Reputation and screening in a noisy environment with irreversible actions

Ekmekci, Mehmet and Maestri, Lucas

Boston College, FGV EPGE

28 August 2019

Reputation and Screening in a Noisy Environment with Irreversible Actions*

Mehmet Ekmekci† and Lucas Maestri ‡

August 28, 2019

Abstract

We introduce a class of two-player dynamic games to study the effectiveness of screening in a principal-agent problem. In every period, the principal chooses either to irreversibly stop the game or to continue, and the agent chooses an action if the principal chooses to continue. The agent's type is his private information, and his actions are imperfectly observed. Players' flow payoffs depend on the agent's action, and players' lump-sum payoffs when the game stops depends on the agent's type. Both players are long-lived and share a common discount factor. We study the limit of the equilibrium outcomes as both players get arbitrarily patient. Nash equilibrium payoff vectors converge to the unique Nash equilibrium payoff vector of an auxiliary, two-stage game with observed mixed actions. The principal learns some but not all information about the agent's type. Any payoff-relevant information revelation takes place at the beginning of the game. We calculate the probability that the principal eventually stops the game, against each type of the agent.

Keywords: Dynamic Games, Screening, Reputation, Imperfect Monitoring.

JEL Codes: C72, C73, D82, D86.

*We thank Alp Atakan, Leandro Gorno, Harry Di Pei, Tristan Tomala and Nicolas Vieille for their comments. Ekmekci thanks the MIT Economics Department for their hospitality. We are thankful for comments by seminar audiences at Yale, ITAM, SITE conference, WUSTL, FGV EPGE, Collegio Alberto, PSE, HEC, and Stony Brook game theory conference.

†Boston College, 140 Commonwealth Ave., Chestnut Hill, MA 02467. Email: ekmekci@bc.edu

‡FGV EPGE, Praia de Botafogo 190, Rio de Janeiro, Brazil. Email: maestri.lucas@gmail.com

Contents

1. Introduction

We introduce a class of two-player dynamic games to study the effectiveness of screening in a principal-agent problem. The principal (player 2, or "she"), at each period, can either take an irreversible action that ends the dynamic game or wait for at least another period. The agent (player 1, or "he") takes one of a finite number of actions in any period in which the game has not yet ended. The agent is either a commitment type, who is committed to playing a fixed mixed action at every period of the game or a normal type, who maximizes his expected discounted payoff. The principal's payoff from stopping the game depends on the agent's type. Therefore, she may choose to delay ending the game if she expects that different types of the agent will play different strategies, which would lead to the revelation of information about the agent's type. She may also choose to not end the game because ending the game may be costly given her beliefs about the agent's type. Both players are long-lived and discount the future using a common discount factor. Moreover, the agent's actions are possibly observed with noise. We ask three main questions: How much does the principal learn about the agent's type? How much does the principal benefit from waiting? How fast does the agent reveal any private information?

We consider two types of situations. In the first scenario, which we call *screening game*, the agent prefers that the irreversible action be taken, while the principal prefers not to take the irreversible action against the normal type of the agent. A prominent example of such situations is that of a relationship between a firm and a worker in a nonflexible labor market. The firm hires the worker using short-term contracts and prefers to offer a long-term contract only if the worker is highly skilled. A low-skilled worker can exert extra effort to imitate a high-skilled worker, and the firm suffers a loss if it offers the long-term contract to a low–skilled worker. This may be because ending a long-term contract is very costly, or simply because it is not an option. Similarly, a firm that subcontracts with a supplier may want to take over the supplier if the latter owns an advanced-technology product line but may prefer subcontracting otherwise. As a result, the subcontractor may undertake costly activities to generate outputs that are similar to that coming from an advanced-technology product line.[1]

In the second scenario, which we call *contracting game*, the agent prefers that the irreversible action is not taken and the principal prefers to take the irreversible action

[1]Examples of screening with irreversible actions are, internships and the decision to hire, dating and the decision to marry, academic tenure-track position and the decision to promote and venture-capital financing. In some of these examples the principal may have limited commitment power. The case of promotion to tenure also comes with the option of firing/rescheduling, and we discuss such a scenario in Subsection 10.1.

against the normal type of the agent. Consider a firm that offers an incentive contract to a worker who may have high or low productivity. If the firm learns that the worker's productivity is high, it would be tempted to change the terms of the contract and reduce the worker's surplus. Therefore, a high-productivity worker has incentives to mimic low productivity workers, even if by doing so he forgoes some surplus. Such situations have been extensively studied in the ratchet effect literature.

The class of games we study resemble an experimentation problem. The principal's payoff from stopping the game depends on her opponent's type, hence she would like to gather a lot of information before taking the irreversible action. On the other side, the information flow is determined by the agent's actions, hence is endogenous. Importantly, the principal and the normal type of the agent have opposing preferences on the irreversible action, which creates the incentives for the agent to pool with a commitment type. This assumption on the payoff functions of the players implies that the principal's optimal strategy against the normal type of the agent is independent of the principal's expectations of what the normal type would do if the principal never took the irreversible action. Hence, throughout the game, if the agent's type is revealed, the equilibrium continuation play is unique. This property in our model helps to eliminate coordination issues that often lead to multiplicity of equilibria in repeated games.

Every Nash equilibrium induces a reduced equilibrium outcome, which specifies for each type and action profile, the expected discounted number of periods in which this action profile is played against this type. Our main result characterizes reduced equilibrium outcomes when players become arbitrarily patient. We show that, reduced equilibrium outcomes converge to the the unique reduced equilibrium outcome of an auxiliary two-stage game. In this auxiliary game player 1 (the agent) chooses a mixed action, and player 2 (the principal), after observing player 1's mixed action, chooses W (wait) or S (stop). Players' payoff functions are equal to the flow payoff functions in the dynamic game if player 2 chooses W, and equal to payoffs when the dynamic game stops, if player 2 chooses S. The distribution over commitment types in the auxiliary game is identical to that in the dynamic game. Our main result implies that equilibrium payoff vectors of the dynamic game converge, as the players become arbitrarily patient, to the unique equilibrium payoff vector of the auxiliary two-stage game.

We further characterize the evolution of the posterior beliefs of player 2 throughout the game. We show that in screening games, when the players are patient, for every belief and action profile, the expected discounted number of periods in which the posterior belief is equal to this belief and the action profile is equal to this action profile converges to the corresponding equilibrium probability in the auxiliary two-stage game. In contracting games, we obtain a similar result for action profiles in which player 2 plays W.

Our characterization results on equilibrium behavior and evolution of beliefs show that any payoff-relevant information revelation takes place almost immediately at the beginning of the game. The equilibrium behavior is akin to having an initial signaling period in which player 1 commits to a stationary strategy, followed by player 2's best response to it given her posterior belief about the agent's type. Hence, reputation building is not gradual. Moreover, our main result facilitates simple and tractable comparative statics of the equilibrium payoffs in the dynamic game with respect to the payoff functions and the distribution over types. Hence, the effectiveness of dynamic screening for the principal can be understood by solving a simple two-stage game.

The intuition for why reputation building is not gradual is as follows. In screening games, the agent may invest in reputation by mimicking a commitment type with a positive probability. If he chooses to do so, then he pools with that commitment type, which eventually leads to the irreversible action. Building reputation gradually would require that after some initial investment, the agent loses his reputation with some probability in a later round. However, an initial investment is not optimal if the agent expects to be indifferent between further investment and losing his reputation later on. In contracting games, once the agent reveals his private information, irreversible action is taken, and this leads to the worst possible outcome for the agent. Hence, the agent finds it optimal to pick which commitment type to mimic at the beginning, and not to reveal any further information.[2]

The main strategic tensions that arise in our model are familiar from the reputation literature with imperfect monitoring (Cripps et al. (2004, 2007)), and from the ratchet effect literature (Laffont and Tirole (1987, 1988, 1993); Gerardi and Maestri (2018), Acharya and Ortner (2017)). In our model, the agent prefers being thought of as one of several commitment types, instead of being thought of as the normal type. If the principal expects to learn much information about the agent's type, and acting on this information worsens the agent's continuation payoff, the agent has strong incentives to imitate a commitment type and build a reputation. This prevents information revelation, even if imitating has a short-run cost to the agent. However, there is a countervailing force that may incentivize the agent to reveal information. If the agent's actions are monitored with noise, and if the agent is not expected to reveal any information for some long duration of play, then

[2] These intuitions are precise only when player 1's actions are observed perfectly. When player 1's actions are observed with noise player 1 will never reveal his type completely.

the principal interprets any signal she observes as pure noise. Hence, the agent may have incentives to play his myopic best response instead of mimicking the commitment type. This contradicts the expectation of no information revelation. The tension between the agent's short-run temptations to reveal information–and long-run benefits from building a reputation–together with the principal's incentives to screen the agent, lead to a nontrivial resolution of equilibrium dynamics.

We further analyze the long-run behavior when the game is a screening game, or when the game is a contracting game and the monitoring structure satisfies the full-support assumption. The characterization result described above does not pin down the long-run behavior. This is however an important question. For instance, in some of our main examples, one would like to know the long-run probability that each type of the agent is eventually promoted or demoted. These predictions can help a modeler who has data only about the aggregate promotion and demotion decisions of a principal against different worker types, and does not have individual specific data about the duration of each relationship. In Theorems 4 and 5, which are our second set of main results, we find the probability that player 2 chooses action S against each type in the long run when players are patient.

These results are stated separately for the two classes of games, because in contracting games long-run behavior depends on whether the monitoring structure has full support or not, while in screening games long-run behavior does not depend on the full-support assumption. Let us start with contracting games. We first observe that throughout the game, the normal type of player 1 reveals information, if the monitoring structure has full support. This insight was provided in Cripps et al. (2004), and we adapt their findings that were provided in the setting of repeated games to our setting. We find that, if players are sufficiently patient, information revelation by player 1 becomes arbitrarily slow when the posterior beliefs about player 1's type are close to a set of beliefs that we call *reputation boundaries*. A belief about player 1's types is a reputation boundary if it puts positive probability only on the normal type and one commitment type, and at that belief, both S and W are myopic best responses to the commitment type's action. We show that whenever player 2 plays S with some delay, then when she plays S, the posterior beliefs are arbitrarily close to a reputation boundary. This finding and the first observation, together with Bayes' rule, allow us to calculate the long-run probability that S is taken against each type.

In screening games, the intuition is different, and the result holds even under perfect monitoring. Our first payoff and behavior characterization result described above imply that in screening games with patient players, if player 2 plays W with a positive probability against the normal type, then player 1's equilibrium behavior is akin to an initial randomization over revealing his type, and mimicking one of the commitment types. If player 1 reveals his type, then player 2 plays W in every future period. If he mimics a commitment type, player 2's posterior belief is close to a reputation boundary. We show that when beliefs reach a reputation boundary, they can only move arbitrarily slowly around this reputation boundary. This is possible only if player 1 pools with the commitment type in the future periods. However, in a screening game, pooling with a commitment type has a short-run cost to player 1, so pooling can be sustained only with expectations of rewards coming in the form of S being played in some not too distant future. Because such expectations have to be sustained at every period, player 2 cannot indefinitely play W, and plays S against the commitment type and the normal type eventually. Because beliefs do not fluctuate much around the reputation boundary with patient players, we can calculate the probability with which S is played against each type in the long run.

Our benchmark models can be further used as building blocks to explore equilibrium behavior and payoffs in more complicated dynamic models. In Section 10, we discuss an example that combines both scenarios: When the principal stops the game, she can take one of two possible irreversible actions. She can either promote the agent, or demote (or reschedule) the agent. The principal prefers to promote the commitment type, and prefers to demote the normal type. We show that screening is not effective in this model. In this section, we also discuss other related scenarios that our model does not cover. We discuss the related literature in more detail in Section 9. All the proofs are in the Appendices.

2. The Model

Two players interact in a dynamic relationship that takes place in discrete time, $t \in \{0, 1, ...\}$. In the beginning of period 0, player 2 chooses whether to take an irreversible action (S) or not (W). If player 2 takes the irreversible action, then the game ends. Otherwise, player 1 takes an action $a \in \mathbb{A} = \{a_0, a_1, ..., a_k\}$. At the end of each period in which S has not been played, a public signal y from a finite set of signals $\mathcal{Y}$ is observed. The probability distribution of the signal depends on player 1's action, and we write $P(y \mid a)$ for this distribution.

After observing the signal y_0 at the end of period 0, player 2 chooses whether to play S or W in the beginning of period 1. If she plays S the game ends. Otherwise a similar dynamic unfolds. In general, if S has not been taken by player 2 before period t, then at period t, first player 2 decides whether to play S or W. If she chooses W, then player 1 takes an action $a_t \in \mathbb{A}$ that leads to the public signal $y_t \in \mathcal{Y}$. If she chooses S the game ends.[3]

For any mixed action of player 1, $\alpha \in \Delta\mathbb{A}$, and any signal $y \in \mathcal{Y}$, let $P(y \mid \alpha) := \sum_{a \in \mathbb{A}} \alpha(a) P(y \mid a)$. We assume that the signal distribution satisfies the following identification assumption:[4]

Assumption 1. *(Identification): $\nexists (\alpha', \alpha'') \in \Delta\mathbb{A} \times \Delta\mathbb{A}$ such that $\alpha' \neq \alpha''$ and*

$$P(\cdot \mid \alpha') = P(\cdot \mid \alpha'').$$

We say monitoring has **full support** if $P(\cdot|a)$ has full support for each $a \in \mathbb{A}$, i.e., if $0 < P(y \mid a) < 1$ for all $a \in \mathbb{A}$ and $y \in \mathcal{Y}$. In the following development, we do not assume full support monitoring assumption unless stated otherwise. At the beginning of the game, nature draws a type for player 1, from a finite set $\Theta := \{\theta_0, \theta_1, ..., \theta_K\}$, according to a distribution $\mu_0 \in \Delta\Theta$,[5] with $\mu_0(\theta_i)$ denoting the prior probability that player 1's type is θ_i. Type θ_0 is a normal type who maximizes his payoffs, which we describe shortly. Each type θ_i for $i \geq 1$ is a commitment type who is committed to playing a fixed (possibly mixed) action $\alpha_i \in \Delta\mathbb{A}$ whenever the game has not ended yet.

Histories and strategies: If an irreversible decision has not been taken by the beginning of period $t \in \mathbb{Z}_+$, a public history contains all the signals observed by the players in periods $\tau < t$. We write $h^0 = \emptyset$ for the initial public history and let $\mathcal{H}^t = \mathcal{Y}^{t-1}$ denote the set of public histories for the beginning of period t, with a typical element identified by h^t. Let $\mathbb{H} := \cup_t \mathcal{H}^t$ be the set of all public histories.

If an irreversible decision has not been made by the beginning of period $t \in \mathbb{Z}_+$, a private history for player 1 also contains all the actions taken in periods $\tau < t$ as well as his type $\theta \in \Theta$. We let $\mathcal{H}_1^t = \mathcal{Y}^{t-1} \times \mathbb{A}^{t-1} \times \Theta$ denote the set of all private histories for player 1 at the beginning of period t, with a typical element identified by h_1^t. Let $\mathbb{H}_1 := \cup_t \mathcal{H}_1^t$ be the set of all private histories for player 1.

An outcome ω of the game consists of the type θ of player 1, $\theta(\omega)$, the stopping time $\mathbb{T}(\omega) \in \mathbb{Z}_+ \cup \{\infty\}$ at which player 2 takes the irreversible action, the realized sequence $(a_\tau(\omega))_{\tau < \mathbb{T}(\omega)}$ of (private) actions taken by player 1 and all signals $(y_\tau(\omega))_{\tau < \mathbb{T}(\omega)}$ observed before the stopping time $\mathbb{T}(\omega)$. We write Ω for the set of all outcomes.

[3] We assume that at the beginning of every period, first player 2 moves, and then player 1 moves. However, our results do not rely on this specification, and essentially all our main results go through even if players moved simultaneously.

[4] This assumption is equivalent to the following assumption: Let M be the matrix where player 1's actions are the rows, and the signals are the columns, and each entry for $(a, y) \in \mathbb{A} \times \mathcal{Y}$ is $P(y \mid a)$. Then, the rows of M are linearly independent. A similar assumption is made in Cripps et al. (2004).

[5] For any set X, ΔX denotes the set of probability distributions on X.

A strategy for player 1 is $\sigma_1 : \mathbb{H}_1 \to \Delta\mathbb{A}$, and a strategy for player 2 is $\sigma_2 : \mathbb{H} \to \Delta\{S, W\}$.[6]

Payoffs: Player 1 receives a flow payoff of $u_1(a)$ in any period where player 2 plays W and player 1 plays a. Player 1 receives a (lump-sum) continuation payoff of V_1 if the irreversible action is taken. Player 2 receives a flow payoff $u_2(a)$ in any period where player 2 plays W and player 1 plays a.[7] Player 2's continuation payoff after she takes the irreversible action depends on the type of her opponent, and $V_2(\theta)$ denotes this payoff. Players have a common discount factor, $\delta \in (0, 1)$.

Every strategy profile $\sigma = (\sigma_1, \sigma_2)$ together with μ_0 induces a probability distribution $\mathbb{P}_{(\mu_0,\sigma)}$ over outcomes. In order to save on notation, we drop the term μ_0 and simply write $\mathbb{P}_{(\sigma)}$ for the probability distribution over outcomes, and we write $\mathbb{E}_{(\sigma)}[\cdot]$ for the induced expectation operator. We write $\mathbb{P}^{\theta}_{(\sigma)}$ for the probability distribution over outcomes when player 1 is type $\theta \in \Theta$, and write $\mathbb{E}^{\theta}_{(\sigma)}[\cdot]$ for the induced expectation operator. We observe that $\mathbb{P}_{(\sigma)}[\cdot] = \sum_{\theta\in\Theta} \mu_0(\theta)\mathbb{P}^{\theta}_{(\sigma)}[\cdot]$ and $\mathbb{E}_{(\sigma)}[\cdot] = \sum_{\theta\in\Theta} \mu_0(\theta)\mathbb{E}^{\theta}_{(\sigma)}[\cdot]$.

Player 1's payoff from the strategy profile σ is:

$$U_1(\sigma) := \mathbb{E}^{\theta_0}_{(\sigma)}\left((1-\delta)\sum_{t=0}^{\mathbb{T}-1}\delta^t u_1(a_t) + \delta^{\mathbb{T}}V_1\right).$$

Player 2's payoff from the strategy profile σ is:

$$U_2(\sigma) := \mathbb{E}_{(\sigma)}\left((1-\delta)\sum_{t=0}^{\mathbb{T}-1}\delta^t u_2(a_t) + \delta^{\mathbb{T}}V_2(\theta)\right).$$

For any $\alpha \in \Delta\mathbb{A}$ and $i = 1, 2$, we use $u_i(\alpha)$ to denote $\sum_{a\in\mathbb{A}} \alpha(a)u_i(a)$ when it does not cause confusion.

Given a strategy profile, we let $\{\mu(h^t)\}_{h^t\in\mathbb{H}}$ denote player's 2 posterior beliefs, where $\mu(h^t) \in \Delta\Theta$ is the belief assigned by player 2 to player 1's type at the history h^t. For any history $h^t \in \mathcal{H}^t$, we let $U_2(h^t;\sigma)$ be the continuation payoff of player 2 at history h^t. For any private history of player 1, h_1^t, let h^t be the associated public history. Notice that the continuation payoff of player 1 at history h_1^t, which we call $U_1(h_1^t;\sigma)$, depends only on the public history and hence $U_1(h_1^t;\sigma) =: U_1(h^t;\sigma)$. Given a strategy profile $\sigma = (\sigma_1, \sigma_2)$, for

[6]The strategy of player 1 conditional on a commitment type, θ_i, is the constant strategy α_i.

[7]This representation of payoff function does not preclude the dependence of payoffs on the observed signals. In particular, if the ex-post payoff of player 2 depended on both player 1's action and the signal at the end of the period, say $v_2(a_t, y)$, then the ex-ante flow payoff function of player 2 would be $u_2(a_t) := \sum_{y\in\mathcal{Y}} v_2(a_t, y)P(y|a_t)$.

every on path public history h^t, and type $\theta \in \Theta$, we define $\sigma_1(h^t;\theta) \in \Delta\mathbb{A}$ as

$$\sigma_1\left(h^t;\theta\right)(a) := \mathbb{P}^{\theta}_{(\sigma)}\left(a_t = a | h^t\right) \text{ for every } a \in \mathbb{A}.$$

The period expected mixed action, $\sigma_1(h^t;\theta)$, corresponds to player 2's expectation about player 1's strategy at public history h^t conditional on type θ.

Reduced outcome: A reduced outcome, $\mathfrak{o} = \{(x(\theta), y(\theta))\}_{\theta\in\Theta}$, is a vector of tuples where $x(\theta) \in [0,1]$, and $y(\theta) \in \Delta(\Delta\mathbb{A})$. The set of all reduced outcomes is $\mathfrak{O}$. Each strategy profile σ induces a reduced outcome $\{(x(\theta), y(\theta))\}_{\theta\in\Theta}$, where for each $\theta \in \Theta$,

$$x(\theta) := 1 - \mathbb{E}^{\theta}_{(\sigma)}\left(\delta^{\mathbb{T}}\right),$$

measures the expected duration of the game against type θ, and for each $\alpha \in \Delta\mathbb{A}$, and $\theta \in \Theta$ with $x(\theta) > 0$,

$$y(\theta)(\alpha) := \frac{\mathbb{E}^{\theta}_{(\sigma)}\left((1-\delta)\sum_{t=0}^{\mathbb{T}-1}\delta^t \mathbb{I}_{\{\sigma_1(h^t;\theta)=\alpha\}}\right)}{x(\theta)},$$

measures the expected discounted number of periods in which player 1's strategy is α conditional on his type being θ, and conditional on histories in which the game is not stopped yet.[8] When $x(\theta) = 0$, then $y(\theta)$ is any arbitrary probability distribution on $\Delta\mathbb{A}$. When $x(\theta) > 0$,

$$\int_{\Delta\mathbb{A}} y(\theta)(\alpha)d[\alpha] = 1,$$

and $y(\theta)(\alpha) \in [0,1]$ for each $\alpha \in \Delta\mathbb{A}$.[9] Therefore, $y(\theta)$ is a probability distribution on $\Delta\mathbb{A}$. Because a commitment type, θ_i, always plays the mixed action α_i, when $x(\theta_i) > 0$, $y(\theta_i)(\alpha_i) = 1$. A direct calculation shows that players' payoffs from a strategy profile σ that induces a reduced outcome $\{(x(\theta), y(\theta))\}_{\theta\in\Theta}$ are given by:

$$U_1(\sigma) = x(\theta_0)\int_{\Delta\mathbb{A}} u_1(\alpha)y(\theta_0)(\alpha)d[\alpha] + (1 - x(\theta_0))V_1, \tag{1}$$

$$U_2(\sigma) = \sum_{i\in\{0,\ldots,K\}} \mu_0(\theta_i)\left(x(\theta_i)\int_{\Delta\mathbb{A}} u_2(\alpha)y(\theta_i)(\alpha)d[\alpha] + (1 - x(\theta_i))V_2(\theta_i)\right). \tag{2}$$

[8] $\mathbb{I}$ is the indicator function.

[9] The function $y(\theta)(\cdot)$ is measurable since its support has countably many elements.

In this paper, we study Nash equilibria of the stochastic game described above. If a reduced outcome is induced by a Nash equilibrium, we call it a reduced equilibrium outcome.

3. Assumptions on Payoffs and Types

We first assume that player 1's myopic best response to W is unique. This assumption is generically satisfied, and is not needed for our results, but makes the exposition of our analysis slightly simpler.

Assumption 2. $\arg max_{a\in\mathbb{A}} u_1(a)$ *has a single element,* a_0.

Our second assumption is the most substantive assumption on the payoff functions. It requires that action S either leads player 1 to his highest payoff, and player 2 to her lowest payoff, or it leads player 1 to his lowest payoff, and player 2 to her highest payoff. Essentially, players disagree on the desirability of the stopping action.

Assumption 3. *(Conflict of desires in stopping decisions) The payoff functions satisfy one of the following two conditions:*

i) (Screening Game) $V_1 > max_{a\in\mathbb{A}} u_1(a)$ *and* $V_2(\theta_0) < min_{a\in\mathbb{A}} u_2(a)$.

ii) (Contracting Game) $V_1 < min_{a\in\mathbb{A}} u_1(a)$ *and* $V_2(\theta_0) > max_{a\in\mathbb{A}} u_2(a)$.

We say that the game is a screening game if the normal type of player 1 prefers player 2 to play action S to any other action profile, and if for player 2 playing S against the normal type is not individually rational. We say that the game is a contracting game if player 1 gets his lowest feasible payoff when player 2 plays S, and player 2 gets her highest feasible payoff against the normal type of player 1 by playing S. If the payoff functions satisfy Assumption 3, and if player 1 is known to be the normal type, the dynamic game has a unique Nash equilibrium outcome. If the game is a screening game, then player 2 never plays S and player 1 always plays a_0. If the game is a contracting game, then player 2 plays S at the beginning of the game.

The next assumption requires the existence of some commitment types that player 1 may find worthwhile to mimic. Player 2's best response against such commitment types coincides with the preferred action of player 1. This assumption ensures that the normal type of player 1 has incentives to establish a reputation by imitating some commitment types.

Assumption 4. *(Existence of commitment types worthwhile to mimic)*

i) If the game is a screening game, then there is a type $s \in \{1, 2, ..., K\}$ *such that* $V_2(\theta_i) > u_2(\alpha_i)$ *for all* $i \in \{1, 2, .., s\}$ *and* $V_2(\theta_i) < u_2(\alpha_i)$ *for all* $i > s$.

ii) If the game is a contracting game, then there is a type $s \in \{1, 2, ..., K\}$ such that $V_2(\theta_i) < u_2(\alpha_i)$ for all $i \in \{1, 2, .., s\}$ and $V_2(\theta_i) > u_2(\alpha_i)$ for all $i > s$.

In the subsequent analysis, "s" is the cutoff index such that all commitment types with an index below s are worthwhile for the normal type to imitate, and all those with an index above s are not. The next assumption requires that the normal type of the agent's myopic best response is different from any commitment action. This assumption ensures that player 1 has a trade-off between revealing information about his type and imitating a commitment type to build a reputation.

Assumption 5. *(Mimicking a commitment type is costly)*
$\alpha_i \neq a_0$ *for every* $i \in \{1, 2, ..., K\}$.[10]

Finally, we impose another assumption that aids the exposition but is not necessary for our findings.

Assumption 6. *$\alpha_i \neq \alpha_j$ for any $i, j \in \{1, ..., K\}$ with $i \neq j$, and $u_1(\alpha_i) > u_1(\alpha_{i+1})$ for every $i \in \{1, 2, ..., s-1\}$.*

This assumption says that commitment types choose distinct strategies, and the flow payoff from imitating a commitment type with a lower index is higher than the flow payoff from imitating a commitment type with a higher index. It implies that $u_1(\alpha_i) \neq u_1(\alpha_j)$ *for every* $(i, j) \in \{1, 2, ..., s\}^2$ with $i \neq j$.

4. Illustrative Examples

We present two numerical examples and an alternative interpretation of our model as an experimentation game. We also preview some of our results in the context of the numerical examples.

Example 1. Screening game: worker promotion

Consider the following example of a dynamic screening game. If player 2 chooses S at some period, then the game ends. In any period where player 2 has never played S, player 1 takes an action $a \in \{a_0, a_1\}$. Player 1's actions are possibly observed with some noise, and the monitoring structure satisfies Assumption 1. Player 1 is either a commitment type, θ_1, who plays a_1 in every period when the game has not ended yet, or a normal type, θ_0, who can choose either action. The prior probability that player 1 is a commitment type is $\mu \in (0, 1)$. The normal type of player 1 receives a flow payoff of 1 if

[10] We use the convention that a_0 denotes both a pure action, and the mixed action that puts probability 1 on the pure action a_0.

he chooses action a_0 ($u_1(a_0) = 1$), and a payoff of 0 if he chooses a_1 ($u_1(a_1) = 0$). Player 2's flow payoff in a period when the game has not ended yet depends only on the action of player 1, and $u_2(a_0) = 0$, and $u_2(a_1) = 1$. If player 2 chooses action S in period t, then the normal type of player 1 receives a lump-sum payoff of $V_1 = 2$ (i.e., player 1 receives a payoff that is equivalent to a payoff stream of 2 at every period $t' \geq t$). Player 2's lump-sum payoff when the game ends depends on the type of her opponent. Specifically, $V_2(\theta_0) = -1$ and $V_2(\theta_1) = 2$. In this game, player 1 prefers the game to end, and player 2 prefers to end the game against the commitment type, but not against the normal type. This example captures a situation where a principal (player 2) would like to promote a high-skilled worker (commitment type), but not a low-skilled one (normal type), who can exert extra effort.

Suppose that player 2 expects the normal type of player 1 to pool with the commitment type, i.e., play a_1 at every period. Then, player 2's best response is to play S if $\mu > \mu^* := \frac{2}{3}$, and to play W if $\mu < \mu^*$. In Theorems 1 and 2, we characterize the players' equilibrium payoffs, and the evolution of the posterior belief process when the players are patient (as $\delta \to 1$). If $\mu > \mu^*$, then player 2 plays S almost immediately against both types, and screening is ineffective. Hence, player 1's payoff converges to V_1, which is equal to 2, and player 2's payoff converges to $\mu V_2(\theta_1) + (1-\mu)V_2(\theta_0) = 2\mu - (1-\mu)$. If $\mu < \mu^*$, then player 1's equilibrium behavior is akin to an initial randomization between mimicking the commitment type or revealing his type. Player 2's posterior belief about player 1's type almost immediately either jumps to μ^* if player 1 mimics the commitment type, or falls to 0 if player 1 reveals his type. Conditional on the former event, the normal type of player 1 continues to pool with the commitment type, and player 2 plays S with some delay. Conditional on the latter event, player 2 never plays S. Moreover, we show in Theorem 4 that S is played only when the posterior belief is in a small neighborhood of μ^*, and is eventually played against the commitment type. Hence, Bayes' rule implies that, she eventually plays S against the normal type with probability $\frac{\mu}{1-\mu}\frac{1-\mu^*}{\mu^*}$.

Example 2. Contracting game: task rescheduling, worker demotion

Consider the following example of a dynamic contracting game. In this game, the action sets of the players, the possible types of player 1, and the flow payoff functions when W is played are the same as in Example 1. However, if player 2 chooses action S, then players' lump-sum payoffs are $V_1 = -1, V_2(\theta_0) = 2$, and $V_2(\theta_1) = -1$. In this game, player 1 prefers the game not to end, and player 2 prefers to not end the game against the commitment type, but prefers to end the game against the normal type. This example captures a situation where a principal (player 2) decides whether to continue hiring an agent (player 1), or to demote/reschedule him. A high ability worker (commitment type) always chooses high effort, while a low ability worker (normal type) can choose either effort.

Suppose that player 2 expects the normal type of player 1 to pool with the commitment type, i.e., play a_1 at every period. Then, player 2's best response is to play W if $\mu > \mu^* := \frac{1}{3}$, and to play S if $\mu < \mu^*$. In Theorems 1 and 3, we characterize the players' equilibrium payoffs, and the evolution of the posterior belief process when the players are patient (as $\delta \to 1$). If $\mu < \mu^*$, then player 2 plays S almost immediately against both types. Hence, player 1's payoff converges to V_1, which is equal to -1, and player 2's payoff converges to $\mu V_2(\theta_1) + (1-\mu)V_2(\theta_0) = -\mu + 2(1-\mu)$. If $\mu > \mu^*$, then player 2 plays W and player 1 plays a_1 for a sufficiently long time, which results in the payoff $u_1(a_1)$ for player 1, and $u_2(a_1)$ for player 2. Moreover, the posterior belief about player 1's type lies close to the set of beliefs $\{\eta \in \Delta\Theta : \eta(\theta_1) \geq 1/3\}$. Hence, equilibrium behavior is akin to the behavior that would result if player 1 committed to pooling with the commitment type, i.e., screening is ineffective.

However, we also show in Theorem 5 that if player 1's actions are observed with noise that has full support, and if $\mu > \mu^*$, then player 1's type is partially revealed. Despite the equilibrium behavior being akin to pooling behavior, pooling at every period is not an equilibrium. Suppose for simplicity that there is a Markov equilibrium where the state variable is the public belief about player 1's type. If player 1 is expected to play a_1 with probability one, then the posterior belief of player 2 about player 1's type is unchanged after the signal realization, and the continuation payoff of player 1 is independent of the signal realization. However, then player 1 would profitably deviate to playing a_0. We show that this argument applies more generally (in all Nash equilibria), and player 2 learns some information about player 1 until the action S is finally taken. This argument, reminiscent of the findings of Cripps, Mailath and Samuelson (2004), shows that imperfect monitoring leads to information revelation. We show that player 2 eventually plays S against the normal type, for any discount factor.

In Theorem 5, we further characterize the long-run behavior when the players are patient (as $\delta \to 1$). In the long run, player 2's posterior belief conditional on her opponent being the normal type eventually falls until reaching an arbitrarily small neighborhood of μ^*, and player 2 plays action S. Conditional on the commitment type, the posterior belief either converges to 1, and player 2 never takes action S, or the posterior belief converges to μ^*, and player 2 plays S. The key finding is that when the posterior belief falls, it moves very slowly around μ^*, and action S is taken only when the posterior belief is close to μ^*. This allows us to calculate the long-run probability of action S against each type using Bayes' rule. Eventually, S is played with probability 1 against the normal type, and with probability $\frac{1-\mu}{\mu}\frac{\mu^*}{1-\mu^*}$ against the commitment type. The total probability that Player 2 eventually plays S is $\frac{1-\mu}{1-\mu^*}$. A key contribution of the paper, which distinguishes it from the insights of Cripps et al. (2004), is that when agents are patient, the posterior belief process has two absorbing states, μ^* and 1, and action S is taken with probability 1 at belief μ^*.

Example 3. An alternative interpretation of the model

Our model can be thought of as a two-armed bandit problem where the outcome distribution of one of the arms can be manipulated by a strategic agent.

More concretely, suppose there are two possible states of the world, $\omega \in \{\alpha, \beta\}$, and two arms, A and B. Each arm produces a random payoff to the principal. The realized payoff is 1 with probabilities $p_A(\omega)$ and $p_B(\omega)$, and the realized payoff is 0 with the remaining probabilities. The principal does not know the state. At the initial period, the principal chooses arm A or arm B. If she chooses arm A, then she observes the realized payoff, and the next period she faces the same situation. If she chooses arm B, then she is stuck with arm B forever, i.e., choosing arm B is an irreversible action. More generally, in any period, if she has never chosen arm B, then she makes a decision between arm A or the irreversible arm B.

There is an agent who has a preference over the principal's arm choices. The agent's flow payoff from the principal choosing arm A is u_A, and his flow payoff from the principal choosing arm B is u_B. The agent observes the state.

Suppose $u_A > u_B$, and the principal prefers arm A in state α and arm B in state β, i.e, $p_A(\alpha) > p_B(\alpha)$ and $p_B(\beta) > p_A(\beta)$. If state is β, each period the agent can pay a cost $c < u_A - u_B$ that changes the probability of a high payoff from arm A to $p_A(\alpha)$ in that period. In state α, the agent cannot manipulate the probability distribution of outcomes. If, in addition, $p_B(\beta) > p_A(\alpha)$, then this experimentation game satisfies all of our assumptions, and is an instance of a contracting game.

Suppose now that $u_B > u_A$, and the principal prefers arm A in state α and arm B in state β, i.e, $p_A(\alpha) > p_B(\alpha)$ and $p_B(\beta) > p_A(\beta)$. If the state is α, each period the agent can pay a cost $c < u_B - u_A$ that changes the probability of a high payoff from arm A to $p_A(\beta)$ in that period. In state β, the agent cannot manipulate the probability distribution of outcomes. If, in addition, $p_A(\beta) > p_B(\alpha)$, then this experimentation game satisfies all of our assumptions, and is an instance of a screening game.

5. Auxiliary Two-stage Game

In this section we introduce an auxiliary two-stage game that is simple to analyze. In this game, player 1 has a single opportunity to signal his type by choosing a mixed action to commit to. Player 2 observes player 1's commitment action, updates her beliefs, and chooses her action. As we argue below, this game generically has a unique Nash equilibrium outcome. Our main result shows that Nash equilibrium payoff vectors of the dynamic game converge to the unique Nash equilibrium payoff vector of the two-stage game when players are patient.

Let G be the dynamic stopping game described in the model that satisfies Assumptions 1-6. We define the corresponding auxiliary two-stage game, $\tilde{G}$ as follows. First, nature picks player 1's type $\theta \in \Theta$ according to the probability distribution $\mu_0 \in \Delta\Theta$, as in G. In the first stage of the game, player 1 chooses a mixed action $\alpha \in \Delta\mathbb{A}$. Player 1 is either a commitment type $\theta_i \in \{\theta_1, ..., \theta_K\}$ who chooses the mixed action α_i with probability one, or a normal type, θ_0, who chooses his action strategically. A strategy for the normal type of player 1 is $\tilde{\sigma}_1 \in \Delta(\Delta\mathbb{A})$. In the second stage, player 2 chooses an action $a_2 \in \{S, W\}$, after observing player 1's mixed action. A strategy for player 2 specifies the probability of playing W expressed as $\tilde{\sigma}_2 : \Delta\mathbb{A} \to [0,1]$. After player 2's action choice, the game ends, and payoffs are realized. Payoff functions of the players in this auxiliary game are closely linked to those in the dynamic game. If player 2 chooses S in the second stage, then player 1's payoff is V_1, and player 2's ex-post payoff is $V_2(\theta)$. If instead, player 2 chooses action W when player 1 has chosen action $\alpha \in \Delta\mathbb{A}$, then players get payoffs $u_1(\alpha)$ and $u_2(\alpha)$, respectively. When it does not cause a confusion, if G is a screening (contracting) game, we also refer to the corresponding auxiliary two-stage game as a screening (contracting) game. We study Nash equilibria of this two-stage game.

