Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 210 (2023) 105668

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Repeated communication with private lying costs

Harry Pei [1]

Department of Economics, Northwestern University, United States of America

Received 28 September 2022; final version received 1 March 2023; accepted 27 April 2023
Available online 5 May 2023

Abstract

I study a repeated communication game between a patient sender and a sequence of myopic receivers. The sender has persistent private information about his cost of lying. Each period, the sender privately observes an i.i.d. state and recommends an action. The receiver takes an action after observing this recommendation and the full history of states and recommendations. I provide conditions under which a sufficiently patient sender can attain his optimal commitment payoff in the repeated game. My results provide justifications for the commitment assumption in Bayesian persuasion models using a repeated communication game without any commitment. My results also imply that the sender's equilibrium payoff is not monotone with respect to the receiver's belief about his lying cost.

JEL classification: C72; C73; D82; D83

Keywords: Repeated communication game; Persistent private information; Bayesian persuasion

1. Introduction

The commitment power of informed experts affects communication outcomes. In Crawford and Sobel (1982)'s seminal model, an expert's temptation to mislead his audience undermines his credibility. When an expert can commit to a disclosure policy, as in the model of Kamenica

E-mail address: harrydp@northwestern.edu.

[1] I thank the associate editor, three anonymous referees, James Best, Eddie Dekel, Piotr Dworczak, Drew Fudenberg, Yingni Guo, Elliot Lipnowski, Daniel Luo, Delong Meng, Wojciech Olszewski, Daniel Quigley, Teck Yong Tan, Matthew Thomas, Eran Shmaya, Joel Sobel, Bruno Strulovici, and Alex Wolitzky for helpful comments. I acknowledge financial support from NSF Grant SES-1947021.

https://doi.org/10.1016/j.jet.2023.105668
0022-0531/

and Gentzkow (2011), his messages are more credible and can have more influence over others' decisions.

In practice, experts often face credibility issues when committing to disclosure policies. For example, real estate agents receive high commissions when their clients purchase overpriced properties, and pharmaceutical lobbyists have incentives to recommend policies that benefit pharmaceutical companies but hurt social welfare. These experts' optimal disclosure policies require them to reveal unfavorable information, such as that the property they are selling is overpriced or the policy they are lobbying for is socially harmful. However, if they are aware of such unfavorable information, honoring their commitment to reveal it is against their own interests.

A plausible foundation for an expert's commitment is that he communicates with *a sequence of receivers*, one in each period. Each receiver observes the optimal actions of previous periods and is able to compare them to the expert's recommendations. However, when the expert's optimal disclosure policy is stochastic, which is the case in the leading example of Kamenica and Gentzkow (2011), the result in Fudenberg et al. (1990) implies that the expert can never attain his optimal commitment payoff in the repeated communication game, no matter how patient he is.

Several insightful papers bypass the above problem by introducing alternative forms of commitment, under which a patient expert can attain his optimal commitment payoff in the repeated game. For example, Best and Quigley (2022) assume that there is a third-party who can commit to a public record system that determines which information the receivers can observe about the past history. Mathevet et al. (2022) assume that with positive probability, the expert is a commitment type who uses his optimal disclosure policy in every period.

I examine the extent to which patient experts can restore their commitment power when it is common knowledge that they *cannot* honor any promise against their own interests. I focus on a repeated version of the leading example in Kamenica and Gentzkow (2011). Each period, a patient sender privately observes an i.i.d. state, which is either *good* or *bad*, and recommends a *good action* or a *bad action*. He may also send a cheap talk message that has no intrinsic meaning together with his action recommendation. A myopic receiver chooses an action after observing the sender's recommendation and message, together with the history of states, recommendations, and messages.

The receiver prefers to match her action with the state. The sender prefers the good action regardless of the state. My modeling innovation is that the sender incurs a psychological cost of lying whenever his action recommendation does not match the realized state.[2] I assume that this lying cost is fixed over time and is the sender's private information, which I call his *type*.

I start from a setting where the lying costs of all types are lower than the sender's benefit from the good action, i.e., every type prefers to mislead the receiver in the bad state when his recommendation is taken at face value. Theorem 1 characterizes the *highest equilibrium payoff* for each sender type in the limit where the sender is arbitrarily patient. This payoff depends only on his own lying cost and *the highest lying cost* in the support of the receivers' prior belief. Intuitively, when a sender type decides which of the other types to imitate, he prefers to imitate types that occur with higher probabilities, since it takes fewer periods to build reputations for behaving

[2] The relevance and significance of lying costs have been established experimentally by Gneezy (2005). They have been incorporated into models of strategic information transmission (Kartik et al., 2007; Chen et al., 2008; Kartik, 2009) and mechanism design (Ortner, 2015). See Sobel (2020) for a detailed discussion of lying costs in economic models. I assume that the sender does not incur any lying cost regardless of the cheap talk message he sends given that cheap talk messages have no intrinsic meaning.

like those types; he also prefers to imitate types that have higher lying costs, since those types are more trustworthy and reputations for behaving like them are more valuable. When the sender is more patient, he puts more weight on his long-term payoff after he has established a reputation and puts less weight on his stage-game payoff in periods where he builds his reputation. As a result, only the sender's true lying cost as well as the lying cost of the highest-cost type affect his highest equilibrium payoff.

My result implies that senders of all types except for the highest-cost one can benefit from their persistent private information, in the sense that their highest equilibrium payoff in the repeated incomplete information game is *strictly* greater than their highest equilibrium payoff in a repeated game where their lying cost is common knowledge. This is because the highest-cost type is the most trustworthy type, so he has no good candidate to imitate in the incomplete information game.

My result also implies that the highest equilibrium payoff for every type converges to his optimal commitment payoff as the highest lying cost converges to the sender's benefit from the good action. This can be viewed as a justification for the commitment assumption in Bayesian persuasion models: Even though it is common knowledge that no type can commit and every type prefers to mislead the receivers in the stage game, all types of the sender can approximately attain their optimal commitment payoffs in the repeated communication game as long as the set of types is rich enough.

Next, I allow for *ethical types* whose lying costs exceed their benefits from the good action. Theorem 2 shows that there exist equilibria where all ethical types attain their optimal commitment payoffs and all non-ethical types attain the payoffs described in Theorem 1. This is because in the first period, the sender can credibly reveal whether he is ethical via a cheap talk message. After he claimed to be ethical, he always recommends the receiver-optimal action and the receivers follow his recommendation. After he claimed to be non-ethical, players coordinate on the sender-optimal equilibrium in the game without ethical types. Theorem 2 implies that all types can approximately attain their optimal commitment payoffs as the set of types becomes rich enough.

In the case where the sender can only recommend actions but *cannot* send cheap talk messages, Theorem 3 shows that all the non-ethical types can attain their optimal commitment payoffs *if and only if* the lying cost of the highest type is below some cutoff. This suggests that the presence of types with high lying costs can lower the non-ethical types' equilibrium payoffs. This observation is driven by a novel *outside option effect*. Suppose there is a type who has an infinite lying cost. In equilibrium, this type will never lie, so the other ethical types can secure their full disclosure payoffs by imitating his behavior. As a result, the non-ethical types cannot lie while pooling with the ethical types and hence cannot obtain their optimal commitment payoffs. More generally, each ethical type's lowest equilibrium payoff increases with the highest lying cost in the support of the receivers' prior belief, and the non-ethical types cannot attain their optimal commitment payoffs when every ethical type's payoff from imitating the highest-cost type exceeds some threshold.

Related literature: My paper is related to the literature on repeated communication games pioneered by Aumann and Maschler (1965), Sobel (1985), and Benabou and Laroque (1992).[3] Best

[3] Recent contributions to the repeated communication game literature include those of Renault et al. (2013), Margaria and Smolin (2018), Meng (2021), and Kuvalekar et al. (2022).

and Quigley (2022) and Mathevet et al. (2022) use this framework to provide justifications for the sender's commitment power in Bayesian persuasion models.[4]

Best and Quigley (2022) show that the sender can attain his optimal commitment payoff if a third party can commit to a particular public record system that determines the receiver's information about the game's history. Mathevet et al. (2022) show that a patient sender can attain his optimal commitment payoff whenever there is a positive probability that he is a *commitment type* who uses his optimal disclosure policy in every period. In contrast to those papers, I provide a justification for the commitment assumption using a repeated communication game where it is *common knowledge* that the sender cannot honor any promise against his own interest.

One limitation of my approach is that it restricts attention to a special class of games. Although my upper bound on the sender's payoff can be extended to more general settings, to establish that these bounds are attainable, one needs to construct equilibria in repeated communication games with persistent private information, which is an intractable problem.[5] Hence, I view my main result as a proof of concept rather than a general lesson. It demonstrates that it is possible for a patient sender to obtain his optimal commitment payoff even when it is common knowledge that (i) he *cannot* commit to disclosure policies and that (ii) he has a strict incentive to deceive the receivers.

I also characterize the extent to which the sender can partially overcome his lack-of-commitment problem in environments where he cannot attain his optimal commitment payoff. This is related to Lipnowski et al. (2022), who characterize the sender's highest equilibrium payoff in a one-shot game when he can commit to a disclosure policy with positive probability. In Guo and Shmaya (2021) and Nguyen and Tan (2021), the sender faces a lying cost, chooses an information structure, and receives information about the state according to that information structure. By contrast, the sender in my model cannot commit to receiving coarser information about the state.

This paper is also related to the literature on repeated games with incomplete information, pioneered by Aumann and Maschler (1965). Hart (1985), Shalev (1994), and Pęski (2014) characterize the set of equilibrium payoffs when all players are patient. When the uninformed player is impatient, Cripps and Thomas (2003) show that every equilibrium payoff must satisfy the conditions in Shalev (1994), although some of the payoffs that satisfy this necessary condition cannot be attained in any equilibrium. By contrast, I characterize a patient informed sender's highest equilibrium payoff when his opponents' discount factor is either zero or close to zero. This paper is also related to Pei (2021) who characterizes a patient seller's highest equilibrium payoff when he has private information about his production cost. I elaborate on the differences between the two papers in Section 6.

2. Model

Time is indexed by $t = 0, 1, 2, \ldots$. A patient sender with discount factor $\delta \in (0, 1)$ interacts with a different short-lived receiver in each period. In period t, the sender privately observes a state $\omega_t \in \Omega \equiv \{g, b\}$, where $\{\omega_t\}_{t \in \mathbb{N}}$ are drawn i.i.d. from a distribution where the probability of

[4] Titova (2022) provides a justification of the commitment assumption in Bayesian persuasion models using a *static* communication game where the sender cannot commit but can disclose verifiable evidence.

[5] The challenges of constructing equilibria in repeated incomplete information games are also explained in Section 4 of Mathevet et al. (2022). That section focuses on the $2 \times 2 \times 2$ example in Kamenica and Gentzkow (2011) and computes players' equilibrium behaviors in a reputation model with a mixed-strategy commitment type.

$\omega_t = g$ is p. The sender makes an action recommendation $r_t \in R \equiv \{g, b\}$ and may also send a cheap talk message $m_t \in M$ that has no intrinsic meaning. The period-t receiver takes an action $y_t \in Y \equiv \{g, b\}$ after observing (r_t, m_t) as well as the public history $h^t \equiv \{y_s, \omega_s, r_s, m_s\}_{s=0}^{t-1}$.

The sender's stage-game payoff is $\mathbf{1}\{y_t = g\} - c \cdot \mathbf{1}\{r_t \neq \omega_t\}$, where $c \in \mathcal{C} \equiv \{c_1, ..., c_n\} \subset [0, 1)$ is his private information (or his *type*), is constant over time, and is independent of $\{\omega_t\}_{t\in\mathbb{N}}$. I interpret c as the sender's cost of lying relative to his benefit from the receiver's good action, which is incurred whenever his action recommendation does not match the realized state. The sender incurs no lying cost regardless of the cheap talk message he sends since those messages have no intrinsic meaning.

The receiver's prior belief about c is $\boldsymbol{\pi} \equiv (\pi_1, ..., \pi_n) \in \Delta(\mathcal{C})$, where π_j is the probability of type c_j. The receiver's payoff is normalized to 0 when $y_t = b$. I assume that the receiver's payoff is 1 when $(y_t, \omega_t) = (g, g)$ and is -1 when $(y_t, \omega_t) = (g, b)$. My results extend as long as the receiver's payoff is strictly positive when $(y_t, \omega_t) = (g, g)$ and is strictly negative when $(y_t, \omega_t) = (g, b)$.

Assumption 1. $n \geq 2$, $\boldsymbol{\pi}$ has full support, $p \in (0, \frac{1}{2})$, and $0 \leq c_n < c_{n-1} < ... < c_2 < c_1$.

Assumption 1 requires that (i) the receivers face uncertainty about the sender's type, (ii) the receivers have no incentive to take action g under their prior belief about ω_t, and (iii) every type has a non-negative lying cost. I distinguish between types with lying costs strictly less than 1, i.e., *non-ethical types*, and types with lying costs greater than 1, i.e., *ethical types*. Intuitively, when the receivers always follow the sender's recommendation, every non-ethical type prefers to deceive the receivers in the bad state and every ethical type prefers to recommend the receiver-optimal action.

The sender's *pure stage-game strategy* is $\mathbf{a}_1 : \Omega \to R \times M$, with $\mathbf{a}_1 \in \mathbf{A}_1$. The receiver's *pure stage-game strategy* is $\mathbf{a}_2 : R \times M \to Y$, with $\mathbf{a}_2 \in \mathbf{A}_2$. Let $\mathbf{a}_{1,t}$ and $\mathbf{a}_{2,t}$ denote the sender's and the receiver's stage-game strategies in period t. Let $u_1(c, \mathbf{a}_1, \mathbf{a}_2)$ denote type-c sender's stage-game payoff. Type-c sender maximizes his discounted average payoff $\sum_{t=0}^{\infty}(1 - \delta)\delta^t u_1(c, \mathbf{a}_{1,t}, \mathbf{a}_{2,t})$. Since there are n types, the sender's *payoff* is an n-dimensional vector $\boldsymbol{v} \equiv (v_1, v_2, ..., v_n) \in \mathbb{R}^n$, where the jth entry v_j represents the discounted average payoff of type c_j. Let $u_2(\mathbf{a}_1, \mathbf{a}_2)$ denote the receiver's stage-game payoff. Let $\mathcal{H}$ denote the set of public histories $h^t \equiv \{y_s, \omega_s, r_s, m_s\}_{s=0}^{t-1}$. The *repeated-game strategy* of type-c sender is $\sigma_c : \mathcal{H} \to \Delta(\mathbf{A}_1)$, and that of the receiver's is $\sigma_2 : \mathcal{H} \to \Delta(\mathbf{A}_2)$.

The solution concept is Nash equilibrium, which is a strategy profile $\sigma \equiv \big((\sigma_c)_{c\in\mathcal{C}}, \sigma_2\big)$. My results extend to Perfect Bayesian equilibrium defined in Fudenberg and Tirole (1991), where players' beliefs satisfy *no-signaling-what-you-don't-know* and respect Bayes rule at every on-path history.

