}(\phi_{\overline{a}_1},\overline{a}_1)+\sum_{\theta\in\Theta^*}q^\sigma_{t+1}(\theta)\mathcal{D}(\theta,\overline{a}_1)+\sum_{\theta\in\Theta^*}(q^\sigma_t(\theta)-q^\sigma_{t+1}(\theta))\mathcal{D}(\theta,\underline{a}_1)\leq 0. \tag{D.5}$$

The LHS of (D.5) is a lower bound on the difference between player 2's expected payoff from playing $\overline{a}_2$ and $\underline{a}_2$ at $\overline{h}^t$. Since ϕ is optimistic and every type in Θ_g plays $\overline{a}_1$ with probability 1 at every $\overline{h}^t$, we know that for every $t\geq t^*$, $\sum_{\theta\in\Theta}q^\sigma_t(\theta)\mathcal{D}(\theta,\overline{a}_1)=\sum_{\theta\in\Theta^*}q^\sigma_t(\theta)\mathcal{D}(\theta,\overline{a}_1)\geq 0$. Therefore, (D.5) implies that:

$$\overline{q}\mathcal{D}(\phi_{\overline{a}_1},\overline{a}_1)\leq\sum_{\theta\in\Theta^*}\big(q^\sigma_t(\theta)-q^\sigma_{t+1}(\theta)\big)\big(\mathcal{D}(\theta,\overline{a}_1)-\mathcal{D}(\theta,\underline{a}_1)\big),$$

or equivalently,

$$\sum_{\theta\in\Theta^*}\big(q^\sigma_t(\theta)-q^\sigma_{t+1}(\theta)\big)\geq C_1\equiv\frac{\overline{q}\mathcal{D}(\phi_{\overline{a}_1},\overline{a}_1)}{\max_{\theta\in\Theta^*}\{\mathcal{D}(\theta,\overline{a}_1)-\mathcal{D}(\theta,\underline{a}_1)\}}. \tag{D.6}$$

The MS condition implies that $\max_{\theta\in\Theta^*}\{\mathcal{D}(\theta,\overline{a}_1)-\mathcal{D}(\theta,\underline{a}_1)\}>0$. Therefore, C_1 is a strictly positive constant, which is independent of δ.

Since $\sum_{\theta \in \Theta_p} q_t^\sigma(\theta) \leq 1$ for all $t \in \mathbb{N}$, the number of periods such that $\overline{a}_2$ is not a strict best reply at $\overline{h}^t$ is no more than $\lceil 1/C_1 \rceil$. Therefore, for every $\theta \in \Theta^*$, if type θ plays $\overline{a}_1$ in every period, his loss relative to $u_1(\theta, \overline{a}_1, \overline{a}_2)$ is no more than $(1-\delta)C_0$ from period 0 to t^*-1, and is no more than $1-\delta^{\lceil 1/C_1 \rceil}$ after period t^*. As a result, his payoff in any regular equilibrium is no less than $u_1(\theta, \overline{a}_1, \overline{a}_2)$ as $\delta \to 1$.

Case 2: ϕ is Pessimistic For every regular equilibrium σ, let

$$\mathcal{I}_t^\sigma \equiv \overline{q}\mathcal{D}(\phi_{\overline{a}_1}, \overline{a}_1) + \sum_{\theta \in \Theta} q_t^\sigma(\theta)\mathcal{D}(\theta, \overline{a}_1). \tag{D.7}$$

Given that $\theta^* \in \Theta_p$, which implies that $\Theta_p \neq \{\varnothing\}$, Lemmas D.1 and D.2 imply that first, $\mathcal{D}(\theta, \overline{a}_1) > 0$ only if $\theta \in \Theta_g$, and second, every type in Θ_g plays $\overline{a}_1$ with probability 1 at every $\overline{h}^t$. Therefore, $\mathcal{I}_t^\sigma$ is *nondecreasing* in t for every regular equilibrium σ. I show the following lemma:

Lemma D.3. *If ϕ is pessimistic and the probability of commitment types is small enough, then in every regular equilibrium σ, $t^* = 1$ and $\mathcal{I}_t^\sigma = 0$ for every $t \geq 1$.*

Proof of Lemma D.3: I proceed in two steps. In step 1, I show that $\mathcal{I}_t^\sigma = 0$ for every $t \geq \max\{1, t^*\}$. In step 2, I show that $t^* = 1$. The two steps together lead to the conclusion of Lemma D.3.

