REPUTATION AND INFORMATION DESIGN

Laurent Mathevet — *University of Arizona*
David Pearce — *New York University*
Ennio Stacchetti — *New York University*

May 6, 2026

Abstract. Can reputation replace legal commitment for an institution making periodic announcements? Near the limiting case of ideal patience, results of Fudenberg and Levine (1992) imply a positive answer in *value* terms: the commitment value can be approximated if the long run institution can imitate a behavioral type whose reliability is close to that to which it would commit if it could. However, because little is known about equilibrium *behavior* in dynamic reputational models, the classical dynamic foundation for commitment in Bayesian persuasion is incomplete. Computational and analytic approaches are combined here to characterize equilibrium behavior in a dynamic reputational cheap talk model. Behavior depends upon which of three reputational regions pertains after a history of play. These characterizations hold even far from the patient limit. But combined with a novel method of calculating average discounted values, they allow us to show behavioral convergence toward the static Bayesian persuasion solution, as the discount factor approaches 1.

†Special thanks are due to Vishal Ashvinkumar, Rumen Kostadinov, Diego Cussen and Nilima Pisharody for exceptional research assistance, to Benjamin Bernard for serving as our discussant, and to the editor and three referees for their detailed and constructive feedback. We are also grateful to Dilip Abreu, Ben Brooks, Jeff Ely, Drew Fudenberg, Ed Green, Emir Kamenica, Elliot Lipnowski, Daniel Monte, Stephen Morris and Debraj Ray for their helpful comments and conversations, as well as to seminar audiences at the University of Arizona, Columbia University, Cornell University, the University of Edinburgh, the University of Illinois, London School of Economics, the University of Michigan, New York University, the University of North Carolina, Pennsylvania State University, Rutgers University, the University of Toronto, the University of California San Diego, and Washington University in St. Louis, as well as participants at the 2017 Latin American Meeting of the Econometric Society, the 2017 Becker Friedman Institute Economic Theory Conference, the 2020 North American Meeting of the Econometric Society, and the 2023 Northwestern conference honoring the memory of Nancy Schwartz.

Introduction

The publication of Kamenica and Gentzkow (2011) (KG, hereafter) saw an explosion of interest in Bayesian persuasion problems. In this static model, a sender commits to sending costless signals to a receiver as a function of information the sender will learn. In this way the sender may influence the receiver in her choice of action. The commitment assumption differentiates this "Bayesian persuasion" exercise from the cheap talk literature (Crawford and Sobel (1982)), allowing the sender to achieve a higher optimal value and theorists to provide powerful characterizations. A drawback is that commitment can be hard to justify in many situations, because of inherent interim opportunities to deviate.

Rayo and Segal (2010) proposed to employ Fudenberg and Levine (1989, 1992, FL hereafter) to justify the static payoff of a sender with ex-ante commitment as the equilibrium payoff of a long-run sender facing a sequence of short-run receivers in a dynamic reputational cheap talk game. Consider, for example, a car dealership selling a new car to a different buyer every period. In this dynamic setting, the reputational cost of dishonesty could potentially harm future sales and mitigate the short-term incentive to misrepresent the quality of a car. FL's results corroborate this notion, implying that a patient sender can secure an average discounted payoff virtually as high as his Bayesian persuasion payoff in any equilibrium. These reputational considerations are part of a broader effort to understand how repetition can substitute for the commitment assumption in Bayesian persuasion (see literature review).

Although the dynamic reputational model offers a classical foundation for Bayesian persuasion *payoffs*, the role of Bayesian persuasion *behavior* within this dynamic reputational narrative remains unclear. Indeed, FL obtain their results by using an imitation strategy that cannot be an equilibrium strategy in many games of interest,[1] and their convergence in value does not a priori imply convergence in behavior, as Bayesian persuasion payoffs could be obtained as average discounted payoffs without players ever adhering closely to Bayesian persuasion behavior. It is crucial that the informed player be able to imitate a behavioral type that randomizes, but little is known about the details of behavior in dynamic reputational models with that feature.[2] (And in standard repeated games, it is the need to randomize that often prevents commitment payoffs from being reached in equilibrium.[3]) Thus, we find ourselves in the unsatisfying and

[1]If imitating a mixing reputational type in every period were an equilibrium, there would be no way of distinguishing the rational type's behavior from the reputational type's, hence reputation would be constant and the sender always trusted. In this context, the rational type would lie with certainty when needed, contradicting said equilibrium behavior.

[2]Important exceptions are included in the literature review that concludes this Section.

[3]In the standard repeated game (without reputation) between a long-run sender and a sequence of short-run receivers, randomization by the sender requires the continuation values to be adjusted to maintain indifference, which can induce a per period cost that does not arise in the static problem with commitment. This cost is often so severe that there is no benefit at all to randomizing compared to telling the truth (see Fudenberg, Kreps and Maskin (1990) and Best and Quigley (2022)).

rather puzzling situation where the classical reputational model lays the foundation for Bayesian persuasion payoffs, but we do not know *how*.

In a dynamic, reputational, cheap talk game with binary actions and states and a single reputational type,[4,5] we study the equilibrium behavior that delivers the largest and the lowest equilibrium values, respectively, to the sender.[6] This analysis describes the dynamics of reputational management by the sender and the evolution of trust by the receiver. It also explains the mechanism by which both sender and receiver can approach their Bayesian persuasion payoffs for high discount factors, despite the sender's need to randomize between honest and dishonest reports.

For any reputation, one can characterize the first-period behavior of an optimal equilibrium starting at that reputation. This turns out to be critical for understanding the evolution of equilibrium behavior over time. This characterization partitions the reputation space into three regions, described by qualitatively uniform behavior within, but providing incentives in different ways across regions. In the very high reputation region, which we call Region 3, the sender capitalizes on his established reputation by lying with certainty, while the receiver trusts any recommendation. Outside that region, the sender faces the choice of investing or disinvesting in reputation, resolved stochastically. On the opposite end, in the low reputation region we call Region 1, the sender's probability of investing decreases with reputation, following a simple monotonic formula. If he lies, his reputation falls, and his continuation payoff will *not* be sender-optimal at that new reputation. Notice that this entails throwing away some surplus in order to punish the dishonest report.

An intermediate region, Region 2, features a notable phenomenon, as the sender's incentives are met by using *optimal* continuation values, whether he was honest or not. Although the sender must be indifferent in the low state between telling the truth or lying, the receiver does *not* randomize to produce that indifference: she always follows his recommendation. Instead, the sender's equilibrium probability of lying creates just enough spread in updated reputations—higher after truth-telling and lower after lying—that the sender-optimal continuation value at the higher reputation exactly compensates for that at the lower one, making him indifferent between reports.

In both regions 1 and 2, the sender is truthful with positive probability, and part of the incentive (in Region 2, all of it) comes from the higher reputation that follows honesty. It is natural to conjecture, then, that equilibrium value is everywhere increasing in reputation. This turns out to be hard to establish in general. Another natural conjecture is that the equilibrium rate of lying by the sender is also increasing in his reputation (this higher exploitation rate is an important part of his preference for having a high reputation). This too turns out to be hard to establish in general.

---

[4]**In Section 8, we consider multiple behavioral types**.

[5]The model captures a canonical yet stylized class of dynamic cheap talk games, such as a seller periodically selling goods of uncertain quality, an infectious disease expert advocating for healthy practices, or KG's celebrated prosecuting attorney model.

[6]This characterization holds for discount factors away from one.

We respond to this difficulty by providing computational support for monotonicity of value and exploitation in reputation.[7] By adapting the APS algorithm (Abreu, Pearce and Stacchetti (1990)) to the current problem, we provide numerical evidence, for a rich set of parameters, that monotonicity of value and exploitation does hold. Several of our lemmas and propositions state that *when* value and exploitation are monotonic in reputation, certain conclusions about behavior and rewards follow. These implications are all proved analytically, the conclusions applying whenever the maintained hypotheses about monotonicity are true. The numerical results assure that there are broad ranges of the exogenous parameters for which the maintained assumptions are indeed satisfied.

We prove *behavioral convergence results* that clarify the role of Bayesian persuasion behavior within the dynamic reputational narrative. If the sender's average discounted value and rate of lying are both increasing in current reputation, then the following holds: for any behavioral type, $\epsilon > 0$ and number T, there exists a discount factor above which, if the initial reputation (which is the receiver's initial belief that the sender is the behavioral type) lies in $[\epsilon, 1-\epsilon]$, then in any sender-optimal equilibrium, play by the sender and the receiver in the first T periods almost coincides with Bayesian persuasion behavior. Note that this is convergence at each reputational level as δ approaches 1, not convergence over time. This behavioral convergence result also holds approximately for values close to the upper frontier. One could be forgiven for guessing that behavior would converge to the reputational type that is available to imitate. But for high discount factors it approximates instead the Bayesian persuasion benchmark.

Obtaining approximately Bayesian persuasion behavior for a large T number of periods requires letting δ grow as much as needed. This raises the concern that although T is large, T periods might be insignificant in *discounted* terms.[8] Happily we can also offer a result in expected discounted terms. For high δ, we show that in all sender-preferred equilibria, sender and receiver behavior is close to that in the Bayesian persuasion benchmark for a high proportion of *δ-relevant time*. How close, and how high a proportion, depends on what behavioral types are available to imitate.

In the course of establishing behavioral convergence, we develop a novel technique called *stationary promising keeping*. When evaluating a player's average discounted payoff in a dynamic game, one can equivalently examine the payoff associated with any of the actions in the support of his current-period mixed strategy. The trouble is that this yields a weighted average of today's immediate payoff and a continuation value from tomorrow onward; often little is known about the latter. Instead, we focus on

[7] If the upper and lower boundaries of the value set are too close together, characterizations of the upper boundary are potentially highly complex, because the lower boundary can "get in the way" of punishing dishonest behavior most effectively. Our numerical evidence establishes that, for a rich set of parameters, the lower boundary is nowhere close to binding. This allows an analytical proof of the monotonicity of the upper boundary.

[8] We thank Emir Kamenica and a referee for encouraging us to think about this question.

what happens if the sender behaved as though he had a target reputation, always lying if his reputation is above the target and telling the truth otherwise. Even though this is not equilibrium behavior, it is optimal behavior and can be used to establish a close relationship between the sender's payoff and his equilibrium rate of exploitation. In this paper, it allows us to prove behavioral convergence and to get a quick understanding of many asymptotic phenomena. We hope that stationary promising keeping may prove useful in other dynamic settings.

The paper is organized as follows. A literature review completes this Section. Section 2 lays out the model. Section 3 presents the value convergence result. Section 4 introduces the essential incentive structure of perfect Bayesian equilibria of the model, and states the generalizations of self generation and the APS algorithm (Abreu, Pearce and Stacchetti (1990) that support our analysis. Using a robust computational result for the parameter ranges we consider, it also proves a monotonicity result for the upper frontier of the value correspondence: the best payoff increases with reputation. Section 5 then shows how initial behavior in any sender-optimal equilibrium varies across three regions of reputation space, and what this means about how equilibrium unfolds over time. Section 6 presents the behavioral convergence results, featuring a new technique we call stationary promise-keeping. The robustness of our computational results is discussed in Section 7. Section 8 allows for an arbitrary finite set of behavioral types. Higher-dimensional reputations present significant new challenges, but we are able to replicate the δ-relevant behavioral convergence result of Section 6. The Appendix gives details of the computational method and contains the proofs. The computer codes that generate our numerical results and figures can be found at https://tinyurl.com/yc6eantw.

Literature Review

Crawford and Sobel (1982) and Kamenica and Gentzkow (2011) are the seminal contributions on cheap talk and Bayesian persuasion (or information design), respectively. These two literatures differ in the sender's ability to commit to a communication protocol. Our work connects them via implicit enforcement. In an earlier related result, Aumann and Maschler (1995) studied optimal disclosure in infinitely repeated zero-sum games between two players maximizing their long-run average payoffs, when only one of them is informed about the state. More recent works include Hörner, Takahashi, and Vieille (2015), Ely (2017), Best and Quigley (2023), and Margaria and Smolin (2018). Our paper shares with Best and Quigley (2023) the goal of understanding how repetition can substitute for the commitment assumption in information design. They look for ways such as "review aggregation" to change the informational conditions and escape the negative implications of Fudenberg, Kreps and Maskin (1990). A well-calibrated review aggregator that creates long delays in reporting on the average veracity of a sender's messages, while keeping receivers (but not the sender) unclear about the timing of the next review, can approximate the sender's payoff with commitment. By contrast, our approach is to accept the existing informational limitations and explore how reputational mechanisms can overcome those limitations.

More broadly, our paper belongs to the classical dynamic reputation literature. In the chain-store model, employing deterministic behavioral types, Kreps and Wilson (1982) and Milgrom and Roberts (1982) explicitly characterize equilibria in which the long-run incumbent predates against early entrants, builds a reputation for toughness, and thereby deters subsequent entry. More generally, Kreps, Milgrom, Roberts and Wilson (1982) show how a small amount of incomplete information about player types can sustain long stretches of cooperative equilibrium play in the finitely repeated prisoners' dilemma.

On the communication side, Sobel (1985), Bénabou and Laroque (1992), and Morris (2001) study privately informed, long-lived senders who interact with sequences of myopic receivers. Sobel (1985) considers a sender who may be a "good" type that always tells the truth or a "bad" type whose preferences are strongly misaligned with the receiver's. The bad type initially builds a reputation for honesty and then times deception for maximal gain. Our work parallels the good type's commitment to relatively trustworthy communication, but allows randomization and an infinite horizon, which generate richer dynamics. Bénabou and Laroque (1992) extend Sobel's framework by modeling a financial journalist who each period receives imperfect information about a traded asset. His type, chosen once and for all, is either a rational, profit-maximizing type who is willing to misinform the market in order to create a trading opportunity for himself, or a compulsively honest reputational type. Because he observes an imperfect signal of the state, the sender can report dishonestly without entirely losing his reputation, something shared by our paper. Attention is limited to Markov perfect equilibrium, where behavior depends only on the sender's current reputation. A key distinction is that the journalist always has a myopic incentive to deceive, whereas in our class of games the sender wishes to inform the receiver correctly in the high state and incorrectly in the low state, which makes the analysis more complex.[9] Morris (2001) instead assumes that the good type shares the receiver's preferences, while the bad type prefers the same action regardless of the state; in equilibrium, the good types attempt to separate collapses into no information transmission. This logic echoes Ely and Välimäki (2003), where reputational concerns also undermine valuable information transmission.

Further papers derive explicit equilibrium dynamics in modified reputation environ-

[9]To formalize this point, we refer the interested reader to Section A.1 of the Appendix. Consider a function-to-function version of correspondence B defined in Section A.1: To any increasing, continuous, continuation payoff function $w : \beta \to [\underline{v}, \bar{v}]$, let b associate another payoff function $b(w) : \beta \to [\underline{v}, \bar{v}]$, which gives the (static Nash) equilibrium payoff of a game whose payoffs are the per-period payoffs augmented with w. Bénabou and Laroque (1992) show that in their model, b is a contraction, which ensures a unique MPE in the class of increasing, continuous solutions. In our model, b is not a contraction, and typically, no MPE exists. We study the set of Perfect Bayesian equilibria (PBE), giving special attention to the sender-optimal solutions.

ments. Mailath and Samuelson (2001) study seller reputation with occasional exit and inheritance of reputation; Phelan (2006) studies erratic betrayal and gradual rebuilding of public trust; Liu (2011) characterizes accumulation, consumption, and restoration of reputation when observing past behavior is costly; Liu and Skrzypacz (2014) analyze reputation bubbles when only limited records are available; and Ekmekci (2011) shows how finite rating systems can sustain payoffs close to the Stackelberg benchmark after every history. These papers modify the classical environment through type turnover, costly observation, limited records, or rating systems. By contrast, we keep the standard long-run/short-run reputation structure and characterize sender-optimal PBE behavior in a mixed behavioral-type sender-receiver game, and then connect that behavior to the relevant commitment benchmark through (pointwise and δ-relevant) behavioral convergence. Kostadinov and Kuvalekar (2022) use the APS algorithm not for computation, but to establish a three-region structure for relational contracts in a model in which firm and worker are learning about the quality of their match.

Additional recent work includes Pei (2023), who studies a patient sender with a privately known cost of lying and identifies conditions under which the sender can achieve the Bayesian persuasion payoff. His model coincides with ours with the major exception that the private information of the sender concerns his cost of lying. A similar result obtains in Pei (2021) in a buyer/seller model in which the seller has private information about his cost of effort. Like the current paper, Pei (2022) allows reputational types to play mixed actions, but the setting is one where the state is constant over time, as opposed to being drawn independently period by period. In reputation games with only pure-strategy commitment types, Li and Pei (2021) provide a full characterization of the long-lived player's discounted action frequencies. Fudenberg, Gao and Pei (2022) give positive results for a novel problem in which a long run sender is partially informed of which actions will be available to him. Before the uncertainty is fully resolved, he announces his intended action; he can build a reputation for carrying out his announced intentions, when those are feasible. At the limits of informational scarcity, Lipnowski and Ramos (2022) study a long-run sender whose messages' honesty is never ascertained, even stochastically, relating the dynamic equilibria to the outcomes of a particular money burning game.

Another strand of work shows that richer communication protocols can partially overcome the inefficiencies of Crawford and Sobel (1982). When the state is fixed once at the outset, long or structured communication can enlarge the equilibrium set or improve information transmission (Aumann and Hart, 2003; Krishna and Morgan, 2004; Goltsman, Hörner, Pavlov and Squintani, 2009; Forges and Koessler, 2008a,b). In fact, Golosov, Skreta, Tsyvinski and Wilson (2014) even obtain full revelation in a dynamic model where intertemporal incentives drive information release. Relatedly, in single-round expert advice with career concerns, Ottaviani and Sørensen (2006a,b) show that reputational motives typically lead to partial rather than full revelation.

Methodologically, our approach illustrates the complementary roles of theory and computation emphasized by Judd (1997). Computation provides an overview of properties away from the patient limit and a foundation for behavioral convergence, thus clearing the path for a study of equilibrium behavior in reputational models. Numerical methods have been crucial in dynamic stochastic environments where analytical solutions are intractable, such as DSGE models (e.g., Gourio, 2012, 2013; Andreasen, 2012; Isoré and Szczerbowicz, 2013, 2015; Petrosky-Nadeau, Zhang, and Kuehn, 2015), dynamic public finance (Golosov et al., 2007), and dynamic asset pricing (Borovicka and Stachurski, 2021). Phelan and Stacchetti (2001) adapt the APS algorithm to study a Ramsey tax model, and Renner and Scheidegger (2020) develop APS-based methods for computing solutions beyond the reach of standard techniques (see also Fernandes and Phelan, 2000). For a broader overview of computational approaches, including dynamic programming with multiple states, nonlinear DSGE models, and constrained environments, see Aldrich, Fernández-Villaverde, Gallant, and Rubio-Ramirez (2010), Fernández-Villaverde and Levintal (2018), and Fernández-Villaverde and Valencia (2018).

2. The Model

A long-lived sender communicates information to a sequence of short-lived receivers about an i.i.d state at discrete time periods $t \in \{0, \dots\}$. The sender can be either a rational type s_R or a behavioral type s_B (as described below), which is private information to the sender. The sender is behavioral with prior probability β_0 and rational with the residual probability.[10]

At the beginning of period t, state θ_t is drawn from $\Theta = \{\ell, h\}$ according to distribution $\mu_0 \in \Delta\Theta$. When there is no confusion, denote $\mu_0 = \mu_0(h)$. The sender observes the realized state and then sends a message $m \in M = \{L, H\}$ to the receiver. Denote the rational sender's strategy at t by $\pi_t : \Theta \to \Delta(M)$ where $\pi_t(\cdot|\theta)$ is the distribution of messages (π_t may also depend on the observed history of play before t). The behavioral type follows the same strategy $\pi_B : \Theta \to \Delta(M)$ in every period.

Knowing the history of state realizations and messages up to period $t-1$, the receiver enters period t with a belief $\beta_t \in [0, 1]$ that the sender is of type s_B, which represents the sender's reputation, and with prior belief μ_0 over the state. Upon receiving message m_t, the receiver updates her belief about the state to $\mathbb{P}[\theta \mid m_t, \beta_t, \pi_t]$. This updated belief depends both on the sender's reputation β_t, which affects the credibility of the message, and on the sender's strategy π_t, which is known in equilibrium. The receiver then chooses an action $a \in A = \{L, H\}$ and receives utility

$$u(\theta, a) = \begin{cases} 1 & \text{if} \quad (\theta, a) \in \{(h, H), (\ell, L)\} \\ 0 & \text{otherwise} \end{cases}.$$

[10]Starting from a reputation $\beta_0 \in [0, 1]$, it is unclear whether all other $\beta \in [0, 1]$ are reachable in some PBE. Therefore, we study a game where all $\beta \in [0, 1]$ will arise in all PBE. Specifically, before the game starts, Nature draws a number β_0 uniformly from $[0, 1]$ and then draws the sender's type such that he is rational with probability β_0. The realization β_0 is common knowledge among the sender and the receivers. The uniform distribution plays no role at all in the analysis.

In each period, the receiver adopts a myopic best response and selects her action to solve

$$\max_{a \in A} \sum_{\theta \in \Theta} \mathbb{P}[\,\theta \mid m_t, \beta_t, \pi_t]\, u(\theta, a).$$

This behavior can be justified in several ways. One interpretation is that receivers are short-lived agents, as a buyer making a purchase decision before being replaced in the subsequent period. Alternatively, anonymity in the large is a standard justification for such a behavior. Because any receiver among the public is anonymous from the view of the long-run sender, her individual actions are not detectable and hence she plays myopically.

Whereas the receiver wants her action to match the current state, the rational sender wants her to always choose the high action, as captured by the following sender's payoff

$$v(a) = \mathbb{1}\{a = H\}.$$

The sender maximizes her average discounted expected payoff, where $\delta \in (0, 1)$ is the discount factor.

Throughout the paper, we assume:

Assumption 1 [Low State Presumption]: $\mu_0 < 1/2$.

Assumption 2 [Bounded Discounting]: $\delta \geq 1/(1 + \mu_0)$.

Assumption 3 [Trusted Behavioral Type]: $\pi_B(H|h) = 1$ and $0 < \pi_B(H|\ell) < \frac{\mu_0}{1-\mu_0}$.

If, contrary to Assumption 1, μ_0 exceeded 1/2, there would be no need for persuasion. Assumption 2 assures that some degree of cooperation is possible even at reputation $\beta = 0$, allowing us to compare that to what is achieved when reputation can evolve. The last assumption states that the behavioral type always tells the truth in the high state. Consequently, for the remainder of the paper, we use the notation

$$\boxed{\pi_B := \pi_B(H|\ell)}$$

to denote the sole parameter of interest concerning the behavioral type's strategy: his probability of lying in the low state. This assumption also places a bound on this probability, implying that a receiver who knows that the sender is s_B would rationally follow his advice, because it corresponds to the most likely state.

Moves by the behavioral type are treated like moves of nature. Accordingly, if a message outside the support of the behavioral type is observed (that is, a report of low when the state is high), reputation is lost permanently, and $\beta = 0$ henceforth. Otherwise, beliefs can always be computed by Bayes' Rule. Thus, all versions of perfect Bayesian equilibrium coincide here: beliefs in any equilibrium are fully determined, and sequential rationality simply has the rational sender play a best response to those beliefs from each period forward.[11]

Let $V(\beta)$ be the set of PBE payoffs for the rational sender in the (continuation) game starting with receiver's belief β. Sometimes we write $V(\beta, \delta)$ to make explicit the dependence on the discount rate δ. A public randomization device ensures that $V(\beta)$ is a convex set: at the beginning of every period, before the state is realized, the players publicly observe the outcome of a draw from a uniform distribution in $[0, 1]$. (The realization of the public randomization device is observed by all current and future players.)

This model admits a range of interpretations as long as the receiver (a worker, buyer, citizen, etc.) finds it worthwhile to choose H (high effort, purchase, precautions) only in the high state, whereas the sender (an employer, seller, medical expert) would like her to choose H regardless of the state. For the sake of illustration, think of a car dealership selling a new car to a different buyer every period. The seller only gets paid when he sells a car, while a buyer prefers to acquire a vehicle of high quality rather than one of inferior quality.

3. Asymptotic Efficiency

This section presents the value convergence result of FL within the specific context of our dynamic sender-receiver game. Given our trusted type assumption, a sender who, in the static version of our game, could commit to following the behavioral type's strategy would achieve an expected value of

$$v^{\pi_B} = \mu_0 + (1 - \mu_0)\pi_B.$$

Suppose instead that the sender could commit to always recommending H in state h, while recommending H with the following probability in state ℓ

$$\boxed{\pi^* := \frac{\mu_0}{1 - \mu_0}}$$

(Note that π^* is well-defined under Assumption 1). Then, the receiver would be indifferent to accepting the H recommendation, and if he did, the sender's expected value in the static problem would become

$$v^* \ = \ \mu_0 + (1 - \mu_0)\pi^* \ = \ 2\mu_0.$$

This is the maximal value the sender can attain with commitment when $\mu_0 < 1/2$ (see Kamenica and Gentzkow (2011)).[12]

[11] At the beginning of each period, all messages and actions in earlier periods are common knowledge. So there is no loss of recursive structure and no complications in using the APS algorithm to compute the full PBE payoff correspondence allowing for unrestricted mixed strategies.

[12] The case $\mu_0 \geq 1/2$ is trivial, because $v^* = 1$.

Proposition 1. [Value Convergence] For each $\beta \in (0,1)$ and $\epsilon > 0$, there exists $\underline{\delta} < 1$ such that $V(\beta, \delta) \subseteq [v^{\pi_B} - \epsilon, v^* + \epsilon]$ for all $\delta \geq \underline{\delta}$.

This result is an immediate corollary of FL. In all PBE, the rational sender receives at least v^{π_B} and at most his Bayesian persuasion payoff, both approximately. When $\pi_B \in [\pi^* - \epsilon, \pi^*)$, a patient sender in the dynamic reputational game is virtually guaranteed, in all PBE, the same value as he would receive in the static Bayesian persuasion model.

The lower payoff bound in this proposition follows from the imitation strategy of FL.[13] In any PBE, the sender could deviate from his actual equilibrium strategy, and behave in the same way as the behavioral type by adopting strategy π_B in every period forever. This deviation would yield a patient sender a payoff of approximately v^{π_B}. Since no unilateral deviation can be strictly profitable in equilibrium, that payoff is a lower bound to the set $V(\beta, \delta)$. Although this imitation strategy is useful in establishing equilibrium payoff properties, it cannot be an equilibrium strategy (refer to Footnote 1 for further details).

In trying to understand what equilibrium behavior can support the conclusion of Proposition 1, one encounters a conundrum. The sender must sometimes recommend H in state ℓ to approach v^*, yet this recommendation must be probabilistic to maintain credibility. This requires an indifference condition in average discounted payoffs between recommending H and L—a condition that paradoxically impedes achieving v^* in the standard repeated games without reputational concerns (refer to Footnote 2 for further details)! Proposition 1 therefore raises further questions about what self-enforcing mechanisms underlie convergence in value. In any case, since v^* could emerge from average discounted payoffs without players ever adhering closely to Bayesian persuasion behavior, the proposition alone offers limited insight into equilibrium behavior.

In Section 6, we present a stronger value convergence result: holding π_B fixed, the value of the sender-optimal equilibrium converges pointwise, as $\delta \to 1$, to a particular value, independent of β. Once again, this alone will *not* imply convergence of behavior across the reputation space. One might expect that behavior converges to π_B. But it does not. Section 6 proves (Proposition 4) that it converges to the Bayesian persuasion behavior π^*.

4. Incentives and Computation

The tools of strategic dynamic programming, in particular self-generation, enable a characterization of equilibrium behavior as a function of the sender's reputation, even for discount factors away from the limit. One feature of our methodology is our application of the APS algorithm (Abreu, Pearce, and Stacchetti, 1990) to reputational

[13] The sender can only temporarily sustain average payoffs above her commitment payoff, as doing so necessarily deteriorates her reputation. For large δ, such episodes have negligible payoff consequences, hence the upper bound.

models, developed in Appendix A. The modified algorithm allows for detailed computational analyses to see what payoffs and behavior look like across diverse combinations of behavioral types and discount factors. We apply the algorithm over a wide grid of model parameters, presented in this Section, to examine properties such as the monotonicity of the equilibrium correspondence. These properties are important building blocks for our theoretical analysis.

4.1. Reputational Incentives

After any history of play leading to a current reputation β, a rational sender who observes today's state must decide whether to report truthfully or to lie. The optimal decision takes account of today's myopic payoff, how his possible reports will affect his reputation, and how those will affect his continuation payoffs from tomorrow onward. The relevant information is summarized in the tuple (π, α, w): π is the rational sender's current strategy, where $\pi(m|\theta)$ is the probability that he sends message m when today's state is θ; $\alpha = (\alpha_H, \alpha_L)$ is the the receiver's current strategy, where $\alpha_m(a)$ is the probability that she takes action a after receiving the message m; and $w = (w_{H,h}, w_{H,\ell}, w_{L,h}, w_{L,\ell})$ is a vector of (average discounted) continuation values. Given today's message m and state θ, tomorrow's reputation is given by

$$\beta_{m,\theta}(\beta, \pi) = \frac{\beta \pi_B(m|\theta)}{\beta \pi_B(m|\theta) + (1-\beta)\pi(m|\theta)}.$$

We often omit the variables (β, π) and simply write $\beta_{m,\theta}$. Since $\pi_B(H|h) = 1$, $\beta_{L,h} = 0$. Continuation values must be feasible, that is, $w_{m,\theta} \in V(\beta_{m,\theta})$ for each (m, θ). In addition, today's and tomorrow's values must provide incentives for the players to take the equilibrium actions. Formally

$$\pi(m|\theta) > 0 \quad \Rightarrow \quad m \in \operatorname*{argmax}_{m' \in M} \, (1-\delta) \sum_a \alpha_{m'}(a) v(a) + \delta w_{m',\theta} \tag{1}$$

$$\alpha_m(a) > 0 \quad \Rightarrow \quad a \in \operatorname*{argmax}_{a' \in A} \, \sum_\theta \mathbb{P}[\theta | m, \pi, \beta] u(\theta, a') \tag{2}$$

These conditions are reflected in the definition of admissibility.

Definition 1. The tuple (π, α, w) is admissible at β if for all $m \in M$, $a \in A$ and $\theta \in \Theta$, $w_{m,\theta} \in V(\beta_{m,\theta})$, $\pi(m|\theta)$ satisfies (1) and $\alpha_m(a)$ satisfies (2).

Technically, here we are talking about admissibility with respect to the PBE value correspondence V. In the Appendix, we will exploit the greater power of admissibility with respect to a general correspondence W (where continuation values can be drawn from W).

4.2 Monotone Optimal Value: Reputation as an Asset

At each reputation value, it is of particular interest to see how to optimize the sender's (average discounted) equilibrium payoff. Define

$$\overline{V}(\beta) = \max V(\beta) \quad \text{and} \quad \underline{V}(\beta) = \min V(\beta) \quad \text{for } \beta \in [0,1]$$

to be the upper and the lower boundaries of the equilibrium value correspondence.[14] The optimization exercise underlying $\overline{V}$ and $\underline{V}$ motivates the subject of optimal admissible tuples.

Definition 2. A tuple (π, α, w) is optimal at β if it is admissible at β and

$$\sum_\theta \mu_0(\theta) \sum_m \pi(m|\theta) \left[(1-\delta) \sum_a \alpha_m(a) v(a) + \delta w_{m,\theta} \right] = \overline{V}(\beta).$$

In optimal tuples, the sender's strategy satisfies several useful properties, established in Section 5.1. First, in the high state the sender always reports truthfully, i.e, $\pi(H|h) = 1$. Consequently, the sender's current strategy in an optimal tuple at reputation β is characterized by a single parameter, the probability the sender lies in the low state:

$$\boxed{\pi(\beta) := \pi(H|\ell)}$$

In Section 5 we provide conditions under which an optimal tuple is uniquely defined.