I present several useful benchmarks. First, suppose the sender *can* commit to a disclosure policy $\boldsymbol{\alpha} : \Omega \to \Delta(R \times M)$ before he observes ω_t. His payoff is p if he commits to always recommend the receiver-optimal action. If he chooses a disclosure policy in order to maximize his payoff, then (i) for every *ethical type*, it is optimal to recommend the receiver-optimal action; (ii) for every *non-ethical type*, it is optimal to recommend the good action in the good state, and to recommend the good action with probability $\rho^* \equiv \frac{p}{1-p}$ in the bad state, which makes the receiver indifferent when the sender recommends the good action. Therefore, type-c_j sender's *optimal commitment payoff* is

$$v_j^{**} \equiv \begin{cases} p + p(1 - c_j) & \text{if } c_j < 1, \\ p & \text{if } c_j \geq 1. \end{cases} \tag{2.1}$$

For future reference, I write the sender's optimal disclosure policy as a distribution over his *pure stage-game strategies*. I define the sender's *honest strategy* $\mathbf{a}_1^H$ and *lying strategy* $\mathbf{a}_1^L$ as:

$$\mathbf{a}_1^H(\omega) \equiv \begin{cases} g & \text{if } \omega = g, \\ b & \text{if } \omega = b, \end{cases} \quad \text{and} \quad \mathbf{a}_1^L(\omega) \equiv \begin{cases} g & \text{if } \omega = g, \\ g & \text{if } \omega = b. \end{cases} \tag{2.2}$$

I define the receiver's *trusting strategy* $\mathbf{a}_2^T$ and *non-trusting strategy* $\mathbf{a}_2^N$ as:

$$\mathbf{a}_2^T(r) \equiv \begin{cases} g & \text{if } r = g, \\ b & \text{if } r = b, \end{cases} \quad \text{and} \quad \mathbf{a}_2^N(r) \equiv \begin{cases} b & \text{if } r = g, \\ b & \text{if } r = b. \end{cases} \tag{2.3}$$

It is optimal for every ethical type to commit to $\mathbf{a}_1^H$. For every non-ethical type, it is optimal to commit to $\rho^*\mathbf{a}_1^L + (1-\rho^*)\mathbf{a}_1^H$, under which the receiver is indifferent between $\mathbf{a}_2^T$ and $\mathbf{a}_2^N$.

Next, in a repeated game where the sender cannot commit and his lying cost is common knowledge, the sender's equilibrium payoff is no more than p no matter how patient he is. This is because when $p < 1/2$, every on-path history of every equilibrium belongs to one of the following two classes:

1. The receiver has a strict incentive to take action b regardless of (r, m).
2. The receiver has an incentive to take action g after observing some (r, m). Since $p < 1/2$, there exists (r', m') such that (i) the sender sends (r', m') with positive probability when the state is b and (ii) the receiver has a strict incentive to take action b upon receiving (r', m').

For any equilibrium, consider a deviation for the sender in which (i) at any history that belongs to the first class, he follows his equilibrium strategy, and (ii) at any history that belongs to the second class, he uses his equilibrium strategy in state g and sends (r', m') with probability 1 in state b. This deviation is the sender's best reply in the repeated game, from which his expected stage-game payoff at every history is no more than p. This implies that the sender's equilibrium payoff is no more than p. As a corollary, when δ is close to 1, the sender's equilibrium payoff cannot exceed p in reputation games with one rational type and one or more pure-strategy commitment types.

Remark: In my model, the sender can only recommend one of the two actions (which have intrinsic meaning and may incur lying costs) and may also send a cheap talk message together with his action recommendation. My assumption on the message space is *not* without loss since it is well-known from Green and Laffont (1986) that the standard revelation principle does not apply in settings where messages are costly. For example, including an additional message m^* where the lying cost of sending m^* in state ω being $c(m^*, \omega)$ may affect the sender's equilibrium payoffs.

3. Upper bound on the sender's equilibrium payoff

This section establishes two necessary conditions that every Nash equilibrium must satisfy. These conditions lead to an upper bound on each type of the patient sender's equilibrium payoff.

First, I show that type c_1's equilibrium payoff *cannot* exceed his highest equilibrium payoff in a repeated game where his lying cost is common knowledge. Intuitively, type c_1 has no good candidate to imitate, so he cannot benefit from having persistent private information about his lying cost.

Proposition 1. *Fix any equilibrium σ under any discount factor δ. For any history h^t that occurs with positive probability under σ, suppose c_i is the highest lying cost in the support of the receiver's belief at h^t, then type-c_i sender's continuation value at h^t is no more than p.*

The proof uses an induction argument on *the number of types in the support of the receiver's belief*. If there is only one type, the conclusion follows from the argument in Fudenberg et al. (1990). I explain how the inductive step works focusing on the case where (i) the sender can only recommend actions but *cannot* send cheap talk messages, and (ii) the sender always recommends the good action in the good state. The general argument can be found in Appendix A.

Suppose that the highest-cost type in the support of the receiver's belief obtains a continuation value of at most p at every history where there are *no more than k types* in that support.

Consider an on-path history h^t such that there are $k+1$ types in the support of the receiver's belief, where the highest lying cost among those $k+1$ types is denoted by c_i. I construct strategy $\widetilde{\sigma}_{c_i}$ based on type c_i's equilibrium strategy σ_{c_i}. For every $h^s \succeq h^t$, if σ_{c_i} recommends the bad action with positive probability when the state is bad, then $\widetilde{\sigma}_{c_i}$ recommends the bad action with probability 1 when the state is bad; if σ_{c_i} recommends the good action for sure when the state is bad, then $\widetilde{\sigma}_{c_i}$ recommends the good action with probability 1 when the state is bad. By construction, $\widetilde{\sigma}_{c_i}$ is type c_i's best reply to the receiver's equilibrium strategy, from which he obtains his equilibrium payoff.

When type c_i uses $\widetilde{\sigma}_{c_i}$, his stage-game payoff is no more than p *except* at histories where (i) he recommends the good action in the bad state, and (ii) the receiver takes the good action with positive probability after he recommends the good action. At every such history, the receiver's incentive to take the good action together with Bayes rule implies that there exist types who recommend the bad action in the bad state. Since type c_i does not recommend the bad action in the bad state, there are at most k types in the support of the receivers' posterior belief after they observe $(\omega_s, r_s) = (b, b)$. By induction hypothesis, the highest type in the support of their posterior, denoted by c_j, receives payoff no more than p. Type c_i's continuation value at h^s is no more than p, since (i) type c_j's continuation value at h^s is no more than p if he deviates to type c_i's equilibrium strategy, and (ii) type c_i's payoff from any strategy is no more than that of type c_j's given that $c_i > c_j$.

Next, I derive an upper bound on the frequency with which the sender plays $\mathbf{a}_1^L$ when the receivers play $\mathbf{a}_2^T$. Fix an equilibrium $\sigma \equiv \big((\sigma_c)_{c\in\mathcal{C}}, \sigma_2\big)$. Let $\gamma^j \equiv \{\gamma^j(\mathbf{a}_1, \mathbf{a}_2)\}_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2}$, where

$$\gamma^j(\mathbf{a}_1, \mathbf{a}_2) \equiv \mathbb{E}^{(\sigma_{c_j},\sigma_2)}\Big[\sum_{t=0}^{\infty}(1-\delta)\delta^t \mathbf{1}\{(\mathbf{a}_{1,t}, \mathbf{a}_{2,t})=(\mathbf{a}_1, \mathbf{a}_2)\}\Big] \text{ for every } (\mathbf{a}_1, \mathbf{a}_2) \in \mathbf{A}_1 \times \mathbf{A}_2. \tag{3.1}$$

Intuitively, $\gamma^j(\mathbf{a}_1, \mathbf{a}_2)$ is the discounted frequency of $(\mathbf{a}_1, \mathbf{a}_2)$ when the sender uses type c_j's equilibrium strategy σ_{c_j} and the receivers play their equilibrium strategy σ_2. By definition, for every $c \in \mathcal{C}$ and $j \in \{1, 2, ..., n\}$, type-c sender's stage-game payoff from γ^j equals his discounted average payoff in the repeated game when he plays σ_{c_j} and the receivers play σ_2. For every $\gamma \in \Delta(\mathbf{A}_1 \times \mathbf{A}_2)$ and $\mathbf{a}_2 \in \mathbf{A}_2$, let $\gamma(\cdot|\mathbf{a}_2) \in \Delta(\mathbf{A}_1)$ be the conditional distribution of the sender's stage-game strategy when the joint distribution is γ and the receiver's stage-game strategy is $\mathbf{a}_2$. Proposition 2 implies that when the sender is patient, he cannot lie too frequently in periods where the receivers play $\mathbf{a}_2^T$.

Proposition 2. *For every $\varepsilon > 0$, there exists $\underline{\delta} \in (0,1)$ such that for every $\delta > \underline{\delta}$, every Nash equilibrium $\big((\sigma_c)_{c\in\mathcal{C}}, \sigma_2\big)$ under δ, every $j \in \{1,2,...,n\}$, and every $\mathbf{a}_2 \in \mathbf{A}_2$, if γ^j assigns probability more than ε to $\mathbf{a}_2$, then $\mathbf{a}_2$ is an ε-best reply to $\gamma^j(\cdot|\mathbf{a}_2)$.*

The proof is in Appendix B. The intuition follows from the payoff upper bound result in Fudenberg and Levine (1992) and Gossner (2011). Since the receivers can observe the history of (ω, r, m), they can statistically identify the sender's stage-game strategy. If the sender plays σ_{c_j}, then in all except for a bounded number of periods, the receivers' predictions about the sender's stage-game strategy are close to type c_j's equilibrium stage-game strategy. Therefore, the receivers will play a best reply to some $\boldsymbol{\alpha}_1 \in \Delta(\mathbf{A}_1)$ that is close to type c_j's equilibrium stage-game strategy. When the receivers' predictions are sufficiently precise, their stage-game strategy is an ε-best reply to type c_j's equilibrium stage-game strategy. This implies that as $\delta \to 1$, periods where the receivers do not play any ε-best reply have negligible impact on the discounted frequencies. Therefore, every $\mathbf{a}_2$ that occurs with probability bounded away from zero in equilibrium must be an ε-best reply to $\gamma^j(\cdot|\mathbf{a}_2)$.

The two necessary conditions lead to an upper bound on every type of the sender's equilibrium payoff. For every j, if type c_1 uses type c_j's equilibrium strategy, then his discounted average payoff is $\sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma^j(\mathbf{a}_1,\mathbf{a}_2)u_1(c_1,\mathbf{a}_1,\mathbf{a}_2)$. Since type c_1's equilibrium payoff is no more than p and he prefers his equilibrium strategy to that of any other type's, we have:

$$\sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma^j(\mathbf{a}_1,\mathbf{a}_2)u_1(c_1,\mathbf{a}_1,\mathbf{a}_2) \leq p, \text{ for every } j \in \{1,2,...,n\}. \tag{3.2}$$

This together with Proposition 2 implies that for every $\varepsilon > 0$, there exists $\underline{\delta} \in (0,1)$ such that when $\delta > \underline{\delta}$, type c_j's equilibrium payoff is no more than the value of the following constrained optimization problem:

$$v_j^\varepsilon \equiv \max_{\gamma^j\in\Delta(\mathbf{A}_1\times\mathbf{A}_2)} \sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma^j(\mathbf{a}_1,\mathbf{a}_2)u_1(c_j,\mathbf{a}_1,\mathbf{a}_2), \tag{3.3}$$

subject to (3.2) and $\mathbf{a}_2$ being an ε-best reply to $\gamma^j(\cdot|\mathbf{a}_2)$ for every $\mathbf{a}_2$ such that $\sum_{\mathbf{a}_1\in\mathbf{A}_1}\gamma^j(\mathbf{a}_1,\mathbf{a}_2) > \varepsilon$.

Let

$$v_j^* \equiv \max_{\gamma^j\in\Delta(\mathbf{A}_1\times\mathbf{A}_2)} \sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma^j(\mathbf{a}_1,\mathbf{a}_2)u_1(c_j,\mathbf{a}_1,\mathbf{a}_2), \tag{3.4}$$

subject to (3.2) and $\mathbf{a}_2$ being a best reply to $\gamma^j(\cdot|\mathbf{a}_2)$ for every $\mathbf{a}_2$ that satisfies $\sum_{\mathbf{a}_1\in\mathbf{A}_1}\gamma^j(\mathbf{a}_1,\mathbf{a}_2) > 0$.

Compared to the constrained optimization problem that defines v_j^*, the one that defines v_j^ε relaxes the best reply constraint by requiring the constraint to hold only for $\mathbf{a}_2$ that occurs with probability greater than ε instead of for every $\mathbf{a}_2$ in the support; and by requiring $\mathbf{a}_2$ be an ε-best reply instead of a best reply. Therefore, $v_j^\varepsilon \geq v_j^*$. The next proposition shows that v_j^ε converges to v_j^* as $\varepsilon \to 0$.

Proposition 3. *For every $j \in \{1,2,...,n\}$, we have $\lim_{\varepsilon\downarrow 0} v_j^\varepsilon = v_j^*$.*

The proof is in Appendix C. Propositions 1, 2, and 3 together imply that for every $j \in \{1,2,...,n\}$, type c_j's equilibrium payoff cannot exceed v_j^* in the limit where $\delta \to 1$.

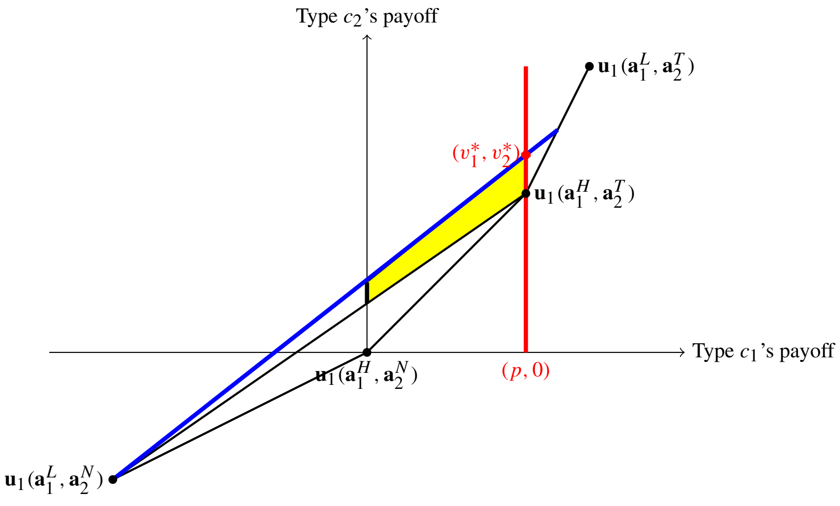

Fig. 1. An example with two types $0 \le c_2 < c_1 < 1$. Let $\mathbf{u}_1(\mathbf{a}_1, \mathbf{a}_2) \equiv \big(u_1(c_1, \mathbf{a}_1, \mathbf{a}_2), u_1(c_2, \mathbf{a}_1, \mathbf{a}_2)\big)$. Constraint (3.2) is depicted in red, and the constraint that $\mathbf{a}_2^T$ best replies to $\gamma^j(\cdot|\mathbf{a}_2^T)$ is depicted in blue. The intersection between the two lines is the sender's highest equilibrium payoff (v_1^*, v_2^*). (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

4. Equilibrium payoff without ethical types

This section examines the case where $0 \le c_n < c_{n-1} < ... < c_1 < 1$. That is, it is common knowledge that the sender is non-ethical and that he has a strict incentive to deceive the receivers when they always follow his action recommendation. I solve for v_j^* defined in (3.4) and obtain:

$$v_j^* \equiv \left\{1 + \frac{c_1 - c_j}{2p + c_1(1-2p)}\right\} p. \tag{4.1}$$

The binding constraints that pin down v_j^* are (i) type c_1's payoff being no more than p and (ii) at histories where the receivers play $\mathbf{a}_2^T$, the discounted frequency with which the sender plays $\mathbf{a}_1^L$ being no more than $\rho^* \equiv \frac{p}{1-p}$. Fig. 1 depicts (v_1^*, v_2^*) in an example with two types $\mathcal{C} = \{c_1, c_2\}$.