Step 1: First, suppose that $\mathcal{I}_t^\sigma < 0$ for some $t \geq \max\{1, t^*\}$. Since player 2's belief is a martingale, there exists $\theta_p \in \Theta_p$ such that (1) $q_t^\sigma(\theta_p) > 0$, and (2) under one of type θ_p's pure strategy best replies to σ_2, $\underline{a}_2$ is player 2's strict best reply in all subsequent periods. If type θ_p plays according to this pure strategy best reply against σ_2, then his continuation payoff at $\overline{h}^t$ is $u_1(\theta, \underline{a}_1, \underline{a}_2)$, which is his minmax payoff. He can profitably deviate at $\overline{h}^{t-1}$ by playing $\underline{a}_1$ in every period, which leads to a strictly higher stage-game payoff in period $t-1$. This leads to a contradiction.

Next, suppose that $\mathcal{I}_t^\sigma > 0$ for some $t \geq \max\{1, t^*\}$. I start from showing that player 2 has a strict incentive to play $\overline{a}_2$ at $\overline{h}^s$ for every $s \geq t$. Suppose towards a contradiction that $\overline{a}_2$ is not a strict best reply at $\overline{h}^{s_0}$ for some $s_0 \geq t$, then there exists $\theta_p \in \Theta_p$ such that type θ_p plays $a_1 \neq \overline{a}_1$ with positive probability at $\overline{h}^{s_0}$. Similar to (D.6), we have:

$$\sum_{\theta \in \Theta_p} \left(q_{s_0}^\sigma(\theta) - q_{s_0+1}^\sigma(\theta)\right) \geq \frac{\mathcal{I}_{s_0}^\sigma}{\max_{\theta \in \Theta_p}\{\mathcal{D}(\theta, \overline{a}_1) - \mathcal{D}(\theta, \underline{a}_1)\}}. \tag{D.8}$$

Since all types in Θ_g play $\overline{a}_1$ with probability 1, type θ_p's payoff at $\overline{h}^{s_0}$ is at most $(1-\delta)u_1(\theta_p, \underline{a}_1, \overline{a}_2) + \delta u_1(\theta_p, \underline{a}_1, \underline{a}_2)$. Since $u_1(\theta_p, \underline{a}_1, \underline{a}_2) < u_1(\theta_p, \overline{a}_1, \overline{a}_2)$, type θ_p has an incentive to play $\underline{a}_1$ instead of $\overline{a}_1$ at $\overline{h}^{s_0}$ only if there exists $s_1 > s_0$ such that $\overline{a}_2$ is not a strict best reply at $\overline{h}^{s_1}$. Iterate this process, one can obtain an *infinite sequence* $\{s_0, s_1, s_2, ...\}$ such that for every $i \in \mathbb{N}$,

$$\sum_{\theta \in \Theta_p} \left(q^{\sigma}_{s_i}(\theta) - q^{\sigma}_{s_i+1}(\theta) \right) \geq \frac{\mathcal{I}^{\sigma}_{s_i}}{\max_{\theta \in \Theta_p} \{ \mathcal{D}(\theta, \overline{a}_1) - \mathcal{D}(\theta, \underline{a}_1) \}}.$$

Since $\mathcal{I}^{\sigma}_s$ is non-decreasing in s, the RHS is bounded away from 0. Since $\sum_{\theta \in \Theta_p} q^{\sigma}_s(\theta) \leq 1$ for every $s \in \mathbb{N}$, this sequence *ends in finite time*. This contradiction implies that $\overline{a}_2$ is a strict best reply at $\overline{h}^s$ for every $s \geq t$.