The receiver follows advice H whenever it conveys that $\theta = h$ with probability at least $1/2$, i.e., whenever $\mathbb{P}[h|H, \beta, \pi] \geq 1/2$. This holds if and only if $\pi(\beta) \leq \bar{\pi}(\beta)$, where

$$\boxed{\bar{\pi}(\beta) := \frac{\pi^* - \pi_B \beta}{1-\beta}}$$

In optimal tuples, this inequality always holds, so the sender's advice remains (at least weakly) trusted throughout. Finally, except at very high reputations β, the sender randomizes in state ℓ between honesty and lying. This requires indifference at $\theta = \ell$ between obtaining a payoff of 1 today with continuation value $w_{H,\ell} \in V(\beta_{H,\ell})$ and a payoff of 0 today with continuation value $w_{L,\ell} \in V(\beta_{L,\ell})$:

$$(1-\delta) + \delta w_{H,\ell} = \delta w_{L,\ell} \quad \Longleftrightarrow \quad w_{L,\ell} - w_{H,\ell} = \frac{1-\delta}{\delta}.$$

[14] Proposition 9 in Section A.1 ensures that $V(\beta)$ is a compact interval for each $\beta \in [0,1]$, so that $\overline{V}(\beta)$ and $\underline{V}(\beta)$ are well defined.

To ensure the sender is incentivized to report truthfully, particularly after observing $\theta = \ell$, it may be necessary to implement punitive measures for dishonesty that result in a continuation value $w_{H,\ell}$ strictly below $\overline{V}(\beta_{H,\ell})$. Even this, in fact, might not suffice to deter dishonest behavior adequately. In such situations, the temptation to cheat would need to be softened by reducing the receiver's likelihood of following advice below 1. However, this introduces significant inefficiencies, as it inadvertently punishes message H even when it is truthful!

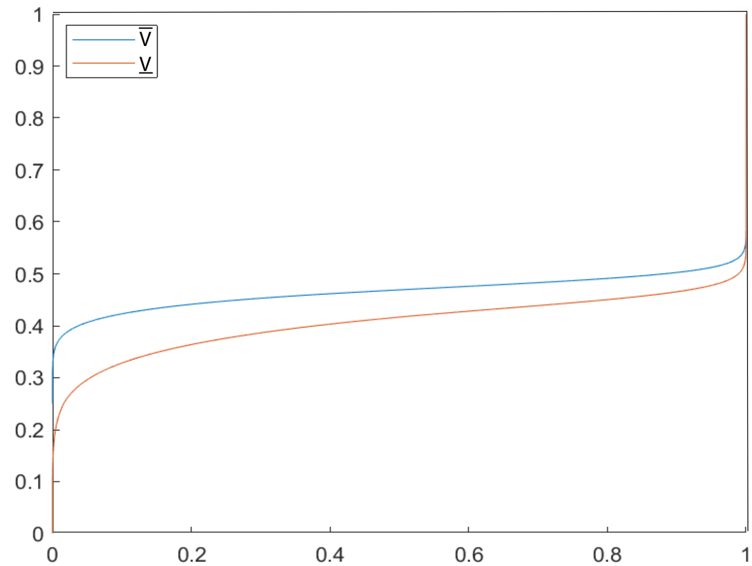

Figure 1: Value Correspondence for $\mu_0 = 0.25$, $\pi_B = 0.2$ and $\delta = 0.98$.

We investigate conditions under which these drastic measures can be avoided in optimal tuples (or in the first period of a sender-optimal PBE at β), relying on the adjustment of continuation values as the only means of providing the desired incentives. Suppose the probability of lying in state ℓ, at current β, is increased to $\bar{\pi}(\beta)$ where the receiver is indifferent to following advice. It can be shown that this will generate updated reputation $\lambda_0\beta$ or $\lambda_1\beta$, respectively, after a dishonest or honest statement, where

$$\lambda_0 = \frac{\pi_B(1-\mu_0)}{\mu_0} < 1 \quad \text{and} \quad \lambda_1 = \frac{(1-\pi_B)(1-\mu_0)}{1-2\mu_0} > 1.$$

The "gap condition" below guarantees that extreme measures, which would require the receiver to probabilistically ignore the sender's recommendations, are unnecessary in the first period of an optimal PBE. Importantly, when the gap condition is satisfied, it not only obviates the need for overly punitive measures, but it also enables us to prove that the sender-optimal equilibrium value is weakly increasing in reputation, so that reputation is an asset for the sender.

Definition 3. [Gap Condition] The equilibrium correspondence V satisfies the gap condition if $\overline{V}(\lambda_1\beta, \delta) - \underline{V}(\lambda_0\beta, \delta) \geq (1-\delta)/\delta$ for all $\beta \in [0, 1/\lambda_1]$.

Using the modified APS algorithm from Appendix A, we compute the value correspondence and associated behavior for all combinations of parameters on the following grid:

$$\text{Grid} = \Big\{(\mu_0, \delta, \pi_B): \quad \mu_0 \in \{.1, .15, \dots, .4\}, \delta \in \{.94, .945, \dots, .995, .999\}$$
$$\pi_B \in \{.06, .08 \dots, .66\} \text{ s.t. Assumptions 2-3}\Big\}.$$

The algorithm lets us look at the value correspondence for any desired combination of parameters, as illustrated in Figure 1. For all parameter combinations within the grid, the algorithm also determines that the gap condition holds, and notably, it does so while leaving untapped at least 35% of the punitive capacity.

Numerical Property 1. [Gap Condition holds] For all (μ_0, δ, π_B) on the Grid and $\beta \in [0, 1/\lambda_1]$, $\overline{V}(\lambda_1\beta, \delta) - \underline{V}(\lambda_0\beta, \delta) \geq (1-\delta)/\delta + 0.35(\overline{V}(\lambda_0\beta, \delta) - \underline{V}(\lambda_0\beta, \delta))$.

Proposition 2. [Monotonicity and Continuity] If V satisfies the gap condition, then $\overline{V}(\beta)$ is weakly increasing and continuous in β.

This result establishes a consistent relationship between maximal value and reputation, something that plays a central role in replacing legal commitment with reputational mechanisms. The result is proved analytically, but its relevance is supported numerically by Numerical Property 1, which shows that the key assumption (the gap condition) holds robustly across the parameter configurations on the Grid. In fact, our computations support strict monotonicity of $\overline{V}$, which ensures a unique optimal tuple for each β.

Numerical Property 2. [Strict Monotonicity] For all (μ_0, δ, π_B) on the Grid, $\overline{V}(\beta, \delta)$ is strictly increasing in reputation β.

What form PBEs take ultimately must depend on the parameters of the model. For example, as δ becomes increasingly small, all intertemporal incentives will eventually collapse. Before reaching that point, presumably drastic forms of punishment will have to be used to supplement more efficient ones. In Section 7, we examine which of the numerical properties established earlier fail when tested on an extended grid—leveraging all available computations, including extreme scenarios such as very low discount factors or behavioral types that seldom lie. Notably, the upper boundary of the value correspondence always remains strictly monotonic, even on the extended grid.

5. Equilibrium Behavior Across Reputation Space

This Section reveals the workings of reputation management for general discount factors. To understand how a sender-optimal equilibrium unfolds over time, it turns out to be crucial to know what optimal tuples look like in different parts of reputation space.[15] The space is divided into three regions with distinctly different characters. For reasons to become clear below, one can think of Regions 1, 2 and 3 (defined in the Introduction) respectively as inefficient provision of incentives for a degree of honesty, efficient provision of those incentives, and exploitation of reputation. Each of these is studied in subsection 5.1, while 5.2 considers tuples that *minimize* the sender's average discounted value (which we use later in the paper). Section 5 concludes by tracing out, in 5.3, the implications of results on optimal tuples for different reputational regions, for the evolution of a sender-optimal equilibrium. The challenging exercise of characterizing limiting behavior as the Sender approaches ideal patience is postponed until Section 6.

Before focusing on particular regions in the interior of reputation space, we discuss some basic properties of optimal tuples, and take care of two cases of reputational extremes.

5.1. Properties of Optimal Tuples

We characterize equilibrium behavior with initial value on the upper boundary by showing that the interior of reputation space can be organized into three regions. But first, consider two reputational extremes: $\beta = 0$ and $\beta = 1$.

When $\beta = 0$, it is common knowledge that the sender is not reputational, but rational. Here the dynamic game reduces to a standard game with one long run player interacting with a short run player each period. There is a babbling equilibrium in which the sender's messages are uninformative, and he is never trusted; its payoff to the sender is 0. Under our assumption on the discount factor, there is also a truthtelling equilibrium, in which the sender is trusted, but if he lies even once, he is never trusted again. This has value μ_0. A value recursion shows that there are no equilibria with higher values than this, no matter how high δ is (Lemma 4). Thus, $V(0) = [0, \mu_0]$. Recall that with the commitment power of Bayesian persuasion, the sender could attain expected payoff of $v^* = 2\mu_0$. This is a dramatic expression of the possible inefficiency of supporting randomization in repeated games (see Fudenberg, Kreps and Maskin, 1990).

If instead $\beta = 1$, the receiver is certain that the sender is behavioral. Since the behavioral sender lies infrequently enough to be trusted, any recommendation to play

[15] Recall that finding an optimal tuple at reputation β involves the same optimization problem as finding behavior and continuation values of a sender-optimal PBE: in either case, one is maximizing the same weighted sum while providing the needed incentives (for sender indifference if β is in Region 1 or 2) and using continuation values from the respective value sets depending on the realization of the state and what was reported by the sender. Thus, behavior in the first period of a sender-optimal PBE with initial reputation β is the same as behavior in the optimal tuple at that β. Both of these coincide with behavior in the best continuation PBE after, say, seven periods of play, if reputation at that point happened to be β. We give conditions under which such behavior is unique; when it is not, the *sets* of optimal behaviors coincide.

H is respected, so the sender, if in fact rational (contrary to the receiver's view), can achieve a payoff of 1 in every period by claiming the state is high, without ever losing reputation. So $V(1) = \{1\}$.

The next proposition summarizes the results of Lemmas 8 and 9 in the Appendix for optimal tuples (or more properly, their translations from general W correspondences with monotonic upper frontiers, to the equilibrium value correspondence V), together with related discussion from the remark on p. 48. We state the proposition formally and then give a relatively informal account of what it says about behavior.

Proposition 3. If $\overline{V}(\beta)$ is strictly increasing in β and the gap condition holds, then the optimal tuple (π, α, w) at β is unique and satisfies

1. $\pi(H|h) = 1$ and $\alpha_L(L) = 1$.
2. $\pi_B \leq \pi(H|\ell) \leq \bar{\pi}(\beta)$ and $\alpha_H(H) = 1$.

First, in the high state, the sender is always truthful, as recommending L instead would result not only in a zero payoff today but also annihilate reputation. This choice must be suboptimal, in line with our monotonicity theorem. It follows that the sender's strategy at reputation β is fully characterized by $\pi(\beta) := \pi(H|\ell)$, as discussed earlier. Moreover, since recommendation L is exclusively associated with state ℓ, it is rational to always trust it, hence $\alpha_L(L) = 1$.

Second, the sender lies more frequently than the behavioral type, $\pi(\beta) \geq \pi_B$, but not so much as to make the receiver unwilling to follow his advice about the state. An immediate implication of $\pi(\beta) \geq \pi_B$ is that, after observing state ℓ, a lie shifts reputation to the left, whereas an honest message shifts it to the right.

Under the maintained gap condition (which is easily satisfied across our entire parameter grid), optimality in fact requires that the receiver follow advice with probability one. The receiver is willing to follow the advice H because $\pi(\beta) \leq \bar{\pi}(\beta)$ implies that $\mathbb{P}[\, h \mid H, \beta, \pi] \geq 1/2$. Thus, she is either indifferent between following and not following H (when $\pi(\beta) = \bar{\pi}(\beta)$) or strictly prefers to follow it (when $\pi(\beta) < \bar{\pi}(\beta)$). In every optimal tuple, she will in fact always follow recommendation H with probability 1, a point we return to when discussing Region 1 below.

As we describe optimal tuples in each of the three reputational regions below, we take advantage of Numerical Property 2, the strict monotonicity of $\overline{V}$.

5.1.1. Region 3: Exploitation

At an extremely high reputation, the sender is no longer willing to invest in his reputation by reporting honestly. His continuation payoffs from lying or telling the truth are both so close to 1 that the difference between them cannot prevent him from grabbing the myopic reward from lying today. We say that such a reputation β is in Region 3.

More generally, if setting $\pi(\beta) = 1$ and using continuation values on the upper frontier leaves the sender at least weakly preferring to lie after observing ℓ, β is in Region 3. Recall that under Numerical Property 2, $\overline{V}(\beta)$ is strictly increasing in β. This ensures that Region 3, the simplest of our regions, is the interval $[\beta^{23}, 1]$ where

$$\beta^{23} := \min\Big\{\beta : \ (\pi, \alpha, w) \text{ is optimal at } \beta \text{ and } \pi(\beta) = 1\Big\}.$$

Since $\pi(\beta^{23}) = 1$, $\beta_{L,\ell}(\beta^{23}, \pi(\beta^{23})) = 1$ because only the behavioral type could ever report L, and hence $\overline{V}(\beta_{L,\ell}) = 1$. Let $V^{23} = \overline{V}(\beta^{23})$. By definition, at β^{23} the sender is indifferent between lying and telling the truth when the receiver expects him to lie for sure. Therefore

$$V^{23} = \mu_0[(1-\delta) + \delta V^{23}] + (1-\mu_0)\delta\overline{V}(1).$$

That is,

$$V^{23} = \frac{\delta - (2\delta - 1)\mu_0}{1 - \delta\mu_0},$$

and $\beta \geq \beta^{23}$ if and only if $\overline{V}(\beta) \geq V^{23}$.

5.1.2. Region 2: Efficient Incentives

For reputations $\beta < \beta^{23}$, where $\overline{V}(\beta) < V^{23}$, the sender must necessarily randomize. If instead he were supposed to send a pure message L after observing ℓ, deviating to the message H would prove him behavioral, giving him a reputation of 1 and continuation value of 1 tomorrow. This would be irresistible, contradicting the fact that L was supposed to be his best response.

The interval $(0, \beta^{23})$ is divided into regions we call Region 1 and Region 2, distinguished by the nature of punishments for lying. In an optimal tuple at reputation β, $w_{H,h} = \overline{V}(\beta)$. If the continuation values $w_{L,\ell}$ and $w_{H,\ell}$ are both on the upper frontier as well, we say β is in Region 2. When the sender lies in state ℓ, he loses reputation, but still receives the highest payoff available at his new lower reputation. In this sense, the punishments for lying are not wasteful: surplus is not thrown away by using continuation values below the upper frontier of V. Instead, as the sender sometimes lies and sometimes advises honestly, he is effectively passing surplus back and forth between his current and future selves, as he draws upon or adds to his stock of reputation. There is room for this mechanism in the current game, exactly because it is a dynamic game rather than a strictly repeated game (the class for which Fudenberg, Kreps and Maskin (1990) produced a uniform inefficiency result).

Think of designing an optimal tuple at β in Region 2. As you gradually increase $\pi(\beta)$, reputation after honesty increases and reputation after a lie decreases. Because $\overline{V}$ is increasing, this growing difference in posterior reputations means a greater penalty

$\delta[\overline{V}(\beta_{L,\ell}) - \overline{V}(\beta_{H,\ell})]$ in terms of continuation values tomorrow for lying today. Stop increasing when $\delta[\overline{V}(\beta_{L,\ell}) - \overline{V}(\beta_{H,\ell})] = (1-\delta)$ holds exactly; this makes the sender who observes ℓ indifferent between lying and telling the truth, and hence willing to randomize as necessary for equilibrium. Since the sender is indifferent, he expects the same total utility whether he lies or tells the truth at ℓ:

$$\overline{V}(\beta) = (1-\delta) + \delta[\mu_0 \overline{V}(\beta) + (1-\mu_0)\overline{V}(\beta_{H,\ell})] \qquad (LPK)$$
$$= (1-\delta)\mu_0 + \delta[\mu_0 \overline{V}(\beta) + (1-\mu_0)\overline{V}(\beta_{L,\ell})]. \qquad (RPK)$$

These equations are referred to respectively as left and right promise keeping.

What if, as you raise $\pi(\beta)$, it hits $\bar{\pi}(\beta)$ before the incentive constraint is satisfied? That is, raising $\pi(\beta)$ enough to satisfy $\delta[\overline{V}(\beta_{L,\ell}) - \overline{V}(\beta_{H,\ell})] = (1-\delta)$ makes the receiver unwilling to follow advice. Then keeping both continuation values on the upper frontier of V is incompatible with satisfying incentives for both the sender and the receiver, and we say β is in Region 1, which we discuss below. Let

$$\beta^{12} = \inf\big\{\beta:\ (\pi,\alpha,w) \text{ is an optimal tuple at } \beta \text{ and } \pi(\beta) < \min\{\bar{\pi}(\beta),1\}\big\}$$

be the infimum of Region 2.

In the interior of Region 2, $\pi(\beta) < \bar{\pi}(\beta)$, and the receiver strictly prefers to follow advice. So $\alpha_H(H) = 1$. Both behavioral and rational types of sender report truthfully when the state is h, and the receiver follows the advice.

Interestingly, since $\alpha_H(H) = 1$ in Region 2, α_H is *not* chosen to achieve the sender's indifference. Instead, the sender randomizes to keep *himself* indifferent.

5.1.3. Region 1: Inefficient Incentives

An optimal tuple delivers incentives to the sender as efficiently as possible. In Region 2, one can have advice always followed and the sender indifferent between honesty and lying in the low state, without using continuation values below the upper frontier of V. By definition, this is not possible in Region 1. Although $\pi(\beta) = \bar{\pi}(\beta)$ to maximize spread, the incentives created for honesty are insufficient. There are two ways to make lying less attractive to the sender: lowering the continuation value after a lie, or having the receiver follow a recommendation of H only probabilistically (thereby reducing the sender's immediate reward today from lying). Intuitively, the second method is the more wasteful: it is poorly targeted, punishing message H whether it is true or not. Lemma 13 proves that if $w_{H,\ell}$ can be dropped below the upper frontier to meet the incentives $(1-\delta) + \delta w_{H,\ell} = \delta\overline{V}(\beta_{L,\ell})$, without violating the lower boundary of V and without setting $\alpha_H(H) < 1$, then this is the optimal configuration. Reducing $\alpha_H(H)$ is a last resort.

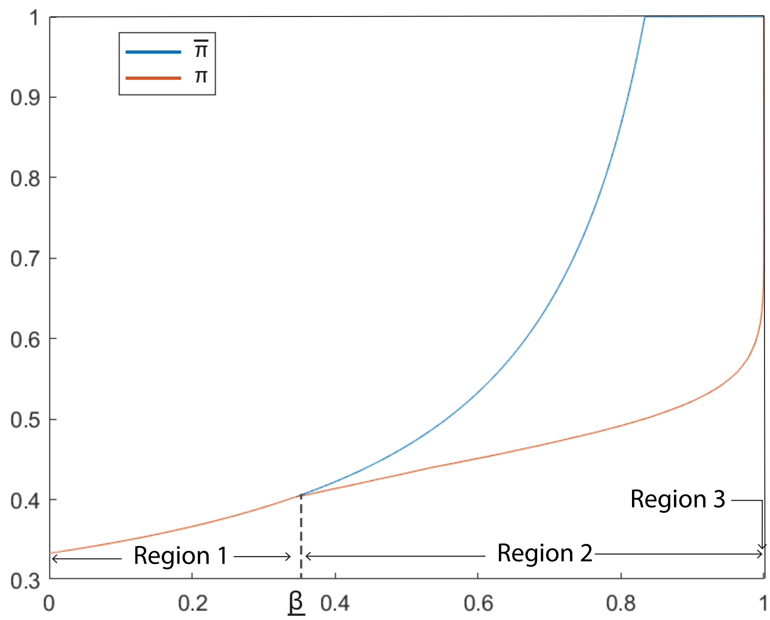

Figure 2: Sender's Equilibrium Strategy ($\mu_0 = 0.25$, $\pi_B = .2$ and $\delta = .98$)

Because of the gap condition from Subsection 4.2, which is satisfied effortlessly across our parameter grid (Numerical Property 1), it is never necessary to set $\alpha_H(H) < 1$ at any reputation β.

Nothing in the analytical results ensures that Regions 1 and 2 are intervals. Alternations between the two are conceivable. Interestingly, however, Numerical Property 4 in Section 7 indicates that, for parameter values on the Grid, each region is in fact an interval.

5.2. Lower Boundary Behavior

We make just a few remarks about the lower boundary, starting with the observation that $\underline{V}$ is weakly increasing in reputation β for all parameter combinations on the maintained Grid, as documented in Numerical Property 5 in Section 7.

The structure of the lower boundary is most transparent when one begins with the sender's behavior after observing the low state. Accordingly, we focus first on $\pi^l(\beta) := \pi^l(H \mid \ell)$, saving for the end of this subsection some remarks on an intriguing incentive issue after the high state.

As with the upper boundary, the lower boundary consists of three regions. If the sender's value $\underline{V}(\beta)$ exceeds the critical value V^{23}, the sender chooses $\pi^l(\beta) = 1$ and the receiver follows his advice with certainty. This is the analogue of the upper boundary's Region 3.

At lower β, $\underline{V}(\beta)$ requires the sender to randomize in the low state and two regions emerge, similarly to Section 5.1. Ideally, the receiver would ignore the sender's message, $\alpha_H(H) = 0$, eliminating any current-period reward. But this (often) requires receiver's indifference, hence $\pi^l(\beta) = \bar{\pi}(\beta)$, which need not minimize value. Indeed, a larger $\pi^l(\beta)$ increases $\beta_{L,\ell}$ and therefore raises $w_{L,\ell} = \underline{V}(\beta_{L,\ell})$, improving the sender's continuation payoff. The resolution of this tradeoff distinguishes Regions 1 and 2 for the lower boundary.

In Region 1, the receiver ignores the sender's message as much as the upper boundary permits, ideally setting $\alpha_H(H) = 0$, while $\pi^l(\beta) = \bar{\pi}(\beta)$ makes the receiver indifferent to doing so. Sender indifference is ensured through appropriate continuation payoffs, $w_{H,\ell} = w_{L,\ell} + (1-\delta)\alpha_H/\delta$. For some reputations, this is the best configuration to hurt the sender.

In Region 2, instead, the most punishing configuration is to have both $w_{H,\ell}$ and $w_{L,\ell}$ on the lower boundary, their difference providing the incentive for the sender to be weakly willing to report truthfully. In this region $\pi^l(\beta) < \bar{\pi}(\beta)$ and $\alpha_H(H) = 1$. This mirrors Region 2 on the upper boundary.

At first glance, it appears obvious that when the sender observes the high state, he will say so: otherwise the receiver will not take the high action today, and tomorrow the sender's reputation will be shot. But consider the following (far-fetched?) possibility. In a PBE in which the sender recommends L despite the state being high, the informational content of messages is diluted: an L report can now occur despite a high state, and an H report becomes less strongly associated with the high state. Since the credibility of message H is reduced, maintaining the receiver's indifference to follow H requires the sender to lie less frequently in the low state. This brings $\beta_{H,\ell}$ and $\beta_{L,\ell}$ closer together and permits equal continuation payoffs to be implemented at a lower level. This matters precisely for the Region 1 configuration, because, by lying in the high state, the sender can meet his incentive constraint in the low state through a harsher continuation value $w_{L,\ell}$.

We allow for this possibility in the algorithm computing the value correspondence. We also allow for configurations that implement the Region 1 regime ($\alpha_H(H) = 0$ and equal continuation values) without relying on indifference, by considering $\pi^l(\beta) > \bar{\pi}(\beta)$. On the maintained Grid, the set of reputations for which the sender lies with strictly positive probability in the high state forms a single connected interval, almost always beginning at the lowest reputation. Moreover, this interval shrinks from above as $\delta \to 1$ and becomes empty in the limit. By contrast, the tuple with $\pi^l(\beta) > \bar{\pi}(\beta)$ is never selected on the maintained Grid. It arises on the extended Grid only for parameter triples with $\mu_0 \in \{0.4, 0.45\}$ and $\delta \leq 0.935$, precisely in cases where $\underline{V}$ fails to be monotone increasing.

5.3 Equilibrium Path

Starting from some reputation β, let us follow a sender-optimal equilibrium across periods, for example starting on the upper boundary. Say β is in Region 2. If the state is h, both types of the sender reveal it by Lemma 8, the receiver chooses H by Lemma 7, and no updating occurs. If instead the state is ℓ, the sender randomly recommends H or L and his reputation declines or improves, respectively. Since the rational sender lies more frequently than does the behavioral type, there is an overall downward drift in reputation.[16] Once reputation enters Region 1 and the sender recommends action H in state ℓ, his continuation value leaves the upper boundary. Either this non-maximal continuation value is interior or on the lower boundary. If it is interior, there are many ways to deliver it: one of them is to do a public randomization and, depending on the result, follow either the upper boundary or the lower boundary equilibrium behavior.

For example, consider the realized path $(\beta_t, \overline{V}(\beta_t))$ in Figure 3, starting from the red point on the upper boundary. A truthful report in state ℓ shifts the point to the right, reflecting an increase in both the sender's reputation and continuation value. Suppose this is followed by five lies in state ℓ, each producing a leftward jump as both reputation and continuation value deteriorate. As long as reputation β_t has never visited Region 1, *all* sender-optimal equilibria coincide and prescribe the same behavior: both players act according to the optimal tuple at β_t, regardless of the history (and hence, of the initial reputation) that has led to β_t. The optimal tuple also pins down the continuation values, which remain on the upper boundary.

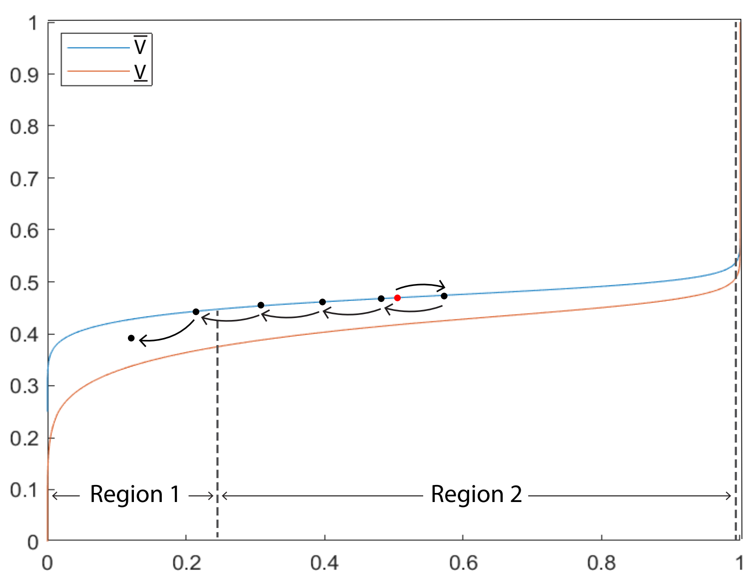

Figure 3: Equilibrium Path

This changes the first time β_t enters Region 1. If, at that point, the sender lies in state ℓ, the continuation value moves from the upper boundary into the interior of

[16]As the model meets the conditions of Cripps, Mailath and Samuelson (2004), the reputation β tends toward 0 in the long run.

the equilibrium value correspondence. There are now multiple ways (i.e., admissible tuples) to deliver the equilibrium continuation value: one has to waste a certain amount of surplus, and naturally there are many ways of doing that. This multiplicity gives rise to distinct sender-optimal equilibria, which all coincide prior to entering Region 1.

6. Behavioral Convergence

We establish a connection between sender-optimal equilibrium behavior in the limit as $\delta \to 1$ and the commitment solution in the Bayesian persuasion literature. Given the intense interest in static Bayesian persuasion and information design, it is encouraging to find theoretical support for its equilibrium behavior in dynamic environments without commitment.

Our results articulate this connection in two complementary ways. The first, developed in Sections 6.1 and 6.2, demonstrates what we term "pointwise behavioral convergence": for any fixed reputation, there exists a discount factor sufficiently close to one such that, in any sender-optimal dynamic equilibrium, both players behave at that reputation in a way that closely approximates the static Bayesian persuasion solution. This notion of convergence is the same as that used in the classical value convergence results of FL (Proposition 1), and is relevant for those interested in present behavior at a given reputation for high δ. As a corollary, starting from any non-extreme reputation and for any number T, for δ high enough both players adopt persuasion-like behavior for at least the next T periods. The argument behind these results relies on a novel way of evaluating the sender's equilibrium payoffs, which we term "stationary promise-keeping" and explain in 6.2.

The second connection, presented in 6.3, addresses the concern that the T Bayesian persuasion-like periods identified above may be negligible in discounted terms. We provide reassurance by showing that at high δ, in all sender-optimal equilibria, sender and receiver behavior is close to that in the Bayesian persuasion benchmark for a high proportion of average discounted time. We call this "δ-relevant behavioral convergence."

6.1. Pointwise Asymptotic Behavior

For all sender-optimal equilibria, initial behavior converges to the *same* limit as $\delta \to 1$, for all but a vanishing set of reputations. That limiting behavior coincides with the Bayesian persuasion solution defined in Section 3. A sequence of demanding lemmas (Lemmas 15-20 in the Appendix) supports these behavioral convergence results, reported in Proposition 4 and Corollary 1.

The result builds on the monotonicity of $\overline{V}$ (supported by Numerical Property 2) and that of the equilibrium exploitation rate $\pi(\beta)$. In Region 1, $\pi(\beta) = \bar{\pi}(\beta)$ is strictly increasing in β, while in Region 3 it has reached 1. Because numerical evidence casts doubt on monotonicity in Region 2 for low discount factors, we instead focus on *asymptotic monotonicity*: there is $\delta' \in (0,1)$ such that for all $\delta \geq \delta'$, $\pi(\beta, \delta)$ is weakly increasing in β. Numerical computations consistently support this property on the Grid:

Numerical Property 3. [Asymptotic Monotonicity] For all (μ_0, π_B) on the Grid, there is $\delta' \leq 0.999$ such that $\pi(\beta, \delta)$ is weakly increasing in β for all $\delta \geq \delta'$.

Proposition 4. [Pointwise Behavioral Convergence] Suppose $\overline{V}(\beta, \delta)$ is strictly increasing in β for every δ and that $\pi(\beta, \delta)$ is asymptotically increasing in β. Then, for every $\eta > 0$, there exists $\underline{\delta} \in (0,1)$ such that for all $\delta \geq \underline{\delta}$, $\pi^* \leq \pi(\beta, \delta) \leq \pi^* + 2\eta$ for all $\beta \in [0, 1-\eta]$ and $\alpha(\beta, \delta) = 1$ for all $\beta \in [2\eta/(\pi^* - \pi_B), 1-\eta]$.[17]

We see that for any fixed interior reputation β, for sufficiently high discount factors, optimal tuples at β specify sender and receiver behavior extremely close to Bayesian persuasion behavior. As Footnote 14 of Section 5 explains, we know that behavior in the first period of any sender-optimal equilibrium with initial reputation β agrees with the optimal tuple at that β. Hence, Proposition 4 is saying that every sender-optimal equilibrium begins with (virtually) Bayesian persuasion behavior.

Corollary 1 adds that for any interior initial β and positive integer T, the discount factor can be chosen large enough that in every sender-optimal equilibrium starting at β, behavior in the first T periods is virtually Bayesian persuasion behavior.

Corollary 1. Suppose $\overline{V}(\beta, \delta)$ is strictly increasing in β for every δ and $\pi(\beta, \delta)$ is asymptotically increasing in β. Then, for any $\epsilon > 0$ and number $T \in \mathbb{N}$, there exists $\underline{\delta} \in (0,1)$ such that for any $\delta \geq \underline{\delta}$, $\beta_0 \in [\epsilon, 1-\epsilon]$ and any realization of $\{\beta_t\}_{t=1}^{T}$ on the (stochastic) equilibrium path, $|\pi(\beta_t, \delta) - \pi^*| < \epsilon$ and $\alpha(\beta_t, \delta) = 1$ for all $t = 0, \ldots, T$.