Every strategy profile together with the prior μ_0 induces a probability distribution, $\tilde{\mathbb{P}}_{(\sigma,\mu_0)}$ over the action profiles $(\alpha_1, a_2) \in \Delta\mathbb{A} \times \{W, S\}$ and player 1's types, which we call an outcome distribution. Each outcome distribution induces a reduced outcome, $\mathfrak{o} = \{x(\theta), y(\theta)\}_{\theta\in\Theta} \in \mathfrak{O} := \{[0,1], \Delta(\Delta\mathbb{A})\}_{\theta\in\Theta}$, where for each $\theta \in \Theta$,

$$x(\theta) := \tilde{\mathbb{P}}_{(\sigma,\mu_0)}\left(a_2 = W|\theta\right),$$

and when $x(\theta) > 0$, for every $B \in \mathbb{B}\left(\Delta\mathbb{A}\right)$,we have[11]

$$y(\theta)\left(B\right) := \tilde{\mathbb{P}}_{(\sigma,\mu_0)}\left(B|a_2 = W, \theta\right).$$

[11]We write $\mathbb{B}\left(\mathbb{X}\right)$ for the Borel sigma-field for the topological space $\mathbb{X}$.

The term $x(\theta)$ is the probability that player 2 plays W against type θ, and $y(\theta)$ is the probability distribution over player 1's actions conditional on player 2 playing W and player 1's type being θ. Similar to the description in the dynamic model, when $x(\theta) = 0$, $y(\theta)$ is any arbitrary probability distribution on $\Delta\mathbb{A}$. For a commitment type, θ_i, if $x(\theta_i) > 0$, then $y(\theta_i)(\alpha_i) = 1$. Finally, players' payoffs from a strategy profile σ that induces a reduced outcome $\{(x(\theta), y(\theta))\}_{\theta\in\Theta}$ are given by equations (1) and (2). We say two reduced outcomes $\{x_1(\theta), y_1(\theta)\}_{\theta\in\Theta}$, and $\{x_2(\theta), y_2(\theta)\}_{\theta\in\Theta}$ are equivalent if for every $\theta \in \Theta$, either $x_1(\theta) = x_2(\theta)$, and $y_1(\theta) = y_2(\theta)$, or $x_1(\theta) = x_2(\theta) = 0$.

This auxiliary game is simple to analyze since it is a two-stage game with observable mixed actions. Hence, all complicated inference problems pervasive in dynamic games with imperfect monitoring are absent in this game. We will show that generically there exists a unique reduced equilibrium outcome.

For $i \in \{1, 2, ..., s\}$, let ${\mu_{\mathbf{i}}}^b := \{\mu \in \Delta\Theta : \mu(\theta_i) = \mu_i^* \text{ and } \mu(\theta_0) = 1 - \mu_i^*\}$ be the belief that makes player 2 indifferent between playing S and W when she believes that her opponent's type is θ_i with probability μ_i^*, and with the remaining probability her opponent is a normal type who picked action α_i. More precisely,

$$\mu_i^* V_2(\theta_i) + (1 - \mu_i^*) V_2(\theta_0) = u_2(\alpha_i).$$

The existence and uniqueness of a threshold $\mu_i^* \in (0, 1)$ follows from Assumptions 3 and 4. Suppose that in the second stage, the support of the posterior belief of player 2 is $\{\theta_0, \theta_i\}$ for some $i \in \{1, ..., s\}$. Then, player 2's optimal action is to play the desirable action for the normal type (S in a screening game, W in a contracting game) whenever her posterior belief attaches a probability more than μ_i^* on type θ_i, and her optimal action is to play the undesirable action for the normal type (W in a screening game, S in a contracting game) when this probability is strictly smaller than μ_i^*. Finally, the structure of the equilibrium will depend on which one of the following two inequalities on the prior belief holds:

$$\sum_{i\in\{1,2,...,s\}} \mu_0(\theta_i) \frac{1 - \mu_i^*}{\mu_i^*} > \mu_0(\theta_0), \tag{3}$$

$$\sum_{i\in\{1,2,...,s\}} \mu_0(\theta_i) \frac{1 - \mu_i^*}{\mu_i^*} < \mu_0(\theta_0). \tag{4}$$

If inequality (3) holds, then for any fixed strategy of player 1, there exists a commitment type θ_i, with $i \leq s$, such that player 2's posterior belief after observing action α_i attaches a probability strictly more than μ_i^* to type θ_i. If inequality (4) holds, then for any fixed strategy of player 1 whose support is contained in the set $\{\alpha_1, ..., \alpha_s\}$, there exists a commitment type θ_i, with $i \leq s$, such that player 2's posterior belief after observing action α_i attaches a probability strictly less than μ_i^* to type θ_i.

5.1. Screening Game

Lemma 1. *Suppose G is a screening game, and let $\tilde{G}$ be the corresponding auxiliary two-stage game.*

1. *If inequality (3) holds, then $\tilde{G}$ has a unique reduced equilibrium outcome. In this outcome, player 2 plays S against types $\theta \in \{\theta_0, \theta_1, ..., \theta_s\}$ and W against types $\theta \in \{\theta_{s+1}, ..., \theta_K\}$.*

2. *If inequality (4) holds, then $\tilde{G}$ has a unique equilibrium outcome distribution, which is generated by the following strategy profile $(\tilde{\sigma}_1, \tilde{\sigma}_2)$:*

   2.1. For $i \in \{1, .., s\}$, $\tilde{\sigma}_1(\alpha_i) = \frac{\mu_0(\theta_i)}{\mu_0(\theta_0)} \frac{1-\mu_i^*}{\mu_i^*}$.

   2.2. $\tilde{\sigma}_1(a_0) = 1 - \sum_{i \in \{1,...,s\}} \frac{\mu_0(\theta_i)}{\mu_0(\theta_0)} \frac{1-\mu_i^*}{\mu_i^*}$.

   2.3. For $i \in \{1, .., s\}$, $\tilde{\sigma}_2(\alpha_i) = \frac{V_1 - u_1(a_0)}{V_1 - u_1(\alpha_i)}$.

   2.4. $\tilde{\sigma}_2(\alpha) = 1$, for every $\alpha \notin \{\alpha_1, ..., \alpha_s\}$.

Lemma 1 shows that generically, two-stage screening games have a unique reduced equilibrium outcome. Below, we provide sketch of the proof, but we skip a formal proof since it is straightforward.

If inequality (3) holds, then for any fixed strategy of player 1 there exists some type $i \in \{1, ..., s\}$, such that observing α_i, player 2's posterior belief that player 1 is type θ_i exceeds μ_i^*, hence plays S with probability 1. This observation is the key for the argument of the first claim.

If inequality (4) holds, the normal type does not obtain a payoff of V_1 in equilibrium. Suppose by way of contradiction that this is not true. Then the support of player 1's equilibrium strategy is a subset of $\{\alpha_1, ..., \alpha_s\}$. However, inequality (4) implies that for any fixed strategy of player 1 with its support contained in $\{\alpha_1, ..., \alpha_s\}$, there exists some type $i \in \{1, ..., s\}$ such that α_i is in the support of player 1's equilibrium strategy and, after observing α_i, player 2's posterior belief that player 1 is type θ_i is strictly below μ_i^*. Therefore, player 2 plays W after observing α_i, and player 1's equilibrium payoff is less than V_1. A similar argument also establishes that the support of player 1's equilibrium strategy must contain a_0, hence player 1's equilibrium payoff is equal to $u_1(a_0)$. This

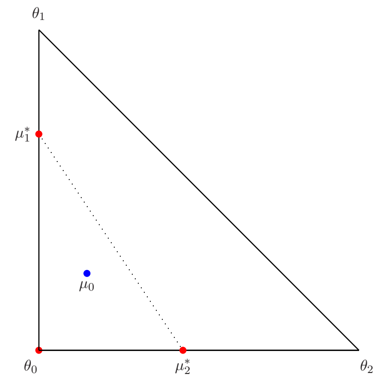

Figure 1: This figure illustrates the equilibrium posterior beliefs of the principal on the belief simplex. Each corner of the triangle represents the degenerate distribution that puts probability one on the type that is labeled at that corner. This is an example of a screening game where $\Theta = \{\theta_0, \theta_1, \theta_2\}$, and $s = 2$. If inequality (4) holds, as illustrated in the figure, then player 2's equilibrium posterior belief is one of the three beliefs: $(1, 0, 0)$, $(1 - \mu_1^*, \mu_1^*, 0)$, or $(1 - \mu_2^*, 0, \mu_2^*)$, illustrated as the red dots. The Martingale property of the beliefs implies a unique distribution over this set of posterior beliefs, and a unique strategy for the normal type of player 1 that induces this distribution.

implies that player 2's posterior belief after observing α_i for any $i \in \{1, ..., s\}$ should attach a probability to type θ_i that does not exceed μ_i^*, and should not be strictly below μ_i^*. In other words, in every equilibrium, player 2's posterior belief after observing α_i should be μ_i^b (see Figure 1). Because player 1's equilibrium payoff is $u_1(a_0)$, and because he is choosing each α_i with a strictly positive probability, player 2 plays W after observing α_i with probability $\tilde{\sigma}_2(\alpha_i) \in (0, 1)$ that satisfies the following indifference condition for player 1:

$$u_1(a_0) = \tilde{\sigma}_2(\alpha_i) u_1(\alpha_i) + (1 - \tilde{\sigma}_2(\alpha_i)). \tag{5}$$

5.2. Contracting Game For contracting games, we make the following genericity assumption about the prior belief μ_0. For every $l \in \{1, ..., s\}$, we have

$$\sum_{0 < i \leq l} \mu_0(\theta_i) \frac{1 - \mu_i^*}{\mu_i^*} \neq \mu_0(\theta_0). \tag{6}$$

Recall Assumption 6 that, $u_1(\alpha_i) > u_1(\alpha_j)$ for all $i < j$, $i, j \in \{1, ..., s\}$.

Lemma 2. *Suppose G is a contracting game. Let $\tilde{G}$ be the corresponding auxiliary two-stage game, and assume μ_0 satisfies inequality (6).*

1. *If inequality (4) holds, then in the unique reduced equilibrium outcome of $\tilde{G}$, player 2 plays S against every type.*

2. If inequality (3) holds, then let j be the smallest element of $\{1, 2, ..., s\}$ such that $\mu_0(\theta_0) < \sum_{0<i\leq j} \mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*}$. $\tilde{G}$ has a unique equilibrium outcome distribution, which is generated by the following strategy profile $(\tilde{\sigma}_1, \tilde{\sigma}_2)$:

2.1. If $j = 1$: $\tilde{\sigma}_1(\alpha_1) = 1$, $\tilde{\sigma}_2(\alpha_i) = 1$ for $i \in \{1, ..., s\}$, $\tilde{\sigma}_2(\alpha) = 0$ for $\alpha \notin \{\alpha_1, ..., \alpha_s\}$.

2.2. If $j \in \{2, 3, ..., s\}$:

2.2.a $\tilde{\sigma}_1(\alpha_i) = \frac{\mu_0(\theta_i)}{\mu_0(\theta_0)}\frac{1-\mu_i^*}{\mu_i^*}$ for $i \in \{1, ..., j-1\}$, $\tilde{\sigma}_1(\alpha_j) = 1-\sum_{i\in\{1,...,j-1\}} \frac{\mu_0(\theta_i)}{\mu_0(\theta_0)}\frac{1-\mu_i^*}{\mu_i^*}$.

2.2.b $\tilde{\sigma}_2(\alpha_i) = 1$ for $i \in \{j, ..., s\}$, $\tilde{\sigma}_2(\alpha_i) = \frac{u_1(\alpha_j)-V_1}{u_1(\alpha_i)-V_1}$ for $i \in \{1, ..., j-1\}$, $\tilde{\sigma}_2(\alpha) = 0$ for $\alpha \notin \{\alpha_1, ..., \alpha_s\}$.

Lemma 2 shows that generically, two-stage contracting games have a unique reduced equilibrium outcome. Below, we provide sketch of the proof, but we skip a formal proof since it is straightforward.

If inequality (4) holds, then player 1's equilibrium payoff is V_1. Suppose by way of contradiction that this is not true. Then, the support of player 1's equilibrium strategy is contained in $\{\alpha_1, ..., \alpha_s\}$. However, inequality (4) implies that there exists a type $i \in \{1, .., s\}$ such that α_i is in the support of player 1's equilibrium strategy, and player 2's posterior belief that player 1's type is θ_i after observing α_i is strictly below μ_i^*. Hence, after observing α_i, player 2 plays S. Therefore, player 1's equilibrium payoff is V_1. Player 1's incentive constraints imply that player 2 plays S against every type.

We now analyze the case when inequality (3) holds. Notice that because we reordered the commitment types, if all the types when mimicked would induce player 2 to play W with probability 1, then the normal type would prefer to mimic type θ_1. Suppose that $\mu_0(\theta_1)\frac{1-\mu_1^*}{\mu_1^*} > \mu_0(\theta_0)$. After observing action α_1, player 2's posterior belief that player 1's type is θ_1 exceeds μ_1^*, regardless of the normal type's strategy. Hence, in equilibrium, player 1 guarantees the payoff $u_1(\alpha_1)$ by mimicking θ_1. Because mimicking another commitment type, or playing any other action, gives a strictly lower payoff than $u_1(\alpha_1)$, in the unique equilibrium outcome he mimics θ_1, which leads to a posterior belief that attaches a probability greater than μ_1^* on type θ_1. Hence, if $j = 1$, a unique equilibrium outcome is obtained.

Let us then assume that $j > 1$. What happens as the prior on the normal type increases slightly so that $\mu_0(\theta_1)\frac{1-\mu_1^*}{\mu_1^*} < \mu_0(\theta_0) < \mu_0(\theta_1)\frac{1-\mu_1^*}{\mu_1^*} + \mu_0(\theta_2)\frac{1-\mu_2^*}{\mu_2^*}$? In this case, for any strategy of the normal type of player 1, one of the actions $\alpha_i \in \{\alpha_1, \alpha_2\}$ must lead to a posterior with support contained in $\{\theta_0, \theta_i\}$, and that attaches a probability strictly higher than μ_i^* on type θ_i. Hence, playing α_i must lead to the payoff $u_1(\alpha_i)$. It then follows that player 1's equilibrium strategy may put positive probability only on α_1 and α_2. Since $u_1(\alpha_1) > u_1(\alpha_2)$, α_2 must lead to a posterior belief on type θ_2 greater than μ_2^*, to which player 2 best responds by playing W, while action α_1 leads to the posterior μ_1^* on type θ_1 and player 2 randomizes, putting probability $\tilde{\sigma}_2(\alpha_1)$ on W that keeps the normal type indifferent between the two actions:

$$u_1(\alpha_2) = \tilde{\sigma}_2(\alpha_1)u_1(\alpha_1) + (1 - \tilde{\sigma}_2(\alpha_1))V_1.$$

In general, when inequality (3) holds, let $0 < j \leq s$ be the smallest integer such that

$$\mu_0(\theta_0) < \sum_{0<i\leq j} \mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} \tag{7}$$

holds. A similar argument implies that the normal type will randomize between actions $\alpha_i \in \{\alpha_1, ..., \alpha_j\}$. The action α_j will lead to a posterior belief that puts a probability more than μ_j^* on type j, and will lead player 2 to best respond by playing W. On the other hand, each action $\alpha_i \in \{\alpha_1, ..., \alpha_{j-1}\}$ will lead to the posterior belief μ_i^b, and it will be followed by a randomization of player 2 which puts probability $\tilde{\sigma}_2(\alpha_i)$ on W, keeping the normal type indifferent between all the actions $\alpha_i \in \{\alpha_1, ..., \alpha_j\}$:

$$u_1(\alpha_j) = \tilde{\sigma}_2(\alpha_i)u_1(\alpha_i) + (1 - \tilde{\sigma}_2(\alpha_i))V_1. \tag{8}$$

Figure 2 illustrates the equilibrium posterior beliefs in an example. Although it is possible to characterize the equilibria of the auxiliary two-stage game and the dynamic game for a proper subset of beliefs which do not satisfy inequality (6), this genericity assumption allows us to express and explain our main results in a simpler way.

6. Equilibrium of the Dynamic Model

Recall that $\mathfrak{O}$ is the set of all reduced outcomes. Take an element $\mathfrak{o}^* = \{(x^*(\theta), y^*(\theta))\}_{\theta\in\Theta} \in \mathfrak{O}$ and, for each $n \in \mathbb{N}$, an element $\mathfrak{o}_n = \{(x_n(\theta), y_n(\theta))\}_{\theta\in\Theta} \in \mathfrak{O}$. We say that the sequence $\{\mathfrak{o}_n\}_{n=1}^{\infty}$ converges to $\mathfrak{o}^*$ if for each $\theta \in \Theta$, $x_n(\theta) \to x^*(\theta)$, and $x^*(\theta) > 0$ implies $y_n(\theta)$ converges weakly to $y^*(\theta)$.[12] Our main result is Theorem 1.

Theorem 1. *Take a sequence $\{\delta_n\}$ of discount factors converging to one and let $\{\sigma_n\}$ be a sequence of Nash equilibria of a sequence of games where along the sequence, the payoff functions and the prior belief μ_0 are fixed, μ_0 satisfies (6), and in the n^{th} game the discount*

[12]Let X be a metric space with its Borel sigma-field $\mathbb{B}(X)$. A bounded sequence of probability measures (P_n) on $(X, \mathbb{B}(X))$ converges weakly to the probability measure P if $\lim P_n(C) = P(C)$ for all continuity sets C of measure P.

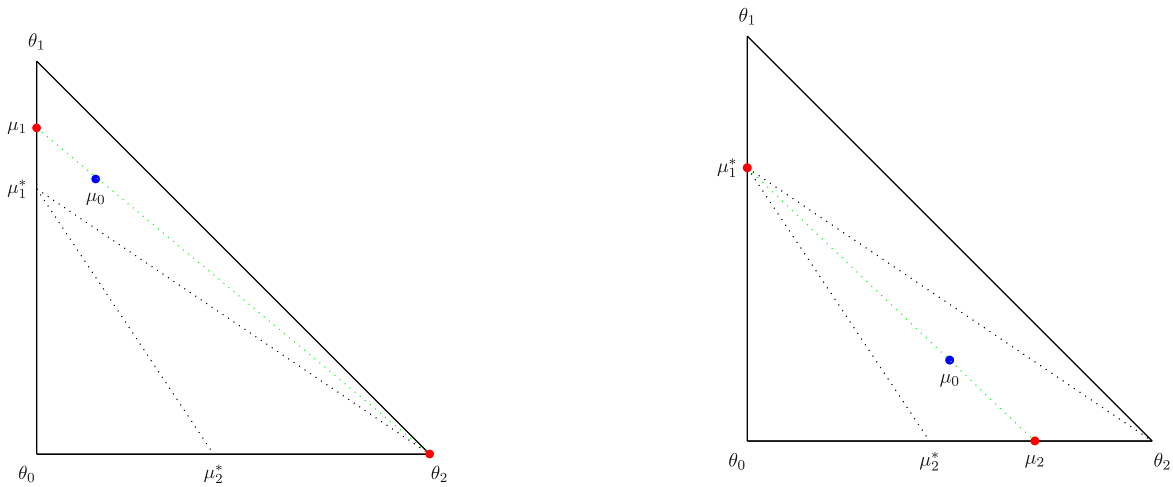

(a) $\mu_0(\theta_1)\frac{1-\mu_1^*}{\mu_1^*} > \mu_0(\theta_0)$

(b) $\mu_0(\theta_1)\frac{1-\mu_1^*}{\mu_1^*} < \mu_0(\theta_0) < \sum_{0<i\leq 2}\mu(\theta_i)\frac{1-\mu_i^*}{\mu_i^*}$

Figure 2: This panel illustrates the equilibrium posterior beliefs of the principal on the belief simplex. This is an example of a contracting game where, $\Theta = \{\theta_0, \theta_1, \theta_2\}$, and $s = 2$ and inequality (3) holds. If $j = 1$, as illustrated in the left figure, then player 2's posterior belief is either $(0, 0, 1)$ or $(1 - \mu_1, \mu_1, 0)$. If $j = 2$, as illustrated in the right figure, then player 2's posterior belief is either $(1 - \mu_1^*, \mu_1^*, 0)$ or $(1 - \mu_2, 0, \mu_2)$. The Martingale property of the beliefs implies a unique distribution over this set of posterior beliefs, and unique strategy for the normal type of player 1 that induces this distribution.

factor is δ_n. Let $\mathfrak{o}_n = \{(x_n(\theta), y_n(\theta))\}_{\theta\in\Theta}$ be a reduced outcome induced by σ_n and let $\mathfrak{o}^* = \{(x^*(\theta), y^*(\theta))\}_{\theta\in\Theta}$ be the unique reduced equilibrium outcome of the corresponding auxiliary two-stage game. The sequence of reduced outcomes $\{\mathfrak{o}_n\}_{n=1}^{\infty}$ converges to $\mathfrak{o}^*$.

Theorem 1 asserts that any sequence of equilibrium-reduced outcomes $\{(x_n(\theta), y_n(\theta))\}_{\theta\in\Theta}$ of dynamic games converges to the unique equilibrium-reduced outcome $\{(x^*(\theta), y^*(\theta))\}_{\theta\in\Theta}$ of the auxiliary two-stage game, when players are patient. This implies that, in the patient limit, the equilibrium payoff vector of the dynamic game is identical to the equilibrium payoff vector of the auxiliary two-stage game. Hence, noise in the monitoring structure does not affect equilibrium payoffs when players are patient.

In light of Lemmas 1 and 2, Theorem 1 implies that dynamic screening is partially effective. First, dynamic screening allows the principal to screen all types that are not worthy to be mimicked by the normal type, i.e., types with index greater than s, and play a best response to them. In screening games, if the prior belief about the commitment types with indices less than s is high relative to the prior belief about the normal type (high reputation case), then player 2 plays S almost immediately. Hence, screening among such types is ineffective. If the prior belief about the commitment types with indices less than s is low relative to the prior belief about the normal type (low reputation case), then player 1 reveals partial information. In contracting games, in the low reputation case, player 2 plays S almost immediately. Hence, screening among such types is ineffective. In the high reputation case, player 1 reveals partial information, and player 2 utilizes this information.

However, the implication of the theorem goes further than the payoff equivalence. The reduced outcome of a strategy profile in the dynamic game captures some information about the behavior of the players. In the unique reduced equilibrium outcome of the auxiliary game, for any type $\theta \in \Theta$, either $x^*(\theta) = 0$, or $y^*(\theta)$ has finite support. The term $x^*(\theta)$ tells us the expected discounted duration of the game against type θ. If $x^*(\theta) = 0$, then player 2 plays S almost immediately against type θ. If $x^*(\theta) > 0$ and $y^*(\theta)(\alpha) = 0$ for some $\alpha \in \Delta\mathbb{A}$, then we do not expect to see type θ playing α for a number of periods that have a non-negligible impact on ex ante payoffs. On the other side, reduced outcomes in the dynamic game do not capture the long run behavior, nor anything about the dynamics of the play. To gain more insight about these aspects of the equilibrium play, we examine the belief and behavior dynamics jointly in sections 6.1.1 and 6.1.2, and the long run behavior in section 7.

6.1. Equilibrium Belief Process and Behavior We now investigate the belief dynamics in conjunction with behavior in equilibrium. We will show that the evolution of player 2's posterior beliefs is closely linked to the belief process in the equilibrium of the auxiliary two-stage game. In particular, the belief process and the associated behavior confirm that player 1's equilibrium behavior is as if the normal type chooses a strategy $\alpha \in \{a_0, \alpha_1, ..., \alpha_s\}$ in the beginning of the game and sticks to it until the end of the game. Hence, the normal type does not build his reputation gradually over time. Rather, it forms it in the beginning through his initial actions and then sticks to it for a long time. In Theorems 2 and 3 below, we formalize and show these findings.

6.1.1. *Behavior in Screening Games* Recall that Lemma 1 and Theorem 1 imply that in screening games, if inequality (3) holds, then player 2 plays S against types $\theta \in \{\theta_0, \theta_1, ..., \theta_s\}$ almost immediately. Therefore, we cannot obtain tight results about the belief dynamics conditional on these types in this case. However, if inequality (4) holds, then player 2 plays W against types $\theta \in \{\theta_0, \theta_1, ..., \theta_s\}$ with positive probability, and in this case, we have a tight result about belief dynamics against such types. For $i \in \{0, 1, ..., K\}$, let $\boldsymbol{\mu_i} \in \Delta\Theta$ be the distribution that puts probability one on type θ_i. Given two elements x, y of a finite dimensional Euclidean space, we define $d(x, y) := \|x - y\|$, where $\|.\|$ stands for the Euclidian norm.

Theorem 2. *Take a sequence $\{\delta_n\}$ of discount factors converging to one and let $\{\sigma_n\}$ be a sequence of Nash equilibria of a sequence of screening games where along the sequence, the payoff functions and the prior belief μ_0 are fixed, and in the n^{th} game the discount factor is δ_n. Let $\tilde{G}$ be the corresponding auxiliary two-stage game.*

1. *For every $i > s$ and for every $\varepsilon > 0$,*
   $\lim \mathbb{E}^{\theta_i}_{(\sigma_n)} \left((1-\delta_n) \sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\{d(\mu_t, \boldsymbol{\mu_i}) < \varepsilon\}} \right) = 1.$

2. *Suppose inequality (4) holds. Let $(\tilde{\sigma}_1, \tilde{\sigma}_2)$ be the strategy profile that leads to the unique equilibrium outcome distribution of $\tilde{G}$. There exists $\varepsilon^* > 0$ such that for every $\varepsilon \in (0, \varepsilon^*)$ the following hold:*

   2.1. $\lim \mathbb{E}^{\theta_i}_{(\sigma_n)} \left((1-\delta_n) \sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\left\{d\left(\mu_t, \mu_i^b\right)\right) < \varepsilon\}} \right) = \tilde{\sigma}_2(\alpha_i)$ *for every $i \in \{1, ..., s\}$.*

   2.2. $\lim \mathbb{E}^{\theta_i}_{(\sigma_n)} \left(\delta_n^{\mathbb{T}} \mathbb{I}_{\left\{d\left(\mu_t, \mu_i^b\right) < \varepsilon\right\}} \right) = 1 - \tilde{\sigma}_2(\alpha_i)$ *for every $i \in \{1, ..., s\}$.*

   2.3. $\lim \mathbb{E}^{\theta_0}_{(\sigma_n)} \left((1-\delta_n) \sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\{d(\mu_t, \boldsymbol{\mu_0}) < \varepsilon, d(\sigma_1(h^t;\theta_0), a_0) < \varepsilon\}} \right) = \tilde{\sigma}_1(a_0).$

   2.4. $\lim \mathbb{E}^{\theta_0}_{(\sigma_n)} \left((1-\delta_n) \sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\left\{d\left(\mu_t, \mu_i^b\right) < \varepsilon, d(\sigma_1(h^t;\theta_0), \alpha_i) < \varepsilon\right\}} \right) = \tilde{\sigma}_1(\alpha_i)\tilde{\sigma}_2(\alpha_i)$ *for every $i \in \{1, ..., s\}$.*

   2.5. $\lim \mathbb{E}^{\theta_0}_{(\sigma_n)} \left(\delta_n^{\mathbb{T}} \mathbb{I}_{\left\{d\left(\mu_t, \mu_i^b\right) < \varepsilon\right\}} \right) = \tilde{\sigma}_1(\alpha_i)\left(1 - \tilde{\sigma}_2(\alpha_i)\right)$ *for every $i \in \{1, ..., s\}$.*

Theorem 2 characterizes the equilibrium outcomes and evolution of posterior beliefs in screening games when the players are patient. Recall that in a screening game, the equilibrium outcomes of the auxiliary two-stage game entail player 2 to play W against types $i > s$. Item 1 of Theorem 2 states that the expected discounted behavior of player 2 against such types is to play W with probability 1, and that the posterior beliefs attach probability close to one on such types. Item 2 of Theorem 2 considers the case in which inequality (4) holds. Recall that, in this case, the equilibria of the auxiliary two-stage game involves player 2 playing W against types $\{\theta_0, ..., \theta_s\}$ with positive probability. In the unique equilibrium outcome, player 1 randomizes among $\{a_0, \alpha_1, ..., \alpha_s\}$ in a unique way such that the posterior belief after observing α_i is equal to μ_i^b. Moreover, player 2 plays W with probability $\tilde{\sigma}_2(\alpha_i) = \frac{V_1 - u_1(a_0)}{V_1 - u_1(\alpha_i)}$ after observing α_i. The first subitem in this case looks at equilibrium behavior and beliefs conditional on a type $\theta \in \{\theta_1, ..., \theta_s\}$. First, the expected discounted number of periods in which W is played, and the posterior belief is close to μ_i^b is equal to $\tilde{\sigma}_2(\alpha_i)$. Second, the expected discounted probability with which *i)* S is played, and *ii)* the posterior belief is close to μ_i^b, is equal to $1 - \tilde{\sigma}_2(\alpha_i)$. The third and the fourth subitems in this case look at the equilibrium behavior and beliefs conditional on type θ_0 until the stopping time. First, the expected discounted number of periods in which player 2 plays W, the posterior belief is close to $\boldsymbol{\mu_0}$, and player 1's strategy is close to a_0 is $\tilde{\sigma}_1(a_0)$. Second, the expected discounted number of periods in which player 2 plays W, posterior belief is close to μ_i^b, and player 1's strategy is close to α_i is $\tilde{\sigma}_1(\alpha_i)\tilde{\sigma}_2(\alpha_i)$. The fifth subitem states that the expected discounted probability with which player 2 plays S, and the posterior belief is close to μ_i^b conditional on type θ_0 is $\tilde{\sigma}_1(\alpha_i)\left(1-\tilde{\sigma}_2(\alpha_i)\right)$.

Theorem 2 implies that when players are patient, player 2's posterior belief process spends all the payoff-relevant time in a small neighborhood of the beliefs $M := \{\mu_i^b\}_{i\in\{1,\ldots,s\}} \cup \{\boldsymbol{\mu_i}\}_{i\in\{0,s+1,\ldots,K\}}$. Because each element of M is an extreme point of the convex hull of M, once the posterior belief gets close to an element of M, m, then the discounted probability with which the posterior belief moves away from m is very small. This follows because the belief process is a martingale, and hence after reaching m, if beliefs further moved away from m, they would spend some time outside of the set M. This means that, player 2's posterior beliefs almost immediately reach close to an element of M, and stay there. Moreover, once the beliefs reach close to an element of M, the behavior of player 1 stays approximately constant. Hence, reputations are not built gradually. They are built or destroyed almost immediately.

6.1.2. *Behavior in Contracting Games* We now present our results about belief dynamics in contracting games. Recall that if inequality (3) holds , then we let j to be the smallest element of $\{1, 2, ..., s\}$ such that $\mu_0(\theta_0) < \sum_{0<i\leq j}\mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*}$. Let

$$\mathcal{M}_j := \{\mu \in \Delta\Theta : \mu(\theta_j) \geq \mu_j^* \text{ and } \mu(\theta_0) = 1-\mu(\theta_j)\}.$$

If inequality (4) holds, then Lemma 2 and Theorem 1 imply that in contracting games, player 2 plays S almost right away at the beginning of the game against all types. Hence, we cannot obtain tight results about the belief dynamics conditional on these types in this case. However, if inequality (3) holds, then player 2 plays W against types $\theta \in \{\theta_0, \theta_1, ..., \theta_s\}$ with positive probability, and in this case, we have a tight result about belief dynamics against such types. Given a set $A \subseteq \mathbb{R}^{K+1}$ and a vector $x \in \mathbb{R}^{K+1}$, we define $d(x, A) := \inf_{\tilde{x}\in A}\|x-\tilde{x}\|$.

Theorem 3. *Take a sequence $\{\delta_n\}$ of discount factors converging to one and let $\{\sigma_n\}$ be a sequence of Nash equilibria of a sequence of contracting games where along the sequence, the payoff functions and the prior belief μ_0 are fixed, μ_0 satisfies inequality (6), inequality (3), and in the n^{th} game the discount factor is δ_n. Let $\tilde{G}$ be the corresponding auxiliary* *two-stage game, and $(\tilde{\sigma}_1, \tilde{\sigma}_2)$ be the strategy profile that leads to the unique equilibrium outcome distribution of $\tilde{G}$. There exists $\varepsilon^* > 0$ such that for every $\varepsilon \in (0, \varepsilon^*)$ the following hold:*

1. *If $j < s$, then for every $j < i \leq s$,*
   $\lim \mathbb{E}^{\theta_i}_{(\sigma_n)} \left((1-\delta_n) \sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\{d(\mu_t, \boldsymbol{\mu_i})<\varepsilon\}}\right) = 1.$
2. *If $j > 1$, then for every $0 < i < j$,*
   $\lim \mathbb{E}^{\theta_i}_{(\sigma_n)} \left((1-\delta_n) \sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\left\{d\left(\mu_t, \mu_i^b\right)<\varepsilon\right\}}\right) = \tilde{\sigma}_2(\alpha_i),$
3. $\lim \mathbb{E}^{\theta_j}_{(\sigma_n)} \left((1-\delta_n) \sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\{d(\mu_t, \mathcal{M}_j)<\varepsilon\}}\right) = 1.$
4. *If $j > 1$, then for every $0 < i < j$,*
   $\lim \mathbb{E}^{\theta_0}_{(\sigma_n)} \left((1-\delta_n) \sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\left\{d\left(\mu_t, \mu_i^b\right)<\varepsilon, d(\sigma_1(h^t;\theta_0),\alpha_i)<\varepsilon\right\}}\right) = \tilde{\sigma}_1(\alpha_i)\tilde{\sigma}_2(\alpha_i),$
5. $\lim \mathbb{E}^{\theta_0}_{(\sigma_n)} \left((1-\delta_n) \sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\{d(\mu_t, \mathcal{M}_j)<\varepsilon, d(\sigma_1(h^t;\theta_0),\alpha_j)<\varepsilon\}}\right) = \tilde{\sigma}_1(\alpha_j).$

Theorem 3 characterizes the equilibrium outcomes and the evolution of posterior beliefs in contracting games when the players are patient. Lemma 2 and Theorem 1 imply that the expected discounted number of periods in which player 2 plays W against any type $\theta_i \in \{\theta_{j+1}, ..., \theta_s\}$ converges to one. Item 1 of Theorem 3 states that the expected discounted number of periods in which the posterior belief is close to $\boldsymbol{\mu_i}$ conditional on type θ_i also converges to one. Lemma 2 and Theorem 1 imply that the expected discounted number of periods in which player 2 plays W against any type $\theta_i \in \{\theta_1, ..., \theta_{j-1}\}$ converges to $\tilde{\sigma}_2(\alpha_i)$. Item 2 of Theorem 3 states that conditional on type θ_i, the expected discounted number of periods in which player 2 plays W, and the posterior belief is close to μ_i^b converges to $\tilde{\sigma}_2(\alpha_i)$. Item 3 states that conditional on type θ_j, the expected discounted number of periods in which player 2 plays W, and the posterior belief is close to the set $\mathcal{M}_j$ converges to $\tilde{\sigma}_2(\alpha_j)$, which is 1. Items 4 and 5 consider the belief process and behavior conditional on type θ_0. Item 4 states that conditional on type θ_0, the expected discounted number of periods in which the posterior belief is close to μ_i^b, player 2 plays W, and player 1 plays a strategy close to α_i converges to the auxiliary two-stage game's equilibrium probability with which player 1 mimics type θ_i, and player 2 plays W after observing α_i, for all $i \in \{1, ..., j-1\}$. Finally, item 5 states that the expected discounted number of periods in which the posterior belief is close to the set of beliefs $\mathcal{M}_j$, player 2 plays W, and player 1 plays a strategy close to α_j converges to the auxiliary two-stage game's equilibrium probability with which player 1 mimics type θ_j.

In contracting games, we cannot rule out the possibility that player 2 plays S with positive probability at some histories in which the posterior belief assigns a positive probability to two or more commitment types in equilibrium. This may be possible only in the very early stages of the game, because commitment types separate from each other very fast. Therefore, Theorem 3 does not characterize the discounted joint probability with which S is played, and the posterior belief is close to μ_i^b for $i \in \{1, ..., j-1\}$. However, the theorem still implies that conditional on the game not being stopped almost immediately, the posterior beliefs reach close to either the set $\{\mu_i^b\}_{i\in\{1,...,j-1\}}$, the convex hull of μ_j^b and $\boldsymbol{\mu_i}$ (i.e., $\mathcal{M}_j$), or the set $\{\boldsymbol{\mu_i}\}_{i\in\{j+1,...,s\}}$. Let $N := \{\mu_i^b\}_{i\in\{1,...,j\}} \cup \{\boldsymbol{\mu_i}\}_{i\in\{j,...,s\}}$. Each element of N is an extreme point of the convex hull of N. Therefore, once the posterior beliefs reach a small neighborhood of an element of $\{\mu_i^b\}_{i\in\{1,...,j-1\}}$ or $\{\boldsymbol{\mu_i}\}_{i\in\{j+1,...,s\}}$, the posterior beliefs can move out of this neighborhood with only a very small probability. Similarly, if the posterior belief reaches a small neighborhood of $\mathcal{M}_j$, then the posterior beliefs can move away from the set $\mathcal{M}_j$ with only a very small probability. Therefore, posterior beliefs (reputations) again are approximately constant throughout the game. This implies that, conditional on the game not being stopped almost right away, the normal type does not reveal further private information that is payoff-relevant for player 2.

7. Long-Run Equilibrium Behavior

Theorems 1, 2 and 3 provide a tight characterization of the equilibrium payoffs and give insights about the equilibrium behavior of the dynamic game when the players are patient. An outsider may not observe how a dynamic strategic situation unfolds completely. Rather, she may observe some coarse information about the frequency of stopping actions, without observing the duration of the relationship. To connect the equilibrium predictions to such observable outcomes, one needs to know the long-run behavior of the game, which is not captured by reduced outcomes.

In contrast to the results in Theorems 1, 2 and 3, the monitoring structure matters for the long-run behavior in contracting games, while in screening games, the long-run behavior is independent of the monitoring structure. We explain this distinction in the discussions that follows the results.

We start with screening games, and assume that the monitoring structure satisfies Assumption 1, i.e., we do not make the full-support assumption. In a screening game, either inequality (3) holds and a patient player 2 plays S almost immediately against all types $\theta \in \{\theta_0, \theta_1, ..., \theta_s\}$, or inequality (4) holds and player 1 randomizes in such a way that player 2's posterior belief is either close to $\mu_{\mathbf{0}}$ or is close to μ_i^b for some $i \in \{1, 2, ..., s\}$. Our long-run results are about the behavior when the game does not stop in the short run, hence we focus on the case in which inequality (4) holds.