Therefore, for every $\theta \in \Theta^*$ and $t \in \mathbb{N}$, by playing $\overline{a}_1$ in every period, type θ's payoff at $\overline{h}^t$ is no less than $u_1(\theta, \overline{a}_1, \overline{a}_2) - (1-\delta)C_0$, where C_0 is the constant defined in (D.4). Since ϕ is pessimistic, $\mathcal{I}^{\sigma}_0 < 0$ when the total probability of commitment types is small enough. Suppose $\mathcal{I}^{\sigma}_t > 0$ for some $t \geq \max\{1, t^*\}$, then there exists $\theta_p \in \Theta_p$ that plays $a_1 \neq \overline{a}_1$ with positive probability at $\overline{h}^s$ for some $s < t$, after which his continuation payoff is $u_1(\theta_p, \underline{a}_1, \underline{a}_2)$. This payoff is strictly less compared to his continuation payoff from playing $\overline{a}_1$ in every period, which leads to a contradiction. This implies that $\mathcal{I}^{\sigma}_t = 0$ for every $t \geq \max\{1, t^*\}$.

Step 2: I show that $t^* = 1$. Suppose towards a contradiction that $t^* > 1$. Then there exists $\theta_n \in \Theta_n$ whose best response to σ_2 is to play $\overline{a}_1$ until period $t^* - 1$. Since $\mathcal{I}^{\sigma}_t = 0$ for all $t \geq t^*$ and $\mathcal{I}^{\sigma}_t$ is non-decreasing in t, then $\mathcal{I}^{\sigma}_t < 0$ for all $t < t^*$. This implies that type θ_n's payoff following his equilibrium strategy is at most $(1 - \delta^{t^*-1}) u_1(\theta_n, \overline{a}_1, \underline{a}_2) + \delta^{t^*-1} u_1(\theta_n, \underline{a}_1, \underline{a}_2)$, which is strictly less than his minmax payoff $u_1(\theta_n, \underline{a}_1, \underline{a}_2)$. This leads to a contradiction. □

Lemma D.3 implies that in every regular equilibrium under a pessimistic ϕ, first, all types in Θ_n play $\underline{a}_1$ in every period. Second, player 2's posterior after observing $\overline{a}_1$ is such that she is indifferent between $\overline{a}_2$ and $\underline{a}_2$ against $\overline{a}_1$. Third, $\mathcal{I}^{\sigma}_t = 0$ for all $t \geq 1$ implies that if player 1 plays $\overline{a}_1$ in period 0 on the equilibrium path, then he plays $\overline{a}_1$ at $\overline{h}^t$ for every $t \in \mathbb{N}$.

According to Lemma D.2, for every regular equilibrium σ, there exists a cutoff state $\theta^* \in \Theta^*$ such that strategic player 1 plays $\overline{a}_1$ in every period with probability 1 when $\theta \succ \theta^*$. When the total probability of commitment types is sufficiently small, and given that player 1's characteristics and the state are *independent*, the cutoff state equals $\theta^*(\phi)$, defined in (3.10), for all regular equilibria under ϕ. According to the definition of $\theta^*(\phi)$, we have

$$\sum_{\theta' \succeq \theta^*(\phi)} \phi(\theta') \mathcal{D}(\theta', \overline{a}_1) < 0 \quad \text{and} \quad \sum_{\theta' \succ \theta^*(\phi)} \phi(\theta') \mathcal{D}(\theta', \overline{a}_1) \geq 0. \tag{D.9}$$

This implies that $\theta^*(\phi) \in \Theta_p$. Since $\mathcal{I}^{\sigma}_1 = 0$, $\sum_{\theta' \succ \theta^*(\phi)} \phi(\theta') \mathcal{D}(\theta', \overline{a}_1) \geq 0$, and $\mathcal{D}(\phi_{\overline{a}_1}, \overline{a}_1) > 0$, the cutoff type $\theta^*(\phi)$ plays $\overline{a}_1$ with strictly positive probability in period 0. This implies that the cutoff type's equilibrium payoff is bounded from above by $u_1(\theta^*(\phi), \overline{a}_1, \overline{a}_2)$.