We make several remarks. First, convergence in behavior is not an immediate consequence of convergence in value (Proposition 1 and FL). In the dynamic model, behavior can change substantially while payoffs remain constant. This is because of

[17] A referee asked how our results would be affected by the following change in the information environment: instead of the true state in period t being revealed publicly at the end of the period, at that time there is a public signal whose distribution depends nontrivially on the state in t and remains constant across periods. Continuation reputation and value would generally be different for each of the possible realizations, but there would still be a Region 3 in which the rational sender inevitably lies, a Region 2 in which all continuation values tomorrow are on the upper frontier of the value correspondence, and Region 1 where one or more continuation payoffs are inefficient. As the discount factor approaches 1, the same asymptotic values and behavior would apply as in the perfect monitoring case. The intuition for this rests on stationary promise-keeping (see Section 6.2). If the sender were to practice reputation maintenance somewhere in Region 2, in the long run the protocol being used is identified, and will look as much like the behavioral type as like the equilibrium type (in terms of sample likelihood). That is exactly the same as with perfect observation, and will generate the same long run rate of lying, and hence the same average payoff.

two features of behavior that are substitutes in value production: a large $\pi(\cdot)$ increases value by inducing more H, but a steep $\pi(\cdot)$ destroys value through adverse reputational dynamics (see Lemma 24 in the Appendix). At first glance, then, it would seem possible to sustain the roughly constant value of $\overline{V}$ that one sees in FL by having $\pi(\cdot)$ increase at the right rate. Proposition 4 shows, however, that any substantial increase in π leads to a feeding frenzy: more and more explosive increases in π are needed at successively higher levels of β, which is unsustainable.

Furthermore, Proposition 4 and its corollary hold for all $\pi_B < \pi^*$, as illustrated in Figure 3. Whereas convergence in value requires π_B to be close to π^*, the analogous result for behavior does not: even when π_B is much lower than π^*, behavior in the sender-optimal PBE converges. But it converges to π^*, not to π_B! And the corresponding value converges to neither the average payoff from always being trusted and cheating according to π^* (which is the Bayesian persuasion value v^*), nor to the average payoff from always being trusted and cheating according to π_B (defined as v^{π_B} in Section 3). Instead, it converges to the value $V^M(\pi^*, \pi_B)$, defined below in (3). We return to the issue of discounted relevant time in Section 6.3, where we prove a different behavioral convergence result.

Proposition 4 has further dramatic consequences: *Regions 1 and 3 vanish asymptotically.* If Region 1 did not vanish, then $\pi(\cdot)$ would be equal to $\bar{\pi}(\cdot)$ for reputations β bounded away from 0 and hence grow to values bounded away from π^*. This would contradict Proposition 4, which asserts that $\pi(\cdot)$ gets arbitrarily close to π^*. Moreover, $\pi(\beta)$ can equal 1 only if $\beta > 1-\epsilon$, hence Region 3 also vanishes asymptotically, leaving Region 2 and its efficient regime to fill the entire space.

6.2. Intuition and Stationary Promise-Keeping

Flatness. We first argue that for any $\pi_B < \pi^*$, $\overline{V}$ becomes almost flat over the entire reputation space. Think of the incentive needed to keep the sender indifferent in Region 1 or 2:

$$\overline{V}(\beta_{L,\ell}, \delta) - \overline{V}(\beta_{H,\ell}, \delta) \leq \frac{1-\delta}{\delta}.$$

The right side vanishes as $\delta \to 1$, because the average discounted value of an extra action H becomes negligible in the long run. Therefore, the vertical step sizes become minuscule while the horizontal step sizes $\beta_{L,\ell} - \beta_{H,\ell}$ are bounded away from zero when $\pi(\beta, \delta) \geq \pi^*$ for all β (except for β extremely close to 1 or 0 where a new message causes little updating). Both effects make for a *very* flat $\overline{V}$ function.

Lemma 1. Suppose that $\pi(\beta, \delta) \geq \pi^*$ for all β. Then, for any $\epsilon > 0$, there exists $\underline{\delta} \in (0, 1)$ such that $\overline{V}(1-\epsilon, \delta) - \overline{V}(\epsilon, \delta) \leq \epsilon$ for all $\delta \geq \underline{\delta}$.

Stationary Promise-Keeping. When $\delta \to 1$, $\overline{V}(\beta, \delta)$ converges to the same value for all $\beta \in (0, 1)$, a value which depends on π_B. What is that value? To answer, we propose a novel way of evaluating the sender's expected payoff at reputation β, called stationary promise keeping. In a sender-optimal equilibrium starting in Region 1 or 2, if the rational sender observes ℓ, both messages H and L are best responses for him. Hence, his expected utility can equivalently be evaluated by right or left promise keeping, defined in 5.1.2 above. These are value decompositions that break the payoff into a weighted average of today's return and the appropriate continuation values. If little is known about the latter, this does not say much about the average value. One could iterate left promise-keeping many times, progressively unpacking the continuation values, but this leads through a nonlinearly changing environment and is hard to evaluate.

Instead, starting at some β_0 in Region 2, the sender could maintain his reputation as steady as possible, always saying L if current reputation is strictly below β_0 and $\theta = \ell$, and always saying H otherwise. This "reputation maintenance" strategy keeps his reputation permanently within the interval $[\beta_H, \beta_L]$, where $\pi_0 = \pi(\beta_0)$, $\beta_H = \beta_{H,\ell}(\beta_0, \pi_0)$ and $\beta_L = \beta_{L,\ell}(\beta_0, \pi_0)$. As long as $[\beta_0, \beta_L]$ is included in Region 2, both messages are optimal, the continuation values remain on the upper boundary, and hence this reputation maintenance strategy has the same payoff as his equilibrium strategy.[18]

The idea now is to approximate this value by approximating the flow rate of actions H it can induce, which is what the sender cares about if he is patient. In the case where π is flat at some π_0 across the interval $[\beta_H, \beta_L]$, it is remarkably easy to find the proportion of H that is possible. Forget about following the details of the path that the sender's reputation will follow when he engages in reputation maintenance, and focus instead on long run likelihoods. The receiver entertains two hypotheses: she is watching a time series generated either by a behavioral type who uses π_B, or a rational type who uses π. Her posterior will be unaffected if the number of H messages makes the observed sequence of messages as likely under the one hypothesis as under the other. Lemmas 20 to 22 in the Appendix do that computation, which is approximately

$$V^M(\pi_0, \pi_B) = \mu_0 + (1 - \mu_0) \frac{\log\left[\frac{1-\pi_B}{1-\pi_0}\right]}{\log\left[\frac{\pi_0(1-\pi_B)}{(1-\pi_0)\pi_B}\right]} \qquad \text{for} \quad \pi_0 \in [\pi^*, 1).$$

At very low reputations β, $\pi(\beta) = \bar{\pi}(\beta)$ is close to π^*. Thus, for δ high enough that $\overline{V}$ is flat around β and that Region 2 begins at very low reputation levels, stationary promise-keeping suggests that (if reputation were maintained within Region 2) $\overline{V}(\beta, \delta)$ should be close to $V^M(\pi^*, \pi_B)$.

[18] Recall that the standard devices of left and right promise-keeping correctly evaluate the sender's equilibrium continuation value, even though neither conforms to equilibrium behavior. The same is true of stationary promise-keeping: holding the receiver's strategy fixed at its equilibrium specification, anything in the support of the sender's equilibrium distribution yields the correct equilibrium continuation value.

6.3. δ-Relevant Asymptotic Behavior

Even though behavior throughout Region 2 closely tracks the Bayesian persuasion solution when δ is large, the inexorable downward reputational drift caused by $\pi(\cdot) \geq \pi_B$, and demonstrated by Cripps, Mailath, and Samuelson (2004), means that the speed of reputational decline ultimately governs how quickly the sender reaches low payoffs. The lower π_B, the faster the decline and the greater the negative impact on the sender's average discounted payoffs throughout the reputation space.

The limitation of Proposition 4 and Corollary 1, therefore, is that although T may be chosen as large as desired in absolute terms, T periods might be an insignificant time span when the discount factor is so close to 1. One would like to know whether the time spent close to Bayesian persuasion behavior is significant in discounted terms. To address this question, consider a metric that measures the δ-relevant distance between a PBE (π, α) and the Bayesian persuasion benchmark:

$$D(\pi, \alpha) := \mathbb{E}\left[(1-\delta)\sum_{t\geq 0}\delta^t\Big(\mu_0\left|1-\alpha_t\pi_{t,h}\right| + (1-\mu_0)\left|\pi^* - \alpha_t\pi_{t,\ell}\right|\Big)\right].$$

In the Bayesian persuasion benchmark, the sender earns 1 in the high state and π^* on average in the low state, while the receiver follows all recommendations. Therefore, $D(\pi, \alpha)$ captures how close the equilibrium products $\alpha_t\pi_{t,h}$ and $\alpha_t\pi_{t,\ell}$ are, in δ-relevant time, to their benchmark levels 1 and π^*. If these products are close in δ-relevant time, then the underlying equilibrium behavior, namely $\alpha = (\alpha_t)$ and $\pi = (\pi_{t,\ell}, \pi_{t,h})$, must also be close to their Bayesian persuasion counterparts.

Proposition 5 shows that, for sufficiently patient senders, this distance can be uniformly bounded above, providing a measure of how closely equilibrium behavior tracks the Bayesian persuasion benchmark.

Proposition 5. [δ-Relevant Behavioral Convergence] Suppose $\overline{V}(\beta, \delta)$ is strictly increasing in β for every δ and that $\pi(\beta, \delta)$ is asymptotically increasing in β. For any $\epsilon > 0$, there exists $\underline{\delta} \in (0,1)$ such that for every $\delta \in [\underline{\delta}, 1)$ and every initial reputation $\beta_0 \in [\epsilon, 1-\epsilon]$, all sender-optimal equilibria (π, α) starting at β_0 satisfy $D(\pi, \alpha) \;\leq\; v^* - V^M(\pi^*, \pi_B) + 5\epsilon$.

This result about behavior does *not* follow easily from value convergence results as $\delta \to 1$ (Proposition 1). The proposition applies to *all* sender-optimal equilibria, and once these reach Region 1 for the first time, they display great variety: any continuation equilibrium after a lie in Region 1 requires some surplus to be wasted, and there are limitless ways to do so. Some PBE may wander back into higher-reputation territory, where the sender can be trusted even if his exploitation rate is much higher than the Bayesian persuasion threshold π^*. In such equilibria, the sender could in principle achieve a value close to the commitment benchmark v^* by sometimes lying with higher probability than π^*, and sometimes much lower. One needs to prove this does not happen.

The key technical input is Lemma 27 in the Appendix. It shows that, when δ is close to 1, the expected discounted mass of histories in which both reputation is high and the lying rate exceeds the threshold π^* is of order $1-\delta$, uniformly across equilibria. The key step is to recast Bayesian updating as a likelihood-ratio process and construct a supermartingale for reputation. As a result, whenever the sender exploits too aggressively at high reputation, reputation declines sufficiently quickly to prevent such episodes from persisting or having a meaningful impact.

This mechanism is related to the impermanence force identified by Cripps, Mailath, and Samuelson (2004). In their setting, for a fixed prior and discount factor, posterior beliefs on behavioral types converge to zero as $t \to \infty$, so that reputation eventually vanishes. We rely on the same Bayesian learning force—persistent statistical separation between rational and commitment behavior erodes reputation—but use it differently. Rather than focusing on the calendar-time tail, we study how much discounted time the process spends in different regions of the reputation space.

Lemma 27 therefore implies that, for δ close to 1, in every PBE almost all δ-relevant time is spent either near reputation 0, where lying significantly more than π^* yields no payoff, or at higher reputations with a lying rate close to π^* or below. From this perspective, impermanence does not preclude discipline: even though reputation eventually vanishes in calendar time, sender-optimal equilibria remain close to the Bayesian persuasion benchmark for most of the δ-relevant horizon when δ is high.

To illustrate the result, take $\mu_0 = 1/4$ and suppose ϵ is negligible. When $\pi_B = .3$, we have $v^* - V^M = .0125$. Thus, the δ-relevant average deviation of the equilibrium products from their Bayesian persuasion levels cannot exceed 1.25%. If the low state alone were responsible for this deviation, $\alpha_t \pi_{t,\ell}$ would be within $.0125/0.75 \approx 1.7\%$ of its KG benchmark in δ-relevant time.

Considering Propositions 4 and 5 together reveals two complementary perspectives on sender-optimal equilibrium behavior. Proposition 4 says that for high discount factors, every sender-optimal equilibrium virtually agrees with Bayesian persuasion behavior in the first period, and indeed for many periods. But it says nothing about the δ-relevant medium run. Proposition 5 says nothing about how a sender-optimal equilibrium starts, but for sufficiently high δ, shows that all sender-optimal equilibria agree closely with Bayesian persuasion behavior virtually all of the δ-relevant time.

7. Robustness and Extension of Numerical Properties

Several arguments in the paper rely on monotonicity properties of equilibrium objects. Section 4 establishes these properties numerically on a rich parameter Grid. In this section we ask how far these conclusions survive when the environment is pushed toward less favorable conditions for the sender.

To that end, we enlarge the parameter domain and consider the following grid:

$$\text{Extended Grid} = \Big\{(\mu_0, \delta, \pi_B): \quad \mu_0 \in \{.1, .15, \dots, .45\}, \delta \in \{.9, .905, \dots, .995, .999\}$$
$$\pi_B \in \{.02, .04, \dots, .8\} \text{ s.t. Assumptions 2–3}\Big\}.$$

Relative to the maintained Grid, this extension lowers discounting, admits smaller behavioral types, and includes $\mu_0 = 0.45$. We also address potential numerical precision concerns that arise in extreme parametrizations.

7.1 Robustness

On the maintained Grid we established three baseline numerical properties:

Numerical Property 1. The gap condition holds.
Numerical Property 2. $\overline{V}$ is increasing in reputation.
Numerical Property 3. π is asymptotically increasing in reputation.

On the Extended Grid the gap condition continues to hold throughout the parameter region satisfying bounded discounting (recall Assumption 2 from Section 2), although the available slack is substantially reduced. In the most demanding cases the computation uses up to 99% of punishment capacity, leaving only about 1% unused. When bounded discounting fails, however, the gap condition can break down—for example at $(\mu_0, \delta) = (0.1, 0.9)$. This is consistent with the fact that when δ is sufficiently small, $\overline{V}$ and $\underline{V}$ coincide at the origin and no surplus can be generated in the standard repeated game. When the gap condition fails at a particular reputation β, the optimal tuple there will involve the receiver following the sender's recommendation with probability less than 1. This is something that happens in extreme cases on the Extended Grid, but never on our maintained Grid.

Across all parameter combinations on both grids, $\overline{V}$ remains increasing in reputation. Thus monotonicity persists even in cases where the gap condition fails. This shows that the gap condition is sufficient but not necessary for monotonicity of $\overline{V}$.

Violations of Numerical Property 3 are rare and occur only at extreme parameterizations. On the maintained Grid there are two isolated cases,

$$(\mu_0, \delta, \pi_B) = (0.1, 0.999, 0.06) \quad \text{and} \quad (0.15, 0.999, 0.16),$$

in which $\pi(\beta)$ exhibits a single downward step at reputations arbitrarily close to one ($\beta \approx 1 - 10^{-199}$). Because these deviations occur precisely where slopes become numerically unstable, and disappear under minor tolerance adjustments, we interpret them as computational artifacts.

By contrast, on the Extended Grid, robust non-monotonicities arise for very small behavioral types ($\pi_B < 0.06$), suggesting that larger discount factors may be required to restore asymptotic monotonicity in that region. In addition, at $\mu_0 = 0.45$ a persistent nonmonotonic case appears at $\pi_B = 0.18$ across the full discount range considered. These failures occur away from the extreme tail and are therefore structurally distinct from the numerical artifacts described above.

7.2 Extensions

Beyond the baseline properties, we examine additional structural features of the equilibrium correspondence.

Numerical Property 4. [Interval Regions] For all (μ_0, δ, π_B) on the Grid, Regions 1, 2, and 3 are intervals.

Numerical Property 5. [Monotone Lower Boundary] For all (μ_0, δ, π_B) on the Grid, $\underline{V}(\beta, \delta)$ is weakly increasing in reputation β.

First, the interval property of Regions 1–3 remains remarkably stable. On the Extended Grid the interval structure persists throughout, with a single isolated exception: at $(\mu_0, \delta, \pi_B) = (0.45, 0.9, 0.18)$ the region labels locally oscillate between Regions 1 and 2 for a small number of adjacent gridpoints.

Second, while the *lower* boundary remains weakly increasing for all $\delta \geq 0.94$, weak monotonicity can fail when $\delta \leq 0.935$ and (μ_0, π_B) lie in the upper part of the grid (specifically $\mu_0 \in \{0.40, 0.45\}$ and $\pi_B \in [0.32, 0.78]$). Thus monotonicity of the lower boundary is specific to the maintained Grid.

One last issue requires separate discussion. While monotonicity of $\overline{V}$ is numerically unambiguous on some reputation ranges, in others the relevant payoff differences are extremely small—well below standard double precision. Thus, a purely computational verification of strict monotonicity could in principle be vulnerable to numerical error. To address this issue we combine computation with a structural argument. The key step is the identification of a "seed" interval inside Region 2 on which strict monotonicity can be verified with comfortable numerical slack.

Numerical Property 6. [Monotone Seed] For all (μ_0, δ, π_B) on the Grid, there exist $\beta_l < \beta_r$ such that (i) $\overline{V}(\beta)$ is strictly increasing on $[\beta_l, \beta_r] \subset R_2$ (with increases of at least 10^{-4} across adjacent reputations), and (ii) $\overline{V}(\beta_r) - \overline{V}(\beta_l) \geq (1-\delta)/\delta$.

The following lemma shows that strict monotonicity on such a seed interval propagates to the entire reputation space.

Lemma 2. Suppose the gap condition holds and Regions R_1, R_2, and R_3 are intervals. If there exist $\beta_l < \beta_r$ such that $\overline{V}(\beta)$ is strictly increasing on $[\beta_l, \beta_r] \subset R_2$ and $\overline{V}(\beta_r) - \overline{V}(\beta_l) \geq (1-\delta)/\delta$, then $\overline{V}(\beta)$ is strictly increasing on $[0, 1]$.

By the numerical properties established on the maintained Grid, all hypotheses of Lemma 2 are satisfied. Consequently, strict monotonicity of $\overline{V}$ holds on $[0,1]$ even in parametrizations where machine precision makes direct pointwise comparison unconvincing.

8. Multiple Behavioral Types

Focusing on a single behavioral type reveals a close connection between the reputational stochastic game and the static commitment problem studied in the Bayesian persuasion literature. It is natural, however, to allow for several behavioral types, corresponding to different reputations the sender might build for information transmission. In this section we extend the analysis to an arbitrary finite set of behavioral types.

The extension has two parts. We first prove an analogue of the δ-relevant behavioral convergence result from Section 6. We then show that, under additional monotonicity and regularity assumptions, the bound in that result can be sharpened, and that much of the equilibrium construction from the single-type model carries over to the multidimensional setting.

Let

$$\Pi^B = \{\pi_1^B, \dots, \pi_N^B\}, \qquad \pi_1^B < \cdots < \pi_N^B,$$

denote the finite set of behavioral types. Each behavioral type reports H when the state is h, and reports H with probability π_i^B when the state is ℓ. Throughout this section we maintain Assumptions 1 and 2, and replace Assumption 3 by:

Assumption 3'. $\pi_N^B < \pi^*$.

The sender's reputation is now represented by a vector

$$\beta = (\beta_1, \dots, \beta_N),$$

where β_i is the current belief that the sender is behavioral type π_i^B. Hence $1 - \sum_i \beta_i$ is the belief that the sender is rational. For any β, let

$$\hat\beta = \sum_{i=1}^N \beta_i, \qquad \hat\pi^B = \frac{1}{\hat\beta}\sum_{i=1}^N \beta_i \pi_i^B$$

whenever $\hat\beta > 0$. Thus $\hat\beta$ is total reputation and $\hat\pi^B$ is the conditional expected behavioral type.

Recall from the single-type model that, if the rational sender always reports truthfully, the receiver follows recommendation H if and only if

$$\pi(H \mid \ell) \le \bar\pi(\beta, \pi^B) := \frac{\pi^* - \beta\pi^B}{1-\beta}, \qquad \beta \in [0,1).$$

In the multidimensional model, assuming truthtelling in the high state, an H recommendation at reputation β is trusted if and only if

$$\pi(H \mid \ell) \leq \bar{\pi}(\hat{\beta}, \hat{\pi}^B).$$

Trust in the sender therefore depends only on the two aggregates $\hat{\beta}$ and $\hat{\pi}^B$. As before, $\pi(\beta, \delta)$ denotes the sender's low-state lying probability in an optimal tuple at β.

In the single-type model, we could rely on an explicit formula for $\overline{V}$ in Region 1 and on the ability to maintain reputation in Region 2. Neither extends verbatim. After an H report in the low state, the updated reputation on type i is proportional to $\beta_i \pi_i^B$; after an L report in the low state, it is proportional to $\beta_i(1 - \pi_i^B)$. These updates change not only total reputation, but also its composition across behavioral types. For instance, if the rational sender lies more than behavioral types on average at β, an H report decreases total reputation and shifts probability mass toward the behavioral types with the highest lying rates, while an L report increases total reputation and shifts probability mass toward the behavioral types with the lowest lying rates. With only the two reports H and L, the sender therefore lacks enough instruments to keep both total reputation and its composition constant. Hence the sender cannot remain indefinitely in the neighborhood of a multidimensional reputation vector. This is an important difference relative to the single-type case.

8.1 δ-Relevant Behavioral Convergence

The next result shows that the δ-relevant behavioral convergence result from Section 6 survives with a weaker bound.

Proposition 6. [δ-Relevant Behavioral Convergence] Let $\widehat{V}^N := \mu_0 + (1 - \mu_0)\pi_N^B$. For any $\epsilon > 0$, there exists $\underline{\delta} \in (0, 1)$ such that for every $\delta \in [\underline{\delta}, 1)$ and every initial reputation $\beta_0 \in [\epsilon, 1 - \epsilon]$, all sender-optimal equilibria (π, α) starting at β_0 satisfy $D(\pi, \alpha) \ \leq \ v^* - \widehat{V}^N + 5\epsilon$.

For sufficiently patient senders, all sender-optimal equilibria spend almost all δ-relevant time behaving nearly as in the persuasion benchmark. In the many-type model, the relevant lower bound is generated by the most exploitative behavioral type, namely π_N^B.

With a single behavioral type, the bound in Proposition 5 was tighter. To formulate an analogue here, we first need a notion of monotonicity for multidimensional reputations. For $\beta, \beta' \in \mathbb{R}_+^N$, write $\beta' \succeq \beta$ if $\beta_i' \geq \beta_i$ for every i, and write $\beta' \succ \beta$ if in addition strict inequality holds for at least one coordinate. Define $\overline{V}(\beta, \delta)$ to be (strictly) increasing in β if it respects this coordinatewise partial order.

Assume from now on that $V(\beta, \delta)$ is strictly increasing, that $\pi(\beta, \delta)$ is continuous in β, and that V is uniformly continuous near the β_N-axis in the following sense: for every $\epsilon > 0$ there exist $r > 0$ and $\bar{\delta} < 1$ such that, for every $\delta \in [\bar{\delta}, 1)$ and every $\beta, \beta' \in \mathbb{R}^N_+$ satisfying

$$\hat{\beta} = \hat{\beta}' < 1 \qquad \text{and} \qquad \frac{\beta_i}{\beta_N} \le r,\ \frac{\beta'_i}{\beta'_N} \le r \quad \text{for all } i < N,$$

we have

$$|V(\beta, \delta) - V(\beta', \delta)| \le \epsilon.$$

Under these assumptions, the benchmark $\widehat{V}^N$ in Proposition 6 can be replaced by $V^M(\pi^*, \pi^B_N)$. The argument is as follows. Fix an interior reputation β^0. Define recursively a backward L-predecessor sequence $\{\beta^k\}_{k\ge 0}$ by requiring that, for each k, a truthful report in the low state at β^{k+1} produces β^k. Then for every $i < N$,

$$\frac{\beta_i^{k+1}}{\beta_N^{k+1}} = \frac{1 - \pi_N^B}{1 - \pi_i^B}\frac{\beta_i^k}{\beta_N^k}.$$

Since $\pi_i^B < \pi_N^B$ for all $i < N$, the multiplier on the right-hand side is strictly below one. Therefore

$$\frac{\beta_i^k}{\beta_N^k} \to 0 \qquad \text{for every } i < N,$$

so the sequence approaches the β_N-axis.

Along this sequence the reputation vectors become arbitrarily close to the β_N-axis, where the many-type model reduces to the single-type model with behavioral type π_N^B. The single-type analysis therefore delivers the benchmark $V^M(\pi^*, \pi^B_N)$ along the axis. Uniform continuity near the β_N-axis implies that $V(\beta^k, \delta)$ becomes arbitrarily close to this benchmark for large k. Since the sequence $\{\beta^k\}$ also moves monotonically toward the β_N-axis, i.e., $\beta^k \succeq \beta^{k+1}$ in the coordinatewise order, we have $\overline{V}(\beta_0, \delta) \ge \overline{V}(\beta^k, \delta)$ for all k, hence

$$\overline{V}(\beta_0, \delta) \ge V^M(\pi^*, \pi_N^B) - \epsilon.$$

8.2 Equilibrium Structure

With the same monotonicity assumption, much of the equilibrium structure from the single-type model survives. Rather than attempt a complete characterization, we describe how the tripartite classification from Section 5 remains useful in the multidimensional setting.

Holding β fixed, every dimension of $\beta_{H,\ell}(\beta, \pi)$ decreases with π, while every dimension of $\beta_{L,\ell}(\beta, \pi)$ increases with π. Hence, if $\overline{V}(\beta, \delta)$ is strictly increasing in β, raising π increases the continuation-value spread

$$\overline{V}(\beta_{L,\ell}(\beta, \pi), \delta) - \overline{V}(\beta_{H,\ell}(\beta, \pi), \delta).$$

This is the multidimensional analogue of the same observation in Section 5.

We begin with Region 3. As in the single-type model, this is the set of reputations at which the sender prefers to lie for sure in state ℓ when continuation values remain on the upper frontier. The same calculation therefore yields the same threshold

$$V^{23} = \frac{\delta - (2\delta - 1)\mu_0}{1 - \delta\mu_0},$$

so Region 3 is characterized by $V(\beta, \delta) \geq V^{23}$.

Now consider a reputation vector β outside Region 3. If there exists $\pi \leq \bar{\pi}(\hat{\beta}, \hat{\pi}^B)$ such that the continuation-value spread ensures sender indifference while both continuation values remain on the upper frontier, we say that β lies in Region 2. In this region incentives are provided efficiently, in the same sense as in the single-type model: no surplus is destroyed after a lie.

All remaining reputation vectors belong to Region 1. There the trust constraint

$$\pi \leq \bar{\pi}(\hat{\beta}, \hat{\pi}^B)$$

binds before enough continuation-value spread can be generated on the upper frontier. The sender can still be induced to randomize in the low state, but only by lowering the continuation value after a lie below the upper frontier. In this sense, Region 1 is again the region of inefficient incentives.

9. Conclusion

A long-lived sender can build, maintain or run down his reputation for a degree of honesty in reporting information that arrives period by period. If he is extremely patient, a powerful result of Fudenberg and Levine (1992) guarantees him an average discounted payoff almost as high as the Bayesian persuasion value of KG for a sender who can commit himself to a particular random information protocol, if he can imitate a behavioral type close to the KG commitment rate. Here, we investigate theoretically and computationally what kind of equilibrium behavior supports this value result. In addition, we study both value and behavior for discount factors distant from 1.

For general discount factors, sender-optimal equilibria have a three-region structure. Usually each of these regions is an interval in reputation space. In Region 3, where reputation is very high, the receiver always acts on the sender's advice, although the (rational) sender always claims the state is high even when it is low. He is "cashing in" on his high reputation. In Region 2, where reputation is more moderate, the receiver again trusts the sender, who randomizes between reporting honestly or dishonestly. The randomization probability at any reputation level is constructed such that lying is punished just enough to keep the sender indifferent between reporting the truth or lying. Whichever he does, at the new prevailing reputation, the continuation value is maximal in the set of values at that reputation, so we say that Region 2 involves efficient provision of incentives. By contrast, at the lower reputation levels of Region 1, the sender's continuation value after a lie is not on the upper boundary of the value correspondence. Here, there is inefficient provision of incentives.

None of the preceding results requires the discount factor δ to approach 1. We were surprised by what happens when that limit exercise is performed. Section 6 undertakes the asymptotic analysis, in the presence of a single behavioral type. Recall that π_B is the probability with which the behavioral type lies when the state is low, and that $\pi_B < \pi^*$ (where π^* is the Bayesian persuasion probability of lying in the low state). As δ approaches 1, the sender's payoff therefore need not converge to the Bayesian persuasion payoff. In spite of that, in regular cases,[19] the sender's equilibrium *behavior* does approach the Bayesian persuasion probability!

At each fixed reputation value in $(0, 1)$, Proposition 4 establishes that the sender-optimal equilibrium starting at that reputation involves the sender's behavior conforming with the Bayesian persuasion probability for at least T periods, where T can be as large as desired if δ is high enough. For any particular δ, this does not yield the Bayesian persuasion value, because π^* exceeds π_B, and in the long run, the sender loses his reputation and receives low payoffs. But a novel technique for calculating equilibrium payoffs, which we call *stationary promise-keeping*, easily computes the true limiting average discounted value. If one chooses π_B close to π^* and then chooses δ sufficiently high, both behavior and value will closely resemble the Bayesian persuasion solution of KG, providing a dynamic interpretation of Bayesian persuasion analysis without resorting to any legal commitment.

Will Bayesian persuasion behavior persist for a length of time that is significant in discounted terms? (Or could the number of periods T in the preceding paragraph be insignificant relative to the time horizon the sender cares about?) Proposition 5 shows that, for sufficiently patient senders, all sender-optimal equilibria stay close to the Bayesian persuasion benchmark in discount-relevant time, with a uniform bound determined by the gap between the persuasion value and the lower benchmark generated by the available behavioral type. Thus the paper links commitment and reputation not only through pointwise or finite-horizon behavior, but through the part of the horizon that matters for discounted payoffs. When the available behavioral type is itself close to the Bayesian persuasion benchmark, the equilibrium is close to that benchmark for almost all discount-relevant time. Section 8 extends this perspective to an arbitrary finite collection of behavioral types.

[19]Here we assume, with computational support, that for sufficiently large δ, equilibrium lying probabilities increase after an honest report.

Appendix

A. Reputational APS Framework

A.1 Self Generation

In equilibrium, continuation values are drawn from the equilibrium value correspondence V. But to employ the methods of strategic dynamic programming, it is useful to imagine those values being drawn from an abstract value correspondence W. Let $\mathcal{C}$ be the collection of correspondences $W : [0,1]^N \rightrightarrows \mathbb{R}_+$ with a compact graph such that $W(\beta)$ is a nonempty and compact interval for each $\beta \in [0,1]^N$. We will often use W to denote the graph of the correspondence, so that $W \subset W'$ means $W(\beta) \subset W'(\beta)$ for all β.

Assume that in the current period the receiver believes the sender is behavioral type π_i^B with probability β_i, $i = 1, \ldots, N$ (and rational with probability $1 - \hat{\beta}$), and expects the rational sender to play strategy $\pi : \Theta \to \Delta(M)$. After receiving message m, the receiver updates her beliefs about the state and chooses a myopic (possibly mixed) best reply $\alpha_m \in \Delta(A)$. Let $\mathrm{BR}^R(\beta, \pi) = \{\alpha = (\alpha_H, \alpha_L) \mid \alpha_m \text{ is a best reply to message } m,\ m \in M\}$.

If the sender expects the receiver to play strategy $\alpha \in \Delta(A)^M$ in the current period and continuation value $w_{m,\theta}$ after sending m in state θ, then his total expected value for strategy π is

$$E(\pi, \alpha, w) = \sum_{\theta \in \Theta} \mu_0(\theta) \sum_{m \in M} \pi(m|\theta) \left[(1-\delta) \sum_{a \in A} v(\theta, a) \alpha_m(a) + \delta w_{m,\theta} \right]$$

where $w = (w_{m,\theta})_{m,\theta}$. The sender's best-reply set is then

$$\mathrm{BR}^S(\alpha, w) = \operatorname{argmax} \{E(\pi, \alpha, w) | \pi : \Theta \to \Delta(M)\}.$$

Definition 4. Given $W \in \mathcal{C}$, the tuple (π, α, w) is admissible for W at $\beta \in [0,1]^N$ if $w_{m,\theta} \in W(\beta_{m,\theta}(\beta, \pi))$ for each $(m, \theta) \in M \times \Theta$ and

$$\pi \in \mathrm{BR}^S(\alpha, w) \quad \text{and} \quad \alpha \in \mathrm{BR}^R(\beta, \pi).$$

A tuple (π, α, w) is admissible if (π, α) is a *static* Nash equilibrium of a game whose payoffs are given by the per-period payoffs augmented with continuation payoffs w.