Theorem 4. *Take a sequence $\{\delta_n\}$ of discount factors converging to one and let $\{\sigma_n\}$ be a sequence of Nash equilibria of a sequence of screening games, where along the sequence the payoff functions and the prior belief μ_0 are fixed, μ_0 satisfies inequality (4), and in the n^{th} game the discount factor is δ_n. Let $\tilde{G}$ be the corresponding auxiliary two-stage game, and $(\tilde{\sigma}_1, \tilde{\sigma}_2)$ be the strategy profile that leads to the unique equilibrium outcome distribution of $\tilde{G}$. There exists $\varepsilon^* > 0$ such that for every $\varepsilon \in (0, \varepsilon^*)$ the following hold:*

1. *For $i \in \{1, ..., s\}$,*

   1.1. $\lim \mathbb{P}^{\theta_i}_{(\sigma_n)} \left(\mathbb{T} < \infty, d\left(\mu_{\mathbb{T}}, \mu_i^b\right) < \varepsilon\right) = 1.$

   1.2. $\lim \mathbb{P}^{\theta_0}_{(\sigma_n)} \left(\mathbb{T} < \infty, d\left(\mu_{\mathbb{T}}, \mu_i^b\right) < \varepsilon\right) = \tilde{\sigma}_1(\alpha_i).$

2. $\lim \mathbb{P}^{\theta_0}_{(\sigma_n)} (\mathbb{T} = \infty) = \tilde{\sigma}_1(a_0).$

3. *For $i \in \{s+1, ..., K\}$,* $\lim \mathbb{P}^{\theta_i}_{(\sigma_n)} (\mathbb{T} < \infty) = 0.$

Theorem 4 considers the long-run behavior in screening games when players are patient. The first sub-item of item 1 states that, conditional on any type $\theta_i \in \{\theta_1, ..., \theta_s\}$, player 2 eventually plays S, and when S is played, the posterior belief is close to μ_i^b. The second sub-item states that, conditional on type θ_0, the probability that player 2 will eventually play S at a posterior belief close to μ_i^b is equal to the auxiliary two-stage game's equilibrium probability that player 1 mimics type θ_i. To see why, recall that by Theorem 2, the posterior beliefs conditional on type θ_i spend most of the time close to μ_i^b. Hence, conditional on the posterior belief reaching to μ_i^b, the normal type of player 1 (type θ_0) either pools with the commitment type, or S is played. If S is not played, then the normal type should expect S to be played in some future period. Because such expectations about the future play of S should be provided in each period where there is pooling, eventually player 2 plays S.

Item 2 states that conditional on type θ_0, the probability that player 2 never plays S is equal to the auxiliary two-stage game's equilibrium probability that player 1 plays a_0. Items 1 and 2 together imply that in the long run, player 2 plays S against the normal type with probability $\frac{1}{\mu_0(\theta_0)} \sum_{i \in \{1,2,...,s\}} \mu_0(\theta_i) \frac{1-\mu_i^*}{\mu_i^*}$. Item 3 states that S is not played against types $\theta \in \{\theta_{s+1}, ..., \theta_K\}$.

The reason why full-support assumption is not needed in screening games is that irreversible action serves as rewards that incentivize player 1 to pool with the commitment types. Hence, in every period in which player 1 pools, he also expects S to be played with a positive probability in some future period. In contrast, in contracting games irreversible action serves as a threat to player 1. If monitoring is perfect, and if the normal type of player 1 reveals his type by playing a_0, then player 2 plays S in the next period. Hence, pooling forever is sustained as the unique equilibrium outcome if there is a single commitment type, and if the prior belief on the commitment type is sufficiently high. However, if monitoring structure satisfies the full-support assumption, then pooling forever cannot be sustained. We now analyze this case.

In contracting games, we provide sharp predictions for the long-run behavior for the set of types that play an action which is followed by W with positive probability in the unique equilibrium outcome of the auxiliary two-stage game. Recall that in this equilibrium outcome, against any type θ_i, for $i > s$, player 2 plays S. Therefore, in the dynamic game Theorem 1 implies that player 2 plays S against these types almost immediately, i.e., the game ends against such types in the short run. Hence, our results will be about the types $\theta_i \in \{\theta_0, ..., \theta_s\}$. Moreover, as stated earlier, the full-support assumption is needed to obtain sharp predictions.

Theorem 5. *Take a sequence $\{\delta_n\}$ of discount factors converging to one and let $\{\sigma_n\}$ be a sequence of Nash equilibria of a sequence of contracting games, where along the sequence the payoff functions and the prior belief μ_0 are fixed, μ_0 satisfies inequality (6), P satisfies the full-support assumption, and in the n^{th} game the discount factor is δ_n. Let $\tilde{G}$ be the corresponding auxiliary two-stage game, and $(\tilde{\sigma}_1, \tilde{\sigma}_2)$ be the strategy profile that leads to the unique equilibrium outcome distribution of $\tilde{G}$. There exists $\varepsilon^* > 0$ such that for every $\varepsilon \in (0, \varepsilon^*)$ the following hold:*

1. *For every n, Player 2 plays S against the normal type eventually in the n^{th} game, i.e., $\mathbb{P}^{\theta_0}_{(\sigma_n)}(\mathbb{T} < \infty) = 1$.*

2. *Suppose that μ_0 satisfies inequality (3). Let j to be the smallest element of $\{1, 2, ..., s\}$ such that $\mu_0(\theta_0) < \sum_{0<i\leq j} \mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*}$. There exists $\varepsilon^* > 0$ such that for every $\varepsilon \in (0, \varepsilon^*)$ the following hold:*

    2.1. *For $i \in \{1, ..., j-1\}$, $\lim \mathbb{P}^{\theta_i}_{(\sigma_n)}(\mathbb{T} < \infty) = 1$.*

    2.2. $\lim \mathbb{P}^{\theta_0}_{(\sigma_n)}\left(\mathbb{T} < \infty, d\left(\mu_{\mathbb{T}}, \mu_j^b\right) < \varepsilon\right) = \tilde{\sigma}_1(\alpha_j)$.

    2.3. $\lim \mathbb{P}^{\theta_j}_{(\sigma_n)}\left(\mathbb{T} < \infty, d\left(\mu_{\mathbb{T}}, \mu_j^b\right) < \varepsilon\right) = \left(\frac{\mu_j^*}{1-\mu_j^*}\right)\left(\frac{\mu_0(\theta_0)}{\mu_0(\theta_j)}\right)\tilde{\sigma}_1(\alpha_j)$.

    2.4. *For $i \in \{j+1, ..., s\}$, $\lim \mathbb{P}^{\theta_i}_{(\sigma_n)}(\mathbb{T} < \infty) = 0$.*

Theorem 5 considers the long-run behavior in contracting games. The first item states that if the monitoring structure has full support, player 2 plays S against the normal type eventually, for every discount factor. The intuition for this result follows closely to that provided by Cripps et al. (2004): The normal type does not pool with a commitment type forever when the monitoring structure has full support.

The long-run behavior against the commitment types are sharp only when the parties are patient. Let us start with a type $\theta_i \in \{\theta_1, ..., \theta_{j-1}\}$. According to Theorem 2, the posterior belief conditional on type θ_i reaches a small neighborhood of $\mu_{\mathbf{i}}{}^b$ rapidly when the parties are patient. Theorem 5, item 2, first sub-item states that eventually player 2 plays S against these types when players are patient. To see why, first notice that when players are patient, player 1's equilibrium payoffs change steeply around the belief $\mu_{\mathbf{i}}{}^b$, by Theorem 1. This is because, if the probability of the commitment type θ_i is above μ_i^*, then player 1's equilibrium payoffs are close to the payoff he gets by committing to mimic the commitment type forever, which is a moderately high payoff, while if this probability is below μ_i^*, then player 1's equilibrium payoff is close to his minmax payoff. This discontinuity in the limit equilibrium payoffs around the belief μ_i^b implies that the posterior beliefs move in very small step sizes if they reach a small neighborhood of μ_i^b. This is because otherwise, player 1 would have strong incentives to mimic the commitment type, hampering any information revelation and any movement in the posterior beliefs. In other words, posterior beliefs get absorbed at μ_i^b. We now argue that, player 2 should be playing S in the long run, which follows from an intuition similar to the one provided by Cripps et al. (2004). To see why, assume on the way to a contradiction that player 2 plays W in the long-run with positive probability. Since the belief $\{\mu_{t\wedge\mathbb{T}}\}$ is a bounded martingale, with positive probability the posterior would have to converge to $\mu_{\mathbf{i}}{}^b$ in the long run. But in this case, player 2 would have to ignore a long sequence of signals that are more likely under a_0 then under α_i, which would imply the existence of profitable deviations for the normal type.

The second sub-item of item 2 states that player 2 plays S against the normal type of player 1 eventually, at posterior beliefs close to μ_j^b, with the same probability that player 1 mimics type θ_j in the unique equilibrium of the auxiliary two-stage game. The third sub-item of item 2 states that player 2 plays S against type θ_j with a probability that is strictly between 0 and 1. This probability is found by Bayes' rule by considering the experiment that either reveals that player 1 is type θ_j, and leads to the belief $\boldsymbol{\mu_j}$, or leads to the belief μ_j^b. In this experiment, with probability $\mu_0(\theta_0)\tilde{\sigma}_1(\alpha_j)$, player 1 is the normal type, and with the remaining probability, he is type θ_j. The intuition for these results follows from our findings regarding the discontinuity of limit equilibrium payoffs at belief μ_j^b, and the insights from Cripps et al. (2004).

The last sub-item of item 2 states that Player 2 never plays S against other types.

8. Sketch of the Proof of Theorem 1 for a Screening Game

We provide a sketch of the proof of Theorem 1 for screening games. The purpose is to explain some of the new technical tools we use and develop in our proofs that may be useful in dynamic games with incomplete information. We provide the sketch through an example. We will state some results for this example in formal lemmas, but we will not provide additional proofs for these lemmas. The appendices contain more general versions of these lemmas, and their proofs. We provide such statements to guide the steps and the logic of the general proof.

Player 1 has two actions, $\mathbb{A} = \{a_0, a_1\}$, and there is only one commitment type, i.e., $\Theta = \{\theta_0, \theta_1\}$. The commitment type plays action a_1 at every period until the game is stopped. In the following development, $\mu \in [0, 1]$ refers to the probability that player 1 is a commitment type, or player 1's reputation, and the prior probability that player 1 is a commitment type is $\mu_0 \in (0, 1)$. The signal distribution is governed by P, and P satisfies the identification assumption. We further assume that Assumptions 2-6 are satisfied, and the dynamic game is a screening game. Slightly abusing the notation, let $V_2(\mu) := \mu V_2(\theta_1) + (1 - \mu)V_2(\theta_0)$. Also, let $\mu^* \in (0, 1)$ be the number that satisfies the equality

$$V_2(\mu^*) = u_1(a_1).$$

We will show the sketch of the proof of the following claim, which is weaker than the implication of Theorem 1 for the example:[13]

Claim 1. Take a sequence of Nash equilibria $\{\sigma_n\}_n$ of a sequence of games along which all the parameters except the discount factor are fixed, $\mu_0 \in [0, 1]$, and $\delta_n \to 1$.

1. If $\mu_0 > \mu^*$, then $\lim \mathbb{E}_{(\sigma_n)}\left(\delta_n^{\mathbb{T}}\right) = 1$, $\lim U_1(\sigma_n) = V_1$, and $\lim U_2(\sigma_n) = V_2(\mu_0)$.
2. If $\mu_0 < \mu^*$, then $\lim U_1(\sigma_n) = u_1(a_0)$, $\lim U_2(\sigma_n) = \frac{\mu_0}{\mu^*}u_2(a_1) + \left(1 - \frac{\mu_0}{\mu^*}\right)u_2(a_0)$, and $\lim \mathbb{E}_{(\sigma_n)}^{\theta_1}\left(\delta_n^{\mathbb{T}}\right) = \frac{u_1(a_1)-u_1(a_0)}{V_1-u_1(a_1)}$.

8.1. Centralized Play and Coarsening

[13] We omit the sketch of the proof for other implications of Theorem 1 for space considerations. The current sketch contains the most important tools we use for the general proof.

8.1.1. *Centralized Play:* A generalized centralized play Γ is an auxiliary static direct mechanism in which player 1, who is a commitment type with probability μ_0, first reports his type to the designer. The designer then garbles the information sent by the commitment and the normal types, and sends a message, which we take to be the posterior belief induced by the message, $\mu \in [0,1]$. Let the distribution over posteriors be governed by (Borel-measurable) probability measures, $\lambda : \{\theta_0, \theta_1\} \to \Delta[0,1]$. We write $\lambda(\theta)(\mu)$ for the probability that message μ is given when player 1 reports his type θ. Note that Bayes' rule puts a further restriction on the probability that message $\mu \in [0,1]$ is sent after each report, given by

$$\mu_0 (1-\mu)\,\lambda(\theta_1)(\mu) = (1-\mu_0)\,\mu\lambda(\theta_0)(\mu). \tag{9}$$

For convenience, we refer to $\lambda(\mu) := \mu_0\lambda(\theta_1)(\mu) + (1-\mu_0)\lambda(\theta_0)(\mu)$ as the probability that the posterior belief is μ, given the garbling used by the mechanism. Observe that equation (9) implies that

$$\begin{aligned} \lambda(\theta_0)(\mu) &= \lambda(\mu)\frac{1-\mu}{1-\mu_0}, \\ \lambda(\theta_1)(\mu) &= \lambda(\mu)\frac{\mu}{\mu_0}. \end{aligned} \tag{10}$$

The mechanism also identifies a $\lambda-$measurable function $Y : \Theta \times [0,1] \to \Delta\{a_0, a_1, S\}$. The term $Y(\theta)(\mu)$ corresponds to the probability distribution over action profiles that the designer chooses to play for type θ at belief μ. The tuple $\Gamma = (\lambda, Y)$ denotes the garbling used by the designer, and the distribution over the action profiles, for each $\mu \in [0,1]$. Fix a prior belief μ_0, discount factor δ, and a strategy profile σ in the dynamic game. We say that a generalized centralized play Γ implements σ if the equations (11), (12), and (13) below hold for each $\theta \in \Theta$:

$$\lambda(\theta)(\mu) = \mathbb{E}^{\theta}_{\sigma}\left(\sum_{h^t \in H}\left((1-\delta)\delta^t \mathbb{I}_{\{\mu(h^t)=\mu, t<\mathbb{T}\}} + \mathbb{I}_{\{\mu(h^t)=\mu, \mathbb{T}=t\}}\delta^t\right)\right), \tag{11}$$

$$\lambda(\theta)(\mu)Y(\theta)(\mu)(a) = \mathbb{E}^{\theta}_{(\sigma)}\left(\sum_{h^t \in H}\left(\mathbb{I}_{\{\mu(h^t)=\mu, t<\mathbb{T}, a_t=a\}}(1-\delta)\delta^t\right)\right) \text{ for } a \in \{a_0, a_1\}, \tag{12}$$

$$\lambda(\theta)(\mu)Y(\theta)(\mu)(S) = \lambda(\theta)(\mu)\left(1 - \sum_{a \in \{a_0,a_1\}} Y(\theta)(\mu)(a)\right) = \mathbb{E}^{\theta}_{(\sigma_n)}\left(\mathbb{I}_{\{\mu(h^t)=\mu, \mathbb{T}=t\}}\delta^t\right). \tag{13}$$

The introduction of a centralized play will be a useful tool to investigate player 2's equilibrium payoff bounds. If Γ implements σ, then the probability of each action $a \in \{a_0, a_1\}$ chosen by the generalized centralized play after the garbling leads to belief μ, is equal to the expected discounted number of periods at which i) player 1 plays action a and ii) the public belief that player 1 is a commitment type equals μ. Note first that because the commitment type plays a_1 at every period, we have $Y(\theta_1)(\mu)(a_0) = 0$ for all $\mu > 0$.[14] Second, because in any Nash equilibrium, the normal type can always mimic the commitment type, a generalized centralized play Γ which implements a Nash equilibrium σ satisfies the following incentive compatibility constraint:

$$\int_{[0,1]} \lambda(\theta_0)(\mu) \left(\sum_{a \in \{a_0, a_1\}} Y(\theta_0)(\mu)(a) \left(u_1(a) - V_1\right) + V_1 \right) d[\mu] \geq \tag{14}$$
$$\int_{[0,1]} \lambda(\theta_1)(\mu) \left(\sum_{a \in \{a_0, a_1\}} Y(\theta_1)(\mu)(a) \left(u_1(a) - V_1\right) + V_1 \right) d[\mu]$$

Moreover, because the belief process is a martingale, we obtain that for every $\theta \in \Theta$,

$$\int_{[0,1]} \mu \lambda(\theta_1)(\mu) d[\mu] = \mu_0,$$
$$\int_{[0,1]} (1-\mu)\, \lambda(\theta_0)(\mu) d[\mu] = 1 - \mu_0.$$

We now introduce a strengthening of the notion of generalized centralized play by restricting that $Y(\theta_0)(\mu) = Y(\theta_1)(\mu)$ for almost every μ, calling such mechanisms centralized plays.

Definition 1. A generalized centralized play $\Gamma = (\lambda, Y)$ is a **centralized play** if for almost every posterior $\mu \in \Delta\Theta$ with respect to the measure λ, $Y(\theta_0)(\mu) = Y(\theta_1)(\mu) =: Y(\mu)$.

In general, a strategy profile may not be implemented by a centralized play. However, when players are sufficiently patient, the generalized play that implements a strategy profile is approximated by a centralized play. Formally, take a sequence of dynamic games along which all the parameters of the game except for the discount factor δ are fixed. Let $\{\sigma_n\}_{n=1,\ldots}$ be a sequence of strategy profiles for this sequence of games, and suppose $\delta_n \to 1$. Let $\{\Gamma_n\}_n$ be the associated sequence of generalized centralized plays where each

[14] The implication is stated for $\mu > 0$, because Bayes' rule implies that $\lambda(\theta_1)(0) = 0$, hence there is no restriction on $Y(\theta_1)(0)$.

Γ_n implements σ_n. Each Γ_n induces a probability measure on the set $(\triangle\Theta \times \triangle\{\mathbb{A}, S\})$. Since the set $(\triangle\Theta \times \triangle\{\mathbb{A}, S\})$ is compact, the family of probability measures over them is relatively compact by Prohorov's Theorem (Billingsley (2013), Theorem 5.1), and thus each sequence $\{\Gamma_n\}$ of probability measure has a (weakly) convergent subsequence. In the Appendix, we show that every limit point $\Gamma = (\lambda, Y)$ of $\{\Gamma_n\}_n$ is a centralized play. To show this result, we utilize a powerful learning lemma shown in Fudenberg and Levine (1992). We assume that the entire sequence $\{\Gamma_n\}$ is convergent for the remainder of this sketch.

8.1.2. *Coarsening:* Fix a belief $\bar{\mu} \in (0,1)$. A centralized play $\tilde{\Gamma} = (\tilde{\lambda}, \tilde{Y})$ is a $\bar{\mu}$-coarse centralized play if the support of $\tilde{\lambda}$ is $\{0, \bar{\mu}, 1\}$. For every centralized play $\Gamma = (\lambda, Y)$, we construct a new centralized play $\tilde{\Gamma} = (\tilde{\lambda}, \tilde{Y})$ which is a $\bar{\mu}$-coarsening of Γ that preserves the probability distribution over action profiles and types generated by Γ almost surely. We use the following martingale-splitting to construct $\tilde{\Gamma}$:

$$\begin{aligned}
\tilde{\lambda}(0) &= \int_{[0,\bar{\mu}]} \left(1 - \frac{\mu}{\bar{\mu}}\right) \lambda\,[d\mu] \\
\tilde{\lambda}(1) &= \int_{(\bar{\mu},1]} \left(\frac{\mu - \bar{\mu}}{1 - \bar{\mu}}\right) \lambda\,[d\mu] \\
\tilde{\lambda}(\bar{\mu}) &= 1 - \tilde{\lambda}(0) - \tilde{\lambda}(1).
\end{aligned}$$

Note that $\tilde{\lambda}(\theta)(\mu)$ for $\theta \in \{\theta_0, \theta_1\}$, and for $\mu \in \{0, \bar{\mu}, 1\}$ are derived from $\tilde{\lambda}(\mu)$ using equation (10).

Likewise, we construct $\tilde{Y}$ using λ and Y as follows: For each $a \in \{a_0, a_1, S\}$ let

$$\begin{aligned}
\tilde{Y}(0)(a) &= \frac{\int_{[0,\bar{\mu}]} Y(\mu)(a) \left(1 - \frac{\mu}{\bar{\mu}}\right) \lambda\,[d\mu]}{\tilde{\lambda}(0)} \text{ if } \tilde{\lambda}(0) > 0, \\
\tilde{Y}(\bar{\mu})(a) &= \frac{\int_{[0,\bar{\mu}]} Y(\mu)(a) \left(\frac{\mu}{\bar{\mu}}\right) \lambda\,[d\mu] + \int_{(\bar{\mu},1]} Y(\mu)(a) \left(\frac{1-\mu}{1-\bar{\mu}}\right) \lambda\,[d\mu]}{\tilde{\lambda}(\bar{\mu})} \text{ if } \tilde{\lambda}(\bar{\mu}) > 0, \\
\tilde{Y}(1)(a) &= \frac{\int_{(\bar{\mu},1]} Y(\mu)(a) \left(\frac{\mu-\bar{\mu}}{1-\bar{\mu}}\right) \lambda\,[d\mu]}{\tilde{\lambda}(1)} \text{ if } \tilde{\lambda}(1) > 0.
\end{aligned}$$

If $\tilde{\lambda}(\mu) = 0$ for some $\mu \in \{0, \bar{\mu}, 1\}$, then we put no further restriction on $\tilde{Y}(\mu)(\cdot)$ except that $\tilde{Y}(\mu)(a) \geq 0$ for every $a \in \{a_0, a_1, S\}$ and $\sum_{a \in \{a_0, a_1, S\}} \tilde{Y}(\mu)(a) = 1$. Essentially, when coarsening a centralized play, we split the beliefs $\mu \in [0, \bar{\mu}]$ to the set of beliefs $\{0, \bar{\mu}\}$, and

(a) $\mu \leq \bar{\mu}$ (b) $\mu > \bar{\mu}$

Figure 3: This panel illustrates martingale-splitting used in the coarsening procedure. If $\mu \leq \bar{\mu}$, as in the left figure, then $\lambda(\mu)$ is split to $\tilde{\lambda}(0)$ and $\tilde{\lambda}(\mu)$ with the weights $1 - \frac{\mu}{\bar{\mu}}$ and $\frac{\mu}{\bar{\mu}}$, respectively. If $\mu > \bar{\mu}$, as in the right figure, then a similar procedure is used to split $\lambda(\mu)$ to $\tilde{\lambda}(1)$ and $\tilde{\lambda}(\mu)$ with the weights $1 - \frac{1-\mu}{1-\bar{\mu}}$ and $\frac{1-\mu}{1-\bar{\mu}}$, respectively.

beliefs $\mu \in (\bar{\mu}, 1]$ to the set of beliefs $\{\bar{\mu}, 1\}$. Figure 3 illustrates this procedure for λ.

Remark 1. Every coarsening of a centralized play preserves the distribution over action profiles type by type, hence it preserves players' payoffs.

8.2. Reputation Boundary We start with the definition of a reputation boundary. Verbally, a belief $\bar{\mu} \in (0, 1)$ is a reputation boundary if whenever player 1's reputation is above the boundary $\bar{\mu}$, a patient player 2 plays S almost immediately in all Nash equilibria (NE).

Definition 2. A belief $\bar{\mu} \in (0, 1)$ is a reputation boundary if for all $\hat{\mu} > \bar{\mu}$,

$$\liminf_{\delta \nearrow 1} \inf_{\tilde{\mu}_0 \geq \hat{\mu}} \left\{ \begin{array}{c} \mathbb{E}_{(\sigma)}\left[\delta^{\mathbb{T}}\right] : \sigma \text{ is a NE for} \\ \text{the game with prior } \tilde{\mu}_0 \text{ and the discount factor } \delta \end{array} \right\} = 1. \tag{15}$$

Lemma 3. *A reputation boundary exists.*

Existence of such a boundary is obtained rather easily using an auxiliary static mechanism-design program. In this program, player 2 chooses a strategy profile against a normal type, and a strategy against the commitment type, that maximize his expected payoff subject to a static incentive constraint that the normal type does not strictly prefer to follow the strategy profile for the commitment type. Any equilibrium strategy profile that obeys (on path) sequential rationality constraints satisfies the static constraints implied by this program. When the initial belief that player 1 is a commitment type is close to 1, the solution is attained uniquely by player 2 playing S immediately. Therefore, a stronger result is obtained for the special case in which there is a single commitment type: there exists a boundary $\bar{\mu} < 1$ such that if $\mu > \bar{\mu}$, then $\mathbb{T} = 0$ in all Nash equilibria, independent of δ. Reputation boundary separates the region of high reputations, where further screening is not sequentially rational for player 2, from the region of lower reputations where some screening may be possible.

8.3. High Reputation Suppose $\mu_0 > \mu^*$. We will now argue that if players are sufficiently patient, then player 2 plays S almost immediately, in all Nash equilibria. In other words, we will show that μ^* is a reputation boundary.

To this end, assume towards a contradiction that the infimum of all reputation boundaries is $\bar{\mu} > \mu^*$, and observe that Definition 2 implies that $\bar{\mu}$ is itself a reputation boundary. We will find $\varepsilon \in (0, \bar{\mu} - \mu^*)$ such that the property highlighted in equation (15) holds when we replace $\bar{\mu}$ with $\bar{\mu} - \varepsilon$, i.e., $\bar{\mu} - \varepsilon$ is a reputation boundary. This establishes the contradiction to the hypothesis that $\bar{\mu}$ is the infimum over all reputation boundaries.

Take a sequence of games in which the prior is μ_0, and the discount factor along the sequence converges to 1. Also, take a sequence of Nash equilibria of the sequence of games, $\{\sigma_n\}$, and generalized centralized plays, $\{\Gamma_n\}$, where each Γ_n implements σ_n. Let Γ be a limit point of Γ. As we discussed previously, Γ is a centralized play. Because each Γ_n satisfies the IC constraint (14), so does Γ. Moreover, since the commitment type never plays a_0,

$$\int_{\mu>0} \lambda(\mu) Y(\mu)(a_0) d[\mu] = 0. \tag{16}$$

Finally, because of the definition of the reputation boundary, as $\delta_n \to 1$, the probability that action S is taken when the posterior belief $\mu_t > \bar{\mu}$ approaches to 1. Hence, we have that

$$\sum_{a\in\{a_0,a_1\}} \int_{\mu>\bar{\mu}} \lambda(\mu) Y(\mu)(a) d[\mu] = 0. \tag{17}$$

We call a centralized play limit-equilibrium-compatible if it satisfies the incentive compatibility constraint (14), and satisfies equalities (16) and (17). We conclude with the following Lemma.

Lemma 4. Γ *is limit-equilibrium-compatible.*

Player 2's behavior when $\mu_0 \in [\bar{\mu} - \varepsilon, \bar{\mu}]$: Suppose now that $\mu_0 \in [\bar{\mu} - \varepsilon, \bar{\mu}]$ for some small $\varepsilon \in (0, \bar{\mu} - \mu^*)$. In light of Lemma 4 and Remark 1, player 2's limit equilibrium payoff along the sequence $\{\sigma_n\}$ converges to the payoff she would get from the limit-equilibrium-compatible centralized play $\Gamma = (\lambda, Y)$. Let $\tilde{\Gamma} = \left(\tilde{\lambda}, \tilde{Y}\right)$ be the $\bar{\mu}$-coarsening of Γ. Player 2's payoff in $\tilde{\Gamma}$ is the same as her payoff in Γ, and is given by

$$\left[\begin{array}{c} \tilde{\lambda}(0)\left(\sum_{a\in\{a_0,a_1\}} \tilde{Y}(0)(a)\left(u_2(a) - V_2(0)\right) + V_2(0)\right) \\ +\tilde{\lambda}(\bar{\mu})\left(\tilde{Y}(\bar{\mu})(a_1)\left(u_2(a_1) - V_2(\bar{\mu})\right) + V_2(\bar{\mu})\right) \\ \tilde{\lambda}(1)V_2(1) \end{array}\right], \tag{18}$$

and the incentive compatibility constraint is

$$\tilde{\lambda}(\theta_0)(0)\left[\sum_{a\in\{a_0,a_1\}}\tilde{Y}(0)(a)\left(u_1(a)-V_1\right)+V_1\right]+\tilde{\lambda}(\theta_0)(\bar{\mu})\left[\tilde{Y}(\bar{\mu})(a_1)\left(u_2(a_1)-V_1\right)+V_1\right]$$
$$\geq \tilde{\lambda}(\theta_1)(\bar{\mu})\left[\tilde{Y}(\bar{\mu})(a_1)\left(u_2(a_1)-V_1\right)+V_1\right]+\tilde{\lambda}(\theta_1)(1)V_1. \tag{19}$$

Let $\Psi(\Gamma')$ be the value achieved by (18) under a limit-equilibrium compatible centralized play $\Gamma'=(\lambda',Y')$. Consider the problem of maximizing $\Psi(\Gamma')$ across the set of all limit-equilibrium-compatible centralized plays. Recall that $V_2(\mu_0)=\mu_0 V_2(\theta_1)+(1-\mu_0)V_2(\theta_0)$ is the payoff of player 2 from playing S at the beginning of the game. We show in the Appendix the following Lemma, which we use to extend the reputation boundary:

Lemma 5. *There exists $\varepsilon>0$ such that if $\mu_0\in[\bar{\mu}-\varepsilon,\bar{\mu}]$, then for every limit-equilibrium-compatible centralized play $\Gamma'=(\lambda',Y')$, $\int_{\mu\in[0,1]}\lambda'(\mu)Y'(\mu)(S)d[\mu]<1$ implies*

$$\Psi(\Gamma')<V_2(\mu_0).$$

Lemma 5 implies that $\bar{\mu}-\varepsilon$ is a reputation boundary. This is because player 2's equilibrium payoffs are bounded below by $V_2(\mu_0)$ for any discount factor. If, along a sequence of equilibria, player 2 does not play S almost immediately against some type, then the limit-centralized play satisfies $\int_{\mu\in[0,1]}\lambda(\mu)Y(\mu)(S)d\mu<1$. Then, Lemma 5 implies that player 2's limit equilibrium payoff is less than $V_2(\mu_0)$, which contradicts player 2's behavior being part of an equilibrium. Hence, we obtain the following result, which states that for any prior $\mu_0>\mu^*$, a sufficiently patient player 2 almost immediately takes the irreversible action in any Nash equilibrium, and her equilibrium payoffs converge to $V_2(\mu_0)$, i.e., μ^* is a reputation boundary. This also completes the sketch of the first item of Claim 1.

Lemma 6. *μ^* is a reputation boundary.*

Corollary 1. *Take a sequence of Nash equilibria $\{\sigma_n\}_n$ of a sequence of games along which all the parameters except the discount factor is fixed, $\mu_0\in[0,1]$, and $\delta_n\to 1$. Let $\{\Gamma_n\}_n$ be the associated sequence of generalized centralized plays where Γ_n implements σ_n, and let $\Gamma=(\lambda,Y)$ be the limit point of $\{\Gamma_n\}_n$. Then, $Y(\mu)(S)=1$ for almost every $\mu\in(\mu^*,1]$ with respect to the measure λ.*

Corollary 1 follows from Lemma 6 for the following reason. Fix a Nash equilibrium, a belief $\mu>\mu^*$, and consider the first on path history h^t at which the posterior belief $\mu(h^t)=\mu$. The continuation strategy profile at h^t is a Nash equilibrium of the game in which the prior is μ. Hence, Lemma 6 implies that S is played almost immediately. But this means that the expected discounted number of periods in which the posterior belief is equal to μ, and W is played, is close to 0. Hence, either $Y(\mu)(S) = 1$, or $\lambda(\mu) = 0$.

8.4. Low Reputation We now sketch the proof of the second item of Claim 1. We start by showing that player 2 does not play S at histories in which the posterior belief is strictly less than μ^*.

Lemma 7. *Take a sequence of Nash equilibria $\{\sigma_n\}_n$ of a sequence of games along which all the parameters except the discount factor is fixed, $\mu_0 \in [0, 1]$, and $\delta_n \to 1$. Let $\{\Gamma_n\}_n$ be the associated sequence of generalized centralized plays where Γ_n implements σ_n, and let $\Gamma = (\lambda, Y)$ be the limit point of $\{\Gamma_n\}_n$. Then,*

$$\int_{\mu \in [0,\mu^*)} \lambda(\mu) Y(\mu)(S) d[\mu] = 0.$$

To show Lemma 7, we first show that for any fixed prior $\mu_0 = \mu < \mu^*$, a sufficiently patient player 2 never plays S at the beginning of the game. To see why this implies Lemma 7, observe that in any equilibrium, and at any on path history h^t with $\mu(h^t) = \mu$, the continuation strategy profile at h^t is a Nash equilibrium of the game in which the prior is μ, and hence S is not played at such histories, delivering the result.

We now explain why for any fixed prior $\mu_0 = \mu < \mu^*$, a sufficiently patient player 2 never plays S at the beginning of the game. Suppose that $u_2(a_0) \geq u_2(a_1)$. Then, because $\mu < \mu^*$, $V_2(\mu) < u_2(a_1) \leq \min\{u_2(a_0), u_2(a_1)\}$, and never playing S gives player 2 a strictly higher payoff than playing S right away. The more challenging case is if $u_2(a_0) < u_2(a_1)$. In this case, consider an alternative strategy for player 2, σ_2', that plays S at the first history when the posterior belief reaches weakly above μ^*, and plays W at all other histories.[15] If player 2 plays S at the beginning of the game with a positive probability, then her equilibrium payoff is $V_2(\mu)$. However, if she plays σ_2', regardless of player 1's strategy, her payoff is bounded below by a number that converges as $\delta \to 1$ to

$$\left(1 - \frac{\mu}{\mu^*}\right) u_2(a_0) + \frac{\mu}{\mu^*} u_2(a_1). \tag{20}$$

Observe that for $\mu < \mu^*$

$$g(\mu) := \left(1 - \frac{\mu}{\mu^*}\right) u_2(a_0) + \frac{\mu}{\mu^*} u_2(a_1) > V_2(\mu),$$

[15] Note that since this strategy may be a deviation, the play may be off-path with respect to an equilibrium if player 2 plays according to σ_2'.

because $g(\cdot)$ and $V_2(\cdot)$ are affine functions, and $g(0) > V_2(0)$, and $g(\mu^*) = V_2(\mu^*)$. Hence, a sufficiently patient player has a profitable deviation from a strategy in which she plays S with a positive probability at the beginning of the game.

We obtain the lower bound on player 2's payoff in the Expression (20) from the strategy σ_2' by considering all strategies of player 1, not only incentive-compatible ones. Let $\{\sigma_n'\}_n$ be a sequence of strategy profiles of a sequence of games where $\delta_n \to 1$, and player 2's strategy in σ_n' is equal to σ_2'. Let $\{\Gamma_n'\}$ be the sequence of generalized centralized plays that implements $\{\sigma_n'\}$ with a limit point Γ'. Let $\tilde{\Gamma}'$ be the μ^*-coarsening of Γ'. Because σ_2' plays S at histories with posterior belief greater than μ^*, and never plays S when this belief is less than μ^*, by the properties of the coarsening procedure, we have $\tilde{Y}'(0)(S) = 0$, $\tilde{Y}'(1)(S) = 1$ and $\tilde{Y}'(\mu^*)(a_0) = 0$. Consider minimizing player 2's payoff across coarse centralized plays, $\tilde{\Gamma}$, that satisfy the aforementioned properties. The inequality $u_2(a_0) < u_2(a_1)$ implies that the objective function is minimized when $\tilde{Y}(0)(a_0) = 1$. Moreover, $V_2(\mu^*) = u_2(a_1)$. Hence, player 2's payoff minimization problem boils down to:

$$\min_{\lambda(1)\in[0,\mu]} \lambda(0)u_2(a_0) + \lambda(\mu^*)V_2\left(\mu^*\right) + \lambda(1)V_2\left(1\right)$$

Subject to:

$$\sum_{\mu\in\{0,\mu^*,1\}} \lambda(\mu) = 1, \text{ and}$$

$$\mu^*\lambda(\mu^*) = \mu - \lambda(1)$$

The first constraint is imposed because the support of λ is $\{0, \mu^*, 1\}$, and the second constraint follows because Bayes' rule implies $\frac{\mu}{1-\mu}\frac{\lambda(\theta_1)(\mu^*)}{\lambda(\theta_0)(\mu^*)} = \frac{\mu^*}{1-\mu^*}$, $\lambda(\theta_0)(1) = 0$, $\lambda(\theta_1)(0) = 0$. After plugging the constraints in the objective function, and rearranging, the coefficient in front of $\lambda(1)$ becomes positive, hence the program is minimized at $\lambda(1) = 0$. This gives us the lower bound in the Expression (20).

Remark 2. Take a sequence of Nash equilibria $\{\sigma_n\}_n$ of a sequence of games along which all the parameters except the discount factor is fixed, $\mu_0 \in [0, 1]$, and $\delta_n \to 1$. Let $\{\Gamma_n\}_n$ be the associated sequence of generalized centralized plays where Γ_n implements σ_n, and let $\Gamma = (\lambda, Y)$ be the limit point of $\{\Gamma_n\}_n$. Then, Lemmata 6, 7 and Γ being a centralized play imply the following:

1. $\int_{\mu>\mu^*} \lambda(\mu)Y(\mu)(S)d[\mu] = \int_{\mu>\mu^*} \lambda(\mu)d[\mu]$.

2. $\int_{\mu\in[0,\mu^*)} \lambda(\mu)Y(\mu)(S)d[\mu] = 0$.

3. $\int_{\mu\in(0,\mu^*)} \lambda(\mu)Y(\mu)(a_1)d[\mu] = \int_{\mu\in(0,\mu^*)} \lambda(\mu)d[\mu]$.

Therefore, if $\tilde{\Gamma}$ is a μ^*-coarsening of Γ, then $\tilde{\lambda}(0)\tilde{Y}(0)(S) = 0$, $\tilde{\lambda}(1)\left(1 - \tilde{Y}(1)(S)\right) = 0$ and $\tilde{\lambda}(\mu^*)\tilde{Y}(\mu^*)(a_0) = 0$.

Our next step is to provide an upper bound on player 1's equilibrium payoffs when his reputation is $\mu < \mu^*$.