For every $a_1^* \in A_1^g \backslash \{\overline{a}_1, \underline{a}_1\}$, after observing player 1 playing a_1^* in period 0, I show that player 2 must be indifferent between $\overline{a}_2$ and $\underline{a}_2$ against a_1^*. This is because if player 2 strictly prefers $\overline{a}_2$, then a similar argument to Step 1 in the proof of Lemma D.3 implies that type $\theta^*(\phi)$'s discounted average payoff by playing a_1^* in every period converges to $u_1(\theta^*(\phi), a_1^*, \overline{a}_2)$ as $\delta \to 1$, which is strictly greater than his payoff from playing $\overline{a}_1$ in every period. This leads to a contradiction. If player 2 strictly prefers $\underline{a}_2$, then according to the definition of A_1^g, there exists $\theta \prec \theta^*(\phi)$ such that strategic type θ plays a_1^* with positive probability in period 0, and his continuation payoff in period 1 is no more than $u_1(\theta, \underline{a}_1, \underline{a}_2)$. Therefore, this type strictly prefers to play $\underline{a}_1$ in period 0, which leads to a contradiction.

Player 2's indifference after observing $a_1^* \in A_1^g \backslash \{\overline{a}_1, \underline{a}_1\}$ implies that in any regular equilibrium, if a strategic type $\theta \preceq \theta^*(\phi)$ plays a_1^* with positive probability in period 0, then he plays a_1^* with probability 1 at every subsequent history such that he has played a_1^* in all previous periods.

When the total probability of commitment types is sufficiently small and given the MS condition on stage-game payoffs, type $\theta^*(\phi)$ plays every action in $A_1^g \backslash \{\overline{a}_1, \underline{a}_1\}$ with positive probability. Since player 2 is indifferent between $\overline{a}_2$ and $\underline{a}_2$ against a_1^* after observing $a_1^* \in A_1^g \backslash \{\overline{a}_1, \underline{a}_1\}$, the probability with which strategic type $\theta^*(\phi)$ plays actions in $A_1^g \backslash \{\overline{a}_1, \underline{a}_1\}$ converges to 0 as the probability of commitment type goes to 0. This together with (4.9) implies that strategic type $\theta^*(\phi)$ plays $\underline{a}_1$ with positive probability in period 0.

When the total probability of commitment types is small enough, player 2 has no incentive to play $\overline{a}_2$ after observing $\underline{a}_1$, which implies that type $\theta^*(\phi)$'s equilibrium payoff equals its minmax payoff. Therefore, playing $\underline{a}_1$ in every period is type $\theta^*(\phi)$'s best reply against σ_2. Lemma D.2 then implies that types lower than $\theta^*(\phi)$ plays $\underline{a}_1$ with probability 1 in every period. The cutoff type's equilibrium payoff pins down the discounted average frequency with which player 2 plays $\overline{a}_2$ conditional on player 1 plays $\overline{a}_1$ in every period, given by $r(\phi)$. This together with types $\theta \succ \theta^*(\phi)$'s on-path behavior pins down every type's payoff in all regular equilibria.

D.4 Analysis of Irregular Equilibria

I show that first, if ϕ is optimistic, then there exists $C \in \mathbb{R}_+$ such that type $\theta \in \Theta^*$'s payoff in any irregular equilibrium is at least $u_1(\theta, \overline{a}_1, \overline{a}_2) - (1-\delta)C$. This together with the conclusion on regular equilibria establishes Theorem 2. Second, if ϕ is pessimistic, then irregular equilibria do not exist. Therefore, the unique payoff and unique on-path behavior in regular equilibria apply to *all* equilibria.