Define the correspondence $B(W)$ to be the set of sender's attainable payoffs given W:

$$B(W)(\beta) = \{E(\pi, \alpha, w) \mid (\pi, \alpha, w) \text{ is admissible for } W \text{ at } \beta\}, \quad \beta \in [0,1]^N,$$

and define $\tilde{B}(W)$ as $\tilde{B}(W)(\beta) = \mathrm{co}(B(W(\beta)))$ for each $\beta \in [0,1]$. The image of B is convexified utilizing public randomization.

Definition 5. W is self-generating if $W \subset \tilde{B}(W)$.

The following two propositions are natural extensions of their counterparts for repeated games; their proofs are omitted.

Proposition 7. If $W \in \mathcal{C}$ is self-generating, then $\tilde{B}(W) \subset V$.

If a payoff correspondence is self-generating, so that each value can be generated by an admissible tuple whose continuation payoff can also be generated by an admissible tuple and so on, then it must be a subset of the PBE correspondence.

Proposition 8. $V = \tilde{B}(V)$. Moreover, V is the largest fixed point of $\tilde{B} : \mathcal{C} \to \mathcal{C}$.

A.2. Algorithm and Existence

For a fixed δ, the PBE correspondence V may be computed iteratively, starting from the set of feasible payoffs of the sender. Let $W^0 \in \mathcal{C}$ be such that $V \subset W^0$. Given W^k, the iterative step constructs $W^{k+1}(\beta)$ as the set of values of all admissible pairs at β, with continuation payoffs respecting W^k. Below we first introduce admissibility and then show that W^{k+1} inherits from W^k the following properties: W^{k+1} has a closed graph, and for each β, $W^{k+1}(\beta)$ is a nonempty and convex set. Since $\{W^k\}$ is a decreasing sequence of nonempty compact sets, $W^\infty(\beta) = \bigcap_{k \in N} W^k(\beta)$ is nonempty, by the finite intersection property. A self generation argument shows that $W^\infty = V$.

The algorithm applies the $\tilde{B}$ map iteratively to a set that is originally large enough to include all possible PBE payoffs. In particular, starting from the correspondence $W^0 \in \mathcal{C}$, where $W^0(\beta) = [0, 1]$ for all $\beta \in [0, 1^N$, the $\tilde{B}$ map reduces its size without ruling out any PBE payoff:

$$V \subset \tilde{B}(W^0) \subset W^0.$$

Since the $\tilde{B}$ map is monotonic, iterative applications from W^0 keep on shrinking the set without ever excluding any PBE. Formally, let $W^{k+1} = \tilde{B}(W^k)$, $k \geq 0$. The sequence $\{W^k\}$ is decreasing: $W^{k+1} \subset W^k$ for all $k \in \mathbb{N}$. Let

$$W^\infty = \lim_{k \to \infty} W^k = \bigcap_{k \in N} W^k.$$

Importantly, $\tilde{B}(W)$ is nonempty valued when $W \in \mathcal{C}$, which is crucial for non-vacuous convergence of the iterative process.

Lemma 3. If $W \in \mathcal{C}$ then $\tilde{B}(W) \in \mathcal{C}$.

The crucial step in the proof of Lemma 3 is demonstrating, for each possible β, the existence of a tuple admissible with respect to W. While this is similar to establishing the existence of Nash equilibrium in a static game, there is the added complication of finding continuation payoffs from the appropriate value sets at respective updated reputations. The proof essentially adds a dummy player whose strategy selects the continuation payoffs. In the resulting three player game, we let Kakutani's fixed point theorem find suitable continuation payoffs along with strategies for the sender and the receiver.

Proof: Let $\underline{w} = \min\{\underline{W}(\beta) \mid \beta \in [0,1]^N\}$ and $\bar{w} = \max\{\overline{W}(\beta) \mid \beta \in [0,1]^N\}$ Clearly $E(\pi, \alpha, w) \in [\delta\underline{w}, (1-\delta) + \delta\overline{w}]$ for each $\beta \in [0,1]^N$ and each tuple (π, α, w) admissible for W at β. Therefore

$$B(W) \subset [0,1] \times [\delta\underline{w}, (1-\delta) + \delta\overline{w}]$$

is a bounded set. Let $\{(\beta^k, x^k)\} \subset B(W)$ be a sequence such that $\beta^k \to \beta$ and $x^k \to x$. We now show that $(\beta, x) \in B(W)$. For each $k \in \mathbb{N}$, there is a tuple (π^k, α^k, w^k) admissible for W at β^k such that $x^k = E(\pi^k, \alpha^k, w^k)$. Since $\Delta(M)^\Theta$, $\Delta(A)^M$ and $[\underline{w}, \overline{w}]^{M\times\Theta}$ are all compact sets, without loss of generality we can assume that $\pi^k \to \pi^\infty$, $\alpha^k \to \alpha^\infty$ and $w^k \to w^\infty$ for some $\pi^\infty \in \Delta(M)^\Theta$, $\alpha^\infty \in \Delta(A)^M$ and $w^\infty \in [\underline{w}, \overline{w}]^{M\times\Theta}$. One can check that $(\pi^\infty, \alpha^\infty, w^\infty)$ is admissible for W at β. Finally,

$$x = \lim_{k\to\infty} x^k = \lim_{k\to\infty} E(\pi^k, \alpha^k, w^k) = E(\pi^\infty, \alpha^\infty, w^\infty).$$

Therefore, $(\beta, z) \in B(W)$. This establishes that $B(W) \subset [0,1]^N \times \mathbb{R}$ is a closed set and hence that $B(W) : [0,1]^N \to [\delta\underline{w}, (1-\delta) + \delta\overline{w}]$ is an upper hemicontinuous correspondence. Since $\tilde{B}(W)(\beta) = \mathrm{co}(B(W(\beta)))$ for each $\beta \in [0,1]^N$, $\tilde{B}(W)(\beta)$ is a compact convex set for each $\beta \in [0,1]^N$ and $\tilde{B}(W) : [0,1]^N \to [\delta\underline{w}, (1-\delta) + \delta\overline{w}]$ is also an upper hemicontinuous correspondence.

It remains to show that $B(W)(\beta) \neq \emptyset$ for each $\beta \in [0,1]^N$. Fix $\beta \in [0,1]^N$ and consider a simultaneous moves auxiliary game with three players: the sender, the receiver and a "dummy player". The sender's payoff function in the auxiliary game is $E(\pi, \alpha, w)$, and his best reply correspondence is $\mathrm{BR}^S(\alpha, w)$. The receiver's payoff is $u(\theta, a)$, as in the component game (and does not depend on the dummy's actions), and his best reply correspondence is $\mathrm{BR}^R(\beta, \pi)$. We do not specify a payoff function for the dummy player. Instead, we specify directly his "best-reply correspondence":

$$\mathrm{BR}^D(\pi, \alpha) = \prod_{(m,\theta)\in M\times\Theta} W(\beta_{m,\theta}(\beta, \pi)).$$

Clearly, for each (m, θ), the function $\pi \to \beta_{m,\theta}(\beta, \pi)$ from $\Delta(M)^\Theta$ to $[0,1]$ is continuous. Since $W \in \mathcal{C}$, the correspondence $\pi \to W(\beta_{m,\theta}(\beta, \pi))$ is upper hemicontinuous with convex and compact values.

As remarked earlier, $\mathrm{BR}^S(\alpha, w)$ and $\mathrm{BR}^R(\beta, \pi)$ are convex and compact sets. By the Maximum Theorem, one can also easily check that BR^S and BR^R are upper hemicontinuous correspondences. Let

$$\mathrm{BR}(\pi, \alpha, w) = \mathrm{BR}^S(\alpha, w) \times \mathrm{BR}^R(\beta, \pi) \times \mathrm{BR}^D(\pi, \alpha).$$

It is easy to see that (π, α, w) is an admissible tuple for W at β if and only if (π, α, w) is a fixed point for BR. By Kakutani's theorem, the correspondence BR from $\Delta(M)^\Theta \times \Delta(A)^M \times [\underline{w}, \overline{w}]^{M\times\Theta}$ into itself has a fixed point (π, α, w). Therefore, $E(\pi, \alpha, w) \in B(W)(\beta)$ and $B(W)(\beta) \neq \emptyset$. □

Proposition 9. $W^\infty \in \mathcal{C}$ and $W^\infty = V$.

As an immediate corollary, a PBE exists because $V \in \mathcal{C}$.

Proof: By definition,

$$W^\infty(\beta) = \lim_{k\to\infty} W^k(\beta) = \bigcap_{k\in\mathbb{N}} W^k(\beta).$$

The intersection of compact and convex sets is a compact convex set, and by the finite intersection property $W^\infty(\beta)$ is non-empty. Thus $W^\infty \in \mathcal{C}$.

We now show that W^∞ is self-generating. Let $\beta \in [0,1]^N$ and $w \in W^\infty(\beta)$. We need to show that $w \in \tilde{B}(W^\infty)(\beta)$. Since $w \in W^\infty(\beta)$, $w \in W^{k+1}(\beta)$ for all k. Therefore, for each $k \in \mathbb{N}$ there exist $\lambda^k \in [0,1]$ and two tuples $(\pi^{1k}, \alpha^{1k}, w^{1k})$ and $(\pi^{2k}, \alpha^{2k}, w^{2k})$ admissible for W^k at β such that $w = \lambda^k E(\pi^{1k}, \alpha^{1k}, w^{1k}) + (1-\lambda^k)E(\pi^{2k}, \alpha^{2k}, w^{2k})$. Again, without loss of generality we can assume that $\lambda^k \to \lambda^\infty$, $\pi^{jk} \to \pi^{j\infty}$, $\alpha^{jk} \to \alpha^{j\infty}$ and $w^{jk} \to w^{j\infty}$ for some $\pi^{j\infty} \in \Delta(M)^\Theta$, $\alpha^{j\infty} \in \Delta(A)^M$ and $w^{j\infty} \in [\underline{w}, \overline{w}]^{M\times\Theta}$. It is easy to check that $w^{j\infty}_{m,\theta} \in W^\infty(\beta_{m,\theta}(\beta, \pi^{j\infty}))$ for each $(m,\theta) \in M \times \Theta$, and hence that $(\pi^{j\infty}, \alpha^{j\infty}, w^{j\infty})$ is an admissible tuple for W^∞ at β, $j = 1, 2$. Moreover,

$$\begin{aligned} w &= \lim_{k\to\infty} \lambda^k E(\pi^{1k}, \alpha^{1k}, w^{1k}) + (1-\lambda^k)E(\pi^{2k}, \alpha^{2k}, w^{2k}) \\ &= \lambda^\infty E(\pi^{1\infty}, \alpha^{1\infty}, w^{1\infty}) + (1-\lambda^\infty)E(\pi^{2\infty}, \alpha^{2\infty}, w^{2\infty}). \end{aligned}$$

Therefore $w \in \tilde{B}(W^\infty)(\beta)$ as was to be shown.

By Proposition 7, $W^\infty \subset V$. Conversely, since $\tilde{B}$ is monotone and $V \subset \tilde{B}(W^0) \subset W^0$, we have $V \subset W^k$ for all $k \geq 0$. Therefore, we also have that $V \subset W^\infty$. □

A.3 Monotone and Continuous Optimal Values: Proofs

Until Section D of the Appendix, we now restrict attention to the model with just one behavioral type ($N = 1$). For any $W \in \mathcal{C}$, let

$$\underline{W}(\beta) = \min W(\beta) \quad \text{and} \quad \overline{W}(\beta) = \max W(\beta) \quad \text{for } \beta \in [0,1].$$

Value correspondences W with a monotone upper boundary $\overline{W}$ play an important role in our analysis. Let

$$\begin{aligned}
\mathcal{C}^+ &= \{W \in \mathcal{C} \mid \overline{W}(\beta) \text{ is weakly increasing in } \beta\} \\
\mathcal{C}^{++} &= \{W \in \mathcal{C} \mid \overline{W}(\beta) \text{ is strictly increasing in } \beta\}. \\
\mathcal{C}_0 &= \{W \in \mathcal{C} \mid W \supset V\}, \\
\mathcal{C}_0^+ &= \{W \in \mathcal{C}^+ \mid W \supset V\} \quad \text{and} \quad \mathcal{C}_0^{++} = \{W \in \mathcal{C}^{++} \mid W \supset V\}.
\end{aligned}$$

Obviously, $\mathcal{C}^{++} \subset \mathcal{C}^+$. Pick any $W \in \mathcal{C}^+$. It is possible that $\overline{W}$ is not continuous, but since W is upper hemicontinuous (as it has a compact graph), $\overline{W}$ is upper semi-continuous. As $\overline{W}$ is a weakly increasing function, $\overline{W}$ must also be right continuous. Define

$$\overline{W}_L(\beta) = \lim_{\beta' \uparrow \beta} \overline{W}(\beta').$$

By upper hemicontinuity, $\overline{W}_L(\beta) \leq \overline{W}(\beta)$ and $[\overline{W}_L(\beta), \overline{W}(\beta)] \subset W(\beta)$.

We begin with the degenerate case where reputation is 0: it is common knowledge that the sender is rational, not behavioral. The value set there includes 0, the sender's payoff in the babbling equilibrium, and also the truthtelling equilibrium, which relies on our assumption $\delta \geq 1/(1+\mu_0)$.

Lemma 4. $V(0) = [0, \mu_0]$ for all $\delta \geq 1/(1+\mu_0)$.

Proof: When $\beta = 0$ the game reduces to an infinitely repeated game with perfect information. In the infinitely repeated game there is a "babbling" equilibrium where the sender's messages are ignored and the receiver chooses action L in every period regardless of the sender's message. Given that the sender's messages are uninformative and that $\mu_0 < 1/2$, L is the unique optimal action for the receiver. This equilibrium has an expected value 0 (for the sender) and is the worst equilibrium for the sender. The repeated game also has an equilibrium where on the outcome path the sender always tells the truth and the receiver accepts all the sender's recommendations. That is, in every period t on the equilibrium path, the sender chooses the strategy π_t where $\pi_t(H|h) = 1$ and $\pi_t(L|\ell) = 1$, and the receiver chooses the action $a_t = H$ if $m_t = H$ and the action $a_t = L$ if $m_t = L$. If the sender ever lies, the players revert to the babbling equilibrium. This strategy has expected value μ_0 for the sender. Given that the receiver accepts all his recommendations, the sender is tempted to send the message $m_t = H$ when $\theta_t = \ell$. But the continuation value when he lies is $(1-\delta) + \delta 0$, while his continuation value from telling the truth is $\delta\mu_0$. This deviation is not profitable. A simple value recursion shows that this is an optimal equilibrium. Therefore $\{0, \mu_0\} \subset V(0)$ and $[0, \mu_0] = V(0)$. $\square$

For the rest of this section we fix $W \in \mathcal{C}_0^{++}$ and let

$$X = \tilde{B}(W).$$

We now focus on the admissible tuples that support the optimal values $\bar{X}(\beta)$, $\beta \in [0,1]$. Without loss of generality, we restrict attention to admissible tuples (π, α, w) where $\alpha_H(H) \geq \alpha_L(H)$; if $\alpha_H(H) < \alpha_L(H)$ we can just reverse the role played by messages H and L in the definition of the tuple.

Lemma 5. $\overline{X}(0) = (1-\delta)\mu_0 + \delta\overline{W}(0) \geq \mu_0$.

Proof: If (π, α, w) is an admissible tuple for W at $\beta = 0$ and $\alpha_H(H) = 0$ (and hence $\alpha_L(H) = 0$ as well), then

$$E(\pi, \alpha, w) = \delta\big[\mu_0[\pi(H|h)w_{H,h} + \pi(L|h)w_{L,h}] + (1-\mu_0)[\pi(H|\ell)w_{H,\ell} + \pi(L|\ell)w_{L,\ell}]\big]$$

and $E(\pi, \alpha, w) \leq \delta\overline{W}(0)$. On the other hand, a tuple (π, α, w) with $\alpha_H(H) > 0$ is admissible only if the receiver's posterior that $\theta = h$ given the message H is greater or equal to $1/2$. That is,

$$\mathbb{P}[h|H] = \frac{\pi(H|h)\mu_0}{\pi(H|h)\mu_0 + \pi(H|\ell)(1-\mu_0)} \geq \frac{1}{2} \quad \text{or} \quad \pi(H|h)\mu_0 - \pi(H|\ell)(1-\mu_0) \geq 0.$$

Since $\mu_0 < 1/2$, this implies that $\pi(H|\ell) < 1$ and

$$\pi(L|h)\mu_0 - \pi(L|\ell)(1-\mu_0) = (2\mu_0 - 1) - [\pi(H|h)\mu_0 - \pi(H|\ell)(1-\mu_0)] < 0.$$

Thus $\mathbb{P}[h|L] < 1/2$. Therefore $\alpha_L(H) = 0$ and $E(\pi, \alpha, w) = \mu_0 v_h + (1-\mu_0)v_\ell$ where

$$v_\theta = \pi(H|\theta)(\alpha_H(1-\delta) + \delta w_{H,\theta}) + (1 - \pi(H|\theta))\delta w_{L,\theta} \quad \text{for } \theta = h, \ell.$$

Note that when $\beta = 0$, $\beta_{m,\theta} = 0$ for all (m, θ). One can check that any tuple (π^o, α^o, w^o) where

$$\begin{aligned} &\pi^o(H|h) = 1, \quad \pi^o(H|\ell) \leq \pi^*, \quad \alpha^o_H(H) = \alpha^o_L(L) = 1 \quad \text{and} \\ &w^o = (w^o_{H,h}, w^o_{L,h}, w^o_{H,\ell}, w^o_{L,\ell}) = (\overline{W}(0), 0, 0, \overline{W}(0)) \end{aligned}$$

has value

$$E(\pi^o, \alpha^o, w^o) = (1-\delta)\mu_0 + \delta\overline{W}(0)$$

and is optimal for W at $\beta = 0$. □

Definition 6. A tuple (π, α, w) is optimal for W at some $\beta \in [0,1]$ if it is admissible at β and $E(\pi, \alpha, w) = \overline{X}(\beta)$.

We will show below that if (π, α, w) is an optimal tuple for W at some $\beta \in (0,1)$, then $\pi(H|h) = 1$. When $\pi(H|h) = 1$, the posteriors

$$\beta_{H,\ell}(\beta, \pi) = \frac{\beta \pi_B}{\beta \pi_B + (1-\beta)\pi(H|\ell)} \qquad \text{and}$$
$$\beta_{L,\ell}(\beta, \pi) = \frac{\beta(1-\pi_B)}{\beta(1-\pi_B) + (1-\beta)(1-\pi(H|\ell))}$$

depend exclusively on $\pi(H|\ell)$. Hereafter we express these posteriors as $\beta_{H,\ell}(\beta, \pi(H|\ell))$ and $\beta_{L,\ell}(\beta, \pi(H|\ell))$ instead. Also, to simplify notation, we sometimes omit the arguments and simply write $\beta_{H,\ell}$ and $\beta_{L,\ell}$.

Let (π, α, w) be an admissible tuple for W at some $\beta \in (0,1)$ such that $\pi(H|h) = 1$. Then, by Bayes' rule, the receiver's posterior belief that $\theta = h$ given message H is

$$\mu(\pi(H|\ell), \beta) = \frac{\mu_0}{\mu_0 + (1-\mu_0)(\beta\pi_B + (1-\beta)\pi(H|\ell))}.$$

Recall

$$\bar{\pi}(\beta) := \frac{\pi^* - \beta\pi_B}{1-\beta} \quad \text{for} \quad \beta \in [0,1],$$

and note that if $\pi(H|\ell) = \bar{\pi}(\beta)$, the receiver is indifferent about following the recommendation H because in this case $\mu(\pi, \beta) = 1/2$. When the sender deceives as often as can in state ℓ without compromising overall trust, that is, when $\pi(H|\ell) = \bar{\pi}(\beta)$, he induces updated reputations $\beta_{H,\ell} = \lambda_0\beta < \beta$ if the dishonest message is realized, and $\beta_{L,\ell} = \lambda_1\beta > \beta$ otherwise, where

$$\lambda_0 = \frac{\pi_B(1-\mu_0)}{\mu_0} \quad \text{and} \quad \lambda_1 = \frac{(1-\pi_B)(1-\mu_0)}{1-2\mu_0}.$$

Note that $\bar{\pi}(\beta_d) = 1$ when

$$\boxed{\beta_d := 1/\lambda_1}$$

In our numerical computations we always observe that V satisfies the gap condition for all the parameter specifications.

As specified in 4.2, the Gap Condition on V is maintained throughout. We show below that under this assumption, $\overline{V}(\beta)$ is weakly increasing in β.

Lemma 6. If $W \in \mathcal{C}_0$, then W satisfies the gap condition.

Proof: For any $\beta \in [0, 1/\lambda_1]$

$$\overline{W}(\lambda_1\beta) - \frac{1-\delta}{\delta} \geq \overline{V}(\lambda_1\beta) - \frac{1-\delta}{\delta} \geq \underline{V}(\lambda_0\beta) \geq \underline{W}(\lambda_0\beta). \qquad \square$$

We now construct an admissible tuple (π^o, α^o, w^o) for each $\beta \in [0,1]$. We will show later that these tuples are optimal. For any $\beta \in (0,1)$ and π such that $\pi(H|\ell) > \pi_B$, $\beta_{H,\ell}(\beta, \pi(H|\ell)) < \beta < \beta_{L,\ell}(\beta, \pi((H|\ell))$. Let

$$\hat{\pi}(\beta) = \sup\{\pi(H|\ell) \in [\pi_B, 1] \mid \delta[\overline{W}(\beta_{L,\ell}(\beta, \pi(H|\ell))) - \overline{W}_L(\beta_{H,\ell}(\beta, \pi(H|\ell)))] \leq 1-\delta\}.$$

Note that $\beta_{L,\ell}(\beta, 1) = 1$ and that $\beta_{H,\ell}(\beta, 1)$ is a continuously increasing function of β. Let

$$\beta_1 = \inf\{\beta \in [0,1] \mid \hat{\pi}(\beta) = 1\} \quad \text{and} \quad \beta^{23} = \min\{\beta_1, 1/\lambda_1\}.$$

Then

$$\overline{W}(1) - (1-\delta)/\delta \in [\overline{W}_L(\beta_{H,\ell}(\beta_1, 1), \overline{W}(\beta_{H,\ell}(\beta_1, 1)]$$

and for all $\beta > \beta_1$, by the monotonicity of $\overline{W}(\beta)$, $\hat{\pi}(\beta) = 1$ and

$$\delta[\overline{W}(1) - \overline{W}_L(\beta_{H,\ell}(\beta, 1)] < 1 - \delta.$$

Moreover, if $\bar{\pi}(\beta) \geq 1$, following an H recommendation from the sender is optimal even if $\pi(H|\ell) = 1$.

For $\beta \in [0, \beta^{23})$, define the tuple (π^o, α^o, w^o) by

1. $\pi^o(H|h) = 1$ and $\pi^o(H|\ell) = \min\{\hat{\pi}(\beta), \bar{\pi}(\beta)\}$
2. $\alpha^o_H(H) = \alpha^o_L(L) = 1$
3. $w^o_{L,\ell} = \overline{W}(\beta_{L,\ell}(\beta, \pi^o(H|\ell)))$ and $w^o_{H,\ell} = w^o_{L,\ell} - (1-\delta)/\delta$,
4. $w^o_{L,h} = \underline{W}(0)$ (for convenience) and $w^o_{H,h} = \overline{W}(\beta)$.

When $\delta[\overline{W}(\beta_{L,\ell}(\beta, \pi^o(H|\ell))) - \overline{W}_L(\beta_{H,\ell}(\beta, \pi^o(H|\ell)))] > 1-\delta$ (because $\pi^o(H|\ell) = \hat{\pi}(\beta) \leq \bar{\pi}(\beta)$ and the upper boundary of W has a vertical segment at $\beta_{H,\ell}(\beta, \pi^o(H|\ell))$ or at $\beta_{L,\ell}(\beta, \pi^o(H|\ell))$), the continuation values $w^o_{L,\ell}$ and $w^o_{H,\ell}$ are maximized while ensuring that

$$w^o_{m,\ell} \in [\overline{W}(\beta_{m,\ell}(\beta, \pi^o(H|\ell))) - \overline{W}_L(\beta_{m,\ell}(\beta, \pi^o(H|\ell)))] \quad m = L, H.$$

For $\beta \in [\beta^{23}, 1]$, define the tuple (π^o, α^o, w^o) by

1. $\pi^o(H|h) = \pi^o(H|\ell) = 1$
2. $\alpha^o_H(H) = \alpha^o_L(L) = 1$

3. $w^o_{L,\ell} = \overline{W}(1)$ and $w^o_{H,\ell} = \overline{W}(\beta_{H,\ell}(\beta, 1))$

4. $w^o_{L,h} = \underline{W}(0)$ (for convenience) and $w^o_{H,h} = \overline{W}(\beta)$

Remark: In all cases, given that $\pi^o(H|\ell) = 1$ and that $\pi_B(H|h) = 1$, $\beta_{L,h}$ is not well defined by Bayes' rule; by continuity we let $\beta_{L,h} = 0$ (as this is the posterior for any $\pi^o(H|h) < 1$).

Lemma 7. For each $\beta \in [0, 1]$, the corresponding (π^o, α^o, w^o) is an admissible tuple for W at β.

Proof: Consider first the case when $\beta \in (0, \beta^{23}]$. Then the continuation values satisfy the incentive constraint:

$$(1 - \delta) + \delta w^o_{H,\ell} = \delta w^o_{L,\ell}, \tag{IC}$$

the sender is indifferent between messages H and L when $\theta = \ell$ and thus any $\pi^o(H|\ell) \in [0, 1]$ is optimal for the sender. Also, since $w^o_{L,h} = \underline{W}(0) < \overline{W}(0) < \overline{W}(\beta) = w^o_{H,h}$, $(1 - \delta) + \delta w^o_{H,h} > \delta w^o_{L,h}$ and the sender strictly prefers sending message H when $\theta = h$, so $\pi^o(H|h) = 1$ is optimal for the sender. Given that $\pi^o(H|h) = \pi_B(H|h) = 1$, after receiving a message L, the receiver believes $\theta = \ell$ for sure and strictly prefers to take action L, so $\alpha^o_L(L) = 1$ is optimal for her. Finally, given that $\pi^o(H|\ell) \leq \bar{\pi}(\beta)$, the receiver (weakly) prefers to take action H after receiving message H and thus $\alpha^o_H(H) = 1$ is optimal for her. Note also that $w^o_{m,\theta} \in W(\beta_{m,\theta})$ for all (m, θ). In particular, when $\pi^0(H|\ell) = \hat{\pi}(\beta)$,

$$w_{H,\ell} \in [\overline{W}_L(\beta_{H,\ell}), \overline{W}(\beta_{H,\ell})],$$

and when $\pi^0(H|\ell) = \bar{\pi}(\beta) < \hat{\pi}(\beta)$, $w_{H,\ell} \in W(\beta_{H,\ell})$ because W satisfies the gap condition.

Now consider the case when $\beta \in (\beta^{23}, 1]$. Then $\hat{\pi}(\beta) = 1$. This implies that even when $\pi(H|\ell) = 1$, so that $\beta_{L,\ell}(\beta, 1) = 1$, $(1-\delta)+\delta\overline{W}(\beta_{H,\ell}(\beta, 1)) \geq \delta\overline{W}(1)$. Therefore, since $\alpha^o_H(H) = 1$, it is optimal for the sender to send the message H when $\theta = \ell$. Also it is strictly optimal for him to send message H when $\theta = h$. That $\pi^o(\beta) = 1$ also implies that $\bar{\pi}(\beta) \geq 1$, so that $\mu(\pi^o, \beta) \geq 1/2$ and $\alpha^o_H = 1$ is optimal for the receiver. As we saw before, since $\pi^o(H|h) = \pi_B(H|h) = 1$, $\alpha^o_L(L) = 1$ is optimal for the receiver. □

Corollary 2. For any $\beta \in (0, 1]$, $\overline{X}(\beta) \geq \overline{X}(0)$

Proof: Fix $\beta \in (0, 1]$ and its corresponding (π^o, α^o, w^o). If $\beta \in (0, \beta^{23})$ then the (IC) is satisfied and the sender is indifferent between messages H and L when $\theta = \ell$. Therefore, the value of this tuple is

$$E(\pi^o, \alpha^o, w^o) = \mu_0[(1 - \delta) + \delta w^o_{H,h}] + (1 - \mu_0)\delta w^o_{L,\ell}. \tag{RPK}$$

Recall that $w^o_{H,h} = \overline{W}(\beta)$ and that $\beta_{L,\ell} > \beta$ so $w^o_{L,\ell} \geq \overline{W}_L(\beta_{L,\ell}) \geq \overline{W}(\beta)$. Thus

$$\overline{X}(\beta) \geq E(\pi^o, \alpha^o, w^o) \geq \mu_0(1-\delta) + \delta\overline{W}(\beta) \geq \overline{X}(0).$$

If $\beta \in [\beta^{23}, 1]$ it is optimal for the sender to send message H when $\theta = \ell$ and thus $v_\ell \geq \delta w_{L,\ell} = \delta\overline{W}(1)$. Therefore

$$E(\pi^o, \alpha^o, w^o) \geq \mu_0[(1-\delta) + \delta w^o_{H,h}] + (1-\mu_0)\delta\overline{W}(1).$$

Since $w^o_{H,h} = \overline{W}(\beta)$,

$$\overline{X}(\beta) \geq E(\pi^o, \alpha^o, w^o) \geq \mu_0[(1-\delta) + \delta\overline{W}(\beta)] + (1-\mu_0)\delta\overline{W}(1) \geq \overline{X}(0). \qquad \square$$

Lemma 8. Let (π, α, w) be an optimal tuple for W at some $\beta \in (0, 1]$. Then $\pi_B \leq \pi(H|\ell) \leq \bar{\pi}(\beta)$ and $\alpha_H(H) = 1$.

Proof: Let (π, α, w) be an admissible tuple for W at β with $\pi(H|\ell) > \bar{\pi}(\beta)$. Then, the receiver's posterior that $\theta = h$ after receiving message H is strictly less than 1/2. Hence, $\alpha_H(H) = 0$. If $\pi(H|h) = 1$, then $\beta_{H,h} = \beta$, $\beta_{H,\ell} < \beta$ (since $\bar{\pi}(\beta) > \pi_B$), and

$$E(\pi, \alpha, w) = \mu_0\delta w_{H,h} + (1-\mu_0)\delta w_{H,\ell} < \delta\overline{W}(\beta).$$

If $\pi(H|h) < 1$, then $\beta_{L,h} = 0$ and it is optimal for the sender to send message L after observing $\theta = h$, so

$$E(\pi, \alpha, w) = \mu_0\delta w_{L,h} + (1-\mu_0)\delta w_{H,\ell} < \delta\overline{W}(\beta).$$

But, for the corresponding (π^o, α^o, w^o) at β, $\beta^o_{L,\ell} = \beta_{L,\ell}(\beta, \pi^o(H|\ell)) > \beta$, and

$$E(\pi^o, \alpha^o, w^o) = \mu_0(1-\delta) + \delta[\mu_0\overline{W}(\beta) + (1-\mu_0)\overline{W}(\beta^o_{L,\ell})] > \mu_0(1-\delta) + \delta\overline{W}(\beta),$$

so $E(\pi^o, \alpha^o, w^o) > E(\pi, \alpha, w)$. Thus, no such admissible tuple (π, α, w) can be optimal.

Now, assume that (π, α, w) is admissible for W at β and that $\pi(H|\ell) \leq \bar{\pi}(\beta)$ and $\alpha_H(H) < 1$. Then

$$E(\pi, \alpha, w) \leq \mu_0\alpha_H(H)(1-\delta) + \delta\overline{W}(\beta) < E(\pi^o, \alpha^o, w^o),$$

so (π, α, w) is not optimal.

Finally, assume that $\pi(H|\ell) < \pi_B$ and $\alpha_H(H) = 1$. Since $\pi(H|\ell) < \pi_B$, $\pi(L|\ell) = 1 - \pi(H|\ell) > 1 - \pi_B > 0$ and recommending L after $\theta = \ell$ is optimal for the sender. Therefore, $\beta_{L,\ell} < \beta$ and $\beta_{H,h} = \beta$ (since $\pi(H|h) = \pi_B(H|h) = 1$), and

$$E(\pi, \alpha, w) \leq \mu_0\alpha_H(H)(1-\delta) + \delta[\mu_0\overline{W}(\beta) + (1-\mu_0)\overline{W}(\beta_{L,\ell})] < \mu_0(1-\delta) + \delta\overline{W}(\beta).$$

But, as we established above, for the corresponding (π^o, α^o, w^o) at β, we have that $E(\pi^o, \alpha^o, w^o) > \mu_0(1-\delta) + \delta\overline{W}(\beta)$. Hence, (π, α, w) is not optimal. $\square$

Lemma 9. Let (π, α, w) be an optimal tuple for W at some $\beta \in (0,1)$. Then $\pi(H|h) = 1$ and $\alpha_L(L) = 1$.