Lemma 8. *Take a sequence of Nash equilibria $\{\sigma_n\}_n$ of a sequence of games along which all the parameters except the discount factor is fixed, $\mu_0 \in [0, \mu^*)$, and $\delta_n \to 1$. Player 1's equilibrium payoffs converge to $u_1(a_0)$.*

Suppose towards a contradiction that Lemma 8 is false. Because player 1 can always guarantee himself a payoff of at least $u_1(a_0)$, the contradiction hypothesis is that we can find a sequence of Nash equilibria $\{\sigma_n\}_n$ for which $\lim U_1(\sigma_n) > u_1(a_0)$. Let $\{\Gamma_n\}_n$ be the associated sequence of generalized centralized plays where Γ_n implements σ_n, and let $\Gamma = (\lambda, Y)$ be the limit point of $\{\Gamma_n\}_n$. Also let $\tilde{\Gamma}$ be the μ^*-coarsening of Γ.

Recall that $\tilde{\lambda}(\theta_0)(1) = \tilde{\lambda}(\theta_1)(0) = 0$ for any centralized play, due to Bayes' rule. Furthermore, because $\mu_0 < \mu^*$, equation (9) implies that $\tilde{\lambda}(\theta_1)(\mu^*) > \tilde{\lambda}(\theta_0)(\mu^*)$. Notice that the payoff of player 1 in $\tilde{\Gamma}$ is given by:

$$U_1 := \tilde{\lambda}(\theta_0)(0)\left(\sum_{a\in\{a_0,a_1\}} \tilde{Y}(0)(a)u_1(a)\right) + \tilde{\lambda}(\theta_0)(\mu^*)\left(\tilde{Y}(\mu^*)(a_1)\left(u_1(a_1) - V_1\right) + V_1\right), \tag{21}$$

where we used $\tilde{\lambda}(\theta_0)(1) = 0$, and Remark 2. The incentive compatibility constraint is

$$U_1 \geq \tilde{\lambda}(\theta_1)(\mu^*)\left(\tilde{Y}(\mu^*)(a_1)\left(u_1(a_1) - V_1\right) + V_1\right) + \tilde{\lambda}(\theta_1)(1)V_1, \tag{22}$$

where we used $\tilde{\lambda}(\theta_1)(0) = 0$, and Remark 2. Now suppose on the way to a contradiction that $U_1 > u_1(a_0)$. Because $\tilde{\lambda}(\theta_0)(0) + \tilde{\lambda}(\theta_0)(\mu^*) = 1$, and because $u_1(a_1) < u_1(a_0)$, $U_1 > u_1(a_0)$ implies that

$$\begin{aligned}
&\tilde{Y}(\mu^*)(a_1)\left(u_1(a_1) - V_1\right) + V_1 > u_1(a_0), \text{ and} \\
&\tilde{Y}(\mu^*)(a_1)\left(u_1(a_1) - V_1\right) + V_1 \geq U_1.
\end{aligned}$$

Because $V_1 \geq Y(\mu^*)(a_1)\left(u_1(a_1) - V_1\right) + V_1$, inequality (22) is satisfied only if $\tilde{\lambda}(\theta_0)(0) = 0$, and $\tilde{\lambda}(\theta_1)(1)\tilde{Y}(\mu^*)(a_1) = 0$. However, if $\tilde{\lambda}(\theta_0)(0) = 0$, then $\tilde{\lambda}(\theta_0)(\mu^*) = 1$, which contradicts to $\tilde{\lambda}(\theta_1)(\mu^*) > \tilde{\lambda}(\theta_0)(\mu^*)$.

We have shown that if $\mu_0 < \mu^*$, the player 1's equilibrium payoffs converge to $u_1(a_0)$, as the players get arbitrarily patient. We now argue that $\tilde{\lambda}(\theta_1)(1) = 0$, and $\tilde{Y}(\mu^*)(a_1)\,(u_1(a_1) - V_1) + V_1 = u_1(a_0)$, continuing with the arguments we used for the sketch of Lemma 8. Suppose on the way to a contradiction that $\tilde{Y}(\mu^*)(a_1)\,(u_1(a_1) - V_1) + V_1 < u_1(a_0)$. Then $U_1 = u_1(a_0)$ implies that $\tilde{\lambda}(\theta_0)(\mu^*) = 0$, which in turn implies by Bayes' rule that $\tilde{\lambda}(\theta_1)(\mu^*) = 0$. Hence, $\tilde{\lambda}(\theta_1)(1) = 1$, which violates the incentive compatibility constraint, (22). Suppose on the way to a contradiction that $\tilde{Y}(\mu^*)(a_1)\,(u_1(a_1) - V_1) + V_1 > u_1(a_0)$. Then, $U_1 = u_1(a_0)$ implies that again the incentive compatibility constraint, (22), is violated. So, we obtain that $\tilde{Y}(\mu^*)(a_1)\,(u_1(a_1) - V_1) + V_1 = u_1(a_0)$. To see that $\tilde{\lambda}(\theta_1)(1) = 0$, we again use the incentive compatibility constraint, (22), and notice that $\tilde{\lambda}(\theta_1)(1) > 0$ violates the constraint.

The implication of $\tilde{\lambda}(\theta_1)(1) = 0$, and $\tilde{Y}(\mu^*)(a_1)\,(u_1(a_1) - V_1) + V_1 = u_1(a_0)$ is that $1 - \tilde{Y}(\mu^*)(a_1) = \lim \mathbb{E}^{\theta_1}_{(\sigma_n)}\left(\delta_n^{\mathbb{T}}\right) = \frac{u_1(a_0) - u_1(a_1)}{V_1 - u_1(a_1)}$. Finally, note that $\mu < \mu^*$ and $\tilde{\lambda}(\theta_1)(1) = 0$ imply by Bayes' rule (see equation (9)) that $\tilde{\lambda}(\theta_0)(0) = 1 - \frac{\mu_0}{1-\mu_0}\frac{1-\mu^*}{\mu^*} > 0$. Hence, $U_1 = u_1(a_0)$, and $\tilde{Y}(\mu^*)(a_1)\,(u_1(a_1) - V_1) + V_1 = u_1(a_0)$ imply that $\tilde{Y}(0)(a_0) = 1$, and we obtain that $\lim U_2(\sigma_n) = \frac{\mu_0}{\mu^*} u_2(a_1) + \left(1 - \frac{\mu_0}{\mu^*}\right) u_2(a_0)$.

9. Relation to the Literature

This paper is closely related to the reputation literature with imperfect monitoring, reputation literature with long-lived players and the literature on dynamic principal-agent models without commitment.

The first strand of literature we relate to is the reputation literature with imperfect monitoring. The closest papers to ours are Fudenberg and Levine (1992) and Cripps et al. (2004). Both papers study reputation effects in repeated games with imperfect monitoring, whereas we study a dynamic game with irreversible actions.

Fudenberg and Levine (1992) study a repeated game played between a long-lived, patient and informed player against a sequence of uninformed, short-lived opponents, in which the informed player's actions are observed with noise. They find that the informed player receives her Stackelberg payoff as long as there is a small but positive probability that the informed player can be a commitment type who plays the Stackelberg action at every period. In our dynamic stopping model, both players have equal discount factors, and the payoff of the uninformed player if she stops the game depends her opponent's type, i.e., there are interdependent values. In our model, the informed player does not get his Stackelberg payoff in equilibrium. We focus on the limit equilibrium behavior for any initial belief, and the equilibrium payoffs and behavior depends on this initial belief even when the players are patient. The most important connection between our model and theirs is that we utilize a learning result shown in Fudenberg and Levine (1992).

Pei (2018) studies a repeated game with a long-lived player and infinitely many short-lived opponents, and assumes payoffs are interdependent. He shows that reputation effects may fail when values are interdependent. In monotone-supermodular games, he shows that all equilibria involve the informed player signaling his type initially by choosing an action, and always playing that action. In our model, the principal's payoff from stopping the game depends on the agent's type, hence in our model payoffs are likewise interdependent. Different from Pei (2018), we study a dynamic game with an irreversible action, the players share a common discount factor, and informed player's actions are allowed to be observed with noise.

Cripps et al. (2004) showed that in repeated games with reputation effects, when monitoring is imperfect and the information structure has full support, eventually the private information of the informed party gets revealed, and the behavior converges to an equilibrium of the repeated game with complete information. We use this insight to characterize the long-run behavior when there is full-support imperfect monitoring in contracting games. We use a similar but different insight to characterize the long-run behavior in screening games, without assuming the full-support assumption. The insight of Cripps et al. (2004) led us to explore whether the inevitability of information revelation will lead the principal to experiment and wait further. In our model, until the game is stopped, the informed party reveals his private information progressively. In a repeated game, progressive information revelation leads to eventually full revelation of the informed player's types. However, in our dynamic game, the uninformed player may stop the game at certain beliefs. This prevents complete learning of the informed player's type. The beliefs at which the game stops depends on the principal's expectation of the speed by which further information will be revealed. We characterize such beliefs at which the game will eventually be stopped, and this allows us to calculate the long-run probability with which the game stops against each type. To do so, we need to focus on the case of patient players, which is unlike Cripps et al. (2004), whose results hold for every discount factor.

Our paper is related to reputation effects in repeated games with long-lived players. In this line of research, the closest papers to ours are Cripps et al. (2005) and Atakan and Ekmekci (2012). The first paper shows a reputation result when the stage game is a simultaneous-move strictly conflicting interest game, and the latter one shows a reputation result when the stage game is a game of perfect information and is a strictly conflicting interest game or a locally non-conflicting interest game. Both papers assume that the action or the move of the informed player is perfectly observed. Atakan and Ekmekci (2015) assumes that the uninformed player's actions are observed with noise, which allows them to study simultaneous-move games and also extend their previous reputation result to a more robust type of uncertainty about commitment types. We borrow some techniques developed by this previous work. Most importantly, in screening games, to obtain the limit equilibrium payoff function we use mechanism design approach iteratively to get better predictions about Player 2's equilibrium behavior. Atakan and Ekmekci (2012, 2015) use a dynamic programming technique iteratively to obtain a reputation result. Apart from these, we introduce the notion of centralized plays, and an auxiliary mechanism-design problem to obtain equilibrium properties. We also use a suitable version of martingale-splitting (coarsening procedure) that was introduced by Aumann and Maschler (1995), to decrease the dimensionality of the auxiliary mechanism-design problem, and solve it.

Our paper is also related to dynamic principal-agent models without commitment. The seminal papers Sobel (1985) and Watson (1999) study models in which a principal starts setting small stakes in a relationship to screen the agent's type, and increases the stakes as his beliefs about the agent being a trustworthy type increases. Similar forces appear in screening games studied in this paper. Hart and Tirole (1988) study a model in which a seller who repeatedly sells a good to a buyer cannot commit not to decrease its price after learning that the buyer has a lower valuation. Similar forces appear in contracting games. Relatedly, Laffont and Tirole (1987, 1988, 1993), Gerardi and Maestri (2018) and Acharya and Ortner (2017) study the ratchet effect in a similar context. In these papers, the principal cannot commit not to change the terms of trade after learning the agent's type.[16] In our models, the set of screening tools available to the principal is coarse, and the principal has an irreversible action. Moreover, we allow for the agent's actions to be observed with noise, whereas all papers we mentioned in this literature analyze situations with perfect monitoring.

Our paper is also related to screening games with Coasian dynamics. Gul et al. (1986), Deneckere and Liang (2006) and Strulovici (2017) study screening problems modeled through a bargaining process. Liu et al. (Forthcoming) studies Coase conjecture when the seller faces multiple buyers every period, and uses an auction with reserve price to screen the buyers. In our model, different from this literature, the principal cannot adjust

[16] A recent strand of this literature has given a new perspective to dynamic screening problems. Malcomson (2016) studies the ratchet effect in relational contracts. In his model, the agent's actions are observed perfectly. Bhaskar and Mailath (2019) study a dynamic moral hazard problem with initial symmetric information, hidden actions and imperfect monitoring. Unobservability of the agent's action leads to asymmetric information in the out-of-equilibrium paths, and lead to the ratchet effect. In these papers, the principal does not have an irreversible action.

the terms of the relationship finely, and the agent's actions are observed with noise.

10. Discussion

10.1. Multiple Stopping Actions In our model, the principal has a single stopping action. Consider a variation of our model in which if player 2 decides to stop the game, she may choose one of two possible irreversible actions. For a concrete example, consider a hybrid model that incorporates the games described in Examples 1 and 2. In any period when the game has not yet stopped, player 2 chooses W, or one of two possible irreversible actions, S_p or S_f (stands for *stop and promote* and *stop and demote,* respectively) which ends the game. In any period when the game has not yet stopped, player 1 chooses his action from the set $\mathbb{A} = \{a_0, a_1\}$. Player 1 is either a normal type (θ_0), or a commitment type (θ_1), who plays a_1 at every period of the game. Let $\mu_0 \in (0, 1)$ be the probability that player 1 is a commitment type. We allow for imperfect monitoring of player 1's actions, and maintain Assumption 1 on the monitoring structure.

Players' payoff function in any period when the game has not stopped yet are identical to those in the examples, i.e., $u_1(a_1) = 0, u_1(a_0) = 1, u_2(a_1) = 1$, and $u_2(a_0) = 0$. The lump-sum continuation payoffs when the stopping action is taken depends on whether S_p or S_f is chosen by player 2. If S_p is chosen, then the payoffs are as in Example 1, i.e., $V_1(S_p) = 2$, $V_2(S_p, \theta_0) = -1$, and $V_2(S_p, \theta_1) = 2$. If S_f is chosen, then the payoffs are as in Example 2, i.e., $V_1(S_f) = -1$, $V_2(S_f, \theta_0) = 2$, and $V_2(S_f, \theta_1) = -1$. Hence, player 2 prefers to play S_p against the commitment type, and S_f against the normal type, while player 1 prefers S_p to S_f.

We first observe that if the normal type of player 1 is expected to pool with the commitment type (i.e., play a_1) at every period of the game when the game has not been stopped yet, then player 2's best response is to play S_f if $\mu_0 < 1/3$, and to play S_p if $\mu_0 > 2/3$, at the beginning of the game. If $\mu_0 \in (1/3, 2/3)$, then her best response would be to never stop the game (along the path generated by player 1's expected behavior). We have verified that, as the players get arbitrarily patient, the equilibrium outcomes measured as the expected discounted number of periods in which each action profile is taken against each type, converge to the unique equilibrium outcome of the auxiliary two-stage game, which is defined in the same fashion as it is defined in our main model. In particular, if $\mu_0 < 1/3$, then player 2 plays S_f almost immediately against both types. If $\mu_0 > 2/3$, then player 2 plays S_p almost immediately against both types. If $\mu_0 \in (1/3, 2/3)$, then the outcome is as if player 2 always plays W, and player 1 plays a_1. In the long run, however, if the monitoring structure satisfies the full-support assumption, then the posterior beliefs eventually reach the set $\{1/3, 2/3\}$, and player 2 eventually stops the game, and plays either S_f when posterior belief is close to $1/3$, or plays S_p when the posterior belief is close to $2/3$.

10.2. Class of Games Assumption 3 restricts the class of games we study. Importantly, it implies that the dynamic game has a unique equilibrium when player 1's type is known. We now discuss two examples for a preliminary exploration of how our results would change if we relaxed Assumption 3.

First, consider a variation of Example 2, in which the only change is $V_2(\theta_0) = 0.5$. With this modification, the game is no longer a contracting game, because $V_2(\theta_0) < u_2(a_1) = 1$. In this game, player 2 would prefer not to play S against the commitment type, or against the normal type if he is expected to pool with the commitment type, and prefers to play S against the normal type if she expects player 1 to play a_0 for a sufficiently long period. In this game, a reputation result holds: For a fixed $\mu_0 > 0$, player 1's Nash equilibrium payoffs are bounded below by a payoff that converges to $u_1(a_1) = 1$ as the parties become arbitrarily patient. To see why, first observe that, for any discount factor δ, there exists a threshold $\mu(\delta) > 0$ such that player 2 does not play S when her posterior belief about player 1's type is above $\mu(\delta)$, in any Nash equilibrium. Moreover, $\lim_{\delta \to 1} \mu(\delta) = 0$. This is because for a patient player 2, the cost of playing W for a number of periods until the posterior belief falls close to 0 is smaller than the cost of playing S against the commitment type. Second, observe that, by mimicking the commitment type, a patient player 1 can guarantee himself a payoff close to $u_1(a_1)$. This is because the probability that player 2's posterior belief never falls below $\mu(\delta)$, conditional on player 1 always playing a_1, converges to 1 when $\mu(\delta)$ converges to zero. These two observations deliver the result. However, when player 1 is known to be the normal type, and when players are patient, there are multiple payoff profiles that can be sustained in equilibrium. Therefore, our techniques do not allow us to obtain a payoff or behavior characterization result in this example.

Second, consider a variation of Example 1, in which the only change is $u_2(a_1) = -2$. With this modification, the game is no longer a screening game, because $V_2(\theta_0) > u_2(a_1) = -2$. In this game, player 2 would prefer to play S against the commitment type, or against the normal type if he pools with the commitment type, and prefers to play W against the normal type if she expects player 1 to play a_0 for a sufficiently long period. In this game, again, a reputation result holds: For a fixed $\mu_0 > 0$, player 1's Nash equilibrium payoffs are bounded below by a payoff that converges to V_1 as the parties become arbitrarily patient. Moreover, a patient player 2 almost immediately plays S against both types. This result is far from trivial, and the proof is similar to the proof that μ^* is a reputation boundary in the sketch we provide. The intuition is similar to that in Cripps et al. (2005)

and Atakan and Ekmekci (2012), who study reputation effects in repeated games with strictly conflicting interests. Notice that, our example is akin to a strictly conflicting interest game: player 1 has an action, a_1, that if he committed to playing forever, then player 2 has a unique best response, which gives player 2 her minmax payoff, and player 1 his highest individually rational payoff. However, the dynamic game we consider is not a repeated game, which makes the analysis easier, allowing us to work with imperfect monitoring on the informed player's actions, which brings new challenges. We leave a more general analysis of this class of games for future work.

References

Acharya, A. and J. Ortner (2017): "Progressive Learning," *Econometrica*, 85, 1965–1990.

Aliprantis, C. D. and K. C. Border (2006): *Infinite Dimensional Analysis: A Hitchhiker's Guide*, Springer Science & Business Media.

Atakan, A. and M. Ekmekci (2012): "Reputation in Long-Run Relationships," *Review of Economic Studies*, 79, 451–480.

——— (2015): "Reputation in the Long-Run with Imperfect Monitoring," *Journal of Economic Theory*, 157, 553–605.

Aumann, R. J. and M. Maschler (1995): *Repeated games with incomplete information.*

Bhaskar, V. and G. J. Mailath (2019): "The curse of long horizons," *Journal of Mathematical Economics.*

Billingsley, P. (2013): *Convergence of probability measures*, John Wiley & Sons.

Cripps, M., E. Dekel, and W. Pesendorfer (2005): "Reputation with Equal Discounting in Repeated Games with Strictly Conflicting Interests," *Journal of Economic Theory*, 121, 259–272.

Cripps, M., G. Mailath, and L. Samuelson (2004): "Imperfect monitoring and impermanent reputations," *Econometrica*, 72, 407–432.

Cripps, M. W., G. J. Mailath, and L. Samuelson (2007): "Disappearing private reputations in long-run relationships," *Journal of Economic Theory*, 134, 287–316.

Deneckere, R. and M.-Y. Liang (2006): "Bargaining with interdependent values," *Econometrica*, 74, 1309–1364.

Fudenberg, D. and D. Levine (1992): "Maintaining a reputation when strategies are imperfectly observed," *The Review of Economic Studies*, 561–579.

Gerardi, D. and L. Maestri (2018): "Dynamic Contracting with Limited Commitment and the Ratchet Effect," Tech. rep., FGV/EPGE Working Paper Series.

Gul, F., H. Sonnenschein, and R. Wilson (1986): "Foundations of dynamic monopoly and the Coase conjecture," *Journal of Economic Theory*, 39, 155–190.

Hart, O. D. and J. Tirole (1988): "Contract renegotiation and Coasian dynamics," *The Review of Economic Studies*, 55, 509–540.

Laffont, J.-J. and J. Tirole (1987): "Comparative statics of the optimal dynamic incentive contract," *European Economic Review*, 31, 901–926.

——— (1988): "The dynamics of incentive contracts," *Econometrica*, 1153–1175.

——— (1993): *A theory of incentives in procurement and regulation*, MIT press.

Liu, Q., K. Mierendorff, X. Shi, and W. Zhong (Forthcoming): "Auctions with Limited Commitment," *American Economic Review*.

Malcomson, J. M. (2016): "Relational incentive contracts with persistent private information," *Econometrica*, 84, 317–346.

Pei, H. D. (2018): "Reputation Effects under Interdependent Values," *Available at SSRN 3064949*.

Sobel, J. (1985): "A theory of credibility," *The Review of Economic Studies*, 52, 557–573.

Strulovici, B. (2017): "Contract Negotiation and the Coase Conjecture: A Strategic Foundation for Renegotiation-Proof Contracts," *Econometrica*, 85, 585–616.

Watson, J. (1999): "Starting small and renegotiation," *Journal of economic Theory*, 85, 52–90.

Appendix

Section A presents a summary and reminder of the notation. Section B presents a learning result. Section C presents the concept of generalized centralized plays and provides results about the limits of sequences of generalized centralized plays. Section D presents the coarsening technique, which uses the martingale-splitting technique appropriately to reduce the dimensionality of centralized plays. Sections E and F provide the proofs of Theorems 1, 2, and 3. All proofs that are not presented in the Appendix are in the Supplementary Appendix.

A. Notation

For $i = 1, 2$, let $\bar{u}_i := \max_{a \in \mathbb{A}} |u_i(a)|$ and $\underline{u}_i := \min_{a \in \mathbb{A}} u_i(a)$.

For every $\varepsilon \in (0, 1)$ and $\delta \in (0, 1)$, let $t_\delta(\varepsilon)$ be the smallest positive integer t such that $\delta^t < 1 - \varepsilon$.

For each $m \in \mathbb{N}$ and δ, let $t_\delta(m)$ be the smallest positive integer t such that $\delta^t < 1 - m^{-1}$.

For every $i \in \{1, ..., K\}$, let $\mathcal{Y}_i := \{y \in \mathcal{Y} : P(y \mid \alpha_i) > 0\}$, and $\underline{P}_i := \min\{P(y_j \mid \alpha_i) : y \in \mathcal{Y}_i\}$.

Given two elements x, y of a finite dimensional Euclidean space, we define $d(x, y) := \|x - y\|$, where $\|.\|$ stands for the Euclidian norm.

Given a set $A \subseteq \mathbb{R}^{\mathrm{N}}$ and a vector $x \in \mathbb{R}^{\mathrm{N}}$, we define $d(x, A) := \inf_{\tilde{x} \in \mathrm{A}^\cdot} \|x - \tilde{x}\|$. For every subset A of $\mathbb{R}^{\mathrm{N}}$ and $\epsilon > 0$ let $A^\epsilon := \left\{x \in \mathbb{R}^{\mathrm{N}} : d(x, A) < \epsilon\right\}$.

Given $z_1, z_2 \in \Delta\mathcal{Y}$, let

$$||z_1, z_2|| := \max_{y \in \mathcal{Y}} |z_1(y) - z_2(y)|.$$

Consider the set $D = \{S\} \cup \triangle\mathbb{A}$. The metric d^* on D is defined as follows. Take two elements $x', x'' \in D$. We have $d^*(x', x'') = 0$ if $x' = x'' = S$, $d^*(x', x'') = 2$ if $(x', x'') \notin \triangle\mathbb{A} \times \triangle\mathbb{A}$, and finally $d^*(x', x'') = \|x' - x''\|$ otherwise.

Recall that $\boldsymbol{\mu_i}$ is the distribution over types that puts probability 1 on type i, $\mu_{\mathbf{i}}{}^b$ is the distribution that puts probability μ_i^* on type θ_i, and remaining probability on the normal type, θ_0. Moreover, recall that $\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}} := \{\mu \in \Delta\Theta : \mu(\theta_i) > 0, \mu(\theta_j) = 0 \text{ for all } j \neq \{0, i\}\}$.

For $\epsilon > 0$, let $\mu_{\mathbf{i}}{}^{b,\epsilon}$ denote the ϵ neighborhood (given the Euclidean metric) of $\mu_{\mathbf{i}}{}^b$. Likewise, let $\mu^b := \cup_{i \in \{1,2,\ldots,s\}} \mu_{\mathbf{i}}{}^b$. Also let $\boldsymbol{\mu_i}^\varepsilon$ be the ε-neighborhood of $\boldsymbol{\mu_i}$.

Let the set of beliefs at least ε-away from the edges be defined by:

$$\tilde{\boldsymbol{\mu}}_{\boldsymbol{\varepsilon}} := \{\mu \in \Delta\Theta : \exists (i, j) \in \{1, ..., K\}^2, i \neq j \text{ such that } \mu(\theta_i) > \varepsilon \text{ and } \mu(\theta_j) > \varepsilon\}.$$

For any finite set A we let $|A|$ represent its cardinality. For any set $A \subseteq X$, we write

A^c for its complement with respect to X. We write $\mathbb{B}(\mathbb{X})$ for the Borel sigma-field of the topological space $\mathbb{X}$. The term $\mu_t(\theta) \in \Delta\Theta$ means the belief held by player 2 at the beginning of period t. We write $h^t \cup \{y\}$ for a public history for period $t+1$ that follows h^t in which player 2 plays W and the public signal is y at the last period. We use $U_{1,\sigma}(h^t)$ and $U_1(h^t;\sigma)$ interchangeably to denote the continuation payoff of the normal type of player 1 at history h^t when strategy profile is σ. We use NE as an abbreviation for Nash equilibrium.

B. Learning Result

This section establishes a learning result which shows that the expected discounted number of periods in which different types of player 1 play different actions and player 2 remains uncertain about player 1's type goes to zero as δ goes to one. The proof of the lemma adapts techniques developed by Fudenberg and Levine (1992) and their proofs are relegated to the Supplementary Appendix, Section G. In the following, fix a sequence of dynamic games along which the discount factor converges to 1, and take a sequence of strategy profiles $\{\sigma_n\}$, where $\sigma_n = (\sigma_{1,n}, \sigma_{2,n})$ for each $n \in \mathbb{N}$.

Lemma 9.

1. *For every $\epsilon > 0$, $i > 0$, we have:*

$$\lim_{n\to\infty} \mathbb{E}^{\theta_0}_{(\sigma_n)} \left((1-\delta) \sum_{t=0}^{\mathbb{T}-1} \delta^t \mu_t(\theta_i) \mathbb{I}_{\{d(\sigma_{1,n}(h^t;\theta_0),\alpha_i)>\epsilon\}} \right) = 0.$$

2. *For any $i, j > 0$ and $i \neq j$, we have:*

$$\lim_{n\to\infty} \mathbb{E}^{\theta_i}_{(\sigma_n)} \left((1-\delta) \sum_{t=0}^{\mathbb{T}-1} \delta^t \mu_t(\theta_j) \right) = 0.$$

3. *For any $i > 0$, we have:*

$$\lim_{n\to\infty} \mathbb{E}^{\theta_i}_{(\sigma_n)} \left((1-\delta) \sum_{t=0}^{\mathbb{T}-1} \delta^t \mu_t(\theta_0) \mathbb{I}_{d(\sigma_{1,n}(h^t;\theta_0),\alpha_i)>\epsilon} \right) = 0.$$

Lemma 9 shows that identifying types is very fast when types are expected to play even slightly different strategies, when players are patient. Item (i) establishes that the normal type separates himself from a commitment type almost immediately if he does not almost completely mimic a commitment type. Item (ii) establishes that two commitment types are separated almost immediately . Item (iii) together with item (ii) establish that a commitment type is identified almost immediately if the normal type is not almost completely mimicking that type.

C. Generalized Centralized Play, Implementation and Incentive Compatibility

A generalized centralized play is a vector $\Gamma = (\Gamma(\theta_0), ..., \Gamma(\theta_K))$, where each element in the vector, $\Gamma(\theta)$, is a probability measure on $((\triangle\Theta \times (\{S\} \cup \triangle\mathbb{A})), \mathbb{B}(\triangle\Theta \times (\{S\} \cup \triangle\mathbb{A})))$. We say a generalized centralized play Γ implements a strategy profile $\sigma$of a dynamic game if for every $\theta \in \Theta$, $\mu \in \Delta\Theta$, the following holds.

$$\Gamma(\theta)(\{\mu\},\{S\}) = \mathbb{E}^{\theta}_{(\sigma)}\left(\delta^{\mathbb{T}}\mathbb{I}_{\{\mu(h^{\mathbb{T}})=\mu\}}\right),$$
$$\Gamma(\theta)(\{\mu\},\{\alpha\}) = \mathbb{E}^{\theta}_{(\sigma)}\left((1-\delta)\sum_{t=0}^{\mathbb{T}-1}\delta^t\mathbb{I}_{\{\mu(h^t)=\mu,\sigma_1(h^t;\theta)=\alpha\}}\right) \text{ for every } \alpha \in \triangle\mathbb{A}.$$

If a generalized centralized play implements a strategy profile σ, then each element in the vector, $\Gamma(\theta)$, has a marginal distribution on its first element, $\lambda(\theta) \in \Delta(\Delta\Theta)$, and a conditional distribution on the actions, $Y(\theta) : \Delta\Theta \to \Delta\{S, \mathbb{A}\}$, given by:

$$\lambda(\theta)(\mu) := \mathbb{E}^{\theta}_{(\sigma)}\left((1-\delta)\sum_{t=0}^{\mathbb{T}-1}\delta^t\mathbb{I}_{\{\mu(h^t)=\mu\}} + \delta^{\mathbb{T}}\mathbb{I}_{\{\mu(h^{\mathbb{T}})=\mu\}}\right) \text{ for every } \mu \in \Delta\Theta,$$

and for every $\mu \in \Delta\Theta$, with $\lambda(\theta)(\mu) > 0$,

$$Y(\theta)(\mu)(S) := \frac{\mathbb{E}^{\theta}_{(\sigma)}\left(\delta^{\mathbb{T}}\mathbb{I}_{\{\mu(h^{\mathbb{T}})=\mu\}}\right)}{\lambda(\theta)(\mu)},$$
$$Y(\theta)(\mu)(a) := \frac{\mathbb{E}^{\theta}_{(\sigma)}\left((1-\delta)\sum_{t=0}^{\mathbb{T}-1}\delta^t\mathbb{I}_{\{\mu(h^t)=\mu,a_t=a\}}\right)}{\lambda(\theta)(\mu)} \text{ for every } a \in \mathbb{A}.$$

Remark 3. Note that if Γ implements σ, and if $\lambda(\theta)(\mu) = 0$ for some $\theta \in \Theta$, $\mu \in \Delta\Theta$, then there is no restriction on $Y(\theta)(\mu)$.

When we refer to a generalized centralized play, sometimes we mean the vector of measures Γ, and other times we mean the vector of tuples $(\lambda(\theta), Y(\theta))_{\theta\in\Theta}$ of marginal distribution on beliefs, and conditional distributions over actions. We will be explicit when the context does not clarify the usage.

We say a generalized centralized play obeys the martingale property if:

$$\lambda(\theta_i)(\mu)\mu_0(\theta_i)\mu(\theta_j) = \lambda(\theta_j)(\mu)\mu_0(\theta_j)\mu(\theta_i) \text{ for every } i, j \in \{0, 1, ..., K\}. \tag{23}$$

Because the public belief process under any strategy profile is a martingale, if a generalized centralized play implements a strategy profile, then equation (23) holds.

For convenience, when it does not cause a confusion, we refer to $\lambda \in \Delta(\Delta\Theta)$ with $\lambda(\mu) := \sum_{\theta\in\Theta} \mu_0(\theta)\lambda(\theta)(\mu)$ for each $\mu \in \Delta\Theta$ as the distribution over posteriors. Observe that for a given $\mu \in \Delta\Theta$, and θ such that $\mu_0(\theta) > 0$, each $\lambda(\mu)$ implies a unique $\lambda(\theta)(\mu)$ obtained from equation (23). Moreover, if a generalized centralized play implements a strategy profile, then for every $\theta \in \Theta$,

$$\int_{\Delta\Theta} \lambda(\mu)\mu(\theta)d[\mu] = \mu_0(\theta).$$

The introduction of a generalized centralized play will be a useful tool to investigate both players' equilibrium payoff bounds. For any type θ, the probability that each action $a \in \mathbb{A}$ is chosen by the generalized centralized play at belief μ for type θ will correspond to the expected discounted number of periods in which player 2 plays W and player 1 plays action a at histories where the posterior belief about player 1's type equals to μ, conditional on type θ. Likewise, the probability that S is chosen by the generalized centralized play at belief μ for type θ will correspond to the expected discounted probability by which player 2 plays S at histories where the posterior belief about player 1's type equals to μ, conditional on type θ.

Finally, observe that any element of $Y(\theta)(\mu) \in \triangle\{\mathbb{A}, S\}$ can be equivalently expressed with a pair $(v(\theta)(\mu), Z(\theta)(\mu))$ where $v(\theta) : \Delta\Theta \to [0,1]$, and $Z(\theta) : \Delta\Theta \to \Delta\mathbb{A}$ with the equalities:

$$v(\theta)(\mu) := Y(\theta)(\mu)(S) \tag{24}$$

and whenever $v(\theta)(\mu) < 1$, for each $a \in \mathbb{A}$,

$$Z(\theta)(\mu)(a) = \frac{Y(\theta)(\mu)(a)}{1 - v(\theta)(\mu)}. \tag{25}$$

Hence, we can equivalently denote a generalized centralized play as $\Gamma = (\lambda, v, Z)$ where $v : \Theta \times \Delta\Theta \to [0,1]$ and $Z : \Theta \times \Delta\Theta \to \Delta\mathbb{A}$. We use these equivalent representations of a generalized centralized play throughout the Appendix. Observe that for every $i \in \{1, ..., K\}$, for almost every $\mu \in \Delta\Theta$ with respect to the measure $\lambda(\theta_i)$, if $v(\theta_i)(\mu) < 1$, then $Z(\theta_i)(\mu) = \alpha_i$.

The next lemma shows that if Γ implements a strategy profile, then $v(\theta_i)(\mu)$ and $v(\theta_j)(\mu)$ must agree for every $i, j \in \{0, 1, ..., K\}$. Hence, from now on, we drop the dependence of v on θ.

Lemma 10. *If a generalized centralized play $\Gamma = (\lambda, v, Z)$ implements some strategy profile σ, then, for almost every posterior $\mu \in \Delta\Theta$ with $\mu(\theta_i), \mu(\theta_j) > 0$, with respect to the measure λ we have $v(\theta_i)(\mu) = v(\theta_j)(\mu)$.*

Proof. For any $h^t \in H$, $\sigma_2(h^t)$ is independent of $\theta \in \Theta$. Therefore, $v(\theta)(\mu) = \frac{\mathbb{E}^{\theta}_{(\sigma)}\left(\mathbb{I}_{\{\mu(h^t)=\mu\}}\delta^{\mathbb{T}}\right)}{\lambda(\theta)(\mu)}$ is constant across all θ for which $\lambda(\theta)(\mu) > 0$. $\square$

Because in any NE, the normal type can always mimic a commitment type, a generalized centralized play Γ which implements a NE σ satisfies the following incentive compatibility constraints:

$$\begin{aligned}\int_{\Delta\Theta} \lambda(\theta_0)(\mu)\left((1-v(\mu))\, u_1\left(Z\left(\theta_0\right)(\mu)\right) + v(\mu)V_1\right) d\mu \geq \\ \int_{\Delta\Theta} \lambda(\theta_i)(\mu)\left((1-v(\mu))\, u_1\left(\alpha_i\right) + v(\mu)V_1\right) d\mu\end{aligned} \tag{26}$$

for every $i \in \{1, ..., K\}$.

C.1. Centralized Plays, Limit Centralized Plays For the following development, fix a sequence of strategy profiles $\{\sigma_n\}$ of a sequence of dynamic games where in the n^{th} game of the sequence, the discount factor is $\delta_n < 1$, and $\delta_n \to 1$. Let $\{\Gamma_n\}$ be the associated sequence of generalized centralized plays where each Γ_n implements σ_n in the n^{th} game of the sequence.

Since the set $(\triangle\Theta \times (\{S\} \cup \triangle\mathbb{A}))$ is compact, the family of probability measures on $\mathbb{B}\left(\triangle\Theta \times (\{S\} \cup \triangle\mathbb{A})\right)$ is relatively compact by Prohorov's Theorem (Billingsley (2013), Theorem 5.1). Thus, each sequence $\{\Gamma_n(\theta)\}$ of probability measures has a (weakly) convergent subsequence. Therefore, the sequence $\{\Gamma_n\}$ has a subsequence (weakly) converging to $K+1$ probability measures $\Gamma = (\Gamma(\theta_0), ..., \Gamma(\theta_K))$ with each element on $\mathbb{B}\left(\triangle\Theta \times (\{S\} \cup \triangle\mathbb{A})\right)$. We call any limit point of the sequence, $\Gamma = (\Gamma(\theta_0), ..., \Gamma(\theta_K))$, a limit centralized play.

Given a generalized centralized play, Γ, recall that $\lambda(\theta)$ is the marginal distribution over posteriors induced by $\Gamma(\theta)$. For each μ, let $\Gamma(\theta \mid \mu)$ be the conditional distribution of $\Gamma(\theta)$ over $\{S\} \cup \triangle\mathbb{A}$.

Definition 3. A **centralized play** Γ is a generalized centralized play that satisfies the following properties:

1. For almost every posterior $\mu \in \Delta\Theta$ with $\mu(\theta_i), \mu(\theta_j) > 0$, with respect to the measure

λ:

$$\Gamma(\theta_i \mid \mu) = \Gamma(\theta_j \mid \mu).$$

2. For almost every posterior $\mu \in \Delta\Theta \backslash \left(\cup_{i=1}^{K} \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\right) \bigcup \{\boldsymbol{\mu_0}\}$ with respect to the measure λ, $\Gamma(\theta_i \mid \mu)(S) = 1$.

3. For every $i \in \{1,..,K\}$, and for almost every posterior $\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}$ with respect to the measure λ, $\Gamma(\theta_i \mid \mu)(\{\alpha_i\}) + \Gamma(\theta_i \mid \mu)(\{S\}) = 1$.

A centralized play satisfies three properties. The first property is that the strategy profile associated with a belief μ is type independent. The second property is that for beliefs that are outside of the edges of the belief simplex that can put positive probability only on the normal type and at most one commitment type, i.e., $\mu \in \Delta\Theta \backslash \left(\cup_{i=1}^{K} \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\right) \bigcup \{\boldsymbol{\mu_0}\}$, the action is S with probability 1. The third property states that for beliefs on the edges of the belief simplex that puts positive probability on exactly one commitment type, either action S is taken, or commitment type's strategy α_i is played. As will be also clear from the proof of Theorem 6, the first property in the definition of a centralized play implies the second and third properties. We keep the definition as it is since it highlights these other properties of a centralized play that we will be using in what follows.