Theorem 1. *Suppose $c_1 < 1$. For every $\varepsilon > 0$, there exists $\underline{\delta} \in (0, 1)$ such that when $\delta > \underline{\delta}$, there exists an equilibrium in which the sender's payoff belongs to an ε-neighborhood of $\boldsymbol{v}^* \equiv (v_1^*, ..., v_n^*)$.*

The proof is in Appendix D. Theorem 1 together with Propositions 1, 2, and 3 implies that when the sender's discount factor is close to one, v_j^* is type c_j's *highest equilibrium payoff* and that the highest equilibrium payoffs of all types can be approximately attained in the same equilibrium.

Theorem 1 implies that first, type c_j's highest equilibrium payoff does not depend on the probabilities of each type or the lying costs of the other types. Intuitively, when a sender type decides which of the other types to imitate, he may face a tradeoff between (i) imitating a type that has higher ex ante probability, since it takes fewer periods to convince the receivers that he will behave like that type, and (ii) imitating a type that has a higher lying cost, since higher-cost types are more trustworthy and a reputation for behaving like those types leads to a higher continuation value. As $\delta \to 1$, the sender puts negligible weight on his stage-game payoffs in periods where he builds his reputation and puts almost all weight on his continuation value after he has established a reputation. This suggests that the patient sender prefers to build a reputation for behaving like the highest-cost type, irrespective of the lying costs of the other types or the probability of each type.

Second, every type's highest equilibrium payoff is strictly less than his optimal commitment payoff. This is because the receivers can observe (ω, r, m), which can statistically identify the sender's stage-game strategy. Therefore, the receivers' predictions about the sender's stage-game strategies must be arbitrarily precise in all except for a bounded number of periods. This implies that when the sender is patient, each type's equilibrium payoff cannot exceed his optimal commitment payoff.

Third, $v_j^* > p$ for every $j \geq 2$. The intuition is that, except for the type that has the highest lying cost, every type can strictly benefit from his persistent private information, in the sense that he can obtain a payoff that is strictly greater than his payoff from full disclosure. The argument in Fudenberg et al. (1990) does not apply in my setting since different types of the sender may use different strategies. Hence, it could be the case that at some history h^t, some types have no incentive to play $\mathbf{a}_1^H$ yet the receiver has an incentive to play $\mathbf{a}_2^T$. This will happen if the receiver's belief assigns positive probability to sender-types who will play $\mathbf{a}_1^H$ with positive probability at h^t.

Therefore, in order for a low-cost type to obtain a payoff bounded above p, he needs to behave differently relative to the high-cost type. This inevitably reveals his lying cost to future receivers. As $\delta \to 1$, the low-cost type needs to reveal information about his lying cost in an unbounded number of periods in order to obtain a *discounted average payoff* that is bounded above p. This is somewhat puzzling since revealing private information undermines the sender's informational advantage, and he cannot benefit from his private information after the receivers learn his type. I explain the idea behind my constructive proof in Section 4.1, which sheds light both on how the sender can benefit from his persistent private information in the long run and the extent to which he can do that.

Fourth, as $c_1 \to 1$, v_j^* converges to type c_j's optimal commitment payoff $v_j^{**} \equiv p + p(1 - c_j)$. This provides a microfoundation for the sender's commitment power in Bayesian persuasion models. It implies that even when *all* sender types have strict incentives to deceive the receivers, as long as there exists one type whose lying cost is close to his benefit from the good action, *all* types can approximately attain their optimal commitment payoffs in the repeated communication game.

When c_1 is bounded below 1, v_j^* is bounded below v_j^{**}, and Theorem 1 characterizes the extent to which a patient sender can *partially* restore his commitment power when he communicates with a sequence of myopic receivers. According to (4.1), every type's highest equilibrium payoff is a continuous function of his own lying cost and the highest lying cost in the support of the receivers' prior belief. It is strictly increasing in the lying cost of the highest cost type c_1, strictly decreasing in his true cost of lying, and strictly increasing in the ex ante probability that the state is good.

4.1. Equilibria that approximately attain the highest equilibrium payoff

In order to provide guidance on how to construct equilibria in which the sender's payoff is approximately $\boldsymbol{v}^*$, I derive a common property of the sender's behavior that applies to all of those equilibria. In particular, I show that no type of the sender plays both $\mathbf{a}_1^H$ and $\mathbf{a}_1^L$ with positive probability at every on-path history. This implies that every type of the sender's behavior must depend non-trivially on the game's history and no type uses his static optimal disclosure policy or anything close to that at every on-path history. This stands in contrast to the commitment type in Mathevet et al. (2022), who is *assumed* to use his optimal disclosure policy at every history.

Proposition 4. *Suppose $c_1 \leq 1$. There exists $\overline{\varepsilon} > 0$ such that for every $\varepsilon \in (0, \overline{\varepsilon})$, there exists $\underline{\delta} \in (0,1)$ such that for every $\delta > \underline{\delta}$ and every Nash equilibrium where the sender attains a payoff within ε of $\boldsymbol{v}^*$, no sender type plays both $\mathbf{a}_1^H$ and $\mathbf{a}_1^L$ with positive probability at every on-path history.*

The proof is in Appendix E. This result extends to a type whose lying cost is exactly 1. Even though this type is indifferent between $\mathbf{a}_1^H$ and $\mathbf{a}_1^L$ in the one-shot communication game, he *cannot* mix between $\mathbf{a}_1^H$ and $\mathbf{a}_1^L$ in every period in any sender-optimal equilibrium of the repeated game.

This is because a rational type benefits from the good action regardless of his lying cost, which stands in contrast to a commitment type who does not care about his payoff. Therefore, even when a sender is indifferent between $\mathbf{a}_1^H$ and $\mathbf{a}_1^L$ in the stage game, his indifference between $\mathbf{a}_1^H$ and $\mathbf{a}_1^L$ at a given history of a repeated game introduces constraints on the receiver's strategies after that history. This further introduces constraints on the incentives and payoffs of the other sender types.

For example, let $\mathcal{C} \equiv \{c_1, c_2\}$ with $c_1 = 1$ and $c_2 = 0$. Suppose type c_1 mixes between $\mathbf{a}_1^H$ and $\mathbf{a}_1^L$ at every history. Then playing $\mathbf{a}_1^H$ in every period and playing $\mathbf{a}_1^L$ in every period are both type c_1's best replies to the receivers' equilibrium strategy, from which his payoff is p. As a result, type c_2's payoff from playing $\mathbf{a}_1^L$ in every period is

$$\underbrace{p}_{\text{type } c_1\text{'s payoff from playing } \mathbf{a}_1^L \text{ in every period}} + \underbrace{1}_{\text{difference in the two types' lying costs}} \times \underbrace{(1-p)}_{\text{probability of the bad state}} = 1,$$

which is strictly greater than his optimal commitment payoff $2p$. This contradicts our earlier conclusion that no type of the sender can obtain any payoff that exceeds his optimal commitment payoff.

4.2. Constructing equilibria that attain $\boldsymbol{v}^*$

Following the guidance provided by Proposition 4, I explain how my constructive proof works using an example with two types, i.e., $\mathcal{C} = \{c_1, c_2\}$. I explain how the sender can use non-stationary strategies to benefit from his persistent private information in the long run.

Let $\boldsymbol{v}^N \equiv \big(-c_1(1-p), -c_2(1-p)\big)$ be the sender's payoff from $(\mathbf{a}_1^L, \mathbf{a}_2^N)$, $\boldsymbol{v}^L \equiv \big(p + (1-c_1)(1-p), p + (1-c_2)(1-p)\big)$ be his payoff from $(\mathbf{a}_1^L, \mathbf{a}_2^T)$, and $\boldsymbol{v}^H \equiv (p, p)$ be his payoff

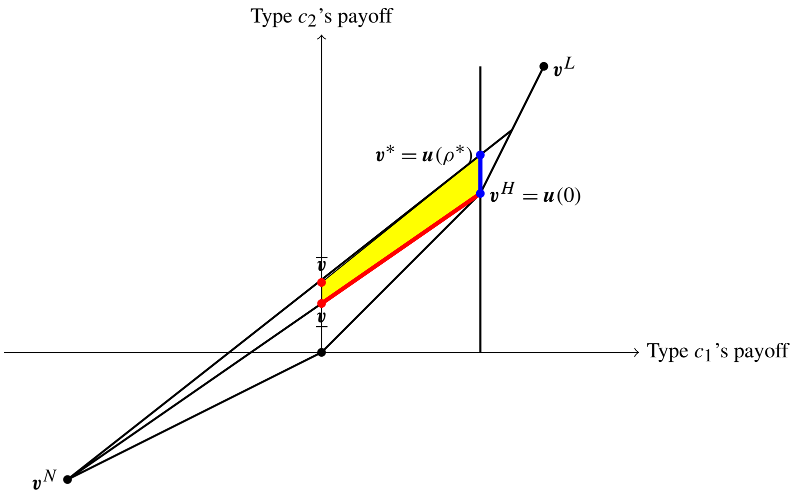

Fig. 2. An example with two types $\mathcal{C} = \{c_1, c_2\}$. The blue line segment is the set of $\boldsymbol{u}(\rho)$ for all $\rho \in [0, \rho^*]$, $\overline{\boldsymbol{v}}$ is the intersection between line segment $[\boldsymbol{v}^*, \boldsymbol{v}^N]$ and the vertical axis, and $\underline{\boldsymbol{v}}$ is the intersection between line segment $[\boldsymbol{v}^H, \boldsymbol{v}^N]$ and the vertical axis.

from $(\mathbf{a}_1^H, \mathbf{a}_2^T)$. Recall that $\rho^* \equiv \frac{p}{1-p}$ as well as the two constraints that pin down $\boldsymbol{v}^*$. I depict the line segment that connects $\boldsymbol{v}^*$ and $\boldsymbol{v}^H$ in Fig. 2. Since each payoff vector on that line belongs to the convex hull of $\{\boldsymbol{v}^H, \boldsymbol{v}^L, \boldsymbol{v}^N\}$ and the first entry of each payoff vector on that line equals p, we have:

$$\boldsymbol{u}(\rho) \equiv \frac{(1-\rho)c_1}{\rho(1-c_1)+c_1}\boldsymbol{v}^H + \frac{\rho c_1}{\rho(1-c_1)+c_1}\boldsymbol{v}^L + \frac{\rho(1-c_1)}{\rho(1-c_1)+c_1}\boldsymbol{v}^N \text{ for } \rho \in [0, \rho^*]. \tag{4.2}$$

By definition, $\boldsymbol{u}(0) = \boldsymbol{v}^H$, $\boldsymbol{u}(\rho^*) = \boldsymbol{v}^*$, and the second entry of $\boldsymbol{u}(\rho)$ is strictly increasing in ρ.

For every $\rho \in [0, \rho^*)$, I construct equilibria such that as $\delta \to 1$, (i) the sender's payoff is approximately $\boldsymbol{u}(\rho)$, (ii) type-c_1 sender's payoff is exactly p, and (iii) the receivers always ignore the sender's cheap talk message. One can use my argument to construct equilibria in which the sender attains any payoff that belongs to the interior of the yellow region of Fig. 2 as $\delta \to 1$.

At every history, *either* the receiver plays $\mathbf{a}_2^T$ and both types of the sender play only $\mathbf{a}_1^H$ and $\mathbf{a}_1^L$ with positive probability, *or* the receiver plays $\mathbf{a}_2^N$ and both types of the sender play $\mathbf{a}_1^L$ for sure. As a result, the sender's stage-game payoff in every period belongs to the convex hull of $\{\boldsymbol{v}^H, \boldsymbol{v}^L, \boldsymbol{v}^N\}$. Since the sender's continuation value in period 0, $\boldsymbol{u}(\rho)$, also belongs in this convex hull, his promised continuation value at every history can be written as a linear combination of $\boldsymbol{v}^H$, $\boldsymbol{v}^L$, and $\boldsymbol{v}^N$.

Let $\eta(h^t) \in [0, 1]$ be the probability of type c_1 at h^t, which I call the sender's *reputation* at h^t. Let $\eta(h^t, (\omega_t, r_t))$ be the probability of type c_1 in period $t+1$ after observing (ω_t, r_t) at h^t.

Active learning phase: The constructed equilibrium starts from an *active learning phase*, in which the receivers play the trusting strategy $\mathbf{a}_2^T$ for sure and both sender types mix between the honest strategy $\mathbf{a}_1^H$ and the lying strategy $\mathbf{a}_1^L$. The low-cost type c_2 plays $\mathbf{a}_1^H$ with lower probability than the high-cost type c_1. This implies that the sender's reputation remains unchanged when the realized state is good, his reputation increases when he recommends the bad action in the bad state, and his reputation decreases when he recommends the good action in the bad state.

In order to make the sender's reputation easy to compute, I construct the sender's mixing probabilities in the active learning phase such that his reputation depends only on *the number of times* that he truthfully reveals the bad state and *the number of times* that he lies about the bad state. The sender's mixing probabilities also need to satisfy the receiver's incentive constraint. Since the receivers play $\mathbf{a}_2^T$, the unconditional probability that the sender plays $\mathbf{a}_1^H$ is no less than $1-\rho^*$.

I show that both objectives can be achieved when type c_1 plays $\mathbf{a}_1^H$ at h^t with probability $\frac{\eta(h^t)-\eta(h^t,(b,g))}{\eta(h^t,(b,b))-\eta(h^t,(b,g))} \cdot \frac{\eta(h^t,(b,b))}{\eta(h^t)}$, and type c_2 plays $\mathbf{a}_1^H$ at h^t with probability $\frac{\eta(h^t)-\eta(h^t,(b,g))}{\eta(h^t,(b,b))-\eta(h^t,(b,g))} \cdot \frac{1-\eta(h^t,(b,b))}{1-\eta(h^t)}$, where $\eta(h^t,(b,g))$ and $\eta(h^t,(b,b))$ can be computed from $\eta(h^t)$ via the following two equations:

$$\eta(h^t,(b,g))-\eta^* = \big(1-\lambda(1-\rho^*)\big)(\eta(h^t)-\eta^*), \tag{4.3}$$

$$\eta(h^t,(b,b))-\eta^* = (1+\lambda\rho^*)(\eta(h^t)-\eta^*), \tag{4.4}$$

for some constant $\eta^*>0$ that is strictly less than the prior probability of type c_1 and some constant $\lambda>0$ that is close to 0. One can verify using Bayes rule that the receivers' posterior beliefs are indeed given by (4.3) and (4.4) when the two types mix according to the probabilities I constructed.

I verify the receiver's incentive to play $\mathbf{a}_2^T$ at h^t. Since the receivers' belief is a martingale, we have $\mathbb{E}[\eta(h^t,(\omega_t,r_t))|h^t]=\eta(h^t)$. Equations (4.3) and (4.4) imply that $\frac{\eta(h^t,(b,b))-\eta(h^t)}{\eta(h^t)-\eta(h^t,(b,g))}=\frac{\rho^*}{1-\rho^*}$. Therefore, the probability that the sender reveals the bad state, i.e., $(\omega_t,r_t)=(b,b)$, divided by the probability that the sender lies about the bad state, i.e., $(\omega_t,r_t)=(b,g)$, is $\frac{1-\rho^*}{\rho^*}$. This implies that the unconditional probability that the sender plays his honest strategy $\mathbf{a}_1^H$ is $1-\rho^*$.