Proof: By contradiction, assume that $\pi(H|h) < 1$. Then, $\pi(L|h) > 0$ and for the tuple to be admissible, it must optimal for the sender to recommend L when $\theta = h$. But $\beta_{L,h} = 0$ and $w_{L,h} \leq \overline{W}(0)$ (since $\pi_B(L|h) = 0$). Therefore, the sender's expected continuation value after observing $\theta = h$ is bounded above by $\delta\overline{W}(0)$. Consider the modified tuple $(\hat{\pi}, \alpha, \hat{w})$ where $\hat{\pi}(m|\ell) = \pi(m|\ell)$ for $m = H, L$, $\hat{\pi}(H,h) = 1$, $\hat{\pi}(L|h) = 0$, $\hat{w}_{m,\ell} = w_{m,\ell}$ for $m = H, L$, $\hat{w}_{H,h} = \overline{W}(\beta) \geq \mu_0$, and $\hat{w}_{L,h} = 0$. Now $\beta_{H,h} = \beta$ and $\hat{w}_{H,h} \in W(\beta_{H,h})$. By previous lemma, $\alpha_H(H) = 1$, so it must be that the receiver's posterior that $\theta = h$ after receiving message H is greater or equal to 1/2. After increasing $\pi(H|h)$ to $\hat{\pi}(H|h) = 1$ this posterior increases and hence $\alpha_H(H) = 1$ remains optimal for the receiver. One can readily verify that this new tuple is also admissible for W at β and the corresponding sender's expected continuation value after observing $\theta = h$ is $\alpha_H(H)(1-\delta) + \delta\overline{W}(\beta) > \delta\mu_0$. Thus $E(\hat{\pi}, \alpha, \hat{w}) > E(\pi, \alpha, w)$, which is a contradiction.

Since $\pi(H|h) = \pi_B(H|h) = 1$, the receiver's posterior that $\omega = \ell$ after receiving the message $m = L$ is $\mathbb{P}[\ell|L] = 1$ and his unique best reply is to take action L, therefore $\alpha_L(L) = 1$. □

Lemma 10. [Exploitation] For any $\beta \in [\beta^{23}, 1]$ the corresponding (π^o, α^o, w^o) is an optimal tuple for W at β and $\pi^o(H|\ell) = 1$.

Proof: By the definition of β^{23}, since $\overline{W}(\beta)$ and $\beta_{H,\ell}(\beta, 1)$ are strictly increasing in β,

$$(1-\delta) + \delta\overline{W}(\beta_{H,\ell}(\beta,1)) \geq \delta\overline{W}(1) \quad \text{for all } \beta \in [\beta^{23}, 1]$$

with strict inequality for $\beta \neq \beta^{23}$.

Let (π, α, w) be an admissible tuple at $\beta \in [\beta^{23}, 1]$ with $\pi(H|h) = 1$. If $\pi(H|\ell) < 1$, then $\beta_{L,\ell} < 1$ and $w_{L,\ell} \leq \overline{W}(\beta_{L,\ell}) \leq \overline{W}(1)$. Since $\bar{\pi}(\beta) \geq 1$, it must be that $\alpha_H(H) = 1$, and since $\pi(L|\ell) = 1 - \pi(H|\ell) > 0$, sending message $m = L$ when $\theta = \ell$ must be optimal for the sender. Therefore

$$E(\pi, \alpha, w) \leq \mu_0[(1-\delta) + \delta\overline{W}(\beta)] + (1-\mu_0)\delta\overline{W}(1),$$

and

$$E(\pi^o, \alpha^o, w^o) = \mu_0[(1-\delta) + \delta\overline{W}(\beta)] + (1-\mu_0)[(1-\delta) + \delta\overline{W}(\beta_{H,\ell}(\beta, 1)) \geq E(\pi, \alpha, w)$$

(with strict inequality for $\beta \neq \beta^{23}$). On the other hand, if $\pi(H|\ell) = 1$, then $w_{H,\ell} \leq \overline{W}(\beta_{H,\ell})$ and

$$E(\pi, \alpha, w) \leq (1-\delta) + \delta[\mu_0\overline{W}(\beta) + (1-\mu_0)\overline{W}(\beta_{H,\ell})] = E(\pi^o, \alpha^o, w^o).$$

Thus (π^o, α^o, w^o) is optimal. □

Lemma 11. [Indifference] For any $\beta \in [0, \beta^{23})$ the corresponding (π^o, α^o, w^o) is an optimal tuple for W at β. Moreover $\pi_B \leq \pi^o(H|\ell) < 1$ and the sender is indifferent between messages H and L when $\theta = \ell$.

Proof: Let (π, α, w) be a tuple such that $\pi(H|\ell) = 1$. We first show that this tuple cannot be admissible at any $\beta \in [0, \beta^{23})$. By the definition of β^{23}, since $\overline{W}(\beta)$ and $\beta_{H,\ell}(\beta, 1)$ are strictly increasing in β,

$$(1-\delta) + \delta\overline{W}(\beta_{H,\ell}(\beta,1)) < \delta\overline{W}(1) \quad \text{for all } \beta \in [0, \beta^{23}).$$

Pick any $\beta \in [0, \beta^{23})$. Since $\alpha_H(H) \leq 1$ and $\pi(H|\ell) = 1$, the sender's expected continuation value when he reports $m = H$ after observing $\theta = \ell$ is less than or equal to

$$(1-\delta) + \delta\overline{W}(\beta_{H,\ell}(\beta,1)).$$

But if the sender reports $m = L$ instead, his continuation value is $\delta\overline{W}(1)$, so π is not a best reply to (α, w) and (π, α, w) is not admissible at β.

Since $X = B(W)$ has a compact graph, for each $\beta \in [0, 1]$ there is an optimal policy. Fix $\beta \in [0, \beta^{23})$ and let (π, α, w) be an optimal policy for W at β and (π^o, α^o, w^o) be the corresponding admissible tuple. By Lemma 8, $\alpha_H(H) = 1$ and $\pi_B \leq \pi(H|\ell) \leq \bar{\pi}(\beta)$, and by the argument above, $\pi(H|\ell) < 1$. Assume by contradiction that $E(\pi, \alpha, w) > E(\pi^o, \alpha^o, w^o)$. Since $\pi(H|\ell) < 1$, sending message L after observing $\theta = \ell$ must be optimal for the sender and thus $(1-\delta) + \delta w_{H,\ell} \leq \delta w_{L,\ell}$ and

$$E(\pi, \alpha, w) = \mu_0(1-\delta) + \delta[\mu_0 w_{H,h} + (1-\mu_0) w_{L,\ell}],$$

where $w_{H,h} \leq \overline{W}(\beta)$, $w_{H,\ell} \leq \overline{W}(\beta_{H,\ell})$ and $w_{L,\ell} \leq \overline{W}(\beta_{L,\ell})$. Increasing $w_{H,h}$ strengthens the incentives for sending message H when $\theta = h$, and thus optimally $w_{H,h} = \overline{W}(\beta)$. We also have that

$$E(\pi^o, \alpha^o, w^o) = \mu_0(1-\delta) + \delta[\mu_0\overline{W}(\beta) + (1-\mu_0) w^o_{L,\ell}],$$

where $w^o_{L,\ell} \in [\overline{W}_L(\beta^o_{H,\ell}), \overline{W}(\beta^o_{H,\ell})]$. Therefore $w_{L,\ell} > w^o_{L,\ell}$, which implies that $\beta_{L,\ell} > \beta^o_{L,\ell}$. Hence, $\pi(H|\ell) > \pi^o(H|\ell)$ and $\beta_{H,\ell} < \beta^o_{H,\ell}$, so $w_{H,\ell} < w^o_{H,\ell}$. By definition, $(1-\delta) + \delta w^o_{H,\ell} = \delta w^o_{L,\ell}$, which implies that $(1-\delta) + \delta w_{H,\ell} < \delta w_{L,\ell}$. But this implies that when $\theta = \ell$, the sender strictly prefers L to H, so (π, α, w), so π with $\pi(L|\ell) = 1 - \pi(H|\ell) \geq 1 - \pi_B > 0$ is not a best reply for the sender and (π, α, w) is not admissible, which is a contradiction. $\square$

Lemma 12. Assume $W \in \mathcal{C}_0^{++}$. Then $\overline{X}(\beta)$ is strictly increasing in $\beta \in [0, 1]$.

Proof: Let $0 < \beta < \hat{\beta}$. We now show that $\overline{X}(\beta) < \overline{X}(\hat{\beta})$. Let (π^o, α^o, w^o) and $(\hat{\pi}^o, \hat{\alpha}^o, \hat{w}^o)$ be the corresponding optimal tuples. Then $\hat{\beta}_{L,\ell} := \beta_{L,\ell}(\hat{\beta}, \hat{\pi}^o) \geq \beta_{L,\ell}(\beta, \pi^o) := \beta_{L,\ell}$. By contradiction, assume that $\hat{\beta}_{L,\ell} < \beta_{L,\ell}$. Then $\hat{\pi}^o(H|\ell) < \pi^o(H|\ell) \leq \bar{\pi}(\beta) < \bar{\pi}(\hat{\beta})$, so $\hat{\beta}_{H,\ell} > \beta_{H,\ell}$ and

$$\overline{W}(\hat{\beta}_{L,\ell}) - \overline{W}_L(\hat{\beta}_{H,\ell}) < \overline{W}(\beta_{L,\ell}) - \overline{W}(\beta_{H,\ell}) \leq (1-\delta)/\delta,$$

contradicting the definition of $\hat{\pi}(H|\ell)$ (by the previous inequality, $\hat{\pi}(H|\ell)$ should be increased to increase $\hat{w}_{L,h}$ and improve $E(\hat{\pi}^o, \hat{\alpha}^o, \hat{w}^o)$). Note that even if $\overline{W}(\beta_{L,\ell}) - \overline{W}_L(\beta_{H,\ell}) > (1-\delta)/\delta$, $\overline{W}(\beta_{L,\ell}) - \overline{W}(\beta_{H,\ell}) \leq (1-\delta)/\delta$ by the definition of $\hat{\pi}(\beta)$.

Previously we also established that $\overline{X}(\beta) \geq \overline{X}(0)$ for all $\beta > 0$. Thus, $\overline{X}(\beta)$ is strictly increasing in $\beta \in [0,1]$. □

Lemma 13. [Monotonicity] $\overline{V}(\beta)$ is weakly increasing in β.

Proof: Construct the equilibrium correspondence V starting with the initial seed

$$W^0 = \{(\beta, w) \mid \beta \in [0,1] \text{ and } w \in [0, 1+\gamma\beta] \text{ for each } \beta \in [0,1]\},$$

where $\gamma > 0$ (small). Clearly $W^0 \in \mathcal{C}_0^{++}$. Let $W^{k+1} = \tilde{B}(W^k)$ for $k = 0, 1, \dots$ Then, $W^k \in \mathcal{C}_0^{++}$ for each $k \geq 1$. Since $V = \lim_{k\to\infty} W^k$, we have that $\overline{V}(\beta)$ is weakly increasing in β. □

Remark: If $W \in \mathcal{C}_0^{++}$, then for each $\beta \in (0,1]$ the corresponding (π^o, α^o, w^o) is essentially the unique optimal tuple. We have chosen for convenience $w^o_{L,h} = \underline{W}(0)$, but typically there is an interval of continuation values $w^o_{L,h} \in W(0)$ that provides the proper incentives for the sender not to lie and send the message $m = H$ when $\theta = h$.

In Lemma 13, for each $\beta \in [0,1]$ denote by (π^k, α^k, w^k) the (unique) optimal tuple for W^k that delivers value $\overline{W}^{k+1}(\beta)$. Since the sequence $\{(\pi^k, \alpha^k, w^k)(\beta)\}_{k\geq 0}$ stays in a compact set, it has a convergent subsequence. Let (π, α, w) be a limit point of the sequence. If $\overline{V}$ failed to be strictly increasing, there might be multiple optimal tuples for a set of reputations β. Nevertheless, one can show that for each $\beta \in [0,1]$, (π, α, w) is an optimal tuple for V at β. This limit tuple preserves many of the properties of (π^k, α^k, w^k). In particular, $\pi(H|h) = 1$, $\alpha_L(L) = 1$, $\pi_B \leq \pi(H|\ell) \leq \bar{\pi}(\beta)$ and $\alpha_H(H) = 1$.

Since V is upper hemicontinuous and $\overline{V}$ is weakly increasing, $\overline{V}$ is right continuous. For any β, define

$$\overline{V}_L(\beta) = \lim_{\beta' \uparrow \beta} \overline{V}(\beta'), \quad \Delta(\beta) = \overline{V}(\beta) - \overline{V}_L(\beta) \quad \text{and} \quad I(\beta) = [\overline{V}_L(\beta), \overline{V}(\beta)].$$

Then, $\Delta(\beta) \geq 0$ and $\overline{V}$ is discontinuous at β if and only if $\Delta(\beta) > 0$. Also, for any increasing sequence $\{\beta_k\}$ that converges to β and any sequence $\{w_k\}$ with $w_k \in I(\beta_k)$, $w_k \to \overline{V}_L(\beta)$.

Lemma 14. If $\overline{V}$ is discontinuous at β such that $\pi(\beta) < 1$, then $\overline{V}$ is discontinuous at $\beta_{L,\ell} = \beta_{L,\ell}(\beta, \pi(\beta))$. Moreover, $\Delta(\beta_{L,\ell}) \geq B\Delta(\beta)$, where $B = [1-\delta\mu_0]/[\delta(1-\mu_0)] > 1$.

Proof: Assume $\overline{V}$ is discontinuous at β and that $(\pi(\beta), \alpha, w)$ is an optimal tuple at β with minimal $\pi(\beta)$ that supports $\overline{V}(\beta)$, where $\pi(\beta) < 1$. If $\pi(\beta) < \bar{\pi}(\beta)$ then $\alpha = 1$, and if $\pi(\beta) = \bar{\pi}(\beta)$, the gap condition ensures that $\alpha = 1$ in this case as well. Moreover,

$$w_{L,h} = 0, \quad w_{H,h} = \overline{V}(\beta), \quad w_{L,\ell} \in I(\beta_{L,\ell}) \quad \text{and} \quad w_{H,\ell} = w_{L,\ell} - \frac{1-\delta}{\delta}.$$

Since $\overline{V}$ is weakly increasing, $\overline{V}(\beta_{L,\ell}(\beta, \hat{\pi})) - \overline{V}_L(\beta_{H,\ell}(\beta, \hat{\pi}))$ is weakly increasing in $\hat{\pi}$. Whether $\pi(\beta) = \bar{\pi}(\beta)$ or $\pi(\beta) < \bar{\pi}(\beta)$, because $\pi(\beta)$ is minimal,

$$\overline{V}(\beta_{L,\ell}(\beta, \hat{\pi})) - \overline{V}_L(\beta_{H,\ell}(\beta, \hat{\pi})) < \frac{1-\delta}{\delta} \qquad \text{for all } \hat{\pi} < \pi(\beta).$$

For any $\beta' < \beta$, let π'_- and π'_+ be such that

$$\beta_{H,\ell}(\beta', \pi'_-) = \beta_{H,\ell}(\beta, \pi(\beta)) \quad \text{and} \quad \beta_{L,\ell}(\beta', \pi'_+) = \beta_{L,\ell}(\beta, \pi(\beta)).$$

One can check that as a function of β' (keeping β fixed),

$$\pi'_-(\beta') = \pi(\beta)\left(\frac{1-\beta}{\beta}\right)\left(\frac{\beta'}{1-\beta'}\right).$$

Since $\beta' < \beta$, $\pi'_- < \pi(\beta) < \pi'_+$. Moreover, $\pi'_- \uparrow \pi(\beta)$ and $\pi'_+ \downarrow \pi(\beta)$ as $\beta' \uparrow \beta$. Let $(\pi(\beta'), \alpha', w')$ be the corresponding optimal tuple at β' that supports $\overline{V}(\beta')$ with minimal $\pi(\beta')$. We now show that $\pi'_- \leq \pi(\beta') \leq \pi'_+$. Since $\pi(\beta) \leq \bar{\pi}(\beta) = (\pi^* - \pi_B)/(1-\beta)$,

$$\pi'_-(\beta') \leq \frac{\beta'}{\beta}\frac{\pi^* - \pi_B}{1-\beta'} < \frac{\pi^* - \pi_B}{1-\beta'} = \bar{\pi}(\beta').$$

For any $\tilde{\pi} < \pi'_-$, there exists $\hat{\pi} \in (\tilde{\pi}, \pi(\beta))$ such that $\beta_{H,\ell}(\beta, \hat{\pi}) = \beta_{H,\ell}(\beta', \tilde{\pi})$. Then, $\beta_{L,\ell}(\beta, \hat{\pi}) > \beta_{L,\ell}(\beta', \tilde{\pi})$ and

$$\overline{V}(\beta_{L,\ell}(\beta', \tilde{\pi})) - \overline{V}_L(\beta_{H,\ell}(\beta', \tilde{\pi})) \leq \overline{V}(\beta_{L,\ell}(\beta, \hat{\pi})) - \overline{V}_L(\beta_{H,\ell}(\beta, \hat{\pi})) < \frac{1-\delta}{\delta}.$$

As $\tilde{\pi} < \pi'_- < \bar{\pi}(\beta')$, this implies that $\tilde{\pi} < \pi(\beta')$. Therefore $\pi(\beta') \geq \pi'_-$.

Finally, recall that $\beta_{L,\ell}(\beta', \pi'_+) = \beta_{L,\ell}(\beta, \pi(\beta))$. Hence, if $\pi(\beta) < \bar{\pi}(\beta)$, then

$$\overline{V}(\beta_{L,\ell}(\beta', \pi'_+)) - \overline{V}_L(\beta_{H,\ell}(\beta', \pi'_+)) \geq \overline{V}(\beta_{L,\ell}(\beta, \pi(\beta))) - \overline{V}_L(\beta_{H,\ell}(\beta, \pi(\beta))) = \frac{1-\delta}{\delta},$$

which implies that $\pi(\beta') \le \pi'_+$. And if $\pi(\beta) = \bar{\pi}(\beta)$, then we also have

$$\pi_+ > \bar{\pi}(\beta) > \bar{\pi}(\beta') \ge \pi(\beta').$$

The inequalities above imply that $\pi(\beta') \to \pi(\beta)$ as $\beta' \uparrow \beta$. Therefore, for any $\beta' < \beta$ with $\beta - \beta'$ sufficiently small, $\pi(\beta') < 1$. Moreover $\alpha' = 1$ and $w'_{L,\ell} \in I(\beta'_{L,\ell})$, where $\beta'_{L,\ell} = \beta_{L,\ell}(\beta', \pi(\beta'))$, and

$$w'_{L,h} = 0, \quad w'_{H,h} = \overline{V}(\beta') \quad \text{and} \quad w'_{H,\ell} = w'_{L,\ell} - \frac{1-\delta}{\delta}.$$

Since $\beta'_{L,\ell} \uparrow \beta_{L,\ell}$, $w'_{L,\ell} \to \overline{V}_L(\beta_{L,\ell}) = \overline{V}(\beta_{L,\ell}) - \Delta(\beta_{L,\ell})$ as $\beta' \uparrow \beta$. We need to show that $\Delta(\beta_{L,\ell}) > 0$. By RPK,

$$\begin{aligned}\overline{V}(\beta) &= \mu_0(1-\delta) + \delta[\mu_0\overline{V}(\beta) + (1-\mu_0)w_{L,\ell}] \\ \overline{V}(\beta') &= \mu_0(1-\delta) + \delta[\mu_0\overline{V}(\beta') + (1-\mu_0)w'_{L,\ell}].\end{aligned}$$

Since $w_{L,\ell} \in I(\beta_{L,\ell})$, $\overline{V}(\beta_{L,\ell}) - \Delta(\beta_{L,\ell}) \ge w_{L,\ell} - \Delta(\beta_{L,\ell})$. Thus, taking the limit of the second identity above as $\beta' \uparrow \beta$, we get

$$\begin{aligned}\overline{V}(\beta) - \Delta(\beta) &= \mu_0(1-\delta) + \delta[\mu_0(\overline{V}(\beta) - \Delta(\beta)) + (1-\mu_0)(\overline{V}(\beta_{L,\ell}) - \Delta(\beta_{L,\ell}))] \\ &\ge \overline{V}(\beta) - \delta\mu_0\Delta(\beta) - \delta(1-\mu_0)\Delta(\beta_{L,\ell}),\end{aligned}$$

which implies that $\Delta(\beta_{L,\ell}) \ge B\Delta(\beta)$. □

Lemma 15. If $\overline{V}$ is discontinuous at some $\beta > \beta_d$, then $\overline{V}$ is discontinuous at $\beta_{H,\ell} = \beta_{H,\ell}(\beta, \pi(\beta))$. Moreover, $\Delta(\beta_{H,\ell}) \ge B\Delta(\beta)$.

Proof: Since $\pi(\beta) \le 1 < \bar{\pi}(\beta)$, β is not in Region 1 and $w_{H,\ell} \in I(\beta_{H,\ell})$.

Let $\beta' < \beta$ and $(\pi(\beta'), \alpha', w')$ be as defined above. As $\bar{\pi}(\beta')$ is continuous in β', $\pi(\beta') < 1$ for $\beta - \beta'$ sufficiently small. Since $\beta'_{H,\ell} \uparrow \beta_{H,\ell}$, $w'_{H,\ell} \to \overline{V}_L(\beta_{H,\ell}) = \overline{V}(\beta_{H,\ell}) - \Delta(\beta_{H,\ell})$ as $\beta' \uparrow \beta$. We need to show that $\Delta(\beta_{H,\ell}) > 0$. By LPK,

$$\begin{aligned}\overline{V}(\beta) &= (1-\delta) + \delta[\mu_0\overline{V}(\beta) + (1-\mu_0)w_{H,\ell}] \\ \overline{V}(\beta') &= (1-\delta) + \delta[\mu_0\overline{V}(\beta') + (1-\mu_0)w'_{H,\ell}].\end{aligned}$$

Since $w_{H,\ell} \in I(\beta_{H,\ell})$, $\overline{V}(\beta_{H,\ell}) - \Delta(\beta_{H,\ell}) \ge w_{H,\ell} - \Delta(\beta_{H,\ell})$. Taking the limit of the second identity above as $\beta' \uparrow \beta$, we get

$$\begin{aligned}\overline{V}(\beta) - \Delta(\beta) &= (1-\delta) + \delta[\mu_0(\overline{V}(\beta) - \Delta(\beta)) + (1-\mu_0)(\overline{V}(\beta_{L,\ell}) - \Delta(\beta_{H,\ell}))] \\ &\ge \overline{V}(\beta) - \delta\mu_0\Delta(\beta) - \delta(1-\mu_0)\Delta(\beta_{H,\ell}),\end{aligned}$$

which implies that $\Delta(\beta_{H,\ell}) \ge B\Delta(\beta)$. □

Lemma 16. [Continuity] $\overline{V}(\beta)$ is continuous in β.

Proof: By contradiction, suppose that $\overline{V}$ has a discontinuity of size Δ_0 at some β_0 outside of Region 3. Consider the infinite sequence $\{\beta_k\}_{k\geq 0}$ defined recursively as follows: if $\beta_k > \beta_d$, then

$$\beta_{k+1} = \beta_{H,\ell}(\beta_k, \pi(\beta_k)),$$

while if $\beta_k \leq \beta_d$, then

$$\beta_{k+1} = \beta_{L,\ell}(\beta_k, \pi(\beta_k)).$$

By Claims 1 and 2, $\overline{V}$ has a discontinuity of size Δ_k at every β_k along this sequence, where

$$\Delta_{k+1} \geq B\Delta_k \qquad \text{for all } k \geq 0.$$

Hence, $\Delta_k \geq B^k \Delta_0 \to \infty$. This is a contradiction since $\overline{V}(\beta_k) \in [0, 1]$ for all k. □

Proof of Proposition 2. Follows directly from Lemmas 13 and 16. □

B. Behavioral Convergence With One Behavioral Type

B.1. Pointwise Behavioral Convergence

This section focuses on the upper boundary and optimal tuples. Recall that we denote by $\pi(\beta, \delta)$ (or simply $\pi(\beta)$ when no confusion arises) the lying probability $\pi(H|\ell)$ associated with an optimal tuple (π, α, w) at reputation β.

Lemma 17. For $\beta \in [0, \beta^{23})$, $\beta_{H,\ell}(\beta, \pi(\beta))$ and $\beta_{L,\ell}(\beta, \pi(\beta))$ are increasing functions of β.

Proof: Let $\beta \in [0, \beta^{23})$ be such that for some $\epsilon > 0$, $\pi(\beta') = \bar{\pi}(\beta')$ for all $\beta' \in [\beta, \beta+\epsilon]$. Then, for any $\beta' \in (\beta, \beta + \epsilon]$

$$\beta_{H,\ell}(\beta, \bar{\pi}(\beta)) = \lambda_0 \beta < \lambda_0 \beta' = \beta_{H,\ell}(\beta', \bar{\pi}(\beta')) \quad \text{and}$$
$$\beta_{L,\ell}(\beta, \bar{\pi}(\beta)) = \lambda_1 \beta < \lambda_1 \beta' = \beta_{L,\ell}(\beta', \bar{\pi}(\beta')).$$

Now consider $\beta \in [0, \beta^{23})$ such that for some $\epsilon > 0$, $\pi(\beta') < \bar{\pi}(\beta')$ for all $\beta' \in (\beta, \beta+\epsilon)$. Assume by contradiction that for some $\beta' \in (\beta, \beta + \epsilon)$ we have that $\beta_{H,\ell}(\beta, \pi(\beta)) \geq \beta_{H,\ell}(\beta', \pi(\beta'))$. This implies that $\pi(\beta') > \pi(\beta)$. But $\beta' > \beta$ and $\pi(\beta') > \pi(\beta)$ imply that $\beta_{L,\ell}(\beta, \pi(\beta)) < \beta_{L,\ell}(\beta', \pi(\beta'))$. By definition of $\pi(\beta)$,

$$\delta[\overline{V}(\beta_{L,\ell}(\beta', \pi(\beta'))) - \overline{V}(\beta_{H,\ell}(\beta', \pi(\beta')))]$$
$$> \delta[\overline{V}(\beta_{L,\ell}(\beta, \pi(\beta))) - \overline{V}(\beta_{H,\ell}(\beta, \pi(\beta)))] = 1 - \delta,$$

a contradiction. Therefore $\beta_{H,\ell}(\beta, \pi(\beta)) < \beta_{H,\ell}(\beta', \pi(\beta'))$ for all $\beta' \in (\beta, \beta+\epsilon)$. Given

$$\delta[\overline{V}(\beta_{L,\ell}(\beta', \pi(\beta'))) - \overline{V}(\beta_{H,\ell}(\beta', \pi(\beta')))] = 1-\delta,$$

by definition of $\pi(\beta')$ it must be that $\overline{V}(\beta_{L,\ell}(\beta', \pi(\beta'))) > \overline{V}(\beta_{L,\ell}(\beta, \pi(\beta)))$ and $\beta_{L,\ell}(\beta', \pi(\beta')) > \beta_{L,\ell}(\beta, \pi(\beta))$. □

Fixing β_0, let $\pi_0 = \pi(\beta_0)$,

$$\beta_H = \beta_{H,\ell}(\beta_0, \pi_0) \quad \text{and} \quad \beta_L = \beta_{L,\ell}(\beta_0, \pi_0).$$

Section 6.2 considers a reputation maintenance strategy based at β_0. By Lemma 17, $\beta_H \le \beta_{H,\ell}(\beta, \pi(\beta)) < \beta$ for all $\beta \in [\beta_0, \beta_L]$ and $\beta < \beta_{L,\ell}(\beta, \pi(\beta)) < \beta_L$ for all $\beta \in [\beta_H, \beta_0)$. Also $\beta_{H,h}(\beta, \pi(\beta)) = \beta$. Hence, when the sender follows this strategy, his reputation stays in the interval $[\beta_H, \beta_L]$ in every period.

In the spirit of dynamic programming, for an arbitrary continuation value function W from $[\beta_H, \beta_L]$ to $[0,1]$, compute for each $\beta \in [\beta_H, \beta_L]$ the value of doing one round of reputation maintenance (that is, play L if $\beta \le \beta_0$ and $\theta = \ell$, play H otherwise), and using continuation values given by W.

Formally, let $\mathbb{W} = [0,1]^{[\beta_H, \beta_L]}$ be the set of all functions W from $[\beta_H, \beta_L]$ to $[0,1]$. Endow $\mathbb{W}$ with the sup norm:

$$||W|| = \sup\{|W(\beta)| \mid \beta \in [\beta_H, \beta_L]\}.$$

Let $T^\pi : \mathbb{W} \to \mathbb{W}$ be the map

$$T^\pi(W)(\beta) = \begin{cases} \mu_0(1-\delta) + \delta[\mu_0 W(\beta) + (1-\mu_0)W(\beta_{L,\ell}(\beta, \pi(\beta)))] & \text{for } \beta \in [\beta_H, \beta_0) \\ (1-\delta) + \delta[\mu_0 W(\beta) + (1-\mu_0)W(\beta_{H,\ell}(\beta, \pi(\beta)))] & \text{for } \beta \in [\beta_0, \beta_L]. \end{cases}$$

One can easily check that T^π is a contraction and therefore it has a unique fixed point. Denote by $V^\pi(\beta, \delta)$ this fixed point to make explicit that it depends on δ as well.

In the next Lemma (and below), if $[\beta_H, \beta_0)$ does not intersect Region 3, L is an optimal recommendation for the sender for any $\beta \in [\beta_H, \beta_0)$. For simplicity, we will assume that $\beta_0 \le \beta^{23}$, thus ensuring that $[\beta_H, \beta_0)$ does not intersect Region 3 (if $\beta_0 = \beta^{23}$, recommending L is also optimal, even if in our equilibrium this is never done).

Lemma 18. Assume $\beta_0 \le \beta^{23}$. Then

1. $V^\pi(\beta, \delta) \ge \overline{V}(\beta, \delta)$ for all $\beta \in [\beta_H, \beta_L]$. Conversely, if $(\beta_0, \beta_L]$ does not intersect Region 1 then $\overline{V}(\beta, \delta) = V^\pi(\beta, \delta)$ for all $\beta \in [\beta_H, \beta_L]$.
2. $V^\pi(\beta, \delta)$ is weakly increasing in $\beta \in [\beta_H, \beta_L]$.

3. Let $\pi_0(\beta) = \pi(\beta_0)$ for all $\beta \in [\beta_H, \beta_L]$. Then $V^{\pi_0}(\beta, \delta) \geq V^{\pi}(\beta, \delta)$ for all $\beta \in [\beta_H, \beta_L]$.

Proof: (1) We first show that if $W(\beta) \geq \overline{V}(\beta, \delta)$ for all $\beta \in [\beta_H, \beta_L]$, then $T^{\pi}(W)(\beta) \geq \overline{V}(\beta, \delta)$ for all $\beta \in [\beta_H, \beta_L]$. If $\beta \in [\beta_H, \beta_0)$, by right promise keeping we have that

$$\begin{aligned}\overline{V}(\beta, \delta) &= \mu_0(1-\delta) + \delta\left[\mu_0\overline{V}(\beta, \delta) + (1-\mu_0)\overline{V}(\beta_{L,\ell}, \delta)\right] \\ &\leq \mu_0(1-\delta) + \delta\left[\mu_0 W(\beta) + (1-\mu_0)W(\beta_{L,\ell})\right] = T^{\pi}(W)(\beta).\end{aligned}$$

If $\beta \in [\beta_0, \beta_L]$, by left promise keeping there exists $w_{H,\ell} \leq \overline{V}(\beta_{H,\ell}, \delta)$ such that

$$\begin{aligned}\overline{V}(\beta, \delta) &= (1-\delta) + \delta\left[\mu_0\overline{V}(\beta, \delta) + (1-\mu_0)w_{H,\ell}\right] \\ &\leq (1-\delta) + \delta\left[\mu_0 W(\beta) + (1-\mu_0)W(\beta_{H,\ell})\right] = T^{\pi}(W)(\beta).\end{aligned}$$

Since $\{W \in \mathbb{W} \mid W(\beta) \geq \overline{V}(\beta, \delta)$ for all $\beta \in [\beta_H, \beta_L]\}$ is a closed set in $(\mathbb{W}, ||\cdot||)$, $V^{\pi}(\beta, \delta) \geq \overline{V}(\beta, \delta)$ for all $\beta \in [\beta_H, \beta_L]$.