Recall that $t^* \in \mathbb{N}$ is the smallest $t \in \mathbb{N}$ such that $q_t^\sigma(\theta) = 0$ for all $\theta \in \Theta_n$. Similar to the analysis of regular equilibria, there exists a constant $C_0 > 0$ such that in every equilibrium for every $\theta \in \Theta^*$, type θ's loss from period 0 to t^* is no more than $(1-\delta)C_0$ relative to $u_1(\theta, \overline{a}_1, \overline{a}_2)$. I show the following lemma:

Lemma D.4. *There exists $C_2 > 0$ such that for every $t \geq t^*$, if $\sum_{\theta \in \Theta^*} q_t^\sigma(\theta)\mathcal{D}(\theta, \overline{a}_1) \geq 0$ and $\overline{a}_2$ is not a* *strict best reply at $\overline{h}^t$, then*

$$\sum_{\theta\in\Theta^*}\Big(q_t^\sigma(\theta)-q_{t+1}^\sigma(\theta)\Big)\geq C_2. \tag{D.10}$$

Proof of Lemma D.4: By definition, when $t\geq t^*$, $q_t^\sigma(\theta)=0$ for every $\theta\notin\Theta^*$. If $\overline{a}_2$ is not a strict best reply at $\overline{h}^t$, then

$$\overline{q}\mathcal{D}(\phi_{\overline{a}_1},\overline{a}_1)+\sum_{\theta\in\Theta^*}q_{t+1}^\sigma(\theta)\mathcal{D}(\theta,\overline{a}_1)+\sum_{\theta\in\Theta^*}(q_t^\sigma(\theta)-q_{t+1}^\sigma(\theta))\mathcal{D}(\theta,\underline{a}_1)\leq 0, \tag{D.11}$$

where the LHS is a lower bound on player 2's relative payoff from playing $\overline{a}_2$ instead of $\underline{a}_2$ at history $\overline{h}^t$. Inequality (D.11) can be rewritten as:

$$\overline{q}\mathcal{D}(\phi_{\overline{a}_1},\overline{a}_1)+\sum_{\theta\in\Theta^*}q_t^\sigma(\theta)\mathcal{D}(\theta,\overline{a}_1)+\sum_{\theta\in\Theta^*}(q_t^\sigma(\theta)-q_{t+1}^\sigma(\theta))\Big(\mathcal{D}(\theta,\underline{a}_1)-\mathcal{D}(\theta,\overline{a}_1)\Big)\leq 0, \tag{D.12}$$

which implies:

$$\sum_{\theta\in\Theta^*}\Big(q_t^\sigma(\theta)-q_{t+1}^\sigma(\theta)\Big)\geq C_2\equiv\frac{\overline{q}\mathcal{D}(\phi_{\overline{a}_1},\overline{a}_1)}{\max_{\theta\in\Theta^*}\{\mathcal{D}(\theta,\overline{a}_1)-\mathcal{D}(\theta,\underline{a}_1)\}}. \tag{D.13}$$

□

Next, I establish a uniform lower bound on player 2's posterior belief about the state at $\overline{h}^t$.

Lemma D.5. *In every irregular equilibrium, $\sum_{\theta\in\Theta^*}q_t^\sigma(\theta)\mathcal{D}(\theta,\overline{a}_1)\geq 0$ for every $t\in\mathbb{N}$.*

Proof of Lemma D.5: The definition of irregular equilibrium implies the existence of $t\in\mathbb{N}$ such that $q_t^\sigma(\theta)=0$ for all $\theta\in\Theta_p\cup\Theta_n$. Let $\bar{t}$ be the smallest of such t. By definition, $\sum_{\theta\in\Theta^*}q_t^\sigma(\theta)\mathcal{D}(\theta,\overline{a}_1)\geq 0$ for every $t\geq\bar{t}$.