The sender's reputation is easy to compute. Recall that η^* is a constant that is strictly less than the sender's reputation in period 0. According to (4.3) and (4.4), if the sender revealed the bad state K_1 times and lied about the bad state K_2 times before reaching h^t, then his reputation at h^t satisfies:

$$\eta(h^t)-\eta^* = \Big(\eta(h^0)-\eta^*\Big)\cdot\Big(1+\lambda\rho^*\Big)^{K_1}\cdot\Big(1-\lambda(1-\rho^*)\Big)^{K_2}. \tag{4.5}$$

The *active learning phase* ends in two circumstances. First, when the sender revealed the bad state many times before, in the sense that according to (4.5), his reputation will be greater than 1 if he reveals the bad state again. Second, when the sender lied about the bad state many times before, in which case the receivers cannot always trust him as they did in the active learning phase. This is because in order to provide the sender an incentive to reveal the bad state, he must be punished for lying and the only way to punish him is by trusting him with lower frequency in the future.

Perfect reputation phase: When $\eta(h^t)$ is close enough to 1 that $(1+\lambda\rho^*)(\eta(h^t)-\eta^*)\geq 1-\eta^*$, i.e., the sender's reputation in the next period will exceed 1 after he reveals the bad state again, then the receiver plays $\mathbf{a}_2^T$, the low-cost type c_2 plays $\mathbf{a}_1^L$ for sure, and the high-cost type c_1 mixes between $\mathbf{a}_1^L$ and $\mathbf{a}_1^H$. Hence, the receiver's posterior belief after the sender reveals the bad state satisfies:

$$\eta(h^t,(b,b))-\eta^* = 1-\eta^* \leq (1+\lambda\rho^*)(\eta(h^t)-\eta^*). \tag{4.6}$$

I choose the probability with which type c_1 plays his lying strategy $\mathbf{a}_1^L$ so that after the receiver observes the sender lying about the bad state at h^t, her belief in period $t+1$ satisfies (4.3).

Hence, the sender's reputation remains unchanged if the realized state is good, his reputation reaches 1 after he truthfully reveals the bad state, and his reputation decreases after he lies about the bad state. Since belief is a martingale, we have $\mathbb{E}[\eta(h^t,(\omega_t,r_t))|h^t]=\eta(h^t)$. According to (4.3) and (4.6), $\frac{\eta(h^t,(b,b))-\eta(h^t)}{\eta(h^t)-\eta(h^t,(b,g))}\leq\frac{\rho^*}{1-\rho^*}$. Hence, the probability that the sender reveals the bad state, i.e., $(\omega_t,r_t)=(b,b)$, divided by the probability with which the sender lies about the bad state, i.e., $(\omega_t,r_t)=(b,g)$, is no less than $\frac{1-\rho^*}{\rho^*}$. This implies that the probability that the sender plays $\mathbf{a}_1^H$ at this history is no less than $1-\rho^*$. Thus, the receiver has an incentive to play $\mathbf{a}_2^T$ at h^t.

If the sender lies about the bad state at such a history, his reputation can be computed via (4.5), which is the same formula as in the active learning phase. If the sender truthfully reveals the bad state at such a history, his reputation reaches 1 and the continuation play consists only of two stage-game strategy profiles: $(\mathbf{a}_1^H,\mathbf{a}_2^T)$ and $(\mathbf{a}_1^L,\mathbf{a}_2^N)$. Under such a continuation play, the receivers' myopic incentive constraints are satisfied, and the sender's continuation value after his reputation reaches 1 belongs to the convex hull of $\{\boldsymbol{v}^N,\boldsymbol{v}^H\}$, which is depicted as the red line in Fig. 2.

Since the sender's continuation value in the active learning phase belongs to the convex hull of $\{\boldsymbol{v}^N,\boldsymbol{v}^L,\boldsymbol{v}^H\}$, there exists a continuation value that belongs to the convex hull of $\{\boldsymbol{v}^N,\boldsymbol{v}^H\}$ so that (i) type c_1 is indifferent between lying about the bad state and truthfully revealing the bad state, and (ii) type c_2 prefers to lie about the bad state. This is because when type c_1 is indifferent between a payoff that belongs to the convex hull of $\{\boldsymbol{v}^N,\boldsymbol{v}^H\}$ and a payoff that belongs to the convex hull of $\{\boldsymbol{v}^N,\boldsymbol{v}^H,\boldsymbol{v}^L\}$, type c_2 prefers the latter since it is higher in the vertical dimension.

Low value phase: The active learning phase also ends if the sender lied about the bad state too many times in the past. In order to determine when to stop the active learning phase, we need to keep track of the sender's continuation value $\boldsymbol{v}(h^t)$. Since the sender's equilibrium payoff belongs to the convex hull of $\{\boldsymbol{v}^N,\boldsymbol{v}^H,\boldsymbol{v}^L\}$ and the sender's stage-game payoff in each period is either $\boldsymbol{v}^N$, $\boldsymbol{v}^H$, or $\boldsymbol{v}^L$, his promised continuation value at h^t, denoted by $\boldsymbol{v}(h^t)$, can be written as a linear combination of $\boldsymbol{v}^H$, $\boldsymbol{v}^L$, and $\boldsymbol{v}^N$. Hence, there exist real numbers $p^H(h^t)$, $p^L(h^t)$, and $p^N(h^t)$ which sum up to one such that:

$$\boldsymbol{v}(h^t)=p^H(h^t)\boldsymbol{v}^H+p^L(h^t)\boldsymbol{v}^L+p^N(h^t)\boldsymbol{v}^N.$$

Therefore, keeping track of the sender's continuation value is equivalent to keeping track of the weights $p^H(h^t)$, $p^L(h^t)$, and $p^N(h^t)$. I set $\boldsymbol{v}(h^0)\equiv\boldsymbol{u}(\rho)$ since $\boldsymbol{u}(\rho)$ is the sender's equilibrium payoff. Equation (4.2) implies that $p^H(h^0)=\frac{(1-\rho)c_1}{\rho(1-c_1)+c_1}$, $p^L(h^0)=\frac{\rho c_1}{\rho(1-c_1)+c_1}$, and $p^N(h^0)=\frac{\rho(1-c_1)}{\rho(1-c_1)+c_1}$.

In the *active learning phase*, the sender's continuation value in period $t+1$ depends on his continuation value in period t and the realization of (ω_t,r_t). I list some requirements that the sender's continuation value in period $t+1$ needs to satisfy. First, since both types mix between $\mathbf{a}_1^H$ and $\mathbf{a}_1^L$, they must be indifferent between revealing and lying about the bad state. Second, when the realized state is good, the sender's continuation value remains unchanged given that his reputation remains unchanged. Third, the sender's stage-game payoff in period t and his continuation value in period $t+1$ must deliver him payoff $\boldsymbol{v}(h^t)$ in period t. These three requirements pin down the sender's continuation value in period $t+1$, which is given by $\boldsymbol{v}(h^t,(g,g))=\boldsymbol{v}(h^t)$,

$$\boldsymbol{v}(h^t,(b,g)) = \frac{p^H(h^t)}{\widehat{\delta}}\boldsymbol{v}^H + \frac{p^L(h^t)-(1-\widehat{\delta})}{\widehat{\delta}}\boldsymbol{v}^L + \frac{p^N(h^t)}{\widehat{\delta}}\boldsymbol{v}^N,$$

$$\boldsymbol{v}(h^t,(b,b)) = \frac{p^H(h^t)-(1-\widehat{\delta})}{\widehat{\delta}}\boldsymbol{v}^H + \frac{p^L(h^t)}{\widehat{\delta}}\boldsymbol{v}^L + \frac{p^N(h^t)}{\widehat{\delta}}\boldsymbol{v}^N,$$

where $\widehat{\delta} \equiv \delta\frac{1-p}{1-\delta p}$. Intuitively, when the sender's lying cost is common knowledge, his equilibrium payoff belongs to the convex hull of $\{\boldsymbol{v}^H, \boldsymbol{v}^N\}$, i.e., the red line segment in Fig. 2. When the sender has private information about his lying cost, he can attain payoffs strictly above the red line, i.e., payoffs where the weight of $\boldsymbol{v}^L$, denoted by $p^L(h^t)$, is bounded away from zero. Each time the sender lies about the bad state, the weight of $\boldsymbol{v}^L$ in his continuation value decreases, and he should not be allowed to lie while obtaining the receiver's trust after the weight on $\boldsymbol{v}^L$ is close to or equal to 0.

Formally, the active learning phase ends at history h^t as long as $p^L(h^t) \leq 1-\widehat{\delta}$. After $p^L(h^t)$ reaches 0,[6] in every period of the continuation game, either the receiver plays $\mathbf{a}_2^T$ and all types of the sender play $\mathbf{a}_1^H$, or the receiver plays $\mathbf{a}_2^N$ and all types of the sender play $\mathbf{a}_1^L$. Whether $(\mathbf{a}_1^H, \mathbf{a}_2^T)$ or $(\mathbf{a}_1^L, \mathbf{a}_2^N)$ will be played in each period depends on the sender's continuation value. In particular, the sender's continuation value in the beginning of this phase is a convex combination of $\boldsymbol{v}^H$ and $\boldsymbol{v}^N$. Since his stage-game payoff is $\boldsymbol{v}^H$ when $(\mathbf{a}_1^H, \mathbf{a}_2^T)$ is played and is $\boldsymbol{v}^N$ when $(\mathbf{a}_1^L, \mathbf{a}_2^N)$ is played, one can determine the play in each period using the algorithm in Fudenberg and Maskin (1991). The idea is that players play $(\mathbf{a}_1^H, \mathbf{a}_2^T)$ if the continuation value is too low and play $(\mathbf{a}_1^L, \mathbf{a}_2^N)$ if the continuation value is too high, so that the sender's continuation value is always close to his continuation value when play first enters the low-value phase.

Verifying incentive constraints: I have verified the receiver's incentive constraints when constructing the sender's mixing probabilities. The construction of the sender's continuation value ensures that his incentive constraints in *recursive form* are satisfied at every history.

What remains to be verified is the *promise-keeping constraint*, that type c_j's discounted average payoff equals the jth entry of $\boldsymbol{u}(\rho)$ when players play according to the constructed strategies. This is implied by the sender's incentive constraints in *sequential form*. In order to verify the sender's incentive constraints in sequential form, I only need to show that the continuation value promised to the sender belongs to the convex hull of $\{\boldsymbol{v}^N, \boldsymbol{v}^H, \boldsymbol{v}^L\}$. When the promised continuation values are bounded, Theorem 9.2 in Stokey et al. (1989) shows that the sender's incentive constraints in recursive form is sufficient for his incentive constraints in sequential form.

Intuitively, in the active learning phase, the sender's promised continuation value increases every time he truthfully reveals the bad state. Two arrangements in my construction prevent the sender's promised continuation value from exploding. First, the sender's continuation value belongs to the convex hull of $\{\boldsymbol{v}^N, \boldsymbol{v}^H\}$ after his reputation reaches 1. Second, the sender's continuation value belongs to the convex hull of $\{\boldsymbol{v}^N, \boldsymbol{v}^H\}$ after he lied about the bad state too frequently in the past. Therefore, in order to make sure that the sender's continuation value is bounded, his reputation needs to reach 1 before his promised continuation value escapes the

[6] In the main text, I only describe players' behaviors at histories that satisfy either $p^L(h^t) \geq 1-\widehat{\delta}$ or $p^L(h^t) = 0$. I refer the technical details to Appendix D which discusses players' strategies at histories where $0 < p^L(h^t) < 1-\widehat{\delta}$, i.e., the remaining convex weight on $\boldsymbol{v}^L$ is strictly positive but is strictly less than the weight of the current period $1-\widehat{\delta}$. Appendix D also discusses technical details when the sender's continuation value reaches the vertical axis, that is, type c_1's individual rationality constraint will be violated if the receiver trusts the sender in the current period.

convex hull of $\{\boldsymbol{v}^N, \boldsymbol{v}^H, \boldsymbol{v}^L\}$. How fast can the sender build his reputation and how fast can his promised continuation value increase?

1. In order to provide the receivers an incentive to trust the sender in the active learning phase, the sender needs to play $\mathbf{a}_1^H$ with probability at least $1-\rho^*$. This implies that the ratio between the magnitude of reputation increase after the sender truthfully reveals the bad state and the magnitude of reputation decrease after the sender lies about the bad state must be no more than $\frac{\rho^*}{1-\rho^*}$. For example, when the sender's reputation in the active learning phase satisfies (4.3) and (4.4) with $\lambda > 0$ small enough, the sender's reputation increases over time when he lies about the bad state with frequency *less than* ρ^*.
2. Suppose the sender's continuation value in period 0 satisfies $\frac{p^L(h^0)}{p^H(h^0)} < \frac{\rho^*}{1-\rho^*}$. As long as the ratio between the frequency of outcome $(\omega_t, r_t) = (b, g)$ and the frequency of outcome $(\omega_t, r_t) = (b, b)$ is more than $\frac{\rho^*}{1-\rho^*}$, the weight of $\boldsymbol{v}^L$ divided by the weight of $\boldsymbol{v}^H$ in the sender's promised continuation value must be no more than $\frac{\rho^*}{1-\rho^*}$, as implied by the sender's incentive constraints in recursive form. Hence, the sender's continuation value decreases over time if he lies about the bad state with frequency *more than* ρ^*.

The above argument also provides a heuristic explanation for why one cannot use my techniques to construct equilibria in which the sender's equilibrium payoff satisfies $\frac{p^L(h^0)}{p^H(h^0)} > \frac{\rho^*}{1-\rho^*}$. This is because when the sender lies about the bad state with frequency slightly above ρ^*, his reputation declines over time but the convex weight of $\boldsymbol{v}^L$ in his continuation value increases. If this is the case, the sender's continuation value explodes, which cannot be delivered by any equilibrium in the continuation game.

Long-run dynamics: The receivers may not learn the sender's type in the long run. This is because when play enters the low value phase, the sender's reputation is strictly between η^* and 1. After that, both types of the sender will use the same strategy in the continuation game. This does not contradict the conclusion of Cripps et al. (2004) since their impermanent reputation result requires *full support monitoring*, which is violated in my model.

The sender's continuation value will belong to the convex hull of $\{\boldsymbol{v}^H, \boldsymbol{v}^N\}$ in the long run, regardless of the frequency with which he lies about the bad state. This is because when he lies with frequency less than ρ^*, his reputation will reach 1, and when he lies with frequency more than ρ^*, the weight of $\boldsymbol{v}^L$ in his continuation value will reach 0. This, together with the fact that the sender's payoff is close to $\boldsymbol{v}^*$, implies that (i) the active learning phase will terminate almost surely in finite time, (ii) only outcomes $(\mathbf{a}_1^H, \mathbf{a}_2^T)$ and $(\mathbf{a}_1^L, \mathbf{a}_2^N)$ occur in the long run and in expectation, outcome $(\mathbf{a}_1^L, \mathbf{a}_2^T)$ will disappear in finite time, and (iii) the length of the active learning phase goes to infinity as $\delta \to 1$, so that the sender can obtain payoff strictly greater than p from outcome $(\mathbf{a}_1^L, \mathbf{a}_2^T)$.

5. Equilibrium payoffs with ethical types

This section allows some types to be ethical. Let $k \in \{1, 2, ..., n\}$ be such that $c_i < 1$ if and only if $i \geq k$. By definition, types $c_k, ..., c_n$ are called *non-ethical types* and types $c_1, ..., c_{k-1}$ are called *ethical types*. Section 5.1 shows that when the sender can send cheap talk messages in addition to recommending an action, there exists an equilibrium in which all sender types approximately attain their optimal commitment payoffs when the set of types is rich enough. Section 5.2 shows that when the sender cannot send cheap talk messages, there exists an equilibrium where the non-ethical types attain their optimal commitment payoffs *if and only if* the ethical types' lying costs are not too high.