In the Supplementary Appendix, Section H, we prove Theorem 6, which shows that every limit centralized play is a centralized play. Unless otherwise stated, we work with centralized plays for the remaining of this Appendix.

Theorem 6. *Every limit centralized play is a centralized play.*

Theorem 6 shows that any limit point of a sequence a generalized centralized plays that implements a sequence of strategy profiles is a centralized play (hence it is a generalized centralized play). Clearly, the limit centralized play obeys the martingale property, equation (23). Recall the representation of a generalized centralized play in terms of its marginal distribution, and conditional distribution over actions, (λ, v, Z). A limit centralized play, then satisfies the following properties:

1. For almost every posterior $\mu \in \Delta\Theta$ with respect to the measure λ, v and Z are type independent.

2. For almost every posterior $\mu \in \Delta\Theta \backslash \left(\cup_{i=1}^{K} \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\right) \bigcup \{\boldsymbol{\mu_0}\}$ with respect to the measure λ, $v(\mu) = 1$.

3. For almost every posterior $\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}$ with respect to the measure λ, $v(\mu) < 1$ implies $Z(\mu) = \alpha_i$.

From hereon, in any centralized play we drop the dependence of v and Z on θ.

C.2. Properties Inherited by Centralized Plays To express the payoff functions of players 1 and 2 in centralized plays we define $U_1 : [0,1] \times \Delta\mathbb{A} \to \mathbb{R}$, and $U_2 : [0,1] \times \Delta\mathbb{A} \times \Theta \to \mathbb{R}$ as:

$$U_1 (v, Z) := (1 - v) u_1(Z) + vV_1$$
$$U_2 (v, Z, \theta) := (1 - v) u_2(Z) + vV_2 (\theta) .$$

Hence if the normal type follows the strategy of type θ_i in a centralized play $\Gamma = (\lambda, Z, v)$, he obtains

$$\int_{\triangle\Theta\times[0,1]\times\triangle\mathbb{A}} U_1 (v, Z) \Gamma (\theta_i) (d\mu, dv, dZ) = \int_{\triangle\Theta} U_1 (v (\mu) , Z (\mu)) \lambda (\theta_i) (\mu) d\mu$$

We can thus write the incentive-compatibility constraints as

$$\int_{\triangle\Theta} U_1 (v (\mu) , Z (\mu)) \lambda (\theta_0) (\mu) d\mu \geq \int_{\triangle\Theta} U_1 (v (\mu) , Z (\mu)) \lambda (\theta_i) (\mu) d\mu$$

for all $i \in \{1, 2, ..., K\}$.

Player 2's payoff in a centralized play $\Gamma = (\lambda, Z, v)$ when the prior is μ_0 can be expressed as:

$$\sum_{\theta\in\Theta} \mu_0 (\theta) \int_{\triangle\Theta\times[0,1]\times\triangle\mathbb{A}} U_2 (v, Z, \theta) \Gamma (\theta) (d\mu, dv, dZ) = \sum_{\theta\in\Theta} \mu_0 (\theta) \int_{\triangle\Theta} U_2 (v (\mu) , Z (\mu) , \theta) \lambda (\theta_i) (\mu) d\mu.$$

The next Theorem uses basic results on weak-convergence of probability measures to establish that if the sequence of strategy profiles $\{\sigma_n\}$ satisfies some particular property, then the limit centralized play $\Gamma = (\lambda, Z, v)$ satisfies a certain limit version of that property. The proof of Theorem 7 is in the Supplementary Appendix, Section H.

Theorem 7. *Take a sequence of strategy profiles $\{\sigma_n\}$ associated with a sequence games where the discount factors converge to 1. Let $\{\Gamma_n\}$ be the associated sequence of generalized centralized plays where Γ_n implements σ_n. Every limit point of $\{\Gamma_n\}$, Γ, is a centralized play and satisfies the following properties.*

1. *(**Γ preserves the IC constraints**) Suppose that each σ_n is a NE of the n^{th} game in the sequence, and hence is incentive compatible. Then the limit centralized play $\Gamma = (\lambda, Z, v)$ is incentive-compatible: F*or every $i \in \{1, 2, ..., K\}$

$$\int_{\triangle\Theta} U_1 (v (\mu) , Z (\mu)) \lambda (\theta_0) (\mu) d\mu \geq \int_{\triangle\Theta} U_1 (v (\mu) , Z (\mu)) \lambda (\theta_i) (\mu) d\mu.$$

2. ***(S happens with probability zero in a subset of beliefs)*** *Take an open subset $B \in \mathbb{B}(\triangle\Theta)$. Suppose $\{\sigma_n\}$ satisfies the "never stop in B" property for some type $\theta \in \Theta$, i.e., for every $\varepsilon > 0$ there is $\delta^* \in (0,1)$ such that for all $\delta_n > \delta^*$, for every on-path history h^t of σ_n satisfying $\mu(h^t) \in B$, the continuation play starting at h^t satisfies $\mathbb{E}^{\theta}{}_{(\sigma_n)}\left(\delta_n^{\mathbb{T}-t}|h^t\right) < \varepsilon$. Then we have*

$$\int_B v(\mu)\,\lambda(\theta)(\mu)\,d\mu = 0.$$

3. ***(S happens with probability one in a subset of beliefs)*** *Take an open subset $B \in \mathbb{B}(\triangle\Theta)$. Suppose $\{\sigma_n\}$ satisfies the "stop in B" property for some type $\theta \in \Theta$, i.e., for every $\varepsilon > 0$ there is $\delta^* \in (0,1)$ such that for all $\delta_n > \delta^*$, for every on-path history h^t of σ_n satisfying $\mu(h^t) \in B$, the continuation play starting at h^t satisfies $\mathbb{E}^{\theta}{}_{(\sigma_n)}\left(\delta_n^{\mathbb{T}-t}|h^t\right) > 1-\varepsilon$. Then we have*

$$\int_B (1-v(\mu))\,\lambda(\theta)(\mu)\,d\mu = 0.$$

D. Coarsening Procedure

Fix a centralized play $\Gamma = (\lambda, Y)$, and a real number $\bar{\kappa} > 0$. We say that a centralized play $\tilde{\Gamma}$ is a $\bar{\kappa}$-coarsening of Γ if $\tilde{\Gamma}$ is constructed from Γ according to the procedure we describe in this section.

For every $\mu \in \Delta\Theta$ with $\mu(\theta_0) > 0$, let $\kappa(\mu) \in \mathbb{R}_+$ be the unique number such that

$$\sum_{i\in\{1,2,\ldots,s\}} \mu(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} = \kappa\mu(\theta_0). \tag{27}$$

If $\mu(\theta_0) = 0$, let $\kappa(\mu) := \infty$.

For $i \in \{1,.2,...,s\}$, and for any $\kappa \in [0,\infty) \cup \{\infty\}$, let $\tilde{\boldsymbol{\mu}}_{\mathbf{i}}^{\boldsymbol{\kappa}} \in \Delta\Theta$ be the distribution that puts probability μ_i^κ on θ_i and $1-\mu_i^\kappa$ on θ_0, where μ_i^κ is the unique solution to

$$\mu_i^\kappa \frac{1-\mu_i^*}{\mu_i^*} = \kappa(1-\mu_i^\kappa),$$

with the convention that $\mu_i^\infty = 1$. Observe that $\mu_i^1 = \mu_i^*$, and μ_i^κ is strictly increasing in κ.

Lemma 11. *Fix $\mu \in \Delta\Theta$.*

1. *If $\kappa(\mu) > 0$, then there are unique numbers $\{\beta_i(\mu)\}_{i=0,1,...,K}$ with $\beta_i(\mu) \geq 0$ for $i = 0, 1, ..., K$, and $\sum_{i=0,1,..,K} \beta_i(\mu) = 1$ such that*

$$\mu = \sum_{1=1,2,...,s} \beta_i(\mu)\tilde{\boldsymbol{\mu}}_{\mathbf{i}}^{\boldsymbol{\kappa}(\boldsymbol{\mu})} + \sum_{i=0,s+1,...,K} \beta_i(\mu)\boldsymbol{\mu_i}.$$

Moreover, $\beta_0(\mu) = 0$.

2. *If $\kappa(\mu) = 0$, then there are unique numbers $\{\beta_i(\mu)\}_{i=0,1,...,K}$ with $\beta_i(\mu) \geq 0$ for $i = 0, 1, ..., K$, $\beta_i(\mu) = 0$ for $i = 1, .., s$, and $\sum_{i=0,..,K} \beta_i(\mu) = 1$ such that*

$$\mu = \sum_{i=1,2,...,s} \beta_i(\mu)\tilde{\boldsymbol{\mu}}_{\mathbf{i}}^{\boldsymbol{\kappa}(\boldsymbol{\mu})} + \sum_{i=0,s+1,...,K} \beta_i(\mu)\boldsymbol{\mu_i}.$$

Hence, $\beta_i(\mu) = \mu(\theta_i)$ for $i = 0, s + 1, ..., K$.

The proof of this lemma is in the Supplementary Appendix, Section I. Lemma 11 shows that each belief $\mu \in \Delta\Theta$ with $\kappa(\mu) > 0$ can be written uniquely as a convex combination of beliefs $\left\{\tilde{\boldsymbol{\mu}}_{\mathbf{i}}^{\boldsymbol{\kappa}(\boldsymbol{\mu})}\right\}_{i\in\{1,2,...,s\}} \cup \{\boldsymbol{\mu_i}\}_{i\in\{s+1,...,K\}}$. When $\kappa(\mu) = 0$, $\boldsymbol{\mu_0} = \tilde{\boldsymbol{\mu}}_{\mathbf{i}}^{\boldsymbol{\kappa}(\boldsymbol{\mu})}$ for every $i \in \{1, 2, ..., s\}$, so a multiplicity occurs in the decomposition. In this case, μ can be written uniquely as a convex combination of beliefs $\{\boldsymbol{\mu_i}\}_{i\in\{0,s+1,...,K\}}$.

For each $i = 1, ..., s$, let $m_i := \frac{\mu_i^*}{1-\mu_i^*}$.

For each $\kappa \geq 0$, let

$$x_i(\kappa) := \max\left\{\frac{(\kappa - \bar{\kappa})\, m_i}{\kappa m_i + 1}, 0\right\},$$

and let

$$y_i(\kappa) := \min\left\{\frac{\bar{\kappa} m_i + 1}{\kappa m_i + 1}, \frac{\kappa}{\bar{\kappa}}\frac{\bar{\kappa} m_i + 1}{\kappa m_i + 1}\right\}.$$

The terms $x_i(\kappa)$, $y_i(\kappa)$ and $1 - y_i(\kappa)$ will be weights we will use to further decompose each belief $\tilde{\boldsymbol{\mu}}_{\mathbf{i}}^{\boldsymbol{\kappa}(\boldsymbol{\mu})}$ (recall these beliefs are on the edge of the simplex of beliefs connecting the normal type and a commitment type θ_i for $i \in \{1, ..., s\}$) to the set of beliefs $\{\boldsymbol{\mu_0}, \boldsymbol{\mu_i}, \tilde{\boldsymbol{\mu}}_{\mathbf{i}}^{\bar{\boldsymbol{\kappa}}}\}$, i.e., the set of beliefs that either puts probability 1 to the normal type (θ_0), or puts probability 1 to type θ_i, or puts probability $\mu_i^{\bar{\kappa}}$ to type θ_i and probability $1 - \mu_i^{\bar{\kappa}}$ to type θ_0. More precisely, If $\kappa(\mu) \geq \bar{\kappa}$, then we split $\tilde{\boldsymbol{\mu}}_{\mathbf{i}}^{\boldsymbol{\kappa}(\boldsymbol{\mu})}$ to $\tilde{\boldsymbol{\mu}}_{\mathbf{i}}^{\bar{\boldsymbol{\kappa}}}$ and $\boldsymbol{\mu_i}$ with weights $y_i(\kappa(\mu)) = \frac{\bar{\kappa} m_i+1}{\kappa(\mu) m_i+1}$ and $x_i(\kappa(\mu)) = \frac{(\kappa(\mu)-\bar{\kappa})m_i}{\kappa(\mu)m_i+1}$, respectively. If $\kappa(\mu) < \bar{k}$, then $\frac{(\kappa(\mu)-\bar{\kappa})m_i}{\kappa(\mu)m_i+1} < 0$, hence $x_i(\kappa(\mu)) = 0$. In this case, we split the belief $\tilde{\boldsymbol{\mu}}_{\mathbf{i}}^{\boldsymbol{\kappa}(\boldsymbol{\mu})}$ to $\tilde{\boldsymbol{\mu}}_{\mathbf{i}}^{\bar{\boldsymbol{\kappa}}}$ and $\boldsymbol{\mu_0}$ with weights $y_i(\kappa(\mu)) = \frac{\kappa(\mu)}{\bar{\kappa}}\frac{\bar{\kappa} m_i+1}{\kappa(\mu)m_i+1}$ and $1 - y_i(\kappa(\mu)) = 1 - \frac{\kappa(\mu)}{\bar{\kappa}}\frac{\bar{\kappa} m_i+1}{\kappa(\mu)m_i+1}$, respectively. Importantly, the following equality, which is easy to verify, allows us to do the martingale-splitting with the specified weights:

$$\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa(\mu)}} = x_i(\kappa(\mu))\boldsymbol{\mu_i} + y_i(\kappa(\mu))\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\bar{\kappa}}} + (1 - x_i(\kappa(\mu)) - y_i(\kappa(\mu)))\,\boldsymbol{\mu_0}.$$

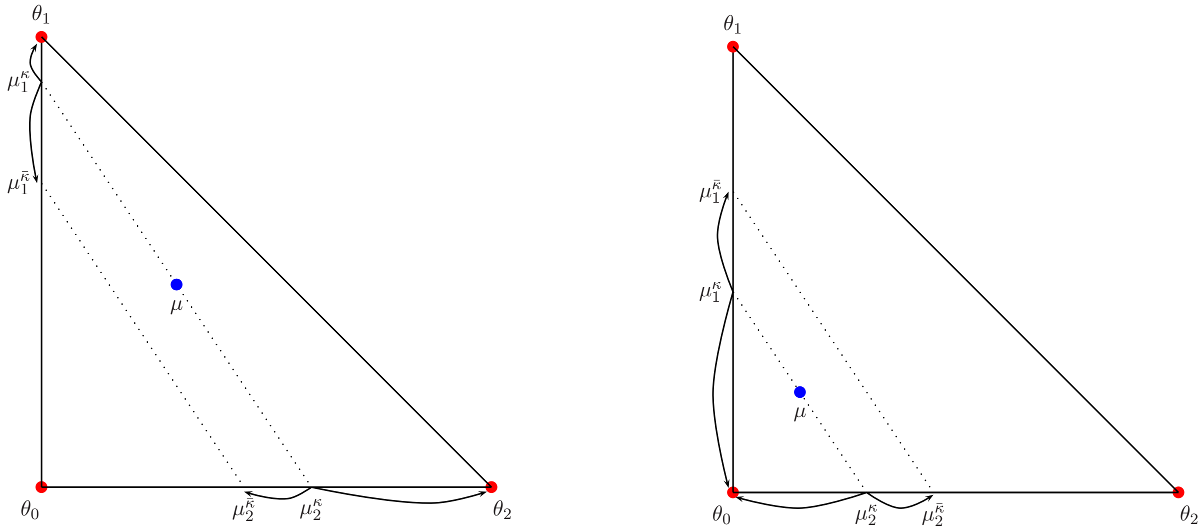

(a) $\kappa(\mu) > \bar{\kappa}$

(b) $\kappa(\mu) < \bar{\kappa}$

Figure 4: This panel illustrates the decomposition of each belief μ in the $\bar{\kappa}$-coarsening procedure.

The figures in panel 4 illustrate the decomposition of beliefs in an example where $\Theta = \{\theta_0, \theta_1, \theta_2\}$, and $s = 2$. For any $\mu \in \Delta\Theta$, we first determine $\kappa(\mu)$. In Figure 4a, $\kappa(\mu) > \bar{\kappa}$, and in figure 4b, $\kappa(\mu) < \bar{\kappa}$. The belief μ is decomposed to the set of beliefs $\left\{\tilde{\boldsymbol{\mu}}_{\mathbf{1}}^{\boldsymbol{\kappa(\mu)}}, \tilde{\boldsymbol{\mu}}_{\mathbf{2}}^{\boldsymbol{\kappa(\mu)}}\right\}$, where $\tilde{\boldsymbol{\mu}}_{\mathbf{1}}^{\boldsymbol{\kappa(\mu)}} = \left(1 - \mu_1^{\kappa(\mu)}, \mu_1^{\kappa(\mu)}, 0\right)$, and $\tilde{\boldsymbol{\mu}}_{\mathbf{2}}^{\boldsymbol{\kappa(\mu)}} = \left(1 - \mu_2^{\kappa(\mu)}, 0, \mu_2^{\kappa(\mu)}\right)$. This is the decomposition in Lemma 11. Then, we further decompose $\tilde{\boldsymbol{\mu}}_{\mathbf{i}}^{\boldsymbol{\kappa(\mu)}}$ to te set of beliefs $\{\boldsymbol{\mu_i}, \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\bar{\kappa}}}\}$ when $\kappa(\mu) > \bar{\kappa}$, with weights $x_i(\kappa(\mu))$ and $y_i(\kappa(\mu))$, respectively. When $\kappa(\mu) < \bar{\kappa}$, we decompose $\tilde{\boldsymbol{\mu}}_{\mathbf{i}}^{\boldsymbol{\kappa(\mu)}}$ to te set of beliefs $\{\boldsymbol{\mu_0}, \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\bar{\kappa}}}\}$, with weights $1 - y_i(\kappa(\mu))$ and $y_i(\kappa(\mu))$, respectively.

Now, we construct the $\bar{\kappa}$-coarsening of $\Gamma = (\lambda, Y)$, $\tilde{\Gamma} = (\tilde{\lambda}, \tilde{Y})$ as follows:

$$\begin{aligned}
\tilde{\lambda}\left(\boldsymbol{\mu_i}\right) &= \int_{\Delta\Theta} \beta_i(\mu) x_i(\kappa(\mu))\lambda[d\mu] \text{ for } i = 1, .., s. \\
\tilde{\lambda}\left(\boldsymbol{\mu_i}\right) &= \int_{\Delta\Theta} \beta_i(\mu)\lambda[d\mu] \text{ for } i > s. \\
\tilde{\lambda}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\bar{\boldsymbol{\kappa}}}\right) &= \int_{\Delta\Theta} \beta_i(\mu) y_i(\kappa(\mu))\lambda[d\mu] \text{ for } i = 1, .., s. \\
\tilde{\lambda}\left(\boldsymbol{\mu_0}\right) &= 1 - \sum_{i>0} \tilde{\lambda}\left(\boldsymbol{\mu_i}\right) - \sum_{i=1,..,s} \tilde{\lambda}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\bar{\boldsymbol{\kappa}}}\right) \\
&= \int_{\Delta\Theta} \left(\beta_0(\mu) + \sum_{i=1,2...,s} \beta_i(\mu)\left(1 - x_i(\kappa(\mu)) - y_i(\kappa(\mu))\right)\right) \lambda[d\mu].
\end{aligned}$$

Once we have $\tilde{\lambda}(\mu)$, we obtain $\tilde{\lambda}(\theta)(\mu)$ for each θ using Bayes' rule, equation (23). For each $a \in \mathbb{A} \cup \{S\}$:

$$\begin{aligned}
\tilde{Y}\left(\boldsymbol{\mu_i}\right)(a) &= \frac{\int_{\Delta\Theta} \beta_i(\mu) x_i(\kappa(\mu)) Y(\mu)(a)\lambda[d\mu]}{\tilde{\lambda}\left(\boldsymbol{\mu_i}\right)} \text{ if } \tilde{\lambda}\left(\mu_i\right) > 0, \text{ for } i = 1, .., s. \\
\tilde{Y}\left(\boldsymbol{\mu_i}\right)(a) &= \frac{\int_{\Delta\Theta} \beta_i(\mu) Y(\mu)(a)\lambda[d\mu]}{\tilde{\lambda}\left(\boldsymbol{\mu_i}\right)} \text{ if } \tilde{\lambda}\left(\mu_i\right) > 0, \text{ for } i > s. \\
\tilde{Y}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\bar{\boldsymbol{\kappa}}}\right)(a) &= \frac{\int_{\Delta\Theta} \beta_i(\mu) y_i(\kappa(\mu)) Y(\mu)(a)\lambda[d\mu]}{\tilde{\lambda}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\bar{\boldsymbol{\kappa}}}\right)} \text{ if } \tilde{\lambda}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\bar{\boldsymbol{\kappa}}}\right) > 0, \text{ for } i = 1, .., s \\
\tilde{Y}\left(\boldsymbol{\mu_0}\right)(a) &= \frac{\int_{\Delta\Theta} \left(\beta_0(\mu) + \sum_{i=1,2...,s} \beta_i(\mu)\left(1 - x_i(\kappa(\mu)) - y_i(\kappa(\mu))\right)\right) Y(\mu)(a)\lambda[d\mu]}{\tilde{\lambda}\left(\boldsymbol{\mu_0}\right)} \text{ if } \tilde{\lambda}\left(\boldsymbol{\mu_0}\right) > 0.
\end{aligned}$$

Finally, for each $i \in \{0, 1, ..., K\}$ and $\mu \in \Delta\Theta$, we set $\tilde{Y}(\theta_i)(\mu) = \tilde{Y}(\mu)$. This completes the description of the construction of $\bar{\kappa}$-coarsening of a centralized play Γ. The term $\left(\tilde{v}(\mu), \tilde{Z}(\mu)\right)$ denotes distribution over action profiles conditional on belief μ in the centralized play $\tilde{\Gamma}$.

Lemma 12. *Suppose $\tilde{\Gamma}$ is a $\bar{\kappa}$-coarsening of a centralized play Γ. Then:*

1. $\int_{\Delta\Theta} \mu\lambda(\mu)[d\mu] = \int_{\Delta\Theta} \mu\tilde{\lambda}(\mu)[d\mu]$.
2. *For each $a \in \mathbb{A} \cup \{S\}$ and each $\theta \in \Theta$:*

$$\int_{\Delta\Theta} Y(\mu)(a)\lambda(\theta)[d\mu] = \int_{\Delta\Theta} \tilde{Y}(\mu)(a)\tilde{\lambda}(\theta)[d\mu].$$

3. *$\tilde{\Gamma}$ is a centralized play.*

4. *If Γ is incentive compatible, then $\tilde{\Gamma}$ is also incentive compatible.*

The proof of Lemma 12 is in the Supplementary Appendix, Section I. It asserts that coarsening operation maps a centralized play to another centralized play, and preserves the action profiles type by type. It also asserts that coarsening preserves incentive compatibility property of Γ.

E. Screening Games

Recall that in screening models, for each $i \in \{1, 2, ..., s\}$, there exists a unique number $\mu_i^* \in (0,1)$ such that $u_2(\alpha_i) = \mu_i^* V_2(\theta_i) + (1 - \mu_i^*) V_2(\theta_0)$. We say that the prior belief $\mu_0 \in \Delta\Theta$ is a high reputation prior if

$$\sum_{i \in \{1,2,...,s\}} \mu_0(\theta_i) \frac{1 - \mu_i^*}{\mu_i^*} > \mu_0(\theta_0).$$

We say that the prior belief μ_0 is a low reputation prior if

$$\sum_{i \in \{1,2,...,s\}} \mu_0(\theta_i) \frac{1 - \mu_i^*}{\mu_i^*} < \mu_0(\theta_0).$$

E.1. High Reputation Priors: We start by providing a lower bound on player 2's limit equilibrium payoffs along all sequences of NE $\{\sigma_n\}$ of a sequence of games where along the sequence all parameters of the game except for the discount factor are fixed, and where in the n^{th} game the discount factor is δ_n, with $\delta_n \to 1$.

Lemma 13. *There exists a sequence of strategies for player 2 that guarantees player 2 a payoff that converges along the sequence of games to*

$$v_2 := \sum_{i \in \{0,1,...,s\}} \mu_0(\theta_i) V_2(\theta_i) + \sum_{i \in \{s+1,...,K\}} \mu_0(\theta_i) u_2(\alpha_i).$$

Proof. For every (small) $\varepsilon \in (0,1)$, consider the following strategy $\sigma_{2,n}(\varepsilon)$ for player 2: she chooses W at every period $t < t_{\delta_n}(\varepsilon)$. At history $h^{t_{\delta_n}(\varepsilon)}$, if $\sum_{i>s} \mu_t(h^{t_{\delta_n}(\varepsilon)})(\theta_i) < \varepsilon$, then player 2 plays S right away. Otherwise she plays W at every period in the rest of the game. Let $\{\sigma_n(\varepsilon)\}$ represent a strategy profile sequence in which player 2's strategy coincides with $\sigma_{2,n}(\varepsilon)$. For each n, let Γ_n be the generalized centralized play that implements $\sigma_n(\varepsilon)$. Taking a subsequence if necessary, Γ_n converges weakly to a centralized play Γ (Theorem 6). Since payoffs are continuous, player 2's payoffs converge to his payoffs from the centralized play Γ. Next notice in the first $t_{\delta_n}(\varepsilon)$ periods, player 2's flow payoffs are at least $\underline{u}_2$. Therefore, using property 2 of centralized plays (For all $\mu \in \Delta\Theta \backslash \left(\cup_{i=1}^{K} \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\right) \bigcup \{\boldsymbol{\mu_0}\}$ we have $v(\mu)=1$), and the learning lemma 9, we obtain that conditional on player 1 being type θ_i, player 2's payoff converges to a number at least as large as $\varepsilon \underline{u}_2 + (1-\varepsilon) V_2(\theta_i)$ if $i \in \{1, 2, ..., s\}$, and to $\varepsilon \underline{u}_2 + (1-\varepsilon) u_2(\alpha_i)$ for every $i > s$. Moreover, against the normal type, player 2's limit payoff is at least $\varepsilon \underline{u}_2 + (1-\varepsilon) V_2(\theta_0)$. Hence a lower bound on player 2's payoff from this strategy converges to

$$\varepsilon \underline{u}_2 + (1-\varepsilon)\left[\sum_{i\in\{0,1,\ldots,s\}} \mu_0(\theta_i) V_2(\theta_i) + \sum_{i\in\{s+1,\ldots,K\}} \mu_0(\theta_i) u_2(\alpha_i)\right].$$

Since ε is arbitrary, we obtain the result. $\square$

E.1.1. *Reputation Boundary:* We now introduce the concept of reputation boundary. Recall the definition of $\kappa(\mu)$ given by equation (27).

Definition 4. A number $\kappa \geq 0$ is a reputation boundary if for every $i \in \{0, 1, 2, ..., s\}$, for every $\hat{\kappa} > \kappa$,

$$\liminf_{\delta \nearrow 1} \left\{ \begin{array}{c} \mathbb{E}^{\theta_i}{}_{(\sigma)}\left[\delta^{\mathbb{T}}\right] : \sigma \text{ is a NE for} \\ \text{the game with prior } \tilde{\mu}_0 \text{ with } \kappa(\tilde{\mu}_0) \geq \hat{\kappa} \text{ and discount factor } \delta \end{array} \right\} = 1. \quad (28)$$

We will show the existence of such a boundary in the following Lemma.

Lemma 14. *There exists a reputation boundary.*

Proof. Fix a prior μ_0. We will show that there is some $\bar{\kappa} > 0$ such that for every $\kappa > \bar{\kappa}$, if $\sum_{i\in\{1,2,\ldots,s\}} \mu_0(\theta_i) \frac{1-\mu_i^*}{\mu_i^*} > \kappa \mu_0(\theta_0)$, then as $\delta \to 1$, across all NE, player 2 takes action S almost immediately against all types θ_i for $i \in \{0, 1, 2, ..., s\}$. Consider the mechanism-design problem in which player 1 first reports his type, and the mechanism chooses a number $Y(\theta_i) \in [0,1]$ for each type $i > 0$ that corresponds to the expected discounted number of periods in which player 2 plays W, and $\{Y_a(\theta_0)\}_{a\in\mathbb{A}}$ where $Y_a(\theta_0) \in [0,1]$ corresponds to the expected discounted number of periods in which normal type of player 1 plays action a and player 2 plays W, with $\sum_{a\in A} Y_a(\theta_0) \leq 1$. Consider the problem of maximizing player 2's payoffs via such a mechanism subject to truth telling constraints for the normal type.

$$\max \bar{U}_2 := \sum_{i>0} \mu_0(\theta_i) \left(Y(\theta_i)(u_2(\alpha_i) - V_2(\theta_i)) + V_2(\theta_i)\right)$$
$$+ \sum_{a\in\mathbb{A}} \mu_0(\theta_0) \left(Y_a(\theta_0)(u_2(a) - V_2(\theta_0)) + V_2(\theta_0)\right)$$

subject to

$$\sum_{a\in\mathbb{A}} (Y_a(\theta_0)(u_1(a) - V_1) + V_1) \geq Y(\theta_i)(u_1(\alpha_i) - V_1) + V_1 \text{ for every } i \in \{1, 2, ..., K\},$$
$$\sum_{a\in\mathbb{A}} (Y_a(\theta_0)(u_1(a) - V_1) + V_1) \geq u_1(a_0).$$

In the optimal mechanism, let $\bar{u} := V_1 + \sum_{a\in\mathbb{A}} Y_a(\theta_0)(u_1(a) - V_1)$. Clearly, if $\bar{u} = V_1$, then the optimal mechanism is unique and $Y(\theta_i) = 0$ for $i \in \{1, 2, ..., s\}$ and $Y(\theta_i) = 1$ for $i > s$.

Suppose towards a contradiction that $\bar{u} < V_1$. Because $\bar{u} \geq u_1(a_0)$ in the optimal mechanism, $Y(\theta_i) = 1$ for every $i > s$, since this maximizes player 2's payoff against such types, and does not violate IC constraints. Moreover, all other IC constraints for $i \in \{1, ..., s\}$ hold with equality. This is because if the IC constraint for type θ_i is slack then the objective function would be improved by a small decrease in $Y(\theta_i)$. Therefore,

$$Y(\theta_i) = \frac{V_1 - \bar{u}}{V_1 - u_1(\alpha_i)} > 0 \text{ for every } i \in \{1, 2, ..., s\}$$

Rewriting the objective function by plugging in $Y(\theta_i)$ for $i > 0$, we have

$$\bar{U}_2 = \sum_{i\leq s} \mu_0(\theta_i)V_2(\theta_i) + \sum_{i>s} \mu_0(\theta_i)u_2(\alpha_i) + \sum_{i\in\{1,2,...,s\}} \mu_0(\theta_i)\frac{V_1 - \bar{u}}{V_1 - u_1(\alpha_i)}(u_2(\alpha_i) - V_2(\theta_i))$$
$$+ \mu_0(\theta_0) \sum_{a\in\mathbb{A}} Y_a(\theta_0) \left(u_2(a) - V_2(\theta_0)\right)$$

Because $\bar{u} < V_1$, $\sum_{a\in\mathbb{A}} Y_a(\theta_0) > 0$. Let

$$\gamma_1 := \min_{a\in\mathbb{A}} \frac{V_1 - u_1(a)}{u_2(a) - V_2(\theta_0)}, \;\; \gamma_2 := \min_{i\in\{1,...,s\}} \frac{V_2(\theta_i) - u_2(\alpha_i)}{V_1 - u_1(\alpha_i)}.$$

Note that $\gamma_1, \gamma_2 > 0$. Because $(u_2(\alpha_i) - V_2(\theta_i)) < 0$ for every $i \in \{1, ..., s\}$, and because

$V_1 - \bar{u} = \sum_{a \in \mathbb{A}} Y_a(\theta_0)(V_1 - u_1(a))$, we have that

$$\sum_{i \in \{1,2,\ldots,s\}} \mu_0(\theta_i) \frac{V_1 - \bar{u}}{V_1 - u_1(\alpha_i)} (u_2(\alpha_i) - V_2(\theta_i)) \leq -\gamma_1 \gamma_2 \left(\sum_{a \in \mathbb{A}} Y_a(\theta_0) \left(u_2(a) - V_2(\theta_0)\right) \right) \sum_{i \in \{1,2,\ldots,s\}} \mu_0(\theta_i).$$

Moreover, because $\sum_{i \in \{1,2,\ldots,s\}} \mu_0(\theta_i) \frac{1-\mu_i^*}{\mu_i^*} > \kappa \mu_0(\theta_0)$, we have

$$\max_{i \in \{1,2,\ldots,s\}} \mu_0(\theta_i) \frac{1-\mu_i^*}{\mu_i^*} \geq \frac{1}{s} \kappa \mu_0(\theta_0).$$

It then follows that for some $j \in \{1, \ldots, s\}$,

$$\sum_{i \in \{1,2,\ldots,s\}} \mu_0(\theta_i) \geq \frac{\kappa}{s} \frac{\mu_j^*}{1-\mu_j^*} \mu_0(\theta_0).$$

Then,

$$\sum_{i \in \{1,2,\ldots,s\}} \mu_0(\theta_i) \frac{V_1 - \bar{u}}{V_1 - u_1(\alpha_i)} (u_2(\alpha_i) - V_2(\theta_i)) + \mu_0(\theta_0) \sum_{a \in \mathbb{A}} Y_a(\theta_0) \left(u_2(a) - V_2(\theta_0)\right) \leq \left(\mu_0(\theta_0) \sum_{a \in \mathbb{A}} Y_a(\theta_0) \left(u_2(a) - V_2(\theta_0)\right) \right) \left(-\gamma_1 \gamma_2 \frac{\kappa}{s} \frac{\mu_j^*}{1-\mu_j^*} + 1 \right).$$

Then, there exists a $\bar{\kappa} < \infty$ such that for all $i \in \{1, .., s\}$, $\kappa > \bar{\kappa}$ implies $-\gamma_1 \gamma_2 \frac{\kappa}{s} \frac{\mu_j^*}{1-\mu_j^*} + 1 < 0$. When $\mu_0(\theta_0) > 0$, because $\sum_{a \in \mathbb{A}} Y_a(\theta_0) > 0$, and because $(u_2(a) - V_2(\theta_0)) > 0$ for every $a \in \mathbb{A}$, we obtain that when $\kappa > \bar{\kappa}$, and $\mu_0(\theta_0) > 0$,

$$\sum_{i \in \{1,2,\ldots,s\}} \mu_0(\theta_i) \frac{V_1 - \bar{u}}{V_1 - u_1(\alpha_i)} (u_2(\alpha_i) - V_2(\theta_i)) + \mu_0(\theta_0) \sum_{a \in \mathbb{A}} Y_a(\theta_0) \left(u_2(a) - V_2(\theta_0)\right) < 0.$$

When $\mu_0(\theta_0) = 0$, because at least one of $\mu_0(\theta_i) > 0$, we again have

$$\sum_{i \in \{1,2,\ldots,s\}} \mu_0(\theta_i) \frac{V_1 - \bar{u}}{V_1 - u_1(\alpha_i)} (u_2(\alpha_i) - V_2(\theta_i)) + \mu_0(\theta_0) \sum_{a \in \mathbb{A}} Y_a(\theta_0) \left(u_2(a) - V_2(\theta_0)\right) < 0.$$

Hence, the payoff of player 2 is uniquely maximized across all such mechanisms by setting $\bar{u} = V_1$ and $Y(\theta_i) = 0$ for every $i \in \{1, 2, \ldots, s\}$ and $Y(\theta_i) = 1$ for every $i > s$ when $\kappa > \bar{\kappa}$.

And in this case, her payoff is

$$\sum_{i\in\{0,1,...,s\}} \mu_0(\theta_i)V_2(\theta_i) + \sum_{i\in\{s+1,...,K\}} \mu_0(\theta_i)u_2(\alpha_i),$$

and her payoff is strictly smaller if $Y(\theta_i) > 0$ for some $i \in \{1, 2, ..., s\}$. Therefore, if

$$\liminf_{\delta\to 1} \mathbb{E}^{\theta_i}{}_{(\sigma)} \left[\delta^{\mathbb{T}}\right] < 1$$

for some $i \in \{1, 2, ..., s\}$, then player 2's equilibrium payoffs along a subsequence would converge to a number strictly smaller than v_2, which contradicts Lemma 13. □

E.2. Showing $\kappa = 1$ is a reputation boundary We now tackle with the problem of extending reputation boundaries arbitrarily close to 1, i.e., showing that $\kappa = 1$ is a reputation boundary. Assume towards a contradiction that the infimum of all reputation boundaries is some $\kappa > 1$ and observe that our definition above implies that κ is itself a reputation boundary. We will find $\varepsilon \in (0, \kappa - 1)$ such that $\kappa - \varepsilon$ is a reputation boundary, which contradicts the hypothesis that κ is the infimum of all reputation boundaries. For some $\varepsilon > 0$ which will be determined later, let $\mu^B := \{\mu \in \Delta\Theta : \kappa(\mu) \in [\kappa - \varepsilon, \kappa]\}$ and let $\{\sigma_n\}$ be a sequence of NE of a sequence of games where in the n^{th} game, the prior is $\mu_{n,0} \in \mu^B$ with $\mu_0 = \lim \mu_{n,0} \in \mu^B$, and the discount factor is δ_n with $\delta_n \to 1$.

Suppose $\kappa > 1$ is a reputation boundary. Let Γ be a limit centralized play of the sequence of generalized centralized plays that implements the sequence of NE. Recall the coarsening procedure in section D. Let $\tilde{\Gamma}$ be the κ-coarsening of Γ. Then, we show in Lemma 15 below that every extreme posterior $\mu \in \{\boldsymbol{\mu_i}\}_{i=1,2,...,s}$ after coarsening leads to S with probability one, i.e., $\tilde{v}(\mu) = 1$.

Lemma 15. *For every $\mu \in \{\boldsymbol{\mu_i}\}_{i=1,2,...,s}$, if $\tilde{\lambda}(\mu) > 0$ then $\tilde{v}(\mu) = 1$.*

Proof. The beliefs projected into $\{\boldsymbol{\mu_i}\}_{i=1,2,...,s}$ come from a subset of $B^\kappa := \{\mu \in \triangle\Theta : \kappa(\mu) > \kappa\}$ and hence above the reputation boundary. By definition of reputation boundaries, once a posterior reaches a reputation boundary the expected discounted number of periods until the game is stopped converges to zero. Applying item iii) from Theorem 7 leads to the result. □

Next, we show that $\kappa = 1$ is a reputation boundary.