Suppose towards a contradiction that there exists $t\leq\bar{t}$ such that $\sum_{\theta\in\Theta^*}q_t^\sigma(\theta)\mathcal{D}(\theta,\overline{a}_1)<0$. Let $\widehat{t}$ be the largest of such t. By definition, $\sum_{\theta\in\Theta^*}q_{\widehat{t}}^\sigma(\theta)\mathcal{D}(\theta,\overline{a}_1)<0$, but $\sum_{\theta\in\Theta^*}q_t^\sigma(\theta)\mathcal{D}(\theta,\overline{a}_1)\geq 0$ for all $t>\widehat{t}$. Consider player 1's incentives at history $\overline{h}^{\widehat{t}}$.

1. By definition, there exists $a_1\neq\overline{a}_1$ that is played with positive probability by some types in Θ^* at $\overline{h}^{\widehat{t}}$, such that player 2's belief at $(\overline{h}^{\widehat{t}},a_1)$ is pessimistic.

2. For every $\theta\in\Theta^*$ such that $q_{\widehat{t}}^\sigma(\theta)>0$, if type θ plays $\overline{a}_1$ in every period, then his continuation payoff at $\overline{h}^{\widehat{t}}$ is at least $u_1(\theta,\overline{a}_1,\overline{a}_2)-(1-\delta)C_0-(1-\delta^{1/C_2})$, which converges to 1 as $\delta\to 1$.

The two parts together imply that when δ is close enough to 1, there exists $\theta\in\Theta^*$ with $q_{\widehat{t}}^\sigma(\theta)>0$, such that type θ receives a strictly higher continuation payoff by playing $\overline{a}_1$ at $\overline{h}^{\widehat{t}}$, compared to playing a_1 at $\overline{a}_1$. This violates his incentive constraint. This leads to a contradiction. □

D.5 Summary & Overview of Full Proof

When ϕ is optimistic, Lemma D.5 implies that $\sum_{\theta\in\Theta^*} q_t^\sigma(\theta)\mathcal{D}(\theta,\overline{a}_1) \geq 0$ for every $t \in \mathbb{N}$ and in every irregular equilibrium. Lemma D.4 then implies that for every $\theta \in \Theta^*$, type θ's guaranteed payoff in any irregular equilibrium is no less than $u_1(\theta,\overline{a}_1,\overline{a}_2)$ as $\delta \to 1$. This together with the analysis of regular equilibria establishes Theorem 2. When ϕ is pessimistic, Lemma D.5 implies that irregular equilibria do not exist and all equilibria are regular. Therefore, player 1's equilibrium payoff and on-path behavior in regular equilibria (see section D.3) are player 1's unique payoff and unique on-path behavior for all equilibria.

When player 2's strategy depends on her predecessors' actions, extra complications arise in the analysis of irregular equilibria, such as the proof of Lemma D.5. This is because conditional on player 1 playing $\overline{a}_1$ in every period, there may not exist a *last history* at which player 2's posterior belief about the state is pessimistic.

To overcome this challenge, I show that every time a switching from a pessimistic to an optimistic belief happens, strategic types in Θ_p must be separating from commitment type $\overline{a}_1$ with an ex ante probability bounded from below. This implies that such switching happens at most in a finite number of times conditional on every realized path of play. On the other hand, strategic types in Θ_p only have incentives to separate at those switching histories when his continuation payoff from imitating type $\overline{a}_1$ is low. This implies that there exists at least another switching following that history, meaning that such switching happens infinitely many times if it happens once. This leads to a contradiction.

References

[1] Aumann, Robert and Michael Maschler (1995) *Repeated Games with Incomplete Information*, MIT Press.

[2] Bar-Isaac, Heski (2003) "Reputation and Survival: Learning in a Dynamic Signalling Model," *Review of Economic Studies*, 70(2), 231-251.

[3] Beamish, Paul and Hari Bapuji (2008) "Toy Recalls and China: Emotion vs Evidence," *Management and Organization Review*, 4(2), 197-209.

[4] Cripps, Martin, George Mailath and Larry Samuelson (2004) "Imperfect Monitoring and Impermanent Reputations," *Econometrica*, 72(2), 407-432.