5.1. Attaining commitment payoff with cheap talk communication

Let

$$v_j^\dagger \equiv \left\{1 + \frac{c_k - c_j}{2p + c_k(1-2p)}\right\} p \quad \text{for every } j \in \{k, ..., n\}. \tag{5.1}$$

Compared to v_j^* defined in (4.1), the formula for $v_j^\dagger$ replaces the highest lying cost c_1 with the highest lying cost among the non-ethical types c_k. The properties of $v_j^\dagger$ are similar to those of v_j^*. First, $v_k^\dagger = p$. Second, $v_j^\dagger > p$ for every $j > k$. Third, $v_j^\dagger$ converges to v_j^{**} as c_k converges to 1.

Theorem 2. *Suppose $|M| \geq 2$. For every $\varepsilon > 0$, there exists $\underline{\delta} \in (0, 1)$ such that for every $\delta > \underline{\delta}$, there exists an equilibrium in which the non-ethical types' payoffs are within an ε-neighborhood of $(v_k^\dagger, ..., v_n^\dagger)$ and the ethical types attain their optimal commitment payoff p.*

Theorem 2 implies that when the set of types is rich enough in the sense that for every $\varepsilon > 0$, there exists a type whose lying cost belongs to the interval $[1-\varepsilon, 1]$, then *all* types of the sender can approximately attain their optimal commitment payoffs as their discount factor approaches 1.

The constructive proof of Theorem 2 follows from the proof of Theorem 1. Consider an equilibrium such that in period 0, the receiver follows the sender's recommendation, and all types of the sender recommend the receiver-optimal action and use the cheap talk message to truthfully communicate whether he is ethical. The receivers ignore the sender's cheap talk messages after period 0.

If the sender reports that he is ethical in period 0, then starting from period 1, players coordinate on a continuation equilibrium in which the sender recommends the receiver's optimal action and the receivers follow the sender's recommendation on the equilibrium path. If the sender's recommendation does not match the realized state, then players coordinate on the equilibrium in which the sender's continuation value is close to $\underline{\boldsymbol{v}}$. A patient sender has no incentive to lie since his equilibrium payoff p is bounded above his continuation value after he deviates.

If the sender reports that he is non-ethical in period 0, then starting from period 1, players coordinate on the equilibrium constructed in Theorem 1 with type c_k being the highest-cost type. Type c_k's payoff is p, which implies that he is indifferent between reporting that he is ethical and reporting that he is non-ethical in period 0. Non-ethical types with lying cost strictly lower than c_k receives a continuation value that is strictly greater than p, which implies that they strictly prefer to report that they are non-ethical. The ethical types have no incentive to report that they are non-ethical since their equilibrium payoff p equals their optimal commitment payoff.

5.2. Attaining commitment payoff without cheap talk messages

This section examines the case where $|M| = 1$, which is equivalent to a model where the sender can only make an action recommendation in each period but cannot send cheap talk messages.

Unlike the case with only non-ethical types, there is no equilibrium in which all types attain their highest equilibrium payoffs when ethical and non-ethical types coexist. Intuitively, the ethical types can attain their highest equilibrium payoff if and only if they always recommend the receiver-optimal action. In every such equilibrium, the non-ethical types cannot lie while pooling with the ethical types, in which case the non-ethical types' payoffs cannot exceed p.

My subsequent analysis focuses on the payoffs of the *non-ethical types*. The motivation for this is twofold. First, the incentives of these types are similar to the sender's incentive in the leading example of Kamenica and Gentzkow (2011): They have a conflict of interest with the receiver and strictly prefer one action regardless of the state, even taking their costs of lying into account. Second, the non-ethical types cannot attain their optimal commitment payoffs in the repeated communication game without any private information about their lying costs. Hence, it is important to provide a foundation for the commitment assumption in Bayesian persuasion models for these types.

My next result provides a necessary and sufficient condition under which *all* non-ethical types can attain their optimal commitment payoffs. I also show that when my condition is violated, no non-ethical type can attain his optimal commitment payoff in any Nash equilibrium.

I start from listing these conditions and explain intuitively why they are necessary. First, if any non-ethical type can attain his optimal commitment payoff, then there *exists* an ethical type $c \in \mathcal{C} \cap [1, +\infty)$ whose payoff is no more than the value of the following constrained optimization problem:

$$\overline{v}(c) \equiv \max_{\gamma \in \Delta(\mathbf{A}_1 \times \mathbf{A}_2)} \sum_{(\mathbf{a}_1, \mathbf{a}_2) \in \mathbf{A}_1 \times \mathbf{A}_2} \gamma(\mathbf{a}_1, \mathbf{a}_2) u_1(c, \mathbf{a}_1, \mathbf{a}_2), \tag{5.2}$$

subject to the constraint that there exists $\widehat{c} \in \mathcal{C} \cap [0, 1)$ such that:

$$\sum_{(\mathbf{a}_1, \mathbf{a}_2) \in \mathbf{A}_1 \times \mathbf{A}_2} \gamma(\mathbf{a}_1, \mathbf{a}_2) u_1(\widehat{c}, \mathbf{a}_1, \mathbf{a}_2) \geq p + p(1 - \widehat{c}). \tag{5.3}$$

Intuitively, Theorem 1 implies hat in a repeated game *without* any ethical type, every non-ethical type's equilibrium payoff must be bounded below his optimal commitment payoff. In order to attain his optimal commitment payoff, every non-ethical type needs to imitate some ethical type, so he can attain his optimal commitment payoff when he plays some ethical type's best reply.

Second, for *every* ethical type $c \in \mathcal{C} \cap [1, +\infty)$, his equilibrium payoff is no less than $\underline{v}(c)$, which is defined as his payoff from imitating the type who has the highest lying cost c_1:

$$\underline{v}(c) \equiv \min_{\gamma \in \Delta(\mathbf{A}_1 \times \mathbf{A}_2)} \sum_{(\mathbf{a}_1, \mathbf{a}_2) \in \mathbf{A}_1 \times \mathbf{A}_2} \gamma(\mathbf{a}_1, \mathbf{a}_2) u_1(c, \mathbf{a}_1, \mathbf{a}_2), \tag{5.4}$$

subject to

$$\sum_{(\mathbf{a}_1, \mathbf{a}_2) \in \mathbf{A}_1 \times \mathbf{A}_2} \gamma(\mathbf{a}_1, \mathbf{a}_2) u_1(c_1, \mathbf{a}_1, \mathbf{a}_2) \geq 0, \tag{5.5}$$

$$\mathbf{a}_2 \in \arg\max_{\mathbf{a}_2' \in \mathbf{A}_2} u_2(\gamma(\cdot|\mathbf{a}_2), \mathbf{a}_2') \text{ for every } \mathbf{a}_2 \in \mathbf{A}_2 \text{ in the support of } \gamma. \tag{5.6}$$

The first constraint is necessary since the ethical type can secure payoff 0 by recommending the receiver-optimal action, regardless of the receiver's response. The second constraint comes from Propositions 2 and 3, that under type c_1's equilibrium strategy, the receiver's stage-game strategy must be a best reply to the conditional distribution over the sender's stage-game strategy.

Therefore, in order for some non-ethical type to attain his optimal commitment payoff in some equilibrium, there must exist an ethical type c such that his payoff is no more than $\overline{v}(c)$ and is no less than $\underline{v}(c)$. Hence, it is necessary that $\overline{v}(c) \geq \underline{v}(c)$ for some $c \in \mathcal{C} \cap [1, +\infty)$. I solve the two constrained optimization problems and obtain that $\overline{v}(c) \geq \underline{v}(c)$ if and only if $c_1(c-1) \leq 2$. Let $c^* \equiv \min\left\{\mathcal{C} \cap [1, +\infty)\right\}$, which is the lowest lying cost among the ethical types. Then $c_1(c-1) \leq 2$ is true for at least one ethical type if and only if $c_1(c^*-1) \leq 2$. The next result shows that $c_1(c^*-1) \leq 2$ is both necessary and sufficient for (i) at least one non-ethical type can attain his optimal commitment payoff and (ii) all non-ethical types can attain their optimal commitment payoffs.

Theorem 3. *Suppose $\mathcal{C}$ contains at least one ethical type and at least one non-ethical type.*

1. *If $c_1(c^*-1) > 2$, then there exist $\eta > 0$ and $\underline{\delta} \in (0,1)$ such that if $\delta > \underline{\delta}$, then in every equilibrium, each non-ethical type $c_k \in \mathcal{C} \cap [0,1)$ obtains a payoff of no more than $v_k^{**} - \eta$.*
2. *If $c_1(c^*-1) \leq 2$, then for every $\varepsilon > 0$, there exists $\underline{\delta} \in (0,1)$ such that when $\delta > \underline{\delta}$, there exists an equilibrium in which every non-ethical type $c_k \in \mathcal{C} \cap [0,1)$ obtains a payoff of at least $v_k^{**} - \varepsilon$.*

Theorem 3 implies that the non-ethical types can attain their optimal commitment payoffs *if and only if* the highest lying cost c_1 is no more than $\frac{2}{c^*-1}$. In the case where there is only one ethical type, i.e., $c_1 \geq 1$ and $c_j < 1$ for every $j \geq 2$, the non-ethical types $c_2, ..., c_n$ can attain their optimal commitment payoffs if and only if $c_1 = c^* \in [1,2]$ (Fig. 3).

This leads to a somewhat interesting observation that the highest equilibrium payoffs of the non-ethical types is *non-monotone* with respect to the receiver's belief about the sender's lying cost. Moreover, introducing another ethical type who has a high lying cost can *lower* the non-ethical types' equilibrium payoffs. For example, suppose there are two types with $c_2 \in (0,1)$ and $c_1 \in (1,2)$. Then Theorem 3 implies that a type-c_2 sender can attain his optimal commitment payoff v_2^{**} in some equilibria. Next, let us introduce another type $\overline{c}_1$ with $\overline{c}_1 > c_1$ and $\overline{c}_1(c_1 - 1) > 2$. According to Theorem 3, the payoff for type c_2 is bounded below v_2^{**} in all equilibria.

The above observation is driven by an *outside option effect* that hinges on the ethical types' incentive constraints. Intuitively, when type $\overline{c}_1$ exists, type c_1 has an *outside option* of deviating to the equilibrium strategy of type $\overline{c}_1$. Since type $\overline{c}_1$ receives at least his minmax payoff 0 from his equilibrium strategy, type c_1's outside option leads to a lower bound on his equilibrium payoff. Since type c_1 is ethical, his payoff from sending message g when $\omega_t = b$ is strictly less than his minmax payoff. This implies that his outside option leads to an upper bound on the equilibrium frequency with which he can lie about the bad state. An increase in $\overline{c}_1$ increases the outside option for type c_1 sender, which reduces the frequency with which he is willing to lie about the bad state in equilibrium. This reduces the equilibrium payoff for type c_2, since it lowers the frequency with which he can lie about the bad state and induces the receiver to choose the good action while pooling with type c_1.

The proof is in a working paper version (Pei, 2022). I explained the proof of the first part when I introduce the linear programs. For the rest of this section, I explain the constructive proof of the second part using an example with two types where type c_2 is non-ethical and type c_1 is ethical.

I construct an equilibrium in which the sender's payoff is $\boldsymbol{u}(\rho) \equiv \rho \boldsymbol{v}^L + (1-\rho)\boldsymbol{v}^H$ for some $\rho \in (0, \rho^*)$. Similar to the proof of Theorem 1, the equilibrium starts from an *active learning phase*. Learning will stop either when the sender's reputation, defined as the probability of the

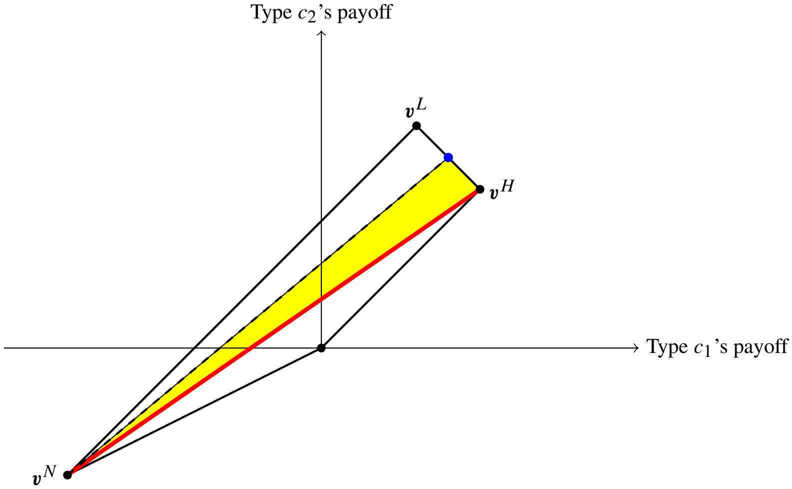

Fig. 3. An example where $0 < c_2 < 1 < c_1$ and $c_1(c_1 - 1) \leq 2$. The goal is to attain the payoff represented by the blue dot in an equilibrium where the sender's stage-game payoff at every history is either $\boldsymbol{v}^L$, $\boldsymbol{v}^H$, or $\boldsymbol{v}^N$.

ethical type, reaches 1, or when his continuation value reaches the red line, i.e., the convex hull of $\{\boldsymbol{v}^N, \boldsymbol{v}^H\}$.

In the active learning phase, the receivers play $\mathbf{a}_2^T$ and both types of the sender mix between $\mathbf{a}_1^H$ and $\mathbf{a}_1^L$. The probability with which type θ_1 plays $\mathbf{a}_1^H$ is *slightly higher* compared to type θ_2. As a result, the sender's reputation remains unchanged when the realized state is good, increases when he reveals the bad state, and decreases when he lies about the bad state. The sender's continuation value is $\boldsymbol{u}(\rho)$ in period 0, does not change if the realized state is good, moves toward $\boldsymbol{v}^H$ if he lies about the bad state, and moves toward $\boldsymbol{v}^L$ if he reveals the bad state.

The active learning phase ends when the sender's continuation value reaches $\boldsymbol{v}^H$, which is the case when he lied about the bad state too many times before. When the sender's continuation value reaches $\boldsymbol{v}^H$, both types of the sender play $\mathbf{a}_1^H$ and the receiver plays $\mathbf{a}_2^T$.

The active learning phase also ends when the sender's reputation reaches 1, after which the sender's continuation value is the projection of his current continuation value to the red line. This ensures that type θ_1 is indifferent between *acquiring a perfect reputation by revealing the bad state* and *remaining in the active learning phase by lying about the bad state*, while type θ_2 strictly prefers to lie about the bad state. After the sender's reputation reaches one, the continuation play in every period is either $(\mathbf{a}_1^H, \mathbf{a}_2^T)$ or $(\mathbf{a}_1^L, \mathbf{a}_2^N)$. The sequence of $(\mathbf{a}_1^H, \mathbf{a}_2^T)$ and $(\mathbf{a}_1^L, \mathbf{a}_2^N)$ is constructed using the algorithm in Fudenberg and Maskin (1991) in order to deliver the promised continuation value when the sender enters the perfect reputation phase.