Lemma 16. *If $\kappa > 1$ is a reputation boundary, then there exists $\varepsilon \in (0, \kappa - 1)$ such that $\kappa - \varepsilon > 1$ is a reputation boundary.*

Proof. We describe the problem of finding an incentive-compatible centralized play Γ for a prior $\mu_0 \in \{\mu : \kappa(\mu) \in [\kappa - \varepsilon, \kappa]\}$ for some $\varepsilon > 0$, that maximizes player 2's payoff, and obeys the property that κ is a reputation boundary. Note that, every equilibrium has to satisfy incentive compatibility, and that κ is a reputation boundary. The solution to this constrained-maximization problem gives an upper bound on player 2's equilibrium payoffs. Let $\tilde{\Gamma}$ denote the κ-coarsening of Γ. Thus, player 2's payoff in incentive compatible and κ-coarse centralized play $\tilde{\Gamma} = (\tilde{\lambda}, \tilde{Z}, \tilde{v})$ when κ is a reputation boundary is given by

$$U_2 = \sum_{\mu \in \{\boldsymbol{\mu_i}\}_{i \geq 0} \cup \{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}\}_{i=1,2,\dots,s}} \tilde{\lambda}(\mu) \left[(1 - \tilde{v}(\mu)) \left(u_2\left(\tilde{Z}(\mu)\right) - V_2(\mu) \right) + V_2(\mu) \right]$$

subject to the following constraints:

$$\begin{aligned}
\bar{u} &= \sum_{\mu \in \boldsymbol{\mu_0} \cup \{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}\}_{i=1,2,\dots,s}} \tilde{\lambda}(\theta_0)(\mu) \left[(1 - \tilde{v}(\mu)) \left(u_1\left(\tilde{Z}(\mu)\right) - V_1 \right) + V_1 \right] \\
\bar{u} &\geq u_1(a_0) \\
\bar{u} &\geq \sum_{\mu \in \{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}, \boldsymbol{\mu_i}\}} \tilde{\lambda}(\theta_i)(\mu) \left[(1 - \tilde{v}(\mu)) \left(u_1\left(\tilde{Z}(\mu)\right) - V_1 \right) + V_1 \right] \text{ for every } i \in \{1, 2, \dots, s\} \\
\bar{u} &\geq (1 - \tilde{v}(\boldsymbol{\mu_i})) \left(u_1\left(\tilde{Z}(\boldsymbol{\mu_i})\right) - V_1 \right) + V_1 \text{ for every } i > s, \\
\tilde{v}(\boldsymbol{\mu_i}) &= 1 \text{ for every } i = 1, \dots, s. \\
\tilde{Z}(\mu) &= \alpha_i \text{ if } \mu(\theta_i) > 0 \text{ for some } i \in \{1, \dots, K\}.
\end{aligned}$$

The first equality defines player 1's payoff in this centralized play. The second constraint arises because $u_1(a_0)$ is player 1's minmax payoff. The third and fourth constraints are player 1's incentive compatibility constraints. The last 2 constraints follow from Lemma 15, and from $\tilde{\Gamma}$ being a centralized play (Lemma 12). Therefore we have for all $\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}$,

$$\left(1 - \tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\right) \left(u_1\left(\tilde{Z}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\right) - V_1 \right) = \left(1 - \tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\right) \left(u_1(\alpha_i) - V_1\right).$$

First notice that to maximize U_2, we can set $\tilde{v}(\boldsymbol{\mu_i}) = 0$ for $i > s$, since this choice maximizes player 2's payoff in the objective function, and does not upset any incentive-compatibility constraint as $u_2(\alpha_i) \geq V_2(\theta_i)$ and $u_1(a_0) \geq u_1(\alpha_i)$ for $i > s$.

Second, we will argue that if $\bar{u} = V_1$, then for every $i \in \{1, \dots, s\}$,

$$\tilde{\lambda}(\theta_i)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}\right) \tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}\right) + \tilde{\lambda}(\theta_i)\left(\boldsymbol{\mu_i}\right) \tilde{v}\left(\boldsymbol{\mu_i}\right) = \tilde{\lambda}(\theta_i)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}\right) + \tilde{\lambda}(\theta_i)\left(\boldsymbol{\mu_i}\right).$$

This implies that, if the coarse centralized play that maximizes player 2's payoff satisfies the property $\bar{u} = V_1$, then player 2 stops against types $i \in \{0, 1, ..., s\}$ almost immediately, i.e, $\kappa - \epsilon$ is a reputation boundary.

To see the claim, note that for $i \in \{1, ..., s\}$, $\tilde{\lambda}(\theta_i)(\boldsymbol{\mu_i})\,\tilde{v}(\boldsymbol{\mu_i}) = \tilde{\lambda}(\theta_i)(\boldsymbol{\mu_i})$. Moreover, if $\bar{u} = V_1$, then from the definition of $\bar{u}$, and from $V_1 > u_1(a)$ for every $a \in \mathbb{A}$, it follows that for each $\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}$, either $\tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}\right) = 1$, or $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}\right) = 0$. If $\tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}\right) = 1$, then the claim is true. If $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}\right) = 0$, because $\bar{\mu}_i^\kappa < 1$, $\tilde{\lambda}(\theta_i)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}\right) = 0$, and the claim is again true.

We will now show that $\bar{u} = V_1$. So now, on the way to a contradiction, suppose that $\bar{u} < V_1$.

We start with the claim that for every $i \in \{1, 2, ..., s\}$ with $\mu_0(\theta_i) > 0$, we have

$$\bar{u} \geq \left(1 - \tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\right)(u_1(\alpha_i) - V_1) + V_1.$$

Notice that, because the contradiction hypothesis is that $\bar{u} < V_1$, $\tilde{\lambda}(\theta_i)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}\right) > 0$. This is because, otherwise $\tilde{\lambda}(\theta_i)(\boldsymbol{\mu_i}) = 1$, and because $\tilde{v}(\boldsymbol{\mu_i}) = 1$, the IC constraint fails. Notice that if $\tilde{\lambda}(\theta_i)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}\right) > 0$, then $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}\right) > 0$. Suppose for some such i, we have $\bar{u} < \left(1 - \tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\right)(u_1(\alpha_i) - V_1) + V_1$. Then, the IC constraint

$$\bar{u} \geq \sum_{\mu \in \{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}, \boldsymbol{\mu_i}\}} \tilde{\lambda}(\theta_i)(\mu)\left[(1 - \tilde{v}(\mu))\left(u_1\left(\alpha_i\right) - V_1\right) + V_1\right]$$

fails. Note that because $\tilde{v}(\boldsymbol{\mu_i}) = 1$ for the set of types in focus now, we have

$$V_1 - \left(\sum_{\mu \in \{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}, \boldsymbol{\mu_i}\}} \tilde{\lambda}(\theta_i)(\mu)\left[(1 - \tilde{v}(\mu))\left(u_1\left(\tilde{Z}(\mu)\right) - V_1\right) + V_1\right]\right) =$$
$$\tilde{\lambda}(\theta_i)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\left(1 - \tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\right)(V_1 - u_1(\alpha_i(a)))\,.$$

We will now construct an alternative κ-coarse centralized play that satisfies the incentive constraints, and increases player 1's payoff by some small number $\eta > 0$, and that improves player 2's payoff. Consider a small increase in $\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})$, $\Delta\tilde{v}_i$, for each i in focus, i.e., $i \in \{1, ..., s\}$ with $\mu_0(\theta_i) > 0$, such that[17]

$$\tilde{\lambda}(\theta_i)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\Delta\tilde{v}_i\,(V_1 - u_1(\alpha_i)) = \eta.$$

If $\mu_0(\theta_i) = 0$, let $\Delta\tilde{v}_i = 0$. Note that such a change in $\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})$ will make player 1's

[17] Note that $\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}) < 1$, since $V_1 > \bar{u} \geq \left(1 - \tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\right)(u_1(\alpha_i) - V_1) + V_1$. Hence, for some $\eta > 0$, such a change is feasible.

payoff from mimicking any type i in focus to increase by (at most) η. Player 1's payoff also increases due to this change, but possibly not as much as η, so we will also change player 1's payoff by compensating him by an amount Δu by increasing $\tilde{v}(\boldsymbol{\mu_0})$ by some amount $\Delta\tilde{v}_0$ that increases $\bar{u}$ by η (hence ensuring that the IC constraints are satisfied). In particular, the expected payoff increase needed at belief $\boldsymbol{\mu_0}$ to ensure $\bar{u}$ increases by η is given by

$$\begin{aligned}\Delta u &:= \eta - \sum_{i\in\{1,2,\dots,s\}} \tilde{\lambda}(\theta_0)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\Delta\tilde{v}_i\left(V_1 - u_1(\alpha_i)\right) \\ &= \eta\left(1 - \sum_{i\in\{1,2,\dots,s\}} \frac{\tilde{\lambda}(\theta_0)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})}{\tilde{\lambda}(\theta_i)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})}\right) = \eta\left(1 - \sum_{i\in\{1,2,\dots,s\}} \frac{1-\bar{\mu}_i^{\kappa}}{\bar{\mu}_i^{\kappa}}\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)}\right).\end{aligned}$$

In the above equalities, the second one follows from $\tilde{\lambda}(\theta_i)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\Delta\tilde{v}_i\left(V_1 - u_1(\alpha_i)\right) = \eta$, and the third one follows from Bayes' rule. Recall that for $i \in \{1, 2, ..., s\}$

$$\bar{\mu}_i^{\kappa}\frac{1-\mu_i^*}{\mu_i^*} = (1-\bar{\mu}_i^{\kappa})\kappa.$$

Therefore,

$$\sum_{i\in\{1,2,\dots,s\}} \frac{1-\bar{\mu}_i^{\kappa}}{\bar{\mu}_i^{\kappa}}\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)} = \frac{1}{\mu_0(\theta_0)\kappa}\sum_{i\in\{1,2,\dots,s\}}\left(\frac{1-\mu_i^*}{\mu_i^*}\mu_0(\theta_i)\right).$$

Because by the assumption that $\kappa(\mu_0) \in [\kappa - \epsilon, \kappa]$, we have

$$\sum_{i\in\{1,2,\dots,s\}}\left(\frac{1-\mu_i^*}{\mu_i^*}\mu_0(\theta_i)\right) \geq (\kappa-\varepsilon)\mu_0(\theta_0).$$

Hence, we obtain that

$$\Delta u \leq \eta\left(1 - \frac{\kappa-\varepsilon}{\kappa}\right) = \eta\frac{\varepsilon}{\kappa}.$$

To calculate the impact of all the changes on player 2's payoff, first observe that there exists $\rho < \infty$, for which we can increase $\tilde{v}(\boldsymbol{\mu_0})$ by $\Delta\tilde{v}_0$ to increase player 1's payoff by Δu, in such a way that player 2's payoff does not decrease by more than $\Delta u\rho$.[18] Hence, a lower bound on the total impact on player 2's payoff is

$$\Delta U_2 \geq \left[\begin{array}{c}\sum_i \mu_0(\theta_i)\tilde{\lambda}(\theta_i)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\Delta\tilde{v}_i\left(V_2(\theta_i) - u_2(\alpha_i)\right) \\ +\sum \mu_0(\theta_0)\tilde{\lambda}(\theta_0)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\Delta\tilde{v}_i\left(V_2(\theta_0) - u_2(\alpha_i)\right) - \mu_0(\theta_0)\eta\frac{\varepsilon}{\kappa}\rho\end{array}\right].$$

[18]In particular, $\rho = \max_{a\in\mathbb{A}} \frac{u_2(a)-V_2(\theta_0)}{V_1-u_1(a)} = \frac{1}{\gamma_1}$.

Observe that because $\kappa > 1$, for every $i \in \{1, ..., s\}$, we have $V_2(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}) > u_2(\alpha_i)$.[19] Hence,

$$\mu_0(\theta_i)\tilde{\lambda}(\theta_i)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\Delta\tilde{v}_i\left(V_2(\theta_i) - u_2(\alpha_i)\right) + \mu_0(\theta_0)\tilde{\lambda}(\theta_0)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\Delta\tilde{v}_i\left(V_2(\theta_0) - u_2(\alpha_i)\right)$$
$$= \tilde{\lambda}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\Delta\tilde{v}_i\left(V_2(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}) - u_2(\alpha_i)\right) > 0.$$

Moreover, because $\tilde{\lambda}(\theta_i)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\Delta\tilde{v}_i\left(V_1 - u_1(\alpha_i)\right) = \eta$, we have

$$\tilde{\lambda}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}) > \mu_0(\theta_i)\tilde{\lambda}(\theta_i)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}) \geq \mu_0(\theta_i)\frac{\eta}{V_1 - u_1(\alpha_i)}.$$

Hence,

$$\Delta U_2 \geq \eta \sum_{i \in \{1,2,...,s\}} \mu_0(\theta_i) \min_{i \in \{1,2,...,s\}} \left\{ \frac{V_2(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}) - u_2(\alpha_i)}{V_1 - u_1(\alpha_i)} \right\} - \mu_0(\theta_0)\eta\frac{\varepsilon}{\underline{\kappa}}\rho$$

Clearly, as $\varepsilon \to 0$, $\Delta U_2 > 0$ for all $\mu_0 \in \{\mu \in \Delta\Theta : \kappa(\mu) \in [\kappa - \varepsilon, \kappa]\}$. So, player 2's payoff maximizing coarse centralized play that satisfies the IC constraints has the property that $\bar{u} = V_1$, and hence $\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}) = 1$ for every $i \in \{1, 2, ..., s\}$ that satisfies $\tilde{\lambda}(\theta_i)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}) > 0$. This implies that, $\kappa - \varepsilon$ is a reputation boundary for some $\varepsilon > 0$. □

Corollary 2. $\kappa = 1$ *is a reputation boundary.*

Proof. Follows from Lemma 16. □

E.3. Low Reputation Priors We now consider prior beliefs μ_0 such that

$$\sum_{i \in \{1,2,...,s\}} \mu_0(\theta_i)\frac{1 - \mu_i^*}{\mu_i^*} < \mu_0(\theta_0).$$

We start by arguing the following claim which argues that a sufficiently patient player 2 never plays S against a type $\theta \in \{\theta_{s+1}, ..., \theta_K\}$, or if her belief is below the reputation boundary.

Lemma 17. *For every $\xi > 0$, there is a $\bar{\delta} < 1$ such that if $\delta > \bar{\delta}$, S is taken with probability 0 at history h^0 (the beginning of the game) across all NE of the game with discount factor δ and prior belief μ_0 where $\sum_{i>s} \mu_0(\theta_i) > \xi$ or $\sum_{i \in \{1,2,...,s\}} \mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} < (1 - \xi)\mu_0(\theta_0)$.*

Proof. For the first part of the claim, suppose towards a contradiction that there is a sequence of equilibria of a sequence of games with discount factors $\delta_n \to 1$ and $\mu_n \to \mu_0$ with $\sum_{i>s} \mu_0(\theta_i) \geq \xi$ and in which player 2 plays S at h^0. Consider the deviation strategy

[19]Recall that $V_2(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}) = \bar{\mu}_i^{\kappa} V_2(\theta_i) + (1 - \bar{\mu}_i^{\kappa})V_2(\theta_0)$, and that $\mu_0(\theta_i)\tilde{\lambda}(\theta_i)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}) = \tilde{\lambda}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\bar{\mu}_i^{\kappa}$, $\mu_\Gamma(\theta_0)\tilde{\lambda}(\theta_0)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}}) = \tilde{\lambda}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{\kappa}})\left(1 - \bar{\mu}_i^{\kappa}\right)$.

by Player 2 that plays S at the first history h^t at which $\sum_{i>s}\mu(h^t)(\theta_i)<\xi/2$. We know from Lemma ii that under the deviation strategy, player 2 plays S at a (random) time t_n such that $\mathbb{E}^{\theta_i}\left(\delta^{t_n}\right)$ converges to one for any type θ_i for $i\in\{1,2,...,s\}$, and plays S against types θ_i for $i>s$ with a total probability bounded above by $\xi/2$. Against the normal type, the worst payoff she can get is $V_2(\theta_0)$, which she gets in the sequence of strategy profiles where player 2 is using the deviation strategies. The deviation strategy is profitable when n is large, because it improves player 2's payoff at least by an amount converging to $\frac{\xi}{2}\left[\min_{j>s}u_2(\alpha_j)-V_2(\theta_j)\right]$ (and hence at least as large as $\frac{\xi}{4}\left[\min_{j>s}u_2(\alpha_j)-V_2(\theta_j)\right]$ for n large enough), which gives the desired contradiction.

For the second part of the claim, for each $m\in\mathbb{N}$, consider the sequence of deviation strategies $(\sigma_{2,m,n})_{n=1}^{\infty}$ that plays S at the first history h^t with $\mu(h^t)(\theta_i)>\mu_i^*$ and $\sum_{j\notin\{0,i\}}\mu(\theta_j)\leq m^{-1}$ for some $i\in\{1,...,s\}$. Let $\Gamma_{m,n}=(\lambda_{m,n},Z_{m,n},v_{m,n})$ be the generalized centralized play that implements the strategy profiles generated by player 1's equilibrium strategy and the deviation strategy $\sigma_{2,m,n}$ of player 2, when the discount factor is δ_n.

Notice that, for every $m\in\mathbb{N}$, the expected discounted number of periods in which W is taken and the belief lies in

$$B_m:=\bigcup_{i=1}^{s}\left\{\mu\in\Delta\Theta:\mu(\theta_i)>\mu_i^*+m^{-1}\right\}$$

is equal to the expected discounted number of periods in which W is taken and the belief lies in

$$\tilde{B}_m:=\bigcup_{i=1}^{s}\left\{\mu\in\Delta\Theta:\mu(\theta_i)>\mu_i^*+m^{-1},\sum_{j\notin\{0,i\}}\mu(\theta_j)>m^{-1}\right\}.$$

Lemma 9 implies this expected discounted number of periods converges to zero as $n\to\infty$. Using this property and a diagonal argument we conclude that $\{\Gamma_{m,n}\}$ contains a subsequence $\left\{\Gamma_{m(r),n(r)}\right\}_{r=1}^{\infty}$ such that $\lim_{r\to\infty}\min\{n\left(r\right),m\left(r\right)\}=\infty$ and for which the expected discounted number of periods in which the belief lies in the set B_r is smaller than r^{-1} in $\Gamma_{m(r),n(r)}$.

Moreover, in light of Theorem Theorem 6, $\left\{\Gamma_{m(r),n(r)}\right\}_{r=1}^{\infty}$ has a subsequence converging to some limit centralized play $\Gamma=(\lambda,Z,v)$. Take this limit centralized play and consider the corresponding 1−coarsening of Γ, $\tilde{\Gamma}=(\tilde{\lambda},\tilde{Z},\tilde{v})$, which decomposes the posteriors in Γ into $\left\{\{\boldsymbol{\mu_i}\}_{i=0,1,...K}\cup\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}}\}_{i=1,2,...,s}\right\}$. That is, we perform the coarsening procedure described previously with $\kappa=1$. We have the following claim about the properties of $\tilde{\Gamma}=(\tilde{\lambda},\tilde{Z},\tilde{v})$:

Claim 2. The following hold:

1. $\tilde{v}\left(\boldsymbol{\mu_i}\right) = 1$ for $\mu \in \boldsymbol{\mu_i}$, for every $i \in \{1, 2, ..., s\}$ if $\tilde{\lambda}\left(\boldsymbol{\mu_i}\right) > 0$.

2. $\tilde{v}\left(\boldsymbol{\mu_i}\right) = 0$ for $\mu \in \boldsymbol{\mu_i}$, for every $i > s$ if $\tilde{\lambda}\left(\boldsymbol{\mu_i}\right) > 0$.

3. $\tilde{v}(\boldsymbol{\mu_0}) = 0$ if $\tilde{\lambda}\left(\boldsymbol{\mu_0}\right) > 0$.

Proof. Item 1 follows from the argument in Lemma 15 and the observation that the expected discounted number of periods in which the belief lies in the set B_r is smaller than r^{-1} in $\Gamma_{m(r),n(r)}$. Item 2 follows because, by construction if S is played at some history h^t, then $\mu(h^t)(\theta_i) \leq \frac{1}{m(r)}$ for $i > s$ in the strategy profile that $\Gamma_{m(r),n(r)}$ implements. Hence the probability that S is played conditional on type $i > s$ converges to zero and thus:

$$\tilde{v}\left(\boldsymbol{\mu_i}\right) = \lim_{r\to\infty} \int_{\mu\in\Delta\Theta} \lambda_{m(r),n(r)}\left(\theta_i\right)(\mu) v_{m(r),n(r)}\left(\mu\right)) d\mu = 0.$$

Item 3 follows, because by construction if S is played at a history h^t, then for some $i \in \{1, 2, ..., s\}$, $\mu(h^t)(\theta_i) > \mu_i^* + m^{-1}(r)$ and $\mu(h^t)(\theta_j) \leq \frac{1}{m(r)}$ for every $j \neq 0, i$ in the strategy profile that $\Gamma_{m(r),n(r)}$ implements. Similar to item 2, and using the feature of the coarsening procedure that for such beliefs, the projection gives 0 weight on posterior $\boldsymbol{\mu_0}$, it is straightforward to verify that the projection of such posteriors on $\{\boldsymbol{\mu_0}\}$ vanishes as $r \to \infty$, which immediately implies that $\tilde{v}(\boldsymbol{\mu_0})\tilde{\lambda}\left(\boldsymbol{\mu_0}\right) = 0$. $\square$

Completion of the proof of Lemma 17:
Recall that $\underline{u}_2 = \min_{a\in\mathbb{A}} u_2(a)$.
The payoff from the equilibrium strategy of playing S at h^0 is equal to

$$\sum_{i>s} \mu_0(\theta_i)V_2(\theta_i) + \sum_{i\leq s} \mu_0(\theta_i)V_2(\theta_i). \tag{29}$$

Player 2's payoff from the deviation strategies converge to the expected payoff from the coarse centralized play $\tilde{\Gamma} = (\tilde{\lambda}, \tilde{Z}, \tilde{v})$, which is at least

$$\sum_{i>s} \mu_0(\theta_i)u_2(\alpha_i) + \sum_{i\in\{1,2,...,s\}} \left(\tilde{\lambda}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})V_2(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}}) + \tilde{\lambda}(\boldsymbol{\mu_i})V_2(\theta_i)\right) + \tilde{\lambda}(\boldsymbol{\mu_0})\underline{u}_2. \tag{30}$$

We will show that player 2's payoff from expression 30 is strictly higher that her payoff from stopping in the beginning of the game, i.e., expression 29. Notice that the payoff in 30 is minimized (over all coarse centralized plays) when $\tilde{\lambda}(\boldsymbol{\mu_i}) = 0$ for all $i \in \{1, 2, ..., s\}$.

This is because, for such θ_i,

$$\mu_i^* V_2(\theta_i) + (1-\mu_i^*)\underline{\mathrm{u}}_2 > \mu_i^* V_2(\theta_i) + (1-\mu_i^*) V_2(\theta_0) = V_2(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}).$$

Therefore, player 2's payoffs are bounded below by

$$\sum_{i>s} \mu_0(\theta_i) u_2(\alpha_i) + \sum_{i\in\{1,2,\dots,s\}} \frac{\mu_0(\theta_i)}{\mu_i^*} V_2(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}) + \left(\mu_0(\theta_0) - \sum_{i\in\{1,2,\dots,s\}} \mu_0(\theta_i) \frac{1-\mu_i^*}{\mu_i^*}\right) \underline{\mathrm{u}}_2. \quad (31)$$

The payoff to player 2 from stopping at the beginning of the game, i.e., expression 29 is equal to

$$\sum_{i>s} \mu_0(\theta_i) V_2(\theta_i) + \sum_{i\in\{1,2,\dots,s\}} \frac{\mu_0(\theta_i)}{\mu_i^*} V_2(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}) + \left(\mu_0(\theta_0) - \sum_{i\in\{1,2,\dots,s\}} \mu_0(\theta_i) \frac{1-\mu_i^*}{\mu_i^*}\right) V_2(\theta_0). \quad (32)$$

To see this, we use the equality $V_2(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}) = \mu_i^* V_2(\theta_i) + (1-\mu_i^*) V_2(\theta_0)$. Because $V_2(\theta_0) < \underline{\mathrm{u}}_2$ and $\left(\mu_0(\theta_0) - \sum_{i\in\{1,2,\dots,s\}} \mu_0(\theta_i) \frac{1-\mu_i^*}{\mu_i^*}\right) > \mu_0(\theta_0)\xi > 0$, and because $V_2(\theta_i) < u_2(\alpha_i)$ for $i > s$, the deviation strategy which delivers a limit payoff at least as much as (31) is strictly higher than the payoff (32). Therefore, the deviation strategy is profitable when δ is larger than some $\bar{\delta} < 1$. □

Lemma 18. *For any prior $\mu_0 \in \Delta\Theta$ such that $\sum_{i\in\{1,2,\dots,s\}} \mu_0(\theta_i) \frac{1-\mu_i^*}{\mu_i^*} < \mu_0(\theta_0)$, player 1's NE payoffs converge to $u_1(a_0)$ as $\delta \to 1$.*

Proof. Take a prior μ_0 that satisfies the condition stated in the Lemma. Player 1's NE payoffs are bounded below by $u_1(a_0)$ since this is the lowest payoff he gets if plays always a_0 in every period (until player 2 plays S). We will show that the limit of his equilibrium payoffs cannot be higher than $u_1(a_0)$. Take a sequence of NE $\{\sigma_n\}$ of a sequence of games with associated discount factors $\delta_n \to 1$. Taking a subsequence if necessary, assume that the sequence of generalized centralized plays that implement the sequence of NE converges to a limit centralized play $\Gamma = (\lambda, Z, v)$. Let $\tilde{\Gamma} = (\tilde{\lambda}, \tilde{Z}, \tilde{v})$ be the 1-coarsening of Γ in which posteriors are decomposed into $\left\{\{\boldsymbol{\mu_i}\}_{i=0,1,\dots K} \cup \{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\}_{i=1,2,\dots,s}\right\}$ (i.e., $\tilde{\Gamma}$ is obtained from Γ by using the coarsening procedure for $\kappa = 1$.). We have the following claim about the properties of $\tilde{\Gamma} = (\tilde{\lambda}, \tilde{Z}, \tilde{v})$:

Claim 3. The following hold:

1. $\tilde{v}\left(\boldsymbol{\mu_i}\right) = 1$ for every $i \in \{1, 2, ..., s\}$ if $\tilde{\lambda}\left(\boldsymbol{\mu_i}\right) > 0$.

2. $\tilde{v}(\boldsymbol{\mu_i}) = 0$ for every $i > s$ if $\tilde{\lambda}(\boldsymbol{\mu_i}) > 0$.

3. $\tilde{v}(\boldsymbol{\mu_0}) = 0$ and $\tilde{\lambda}(\boldsymbol{\mu_0}) > 0$.

Proof. Item 1 follows because $\kappa = 1$ is a reputation boundary and by Lemma 17. For items 2 and 3, first take $\varepsilon > 0$ and consider the (open) set

$$D_\varepsilon := \{\mu \in \Delta\Theta : \mu(\{\theta_{s+1}, ..., \theta_K\}) > \varepsilon\} \bigcup \left\{\mu \in \Delta\Theta : \sum_{i\in\{1,2,...,s\}} \mu(\theta_i) \frac{1-\mu_i^*}{\mu_i^*} < (1-\varepsilon)\mu(\theta_0)\right\}$$

Lemma 17 implies that the probability that S is taken at a history h^t such that $\mu(h^t)$ belongs to D_ε converges to zero. Theorem 7 (item ii) together with the coarsening procedure implies that $\int_{\mu\in D_\varepsilon} \tilde{\lambda}(\mu)\tilde{v}(\mu)d\mu = 0$. Since D_ε increases to

$$D := \{\mu \in \Delta\Theta : \mu(\{\theta_{s+1}, ..., \theta_K\}) > 0\} \bigcup \left\{\mu \in \Delta\Theta : \sum_{i\in\{1,2,...,s\}} \mu(\theta_i) \frac{1-\mu_i^*}{\mu_i^*} < \mu(\theta_0)\right\}$$

as ε decreases to zero, the monotone convergence theorem implies $\int_{\mu\in D} \tilde{\lambda}(\mu)\tilde{v}(\mu)d\mu = 0$, which then implies item 2 that $\tilde{v}(\boldsymbol{\mu_i}) = 0$ for every $i > s$ if $\tilde{\lambda}(\boldsymbol{\mu_i}) > 0$. It also implies the first part of item 3 that $\tilde{v}(\boldsymbol{\mu_0}) = 0$ if $\tilde{\lambda}(\boldsymbol{\mu_0}) > 0$. The conclusion that $\tilde{\lambda}(\boldsymbol{\mu_0}) > 0$ follows, because $\sum_{i\in\{1,2,...,s\}} \mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} < \mu_0(\theta_0)$. □

Completion of the proof of Lemma 18: Assume towards a contradiction that we can find a sequence of NE $\{\sigma_n\}$ of the sequence of games with associated discount factors $\delta_n \to 1$ and in which the normal type obtains a limit payoff that is strictly larger than $u_1(a_0)$.

Notice that if the normal type imitates type θ_i for some $i \in \{1, ..., s\}$ he obtains

$$\begin{aligned} U_i \quad &: \quad = \tilde{\lambda}(\theta_i)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})\left(\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})(V_1 - u_1(\alpha_i)) + u_1(\alpha_i)\right) + \tilde{\lambda}(\theta_i)(\boldsymbol{\mu_i})V_1 && (33)\\ &\geq \quad \tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})(V_1 - u_1(\alpha_i)) + u_1(\alpha_i), && (34)\end{aligned}$$

where the inequality follows because $V_1 > u_1(\alpha_i)$.

Let $v_1 := u_1\left(\tilde{Z}(\boldsymbol{\mu_0})\right)$ be player 1's payoff when the centralized play announces $\boldsymbol{\mu_0}$ (recall property 3 of the claim that, $\tilde{v}(\boldsymbol{\mu_0}) = 0$).

Because $\tilde{v}(\boldsymbol{\mu_0}) = 0$, $v_1 < V_1$. Player 1's payoff in the centralized play is then

$$U_0 := \sum_{i\in\{1,2,...,s\}} \tilde{\lambda}(\theta_0)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})\left(\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})(V_1 - u_1(\alpha_i)) + u_1(\alpha_i)\right) + \tilde{\lambda}(\theta_0)(\boldsymbol{\mu_0})v_1. \quad (35)$$

Suppose now on the way to a contradiction that $U_0 > u_1(a_0)$. Because $v_1 \leq u_1(a_0)$, there is some $i \in \{1, 2, .., s\}$ with $\tilde{\lambda}(\theta_0)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}}) > 0$ and $\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})\,(V_1 - u_1(\alpha_i)) + u_1(\alpha_i) > u_1(a_0)$. Let $U^* = \max_{i \in \{1,\ldots,s\}} \tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})\,(V_1 - u_1(\alpha_i)) + u_1(\alpha_i)$, and suppose U^* is attained for type θ_i. Then:

$$U_i = \tilde{\lambda}(\theta_i)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})\left(\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})\,(V_1 - u_1(\alpha_i)) + u_1(\alpha_i)\right) + \tilde{\lambda}(\theta_i)(\boldsymbol{\mu_i})V_1 \geq U^* > U_0,$$

where the first inequality follows from inequality (33), and the second inequality follows because $v_1 \leq u_1(a_0) < U^*$, $U^* = \max_{i \in \{1,\ldots,s\}} \tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})\,(V_1 - u_1(\alpha_i)) + u_1(\alpha_i)$, and $\tilde{\lambda}(\theta_0)(\boldsymbol{\mu_0}) > 0$. The last inequality that $\tilde{\lambda}(\theta_0)(\boldsymbol{\mu_0}) > 0$ follows from property 3 of the claim. But then $U_0 < U_i$, which is a violation of the IC constraint that type θ_0 does not strictly prefer to imitate type θ_i. This gives the desired contradiction. □

The proof of Lemma 18 delivers the following corollary:

Corollary 3. *Assume that* $\sum_{i \in \{1,2,\ldots,s\}} \mu_0(\theta_i) \frac{1-\mu_i^*}{\mu_i^*} < \mu_0(\theta_0)$ *and let* $\tilde{\Gamma}$ *be the 1-coarsening of a limit centralized play* Γ *of an equilibrium sequence of a sequence of games where* $\delta_n \to 1$*. The following properties are true:*

1. $\tilde{v}\,(\boldsymbol{\mu_i}) = 1$ *for every* $i \in \{1, 2, ..., s\}$ *if* $\tilde{\lambda}\,(\boldsymbol{\mu_i}) > 0$.
2. $\tilde{v}\,(\boldsymbol{\mu_i}) = 0$ *for every* $i > s$.
3. $\tilde{v}(\boldsymbol{\mu_0}) = 0$ *and* $\tilde{\lambda}\,(\boldsymbol{\mu_0}) > 0$.

Proof. Item 2. follows from the fact that given $\lambda\,(\theta_i)$ for $i > s$, the coarsening procedure projects every posterior on $\boldsymbol{\mu_i}$ and from the first statement of Lemma 17. Items 1. and 3. were proven in claim 3 located with in the proof of Lemma 18. □

Lemma 19. *Assume that* $\sum_{i \in \{1,2,\ldots,s\}} \mu_0(\theta_i) \frac{1-\mu_i^*}{\mu_i^*} < \mu_0(\theta_0)$ *and let* $\tilde{\Gamma}$ *be the 1-coarse centralized play of the limit centralized play* Γ *of an equilibrium sequence. The following properties are true:*

1. $\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})\,(V_1 - u_1(\alpha_i)) + u_1(\alpha_i) = u_1\,(a_0)$ *for every* $i \in \{1, ..., s\}$.
2. $\tilde{\lambda}(\theta_i)(\boldsymbol{\mu_i}) = 0$ *for every* $i \in \{1, ..., s\}$.
3. $\tilde{Z}(\boldsymbol{\mu_0}) = a_0$.

Proof. Items 1 and 3: If $\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}})\left(V_1 - u_1(\alpha_i)\right) + u_1(\alpha_i) > u_1(a_0)$ for some $i \in \{1, ..., s\}$ then (33) implies that the normal type can profitably deviate by imitating type θ_i, and obtains a payoff strictly higher than $u_1(a_0)$, a contradiction. So we have for every $i \in \{1, ..., s\}$

$$\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}})\left(V_1 - u_1(\alpha_i)\right) + u_1(\alpha_i) \leq u_1(a_0)\,. \tag{36}$$

Observe that if $\tilde{\lambda}(\theta_i)(\boldsymbol{\mu_i}) = 1$ for some $i \in \{1, ..., s\}$, then $U_i > U_0$, which violates IC constraint. Then $\tilde{\lambda}(\theta_i)(\boldsymbol{\mu_i}) < 1$ for every such i, which thus implies $\tilde{\lambda}(\theta_0)(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}) > 0$ for every such i. Because $\tilde{v}(\boldsymbol{\mu_0}) = 0$ by property 3 of Corollary 3, $v_1 \leq u_1(a_0)$. Moreover, because U_0 is a weighted average of terms less than or equal to $u_1(a_0)$, and because $U_0 = u_1(a_0)$, we have that for every $i \in \{1, ..., s\}$, $\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}})\left(V_1 - u_1(\alpha_i)\right) + u_1(\alpha_i) = u_1(a_0)$, and $v_1 = u_1\left(\tilde{Z}(\boldsymbol{\mu_0})\right) = u_1(a_0)$. Hence, $\tilde{Z}(\boldsymbol{\mu_0}) = a_0$.

Item 2: Since $\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}})\left(V_1 - u_1(\alpha_i)\right) + u_1(\alpha_i) = u_1(a_0)$, which follows from item 1, if $\tilde{\lambda}(\theta_i)(\boldsymbol{\mu_i}) > 0$ then $U_i > u_1(a_0)$, which violates the IC constraint. □

The next lemma shows that the limit centralized play Γ of an equilibrium sequence is essentially equal to the 1-coarse centralized play $\tilde{\Gamma}$ with some additional properties.[20]

Lemma 20. *Assume that $\sum_{i \in \{1,2,...,s\}} \mu_0(\theta_i) \frac{1-\mu_i^*}{\mu_i^*} < \mu_0(\theta_0)$. Let Γ be a limit centralized play of an equilibrium sequence of a sequence of games where $\delta_n \to 1$. The following properties are true:*

1. $\int_{\mu = \boldsymbol{\mu_i}} \lambda(\mu)\left(\theta_i\right)[d\mu] = 1$ *for* $i > s$.
2. $\int_{\mu \in \left\{\boldsymbol{\mu_0}, \tilde{\boldsymbol{\mu}}_{\mathbf{1}}^{\mathbf{1}}, ..., \tilde{\boldsymbol{\mu}}_{\boldsymbol{s}}^{\mathbf{1}}\right\}} \lambda(\mu)\left(\theta_0\right)[d\mu] = 1.$
3. $\int_{\mu = \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}} \lambda(\mu)\left(\theta_i\right)[d\mu] = 1$ *for* $i \in \{1, ..., s\}$.

Proof. We start by proving item 1:

Fix $i > s$. Recall that in a limit centralized play the following holds:

$$\int_{\{\mu:\mu(\theta_i)>0,\mu(\theta_j)>0 \text{ for some } j\neq\{0,i\}\}} v(\mu)\lambda(\mu)[d\mu] = \int_{\{\mu:\mu(\theta_i)>0,\mu(\theta_j)>0 \text{ for some } j\neq\{0,i\}\}} \lambda(\mu)[d\mu].$$

Recall that by corollary 3, $\tilde{v}(\boldsymbol{\mu_i}) = 0$ for every $i > s$. Hence, $\int_{\{\mu:\mu(\theta_i)>0,\mu(\theta_j)>0 \text{ for some } j\neq\{0,i\}\}} \lambda(\mu)[d\mu] = 0$. Therefore, $\int \lambda(\mu)\left(\theta_i\right)[d\mu] = \int_{\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}} \lambda(\mu)\left(\theta_i\right)[d\mu]$, and $\int_{\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}} v(\mu)\lambda(\mu)\left(\theta_i\right)[d\mu] = 0$. Note that $a_0 = \tilde{Z}(\boldsymbol{\mu_0})$ from Lemma 19, and the beliefs $\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}$ are decomposed on $\boldsymbol{\mu_0}$ and $\boldsymbol{\mu_i}$ in the coarsening procedure. Because $Z(\mu) = \alpha_i$ for all $\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}$, we obtain that $\int_{\{\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}:\mu(\theta_i)<1\}} \lambda(\mu)\left[d\mu\right] = 0$. Therefore, $\int_{\mu = \boldsymbol{\mu_i}} \lambda(\mu)\left(\theta_i\right)[d\mu] = 1$.