Now, assume that $(\beta_0, \beta_L]$ does not intersect Region 1. Then $w_{H,\ell} = \overline{V}(\beta_{H,\ell}, \delta)$ and $\overline{V}(\beta, \delta) = T^{\pi}(\overline{V}(\cdot, \delta))(\beta)$ for all $\beta \in [\beta_H, \beta_L]$. That is $\overline{V}(\cdot, \delta)$ is a fixed point of T^{π}, and therefore $\overline{V} = V^{\pi}$.

(2) Let

$$\mathbb{W}_I = \{W \in \mathbb{W} \mid W(\beta) \text{ is weakly increasing and } W(\beta_L) - W(\beta_H) \leq (1-\delta)/\delta\}.$$

We now show that if $W \subset \mathbb{W}_I$ then $T^{\pi}(W) \in \mathbb{W}_I$. Assume that $\beta_L \leq \beta < \hat{\beta} < \beta_0$. Then $\beta_{L,\ell} < \hat{\beta}_{L,\ell}$, and therefore

$$\begin{aligned}T^{\pi}(W)(\beta) &= \mu_0(1-\delta) + \delta\left[\mu_0 W(\beta) + (1-\mu_0)W(\beta_{L,\ell})\right] \\ &\leq \mu_0(1-\delta) + \delta\left[\mu_0 W(\hat{\beta}) + (1-\mu_0)W(\hat{\beta}_{L,\ell})\right] = T^{\pi}(W)(\hat{\beta}).\end{aligned}$$

The case $\beta_0 \leq \beta < \hat{\beta} \leq \beta_L$ is similar. Finally, assume that $\beta_L \leq \beta < \beta_0 \leq \hat{\beta} \leq \beta_L$. Then, $(1-\delta) + \delta W(\hat{\beta}_{H,\ell}) \geq W(\beta_{L,\ell})$ and therefore

$$\begin{aligned}T^{\pi}(W)(\beta) &= \mu_0(1-\delta) + \delta\left[\mu_0 W(\beta) + (1-\mu_0)W(\beta_{L,\ell})\right] \\ &\leq (1-\delta) + \delta\left[\mu_0 W(\hat{\beta}) + (1-\mu_0)W(\hat{\beta}_{H,\ell})\right] = T^{\pi}(W)(\hat{\beta}).\end{aligned}$$

Since $\mathbb{W}_I$ is a closed set in $(\mathbb{W}, ||\cdot||)$, $V^{\pi} \in \mathbb{W}_I$.

(3) Let

$$\mathbb{W}_0 = \{W \in \mathbb{W} \mid W(\beta) \geq V^{\pi}(\beta, \delta) \text{ for all } \beta \in [\beta_H, \beta_L]\}.$$

Note that for any $\beta \in [\beta_H, \beta_0)$, $\pi(\beta) \leq \pi(\beta_0)$ and thus $\beta_{L,\ell} \leq \beta_{L,\ell}(\beta, \pi_0(\beta))$, and similarly, for any $\beta \in [\beta_0, \beta_L]$, $\pi(\beta) \geq \pi(\beta_0)$ and thus $\beta_{H,\ell} \leq \beta_{L,\ell}(\beta, \pi_0(\beta))$. Thus, it is easy to check that $T^{\pi_0}(W)(\beta) \geq W(\beta)$ for all $\beta \in [\beta_H, \beta_L]$. Since $\mathbb{W}_0$ is a closed set in $(\mathbb{W}, ||\cdot||)$, $V^{\pi_0} \in \mathbb{W}_0$. □

Similarly, one can easily show that if W is continuous, $T^{\pi}(W)$ is continuous. Hence, $V^{\pi}(\beta, \delta)$ is continuous in β.

Value Maintenance: For any $V \in (\mu_0, 1)$, let V^+ denote the solution of the RPK equation

$$V = \mu_0(1-\delta) + \delta[\mu_0 V + (1-\mu_0)V^+];$$

and V^- denote the solution of the LPK equation

$$V = (1-\delta) + \delta[\mu_0 V + (1-\mu_0)V^-].$$

Pick any $\beta_0 \in (0,1)$ and let $V_0 = \overline{V}(\beta_0)$. The value maintenance strategy at value V_0 is defined sequentially in every period t as follows: starting at $t = 0$, if $V_t < V_0$ and $\theta_t = \ell$, recommend L and set $V_{t+1} = V_t^+$; if $V_t \geq V_0$ and $\theta_t = \ell$, recommend H and set $V_{t+1} = V_t^-$; and if $\theta_t = h$, recommend H and set $V_{t+1} = V_t$. The strategy generates a random path $\{V_t\}$ of continuation values. Note that $V_t \in [V_0^-, V_0^+]$ for all t.

Lemma 19. For $\beta_0 \in (0,1)$ let β_H and β_L be as defined above and $V_0 = \overline{V}(\beta_0)$. Assume $(\beta_0, \beta_L]$ does not intersect Region 1 and $\beta_0 \leq \beta^{23}$. Then the value maintenance strategy and the reputation maintenance strategy coincide. That is, if the value maintenance strategy at value V_0 generates the value path $\{V_t\}$ and the reputation maintenance strategy starting at β_0 generates the path of posteriors $\{\beta_t\}$, then $V_t = \overline{V}(\beta_t)$ for all t. Moreover, $V_t \geq (<) V_0$ if and only if $\beta_t \geq (<) \beta_0$, so the two strategies generate the same recommendations in each period t.

Proof: We prove the result by induction. By definition, $V_0 = \overline{V}(\beta_0)$. Assume $V_t = \overline{V}(\beta_t)$. Since $\overline{V}(\beta)$ is strictly increasing in β, $V_t \geq (<) V_0$ if and only if $\beta_t \geq (<)\beta_0$, and the two strategies generate the same recommendations in period t. By Proposition 3, all recommendations (H or L) are accepted for sure by the receiver. If $\beta_{t+1} < \beta_t$ (or, equivalently, if $V_{t+1} < V_t$), it must be that $\beta_t \geq \beta_0$ ($V_t \geq V_0$), $\theta_t = \ell$ and the sender lies, so $\beta_{t+1} = \beta_{H,\ell}(\beta_t, \pi(\beta_t))$ and $V_{t+1} = V_t^-$. Since $(\beta_0, \beta_L]$ does not intersect Region 1 and V_t^- satisfies LPK, $V_{t+1} = \overline{V}(\beta_{t+1})$. If $\beta_{t+1} > \beta_t$ ($V_{t+1} > V_t$), it must be that $\beta_t < \beta_0$ ($V_t < V_0$), $\theta_t = \ell$ and the sender reports truthfully, so $\beta_{t+1} = \beta_{L,\ell}(\beta_t, \pi(\beta_t))$ and $V_{t+1} = V_t^+$. Again, since V_t^+ satisfies RPK, $V_{t+1} = \overline{V}(\beta_{t+1})$. □

Given the previous Lemma, if the sender follows the value maintenance strategy starting at $V_0 = \overline{V}(\beta_0)$, the random path of posteriors $\{\beta_t\}$ can be subsequently constructed as follows: if $V_{t+1} > V_t$, let $\beta_{t+1} = \beta_{L,\ell}(\beta_t, \pi_t)$; if $V_{t+1} < V_t$, let $\beta_{t+1} = \beta_{H,\ell}(\beta_t, \pi_t)$; and if $V_{t+1} = V_t$, let $\beta_{t+1} = \beta_t$, where $\pi_t = \pi(\beta_t)$.

Definition 7. Let $f : [\pi^*, 1) \to (0,1)$ be the function

$$f(\pi_0) = \frac{\log\left[\frac{1-\pi_B}{1-\pi_0}\right]}{\log\left[\frac{\pi_0(1-\pi_B)}{(1-\pi_0))\pi_B}\right]}.$$

One can easily check that for any $\pi_0 \in [\pi^*, 1)$, $F = f(\pi_0)$ if and only if $F \in (0,1)$ is the unique solution to the equation

$$\left(\frac{\pi_0}{\pi_B}\right)^F \left(\frac{1-\pi_0}{1-\pi_B}\right)^{1-F} = 1.$$

Lemmas 20 and 21 below make the counterfactual assumption that $\pi(\beta)$ remains constant in $[\beta_H, \beta_L]$. Lemma 22 considers the more general case where $\pi(\beta)$ is weakly increasing in β.

Lemma 20. For $\beta_0 \in (0,1)$ let β_H and β_L be as defined above. Assume that $\pi(\beta) = \pi_0$ for all $\beta \in [\beta_H, \beta_L]$, that $(\beta_0, \beta_L]$ does not intersect Region 1, and that $\beta_0 \leq \beta^{23}$. The lying frequency induced by the reputation maintenance strategy is $F = f(\pi_0)$.

Proof: Let $\{\beta_t\}$ be the sequence of posteriors generated by the reputation maintenance strategy starting at β_0. As in the previous lemma, if after t periods the sender has lied n times and told the truth m times, then

$$\beta_t = \frac{\beta_0 \pi_B^n (1-\pi_B)^m}{\beta_0 \pi_B^n (1-\pi_B)^m + (1-\beta_0)\pi_0^n (1-\pi_0)^m} = \frac{\beta_0}{\beta_0 + (1-\beta_0)L_{n,m}}$$

$$\text{where} \quad L_{n,m} = \frac{\pi_0^n (1-\pi_0)^m}{\pi_B^n (1-\pi_B)^m} \quad \text{is the likelihood ratio.}$$

Let

$$R = \log\left[\frac{1-\pi_B}{1-\pi_0}\right] \Big/ \log\left[\frac{\pi_0}{\pi_B}\right].$$

Then $F = R/[R+1]$. Since $\beta_t \in [\beta_H, \beta_L]$,

$$\frac{\pi_0}{\pi_B} = L_{1,0} \geq L_{n,m} \geq L_{0,1} = \frac{1-\pi_0}{1-\pi_B} \iff \frac{n-1}{m} \leq R \leq \frac{n}{m-1}$$
$$\iff \frac{n-1}{n+m-1} \leq F \leq \frac{n}{n+m-1}.$$

Therefore,

$$F \leq \frac{n}{n+m-1} \leq F + \frac{1}{n+m-1}.$$

As $t \to \infty$, $n \to \infty$ and $m \to \infty$, and thus the frequency of lies $n/[n+m]$ up to period t converges to the lying frequency and $F = f(\pi_0)$. $\square$

Definition 8. Let

$$\varphi(v) = \frac{v - \mu_0}{1 - \mu_0} \qquad \text{for} \quad v \in [\mu_0, 1].$$

The previous lemma establishes a direct relationship between the equilibrium strategy and the frequency of lying. The next lemma establishes a similar relationship between equilibrium value and frequency of lying. The lemma shows that the frequency F with which the rational sender lies in the periods where $\theta_t = \ell$ when he follows the value maintenance strategy at value v is given (approximately) by $\varphi(v)$.

Lemma 21. For $\beta_0 \in (0,1)$ let β_H and β_L be as defined above. Assume that $\pi(\beta) = \pi_0$ for all $\beta \in [\beta_H, \beta_L]$, that $(\beta_0, \beta_L]$ does not intersect Region 1, and that $\beta_0 \leq \beta^{23}$. Let F_t be the sender's lying frequency up to period t when he follows the value maintenance strategy at value $V_0 = \overline{V}(\beta_0)$. Then, for any $\delta \geq 2/(3 - V_0)$,

$$|F_t - \varphi(V_0)| \leq \frac{2}{1 - V_0}\frac{1-\delta}{\delta} \qquad \text{and} \qquad |\mu_0 + (1-\mu_0)F_t - V_0| \leq \frac{2}{1 - V_0}\frac{1-\delta}{\delta}.$$

Proof: By LPK and RPK, each time $\theta_t = \ell$, $V_{t+1} = V_t - Ad_t$ when the sender reports H, and $V_{t+1} = V_t + Au_t$ when he reports L, where

$$d_t = 1 - V_t, \quad u_t = V_t - \mu_0 \qquad \text{and} \qquad A = \frac{1}{1-\mu_0}\left[\frac{1-\delta}{\delta}\right].$$

Let $V_0^+ = V_0 + Au_0$ and $V_0^- = V_0 - Ad_0$. Note that $V_0^+ < V_0 + \frac{1-\delta}{\delta}$ and $V_0^- > V_0 - \frac{1-\delta}{\delta}$. By construction, $V_t \in [V_0^-, V_0^+]$ for all t. Therefore $u_t \in [V_0 - \mu_0 - \frac{1-\delta}{\delta}, V_0 - \mu_0]$ and $d_t \in [1 - V_0 - \frac{1-\delta}{\delta}, 1 - V_0]$ for all t.

Let

$$s_t = \sum_{\{j<t, \theta_j = \ell, V_j < V_0\}} u_j - \sum_{\{j<t, \theta_j = \ell, V_j \geq V_0\}} d_j.$$

Then, $V_t = V_0 + As_t$. Since $V_t \in [V_0^-, V_0^+]$, it must be that $s_t \in [-(1-V_0), V_0 - \mu_0]$ for all t.

Assume after t periods, there have been n periods where the sender lied and m periods where the sender told the truth. Then

$$m\left(V_0 - \mu_0 - \frac{1-\delta}{\delta}\right) - n(1 - V_0) < s_t < m(V_0 - \mu_0) - n\left(1 - V_0 - \frac{1-\delta}{\delta}\right),$$

and therefore

$$m(V_0 - \mu_0) - n(1 - V_0) - m\frac{1-\delta}{\delta} < V_0 - \mu_0 \qquad \text{and}$$
$$m(V_0 - \mu_0) - n(1 - V_0) + n\frac{1-\delta}{\delta} > -(1 - V_0).$$

Recall that $\varphi(V_0) = (V_0 - \mu_0)/(1 - \mu_0)$. Dividing these inequalities by $m(1 - V_0)$ we get

$$-\frac{1}{m} - \frac{1}{1 - V_0}\frac{n}{m}\frac{1-\delta}{\delta} < \frac{V_0 - \mu_0}{1 - V_0} - \frac{n}{m} < \frac{1}{m}\frac{V_0 - \mu_0}{1 - V_0} + \frac{1}{1 - V_0}\frac{1-\delta}{\delta}.$$

The left inequality implies that when $m \to \infty$ (as $t \to \infty$),

$$\frac{n}{m}\left[1 - \frac{1}{1 - V_0}\frac{1-\delta}{\delta}\right] \leq \frac{V_0 - \mu_0}{1 - V_0} < 1.$$

The term in square brackets is greater or equal to $1/2$ when $\delta > 2/(3 - V_0)$. Hence $\frac{n}{m} \leq 2$ as $m \to \infty$ for all $\delta \geq 2/(3 - V_0)$. Let $R_t = n/m$. Then, for t sufficiently large (so m is large too),

$$\left|\frac{V_0 - \mu_0}{1 - V_0} - R_t\right| \leq \frac{2}{1 - V_0}\frac{1-\delta}{\delta} \quad \text{for all} \quad \delta \geq \frac{2}{3 - V_0}.$$

The lying frequency up to period t is $F_t = n/(n+m) = \xi(R_t)$, where $\xi(x) = x/(x+1)$. Since $|\xi(x) - \xi(y)| \leq |x - y|$ for all $x, y \in \mathbb{R}_+$, the previous inequality implies that for all $\delta \geq \frac{2}{3 - V_0}$,

$$\left|\frac{V_0 - \mu_0}{1 - \mu_0} - F_t\right| \leq \frac{2}{1 - V_0}\frac{1-\delta}{\delta} \quad \text{and} \quad |\mu_0 + (1 - \mu_0)F_t - V_0| \leq \frac{2}{1 - V_0}\frac{1-\delta}{\delta}. \qquad \square$$

The previous lemma does not establish that $\{F_t\}$ converges. Even if it did, its limit is not necessarily $\varphi(V_0)$. The value $\varphi(V_0)$ is itself the limit frequency of lying that attains as $\delta \to 1$, as can be deduced by the fact that $F_t \in [F - \epsilon(\delta), F + \epsilon(\delta)]$, where $\epsilon(\delta) = 2(1 - \delta)/[\delta(1 - V_0)] \to 0$ as $\delta \to 1$.

Definition 9. The *reputation maintenance value function* is defined by

$$V^M(\pi_0) = \mu_0 + (1 - \mu_0)f(\pi_0) \qquad \pi_0 \in [\pi^*, 1).$$

Sometimes we may want to make explicit that this function also depends on π_B, and write $f(\pi_0, \pi_B)$ and $V^M(\pi_0, \pi_B)$ instead of $f(\pi_0)$ and $V^M(\pi_0)$.

If $\pi(\beta) = \pi(\beta_0)$ for all $\beta \in [\beta_H, \beta_L]$, $\overline{V}(\beta_0) = V^\pi(\beta_0)$ and by Lemmas 20 and 21,

$$|\overline{V}(\beta_0) - V^M(\pi_0)| \leq \frac{2}{1 - V_0}\frac{1-\delta}{\delta}.$$

Definition 10. Let $F^* = f(\pi^*)$. For any $F \in [F^*, 1)$ let $p_F : [0,1] \to [0,1]$ be the function $p_F(x) = x^F(1-x)^{(1-F)}$.

For any $F \in [F^*, 1)$, the function p_F is quasiconcave and $p_F(0) = p_F(1) = 0$. Moreover, p_F attains its maximum at $x = F$. So $p_F(x)$ is strictly increasing for $x \in [0, F]$ and strictly deceasing for $x \in [F, 1]$.

Define f's inverse function $g : [F^*, 1) \to [\pi^*, 1)$ as follows: $x = g(F)$ is the unique solution $x \in [F, 1)$ of the equation $p_F(x) = p_F(\pi_B)$. That is

$$\left[\frac{x}{\pi_B}\right]^F \left[\frac{1-x}{1-\pi_B}\right]^{(1-F)} = 1.$$

By definition, $\pi^* = g(F^*)$. One can easily check that the function g is strictly increasing and continuous, and by definition, $g(F) > F$ for all $F \in (\pi_B, 1)$. In particular, $\pi^* > F^*$.

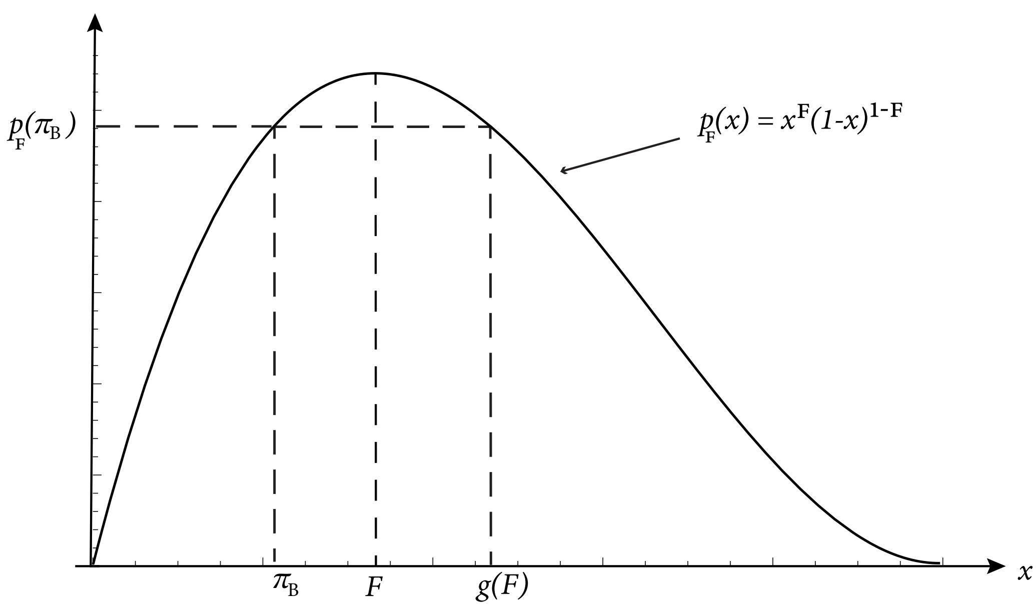

Typically, for any β_0, the equilibrium strategy $\pi(\beta)$ in $[\beta_H, \beta_L]$ is not constant. The following lemma investigates the relationship between equilibrium value and equilibrium strategy when the equilibrium strategy is increasing in the interval $[\beta_H, \beta_L]$.

Lemma 22. Assume $\pi(\beta, \delta)$ is weakly increasing in β. For $\beta_0 \in (0,1)$ let β_H and β_L be as defined above, and assume that $(\beta_0, \beta_L]$ does not intersect Region 1 and that $\beta_0 < \beta^{23}$. Let $V_0 = \overline{V}(\beta_0)$ and $F = \varphi(V_0)$ be the lying frequency induced by the value maintenance strategy at value V_0. Assume that $V_0 > V^M(\pi^*)$. Define $\pi_F = g(F)$ and

$$\gamma = \frac{\pi_F}{1-\pi_F}\frac{1-F}{F}.$$

Note that $\pi_F > F$ implies that $\gamma > 1$. If $\pi(\beta_H) > \pi_F$ then $\pi(\beta_L) \geq \pi_F + \gamma(\pi(\beta_H) - \pi_F)$.

Proof: Let $\{\beta_t\}$ be the random path of posteriors generated when the sender follows the value maintenance at V_0 starting from β_0. By Lemma 21,

$$\mu_0 + (1-\mu_0)F + \frac{2}{1-V_0}\frac{1-\delta}{\delta} \geq V_0 > V^M(\pi^*) = \mu_0 + (1-\mu_0)F^*.$$

Hence, $F > F^*$ for all $\delta > 2/[2 + (1-V_0)(V_0 - V^M(\pi^*))]$.

Let

$$J_{t,H} = \{j < t \mid \theta_j = \ell \text{ and } V_j \geq V_0\}, \quad J_{t,L} = \{j < t \mid \theta_j = \ell \text{ and } V_j < V_0\} \quad \text{and}$$
$$p_t = \left[\prod_{j\in J_{t,H}} \pi(\beta_j)\right] \times \left[\prod_{j\in J_{t,L}} (1-\pi(\beta_j))\right].$$

The sender lies in $n = |J_{t,H}|$ periods and tells the truth in $m = |J_{t,L}|$ periods. By Bayes' rule,

$$\beta_t = \frac{\beta_0}{\beta_0 + (1-\beta_0)L_t} \quad \text{where} \quad L_t = \frac{p_t}{\pi_B^n(1-\pi_B)^m}.$$

By Lemma 20, value maintenance and reputation maintenance coincide and therefore $\beta_t \in [\beta_H, \beta_L]$ for all t. Moreover, since $\beta_0 < \beta^{23}$, $\beta_L < 1$. (Lemma 20 allows the possibility that $\beta_0 = \beta^{23}$. Here we make the stronger assumption that $\beta_0 < \beta^{23}$ precisely to ensure that $\beta_L < 1$.)

Because $\pi(\beta, \delta)$ is weakly increasing in β,

$$p_t \leq \left[\prod_{j\in J_{t,H}} \pi(\beta_L)\right] \times \left[\prod_{j\in J_{t,L}} (1-\pi(\beta_H))\right] = [\pi(\beta_L)]^n[1-\pi(\beta_H)]^m$$
$$L_t \leq \frac{[\pi(\beta_L)]^n[1-\pi(\beta_H)]^m}{\pi_B^n[1-\pi_B]^m}.$$

Let $\bar{p}_t = [\pi(\beta_L)]^F[1-\pi(\beta_H)]^{1-F}$. When t is large (so that n and m are large), $n \approx (n+m)F$, $m \approx (n+m)F$, and $L_t \leq [\bar{p}_t/p_F(\pi_B)]^{n+m}$. Assume $\bar{p}_t < p_F(\pi_B)$. Then $L_t \to 0$ and $\beta_t \to 1$, which is a contradiction. Therefore, it must be that $\bar{p}_t \geq p_F(\pi_B)$.

For any $x \in [\pi_F, 1)$, consider the equation in y

$$y^F(1-x)^{(1-F)} = p_F(\pi_B).$$

This is equation is solved by

$$y = h_F(x) := \left[\frac{p_F(\pi_B)}{(1-x)^{(1-F)}}\right]^{\frac{1}{F}}.$$

Since

$$h'_F(x) = p_F(\pi_B)^{\frac{1}{F}} \left[\frac{1-F}{F}\right](1-x)^{\frac{-1}{F}} > 0 \quad \text{and}$$
$$h''_F(x) = p_F(\pi_B)^{\frac{1}{F}} \left[\frac{1-F}{F^2}\right](1-x)^{\frac{-(1+F)}{F}} > 0,$$

the function h_F is increasing and convex. Moreover, $h_F(\pi_F) = \pi_F$ and

$$h'_F(\pi_F) = \frac{\pi_F}{1-\pi_F}\frac{1-F}{F} := \gamma > 1.$$

Therefore, $\bar{p}_t \geq p_F(\pi_B)$ only if

$$\pi(\beta_L) \geq h_F(\pi(\beta_H)) \geq h_F(\pi_F) + h'_F(\pi_F)(\pi(\beta_H) - \pi_F) = \pi_F + \gamma(\pi(\beta_H) - \pi_F). \qquad \square$$

Define

$$\kappa = \frac{1-\mu_0\delta}{(1-\mu_0)\delta}, \quad \rho = \frac{\log(\kappa)}{\log(\lambda_1)}, \quad \ell_i = \log(\lambda_i) \quad i = 0, 1,$$
$$V^{12}(\pi_B, \delta) = \mu_0 + \frac{1-\delta}{\delta[\kappa - \kappa^{\ell_0/\ell_1}]}.$$

Below, we will usually omit the variables in V^{12}; similarly we have omitted the variables in the definitions of κ and ρ. As stated in Lemma 24 below, $V^{12}(\pi_B, \delta)$ is approximately the value of $\overline{V}(\beta)$ when the posterior β is at the boundary between Region 1 and Region 2.

Lemma 23. There exists a continuous function $a : [0, \lambda_1\beta^{12}] \to \mathbb{R}$ with the cyclical property that $a(\lambda_1\beta) = a(\beta)$ such that

$$\overline{V}(\beta) = \mu_0 + a(\beta)\beta^\rho \quad \text{for all } \beta < \lambda_1\beta^{12}.$$

Proof: Assume $\beta_k := \lambda_1^k\beta_0$ is in Region 1 for $k = 1, \dots, K$. Denote $v_k = \overline{V}(\beta_k)$. The right promising keeping constraint then becomes

$$v_k = \mu_0[(1-\delta) + \delta v_k] + (1-\mu_0)\delta v_{k+1} \quad \text{or} \quad v_{k+1} = \kappa v_k - \frac{\mu_0(1-\delta)}{(1-\mu_0)\delta}.$$

The solution to this linear difference equation is

$$v_k = \mu_0 + \hat{a}\kappa^k$$

for some constant $\hat{a} > 0$ ($\hat{a} = v_0 - \mu_0$). Note that $k = \log[\beta_k/\beta_0]/\ell_1$. Therefore

$$\overline{V}(\beta_k) = \mu_0 + a\beta_k^\rho,$$

where $a > 0$ is constant (determined by the initial condition $\overline{V}(\beta_0)$).

If starting at a different posterior $\tilde{\beta}_0$ all the points $\tilde{\beta}_k = \lambda_1^k \tilde{\beta}_0$ for $k = 1, \dots, \tilde{K}$ are in Region 1, then there exists another constant $\tilde{a}$ such that $\overline{V}(\tilde{\beta}_k) = \mu_0 + \tilde{a}\tilde{\beta}_k^\rho$. Since $\overline{V}(\beta)$ is a continuous and increasing function of β,

$$\overline{V}(\beta) = \mu_0 + a(\beta)\beta^\rho \quad \text{for all } \beta \in [0, \lambda_1 \beta^{12}],$$

where $a(\beta)$ is a continuous function of β such that $a(\lambda_1\beta) = a(\beta)$ for all β. □

Lemma 24. Let $\hat{\beta}$ be such that $\overline{V}(\hat{\beta}) = V^{12}(\pi_B, \delta)$. Then

1. there is $\beta^{12} \in [\hat{\beta}/\lambda_1, \lambda_1\hat{\beta}]$ such that $\overline{V}(\lambda_1\beta^{12}) - \overline{V}(\lambda_0\beta^{12}) = (1-\delta)/\delta$;
2. the interval $[0, \beta^{12})$ is contained in Region 1 and β^{12} is the first reputation in Region 2: $\overline{V}(\lambda_1\beta) - \overline{V}(\lambda_0\beta) < (1-\delta)/\delta$ for all $\beta < \beta^{12}$;
3. for any $\delta \geq 1/(1+\mu_0)$, $|\overline{V}(\beta^{12}) - V^{12}| \leq 3(1-\delta)$.

Proof: For any fixed $\hat{a} > 0$, let $\hat{V}(\beta) = \mu_0 + \hat{a}\beta^\rho$ and $\Delta(\beta) = \hat{V}(\lambda_1\beta) - \hat{V}(\lambda_0\beta)$. Then, (1) $\Delta(0) = 0$; (2) $\Delta(\beta)$ is increasing in β; and (3) $\Delta(\beta) = (1-\delta)/\delta$ if and only if $\hat{V}(\beta) = V^{12}(\pi_B, \delta)$.

Let $\hat{\beta}$ be such that $\overline{V}(\hat{\beta}) = V^{12}$, and define $\hat{a} = a(\hat{\beta})$ (defined in Lemma 23). If $a(\lambda_0\hat{\beta}) = \hat{a}$, then $\overline{V}(\lambda_0\hat{\beta}) = \hat{V}(\lambda_0\hat{\beta})$ and $\overline{V}(\lambda_1\hat{\beta}) - \overline{V}(\lambda_0\hat{\beta}) = \hat{V}(\lambda_1\hat{\beta}) - \hat{V}(\lambda_0\hat{\beta}) = (1-\delta)/\delta$, since $a(\lambda_1\hat{\beta}) = \hat{a}$. But typically $a(\lambda_0\hat{\beta}) \neq \hat{a}$. Assume that $\overline{V}(\lambda_0\hat{\beta}) < \hat{V}(\lambda_0\hat{\beta})$. Then $\overline{V}(\lambda_1\hat{\beta}) - \overline{V}(\lambda_0\hat{\beta}) > \hat{V}(\lambda_1\hat{\beta}) - \hat{V}(\lambda_0\hat{\beta}) = (1-\delta)/\delta$ and $\hat{\beta}$ is already in the interior of Region 2. Since $a(\beta)$ is continuous and makes a full "cycle" in the interval $[\hat{\beta}/\lambda_1, \hat{\beta}]$, there exists $\tilde{\beta} \in (\hat{\beta}/\lambda_1, \hat{\beta})$ such that $\overline{V}(\lambda_0\tilde{\beta}) = \tilde{V}(\lambda_0\tilde{\beta})$. Since $\tilde{\beta} < \hat{\beta}$, $\overline{V}(\tilde{\beta}) < V^*$ and $\overline{V}(\lambda_1\tilde{\beta}) - \overline{V}(\lambda_0\tilde{\beta}) = \tilde{V}(\lambda_1\tilde{\beta}) - \tilde{V}(\lambda_0\tilde{\beta}) < (1-\delta)/\delta$, and $\tilde{\beta}$ is in Region 1. That is, the transition between Region 1 and Region 2 (when $\overline{V}(\lambda_1\beta) - \overline{V}(\lambda_0\beta) = (1-\delta)/\delta$) must occur at some $\beta \in [\hat{\beta}/\lambda_1, \hat{\beta}]$. Similarly, when $\overline{V}(\lambda_0\hat{\beta}) > \hat{V}(\lambda_0\hat{\beta})$, the transition must occur at some $\beta \in [\hat{\beta}, \lambda_1\hat{\beta}]$.

In summary, there exists $\beta^{12} \in [\hat{\beta}/\lambda_1, \lambda_1\hat{\beta}]$, where the first transition between Region 1 and Region 2 occurs: $\overline{V}(\lambda_1\beta^{12}) - \overline{V}(\lambda_0\beta^{12}) = (1-\delta)/\delta$.