One thing to note is that in the active learning phase, (i) the non-ethical type is rewarded when he truthfully reveals the bad state and is punished when he lies about the bad state, and (ii) the ethical type is penalized for truthfully revealing the bad state and is rewarded for lying about the bad state. This is inevitable in every equilibrium where the non-ethical type attains his optimal commitment payoff. The intuition is that the ethical type prefers to reveal the bad state, but the non-ethical type can obtain payoff strictly greater than p only when he pools with the ethical type while the ethical type lies. Hence, the ethical types need to lie at some histories, and punishments after revealing the bad state is required in order to provide them an incentive to lie.

6. Conclusion & discussions

This paper provides a justification for the commitment assumption in Bayesian persuasion models using a repeated communication game where it is common knowledge that the sender cannot commit to disclosure policies. When the sender has persistent private information about his cost of lying, I show that a sufficiently patient sender can approximately attain his optimal commitment payoff as long as the receivers' prior belief assigns positive probability to a sender-type whose lying cost is close to the sender's private benefit from the good action.

This paper is related to Pei (2021) who studies a repeated incomplete information game between a patient seller and a sequence of myopic consumers where the seller has persistent private information about his cost of effort. There are several differences between the current paper and Pei (2021), which introduce new challenges to characterize the patient sender's equilibrium payoff.

First, the sender privately observes an i.i.d. state in the current model. There is *imperfect monitoring* since the receiver cannot observe the sender's behavior in the good state if the realized state is bad and vice versa. This introduces new challenges in constructing continuation values since they depend not only on the sender's stage-game strategies but also on the stochastic state. Second, in Pei (2021)'s model, there are two pure action profiles that are incentive compatible for the consumers, which are (high effort, trust) and (low effort, no trust), both of which lead to a payoff that is strictly individually rational for the seller. By contrast, one of the incentive compatible pure action profiles in the current model $(\mathbf{a}_1^L, \mathbf{a}_2^N)$ is not individually rational for the sender. As a result, I need to include a *rebounding phase* in my constructive proof, which ensures that the sender's continuation value will return to something positive once his promised continuation value approaches the vertical axis. Third, the lack-of monotone-supermodularity in players' stage-game payoffs introduces new challenges to show Proposition 4. For example, it is not true that when a high-cost type finds it optimal to play $\mathbf{a}_1^L$ in every period, then a low-cost type will play $\mathbf{a}_1^L$ with probability one at every on-path history. This is because there are other strategies of the patient player, such as a strategy that lies in every state, that incur higher costs of lying. Fourth, the current paper examines the effects of ethical types on the non-ethical types' equilibrium payoffs and shows that the non-ethical types' highest equilibrium payoffs are not monotone with respect to the receiver's belief about his lying cost. I identify a novel *outside option effect*, which is absent in Pei (2021).

Declaration of competing interest

I disclose that my research is funded by NSF Grant SES-1947021. I have no other conflict of interest.

Data availability

No data was used for the research described in the article.

Appendix A. Proof of Proposition 1

Fix any Nash equilibrium $\sigma \equiv \big((\sigma_c)_{c\in\mathcal{C}}, \sigma_2\big)$. Let $\mathcal{H}^\sigma$ be the set of histories that occur with positive probability under σ. Let $\mathcal{C}(h^t)$ be the support of the receiver's posterior belief after observing history h^t but before observing (r_t, m_t). I show Proposition 1 using an induction argument on $|\mathcal{C}(h^t)|$.

Let $\sigma_c(h^t)(\omega) \in \Delta(R \times M)$ be the equilibrium distribution over type c's recommendations and messages at history h^t when the realized state is ω. Since $p < 1/2$, there exists (r, m) that is sent with positive probability by at least one type such that the receiver has a strict incentive to play b after receiving $(r.m)$. For every $c \in \mathcal{C}$, I define $\widetilde{\sigma}_c$ based on type c's equilibrium strategy σ_c:

$$\widetilde{\sigma}_c(h^s)(\omega) \equiv \begin{cases} (r, m) & \text{if } \sigma_c(h^s)(\omega) \text{ assigns positive probability to } (r, m) \\ & \text{ and } \sigma_2(h^s)(r, m) = b \\ \sigma_c(h^s)(\omega) & \text{otherwise.} \end{cases} \tag{A.1}$$

By definition, $\widetilde{\sigma}_c$ is type c's best reply to σ_2, so his payoff from $(\widetilde{\sigma}_c, \sigma_2)$ equals his equilibrium payoff.

Step 1: Suppose $|\mathcal{C}(h^t)| = 1$. Then $\mathcal{C}(h^s) = \mathcal{C}(h^t)$ for every $h^s \in \mathcal{H}^\sigma$ that satisfies $h^s \succeq h^t$. Let c be the only type in $\mathcal{C}(h^t)$. If type c plays $\widetilde{\sigma}_c$ and the receivers play σ_2, then by construction, type c's stage-game payoff at every on-path history after h^t cannot exceed p. This implies that his continuation value at h^t cannot exceed p.

Step 2: I show that for every $k \in \{1, 2, ..., n\}$, if the conclusion holds for every history h^τ that satisfies $|\mathcal{C}(h^\tau)| \le k$, then it also holds for every h^t that satisfies $|\mathcal{C}(h^t)| = k + \xi$, where $\xi \in \{1, ..., n\}$ is the smallest positive integer such that there exists an on-path history h^t where $|\mathcal{C}(h^t)| = k + \xi$. Since there are n types, there exists a unique ξ, which is at most $n - 1$ and is at least 1.

Let c_i be the highest lying cost in $\mathcal{C}(h^t)$. Suppose by way of contradiction that type c_i's continuation value at h^t, denoted by v, is strictly greater than p. Since $\widetilde{\sigma}_{c_i}$ is type c_i's best reply, his discounted average payoff from $(\widetilde{\sigma}_{c_i}, \sigma_2)$ is v at h^t. Let $\mathcal{H}^{(\widetilde{\sigma}_{c_i}, \sigma_2)}$ be the set of histories that occur with positive probability under $(\widetilde{\sigma}_{c_i}, \sigma_2)$. Type c_i's stage-game payoff from $(\widetilde{\sigma}_{c_i}, \sigma_2)$ is no more than p until play reaches history $h^s (\succeq h^t)$ where there exists $(r^*, m^*) \in R \times M$ that satisfies:

- under strategy σ_2, the receiver plays g with positive probability after receiving (r^*, m^*) at h^s,
- $\widetilde{\sigma}_{c_i}(h^s)$ sends (r^*, m^*) for sure at history h^s.

I call such histories *special histories*. Since $p < 1/2$, there exists $(r, m) \neq (r^*, m^*)$ such that $(r.m)$ is sent with positive probability at $(h^s, \omega = b)$ and the receiver strictly prefers b upon receiving (r, m). Since $\widetilde{\sigma}_{c_i}(h^s)(\omega = b) = (r^*, m^*)$, type c_i sends (r, m) with zero probability at $(h^s, \omega = b)$ under his equilibrium strategy σ_{c_i}. Hence, $|\mathcal{C}(h^s, r, m, \omega = b)| \le k$. Let c_j denote the highest lying cost in $\mathcal{C}(h^s, r, m, b)$. By definition, $c_j < c_i$ and c_j sends (r, m) with positive probability at h^s when the state is b. The induction hypothesis implies that type c_j's continuation value after observing b at h^s is no more than δp. Since $c_i > c_j$, type c_i's payoff from any strategy is weakly lower than type c_j's payoff from the same strategy. This implies that type c_i's continuation value after observing state b at h^s is no more than δp. Hence, type c_i's continuation value at h^s *before* observing ω_s satisfies:

$$V_{c_i}(h^s) \le p\Big((1 - \delta) + \delta V_{c_i}(h^s, \omega_s = g)\Big) + (1 - p)\delta p = p(1 - \delta p) + \delta p V_{c_i}(h^s, \omega_s = g).$$

(A.2)

If $V_{c_i}(h^t) = v > p$, then there exists a special history $h^s \succeq h^t$ such that

$$V_{c_i}(h^s) \geq p + (v - p)\delta^{t-s}. \tag{A.3}$$

In order to satisfy (A.3), inequality (A.2) implies that there exists a *special history* h^τ that succeeds $(h^s, \omega_s = g)$ such that type c_i's continuation value at h^τ is at least $p + (v - p)\delta^{t-s-1}$. Iterate this process, the required continuation value at special histories will exceed 1. This leads to a contradiction since the sender's highest feasible payoff in the stage game is 1. This contradiction implies that type c_i's continuation value at h^t is no more than p if there are $k+\xi$ types in the support of the receivers' belief at h^t. This together with the induction hypothesis establishes Proposition 1.

Appendix B. Proof of Proposition 2

For every $h^\tau \in \mathcal{H}^\sigma$, let $\sigma_{c_j}(h^\tau) \in \Delta(\mathbf{A}_1)$ be the distribution of sender stage-game strategies prescribed by σ_{c_j} at history h^τ, and let $\boldsymbol{\alpha}_1(h^\tau) \in \Delta(\mathbf{A}_1)$ be the receiver's belief about the sender's stage-game strategy at h^τ. Theorem 1 in Gossner (2011) implies the following inequality:

$$\mathbb{E}^{(\sigma_{c_j},\sigma_2)}\Big[\sum_{\tau=0}^{+\infty} d(\sigma_{c_j}(h^\tau)||\boldsymbol{\alpha}_1(h^\tau))\Big] \leq -\log \pi_j, \tag{B.1}$$

where $d(\cdot||\cdot)$ is the Kullback-Leibler divergence and π_j is the prior probability of type c_j.

Inequality (B.1) implies that for every $\xi > 0$, the expected number of periods where $d(\sigma_{c_j}(h^\tau)||\alpha_1(h^\tau)) > \xi$ is no more than $T(\xi) \equiv \left\lceil \frac{-\log \pi_j}{\xi} \right\rceil$. Let $\sigma_2(h^\tau) \in \Delta(\mathbf{B})$ be the distribution over the receiver's stage-game strategies prescribed by σ_2 at h^τ. Let $\mathbf{A}_2^{\sigma_2}(h^\tau)$ be the support of $\sigma_2(h^\tau)$. Since the receiver plays a stage-game best reply to her expectation over the sender's stage-game strategy, we have:

$$\begin{aligned}
&\mathbb{E}^{(\sigma_{c_j},\sigma_2)}\Big[\sum_{\tau=0}^{\infty}(1-\delta)\delta^\tau \mathbf{1}\Big\{\mathbf{a}_2 \in \mathbf{A}_2^{\sigma_2}(h^\tau) \text{ but } \mathbf{a}_2 \text{ does not best reply to any } \boldsymbol{\alpha}_1 \\
&\qquad\qquad \text{with } ||\boldsymbol{\alpha}_1 - \sigma_{c_j}(h^\tau)|| \leq \sqrt{2\xi}\Big\}\Big] \\
&\leq \mathbb{E}^{(\sigma_{c_j},\sigma_2)}\Big[\sum_{\tau=0}^{\infty}(1-\delta)\delta^\tau \mathbf{1}\Big\{\mathbf{b} \in \mathbf{B}^{\sigma_2}(h^\tau) \text{ but } \mathbf{a}_2 \text{ does not best reply to any } \boldsymbol{\alpha}_1 \\
&\qquad\qquad \text{with } d(\sigma_{c_j}(h^\tau)||\boldsymbol{\alpha}_1) \leq \xi\Big\}\Big] \\
&\leq \mathbb{E}^{(\sigma_{c_j},\sigma_2)}\Big[\sum_{\tau=0}^{\infty}(1-\delta)\delta^\tau \mathbf{1}\Big\{d(\sigma_{c_j}(h^\tau)||\boldsymbol{\alpha}_1(h^\tau)) > \xi\Big\}\Big] \leq 1 - \delta^{T(\xi)}.
\end{aligned} \tag{B.2}$$

The first inequality comes from the Pinsker's inequality. The second inequality holds since $\mathbf{a}_2$ best replies to $\alpha_1(h^\tau)$, and the third inequality comes from (B.1).

Recall the definition of $\gamma^j \in \Delta(\mathbf{A}_1 \times \mathbf{A}_2)$ in (3.1). Let $\boldsymbol{\alpha}_2^j$ be the marginal distribution of γ^j on $\mathbf{A}_2$, and let $\gamma^j(\cdot|\mathbf{a}_2)$ be the distribution over $\mathbf{A}_2$ conditional on $\mathbf{a}_2$ when the joint distribution is γ^j. Let $\mathcal{A}_1(\mathbf{a}_2) \subset \Delta(\mathbf{A}_1)$ be the set of sender's mixed stage-game strategies that $\mathbf{a}_2$ best replies to. Consider any $\mathbf{a}_2 \in \mathbf{A}_2$ with the property that the Hausdorff distance between $\gamma^j(\cdot|\mathbf{a}_2)$ and

$\mathcal{A}_1(\mathbf{a}_2)$ is more than ε. I denote this distance by D. For every $\eta > 0$, let $\mathcal{A}_1^{\eta}(\mathbf{a}_2)$ be the set of elements in $\Delta(\mathbf{A}_1)$ whose Hausdorff distance to $\mathcal{A}_1(\mathbf{a}_2)$ is no more than η. Since the Hausdorff distance between any two points in $\Delta(\mathbf{A}_1)$ is at most 1, for any $\rho \in \Delta(\Delta(\mathbf{A}_1))$ that has countable support $\{\boldsymbol{\alpha}_1^i\}_{i\in\mathbb{N}}$, and satisfies $\sum_{i\in\mathbb{N}} \rho(\boldsymbol{\alpha}_1^i)\boldsymbol{\alpha}_1^i = \boldsymbol{\alpha}_1^j(\cdot|\mathbf{a}_2)$, we have $\sum_{\boldsymbol{\alpha}_1^i \notin \mathcal{A}_1^{\eta}(\mathbf{a}_2)} \rho(\boldsymbol{\alpha}_1^i) \geq \frac{D-\eta}{1+D-\eta}$. This implies that

$$\mathbb{E}^{(\sigma_{c_j},\sigma_2)}\Big[\sum_{\tau=0}^{\infty}(1-\delta)\delta^{\tau}\mathbf{1}\Big\{\mathbf{a}_2 \in \mathbf{A}_2^{\sigma_r}(h^{\tau}) \text{ but } \mathbf{a}_2 \text{ doesn't best reply to } \boldsymbol{\alpha}_1 \text{ with } ||\boldsymbol{\alpha}_1 - \sigma_{c_j}(h^{\tau})|| \leq \eta\Big\}\Big]$$

$$\geq \frac{\boldsymbol{\alpha}_2^j(\mathbf{a}_2)(D-\eta)}{1+D-\eta}. \tag{B.3}$$

Let $\eta \equiv \frac{D}{2}$ and $\xi \equiv \frac{D^2}{8}$, we have $\sqrt{2\xi} = \frac{D}{2}$. Therefore, (B.2) and inequality (B.3) together imply that for every strategy profile that is a Nash equilibrium under discount factor δ, we have:

$$\boldsymbol{\alpha}_2^j(\mathbf{a}_2) \leq \Big(1-\delta^{T(\frac{D^2}{8})}\Big)\frac{1+D/2}{D/2}. \tag{B.4}$$

Since $D \geq \varepsilon$, there exists $\underline{\delta} \in (0,1)$ such that the right-hand-side of (B.4) is less than ε for every $\delta \in (\underline{\delta}, 1)$. That is to say, for every $\mathbf{a}_2 \in \mathbf{A}_2$ such that $\mathbf{a}_2$ is not an ε-best reply to $\gamma^j(\cdot|\mathbf{a}_2)$, the marginal distribution $\boldsymbol{\alpha}_2^j$ assigns probability less than ε to $\mathbf{a}_2$.