[20] They may differ in sets in which Γ puts zero measure.

Now we prove item 3:

Take $i \in \{1, ..., s\}$. By item 2 of Lemma 19, $\tilde{\lambda}(\theta_i)(\boldsymbol{\mu_i}) = 0$.

Assume towards a contradiction that $\int_{\left\{\mu\in\Delta\Theta\setminus\left(\cup_{i=1}^{K}\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\right)\bigcup\{\boldsymbol{\mu_0}\}\right\}} \lambda(\mu)\left(\theta_i\right)[d\mu] > 0$. Next define, for each $m \in \mathbb{N}$, the open set $B_m := \left\{\mu \in \Delta\Theta{:}d\left(\mu, \left(\cup_{i=1}^{K}\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\right)\bigcup\{\boldsymbol{\mu_0}\}\right) > m^{-1}\right\}$ and notice that $B_m \uparrow \Delta\Theta\setminus\left(\cup_{i=1}^{K}\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\right)\bigcup\{\boldsymbol{\mu_0}\}$. Hence the contradiction assumption implies that there exists $m^* \in \mathbb{N}$ for which $\xi := \int_{B_{m^*}} \lambda(\mu)\left(\theta_i\right)[d\mu] > 0$. Thus Portmanteau theorem (Billingsley (2013), Theorem 2.1) implies

$$\operatorname{liminf}_{n\to\infty} \int_{B_{m^*}} \lambda_n(\mu)\left(\theta_i\right)[d\mu] \geq \int_{B_{m^*}} \lambda(\mu)\left(\theta_i\right)[d\mu] > 0. \tag{37}$$

Next, for each $l\in \mathbb{N}$, consider the following sequence of deviating strategies $(\sigma_{1,l,n})_{n=1}^{\infty}$ for the normal type that prescribe α_i until period $t_{\delta_n}(l)$ and the action a_0 in each future period. But then an application of Lemma 9 implies that there exists $n_l^* \in \mathbb{N}$ such that for all $n \geq n_l^*$, the normal type is stopped before period $t_{\delta_n}(l)$ with probability at least as large as $\frac{\xi}{2}$ under $\sigma_{2,n}$. It then follows that (for $n \geq n_l^*$) this deviating strategy leads to a payoff at least as large as

$$\frac{\xi}{2}\left[\left(2l^{-1}\right)u_1\left(\alpha_1\right) + (1-2l^{-1})V_1\right] + \left(1-\frac{\xi}{2}\right)\left[\left(2l^{-1}\right)u_1\left(\alpha_1\right) + (1-2l^{-1})u_1\left(a_0\right)\right],$$

which is greater than $u_1\left(a_0\right)$ for l^* sufficiently large. Therefore, since the limit payoff from the equilibrium strategies $(\sigma_{1,n})_{n=1}^{\infty}$ is $u_1\left(a_0\right)$, we can conclude that $\sigma_{1,l^*,n}$ leads to a profitable deviation whenever n is sufficiently large, a contradiction.

Next assume towards a contradiction that $\int_{\left\{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}:\mu(\theta_i)>\mu_i^*\right\}} \lambda(\mu)\left(\theta_i\right)[d\mu] > 0$ and notice that the coarsening procedure would imply $\tilde{\lambda}(\theta_i)(\boldsymbol{\mu_i}) > 0$ in this case, which would contradict Lemma 19.

Finally, assume towards a contradiction that $\int_{\left\{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}:\mu(\theta_i)<\mu_i^*\right\}} \lambda(\mu)\left(\theta_i\right)[d\mu] > 0$ and notice that the second statement in Lemma 17 implies that

$$\int_{\left\{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}:0<\mu(\theta_i)<\mu_i^*\right\}} \lambda(\mu)\left(\theta_i\right)[d\mu] = \int_{\left\{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}:0<\mu(\theta_i)<\mu_i^*\right\}} (1-v\left(\mu\right))\lambda(\mu)\left(\theta_i\right)[d\mu].$$

Moreover, since $Z\left(\mu\right) = \alpha_i$ in $\{\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}} : 0 < \mu\left(\theta_i\right) < \mu_i^*\}$, letting $\left(\frac{d\lambda(\theta_i)}{d\lambda(\theta_0)}\right)$ be the Radon-Nikodym derivative of $\lambda\left(\theta_i\right)$ with respect to $\lambda\left(\theta_0\right)$, we get:

$$\int_{\left\{\mu\in\tilde{\boldsymbol{\mu_i}}:0<\mu(\theta_i)<\mu_i^*\right\}}(1-v(\mu))\lambda(\mu)(\theta_i)[d\mu]=\int_{\left\{\mu\in\tilde{\boldsymbol{\mu_i}}:0<\mu(\theta_i)<\mu_i^*,Z(\mu)=\alpha_i\right\}}(1-v(\mu))\lambda(\mu)(\theta_i)[d\mu]$$
$$=\int_{\left\{\mu\in\tilde{\boldsymbol{\mu_i}}:0<\mu(\theta_i)<\mu_i^*,Z(\mu)=\alpha_i\right\}}(1-v(\mu))\lambda(\mu)(\theta_0)\left(\frac{d\lambda(\theta_i)}{d\lambda(\theta_0)}\right)(\mu)[d\mu].$$

But this implies that a positive measure sets of beliefs in which $Z(\mu)=\alpha_i$ is projected into $\boldsymbol{\mu_0}$, which contradicts our finding that $a_0=\tilde{Z}(\boldsymbol{\mu_0})$ from Lemma 19.

Finally, notice that item 2 follows immediately from items 1 and 3 above. □

Corollary 4. *Assume that* $\sum_{i\in\{1,2,...,s\}}\mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*}<\mu_0(\theta_0)$*. Let* Γ *be the limit centralized play of an equilibrium sequence. For* $i\in\{1,...,s\}$*, we have* $\lambda(\theta_0)\left(\left\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}}\right\}\right)=\left(\frac{1-\mu_i^*}{\mu_i^*}\right)\left(\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)}\right)$*. Moreover, we have* $\lambda(\theta_0)(\{\boldsymbol{\mu_0}\})=1-\sum_{i=1}^{s}\left(\frac{1-\mu_i^*}{\mu_i^*}\right)\left(\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)}\right)$.

Proof. The fact that $\lambda(\theta_0)\left(\left\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}}\right\}\right)=\left(\frac{1-\mu_i^*}{\mu_i^*}\right)\left(\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)}\right)$ follows from our finding from Lemma 20 that $\lambda(\mu)(\theta_i)\left(\left\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}}\right\}\right)=1$ and Bayes' rule. The fact that $\lambda(\theta_0)(\{\boldsymbol{\mu_0}\})=1-\sum_{i=1}^{s}\left(\frac{1-\mu_i^*}{\mu_i^*}\right)\left(\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)}\right)$ follows from the last finding and our finding $\int_{\mu\in\left\{\boldsymbol{\mu_0},\tilde{\boldsymbol{\mu}}_{\boldsymbol{1}}^{\boldsymbol{1}},...,\tilde{\boldsymbol{\mu}}_{\boldsymbol{s}}^{\boldsymbol{1}}\right\}}\lambda(\mu)(\theta_0)[d\mu]=1$ from Lemma 20. □

Corollary 5. *Assume that* $\sum_{i\in\{1,2,...,s\}}\mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*}<\mu_0(\theta_0)$*. Let* Γ *be the limit centralized play of an equilibrium sequence. We have* $v(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})=\frac{u_1(a_0)-u_1(\alpha_i)}{V_1-u_1(\alpha_i)}$ *for every* $i=1,...,s$.

Proof. Our finding in Lemma 20 that $\lambda(\theta_i)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}}\right)=\tilde{\lambda}(\theta_i)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}}\right)$ for every $i=1,...,s$ implies that $v(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})=\tilde{v}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})$. Therefore item 1 from Lemma 19 implies $v(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}})=\frac{u_1(a_0)-u_1(\alpha_i)}{V_1-u_1(\alpha_i)}$. □

E.4. Proofs of Theorems 1 and 2 for Screening Games We show Theorems 1 and 2 together for each of the cases, i.e., when we have high-reputation priors $\sum_{i\in\{1,2,...,s\}}\mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*}>\mu_0(\theta_0)$ or low-reputation priors $\sum_{i\in\{1,2,...,s\}}\mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*}<\mu_0(\theta_0)$:

E.4.1. *Case 1: Inequality (3) holds* $\left(\sum_{i\in\{1,2,...,s\}}\mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*}>\mu_0(\theta_0)\right)$.

Claim 4. $\mathbb{E}_{(\sigma_n)}^{\theta_i}\left(\delta_n^{\mathbb{T}}\right)\to 1$ for every $i\in\{0,...,s\}$. Player 1's equilibrium payoff converges to V_1.

Proof. Follows from Lemma 2. □

Claim 5. The equilibrium payoff of player 2 converges to

$$v_2:=\sum_{i\in\{0,1,...,s\}}\mu_0(\theta_i)V_2(\theta_i)+\sum_{i\in\{s+1,...,K\}}\mu_0(\theta_i)u_2(\alpha_i).$$

Proof. By claim 4, player 2's payoff against each type $\theta \in \{\theta_0, ..., \theta_s\}$ converges to $V_2(\theta)$. Player 2's payoff against each type $\theta \in \{\theta_{s+1}, ..., \theta_K\}$ is weakly below $u_2(\alpha_i)$, and is achieved by playing W against such types at every period. Hence, player 2's limit equilibrium payoff is bounded above by v_2. By Lemma 13, player 2's limit equilibrium payoff is bounded below by v_2. Therefore, her limit equilibrium payoff is v_2. □

Claim 6. $\mathbb{E}^{\theta_i}_{(\sigma_n)}\left(\delta_n^{\mathbb{T}}\right) \to 0$ for every $i > s$.

Proof. Claims 5 and 4 together imply that player 2's limit equilibrium payoff against each type $\theta \in \{\theta_{s+1}, ..., \theta_K\}$ is equal to $u_2(\alpha_i)$. This payoff is attainable only if $\mathbb{E}^{\theta_i}_{(\sigma_n)}\left(\delta_n^{\mathbb{T}}\right) \to 0$. □

Claim 7. For every $i > s$, and for every $\varepsilon > 0$, $\lim \mathbb{E}^{\theta_i}_{(\sigma_n)}\left((1-\delta_n)\sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\{d(\mu_t, \{\boldsymbol{\mu_i}\})<\varepsilon\}}\right) = 1$.

Proof. Follows from Claims 6 and 4, and Lemma 9. □

Proof of Theorem 1 for Case 1: In the auxiliary game the unique equilibrium outcome is that player 2 plays S with probability 1 against each type $\{\theta_0, ..., \theta_s\}$, and plays W against each type $\{\theta_{s+1}, ..., \theta_K\}$. Claims 4 and 6 establish that reduced equilibrium outcomes of the NE converge to the unique equilibrium reduced outcome fo the auxiliary game.. □

Proof of Theorem 2 for Case 1: The only claim in the Theorem for high-reputation priors is for types in $\{\theta_{s+1}, ..., \theta_K\}$. Claim 7 shows that the claim is true in this case. □

E.4.2. *Case 2: Inequality (4) holds* $\left(\sum_{i\in\{1,2,...,s\}} \mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} < \mu_0(\theta_0)\right)$.

Claim 8. $\mathbb{E}^{\theta_i}_{(\sigma_n)}\left(\delta_n^{\mathbb{T}}\right) \to 0$ for every $i > s$.

Proof. Follows from claim 3, item 2 and Lemma 20, item 1. □

Claim 9. For every $i > s$ and for every $\varepsilon > 0$, $\lim \mathbb{E}^{\theta_i}_{(\sigma_n)}\left((1-\delta_n)\sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\{d(\mu_t, \{\boldsymbol{\mu_i}\})<\varepsilon\}}\right) = 1$.

Proof. Follows from Lemma 20, item 1 and claim 8. □

Claim 10. There exists $\varepsilon^* > 0$ such that for every $\varepsilon \in (0, \varepsilon^*)$ and for every $i \in \{1, ..., s\}$:

1. $\lim \mathbb{E}^{\theta_i}_{(\sigma_n)}\left((1-\delta_n)\sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\left\{d\left(\mu_t, \{\tilde{\boldsymbol{\mu}}_i^1\}\right)<\varepsilon\right\}}\right) = \tilde{\sigma}_2(\alpha_i) = 1 - \frac{u_1(a_0)-u_1(\alpha_i)}{V_1-u_1(\alpha_i)}$.
2. $\lim \mathbb{E}^{\theta_i}_{(\sigma_n)}\left(\delta_n^{\mathbb{T}} \mathbb{I}_{\left\{d\left(\mu_{\mathbb{T}}, \{\tilde{\boldsymbol{\mu}}_i^1\}\right)<\varepsilon\right\}}\right) = 1 - \tilde{\sigma}_2(\alpha_i) = \frac{u_1(a_0)-u_1(\alpha_i)}{V_1-u_1(\alpha_i)}$.

Proof. Follows from Lemma 20, item 3 and Corollary 5. □

Claim 11. There exists $\varepsilon^* > 0$ such that for every $\varepsilon \in (0, \varepsilon^*)$:

1. $\lim \mathbb{E}^{\theta_0}_{(\sigma_n)} \left((1-\delta_n) \sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\{d(\mu_t, \boldsymbol{\mu_0})<\varepsilon, d(\sigma_1(h^t;\theta_0),a_0)<\varepsilon\}} \right) = \tilde{\sigma}_1(a_0)$.

2. $\lim \mathbb{E}^{\theta_0}_{(\sigma_n)} \left((1-\delta_n) \sum_{t<\mathbb{T}} \delta_n^t \mathbb{I}_{\left\{d\left(\mu_t, \mu_i^b\right)<\varepsilon, d(\sigma_1(h^t;\theta_0),\alpha_i)<\varepsilon\right\}} \right) = \tilde{\sigma}_1(\alpha_i)\tilde{\sigma}_2(\alpha_i)$ for every $i \in \{1, ..., s\}$.

3. $\lim \mathbb{E}^{\theta_0}_{(\sigma_n)} \left(\delta_n^{\mathbb{T}} \mathbb{I}_{\left\{d\left(\mu_{\mathbb{T}}, \mu_i^b\right)<\varepsilon\right\}} \right) = \tilde{\sigma}_1(\alpha_i)\left(1-\tilde{\sigma}_2(\alpha_i)\right)$ for every $i \in \{1, ..., s\}$.

Proof. Item 1 follows from Corollary 4 and Lemma 19. Items 2 and 3 follow from Corollary 4, Corollary 5 and Theorem 6. □

Proofs of Theorems 1 and 2 for Case 2: In the auxiliary game, in the unique equilibrium outcome, player 2 plays W with probability 1 against each type $\{\theta_{s+1}, ..., \theta_K\}$. This is also the case in the limit outcome behavior of the dynamic game. This follows from Claim 8. Moreover, the first item of Theorem 2 holds from Claim 9. In the unique equilibrium of the auxiliary game, type θ_0 randomizes among $\{a_0, \alpha_1, ..., \alpha_s\}$ with the distribution $\tilde{\sigma}_1(\alpha_i) = \left(\frac{1-\mu_i^*}{\mu_i^*}\right)\left(\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)}\right)$ for every $i \in \{1, .., s\}$, and chooses a_0 with the remaining probability. Moreover, if the choice is a_0, player 2 plays W, and if the choice is α_i, player 2 plays W with probability $\tilde{\sigma}_2(\alpha_i) = 1 - \frac{u_1(a_0)-u_1(\alpha_i)}{V_1-u_1(\alpha_i)}$. Claims 10 and 11 show items 2 and 3 of Theorem 2. These also imply that the discounted strategy profile distribution conditional on each type $i \in \{1, .., s\}$ converges to the strategy profile distribution against such a type in the equilibrium of the auxiliary game. For the normal type, the behavior convergence follows from Corollaries 4 and 5, and from Theorem 6, as shown in Claim 11. □

F. Contracting Games

We order types $i \in \{1, ..., s\}$ so that $j < s$ implies that $u_1(\alpha_j) > u_1(\alpha_{j+1})$. Recall also that $u_1(\alpha_i) \neq u_1(\alpha_j)$, for any $i \neq j \in \{1, .., s\}$.

F.1. Reputation Boundaries Let the reputation boundary be the following set:

$$\mu^b := \{\mu \in \Delta\Theta : \sum_{i \in \{1,2,...,s\}} \mu(\theta_i) \frac{1-\mu_i^*}{\mu_i^*} = \mu(\theta_0)\}.$$

$\mu \in \Delta\Theta$ is above the reputation boundary if $\mu \in \mu^{+b} := \{\mu \in \Delta\Theta : \sum_{i\in\{1,2,...,s\}} \mu(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} > \mu(\theta_0)\}$ and below the reputation boundary if $\mu \in \mu^{-b} := \{\mu \in \Delta\Theta : \sum_{i\in\{1,2,...,s\}} \mu(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} < \mu(\theta_0)\}$.

Recall that $\boldsymbol{\mu_i}$ is the distribution over types that puts probability 1 on type i, ${\mu_{\mathbf{i}}}^b = \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}$ is the distribution that puts probability μ_i^* on type i, and remaining probability on the normal type, θ_0. Recall that $\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}} := \{\mu \in \Delta\Theta : \mu(\theta_i) > 0, \mu(\theta_j) = 0 \text{ for all } j \neq \{0, i\}\}$. Finally, recall the coarsening procedure in Section D.

F.2. Auxiliary Results

Claim 12. For every $i > s$ and $\varepsilon > 0$, there exists $\delta^* \in (0,1)$ such that, for all $\delta > \delta^*$, and any NE, σ, we have $\mathbb{P}_{(\sigma)}\left(h^{t_\delta(\varepsilon)} : \mu\left(h^{t_\delta(\varepsilon)}\right)(\theta_i) > \varepsilon \text{ and } S \text{ has not been played}\right) < \varepsilon$.

Proof. Assume towards a contradiction that we can find $\varepsilon > 0$, $i \in \{s+1, ..., K\}$, a sequence of NE $\{\sigma_n\}$ associated with a sequence of games where $\delta_n \to 1$ such that

$$\lim \mathbb{P}_{(\sigma_n)}\left(h^{t_{\delta_n}(\varepsilon)} : \mu\left(h^{t_{\delta_n}(\varepsilon)}\right)(\theta_i) > \varepsilon,\ S \text{ has not been played}\right) \geq \varepsilon.$$

Thus, taking a subsequence if necessary, we find that there exists $\eta > 0$ such that

$$\lim \mathbb{P}_{(\sigma_n)}\begin{pmatrix} h^{t_\delta\left(\frac{\varepsilon}{2}\right)} : \mu\left(h^{t_\delta\left(\frac{\varepsilon}{2}\right)}\right)(\theta_i) > \eta,\ S \text{ has not been played}, \\ \mathbb{E}_{(\sigma_n, \mu_n)}\left(\delta_n^{\mathbb{T}-t_{\delta_n}(\eta)} \mid h^{t_\delta\left(\frac{\varepsilon}{2}\right)}\right) < 1-\eta \end{pmatrix} \geq \eta.$$

But notice that Lemma 9 implies the existence of a sequence of positive constants $\xi_n \downarrow 0$ such that

$$\lim \mathbb{P}_{(\sigma_n)}\begin{pmatrix} h^{t_\delta\left(\frac{\varepsilon}{2}\right)} : \mu\left(h^{t_\delta\left(\frac{\varepsilon}{2}\right)}\right)(\theta_i) > \eta,\ S \text{ has not been played}, \\ \mathbb{E}_{(\sigma_n, \mu_n)}\left(\delta_n^{\mathbb{T}-t_{\delta_n}(\eta)} \mid h^{t_\delta\left(\frac{\varepsilon}{2}\right)}\right) < 1-\eta, \sum_{j \neq 0, i} \mu\left(h^{t_\delta\left(\frac{\varepsilon}{2}\right)}\right)(\theta_j) < \xi_n \end{pmatrix} \geq \frac{\eta}{2}.$$

But notice also that S is a strict best response for player 2 conditional on any type $\theta \in \{\theta_0, \theta_i\}$, which implies that player 2 has a profitable deviation at the histories $h^{t_\delta\left(\frac{\varepsilon}{2}\right)}$ when n is sufficiently large. □

Corollary 6. *Take a sequence of Nash equilibria $\{\sigma_n\}$ of a sequence games where all parameters except the discount factor are fixed, and the discount factors converge to one. For every $i > s$, if $\mu_0(\theta_i) > 0$, then* $\lim \mathbb{E}^{\theta_i}_{(\sigma_n)}\left(\delta^{\mathbb{T}}\right) = 1$.

Proof. Follows from Claim 12, and Bayes' rule. □

Claim 13. For every $\varepsilon > 0$, there exists $\eta > 0$ such that in any NE, σ, of a game with a prior $\mu_0 \in \boldsymbol{\mu_0}^\eta$, we have $\mathbb{E}_{(\sigma)}\left(\delta^{\mathbb{T}}\right) > 1 - \varepsilon$.

Proof. Suppose to the contrary. First notice that $\mathbb{E}_{(\sigma)}\left(\delta^{\mathbb{T}}\right) < 1-\varepsilon$ implies that $\mathbb{E}^{\theta_0}_{(\sigma)}\left(\delta^{\mathbb{T}}\right) < 1 - \frac{\varepsilon}{2}$ for η small enough. Second, notice that S is a strict best response for player 2 conditional on type θ_0. Hence, player 2 can profitably deviate by playing S at h^0 when η is sufficiently small. □

Claim 14. Take $i \in \{1, ..., s\}$. For every $\varepsilon > 0$, there exists $\eta > 0$ and $\delta^* \in (0,1)$ such that, for all $\delta > \delta^*$, and any NE , σ, of a game with a prior $\mu_0 \in \{\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}} : \mu(\theta_i) < \mu_i^* - \varepsilon\}^\eta$ we have $\mathbb{E}_{(\sigma)}\left(\delta^{\mathbb{T}}\right) > 1 - \varepsilon$.

Proof. For every $m \in \mathbb{N}$, take a sequence $(\delta_n) \to 1$ and an associated sequence of equilibria $(\sigma_{n,m})$ in which the prior beliefs $\mu_{0,m} \in \{\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}} : \mu(\theta_i) < \mu_i^* - \varepsilon\}^{\frac{1}{m}}$ such that $\lim \mu_{0,m} = \mu_0$. Let $\Gamma_{m,n}$ be the generalized centralized play that implements $\sigma_{n,m}$. Using a diagonal argument, we can find a subsequence of equilibria $\left\{\sigma_{n(r),m(r)}\right\}_r$ (with $\lim_{r\to\infty}\min\{n(r), m(r)\} = \infty$) for which the sequence of corresponding generalized centralized plays converges (weakly) to a limit centralized play Γ. Consider simple coarsening of Γ to obtain a centralized play $\tilde{\Gamma}$ which is obtained by a simple martingale split of beliefs on $\{\boldsymbol{\mu_j}\}_{j=0,1,..,K}$.[21] Since $\tilde{\Gamma}$ is a centralized play, we have $\tilde{Z}(\boldsymbol{\mu_i}) = \alpha_i$, and Claim 13 and the assumption that $\mu_{0,m} \in \{\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}} : \mu(\theta_i) < \mu_i^* - \varepsilon\}^{\frac{1}{m}}$ imply $\tilde{Z}(\boldsymbol{\mu_0}) = \alpha_i$ whenever $\tilde{v}(\boldsymbol{\mu_0}) < 1$. To see why, notice that if we had $\tilde{v}(\boldsymbol{\mu_0}) < 1$ and $\tilde{Z}(\boldsymbol{\mu_0}) \neq \alpha_i$, then our learning lemma 9 would imply the existence of a constant $\varphi > 0$ such that for every $\epsilon > 0$ we could find $r^* \in \mathbb{N}$ such that, for every $r \geq r^*$ we would have

$$\mathbb{E}^{\theta_0}_{(\sigma_{n(r),m(r)})}\left((1 - \delta_{n(r)}) \sum_{t<\mathbb{T}} \delta^t_{n(r)} \mathbb{I}_{d(\mu_t, \boldsymbol{\mu_0})<\epsilon}\right) \geq \varphi,$$

contradicting Claim 13. Since $\tilde{\Gamma}$ is incentive compatible (see item i of Theorem 7), we have $(1 - \tilde{v}(\boldsymbol{\mu_0})) \geq (1 - \tilde{v}(\boldsymbol{\mu_j}))$. Next consider the mechanism-design problem of maximizing

$$\mu_0(\theta_0)(V_2(\theta_0) + Y_0(u_2(\alpha_i) - V_2(\theta_0))) + \mu_0(\theta_i)(V_2(\theta_i) + Y_i(u_2(\alpha_i) - V_2(\theta_i)))$$

$$\text{subject to: } 1 \geq Y_0 \geq Y_i \geq 0.$$

Clearly any solution of the problem above involves $Y_0 = Y_i = 0$, whenever $\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)} < \frac{\mu_i^* - \varepsilon}{1 - \mu_i^* + \varepsilon}$. Hence a continuity argument implies that if $\tilde{v}(\boldsymbol{\mu_0}) < 1$ then player 2 would have a profitable deviation by playing S at h^0 at the equilibrium $\left(\sigma_{n(r),m(r)}\right)$ for all large r. Therefore we have $\tilde{v}(\boldsymbol{\mu_0}) = 1$. Moreover, since $(1 - \tilde{v}(\boldsymbol{\mu_0})) \geq (1 - \tilde{v}(\boldsymbol{\mu_j}))$, we also have $\tilde{v}(\boldsymbol{\mu_i}) = 1$, which implies the claim. □

Recall the coarsening procedure from Section D. We will be exclusively using the coarsening procedure for $\kappa = 1$ to obtain from a centralized play Γ the 1-coarsening of Γ, which we denote by $\tilde{\Gamma}$.

[21] This is a different coarsening procedure than the one in Section D. In particular, the simple procedure we use here projects each belief μ to the extreme points of the belief simplex.

Lemma 21. *For any $\xi > 0$ there exists a $\delta^* < 1$ such that if $\delta > \delta^*$, then in any NE, σ, player 2 does not play S at any on path history h^t in which $\sum_{i\in\{1,2,...,s\}} \mu_t(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} >$ $\max\{(1+\xi)\mu_t(\theta_0), \xi\}$.*

Proof. Suppose the claim is not true. Then we can find a sequence if discount factors $(\delta_n) \to 1$ and an associated sequence of NE $\{\sigma_n\}$ with $\mu_{0,n}$ such that $\sum_{i\in\{1,2,...,s\}} \mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} >$ $\max\{(1+\xi)\mu_0(\theta_0), \xi\}$, and in which player 2 plays S at h^0. For each $m \in \mathbb{N}$, consider the sequence of deviation strategies $(\tilde{\sigma}_{2,n,m})_{n=1}^{\infty}$ that plays S at the first history h^t such that : i) $\sum_{i\in\{1,2,...,s\}} \mu(h^t)(\theta_i) \leq \frac{1}{m}$ or ii) $\sum_{i\in\{1,2,...,s\}} \mu(h^t)(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} \leq \mu(h^t)(\theta_0)$ and $\sum_{i>s} \mu(h^t)(\theta_i) \leq \frac{1}{m}$. Let $\tilde{\sigma}_{n,m}$ be the strategy profile that coincides with the sequence of NE at all components except player 2's strategy, which corresponds to the deviation strategies $(\tilde{\sigma}_{2,n,m})_{n=1}^{\infty}$. Let also $\Gamma_{n,m}$ be the associated generalized centralized play that implements $\tilde{\sigma}_{n,m}$. Notice that, for each $m\in \mathbb{N}$, we have

a) $\lim_{n\to\infty}\mathbb{E}^{\theta_i}_{(\tilde{\sigma}_{n,m})}\left(\delta_n^{\mathbb{T}}\right) = 1$ for every $i > s$.[22]

b) $\lim_{n\to\infty}\mathbb{E}_{(\tilde{\sigma}_{n,m})}\left((1-\delta_n)\sum_{t<\mathbb{T}}\delta_n^t\mathbb{I}_{\left\{\sum_{i\in\{1,2,...,s\}}\mu(h^t)(\theta_i)\frac{1-\mu_i^*}{\mu_i^*}\leq\mu(h^t)(\theta_0),\,\mu(h^t)(\theta_0)\geq\frac{1}{m}\right\}}\right) = 0$ for $i \in \{1, ..., s\}$.

Hence, using a diagonal argument we can find a subsequence $\left(\Gamma_{n(k),m(k)}\right)$ with $\lim_{k\to\infty}\min\{n(k), m(k)\} = \infty$ such that property a) holds and, for each $\varepsilon > 0$,

$$\lim_{k\to\infty}\mathbb{E}_{(\tilde{\sigma}_{n,m})}\left(\left(1-\delta_{n(k)}\right)\sum_{t<\mathbb{T}}\delta_{n(k)}^t\mathbb{I}_{\left\{\sum_{i\in\{1,2,...,s\}}\mu(h^t)(\theta_i)\frac{1-\mu_i^*}{\mu_i^*}\leq\mu(h^t)(\theta_0),\,\mu(h^t)(\theta_0)\geq\varepsilon\right\}}\right) = 0. \tag{38}$$

By Theorem 6, taking a subsequence (without relabelling) we may assume that $\left(\Gamma_{n(k),m(k)}\right)$ converges to a centralized play Γ. Let $\tilde{\Gamma}$ be the 1-coarsening of Γ. Note that along the sequence of strategy profiles, the probability that player 2 plays S at histories h^t with $\mu(h^t) \in \mu^{+b}$ converges to zero. To see why, notice that μ^{+b} is equal to the following countable union of open sets $\mu^{+b} = \cup_{l=1}^{\infty} B_l$, where

$$B_l := \left\{\mu \in \mu^{+b} : \sum_{i\in\{1,2,...,s\}} \mu(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} > \mu(\theta_0) + l^{-1}, \sum_{i\in\{1,2,...,s\}} \mu(\theta_i) > l^{-1}\right\}.$$

It is immediate to verify that $\lim_{k\to\infty}\mathbb{E}_{(\tilde{\sigma}_{n,m})}\left(\delta_n^{\mathbb{T}}\mathbb{I}_{\{\mu(h^{\mathbb{T}})\in B_l\}}\right) = 0$. Hence, applying The-

[22] This property follows from two observations: First, conditional on type θ_i, for every $\epsilon > 0$, the joint probability that S is not played by time $t_{\delta_n(\epsilon)}$ and the posterior belief at period $t_{\delta_n(\epsilon)}$ is outside of $\{\boldsymbol{\mu_0}\}^{\epsilon}$ converges to zero with n. Second, the property i) used in the stopping decision of the deviation strategies implies that S is played when the posterior belief reaches $\{\boldsymbol{\mu_0}\}^{\epsilon}$, for small ϵ.

orem 7 item ii, we obtain that $\int_{B_l} \lambda(\mu)v(\mu)\,[d\mu] = 0$, and hence $\int_{\mu^{+b}:} \lambda(\mu)v(\mu)\,[d\mu] = 0$ by monotone convergence. Because in the coarsening procedure, the only beliefs that are projected with positive probability to $\boldsymbol{\mu_i}$ for $i \in \{1, .., s\}$ are beliefs $\mu \in \mu^{+b}$, we get the following property of $\tilde{\Gamma}$:

Property 1: $\tilde{v}(\boldsymbol{\mu_i}) = 0$ for all $i \in \{1, ..., s\}$.

Likewise, along the sequence of strategy profiles $\{\tilde{\sigma}_{n,m}\}$, the probability that player 2 plays S at histories h^t with $\mu(h^t) \in \mu^{-b}$ converges to one. To see why, notice that μ^{-b} is equal to the following countable union of open sets $\cup_{l=1}^{\infty} D_l$, with

$$D_l := \left\{ \mu \in \mu^{-b} : \sum_{i \in \{1,2,\ldots,s\}} \mu(\theta_i) \frac{1 - \mu_i^*}{\mu_i^*} < \mu(\theta_0) - l^{-1}, \mu(\theta_0) > l^{-1} \right\}.$$

Using (38) and applying Theorem 7 item iii, we obtain $\int_{D_l:} \lambda(\mu)v(\mu)\,[d\mu] = \int_{D_l:} \lambda(\mu)\,[d\mu]$ and hence $\int_{\mu^{-b}} \lambda(\mu)v(\mu)\,[d\mu] = \int_{\mu^{-b}} \lambda(\mu)\,[d\mu]$ again by monotone convergence. Because in the coarsening procedure, the only beliefs that are projected with positive probability to $\boldsymbol{\mu_0}$ are those in μ^{-b}, we get that following property of $\tilde{\Gamma}$:

Property 2: $\tilde{v}(\boldsymbol{\mu_0}) = 1$.

Finally, because the sequence satisfies property a) above we immediately obtain

Property 3: $\tilde{v}(\boldsymbol{\mu_i}) = 1$ for all $i > s$.

Consider now the problem of calculating player 2's payoffs across all $1-$ coarse centralized plays that satisfy the 3 properties above. We will show that this payoff is strictly larger than the payoff player 2 would get if she played S at time 0.

First, we calculate player 2's payoff from playing S right away when $\sum_{i \in \{1,2,\ldots,s\}} \mu_0(\theta_i) \frac{1-\mu_i^*}{\mu_i^*} > \mu_0(\theta_0)$. Player 2's payoff from playing S at period 0 is:

$$v_2 := \sum_{\theta \in \Theta} \mu_0(\theta) V_2(\theta).$$

Recall that in any centralized play, for any $\mu \in \boldsymbol{\tilde{\mu}_i}$, $\tilde{Z}(\mu) = \alpha_i$. Moreover, $V_2(\boldsymbol{\tilde{\mu}_i^1}) = \mu_i^* V_2(\theta_i) + (1 - \mu_i^*) V_2(\theta_0)$, and by the definition of μ_i^*, $V_2(\boldsymbol{\tilde{\mu}_i^1}) = u_2(\alpha_i)$. Hence, player 2's payoff from a fixed $1-$ coarse centralized play satisfying the 3 properties above is:

$$v_2' := \sum_{i \in 1,2,\ldots s} \left(\tilde{\lambda}(\boldsymbol{\mu_i}) u_2(\alpha_i) + \tilde{\lambda}(\boldsymbol{\tilde{\mu}_i^1}) V_2(\boldsymbol{\tilde{\mu}_i^1}) \right) + \tilde{\lambda}(\boldsymbol{\mu_0}) V_2(\theta_0) + \sum_{i > s} \tilde{\lambda}(\boldsymbol{\mu_i}) V_2(\theta_i),$$

where we used property 1 and $V_2(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}) = u_2(\alpha_i)$ in the first line, property 2 in the second line, and property 3 in the third line.

First observe that for $i > s$, $\tilde{\lambda}(\boldsymbol{\mu_i}) = \mu_0(\theta_i)$. Second, observe that

$$\sum_{i\in 1,2,\ldots s}\left(\tilde{\lambda}(\boldsymbol{\mu_i})u_2(\alpha_i) + \tilde{\lambda}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}})V_2(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}})\right) + \tilde{\lambda}(\boldsymbol{\mu_0})V_2(\theta_0) =$$

$$\sum_{i\in 1,2,\ldots s}\left(\tilde{\lambda}(\boldsymbol{\mu_i})u_2(\alpha_i) + \mu_i^*\tilde{\lambda}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}})V_2(\theta_i)\right) + \tilde{\lambda}(\boldsymbol{\mu_0})V_2(\theta_0) + \sum_{i\in 1,2,\ldots s}\left(\tilde{\lambda}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}})(1-\mu_i^*)V_2(\theta_0)\right).$$

Observe that Bayes' rule implies that $\tilde{\lambda}(\boldsymbol{\mu_0})+\sum_{i\in 1,2,\ldots s}\left(\tilde{\lambda}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}})(1-\mu_i^*)\right) = \mu_0(\theta_0)$. Therefore,

$$v_2' = \sum_{\theta\in\{\theta_0,\theta_{s+1}\ldots,\theta_K\}}\mu_0(\theta)V_2(\theta) + \sum_{i\in 1,2,\ldots s}\left(\tilde{\lambda}(\boldsymbol{\mu_i})u_2(\alpha_i) + \mu_i^*\tilde{\lambda}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}})V_2(\theta_i)\right).$$

Note also that by Bayes' rule, $\tilde{\lambda}(\boldsymbol{\mu_i}) + \mu_i^*\tilde{\lambda}(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}) = \mu(\theta_i)$, and $u_2(\alpha_i) > V_2(\theta_i)$ for $i \in \{1,...,s\}$. Therefore, $v_2' \geq v_2$. Moreover, again by Bayes' rule, $\sum_{i\in\{1,2,\ldots,s\}}\mu_0(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} > \mu_0(\theta_0)$ implies that there exists $\epsilon > 0$ and an $i \in \{1,2,...,s\}$ with $\tilde{\lambda}(\boldsymbol{\mu_i}) > \epsilon$. Hence, $v_2' > v_2+\epsilon'$ for some ϵ' that does not depend on $\tilde{\Gamma}$. Because this is true for every $1-$coarse centralized play that satisfies the three properties above, we conclude that the deviation strategies give player 2 a strictly higher payoff than playing S right away, no matter what the normal type of Player 1 does. □

Since the set $\mu^{+b} := \{\mu \in \Delta\Theta : \sum_{i\in\{1,2,\ldots,s\}}\mu(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} > \mu(\theta_0)\}$ can be written as a countable union of open sets, $\mu^{+b} = \cup_{l=1}^{\infty}B_l$ with

$$B_l := \left\{\sum_{i\in\{1,2,\ldots,s\}}\mu_t(\theta_i)\frac{1-\mu_i^*}{\mu_i^*} > \max\left\{(1+l^{-1})\mu_t(\theta_0), l^{-1}\right\}\right\},$$

Lemma 21 and Theorem 7, item ii imply:

Corollary 7. *If $\Gamma = (\lambda, v, Z)$ is a centralized play that is a limit of a sequence of generalized centralized plays that implements a sequence of NE of a sequence of games with δ_n converging to 1, then $\int_{B_l}\lambda(\mu)v(\mu)d\mu = 0$ and hence by monotone convergence theorem*

$$\int_{\mu^{+b}}\lambda(\mu)v(\mu)\,[d\mu] = 0.$$

For the following development, we fix a sequence of games along which all parameters except for the discount factors are fixed, and the discount factors converge to 1. Take a sequence of NE of the sequence of games, and let $\{\Gamma_n\}$ be a sequence of generalized centralized plays where each element implements the corresponding Nash equilibrium of the corresponding game, and let $\Gamma = (\lambda, Z, v)$ be a limit centralized play that this sequence converges to weakly.