[5] Darby, Michael and Edi Karni (1973) "Free Competition and the Optimal Amount of Fraud," *Journal of Law and Economics*, 16(1), 67-88.

[6] Deb, Joyee and Yuhta Ishii (2019) "Reputation Building under Uncertain Monitoring," Working Paper.

[7] Ekmekci, Mehmet (2011) "Sustainable Reputations with Rating Systems," *Journal of Economic Theory*, 146(2), 479-503.

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

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

[10] Faingold, Eduardo and Yuliy Sannikov (2011) "Reputation in Continuous-Time Games," *Econometrica*, 79(3), 773-876.

[11] Freedman, Seth, Melissa Kearney and Mara Lederman (2012) "Product Recalls, Imperfect Information, and Spillover Effects: Lessons from the Consumer Response to Toy Recalls," *Review of Economics and Statistics*, 94(2), 499-516.

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

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

[14] Fudenberg, Drew and David Levine (1992) "Maintaining a Reputation when Strategies are Imperfectly Observed," *Review of Economic Studies*, 59(3), 561-579.

[15] Fudenberg, Drew and David Levine (2009) "Repeated Games with Frequent Signals," *Quarterly Journal of Economics*, 124(1), 233-265.

[16] Gossner, Olivier (2011) "Simple Bounds on the Value of a Reputation," *Econometrica*, 79(5), 1627-1641.

[17] Hart, Sergiu (1985) "Nonzero-Sum Two-Person Repeated Games with Incomplete Information," *Mathematics of Operations Research*, 10(1), 117-153.

[18] Hansman, Christopher, Jonas Hjort, Gianmarco León, and Matthieu Teachout (2018) "Vertial Integration, Supplier Behavior, and Quality Upgrading Among Exporters," NBER Working Paper 23949.

[19] Helpman, Elhanan (2006) "Trade, FDI, and the Organization of Firms," *Journal of Economic Literature*, 44(3), 589-630.

[20] Hoffer, George, Stephen Pruitt, and Robert Reilly (1987) "Antomotive Recalls and Informational Efficiency," *The Financial Review*, 22(4), 433-442.

[21] Hörner, Johannes and Stefano Lovo (2009) "Belief-Free Equilibria in Games with Incomplete Information," *Econometrica*, 77(2), 453-487.

[22] Jarrell, Gregg and Sam Peltzman (1985) "The Impact of Product Recalls on the Wealth of Sellers," *Journal of Political Economy*, 93(3), 512-536.

[23] Kaya, Ayça (2009) "Repeated Singalling Games," *Games and Economic Behavior*, 66, 841-854.

[24] Keller, Godfrey and Sven Rady (2010) "Breakdowns," *Theoretical Economics*, 10, 175-202.

[25] Kreps, David and Robert Wilson (1982) "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253-279.

[26] Lee, Jihong and Qingmin Liu (2013) "Gambling Reputation: Repeated Bargaining with Outside Options," *Econometrica*, 81(4), 1601-1672.

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

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

[29] Liu, Shuo and Harry Pei (2020) "Monotone Equilibria in Signalling Games," *European Economic Review*, forthcoming.

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

[31] Milgrom, Paul and John Roberts (1982) "Predation, Reputation and Entry Deterence," *Journal of Economic Theory*, 27, 280-312.

[32] Nelson, Phillip (1970) "Information and Consumer Behavior," *Journal of Political Economy*, 78(2), 311-329.

[33] Ni, John, Barbara Flynn and Robert Jacobs (2016) "The Effect of a Toy Industry Recall Annoucement on Shareholder Wealth," *International Journal of Production Research*, 54(18), 5404-5415.

[34] Pei, Harry (2018) "Essays on Dynamic Games and Reputations," Doctoral Thesis, Massachusetts Institute of Technology.

[35] Pęski, Marcin (2014) "Repeated Games with Incomplete Information and Discounting," *Theoretical Economics*, 9, 651-694.