Appendix C. Proof of Proposition 3

Recall that

$$v_j^* \equiv \max_{\gamma^j \in \Delta(\mathbf{A}_1\times\mathbf{A}_2)} \sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma^j(\mathbf{a}_1,\mathbf{a}_2)u_1(c_j,\mathbf{a}_1,\mathbf{a}_2),$$

subject to

$$\sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma^j(\mathbf{a}_1,\mathbf{a}_2)u_1(c_1,\mathbf{a}_1,\mathbf{a}_2) \leq p \text{ for every } j \in \{1,2,...,n\}, \tag{C.1}$$

and

$$\mathbf{a}_2 \in \arg\max_{\mathbf{a}_2'\in\mathbf{A}_2} u_2(\gamma^j(\cdot|\mathbf{a}_2),\mathbf{a}_2') \text{ for every } \mathbf{a}_2 \in \mathbf{A}_2 \text{ in the support of } \gamma^j. \tag{C.2}$$

I only need to show that $\limsup_{\varepsilon\downarrow 0} v_j^{\varepsilon} \leq v_j^*$. This is trivial for the ethical types since $v_j^{\varepsilon} = v_j^* = p$ for every $c_j \geq 1$. I focus on the non-ethical types for the rest of this proof.

Let Γ^{ε} be the set of $\gamma \in \Delta(\mathbf{A}_1\times\mathbf{A}_2)$ that satisfies constraint (C.1) and the ε-relaxed version of constraint (C.2). Let Γ be the set of $\gamma \in \Delta(\mathbf{A}_1\times\mathbf{A}_2)$ that satisfies constraints (C.1) and (C.2). I show that for every $\eta > 0$, there exists $\varepsilon > 0$, such that for every $\gamma^{\varepsilon} \in \Gamma^{\varepsilon}$, there exists $\gamma \in \Gamma$ that belongs to an η-neighborhood of γ^{ε}. This will imply that $\limsup_{\varepsilon\downarrow 0} v_j^{\varepsilon} \leq v_j^*$.

First, since the number of pure stage-game strategies is finite, for every $\eta > 0$, there exists $\varepsilon > 0$, such that for every $\boldsymbol{\alpha}_1 \in \Delta(\mathbf{A}_1)$ and $\mathbf{a}_2 \in \mathbf{A}_2$ satisfying $\mathbf{a}_2$ being an ε-best reply to $\boldsymbol{\alpha}_1$, there exists $\boldsymbol{\alpha}_1' \in \Delta(\mathbf{A}_1)$ that belongs to an η-neighborhood of $\boldsymbol{\alpha}_1$ such that $\mathbf{a}_2$ best replies to $\boldsymbol{\alpha}_1'$.

Second, for every $\gamma^{\varepsilon} \in \Gamma^{\varepsilon}$, let $\mathbf{A}_2^* \equiv \big\{\mathbf{a}_2 \in \mathbf{A}_2 \big| \, \mathbf{a}_2 \text{ best replies to } \gamma^{\varepsilon}(\cdot|\mathbf{a}_2)\big\}$. By definition, the marginal distribution of γ^{ε} on $\mathbf{A}_2$, denoted by $\boldsymbol{\alpha}_2^{\varepsilon}$, assigns probability at most ε to every $\mathbf{a}_2 \notin \mathbf{A}_2^*$. Let us consider another joint distribution $\gamma' \in \Delta(\mathbf{A}_1\times\mathbf{A}_2)$, which is constructed according to:

1. For every $\mathbf{a}_2 \in \mathbf{A}_2^*$, there exists γ^* that belongs to an η-neighborhood of $\gamma^\varepsilon(\cdot|\mathbf{a}_2)$, with $\mathbf{a}_2$ best replies to $\gamma^*(\cdot|\mathbf{a}_2)$.
2. The marginal distribution of γ' on $\mathbf{A}_2$ assigns probability $\frac{\boldsymbol{\alpha}_2^\varepsilon(\mathbf{a}_2)}{\boldsymbol{\alpha}_2^\varepsilon(\mathbf{A}_2^*)}$ to $\mathbf{a}_2$, and the distribution of the sender's stage-game strategies is $\gamma^*(\cdot|\mathbf{a}_2)$ conditional on $\mathbf{a}_2$.

Since $\sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma^\varepsilon(\mathbf{a}_1,\mathbf{a}_2)u_1(c_1,\mathbf{a}_1,\mathbf{a}_2) \leq p$, and $\boldsymbol{\alpha}_2^\varepsilon$ assigns probability less than ε to every $\mathbf{a}_2 \notin \mathbf{A}_2^*$, there exists $X : [0,1] \to \mathbb{N}$ with $\lim_{\eta\to 0} X(\eta) = 0$ such that

$$\sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma'(\mathbf{a}_1,\mathbf{a}_2)u_1(c_1,\mathbf{a}_1,\mathbf{a}_2) \leq p + X(\eta) \tag{C.3}$$

$$\begin{aligned} v_j^* + X(\eta) &\geq \sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma'(\mathbf{a}_1,\mathbf{a}_2)u_1(c_j,\mathbf{a}_1,\mathbf{a}_2) + X(\eta) \\ &\geq \sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma^\varepsilon(\mathbf{a}_1,\mathbf{a}_2)u_1(c_j,\mathbf{a}_1,\mathbf{a}_2) = v_j^\varepsilon. \end{aligned} \tag{C.4}$$

Consider two cases separately,

1. If $\sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma'(\mathbf{a}_1,\mathbf{a}_2)u_1(c_1,\mathbf{a}_1,\mathbf{a}_2) \leq p$, then γ' satisfies constraints (C.1) and (C.2), and attains payoff within $X(\eta)$ of v_j^ε.
2. If $\sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma'(\mathbf{a}_1,\mathbf{a}_2)u_1(c_1,\mathbf{a}_1,\mathbf{a}_2) > p$, then let $\gamma'' \in \Delta(\mathbf{A}_1 \times \mathbf{A}_2)$ be a convex combination of γ' and the Dirac measure on $(\mathbf{a}_1^L, \mathbf{a}_2^N)$, with the convex weight on γ' equals

$$\frac{p}{\sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma'(\mathbf{a}_1,\mathbf{a}_2)u_1(c_1,\mathbf{a}_1,\mathbf{a}_2)}.$$

Since all types' stage-game payoffs are no more than 0 under $(\mathbf{a}_1^L, \mathbf{a}_2^N)$, γ'' satisfies constraint (C.1). Since $\mathbf{a}_2^N$ best replies to $\mathbf{a}_1^L$, γ'' satisfies constraint (C.2). The definition of v_j^* implies that

$$v_j^* \geq \sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma''(\mathbf{a}_1,\mathbf{a}_2)u_1(c_j,\mathbf{a}_1,\mathbf{a}_2). \tag{C.5}$$

According to (C.3) and (C.4),

$$\begin{aligned} &\sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma''(\mathbf{a}_1,\mathbf{a}_2)u_1(c_j,\mathbf{a}_1,\mathbf{a}_2) \\ &\geq \frac{p}{p+X(\eta)}\Big(\sum_{(\mathbf{a}_1,\mathbf{a}_2)\in\mathbf{A}_1\times\mathbf{A}_2} \gamma'(\mathbf{a}_1,\mathbf{a}_2)u_1(c_j,\mathbf{a}_1,\mathbf{a}_2) - X(\eta)\Big) \\ &\geq \frac{p}{p+X(\eta)}\Big(v_j^\varepsilon - X(\eta)\Big). \end{aligned} \tag{C.6}$$

The expression on the right-hand-side of (C.6) implies that for every $\rho > 0$, there exists $\eta > 0$ such that once we pick ε according to η, we have $v_j^* \geq \frac{p}{p+X(\eta)}\Big(v_j^\varepsilon - X(\eta)\Big) \geq v_j^\varepsilon - \rho$.

Appendix D. Proof of Theorem 1

Let $\boldsymbol{v}^H \equiv (p, ..., p)$, $\boldsymbol{v}^N \equiv \Big(-c_1(1-p), -c_2(1-p), ..., -c_n(1-p)\Big)$, and $\boldsymbol{v}^L \equiv \Big(p+(1-c_1)(1-p), p+(1-c_2)(1-p), ..., p+(1-c_n)(1-p)\Big)$, which are the sender's stage-game payoffs from stage-game strategy profiles $(\mathbf{a}_1^H, \mathbf{a}_2^T)$, $(\mathbf{a}_1^L, \mathbf{a}_2^N)$, and $(\mathbf{a}_1^L, \mathbf{a}_2^T)$, respectively. Let

$$\boldsymbol{u}(\rho) \equiv \frac{(1-\rho)c_1}{\rho(1-c_1)+c_1}\boldsymbol{v}^H + \frac{\rho c_1}{\rho(1-c_1)+c_1}\boldsymbol{v}^L + \frac{\rho(1-c_1)}{\rho(1-c_1)+c_1}\boldsymbol{v}^N \text{ for every } \rho \in [0, \rho^*]. \tag{D.1}$$

Theorem 1 is implied by Proposition 5, which I show in the rest of this appendix.

Proposition 5. *For every $\varepsilon > 0$ and $\rho \in [0, \rho^*)$, there exists $\underline{\delta} \in (0, 1)$ such that for every $\boldsymbol{\pi} \in \Delta(\mathcal{C})$ with $\pi_1 \geq \varepsilon$ and $\delta > \underline{\delta}$, there exists an equilibrium in which the sender's payoff is $\boldsymbol{u}(\rho)$.*

I construct equilibria in which (i) the receiver ignores the sender's cheap talk message at all histories (therefore, I omit m_t in the subsequent part of my proof), and (ii) the distribution of stage-game strategy profiles at every history is supported in $\big\{(\mathbf{a}_1^H, \mathbf{a}_2^T), (\mathbf{a}_1^L, \mathbf{a}_2^T), (\mathbf{a}_1^L, \mathbf{a}_2^N)\big\}$. This implies that the sender's stage-game payoff and his promised continuation value at every history h^t are convex combinations of $\boldsymbol{v}^H$, $\boldsymbol{v}^N$, and $\boldsymbol{v}^L$. Later on, I verify that the sender's continuation value at every history belongs to V^*, which is the convex hull of $\{\boldsymbol{v}^H, \boldsymbol{v}^*, \underline{\boldsymbol{v}}, \overline{\boldsymbol{v}}\}$, where $\underline{\boldsymbol{v}} \equiv p^*\boldsymbol{v}^H + (1-p^*)\boldsymbol{v}^N$ with

$$p^* \equiv \frac{c_1(1-p)}{p+c_1(1-p)}, \tag{D.2}$$

and $\overline{\boldsymbol{v}} \equiv q^*(\rho^*\boldsymbol{v}^L + (1-\rho^*)\boldsymbol{v}^H) + (1-q^*)\boldsymbol{v}^N$ where $q^* \in [0, 1]$ is pinned down by the condition that the first entry of vector $\overline{\boldsymbol{v}}$ equals 0. I depict these payoff vectors in Fig. 2. By definition,

1. $\overline{\boldsymbol{v}}$ is the intersection of the vertical axis and the line segment between $\boldsymbol{v}^*$ and $\boldsymbol{v}^N$,
2. $\underline{\boldsymbol{v}}$ is the intersection of the vertical axis and the line segment between $\boldsymbol{v}^H$ and $\boldsymbol{v}^N$.

The equilibrium play keeps track of the following state variables:

1. The probability of type c_1 in the receivers' belief at h^t, denoted by $\eta(h^t) \in [0, 1]$, which I call the sender's *reputation*. For future reference, let $\eta(h^t, (\omega_t, r_t))$ be the probability the receivers' belief assigns to type c_1 after observing (ω_t, r_t) at h^t.
2. The sender's promised continuation value at h^t, denoted by $\boldsymbol{v}(h^t) \in \mathbb{R}^n$. Since $\boldsymbol{v}(h^0) \equiv \boldsymbol{u}(\rho)$ is a convex combination of $\boldsymbol{v}^H$, $\boldsymbol{v}^N$ and $\boldsymbol{v}^L$, and the sender's stage-game payoff is also a convex combination of $\boldsymbol{v}^H$, $\boldsymbol{v}^N$ and $\boldsymbol{v}^L$, so his continuation value at every history must also be a convex combination of $\boldsymbol{v}^H$, $\boldsymbol{v}^N$ and $\boldsymbol{v}^L$ in order to satisfy his promise-keeping constraint. Therefore, $\boldsymbol{v}(h^t)$ can be written as $\boldsymbol{v}(h^t) = p^H(h^t)\boldsymbol{v}^H + p^N(h^t)\boldsymbol{v}^N + p^L(h^t)\boldsymbol{v}^L$.
3. The lowest lying cost in the support of the receivers' posterior belief $\underline{c}(h^t)$ and its probability.

The initial values of these state variables are $\eta(h^0) = \pi_1$, $\underline{c}(h^0) = c_n$,

$$p^H(h^0) = \frac{(1-\rho)c_1}{\rho(1-c_1)+c_1}, \ p^L(h^0) = \frac{\rho c_1}{\rho(1-c_1)+c_1}, \text{ and } p^N(h^0) = \frac{\rho(1-c_1)}{\rho(1-c_1)+c_1}.$$

Since c_n is the lowest cost type, the prior probability of type $\underline{c}(h^0)$ is π_n.

Defining useful constants: I start from defining several useful constants. Recall that $\rho^* \equiv \frac{p}{1-p}$. For every $j \geq 3$, let k_j be the smallest integer $k \in \mathbb{N}$ such that:

$$\Big(1-(1-\rho^*)\pi_1\Big)\frac{(\pi_j/k)}{\sum_{\tau=2}^{j-1}\pi_\tau+(\pi_j/k)} \leq \rho^*. \tag{D.3}$$

The value of k_j will later be used to in the formulas for the sender's mixing probabilities. Let $\eta^* \in [(1-\rho^*)\pi_1, \pi_1)$ be large enough such that for every $\eta \in [\eta^*, \pi_1]$, we have:

$$\frac{\pi_1-\eta}{\pi_1(1-\eta)} \leq \min_{j\in\{3,\ldots,n\}}\left\{\frac{\pi_j/k_j}{\pi_2+\ldots+\pi_j}\right\} \tag{D.4}$$

For every $\rho \in (0,\rho^*)$, there exist $\widehat{\tau}, \widehat{l} \in \mathbb{N}$ such that $\widehat{\tau}/\widehat{l}$ is strictly between ρ and ρ^*. By construction, there exists a large enough integer $\chi \in \mathbb{N}$ such that $\rho^* > \frac{\widehat{\tau}}{\widehat{l}} = \frac{\widehat{\tau}\chi}{\widehat{l}\chi} > \frac{\widehat{\tau}\chi}{\widehat{l}\chi+1} > \rho$. Let $\tau \equiv \widehat{\tau}\chi$ and $l \equiv \widehat{l}\chi$. Let

$$\widetilde{\rho} \equiv \frac{1}{2}\Big(\frac{\tau}{l}+\frac{\tau}{l+1}\Big) \quad \text{and} \quad \widehat{\rho} \equiv \frac{1}{2}\Big(\frac{\tau}{l}+\rho^*\Big). \tag{D.5}$$

Let $\underline{\delta} \in (0,1)$ to be large enough such that for every $\delta > \underline{\delta}$,

$$\frac{\delta+\delta^2+\ldots+\delta^\tau}{\delta+\delta^2+\ldots+\delta^l} > \widetilde{\rho} > \frac{\delta^{l-\tau+1}(\delta+\delta^2+\ldots+\delta^\tau)}{\delta+\delta^2+\ldots+\delta^{l+1}}. \tag{D.6}$$

Let $\widehat{\delta} \equiv \delta\frac{1-p}{1-\delta p}$ be the sender's effective discount factor. I require the sender's discount factor δ to be large enough such that $\widehat{\delta} > \underline{\delta}$. In the subsequent proof, I will introduce additional requirements on δ. These additional requirements are compatible with (D.6) since the number of restrictions is finite and all restrictions require δ to be close enough to 1. By construction, $\rho^* > \widehat{\rho} > \frac{\tau}{l} > \widetilde{\rho} > \frac{\tau}{l-1} > \rho$. Let $\lambda > 0$ be small enough such that:

$$(1+\lambda\rho^*)^{1-\widehat{\rho}}(1-\lambda(1-\rho^*))^{\widehat{\rho}} > 1. \tag{D.7}$$

Since $\rho^* > \widehat{\rho}$, such a λ exists according to the Taylor's expansion theorem.