Lemma 22. $\int_{\mu^{+b}\setminus\cup_{i\in\{1,2,..,K\}}\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}} \lambda(\mu)\,[d\mu] = 0.$

Proof. Notice that

$$\begin{aligned}\int_{\mu^{+b}\setminus\cup_{i\in\{1,2,..,K\}}\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}} \lambda(\mu)\,[d\mu] &= \int_{\mu^{+b}\setminus\cup_{i\in\{1,2,..,K\}}\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}} (v(\mu) + (1 - v(\mu)))\,\lambda(\mu)\,[d\mu] \\ &= \int_{\mu^{+b}\setminus\cup_{i\in\{1,2,..,K\}}\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}} \lambda(\mu)v(\mu)\,[d\mu] = 0,\end{aligned}$$

where we used property 2. of centralized plays that asserts that $v(\mu) = 1$ for almost every posterior in $\Delta\Theta\setminus\left(\cup_{i=1}^{K}\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\right)\bigcup\{\boldsymbol{\mu_0}\}$ and the last equality follows from Corollary 7. $\square$

Lemma 23. *For every $i \in \{1, 2, ..., s\}$,*

$$\int_{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\cap\mu^{+b}} \lambda(\mu)\,(1 - v(\mu))\,\mathbb{I}_{\{Z(\mu)\,=\,\alpha_i\}}\,[d\mu] = \int_{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\cap\mu^{+b}} \lambda(\mu)\,[d\mu]\,.$$

Proof. We have

$$\int_{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\cap\mu^{+b}} \lambda(\mu)\,[d\mu] = \int_{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\cap\mu^{+b}} (v(\mu) + (1 - v(\mu)))\,\lambda(\mu)\,[d\mu] = \int_{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\cap\mu^{+b}} \lambda(\mu)\mathbb{I}_{\{Z(\mu)\,=\,\alpha_i\}}\,[d\mu]$$

where we first use Corollary 7 and then we use the property 3. of centralized plays that asserts $Z(\mu) = \alpha_i$ when $\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}$ and $v(\mu) < 1$. $\square$

Lemma 24. *For every $i \in \{1, 2, ..., s\}$,*

$$\int_{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\cap\mu^{-b}} \lambda(\mu)v(\mu)\,[d\mu] + \lambda\left(\{\boldsymbol{\mu_0}\}\right)v\left(\{\boldsymbol{\mu_0}\}\right) = \int_{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\cap\mu^{-b}} \lambda(\mu)\,[d\mu] + \lambda\left(\{\boldsymbol{\mu_0}\}\right).$$

Proof. Take $\varepsilon > 0$. Notice that for every $m \in \mathbb{N}$, we have

$$\lambda\left(\{\boldsymbol{\mu_0}\}\right)\left(1 - v\left(\{\boldsymbol{\mu_0}\}\right)\right) \leq \int_{\{\boldsymbol{\mu_0}\}^{m^{-1}}} \lambda(\mu)\,(1 - v(\mu))\,[d\mu]\,.$$

It follows from Claim 13 and the fact that Γ_n converges weakly to Γ that we can find $m^* \in \mathbb{N}$ such that $\int_{\{\boldsymbol{\mu_0}\}^{m^{-1}}} \lambda(\mu)\,(1 - v(\mu))\,[d\mu] < \varepsilon$, implying that $\lambda\left(\{\boldsymbol{\mu_0}\}\right)\left(1 - v\left(\{\boldsymbol{\mu_0}\}\right)\right) < \varepsilon$.

Take $\varepsilon > 0$. Next notice that by monotone convergence $\int_{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\cap\mu^{-b}} \lambda(\mu)\,(1 - v(\mu))\,[d\mu] = \lim_{m\to\infty}\int_{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}:\mu(\theta_i)<\mu_i^*-m^{-1}} \lambda(\mu)(1 - v(\mu))\,[d\mu]$. Take any $m \in \mathbb{N}$. It follows by Claim 14 and the fact that Γ_n converges weakly to Γ that we can find $\eta > 0$ such that

$$\int_{\{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}:\mu(\theta_i)<\mu_i^*-m^{-1}\}^\eta} \lambda(\mu)\left(1-v(\mu)\right)[d\mu] < \varepsilon,$$

which implies $\int_{\mu\in\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}:\mu(\theta_i)<\mu_i^*-m^{-1}} \lambda(\mu)\left(1-v(\mu)\right)[d\mu] < \varepsilon$. □

Consider now a 1-coarsening of Γ, which we call $\tilde{\Gamma}$. The following Lemma characterizes properties of $\tilde{\Gamma}$.

Lemma 25. *If $\tilde{\Gamma} = (\tilde{\lambda}, \tilde{v}, \tilde{Z})$ is a 1-coarsening of Γ, then:*

1. *For all $i \in \{0, s+1, ..., K\}$, $\tilde{v}(\boldsymbol{\mu_i}) = 1$ whenever $\tilde{\lambda}(\boldsymbol{\mu_i}) > 0$.*
2. *For all $i \in \{1, ..., s\}$, $\tilde{Z}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}}\right) = \alpha_i$ whenever $\tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}}\right) < 1$, and $\tilde{\lambda}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\boldsymbol{1}}\right) > 0$.*
3. *For all $i \in \{1, ..., s\}$, $\tilde{v}(\boldsymbol{\mu_i}) = 0$ and $\tilde{Z}(\boldsymbol{\mu_i}) = \alpha_i$ whenever $\tilde{\lambda}(\boldsymbol{\mu_i}) > 0$.*

Proof. We start with the first claim, for $i = 0$. Recall that in the coarsening procedure, the only beliefs μ that are projected to $\boldsymbol{\mu_0}$ with positive probability belong to the set μ^{-b}. Because Γ is a centralized play, Property 2 of the centralized plays imply

$$\int_{\mu\in\mu^{-b}\setminus\{\cup_{i\in\{1,...,K\}}\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\cup\{\boldsymbol{\mu_0}\}\}} \lambda(\mu)v(\mu)\,[d\mu] = \int_{\mu\in\mu^{-b}\setminus\{\cup_{i\in\{1,...,K\}}\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\cup\{\boldsymbol{\mu_0}\}\}} \lambda(\mu)\,[d\mu]\,.$$

Combining this with Lemma 24 and claim 12, we obtain that

$$\int_{\mu\in\mu^{-b}} \lambda(\mu)v(\mu)\,[d\mu] = \int_{\mu\in\mu^{-b}} \lambda(\mu)v(\mu)\,[d\mu]\,.$$

Therefore, $\tilde{v}(\boldsymbol{\mu_0}) = 1$ when $\tilde{\lambda}\left(\boldsymbol{\mu_0}\right) > 0$.

If $i > s$, then in the coarsening procedure, the only beliefs μ that are projected to $\boldsymbol{\mu_i}$ with positive probability belong to the set $\{\mu \in \Delta\Theta : \mu(\theta_i) > 0\}$. But claim 12 implies that

$$\int_{\{\mu:\mu(\theta_i)>0\}} \lambda(\mu)v(\mu)\,[d\mu] = \int_{\{\mu:\mu(\theta_i)>0\}} \lambda(\mu)\,[d\mu]\,.$$

Hence, $\tilde{v}(\boldsymbol{\mu_i}) = 1$ whenever $\tilde{\lambda}(\boldsymbol{\mu_i}) > 0$.

We now prove the second claim. This follows from $1-$ coarsening of a centralized play itself being a centralized play, given by Lemma 12, and the definition of a centralized play.

We now prove the third claim. Note that in the coarsening procedure, the only beliefs μ that are projected to $\boldsymbol{\mu_i}$ with positive probability belong to the set $\mu^{+b}\cap\{\mu \in \Delta\Theta : \mu(\theta_i) > 0\}$. Lemmata 22 and 23 imply that $\tilde{v}(\boldsymbol{\mu_i}) = 0$ and $\tilde{Z}(\boldsymbol{\mu_i}) = \alpha_i$ whenever $\tilde{\lambda}(\boldsymbol{\mu_i}) > 0$. □

Let

$$v_i := \left(1 - \tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right)\right) u_1(\alpha_i) + \tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) V_1 \geq V_1. \tag{39}$$

Then, using Lemma 25, we obtain the following IC constraints for $i \in \{1, 2, ..., s\}$ in $\tilde{\Gamma}$:

$$v \quad : \quad = \sum_{i\in\{1,2,...,s\}} \tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) v_i + \left(1 - \sum_{i\in\{1,2,...,s\}} \tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right)\right) V_1 \tag{40}$$

$$\geq \quad \tilde{\lambda}(\theta_i)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) v_i + (1 - \tilde{\lambda}(\theta_i)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right)) u_1(\alpha_i). \tag{41}$$

Lemma 26. *If $v > V_1$, then the following hold:*

1. *For every $j \in \{1, 2, ..., s\}$, $u_1(\alpha_j) > v$ and $\mu(\theta_j) > 0$ implies:*
   1.1. $\tilde{\lambda}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) > 0$.
   1.2. $v_j = v$.
   1.3. $\tilde{\lambda}(\boldsymbol{\mu_j}) = 0$.
2. $\sum_{i\in\{1,2,...,s\}} \tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) = 1$.
3. $\tilde{\lambda}(\boldsymbol{\mu_0}) = 0$.

Proof. On the way to a contradiction to (1.1), suppose that there is a $j \in \{1, 2, ..., s\}$ with $u_1(\alpha_j) > v$, and $\tilde{\lambda}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) = 0$. Then, $\tilde{\lambda}(\theta_j)(\boldsymbol{\mu_j}) = 1$. But then, the IC constraint for type j is violated in inequality (40), which is a contradiction.

To prove (1.2), first notice that the incentive-compatibility constraint (40) implies that $v_j \leq v$ for all $j \in \{1, 2, ..., s\}$ which satisfy $u_1(\alpha_j) > v$. This is because, the right hand side of inequality (40) is a weighted average of $u_1(\alpha_j)$ and v_j, with weights summing up to 1. Second, suppose on the way to a contradiction that $v_j < v$. Note that, v is a convex combination of terms v_i for which $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) > 0$ and the term $V_1 < v$. Then, there exists $v_i > v$ for some $i \neq j$ for which $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) > 0$. Moreover, $v_i > v$ implies $u_1(\alpha_i) > v$, because by equation (39), v_i is a convex combination of $u_1(\alpha_i)$ and V_1, and $V_1 < v$. Hence, there exists $i \in \{1, ..., s\}$ with $u_1(\alpha_i) > v$, and $v_i > v$, which contradicts the finding in the first sentence of this paragraph. Hence, $v_j = v$ for all $j \in \{1, 2, ..., s\}$ which satisfy $u_1(\alpha_j) > v$.

To prove (1.3), notice that $v = v_j$ together with the IC constraint (40) for type j implies $v_j \geq \tilde{\lambda}(\theta_j)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) v_j + (1 - \tilde{\lambda}(\theta_j)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right)) u_1(\alpha_j)$. Because $u_1(\alpha_j) > v = v_j$, we have implies $\tilde{\lambda}(\theta_j)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) = 1$, and $\tilde{\lambda}(\theta_j)(\boldsymbol{\mu_j}) = 0$. Bayes' rule then also implies $\tilde{\lambda}(\boldsymbol{\mu_j}) = 0$.

To prove (2), notice that if $u_1(\alpha_i) \leq v$, then $v_i \leq v$, because $v > V_1$. Because v is a weighted average of terms that are weakly below v, the terms that are strictly less than v must have a weight equal to zero. Hence, $1 - \sum_{i \in \{1,2,\ldots,s\}} \tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) = 0$.

Finally, (3) follows from (2) and Bayes' rule. □

Recall that we have ordered types $i \in \{1, \ldots, s\}$ so that $j < s$ implies that $u_1(\alpha_j) > u_1(\alpha_{j+1})$.

Corollary 8. *If $v > V_1$, then there is $j^* \in \{1, 2, \ldots, s\}$ such that $v_j = v$ and $\tilde{\lambda}(\boldsymbol{\mu_j}) = 0$ if $j < j^*$, and $\tilde{\lambda}(\theta_j)(\boldsymbol{\mu_j}) = 1$ if $j > j^* + 1$ and $j \leq s$.*

Proof. Clearly, $v \leq u_1(\alpha_1)$. First assume that $v = u_1(\alpha_1)$. In this case, (40) implies that $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\mathbf{1}}^{\mathbf{1}}\right) = 1$ and $\tilde{\lambda}(\theta_j)(\boldsymbol{\mu_j}) = 1$ for all $j > 1$, which implies the result. For the remainder of this proof assume that $v < u_1(\alpha_1)$. Hence the set $\mathfrak{J} := \{j \in \{1, 2, \ldots, s\} : v < u_1(\alpha_j)\}$ is nonempty.

If $\mathfrak{J} = \{1, 2, \ldots, s\}$ then Lemma 26 implies that $v_j = v$ and $\tilde{\lambda}(\boldsymbol{\mu_j}) = 0$ for every $j \in \mathfrak{J}$. We can thus set $j^* = s$, which implies the result.

If $\mathfrak{J} = \{1, 2, \ldots, k\}$ for some $k < s$, let $j^* = k + 1$ and hence Lemma 26 implies that $\tilde{\lambda}(\boldsymbol{\mu_j}) = 0$ for every $j \in \mathfrak{J}$ (hence for every $j < j^*$). Assume towards a contradiction that there is $r \in \{j^* + 1, \ldots, s\}$ such that $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{r}}^{\mathbf{1}}\right) > 0$. Then, v is a convex combinations of elements from $\{v_1, \ldots, v_{j*}\}$ and from a nonempty subset of $\{v_{j^*+1}, \ldots, v_s\}$. Lemma 26 implies that $v_i = v$ for all $i \in \{1, \ldots, j^* - 1\}$. Moreover, by definition of j^* we have $v \geq v_{j*}$ and hence the weighted average of v_i's over $\{v_1, \ldots, v_{j*}\}$ is no more than v. But since $v > u_1(\alpha_z)$ if $z > j^*$ and $z \leq s$, we conclude that the weighted average of v_i's over $\{v_{j^*+1}, \ldots, v_s\}$ is strictly less than v. This is a contradiction, because v cannot be a weighted average of terms less than or equal to v, where at least one term has a strictly positive weight, and is strictly less than v. □

Lemma 27. *If $v > V_1$ and if $\tilde{\lambda}(\boldsymbol{\mu_j}) = 0$ for some $j \in \{1, 2, \ldots, s\}$, then $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) = \frac{\mu(\theta_j)}{\mu(\theta_0)} \frac{1-\mu_j^*}{\mu_j^*}$. If $\tilde{\lambda}(\boldsymbol{\mu_j}) > 0$, then $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) < \frac{\mu(\theta_j)}{\mu(\theta_0)} \frac{1-\mu_j^*}{\mu_j^*}$.*

Proof. If $\tilde{\lambda}(\boldsymbol{\mu_j}) = 0$, then $\tilde{\lambda}(\theta_j)\left(\tilde{\boldsymbol{\mu}}_{\mathbf{1}}^{\mathbf{1}}\right) = 1$. By Bayes' rule, we have $\frac{1-\mu_j^*}{\mu_j^*} = \frac{\mu(\theta_0)}{\mu(\theta_j)} \frac{\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\mathbf{1}}^{\mathbf{1}}\right)}{\tilde{\lambda}(\theta_j)\left(\tilde{\boldsymbol{\mu}}_{\mathbf{1}}^{\mathbf{1}}\right)}$. Plugging in $\tilde{\lambda}(\theta_j)\left(\tilde{\boldsymbol{\mu}}_{\mathbf{1}}^{\mathbf{1}}\right) = 1$ and rearranging delivers the result. The second claim follows because if $\tilde{\lambda}(\boldsymbol{\mu_j}) > 0$, then $\tilde{\lambda}(\theta_j)\left(\tilde{\boldsymbol{\mu}}_{\mathbf{1}}^{\mathbf{1}}\right) < 1$, hence the equality $\frac{1-\mu_j^*}{\mu_j^*} = \frac{\mu(\theta_0)}{\mu(\theta_j)} \frac{\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\mathbf{1}}^{\mathbf{1}}\right)}{\tilde{\lambda}(\theta_j)\left(\tilde{\boldsymbol{\mu}}_{\mathbf{1}}^{\mathbf{1}}\right)}$ delivers the result. □

Lemma 28. *If $\mu_0 \in \mu^{-b}$, then $v = V_1$.*

Proof. Suppose on the way to a contradiction that $v > V_1$. First assume that $v = u_1(\alpha_1)$. Then, $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) = 0$ for all $j \geq 2$. By Lemma 27, $\tilde{\lambda}(\theta_0)(\boldsymbol{\mu}_0) = 1 - \tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) \geq 1 - \sum_{i\in\{1,2,\ldots,s\}} \frac{\mu_0(\theta_j)}{\mu_0(\theta_0)} \frac{1-\mu_j^*}{\mu_j^*} > 0$, because $\mu \in \mu^{-b}$. Thus,

$$v := \sum_{i\in\{1,2,\ldots,s\}} \tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) v_i + \left(1 - \sum_{i\in\{1,2,\ldots,s\}} \tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right)\right) V_1$$

implies that $v < u(\alpha_1)$, a contradiction. Next assume that $v < u(\alpha_1)$. Again by Lemma 27 we have$\tilde{\lambda}(\theta_0)(\boldsymbol{\mu}_0) = 1 - \sum_{i\in\{1,2,\ldots,s\}} \tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) \geq 1 - \sum_{i\in\{1,2,\ldots,s\}} \frac{\mu_0(\theta_j)}{\mu_0(\theta_0)} \frac{1-\mu_j^*}{\mu_j^*}$. If $\mu_0 \in \mu^{-b}$, then $\sum_{i\in\{1,2,\ldots,s\}} \frac{\mu_0(\theta_j)}{\mu_0(\theta_0)} \frac{1-\mu_j^*}{\mu_j^*} < 1$, hence $\tilde{\lambda}(\theta_0)(\boldsymbol{\mu_0}) > 0$. But this contradicts Lemma 26, item (3). Hence, $v \leq V_1$. Moreover, $v \geq V_1$, because each $v_i \geq V_1$. Therefore, $v = V_1$. □

Lemma 29. *If $\mu_0 \in \mu^{+b}$, then $v > V_1$.*

Proof. Note again that $v \geq V_1$. Suppose on the way to a contradiction that $v = V_1$. Then IC constraints 40 imply that $\tilde{\lambda}(\boldsymbol{\mu_i}) = 0$ for all $i \in \{1, 2, ..., s\}$. But then Lemma 27 implies that $\sum_{i\in\{1,2,\ldots,s\}} \tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) = \sum_{i\in\{1,2,\ldots,s\}} \frac{\mu_0(\theta_j)}{\mu_0(\theta_0)} \frac{1-\mu_j^*}{\mu_j^*} > 1$, because $\mu \in \mu^{+b}$, which contradicts that $\tilde{\lambda}(\theta_0)$ is a probability measure on $\Delta\Theta$. □

We define for $j \in \{1, ..., s\}$ (with the convention that if $j = 1$ then $\sum_{i\in\{1,2,\ldots,j-1\}} \mu_0(\theta_i) \frac{1-\mu_i^*}{\mu_i^*} = 0$)

$$\mu^{\mathbf{j}} := \left\{\mu \in \Theta : \sum_{i\in\{1,2,\ldots,j\}} \mu(\theta_i) \frac{1-\mu_i^*}{\mu_i^*} > \mu(\theta_0) > \sum_{i\in\{1,2,\ldots,j-1\}} \mu(\theta_i) \frac{1-\mu_i^*}{\mu_i^*}\right\}.$$

Lemma 30. *If $\mu_0 \in \mu^{\mathbf{j}}$ for some $j \in \{1, 2, ..., s\}$, then $v = u_1(\alpha_j)$, $\tilde{\lambda}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) > 0$, $\tilde{\lambda}(\boldsymbol{\mu_j}) > 0$ and $\tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) = 0$.*

Proof. Notice that if $\mu_0 \in \mu^{\mathbf{j}}$, then $\mu_0 \in \mu^{+b}$ and hence we have $v > V_1$ by Lemma 29. Therefore by corollary 8, there is a j^* such that $\tilde{\lambda}(\boldsymbol{\mu_i}) = 0$ for all $i < j^*$, and $\tilde{\lambda}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) = 0$ if $i > j^*$ and $i \leq s$. Hence, for all $i < j^*$, we have $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) = \frac{\mu_0(\theta_i)}{\mu_0(\theta_0)} \frac{1-\mu_i^*}{\mu_i^*}$, and for $i = j^*$, we have $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) \leq \frac{\mu_0(\theta_i)}{\mu_0(\theta_0)} \frac{1-\mu_i^*}{\mu_i^*}$, by Lemma 27. Thus, since $\sum_{i\in\{1,2,\ldots,s\}} \tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) = 1$, we conclude that $j = j^*$ and hence $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) = \frac{\mu_0(\theta_i)}{\mu_0(\theta_0)} \frac{1-\mu_i^*}{\mu_i^*}$ for all $i < j$. Therefore we have $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) = \frac{\tilde{\lambda}(\theta_j)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right)\mu_0(\theta_i)}{\mu_0(\theta_0)} \frac{1-\mu_j^*}{\mu_j^*}$ and $\mu_0 \in \mu^{\mathbf{j}}$ imply $\tilde{\lambda}(\boldsymbol{\mu_j}) > 0$ and $\tilde{\lambda}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) > 0$.

Now we show that $v = u_1(\alpha_j)$. First we note that because $j = j^*$, Lemma 26 implies that $v_1 = ... = v_{j-1} = v$. Moreover, v is a convex combination of v and v_j and $u_1(\alpha_j)$. Clearly $v_j \leq u_1(\alpha_j)$. If $v_j < u_1(\alpha_j)$, the fact that $\tilde{\lambda}(\boldsymbol{\mu_j}) > 0$ implies that the IC constraint

40 is violated. Therefore $v_j = u_1(\alpha_j)$. Then, v is a convex combination of v and v_j, which delivers $v = u_1(\alpha_j)$. Finally, the definition of v_j, $v_j = u_1(\alpha_j)$,$u_1(\alpha_j) > V_1$, and $\tilde{\lambda}(\boldsymbol{\mu_j}) > 0$ imply that $\tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) = 0$. □

Corollary 9. *If $\mu_0 \in \mu^{\mathbf{j}}$ for some $j \in \{1, 2, ..., s\}$, then for all $i \in \{1, ..., j-1\}$, $\tilde{\lambda}(\boldsymbol{\mu_i}) = 0$, and for all $i \in \{j+1, ..., s\}$, $\tilde{\lambda}(\boldsymbol{\mu_i}) = 1$.*

Proof. Follows from Corollary 8, Lemma 26 and Lemma 30. □

Lemma 31. *If $v = V_1$, then for every $i \in \{1, 2, ..., s\}$, $\tilde{\lambda}(\boldsymbol{\mu}_i) = 0$, and $v_i = V_1$, i.e., $\tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) = 1$.*

Proof. If $v = V_1$, then the IC constraint 40 for type i is satisfied only if $\tilde{\lambda}(\boldsymbol{\mu_i}) = 0$ and $\tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) = 1$, because $u_1(\alpha_i) > V_1$. □

F.3. Proofs of Theorems 1 and 3 for Contracting Games Lemma 32 below proves Theorem 1 for contracting games for the case in which $\mu_{n,0} = \mu_0 \in \mu^{-b}$.

Lemma 32. *For any sequence of NE $\{\sigma_n\}$ of a sequence of games where $\mu_{n,0} = \mu_0 \in \mu^{-b}$ for all n, and in the n^{th} game the discount factor is δ_n and $\delta_n \to 1$, we have $\lim \mathbb{E}_{(\sigma_n)}(\delta^{\mathbb{T}}) = 1$.*

Proof. Assume towards a contradiction that we can find a sequence of NE $\{\sigma_n\}$ such that (taking a subsequence if necessary) $\lim\mathbb{E}_{(\sigma_n)}(\delta^{\mathbb{T}}) < 1$. Claim 12 implies that for $i > s$, $\lim\mathbb{E}^{\theta_i}_{(\sigma_n)}(\delta^{\mathbb{T}}) = 1$. Hence, there is $i \in \{0, ..., s\}$ such that $\lim\mathbb{E}^{\theta_i}{}_{(\sigma_n)}(\delta^{\mathbb{T}}) < 1$. Now recall that Lemma 28 implies $v = V_1$, i.e., player 1's equilibrium payoffs converge to V_1. Hence, there is $i \in \{1, ..., s\}$ such that $\lim\mathbb{E}^{\theta_i}{}_{(\sigma_n)}(\delta^{\mathbb{T}}) < 1$. But then, player 1 can obtain a limit payoff strictly greater than V_1 by playing α_i at every period, which is a contradiction to his equilibrium payoff converging to V_1. □

Lemma 33. *Assume that $\mu_0 \in \mu^{\mathbf{j}}$ for some $j \in \{1, 2, ..., s\}$. We have:*

1. *For every $i > s$, $\mathbb{E}^{\theta_i}_{(\sigma_n)}\left(\delta^{\mathbb{T}}\right) \to 1$.*
2. *For every $i \in \{j, j+1, ..., s\}$, $\mathbb{E}^{\theta_i}_{(\sigma_n)}\left(\delta^{\mathbb{T}}\right) \to 0$.*
3. *For every $i \in \{1, ..., j-1\}$, $\mathbb{E}^{\theta_i}_{(\sigma_n)}\left(\delta^{\mathbb{T}}\right) \to \frac{u_1(\alpha_i) - u_1(\alpha_j)}{u_1(\alpha_i) - V_1}$.*
4. *$\mathbb{E}^{\theta_j}_{(\sigma_n)}\left(\delta^{\mathbb{T}}\right) \to 0$.*
5. *$\mathbb{E}^{\theta_0}_{(\sigma_n)}\left(\delta^{\mathbb{T}}\right) \to \sum_{i<j}\left(\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)}\frac{1-\mu_i^*}{\mu_i^*}\right)\left(\frac{u_1(\alpha_i) - u_1(\alpha_j)}{u_1(\alpha_i) - V_1}\right) + \left(1 - \sum_{i<j}\left(\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)}\frac{1-\mu_i^*}{\mu_i^*}\right)\right)$.*

Proof. Item 1: Follows from the first item of Lemma 25 and Corollary 25.

Item 2: Follows from the fact that $\tilde{\lambda}(\boldsymbol{\mu_i}) = 1$ for every $i \in \{j+1, ..., s\}$ from Corollary 9, the fact $\tilde{v}(\boldsymbol{\mu_i}) = 0$ derived in Lemma 25, and that $\tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) = 0$ from Lemma 30.

Item 3: First notice that Corollary 9 implies that $\tilde{\lambda}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) = 1$ for every $i \in \{1, ..., j-1\}$. Next notice that Lemma 30 implies that $v = u_1(\alpha_j)$. Hence,

$$u_1(\alpha_j) = \tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) V_1 + \left(1 - \tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right)\right) u_1(\alpha_i),$$

implying that $\tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) = \left(\frac{u_1(\alpha_i) - u_1(\alpha_j)}{u_1(\alpha_i) - V_1}\right)$, which leads to the result.

Item 4: Lemma 30 implies $\tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) = 0$, while Lemma 25 implies $\tilde{v}(\boldsymbol{\mu_i}) = 0$. Hence, the result follows.

Item 5: First notice that Lemma 27 implies that $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) = \frac{\mu_0(\theta_i)}{\mu_0(\theta_0)} \frac{1-\mu_i^*}{\mu_i^*}$ for every $i < j$, Lemma 26 implies that $\tilde{\lambda}(\theta_0)(\boldsymbol{\mu_0}) = 0$ and hence $\tilde{\lambda}(\theta_0)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) = \left(1 - \sum_{i<j}\left(\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)} \frac{1-\mu_i^*}{\mu_i^*}\right)\right)$. Since $\tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) = \left(\frac{u_1(\alpha_i) - u_1(\alpha_j)}{u_1(\alpha_i) - V_1}\right)$ for $i < j$ and $\tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) = 0$, the claim follows. □

Lemma 34. *Assume that $\mu_0 \in \mu^{\mathbf{j}}$ for some $j \in \{1, 2, ..., s\}$.*

1. *Take $i \in \{1, 2, ..., j-1\}$. For every $\epsilon > 0$,*

$$\mathbb{E}_{(\sigma_n)}^{\theta_i}\left(\sum_{t \geq 0}(1-\delta)\delta^t \mathbb{I}_{\{t<\mathbb{T}, \mu_t \in \mu_i^{b,\epsilon}\}}\right) \to \left(1 - \frac{u_1(\alpha_i) - u_1(\alpha_j)}{u_1(\alpha_i) - V_1}\right).$$

2. *For every $\epsilon > 0$, $\mathbb{E}_{(\sigma_n)}^{\theta_j}\left(\sum_{t\geq 0}(1-\delta)\delta^t \mathbb{I}_{\left\{t<\mathbb{T}, \left\{\mu_t \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{j}} : \mu_t(\theta_0) \leq 1-\mu_j^*\right\}^\epsilon\right\}}\right) \to 1$.*

3. *Take $i \in \{j+1, ..., s\}$. For every $\epsilon > 0$, $\mathbb{E}_{(\sigma_n)}^{\theta_i}\left(\sum_{t<\mathbb{T}}(1-\delta)\delta^t \mathbb{I}_{\{\mu_t \in \boldsymbol{\mu}_{\boldsymbol{i}}^{\boldsymbol{\epsilon}}\}}\right) \to 1$.*

4. *Take $i \in \{1, 2, ..., j-1\}$. There exists $\underline{\epsilon} > 0$ such that, for every $\epsilon \in (0, \underline{\epsilon})$,*

$$\mathbb{E}_{(\sigma_n)}^{\theta_0}\left(\sum_{t<\mathbb{T}}(1-\delta)\delta^t \mathbb{I}_{\{\mu_t \in \mu_i^{b,\epsilon}\}}\right) \to \left(\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)} \frac{1-\mu_i^*}{\mu_i^*}\right)\left(1 - \frac{u_1(\alpha_i) - u_1(\alpha_j)}{u_1(\alpha_i) - V_1}\right).$$

5. *Take $i = j$. There exists $\underline{\epsilon} > 0$ such that, for every $\epsilon \in (0, \underline{\epsilon})$,*

$$\mathbb{E}_{(\sigma_n)}^{\theta_0}\left(\sum_{t<\mathbb{T}}(1-\delta)\delta^t \mathbb{I}_{\{\mu_t \in \boldsymbol{\mu}_{\boldsymbol{i}} : \mu_t(\theta_i) \geq \mu_i^*\}^\epsilon}\right) \to \left(1 - \sum_{i<j}\left(\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)} \frac{1-\mu_i^*}{\mu_i^*}\right)\right).$$

Proof. Item 1: Corollary 9 implies $\tilde{\lambda}(\boldsymbol{\mu_i}) = 0$ for $i \in \{1, ..., j-1\}$. Hence, $\tilde{\lambda}(\theta_i)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right) = 1$. The coarsening procedure implies

$$\tilde{\lambda}(\theta_i)\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right)\left(1 - \tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right)\right) = \int_{\Delta\Theta} (1 - v(\mu))\lambda\left(\theta_i\right)(\mu)\left[d\mu\right].$$

Moreover, since Γ is a centralized play, we have

$$\int_{\Delta\Theta} (1 - v(\mu))\lambda\left(\theta_i\right)(\mu)\left[d\mu\right] = \int_{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}} (1 - v(\mu))\lambda\left(\theta_i\right)(\mu)\left[d\mu\right].$$

Finally, $\tilde{\lambda}(\boldsymbol{\mu_i}) = 0$ and $\tilde{\lambda}(\boldsymbol{\mu_0}) = 0$ from Lemma 26 imply that

$$\int_{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}} (1 - v(\mu))\lambda\left(\theta_i\right)(\mu)\left[d\mu\right] = \int_{\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\}} (1 - v(\mu))\lambda\left(\theta_i\right)(\mu)\left[d\mu\right].$$

Hence, we obtain

$$\left(1 - \tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right)\right) = \int_{\{\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\}} (1 - v(\mu))\lambda\left(\theta_i\right)(\mu)\left[d\mu\right].$$

A straightforward application of Portmanteau theorem (Billingsley (2013), Theorem 2.1) implies $\mathbb{E}_{(\sigma_n)}^{\theta_i}\left(\sum_{t\geq 0}(1-\delta)\delta^t \mathbb{I}_{\{t<\mathbb{T}, \mu_t \in \mu_i^{b,\epsilon}\}}\right) \to \left(1 - \tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}^{\mathbf{1}}\right)\right)$. The last term was calculated in item 3 from Lemma 33, implying the result.

Item 2: Notice that item 2 from Lemma 33 implies that $\tilde{v}\left(\tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}^{\mathbf{1}}\right) = \tilde{v}(\boldsymbol{\mu_j}) = 0$. Next we claim that

$$\int_{\left\{\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{j}} : \mu(\theta_0) > 1 - \mu_j^*\right\}} \lambda\left(\theta_j\right)(\mu)\left[d\mu\right] = 0.$$

Indeed, notice that our coarsening procedure projects posteriors in $\left\{\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{j}} : \mu\left(\theta_0\right) > 1 - \mu_j^*\right\}$ into $\boldsymbol{\mu_0}$ with a strictly positive probability. Because $\tilde{\lambda}(\boldsymbol{\mu_0}) = 0$ from Lemma 26, the claim is true. Moreover, since every centralized play has the property that for almost every posterior $\mu \in \Delta\Theta \backslash \left(\cup_{i=1}^{K} \tilde{\boldsymbol{\mu}}_{\boldsymbol{i}}\right) \bigcup \{\boldsymbol{\mu_0}\}$ with respect to the measure λ, we have $v(\mu) = 1$, and $\int_{\Delta\Theta \backslash \tilde{\boldsymbol{\mu}}_{\boldsymbol{j}}} \lambda\left(\theta_j\right)(\mu)\left[d\mu\right] = 0$. Therefore,

$$\int_{\left\{\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{j}} : \mu(\theta_0) \leq 1 - \mu_j^*\right\}} (1 - v(\mu))\lambda\left(\theta_j\right)(\mu)\left[d\mu\right] = \int_{\left\{\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{j}} : \mu(\theta_0) \leq 1 - \mu_j^*\right\}} \lambda\left(\theta_j\right)(\mu)\left[d\mu\right] = 1.$$

Hence, Portmanteau theorem implies

$$\liminf \mathbb{E}_{(\sigma_n)}^{\theta_j}\left(\sum_{t\geq 0}(1-\delta)\delta^t \mathbb{I}_{\left\{t<\mathbb{T}, \left\{\mu_t \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{j}} : \mu_t(\theta_0) \leq 1 - \mu_j^*\right\}^{\epsilon}\right\}}\right) = \int_{\left\{\mu \in \tilde{\boldsymbol{\mu}}_{\boldsymbol{j}} : \mu(\theta_0) \leq 1 - \mu_j^*\right\}} (1 - v(\mu))\lambda\left(\theta_j\right)(\mu)\left[d\mu\right] = 1.$$

Item 3: Notice that Corollary 9 implies $\tilde{\lambda}(\boldsymbol{\mu_i}) = 1$ for all $i \in \{j+1, ..., s\}$, and Lemma 25 implies $\tilde{v}(\boldsymbol{\mu_i}) = 0$. This implies that

$$\int_{\boldsymbol{\mu_i}} (1 - v(\mu))\lambda(\theta_i)(\mu)[d\mu] = 1,$$

in which case a straightforward application of Portmanteau theorem implies the result.

Items 4 and 5: Using Lemma 26, the arguments in items 1 and 2 above, and Bayes' rule we see that

$$\begin{gathered}\int_{\Delta\Theta}(1 - v(\mu))\lambda(\theta_0)(\mu)[d\mu] \\ = \sum_{i<j} \int_{\mu_\mathbf{i}{}^b}(1 - v(\mu))\lambda(\theta_0)(\mu)[d\mu] + \int_{\{\mu \in \boldsymbol{\mu_j}: \mu(\theta_j) \geq \mu_j^*\}}(1 - v(\mu))\lambda(\theta_0)(\mu)[d\mu],\end{gathered}$$

$$\int_{\{\mu \in \boldsymbol{\mu_i}: \mu(\theta_j) \geq \mu_j^*\}} (1 - v(\mu))\lambda(\theta_0)(\mu)[d\mu] = \left(1 - \sum_{i<j}\left(\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)}\frac{1 - \mu_i^*}{\mu_i^*}\right)\right),$$

and, for every $i < j$,

$$\int_{\mu_\mathbf{i}{}^b} (1 - v(\mu))\lambda(\theta_0)(\mu)[d\mu] = \left(\frac{\mu_0(\theta_i)}{\mu_0(\theta_0)}\frac{1 - \mu_i^*}{\mu_i^*}\right)\left(1 - \frac{u_1(\alpha_i) - u_1(\alpha_j)}{u_1(\alpha_i) - V_1}\right).$$

Take $\underline{\epsilon}$ so that every pair of sets in the collection $\left\{\mu_1^{b,\underline{\epsilon}}, ..., \mu_{j-1}^{b,\underline{\epsilon}}, \{\mu \in \boldsymbol{\mu_i} : \mu(\theta_i) \geq \mu_i^*\}^{\underline{\epsilon}}\right\}$ has empty intersection. Using these findings, for every $\epsilon < \underline{\epsilon}$, a straightforward application of Portmanteau theorem implies items 4 and 5. □

Proofs of Theorems 1 and 3 when $\mu_0 \in \mu^{+b}$***:*** Lemma 33 and Lemma 34 will be used to complete the proof of Theorems 1 and 3 for contracting games for the case in which $j \in \{1, 2, ..., s\}$. We start with Theorem 1. Lemma 33 establishes the the convergence of the equilibrium reduced outcomes of the dynamic games to the equilibrium reduced outcome of the auxiliary game for every commitment type, $\theta_i, i > 0$. For type θ_0, the result follows from properties 4 and 5 of Lemma 34 and Theorem 6.

Now we turn to Theorem 3. For item 1, assume as in the claim that $j < s$ and take $i \in \{j+1, ..., s\}$. Property 2 of Lemma 33 and property 3 of lemma 34 imply the result. For item 2, assume as in the claim that $j > 1$ and take $i \in \{1, ..., j-1\}$. The result follows from property 3 of Lemma 33 and property 1 of Lemma 34. Next notice that item 3 follows from property 4 of Lemma 34 and property 2 of Lemma 34. Consider now item 4. This follows from property 3 of Lemma 33, property 4 of Lemma 34, Bayes' rule and Theorem 6. Finally notice that item 5 follows from property 5 of Lemma 34 and Theorem 6. □