Let $\hat{a} = a(\beta^{12})$. Since $\beta^{12} \geq \hat{\beta}/\lambda_1$, $\overline{V}(\beta^{12}) \geq \overline{V}(\hat{\beta}/\lambda_1)$. Since $\lambda_1^\rho = \kappa = [1 - \mu_0\delta]/[(1-\mu_0)\delta]$, we have that

$$\begin{aligned}
V^{12} = \overline{V}(\hat{\beta}) &= \mu_0 + \hat{a}\hat{\beta}^\rho \quad \text{and} \\
\overline{V}(\hat{\beta}/\lambda_1) = \mu_0 + \hat{a}(\hat{\beta}/\lambda_1)^\rho &= V^{12} - \left[V^{12} - \mu_0\right]\left[1 - \frac{1}{\lambda_1^\rho}\right] \\
&= V^{12} - \left[V^{12} - \mu_0\right]\frac{1-\delta}{1-\mu_0\delta},
\end{aligned}$$

which establishes the lower bound on $\overline{V}(\beta^{12})$; the upper bound is proved similarly.
□

Let

$$V^{12}(\pi_B, 1) = \lim_{\delta \uparrow 1} V^{12}(\pi_B, \delta).$$

We now show that

$$V^{12}(\pi_B, 1) = V^M(\pi^*, \pi_B) = \mu_0 + (1 - \mu_0) f(\pi^*, \pi_B),$$

which we will return to interpret at the end of the Appendix. Let $D(\delta) = \delta[\kappa(\delta) - \kappa(\delta)^{\ell_0/\ell_1}]$. By L'Hopital

$$\lim_{\delta \uparrow 1} [V^{12} - \mu_0] = \frac{-1}{\lim\ D'(\delta)}.$$

Since

$$\kappa(\delta) \to 1 \quad \text{and} \quad \kappa'(\delta) = -\frac{1}{(1-\mu_0)\delta^2} \to -\frac{1}{1-\mu_0} \quad \text{we have}$$

$$D'(\delta) = [\kappa - \kappa^{\ell_0/\ell_1}] + \delta\kappa'[1 - \kappa^{\ell_0/\ell_1 - 1}\ell_0/\ell_1] \to -\frac{1}{1-\mu_0}[1 - \ell_0/\ell_1] \quad \text{and}$$

$$\lim_{\delta \uparrow 1} V^{12} = \mu_0 + (1-\mu_0)\left[\frac{\ell_1}{\ell_1 - \ell_0}\right] = \mu_0 + (1-\mu_0)\frac{\log\left[\frac{(1-\pi_B)(1-\mu_0)}{1-2\mu_0}\right]}{\log\left[\frac{1-\pi_B}{\pi_B}\frac{\mu_0}{1-2\mu_0}\right]}.$$

Finally recall that $\pi^* = \mu_0/(1-\mu_0)$ and $1 - \pi^* = (1-2\mu_0)/(1-\mu_0)$.

Recall that Region 3 is the interval $[\beta_{2,3}, 1]$ and that $\beta \in [\beta^{23}, 1]$ if and only if $\overline{V}(\beta, \delta) \geq V^{23}(\delta)$ Fix $\epsilon > 0$ (small). Since $V^{23}(\delta) \to 1$ as $\delta \to 1$, there exists $\underline{\delta} \in (0, 1)$ such that $V^{23}(\delta) > V^M(\pi^*, \pi_N) + 2\epsilon$ for all $\delta \geq \underline{\delta}$.

Proof of Lemma 1: Fix $\epsilon \in (0, 1/2)$ and pick any $\delta > \delta_1 := [2(1-\mu_0) - \epsilon]/[2(1-\mu_0) - \epsilon\mu_0]$. Then

$$\overline{V}(\beta^{23}, \delta) - (1-\epsilon) = \frac{\delta - \mu_0(2\delta - 1)}{1 - \delta\mu_0} - (1-\epsilon) > \epsilon/2.$$

If $\overline{V}(\epsilon, \delta) > 1 - \epsilon$, then $\overline{V}(1-\epsilon, \delta) - \overline{V}(\epsilon, \delta) < 1 - (1-\epsilon) = \epsilon$ and we are done. Hereafter, assume that $\overline{V}(\epsilon, \delta) \leq 1 - \epsilon$.

Set $\beta_0 = \epsilon$ and inductively define $v_k = \overline{V}(\beta_k, \delta)$ and $\beta_{k+1} = \beta_{L,\ell}(\beta_k, \pi(\beta_k))$, $k \geq 0$. Then

$$\beta_k = \frac{\beta_0}{\beta_0 + (1-\beta_0)L_k} \quad \text{where} \quad L_k = \frac{1}{(1-\pi_B)^k}[(1-\pi(\beta_0)) \cdots (1-\pi(\beta_{k-1}))].$$

Since $\pi(\beta) \geq \pi^* > \pi_B$ for all $\beta \in (0, 1]$, $(1-\pi(\beta))/(1-\pi_B) < (1-\pi^*)/(1-\pi_B) < 1$ and

$$L_k < \left[\frac{1-\pi^*}{1-\pi_B}\right]^k.$$

Let K be the smallest integer such that $L_K \leq [\epsilon/(1-\epsilon)]^2$. Then $\beta_K > 1-\epsilon$. Since by assumption $v_0 = \overline{V}(\epsilon,\delta) \leq 1-\epsilon < \frac{\delta-\mu_0(2\delta-1)}{1-\delta\mu_0} = \overline{V}(\beta^{23},\delta)$, we have that $\beta_0 < \beta^{23}$. For as long as $\beta_{k-1} < \beta^{23}$, by right promise keeping, we have that

$$v_k < v_{k-1} + \frac{1-\delta}{\delta} \quad \text{so} \quad v_k < v_0 + k\frac{1-\delta}{\delta}.$$

Let δ_2 be such that $K(1-\delta_2)/\delta_2 = \epsilon/2$. Then, for any $\delta \geq \hat{\delta} = \max\{\delta', \delta_1, \delta_2\}$, $v_K - v_0 \leq \epsilon/2$. This implies that $v_k < \overline{V}(\beta^{23},\delta)$, and hence $\beta_k < \beta^{23}$, for all $k \leq K$. Therefore $\overline{V}(1-\epsilon,\delta) - \overline{V}(\epsilon,\delta) < v_K - v_0 \leq \epsilon/2$. □

Lemma 25. For any $\epsilon \in (0,1)$ there exists $\underline{\delta} \in (0,1)$ such that

$$\overline{V}(\beta,\delta) \leq V^M(\pi^*) + \epsilon \quad \text{for all} \quad \beta \in [0, 1-\epsilon] \quad \text{and} \quad \delta \geq \underline{\delta}.$$

Proof: Since $\bar{\pi}(\beta)$ is convex in $\beta \in [0,1)$, and $\bar{\pi}(0) = \pi^*$ and $\bar{\pi}(1/2) < 2\pi^*$, $\bar{\pi}(\beta) \leq (1+2\beta)\pi^*$ for all $\beta \in [0, \frac{1}{2}]$. Therefore, for any $\hat{\epsilon} \in (0, 1/2]$,

$$\pi(\beta) \leq \bar{\pi}(\beta) \leq (1+2\hat{\epsilon})\pi^* \quad \text{for all} \quad \beta \in [0,\hat{\epsilon}].$$

For any $\beta_0 \in [0,\hat{\epsilon}]$, let $\pi_0(\beta) = \pi(\beta_0)$ for all $\beta \in [\beta_H, \beta_L]$, where as usual $\beta_H = \beta_{H,\ell}(\beta_0, \pi(\beta_0))$ and $\beta_L = \beta_{L,\ell}(\beta_0,\pi(\beta_0))$. By Lemma 18, $\overline{V}(\beta_0,\delta) \leq V^{\pi}(\beta_0,\delta) \leq V^{\pi_0}(\beta_0,\delta)$, and by Lemmas 20 and 21, $V^{\pi_0}(\beta_0,\delta) = V^M(\pi(\beta_0)) + O((1-\delta))$. Therefore, since V^M is an increasing function,

$$\overline{V}(\beta_0,\delta) \leq V^M(\pi^* + 2\pi^*\hat{\epsilon}) + O((1-\delta)) \qquad \text{for all} \quad \beta_0 \in [0,\hat{\epsilon}].$$

By Lemma 1, for any $\hat{\epsilon} \in (0,\epsilon/2]$ there exists $\delta_1(\hat{\epsilon})$ such that for any $\delta \geq \delta_1(\hat{\epsilon})$, $\overline{V}(1-\hat{\epsilon},\delta) - \overline{V}(\hat{\epsilon},\delta) < \hat{\epsilon}$. Thus

$$\begin{aligned}\overline{V}(1-\hat{\epsilon},\delta) - V^M(\pi^*) &= [\overline{V}(1-\hat{\epsilon},\delta) - \overline{V}(\hat{\epsilon},\delta)] + [\overline{V}(\hat{\epsilon},\delta) - V^M(\pi^*)] \\ &< \hat{\epsilon} + V^M(\pi^* + 2\pi^*\hat{\epsilon}, \pi_B) - V^M(\pi^*,\pi_B) + O((1-\delta)).\end{aligned}$$

Since V^M is continuous, we can choose $\hat{\epsilon} > 0$ small enough such that $V^M(\pi^* + 2\pi^*\hat{\epsilon}, \pi_B) - V^M(\pi^*,\pi_B) < \epsilon/4$ and then choose $\underline{\delta} > \delta_1(\hat{\epsilon})$ so that $O((1-\delta)) < \epsilon/4$ for all $\delta > \underline{\delta}$. Then, for all $\delta > \underline{\delta}$,

$$\overline{V}(1-\epsilon,\delta) - V^M(\pi^*) \leq \overline{V}(1-\hat{\epsilon},\delta) - V^M(\pi^*) \leq \hat{\epsilon} + \frac{\epsilon}{4} + \frac{\epsilon}{4} \leq \epsilon.$$ □

Recall that $\beta_d = 1/\lambda_1$, $\bar{\pi}(\beta_d) = 1$ and $\bar{\pi}(\beta) \geq 1$ for all $\beta \in [\beta_d, 1]$. Define $\beta_d^+ \in (0,1)$ to be the solution of $\beta_{H,\ell}(\beta_d^+, 1) = \beta_d$. Clearly $\beta_d^+ > \beta_d$. Moreover, for any $\beta \in [\beta_d^+, 1]$, $\beta_{H,\ell}(\beta, \pi(\beta)) \geq \beta_d$. That is, if the sender's current reputation β is above β_d^+, his reputation when he lies remains above β_d even when in equilibrium he is expected to lie with probability 1.

Lemma 26. For all $\epsilon > 0$, there exists $\underline{\delta} \in (0,1)$ such that $V^M(\pi_B,\delta) - 3\epsilon \leq \overline{V}(\beta,\delta)$ for all $\beta \in [\epsilon, 1]$ and all $\delta \geq \underline{\delta}$.

Proof: Pick any $\epsilon \in (0, 1-\beta_d)$. Since $\lim_{\delta\uparrow 1} V^{12}(\pi_B, \delta) = V^M(\pi^*, \pi_B)$, there exists $\delta_M \in (0,1)$ such that $V^{12}(\pi_B, \delta) \geq V^M(\pi^*, \pi_B) - \epsilon$ for all $\delta \geq \delta_M$.

By Lemma 1, there exists $\delta_1 \in (1-\epsilon/3, 1)$ such that $\overline{V}(1-\epsilon, \delta) - \overline{V}(\epsilon, \delta) \leq \epsilon$ for all $\delta \geq \delta_1$. By Lemma 24, for any $\delta \geq \delta_1$, $|\overline{V}(\underline{\beta}, \delta) - V^{12}(\pi_B, \delta)| \leq 3(1-\delta) \leq \epsilon$, where $\underline{\beta}$ is the first reputation in Region 2. By definition $1 \geq \pi(\beta) = \bar{\pi}(\beta)$ for all $\beta < \underline{\beta}$. This implies that $\underline{\beta} \leq \beta_d < 1-\epsilon$. Let $\underline{\delta} = \min\{\delta_1, \delta_M\}$. If $\underline{\beta} < \epsilon$, $V^{12}(\pi_B, \delta) - \epsilon \leq \overline{V}(\underline{\beta}, \delta) \leq \overline{V}(\beta, \delta)$ for all $\beta \geq \underline{\beta}$. If $\underline{\beta} \geq \epsilon$, then $\underline{\beta} \in [\epsilon, 1-\epsilon]$ and $\overline{V}(\underline{\beta}, \delta) - \overline{V}(\epsilon, \delta) \leq \epsilon$. Therefore, $\overline{V}(\beta, \delta) \geq \overline{V}(\epsilon, \delta) \geq \overline{V}(\underline{\beta}, \delta) - \epsilon \geq V^{12}(\pi_B, \delta) - 2\epsilon$ for all $\beta \in [\epsilon, 1]$. Either way, we have that

$$\overline{V}(\beta, \delta) \geq V^{12}(\pi_B, \delta) - 2\epsilon \geq V^M(\pi^*, \pi_B) - 3\epsilon$$

for all $\beta \in [\epsilon, 1]$ and $\delta \geq \underline{\delta}$. □

Proof of Proposition 4: Pick any $\eta > 0$. For any $\epsilon > 0$ let $F = \varphi(V^M(\pi^*) + \epsilon) = F^* + \epsilon/(1-\mu_0)$. Since g is continuous, we can choose $0 < \epsilon < \min\{\eta, 1-\beta_d^+\}$ such that $g(F) - \pi^* = g(F) - g(F^*) < \eta$. Let Q be the smallest integer greater or equal to $-\log(\eta)/\log(\gamma)$ (that is, $\gamma^Q \eta \geq 1$).

Since $V^{23}(\delta) \to 1$ as $\delta \to 1$ and by Lemma 25, we can pick $\underline{\delta}_1 \geq \delta_I$ such that for all $\delta \geq \underline{\delta}_1$, $V^{23}(\delta) \geq V^M(\pi^*) + \epsilon$ and $\overline{V}(\beta, \delta) \leq V^M(\pi^*) + \epsilon/3$ for all $\beta \in [0, 1-\epsilon/3]$. Define the integer M by

$$M = \left\lceil \frac{1 - V^M(\pi^*)}{V^M(\pi^*) - \mu_0} \right\rceil + 1.$$

Let $\underline{\delta}_2$ be such that $(M+1)Q(1-\underline{\delta}_2)/\underline{\delta}_2 \leq \epsilon/3$. Define $\underline{\delta} = \max\{\underline{\delta}_1, \underline{\delta}_2\}$.

Fix $\delta \geq \underline{\delta}$. Assume that $\pi(\beta_0) > \pi^* + 2\eta$ for some $\beta_0 \in (0, 1-\epsilon]$ and let $V_0 = \overline{V}(\beta_0)$. Note that $\bar{\pi}(\beta)$ is continuous and $\lim_{\beta\to 0} \bar{\pi}(\beta) = \pi^*$. In equilibrium $\pi(\beta) \leq \bar{\pi}(\beta)$ for all $\beta \in [0,1]$, so $\pi(\beta_0) > \pi^* + 2\eta$ implies that $\beta_0 > 0$. Also, $\beta_0 < \beta^{23}$ for otherwise $\overline{V}(\beta_0) > V^{23}(\delta)$, which contradicts the assumption that $\overline{V}(\beta) \leq V^M(\pi^*) + \epsilon/3 < V^{23}(\delta)$ for all $\beta \in [0, 1-\epsilon/3]$. If $\beta_0 \geq \beta_d^+$ skip Step 1 below.

Step 1: If $\beta_0 < \beta_d^+$, we now find another posterior β_K such that $\pi(\beta_K) > \pi^* + 2\gamma$ and $\beta_d^+ \leq \beta_K < \beta^{23}$. For $k \geq 0$, inductively define V_{k+1} by RPK and $\beta_{k+1} = \beta_{L,\ell}(\beta_k, \pi(\beta_k))$. Then (see Lemma 23)

$$V_{k+1} = \frac{1-\delta\mu_0}{(1-\mu_0)\delta} V_k - \frac{\mu_0}{1-\mu_0}\frac{1-\delta}{\delta} \qquad \text{and}$$
$$V_k = \mu_0 + (V_0 - \mu_0)\kappa^k \quad \text{where} \quad \kappa = \frac{1-\delta\mu_0}{(1-\mu_0)\delta}.$$

Therefore,

$$\frac{1-\delta}{\delta} > V_{k+1} - V_k = \frac{(V_k - \mu_0)}{(1-\mu_0)}\frac{(1-\delta)}{\delta} \geq \frac{(V_0 - \mu_0)}{(1-\mu_0)}\frac{(1-\delta)}{\delta}.$$

Since $\beta_d^+ < 1-\epsilon$, $\overline{V}(\beta_d^+) \leq V^M(\pi^*) + \epsilon/3$. Let K be the smallest integer such that $V_K \geq \overline{V}(\beta_d^+)$. Then $V_K < \overline{V}(\beta_d^+) + (1-\delta)/\delta < V^M(\pi^*) + 2\epsilon/3 < V^{23}(\delta)$. This implies that $\beta_k < \beta^{23}$ for all $k = 0, \dots, K-1$. Therefore, by RPK, $V_k = \overline{V}(\beta_k)$ for all $k = 0, \dots, K$. Since π is increasing, $\pi(\beta_K) \geq \pi(\beta_0) > \pi^* + 2\eta$. With a slight abuse of notation, rename V_K and β_K to be V_0 and β_0 respectively.

After the completion of Step 1 (if necessary), we have that $\beta_d^+ \leq \beta_0$ and $\overline{V}(\beta_0) \leq V^M(\pi^*) + 2\epsilon/3 < V^M(\pi^*) + \epsilon \leq \overline{V}(\beta^{23})$.

Step 2: Starting with β_0, we now construct an increasing sequence $\{\beta_k\}$ such that (1) $\pi(\beta_k) > g(F) + \gamma^k\eta$ and (2) $\overline{V}(\beta_d^+) \leq \overline{V}(\beta_k) < V^{23}(\delta)$ for all $k = 0, \dots, Q-1$.

For each k, starting from β_k, let $x_{k,m}$ be the resulting posterior after the sender tells the truth m times in a row, and $x_{k,m}^-$ be the posterior that attains from prior $x_{k,m}$ when the sender lies once. Denote $v_{k,m} = \overline{V}(x_{k,m})$ and $v_{k,m}^- = \overline{V}(x_{k,m}^-)$. By definition $v_{k,0} = \overline{V}(\beta_k)$, and $v_{k,m}$ is increasing in m. By RPK, from $x_{k,m}$ to $x_{k,m+1}$ the value increases

$$v_{k,m+1} - v_{k,m} = \frac{1}{1-\mu_0}\frac{1-\delta}{\delta}(v_{k,m} - \mu_0) \geq \frac{1}{1-\mu_0}\frac{1-\delta}{\delta}(v_{k,0} - \mu_0)$$

And by LPK, from $x_{k,m}$ to $x_{k,m}^-$ the value decreases

$$v_{k,m} - v_{k,m}^- = \frac{1}{1-\mu_0}\frac{1-\delta}{\delta}(1 - v_{k,m}) \leq \frac{1}{1-\mu_0}\frac{1-\delta}{\delta}(1 - v_{k,0}).$$

Let M_k be the smallest integer such that $v_{k,M_k}^- \geq V(\beta_k)$. It is easy to see that M_k is no larger than $\lceil (1-v_{k,0})/(v_{K,0} - \mu_0) \rceil$ (where $\lceil z \rceil$ denotes the smallest integer greater or equal to z). Since $v_{k,0} = \overline{V}(\beta_k)$ and, as we will see, $\overline{V}(\beta_k) - \overline{V}(\beta_0)$ is very small for each k, we can bound

$$M_k \leq \left\lceil \frac{1 - V^M(\pi^*)}{V^M(\pi^*) - \mu_0} \right\rceil + 1 = M \qquad \text{for all } k \text{ (in the relevant range).}$$

Let $\beta_{k+1} = x_{k,M}^+$. Then, $\overline{V}(\beta_k) \leq \overline{V}(x_{k,M}^-) \leq \overline{V}(x_{k,M}^+) = \overline{V}(\beta_{k+1})$ and

$$\overline{V}(\beta_{k+1}) - \overline{V}(\beta_k) < (M+1)\left[\frac{1-\delta}{\delta}\right] < (M+1)Q\left[\frac{1-\delta}{\delta}\right] \leq \epsilon/3,$$

and $\overline{V}(\beta_k) - \overline{V}(\beta_0) < \epsilon/3$ for all k, as claimed above. Since $\overline{V}(\beta_0) < V^M(\pi^*) + 2\epsilon/3$, $\overline{V}(\beta_k) < V^M(\pi^*) + \epsilon < V^{23}(\delta)$ for all $k = 0, \dots, Q$. Also, $\pi(x_{k,M}^-) \geq \pi(\beta_k) \geq g(F) + \gamma^k\eta$. In the context of Lemma 22, $x_{k,M}$ is the center of an interval $[\beta_H, \beta_L]$, where $\beta_H = x_{k,M}^- \geq \beta_k$ and $\beta_L = x_{k,M}^+ = \beta_{k+1}$. By construction, $\pi(\beta_H) \geq g(F) + \gamma^k\eta$. Starting at $x_{k,M}$, the value maintainance strategy at value $\overline{V}(x_{k,M})$ induces a lying frequency $F \leq \varphi(V^M(\pi^*) + \epsilon) = F^* + \epsilon/(1-\mu_0)$. Therefore, by Lemma 22, it must be that $\pi(\beta_{k+1}) = \pi(\beta_L) > g(F) + \gamma^{k+1}\eta$.

We have reached a contradiction. When $k = Q-1$, $g(F) + \gamma^{k+1}\eta \geq 1$. That is, $\pi(\beta_L) = \pi(x^+_{k,M})$ would have to be increased above 1 to maintain value $\overline{V}(x_{k,M}) \leq V^M(\pi^*) + \epsilon$, which is not possible. Therefore, $\pi(\beta_0) \leq \pi^* + 2\eta$ for all $\beta \in [0, 1-\epsilon] \supset [0, 1-\eta]$

We finally show that $\alpha(\beta) = 1$ for all $\beta \in [2\eta/(\pi^* - \pi_b), 1-\eta]$. Since $\bar{\pi}$ is convex, $\bar{\pi}(\beta) \geq \bar{\pi}(0) + \bar{\pi}'(0)\beta = \pi^* + (\pi^* - \pi_B)\beta$. Also, $\pi^* + (\pi^* - \pi_B)\beta > \pi^* + 2\eta \geq \pi(\beta)$ for all $\beta \in [2\eta/(\pi^* - \pi_b), 1-\eta]$. When $\pi(\beta) < \bar{\pi}(\beta)$ the receiver strictly prefers to accept an H recommendation and thus $\alpha(\beta) = 1$ □

Proof of Corollary 1: Pick any $\epsilon > 0$ and integer number T. For any given $\delta \in (0,1)$, let β^H_t be the reputation value obtained from $\beta^H_0 = \epsilon$ after t occurrences of $(\theta, m) = (\ell, H)$. Since $\pi(\beta) \leq \bar{\pi}(\beta)$ for all $\beta \in (0,1)$,

$$\underline{\beta}_\epsilon := \lambda_0^T \epsilon = \lambda_0^T \beta_0^H \leq \beta_t^H \quad \text{for all } t = 0, \dots, T.$$

Similarly, let β^L_t be the reputation value obtained from $\beta^L_0 = 1-\epsilon$ after t occurrences of $(\theta, m) = (\ell, L)$. Temporarily, assume that $\pi(\beta^L_t) \leq \pi^* + \epsilon$ for all $t = 0, \dots, T-1$. Let $\bar{\beta}_\epsilon$ be the solution of

$$\frac{1-\bar{\beta}_\epsilon}{\bar{\beta}_\epsilon} = \frac{\epsilon}{1-\epsilon}\left[\frac{1-\pi^*-\epsilon}{\pi_B}\right]^T.$$

Then

$$\frac{1-\beta^L_t}{\beta^L_t} \geq \frac{1-\beta^L_0}{\beta^L_0}\left[\frac{1-\pi^*-\epsilon}{\pi_B}\right]^t \geq \frac{1-\bar{\beta}_\epsilon}{\bar{\beta}_\epsilon},$$

or equivalently, $\beta^L_t \leq \bar{\beta}_\epsilon$ for all $t = 0, \dots, T$. Since $\lambda_0 < 1$ and $\pi_B < \pi^* < \pi^* + \epsilon$, $\underline{\beta}_\epsilon < \epsilon$ and $\bar{\beta}_\epsilon > 1-\epsilon$.

We now show that $\pi(\beta^L_t) \leq \pi^* + \epsilon$ for all $t = 0, \dots, T-1$ is satisfied for δ sufficiently high. Let $\bar{\epsilon} = \bar{\pi}(\underline{\beta}_\epsilon) - \pi^*$. Since $\bar{\pi}(0) = \pi^*$ and $\bar{\pi}$ is a strictly increasing function, $\bar{\epsilon} > 0$. Pick any $0 < \eta < \frac{1}{2}\min\{\underline{\beta}_\epsilon, 1-\bar{\beta}_\epsilon, \bar{\epsilon}\}$. By Proposition 4, there exists $\underline{\delta} \in (0,1)$ such that $\pi^* \leq \pi(\beta) \leq \pi^* + \eta < \pi^* + \epsilon$ for all $\beta \in [0, 1-\eta/2]$ and $\delta \geq \underline{\delta}$. Fix $\delta \geq \underline{\delta}$. Then, for all $\beta \in [\underline{\beta}_\epsilon, 1]$, $\bar{\pi}(\beta) \geq \pi^* + \bar{\epsilon} > \pi^* + \eta \geq \pi(\beta)$. Since $1-\eta/2 > 1-\bar{\beta}_\epsilon$, the assumption made earlier that $\beta^L_t \leq \bar{\beta}_\epsilon$ for all $t = 0, \dots, T$ is satisfied.

Finally, since $\pi(\beta)$ is an increasing function in $[0, 1-\eta/2] \supset [\underline{\beta}_\epsilon, \bar{\beta}_\epsilon]$, for any $\underline{\beta}_\epsilon < \beta < \beta' < \bar{\beta}_\epsilon$ we have that $\beta_{H,\ell} < \beta'_{H,\ell}$ and $\beta_{L,\ell} < \beta'_{L,\ell}$. Therefore, for any $\beta_0 \in [\epsilon, 1-\epsilon]$, $\beta_t \in [\underline{\beta}_\epsilon, \bar{\beta}_\epsilon] \subset [\eta/2, 1-\eta/2]$ for all $t = 1, \dots, T$, and thus $\pi^* < \pi(\beta_t) < \pi^* + \epsilon$ for all $t = 0, \dots, T$. Also $\pi^* + \epsilon \leq \bar{\pi}(\beta_t)$, so $\alpha(\beta_t) = 1$ for all $t = 0, \dots, T$. □

B.2. δ-relevant Behavioral Convergence

Proof of Proposition 5: this result follows directly from the proof of Proposition 6, which applies to any finite number of behavioral types. □

C. Numerical Robustness

Proof of Lemma 2: Let

$$S(\beta_l, \beta_r) = \{(\beta, \overline{V}(\beta)) \mid \beta \in [\beta_l, \beta_r]\}.$$

We show that $S(\beta_l, \beta_r)$ is sufficient to reconstruct the whole upper boundary

$$S(0,1) = \{(\beta, \overline{V}(\beta)) \mid \beta \in [0,1]\}$$

together with $\{\pi(\beta) \mid \beta \in [0,1]\}$, and that $\overline{V}(\beta)$ is strictly increasing on $[0,1]$.

We will use the change of variables for reputation

$$z = \log\left(\frac{\beta}{1-\beta}\right)$$

and let

$$z_l = \log\left(\frac{\beta_l}{1-\beta_l}\right) \qquad \text{and} \qquad z_r = \log\left(\frac{\beta_r}{1-\beta_r}\right).$$

If, at reputation z, the sender lies with probability π, then his posterior beliefs become

$$z_+(z,\pi) = z + \log\left(\frac{1-\pi_B}{1-\pi}\right) \quad \text{or} \quad z_-(z,\pi) = z - \log\left(\frac{\pi}{\pi_B}\right),$$

after an honest report or a lie in state ℓ, respectively.

Let

$$A = \frac{1-\delta}{(1-\mu_0)\delta}.$$

By assumption, $\delta > 1/(1+\mu_0)$ and $\mu_0 < 1/2$, and therefore $A \in (0,1)$. By Proposition 2, $\overline{V}$ is continuous. Therefore, for any $z \in R_2$, by RPK and LPK, respectively, $\pi(z)$ is such that

$$\overline{V}(z_+(z,\pi(z))) = \overline{V}(z) + A(\overline{V}(z) - \mu_0) \qquad \text{and} \qquad \overline{V}(z_-(z,\pi(z))) = \overline{V}(z) - A(1-\overline{V}(z)).$$

Subtracting the two equalities gives, for every $z \in R_2$,

$$\overline{V}(z_+(z,\pi(z))) - \overline{V}(z_-(z,\pi(z))) = A(1-\mu_0) = \frac{1-\delta}{\delta}.$$

Let $z_1 = z_r$ and define

$$v_0 \;=\; \frac{\overline{V}(z_1) + A\mu_0}{1+A}.$$

Since (by assumption)

$$\overline{V}(z_r) - \overline{V}(z_l) \;\geq\; \frac{1-\delta}{\delta},$$

we have $\overline{V}(z_1) - \overline{V}(z_l) \geq A(1-\mu_0) \geq A(\overline{V}(z_l) - \mu_0)$, and thus $v_0 \in [\overline{V}(z_l), \overline{V}(z_1)]$. Because $\overline{V}$ is continuous and strictly increasing on $[z_l, z_1] \subset R_2$, there exists a unique $z_0 \in [z_l, z_1]$ such that $\overline{V}(z_0) = v_0$. Now set

$$v_{-1} \;=\; \overline{V}(z_0) - A(1 - \overline{V}(z_0)) \;=\; \overline{V}(z_1) - \frac{1-\delta}{\delta} \;\geq\; \overline{V}(z_l).$$

Again by strict monotonicity on $[z_l, z_1]$ there exists $z_{-1} \in [z_l, z_1]$ with $\overline{V}(z_{-1}) = v_{-1}$. Consequently $z_l \leq z_{-1} < z_0 < z_1 = z_r$ and

$$\overline{V}(z_1) = \overline{V}(z_0) + A(\overline{V}(z_0) - \mu_0), \qquad \overline{V}(z_{-1}) = \overline{V}(z_0) - A(1 - \overline{V}(z_0)).$$

Extending $S(\beta_l, \beta_r)$ to the right: By assumption and the construction above, there exist points z_0 and z_{-1} such that $z_l \leq z_{-1} < z_0 < z_1 = z_r$ and such that $z_l \leq z_{-1} < z_0 < z_1 = z_r$ and

$$\overline{V}(z_1) = \overline{V}(z_0) + A\big(\overline{V}(z_0) - \mu_0\big) \quad \text{and} \quad \overline{V}(z_{-1}) = \overline{V}(z_0) - A\big(1 - \overline{V}(z_0)\big).$$

We now successively extend the upper boundary from the interval $[z_{-1}, z_k]$ to $[z_{-1}, z_{k+1}]$ for $k = 1, 2, \ldots$, where $z_{k+1} > z_k$. In equilibrium, $S(z_{k-2}, z_k)$ uniquely determines $S(z_k, z_{k+1})$ through the following procedure: for any $z \in [z_{k-1}, z_k]$, find z_H such that

$$\overline{V}(z_H) = \overline{V}(z) - A(1 - \overline{V}(z)) \qquad \text{and define} \qquad \pi(z) = \pi_B e^{z - z_H}.$$

Then $z_-(z, \pi(z)) = z_H$. Since $\overline{V}(z_{k-1}) = \overline{V}(z_k) - A(1 - \overline{V}(z_k))$ and $z \geq z_k$ (so $\overline{V}(z) \geq \overline{V}(z_k)$), $z_H \in [z_{k-2}, z]$. Let

$$z_L = z_+(z, \pi(z)) \qquad \text{and} \qquad \overline{V}(z_L) = \overline{V}(z) + A(\overline{V}(z) - \mu_0).$$

This procedure maps $(z, \overline{V}(z))$ to $\big(z_+(z, \pi(z)), \overline{V}(z_+(z, \pi(z)))\big)$, and in particular, $(z_k, \overline{V}(z_k))$ to $(z_{k+1}, \overline{V}(z_{k+1}))$, where $z_{k+1} = z_+(z_k, \pi(z_k))$. It also determines the value of $\pi(z)$ for each $z \in [z_{k-1}, z_k]$.

One can verify inductively that if $\overline{V}(z)$ is strictly increasing on $[z_{-1}, z_k]$, then it is also strictly increasing on $[z_{-1}, z_{k+1}]$. Indeed, assume $\overline{V}(z)$ is strictly increasing on $[z_{-1}, z_k]$. By continuity of $\overline{V}$, it follows that $z_-(z, \pi(z))$, $z_+(z, \pi(z))$, and $\overline{V}\big(z_+(z, \pi(z))\big) \;=\; \overline{V}(z) + A\big(\overline{V}(z) - \mu_0\big)$ are strictly increasing in $z \in [z_{k-1}, z_k]$.

By construction, the sequence $\{w_k\}$ where $w_k = \overline{V}(z_k)$ satisfies the difference equation

$$w_{k+1} - (1 + A)w_k + A\mu_0 = 0.$$

The solution of this equation is $w_k = (V_0 - \mu_0)(1 + A)^k + \mu_0$, and thus, there is K such that

$$w_{K-1} \leq V^{23} < w_K.$$

The extension procedure ends in round K, as there exists $z^{23} \in [z_{K-1}, z_K]$ such that $\overline{V}(z^{23}) = V^{23}$. Since $V^{23} + A(V^{23} - \mu_0) = 1$, $\pi(z^{23}) = 1$ and $S(z_{K-1}, z^{23})$ maps into $S(z_K, \infty)$ [recall that when $\beta = 1$, $z = \infty$].

Extending $S(\beta_l, \beta_r)$ to the left: Set $z_{-1} = z_l$. By continuity of $\overline{V}$, there exist points z_0 and z_1 such that

$$\overline{V}(z_{-1}) = \overline{V}(z_0) - A\big(1 - \overline{V}(z_0)\big) \qquad \text{and} \qquad \overline{V}(z_1) = \overline{V}(z_0) + A\big(\overline{V}(z_0) - \mu_0\big).$$

We now extend the upper boundary from the interval $[z_{-1}, z_1]$ to $[0, z_1]$. We do this in two stages. In the first stage, while $[z_{k-1}, z_1] \subset R_2$, the upper boundary is successively extended from the interval $[z_k, z_1]$ to $[z_{k-1}, z_1]$ for $k = -1, -2, \dots$, where $z_{k-1} < z_k$, in a similar fashion as we extended above the upper boundary to the right.

Stage 1: For each $z \in [z_{k-1}, z_k]$ (from right to left), and by continuity of $\overline{V}$, find z_L such that

$$\overline{V}(z_L) = \overline{V}(z) + A(\overline{V}(z) - \mu_0).$$

Then $\pi(z) = 1 - e^{z - z_L}(1 - \pi_B)$. Set

$$z_H = z_-(z, \pi(z)) \qquad \text{and} \qquad \overline{V}(z_H) = \overline{V}(z) - A(1 - \overline{V}(z)).$$

This maps $(z, \overline{V}(z))$ to $(z_-(z, \pi(z)), \overline{V}(z_-(z, \pi(z))))$. Continue to decrease z from z_k until $z = z_{k-1}$ or $\pi(z) = \bar{\pi}(z)$. If $\pi(z) = \bar{\pi}(z)$, $z = z^{12}$ and we have reached the beginning of R_1. Since R_1 is an interval, $R_1 = [0, z^{12}]$. In this case we reset $z_{k-1} = z$ and we move to stage 2. Otherwise, set $z_{k-2} = z_-(z_{k-1}, \pi(z_{k-1}))$.

One can verify inductively that if $\overline{V}(z)$ is strictly increasing on $[z_{k-1}, z_1]$, then it is also strictly increasing on $[z_{k-2}, z_1]$. Indeed, assume $\overline{V}(z)$ is strictly increasing on $[z_{k-1}, z_1]$. By continuity of $\overline{V}$, it follows that $z_-(z, \pi(z))$, $z_+(z, \pi(z))$, and $\overline{V}\big(z_-(z, \pi(z))\big) = \overline{V}(z) + A\big(\overline{V}(z) - \mu_0\big)$ are strictly increasing in $z \in [z_{k-1}, z_k]$.

Stage 2: It is more convenient to carry the last step with variable β instead of z. Let

$$\beta^{12} = \frac{e^{z^{12}}}{1 + e^{z^{12}}} \qquad \text{and} \qquad \beta_0 = \frac{e^{z_0}}{1 + e^{z_0}}.$$

Since $\pi(\beta^{12}) = \bar{\pi}(\beta^{12})$, $\beta_-(z^{12}, \pi(\beta^{12})) = \lambda_0\beta^{12}$ and by stage 1, the upper boundary has been extended to the segment $S(\lambda_0\beta^{12}, \beta_0)$. The following step constructs the remaining segment $S(0, \lambda_0\beta^{12})$. For any $\pi_B < \pi^*$, $\lambda_0\lambda_1 < 1$. Hence $\lambda_0\beta^{12} < \beta^{12}/\lambda_1$. Pick any $\beta \in [\lambda_0\beta^{12}, \lambda_0\lambda_1\beta^{12}]$. For any $k \geq 1$, $\beta/\lambda_1^k \in [0, \lambda_0\beta^{12}] \subset [0, \beta^{12}] = R_2$. Therefore, $\pi(\beta/\lambda_1^k) = \bar{\pi}(\beta/\lambda_1^k)$ so $\beta_+(\beta/\lambda_1^k, \pi(\beta)) = \beta/\lambda_1^{k-1}$, and by RPK

$$\overline{V}(\beta/\lambda_1^{k-1}) = \overline{V}(\beta/\lambda_1^k) + A(\overline{V}(\beta/\lambda_1^k) - \mu_0).$$

Therefore

$$\overline{V}(\beta/\lambda_1^k) = \frac{\overline{V}(\beta)}{(1+A)^k} + \mu_0.$$

For any $\beta \in (0, \lambda_0\beta^{12})$ there exist a unique $\hat{\beta} \in [\lambda_0\beta^{12}, \lambda_0\lambda_1\beta^{12}]$ and $k \geq 1$ such that $\beta = \hat{\beta}/\lambda_1^k$. Thus, by the previous argument,

$$\overline{V}(\beta) = \frac{\overline{V}(\hat{\beta})}{(1+A)^k} + \mu_0.$$

Note that as $\beta \to 0$, $k \to \infty$ and $\overline{V}(\beta) \to \mu_0$. For $\beta = 0$ we let $\overline{V}(0) = \mu_0$. This completes the extension of the upper boundary for all $\beta \in [0, 1]$. $\square$

D. Behavioral Convergence With Many Types

For $\delta \in (0, 1)$ and $\beta^0 \in (0, 1)^N$, consider any equilibrium (π, α) for the game where the sender's initial reputation is β^0.

For any $0 \leq x \leq y \leq 1$, let

$$B_{[x,y]} = \{\beta \in \mathbb{R}^N_+ \mid x \leq \hat{\beta} \leq y\}$$

be the subset of reputation vectors whose total reputation is in between x and y. When $y = x$, we write B_x instead of $B_{[x,x]}$.

The strategy (π, α) is adapted to the history of the game. For the sender, for example, π selects the probabilities $\pi_{t,h}$ and $\pi_{t,\ell}$ of sending the message $m_t = H$ when the state is $\omega_t = h$ or $\omega_t = \ell$, respectively, as a function of the public history $\{\mathcal{F}_{t-1}\}$ of states $\{\omega_n\}_{n<t}$, previous messages $\{m_n\}_{n<t}$ and previous actions $\{a_n\}_{n<t}$. Denote by $\{\beta^t\}$ the (random) path of reputations generated by (π, α). For any $a \in (0, 1)$ let

$$T_a(\pi, \alpha, \beta^0, \delta) = \mathbb{E}\left[(1-\delta)\sum_{t\geq 0} \delta^t \mathbb{I}[\hat{\beta}^t \geq a]\mathbb{I}[\pi_{t,\ell} \geq \bar{\pi}(a)]\right]$$

be the expected number of periods $t \geq 1$ along the outcome path where the total reputation is above a and the sender lies with probability at least $\bar{\pi}(a) \geq \pi^*$ when $\omega_t = \ell$.

Lemma 27. For any $a \in (0,1)$ and $\epsilon > 0$ there exists $\bar{S}(a,\epsilon) < \infty$ such that for any $\delta \in (0,1)$, $\beta^0 \in B_{[0,1-\epsilon]}$ and equilibrium (π,α), $T_a(\pi,\alpha,\beta^0,\delta) \leq \bar{S}(a,\epsilon)(1-\delta)$.

Proof: Fix $(\pi,\alpha,\beta^0,\delta)$. For $t \geq 0$ we have that

$$\hat{\beta}^{t+1} = \frac{\hat{\beta}^t}{\hat{\beta}^t + (1-\hat{\beta}^t)\ell_t} \quad \text{where} \quad \hat{\pi}_t^B = [1/\hat{\beta}^t]\sum \beta_i^t \pi_i^B \quad \text{and}$$

$$\ell_t = \begin{cases} \pi_{t,h} & \text{if} \quad (m_t,\omega_t) = (H,h) \\ \infty & \text{if} \quad (m_t,\omega_t) = (L,h) \\ \pi_{t,\ell}/\hat{\pi}_t^B & \text{if} \quad (m_t,\omega_t) = (H,\ell) \\ (1-\pi_{t,\ell})/(1-\hat{\pi}_t^B) & \text{if} \quad (m_t,\omega_t) = (L,\ell). \end{cases}$$

Let $L_t = \ell_0 \times \cdots \times \ell_t$. Then, $\hat{\beta}^{t+1} \geq a$ if and only if

$$L_t \leq c \qquad \text{or} \qquad \frac{1}{\sqrt{L_t}} \geq \frac{1}{\sqrt{c}}, \quad \text{where} \quad c := \frac{(1-a)\hat{\beta}^0}{a(1-\hat{\beta}^0)}.$$

The function $f(x,y) = \sqrt{xy} + \sqrt{(1-x)(1-y)}$, $(x,y) \in [0,1]^2$, is quasiconcave in x for each y (and quasiconcave in y for each x) and attains its maximum at $x = y$ where $f(y,y) = 1$.

For any $a \in (0,1)$, define

$$\bar{\pi}(a) := \bar{\pi}(a, \pi_N^B)$$

and let

$$\bar{f} := \max\{f(x,y) \mid (x,y) \in [0,1]^2,\ x - y \geq \bar{\pi}(a) - \pi_N^B\}.$$

By continuity and since $f(x,y)$ attains its maximum if only if $x = y$, we have that $\bar{f} < 1$.

Note that

$$\begin{aligned} \mathbb{E}\left[\frac{1}{\sqrt{\ell_t}}\middle|\mathcal{F}_{t-1}\right] &= \mathbb{E}\left[\mu_0\sqrt{\pi_{t,h}} + (1-\mu_0)\left(\sqrt{\pi_{t,\ell}\hat{\pi}_t^B} + \sqrt{(1-\pi_{t,\ell})(1-\hat{\pi}_t^B)}\right) \,\middle|\, \mathcal{F}_{t-1}\right] \\ &\leq \mu_0 + (1-\mu_0) f(\pi_{t,\ell},\hat{\pi}_t^B) \leq 1. \end{aligned}$$

Let

$$\gamma = \mu_0 + (1-\mu_0)\bar{f} < 1 \qquad \text{and} \qquad I_t = \mathbb{I}[\pi_{t,\ell} \geq \bar{\pi}(a)].$$

Since $\hat{\pi}_t^B$ is a convex combination of $\{\pi_i^B\}$, $\hat{\pi}_t^B \leq \pi_N^B$. If $I_t = 1$ then $\pi_{t,\ell} \geq \bar{\pi}(a)$ and $\pi_{t,\ell} - \hat{\pi}_t^B \geq \bar{\pi}(a) - \pi_N^B$ and $\mu_0 + (1-\mu_0) f(\pi_{t,\ell},\hat{\pi}_t^B) \leq \gamma$. Therefore

$$\mathbb{E}\left[\frac{1}{\sqrt{L_t}}\,\middle|\,\mathcal{F}_{t-1}\right] = \frac{1}{\sqrt{L_{t-1}}}\mathbb{E}\left[\frac{1}{\sqrt{\ell_t}}\middle|\mathcal{F}_{t-1}\right] \leq \frac{\gamma^{I_t}}{\sqrt{L_{t-1}}}.$$

Define

$$N_t = \sum_{n \le t} I_n \qquad \text{and} \qquad M_t = \frac{1}{\sqrt{L_t}\gamma^{N_t}}.$$

N_t is the number of periods $n \le t$ where $\pi_{n,\ell} \ge \bar{\pi}(a)$. Since $N_t = N_{t-1} + I_t$, $\{M_t\}$ is a supermartingale:

$$\mathbb{E}[M_t|\mathcal{F}_{t-1}] = \frac{1}{\gamma^{N_{t-1}+I_t}}\mathbb{E}\left[\frac{1}{\sqrt{L_t}}\middle|\mathcal{F}_{t-1}\right] \le \frac{1}{\gamma^{N_{t-1}+I_t}}\frac{\gamma^{I_t}}{\sqrt{L_{t-1}}} = M_{t-1}.$$

Let τ_n be the period where $\pi_{t,\ell} \ge \bar{\pi}(a)$ for the n-th time:

$$\tau_1 = \inf\{t \ge 1 \mid I_t = 1\}, \qquad \tau_{n+1} = \inf\{t > \tau_n \mid I_t = 1\}.$$

Note that I_t is $\mathcal{F}_{t-1}$-measurable, so τ_n is $\mathcal{F}_{\tau_n-1}$-measurable and $\tau_n - 1$ is a stopping time. When $\tau_n < \infty$, $N_{\tau_n} = n$ and $\hat{\beta}^{\tau_n} \ge a$ if and only if $M_{\tau_n-1} \ge \frac{1}{\sqrt{c}\gamma^{n-1}}$. Fix $K \in \mathbb{N}$ and define $\tau_n^K = \min\{\tau_n, K\}$. On the event $[\tau_n \le K, \hat{\beta}^{\tau_n} \ge a]$ we have that $\tau_n^K = \tau_n$. Therefore,

$$[\tau_n \le K, \hat{\beta}^{\tau_n} \ge a] \subset \left[M_{\tau_n^K-1} \ge \frac{1}{\sqrt{c}\gamma^{n-1}}\right].$$

By Markov's inequality,

$$\mathbb{P}\left[M_{\tau_n^K-1} \ge \frac{1}{\sqrt{c}\gamma^{n-1}}\right] \le \sqrt{c}\gamma^{n-1}\mathbb{E}[M_{\tau_n^K-1}].$$

Note that

$$M_0 = \frac{1}{\sqrt{\ell_0}\gamma^{I_0}} \qquad \text{so} \qquad \mathbb{E}[M_0] \le 1.$$

Since $\{M_t\}$ is a supermartingale and τ_n^K is bounded, optimal sampling implies that $\mathbb{E}[M_{\tau_n^K-1}] \le \mathbb{E}[M_0] \le 1$. Letting $K \to \infty$ we obtain

$$\mathbb{P}[\tau_n < \infty,\ \hat{\beta}^{\tau_n} \ge a] \le \sqrt{c}\gamma^{n-1}.$$

Let

$$S(a) = \sum_{t \ge 0} \mathbb{I}[\hat{\beta}^t \ge a] I_t \qquad \text{and} \qquad \bar{S}(a,\epsilon) = 1 + \frac{1}{1-\gamma}\sqrt{\frac{(1-a)(1-\epsilon)}{a\epsilon}}.$$

For each $t \ge 1$ such that $I_t = 1$ there exists exactly one n such that $\tau_n = t$. Thus

$$S(a) = \mathbb{I}[\hat{\beta}^0 \ge a]I_0 + \sum_{n \ge 1}\mathbb{I}[\tau_n < \infty,\ \hat{\beta}^{\tau_n} \ge a] \qquad \text{and}$$

$$\mathbb{E}[S(a)] \le 1 + \sum_{n \ge 1}\sqrt{c}\gamma^{n-1} = 1 + \frac{\sqrt{c}}{1-\gamma} \le \bar{S}(a,\epsilon).$$

Therefore,

$$T_a(\pi, \alpha, \beta^0, \delta) = \mathbb{E}\left[(1-\delta)\sum_{t\geq 0}\delta^t \mathbb{I}[\hat{\beta}^t \geq a]\mathbb{I}[\pi_{t,\ell} \geq \bar{\pi}(a)]\right]$$
$$\leq \mathbb{E}\left[(1-\delta)\sum_{t\geq 0}\mathbb{I}[\hat{\beta}^t \geq a] I_t\right] = (1-\delta)\mathbb{E}[S(a)] \leq \bar{S}(a,\epsilon)(1-\delta). \quad \square$$

For a given δ, initial reputation β^0 and equilibrium (π, α), denote by $\{\pi_t\}$ and $\{\alpha_t\}$ the random paths generated by the equilibrium, and let $x_t = \alpha_t[\mu_0\pi_{t,h} + (1-\mu_0)\pi_{t,\ell}]$. Define

$$R_1 = [0,a) \times [0, \bar{\pi}(a)), \quad R_2 = [a,1] \times [0, \bar{\pi}(a)), \quad R_{12} = R_1 \cup R_2,$$
$$R_3 = [0,a) \times [\bar{\pi}(a), 1] \quad \text{and} \quad R_4 = [a,1] \times [\bar{\pi}(a), 1]$$

$$T_i = \mathbb{E}[(1-\delta)\sum_{t\geq 0}\delta^t \mathbb{I}[(\beta^t, \pi_{t,\ell}) \in R_i] \quad i = 1, \ldots, 4.$$

For the proposition below we consider two cases: a single behavioral type and multiple behavioral types. Assume that there exists $\underline{\delta} \in (0,1)$ such that

$$\overline{V}(\beta, \delta) \geq \hat{V} - \epsilon \quad \text{for all} \quad \beta \in B_{[\epsilon,1]} \quad \text{and} \quad \delta \in [\underline{\delta}, 1). \tag{P}$$

Recall that $\pi^* = \mu_0/(1-\mu_0)$ and that $v^* = \mu_0 + (1-\mu_0)\pi^* = 2\mu_0$. In the many-type model (P) holds with

$$\hat{V} = \mu_0 + (1-\mu_0)\pi_N^B,$$

by FL. In particular, it therefore holds in the single-type model as well. In fact, Lemma 26 shows that in the single-type model it also holds with $\hat{V} = V^M(\pi^*, \pi_B)$.

Proof of Proposition 6. Fix $\epsilon > 0$. Let $a \in (0,1)$ be such that $\bar{\pi}(a) = \pi^* + \epsilon$. Choose $\underline{\delta}$ close enough to 1 so that for all $\delta \in [\underline{\delta}, 1)$, $\bar{S}(a,\epsilon)(1-\delta) \leq \epsilon$ and $V(\beta, \delta) \geq \hat{V} - \epsilon$ for all $\beta \in B_{[\epsilon,1]}$.

Fix $\delta \in [\underline{\delta}, 1)$, $\beta_0 \in B_{[\epsilon,1-\epsilon]}$, and an equilibrium (π, α). Let $w_t = (1-\delta)\delta^t$.
Use the identity

$$|x - y| = (x - y) + 2(y - x)_+,$$

with $x = \pi^*$ and $y = \alpha_t \pi_{t,\ell}$ to get

$$\left|\pi^* - \alpha_t\pi_{t,\ell}\right| = (\pi^* - \alpha_t\pi_{t,\ell}) + 2(\alpha_t\pi_{t,\ell} - \pi^*)_+.$$

Since $\alpha_t \pi_{t,h} \leq 1$,

$$\begin{aligned} D(\pi,\alpha) &= \mu_0 \left[1 - \mathbb{E}\left[\sum_{t\geq 0} w_t \alpha_t \pi_{t,h} \right] \right] \\ &\qquad + (1-\mu_0) \left[\pi^* - \mathbb{E}\left[\sum_{t\geq 0} w_t \alpha_t \pi_{t,\ell} \right] \right] + 2(1-\mu_0) P^+ \\ &= v^* - \mathbb{E}\left[\sum_{t\geq 0} w_t \alpha_t (\mu_0 \pi_{t,h} + (1-\mu_0)\pi_{t,\ell}) \right] + 2(1-\mu_0) P^+ \\ &= v^* - V(\beta_0,\delta) + 2(1-\mu_0)P^+. \end{aligned}$$

where

$$P^+ := \mathbb{E}\left[\sum_{t\geq 0} w_t (\alpha_t \pi_{t,\ell} - \pi^*)_+ \right].$$

It remains to bound P^+.

For $(\hat{\beta}_t, \pi_{t,\ell}) \in R_{12}$ we have $\pi_{t,\ell} < \bar{\pi}(a) = \pi^* + \epsilon$, so $\alpha_t \pi_{t,\ell} \leq \pi_{t,\ell} \leq \pi^* + \epsilon$ and thus $(\alpha_t \pi_{t,\ell} - \pi^*)_+ \leq \epsilon$. For $(\hat{\beta}_t, \pi_{t,\ell}) \in R_3$ we have $\alpha_t = 0$, so $(\alpha_t \pi_{t,\ell} - \pi^*)_+ = 0$. Finally, for $(\hat{\beta}_t, \pi_{t,\ell}) \in R_4$, $(\alpha_t \pi_{t,\ell} - \pi^*)_+ \leq 1$. Therefore, by previous lemma,

$$P^+ \leq \epsilon T_{12} + T_4 \leq \epsilon + T_4 \leq \epsilon + S(a,\epsilon)(1-\delta) \leq 2\epsilon.$$

Hence

$$D(\pi,\alpha) \leq v^* - V(\beta_0,\delta) + 4(1-\mu_0)\epsilon \leq v^* - (\hat{V} - \epsilon) + 4(1-\mu_0)\epsilon \leq (v^* - \hat{V}) + 5\epsilon. \qquad \square$$

References

Abreu, Dilip and Faruk Gul (2000), “Bargaining and Reputation,” *Econometrica*, **68**(1), 85–117.

Abreu, Dilip, Paul Milgrom and David Pearce (1991), “Information and Timing in Repeated Partnerships,” *Econometrica*, **59**(6), 1713–1733.

Abreu, Dilip, David Pearce and Ennio Stacchetti (1990), “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, **58**(5), 1041–1064.

Abreu, Dilip and David Pearce (2007), “Bargaining, Reputation, and Equilibrium Selection in Repeated Games with Contracts,” *Econometrica*, **75**(3), 653–710.

Aldrich, Eric M., Jesús Fernández-Villaverde, A. Ronald Gallant and Juan F. Rubio-Ramirez (2011), “Tapping the Supercomputer under Your Desk: Solving Dynamic Equilibrium Models with Graphics Processors,” *Journal of Economic Dynamics and Control*, **35**(3), 386–393.

Andreasen, Martin (2012), “On the Effects of Rare Disasters and Uncertainty Shocks for Risk Premia in Non-Linear DSGE Models,” *Review of Economic Dynamics*, **15**(3), 295–316.

Aumann, Robert and Sergiu Hart (2003), “Long Cheap Talk,” *Econometrica*, **71**(6), 1619–1660.

Aumann, Robert and Michael Maschler (1995), “Repeated Games with Incomplete Information,” MIT Press, Chapter 1.

Backus, David and Edward Driffill (1985), “Inflation and Reputation,” *American Economic Review*, **75**, 530–538.

Bénabou, Roland and Guy Laroque (1992), “Using Privileged Information to Manipulate Markets: Insiders, Gurus, and Credibility,” *Quarterly Journal of Economics*, **107**, 921–958.

Benoît, Jean-Pierre and Juan Dubra (2011), “Apparent Overconfidence,” *Econometrica*, **79**(5), 1591–1625.

Bergemann, Dirk and Stephen Morris (2016), “Bayes Correlated Equilibrium and the Comparison of Information Structures in Games,” *Theoretical Economics*, **11**, 487–522.

Bergemann, Dirk and Stephen Morris (2019), “Information Design: A Unified Perspective,” *Journal of Economic Literature*, **57**(1), 44–95.

Best, James and Daniel Quigley (2023), “Persuasion for the Long Run,” Forthcoming in *Journal of Political Economy.*

Borovicka, Jaroslav and John Stachurski (2020), "Necessary and Sufficient Conditions for Existence and Uniqueness of Recursive Utilities," *Journal of Finance*, **75**(3), 1457–1493.

Crawford, Vincent and Joel Sobel (1982), "Strategic Information Transmission," *Econometrica*, **50**(6), 1431–1451.

Cripps, Martin W., George J. Mailath and Larry Samuelson (2004), "Imperfect Monitoring and Impermanent Reputations," *Econometrica*, **72**(2), 407–432.

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

Ely, Jeffrey (2017), "Beeps," *American Economic Review*, **107**(1), 31–53.

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

Fanning, Jack (2016), "Reputational Bargaining and Deadlines," *Econometrica*, **84**(3), 1131–1179.

Fernandes, Ana and Christopher Phelan (2000), "A Recursive Formulation for Repeated Agency with History Dependence," *Journal of Economic Theory*, **91**, 223–247.

Fernández-Villaverde, Jesús and Oren Levintal (2018), "Solution Methods for Models with Rare Disasters," *Quantitative Economics*, **9**(2), 903–944.

Fernández-Villaverde, Jesús and David Zarruk Valencia (2018), "A Practical Guide to Parallelization in Economics," Working paper.

Fudenberg, Drew, Ying Gao and Harry Pei (2022), "A Reputation for Honesty," *Journal of Economic Theory*, **204**, 105508.

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

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

Fudenberg, Drew and David Levine (1994), "Efficiency and Observability with Long-Run and Short-Run Players," *Journal of Economic Theory*, **62**(1), 103–135.

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

Fudenberg, Drew, David Levine and Eric Maskin (1994), "The Folk Theorem with Imperfect Public Information," *Econometrica*, **62**(5), 997–1039.

Golosov, Mikhail, Vasiliki Skreta, Aleh Tsyvinski and Andrea Wilson (2014), "Dynamic Strategic Information Transmission," *Journal of Economic Theory*, **151**, 304–341.

Golosov, Mikhail, Aleh Tsyvinski and Ivan Werning (2007), "New Dynamic Public Finance: A User's Guide," NBER Chapters, in *NBER Macroeconomics Annual 2006*, **21**, 317–388.

Goltsman, Maria, Johannes Hörner, Gregory Pavlov and Francesco Squintani (2009), "Mediation, Arbitration and Negotiation," *Journal of Economic Theory*, **144**(4), 1397–1420.

Gossner, Olivier (2011), "Simple Bounds on the Value of Reputation," *Econometrica*, **79**(5), 1627–1641.

Gourio, François (2012), "Disaster Risk and Business Cycles," *American Economic Review*, **102**(6), 2734–2766.

Gourio, François (2013), "Credit Risk and Disaster Risk," *American Economic Journal: Macroeconomics*, **5**(3), 1–34.

Gorno, Leandro and Felipe S. Iachan (2020), "Competitive Real Options under Private Information," *Journal of Economic Theory*, **185**, 104945.

Hörner, Johannes, Satoru Takahashi and Nicolas Vieille (2015), "Truthful Equilibria in Dynamic Bayesian Games," *Econometrica*, **83**(5), 1795–1848.

Isoré, Marlène and Urszula Szczerbowicz (2017), "Disaster Risk and Preference Shifts in a New Keynesian Model," *Journal of Economic Dynamics and Control*, **79**, 97–125.

Kamenica, Emir and Matthew Gentzkow (2011), "Bayesian Persuasion," *American Economic Review*, **101**(6), 2590–2615.

Kamenica, Emir (2019), "Bayesian Persuasion and Information Design," *Annual Review of Economics*, **11**, 249–272.

Kandori, Michihiro and Ichiro Obara (2006), "Efficiency in Repeated Games Revisited: The Role of Private Strategies," *Econometrica*, **74**(2), 499–519.

Koessler, Frédéric and Françoise Forges (2008a), "Multistage Communication with and without Verifiable Types," *International Game Theory Review*, **10**, 145–164.

Koessler, Frédéric and Françoise Forges (2008b), "Transmission Stratégique de l'Information et Certification," *Annales d'Economie et Statistiques*, **89**, 3–61.

Kostadinov, Rumen and Aditya Kuvalekar (2022), "Learning in Relational Contracts," *American Economic Journal: Microeconomics*, **14**(1), 284–329.

Kreps, David and Bob Wilson (1982), "Reputation and Imperfect Information," *Journal of Economic Theory*, **27**(2), 253–279.

Kreps, David, Paul Milgrom, John Roberts and Robert Wilson (1982), "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, **27**(2), 245–252.

Krishna, Vijay and John Morgan (2004), “The Art of Conversation: Eliciting Information from Experts Through Multi-Stage Communication,” *Journal of Economic Theory*, **117**(2), 147–179.

Kuvalekar, Aditya, Elliot Lipnowski and João Ramos (2022), “Goodwill in Communication,” *Journal of Economic Theory*, **203**, 105467.

Li, Yingkai and Harry Pei (2021), “Equilibrium Behaviors in Repeated Games,” *Journal of Economic Theory*, **193**, 105222.

Liu, Qingmin (2011), “Information Acquisition and Reputation Dynamics,” *Review of Economic Studies*, **78**(4), 1400–1425.

Liu, Qingmin and Andrzej Skrzypacz (2014), “Limited Records and Reputation Bubbles,” *Journal of Economic Theory*, **151**, 2–29.

Mailath, George and Larry Samuelson (2001), “Who Wants a Good Reputation?,” *Review of Economic Studies*, **68**, 415–441.

Mailath, George and Larry Samuelson (2006), *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press.

Margaria, Chiara and Alexander Smolin (2015), “Dynamic Communication with Biased Senders,” *Games and Economic Behavior*, **110**, 330–339.

Mathevet, Laurent, Jacopo Perego and Ina Taneva (2019), “On Information Design in Games,” *Journal of Political Economy*, **128**(4), 1370–1404.

Milgrom, Paul and John Roberts (1982), “Predation, Reputation and Entry Deterrence,” *Journal of Economic Theory*, **27**(2), 280–312.

Morris, Stephen (2001), “Political Correctness,” *Journal of Political Economy*, **109**(2), 231–265.

Ottaviani, Marco and Peter Sorensen (2006a), “Professional Advice,” *Journal of Economic Theory*, **126**(1), 120–142.

Ottaviani, Marco and Peter Sorensen (2006b), “Reputational Cheap Talk,” *RAND Journal of Economics*, **37**(1), 155–175.

Pei, Harry (2021), “Trust and Betrayals: Reputational Payoffs and Behaviors without Commitment,” *Theoretical Economics*, **16**(2), 449–475.

Pei, Harry (2022), “Reputation for Playing Mixed Actions: A Characterization Theorem,” *Journal of Economic Theory*, **201**, 105438.

Pei, Harry (2023), “Repeated Communication with Private Lying Costs,” *Journal of Economic Theory*, **210**, 105668.

Petrosky-Nadeau, Nicolas, Lu Zhang and Lars-Alexander Kuehn (2018), “Endogenous Disasters,” *American Economic Review*, **108**(8), 2212–2245.

Phelan, Christopher (2006), "Public Trust and Government Betrayal," *Journal of Economic Theory*, **130**, 27–43.

Phelan, Christopher and Ennio Stacchetti (2001), "Sequential Equilibria in a Ramsey Tax Model," *Econometrica*, **69**(6), 1491–1518.

Radner, Roy (1985), "Repeated Principal-Agent Games with Discounting," *Econometrica*, **53**(5), 1173–1198.

Rayo, Luis and Ilya Segal (2010), "Optimal Information Disclosure," *Journal of Political Economy*, **118**(5), 949–987.

Renner, Philipp and Simon Scheidegger (2020), "Machine Learning for Dynamic Incentive Problems," Working paper.

Shamsie, Jamal (2003), "The Context of Dominance: An Industry-Driven Framework for Exploiting Reputation," *Strategic Management Journal*, **24**(3), 199–215.

Sobel, Joel (1985), "A Theory of Credibility," *Review of Economic Studies*, **52**(4), 557–573.

Taneva, Ina (2019), "Information Design," *American Economic Journal: Microeconomics*, **11**(4), 151–185.

Wolitzky, Alexander (2012), "Reputational Bargaining with Minimal Knowledge of Rationality," *Econometrica*, **80**(5), 2047–2087.