Active learning phase: Play starts from the *active learning phase*, and belongs to this phase as long as $p^L(h^t) > 0$ and the first entry of the following n-dimensional vector is non-negative:

$$\frac{p^L(h^t)-(1-\widehat{\delta})}{\widehat{\delta}}\boldsymbol{v}^L + \frac{p^H(h^t)}{\widehat{\delta}}\boldsymbol{v}^H + \frac{p^N(h^t)}{\widehat{\delta}}\boldsymbol{v}^N. \tag{D.8}$$

That is, type c_1's promised continuation value is non-negative. This is because every type can secure payoff 0 by truthfully revealing both states in every period. The receiver plays $\mathbf{a}_2^T$ and the sender's behavior depends on the sign of $p^L(h^t)-(1-\widehat{\delta})$.

1. If $p^L(h^t) \geq 1-\widehat{\delta}$, then the senders of types c_2 to c_n play the same stage-game strategy, and type c_1 behaves differently. All types send message g with probability 1 when $\omega_t = g$, which implies that $\eta(h^t,(g,g)) = \eta(h^t)$. The mixing probabilities when $\omega_t = b$ are pinned down by the receiver's posterior beliefs after observing $\omega_t = b$ and the sender's message m_t, which are given by:

$$\eta(h^t,(b,g))-\eta^*=\Big(1-\lambda(1-\rho^*)\Big)(\eta(h^t)-\eta^*), \tag{D.9}$$

and

$$\eta(h^t,(b,b))-\eta^*=\min\Big\{1-\eta^*,(1+\lambda\rho^*)(\eta(h^t)-\eta^*)\Big\}, \tag{D.10}$$

where $\eta^*\in(0,\pi_1)$ and $\lambda>0$ are constants defined in (D.4) and (D.7). One can compute the sender's mixing probabilities from (D.9) and (D.10) according to Bayes rule.

Intuitively, η^* is a lower bound on the sender's reputation in the active learning phase, and λ is a small positive number that measures the speed with which the receivers learn about the sender's type. The sender's reputation remains unchanged when the state is good, improves when he truthfully reveals the bad state, and deteriorates when he lies about the bad state. All types mix in the bad state unless the sender's reputation $\eta(h^t)$ is close enough to 1, in which case only type c_1 mixes while the other types recommend the good action for sure.

Next, I construct the sender's continuation values in order to satisfy his incentive constraints and promise keeping constraints in recursive form. If h^t is such that the sender's reputation does not reach one after truthfully revealing the bad state, i.e., $\eta(h^t,(b,b))<1$, then the sender's continuation values are $\boldsymbol{v}(h^t,(g,g))=\boldsymbol{v}(h^t)$, $\boldsymbol{v}(h^t,(g,b))=\underline{\boldsymbol{v}}$,

$$\boldsymbol{v}(h^t,(b,g))=\frac{p^H(h^t)}{\widehat{\delta}}\boldsymbol{v}^H+\frac{p^L(h^t)-(1-\widehat{\delta})}{\widehat{\delta}}\boldsymbol{v}^L+\frac{p^N(h^t)}{\widehat{\delta}}\boldsymbol{v}^N, \tag{D.11}$$

and

$$\boldsymbol{v}(h^t,(b,b))=\frac{p^H(h^t)-(1-\widehat{\delta})}{\widehat{\delta}}\boldsymbol{v}^H+\frac{p^L(h^t)}{\widehat{\delta}}\boldsymbol{v}^L+\frac{p^N(h^t)}{\widehat{\delta}}\boldsymbol{v}^N. \tag{D.12}$$

One can verify that under these promised continuation values in period $t+1$, the sender's expected payoff from playing $\mathbf{a}_1^H$ and $\mathbf{a}_1^L$ at h^t are both $\boldsymbol{v}(h^t)$. Therefore, all types are indifferent when $\omega_t=b$, and strictly prefer to send message g when $\omega_t=g$.

If the sender's reputation reaches one after truthfully revealing the bad state, i.e., $\eta(h^t,(b,b))=1$, then his continuation values are $\boldsymbol{v}(h^t,(g,g))=\boldsymbol{v}(h^t)$, $\boldsymbol{v}(h^t,(g,b))=\underline{\boldsymbol{v}}$, $\boldsymbol{v}(h^t,(b,g))=\frac{p^H(h^t)}{\widehat{\delta}}\boldsymbol{v}^H+\frac{p^L(h^t)-(1-\widehat{\delta})}{\widehat{\delta}}\boldsymbol{v}^L+\frac{p^N(h^t)}{\widehat{\delta}}\boldsymbol{v}^N$, and

$$\boldsymbol{v}(h^t,(b,b))=q(h^t)\boldsymbol{v}^H+(1-q(h^t))\boldsymbol{v}^N, \tag{D.13}$$

where $q(h^t)\in[0,1]$ is such that the first entry of $\boldsymbol{v}(h^t,(b,b))$ equals the first entry of vector:

$$\frac{p^H(h^t)-(1-\widehat{\delta})}{\widehat{\delta}}\boldsymbol{v}^H+\frac{p^L(h^t)}{\widehat{\delta}}\boldsymbol{v}^L+\frac{p^N(h^t)}{\widehat{\delta}}\boldsymbol{v}^N. \tag{D.14}$$

One can verify that under these promised continuation values in period $t+1$, type-c_1 sender is indifferent between $\mathbf{a}_1^H$ and $\mathbf{a}_1^L$ at h^t since his continuation value after truthfully revealing the bad state equals the first entry of (D.14). Types c_2 to c_n strictly prefer to lie about the bad state at h^t. Moreover, the sender's expected payoff by playing $\mathbf{a}_1^L$ at h^t equals $\boldsymbol{v}(h^t)$.

2. Suppose $p^L(h^t)\in(0,1-\widehat{\delta})$. All types of sender in the support of the receiver's belief except for type-$\underline{c}(h^t)$ sender play stage-game strategy $\mathbf{a}_1^H$. Only type-$\underline{c}(h^t)$ sender may mix between $\mathbf{a}_1^L$ and $\mathbf{a}_1^H$. I specify type-$\underline{c}(h^t)$ sender's mixing probabilities later on, which depend on $\underline{c}(h^t)$ and the probability of type $\underline{c}(h^t)$ in the receiver's belief at h^t.

   Under the above strategy for the sender, the receiver's belief remains unchanged when $\omega_t=g$, the receiver partially rules out type $\underline{c}(h^t)$ when the sender truthfully reveals the bad state, and the receiver believes that the sender's type is $\underline{c}(h^t)$ when the sender lies about the bad state.

The sender's continuation values are such that $\boldsymbol{v}(h^t,(g,g))=\boldsymbol{v}(h^t)$, $\boldsymbol{v}(h^t,(g,b))=\underline{\boldsymbol{v}}$,

$$\boldsymbol{v}(h^t,(b,g)) \equiv \frac{Q(h^t)}{\widehat{\delta}}\boldsymbol{v}^H + \frac{\widehat{\delta}-Q(h^t)}{\widehat{\delta}}\boldsymbol{v}^N, \tag{D.15}$$

where

$$Q(h^t) \equiv p^H(h^t) - (1-\widehat{\delta}) + \frac{p^L(h^t)}{p+(1-p)\underline{c}(h^t)}. \tag{D.16}$$

The sender's continuation value after sending message b after observing $\omega_t = b$ depends on whether $\eta(h^t,(b,b))$ equals 1. If $\eta(h^t,(b,b)) < 1$, then the sender's continuation value at $(h^t,(b,b))$ is

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

If $\eta(h^t,(b,b)) = 1$, then the sender's continuation value at $(h^t,(b,b))$ is

$$\boldsymbol{v}(h^t,(b,b)) = q(h^t)\boldsymbol{v}^H + (1-q(h^t))\boldsymbol{v}^N,$$

where $q(h^t) \in [0,1]$ is such that the first entry of $\boldsymbol{v}(h^t,(b,b))$ equals the first entry of payoff vector:

$$\frac{p^H(h^t)-(1-\widehat{\delta})}{\widehat{\delta}}\boldsymbol{v}^H + \frac{p^L(h^t)}{\widehat{\delta}}\boldsymbol{v}^L + \frac{p^N(h^t)}{\widehat{\delta}}\boldsymbol{v}^N.$$

Under these continuation values, type $\underline{c}(h^t)$ prefers to send message g when the state is b, while other types strictly prefer to send message b when the state is b.

I specify the sender's mixing probabilities in the active learning phase when $p^L(h^t) \in (0, 1-\widehat{\delta})$. Then I verify the receiver's incentive to play $\mathbf{a}_2^T$ at Class 2 histories given the sender's strategy. Recall that all types in the support of the receiver's belief except for the lowest-cost type $\underline{c}(h^t)$ play $\mathbf{a}^H$ with probability 1, and the only type that may mix is type $\underline{c}(h^t)$. Let

$$l(h^t) \equiv \#\Big\{h^s \Big| h^s \prec h^t, h^s \text{ belongs to Class 2}, \omega_s = b, \text{ and } \underline{c}(h^s) = \underline{c}(h^t)\Big\} \tag{D.17}$$

be the number of history h^s that (1) strictly precedes h^t, and (2) the lowest lying cost type in the support of the receiver's belief is $\underline{c}(h^t)$, and (3) the realized state at history h^s is $\omega_s = b$.

1. If $\underline{c}(h^t) = c_j$ with $j \geq 3$, then type-$\underline{c}(h^t)$ sender plays $\mathbf{a}_1^L$ at h^t with probability $\frac{1}{k_j - l(h^t)}$ and $\mathbf{a}_1^H$ with complementary probability, where k_j is the integer defined in (D.3). The definition of k_j in (D.3) implies that the sender plays $\mathbf{a}_1^H$ with probability at least $1-\rho^*$ at every such history, which implies that the receiver has an incentive to play $\mathbf{a}_2^T$.
2. If $\underline{c}(h^t) = c_2$, then type-c_2 sender plays $\mathbf{a}_1^L$ at h^t with probability $\min\{1, \frac{\rho^*}{1-\eta(h^t)}\}$ and plays $\mathbf{a}_1^H$ with complementary probability. Under these mixing probabilities, the sender plays $\mathbf{a}_1^H$ with probability at least $1-\rho^*$, which implies that the receiver has an incentive to play $\mathbf{a}_2^T$.

Rebounding phase: Play reaches the *rebounding phase* if $p^L(h^t) \neq 0$ and the first entry of (D.8) is negative. At those histories, the receiver plays $\mathbf{a}_2^N$ and all types of sender play $\mathbf{a}_1^L$. Hence, the sender's message reveals no information about his type and the receiver's incentive constraints are satisfied. The sender's continuation value after observing $(\omega_t, r_t) = (g, g)$ at h^t equals $\boldsymbol{v}(h^t)$, his continuation value after observing $(\omega_t, r_t) = (g, b)$ at h^t equals $\boldsymbol{v}(h^t)$, his continuation value after observing $(\omega_t, r_t) = (b, b)$ at h^t equals $\underline{\boldsymbol{v}}$, and his continuation value after observing $(\omega_t, r_t) = (b, g)$ at h^t equals

$$\boldsymbol{v}(h^t, (b, g)) \equiv \frac{p^H(h^t)}{\widehat{\delta}}\boldsymbol{v}^H + \frac{p^L(h^t)}{\widehat{\delta}}\boldsymbol{v}^L + \frac{p^N(h^t) - (1-\widehat{\delta})}{\widehat{\delta}}\boldsymbol{v}^N. \tag{D.18}$$

Under these continuation values, each type of sender has an incentive to send message g when the state is g. This implies that

$$\boldsymbol{v}(h^t) = (1-p)\Big\{(1-\delta)\mathbf{0} + \delta \underbrace{\boldsymbol{v}(h^t, (g, g))}_{=\boldsymbol{v}(h^t)}\Big\} + p\Big\{-(1-\delta)\mathbf{c} + \delta \boldsymbol{v}(h^t, (b, b))\Big\},$$

where $\mathbf{0} \equiv (0, ..., 0)$ and $\mathbf{c} \equiv (c_1, ..., c_n)$. Since $\boldsymbol{v}(h^t) \geq 0$, we have $\boldsymbol{v}(h^t) \geq (1-\delta)\mathbf{0} + \delta\boldsymbol{v}(h^t)$. Hence, $-(1-\delta)\mathbf{c} + \delta\boldsymbol{v}(h^t, (b, g)) > \boldsymbol{v}(h^t) \geq \underline{\boldsymbol{v}}$, so the sender has an incentive to send g when the state is b.

Absorbing phase: Play reaches the absorbing phase when $p^L(h^t) = 0$ after which play remains in the absorbing phase forever and learning about the sender's type stops. I specify the continuation play after reaching this phase by constructing for any $p^H \geq p^*$ and under any belief, a continuation equilibrium in which there is no learning and the sender's continuation value is $\boldsymbol{v} = p^H\boldsymbol{v}^H + p^N\boldsymbol{v}^N$ where $p^N \equiv 1 - p^H$. Recall the definition of p^* in (D.2).

1. If $\frac{p^H(h^t)-(1-\delta)}{\delta} \geq \frac{p^*+1}{2}$, then all types of sender communicate honestly by playing $\mathbf{a}_1^H$, and his continuation values are given by $\boldsymbol{v}(h^t, (g, g)) = \boldsymbol{v}(h^t)$,
$$\boldsymbol{v}(h^t, (b, b)) = \frac{p^H(h^t) - (1-\delta)}{\delta}\boldsymbol{v}^H + \frac{p^N(h^t)}{\delta}\boldsymbol{v}^N, \tag{D.19}$$
and $\boldsymbol{v}(h^t, (b, g)) = \boldsymbol{v}(h^t, (g, b)) = \underline{\boldsymbol{v}}$. When δ is large enough, every type of sender prefers to conform at those histories.
2. If $\frac{p^H(h^t)-(1-\delta)}{\delta} < \frac{p^*+1}{2}$, then all types of sender play $\mathbf{a}_1^L$, and the sender's continuation value after sending message g is $\frac{p^H(h^t)}{\delta}\boldsymbol{v}^H + \frac{p^N(h^t)-(1-\delta)}{\delta}\boldsymbol{v}^N$ and his continuation value after sending message b is $p^*\boldsymbol{v}^H + (1-p^*)\boldsymbol{v}^L$. One can verify that type-c_1 sender weakly prefers to conform and types c_2 to c_m strictly prefer to conform. This continuation equilibrium is incentive compatible regardless of the receiver's belief about the sender's type.

Verifying incentive constraints: I verify that the sender's promised continuation value at every history belongs to $V^* \equiv \text{co}\{\boldsymbol{v}^H, \boldsymbol{v}^*, \underline{\boldsymbol{v}}, \overline{\boldsymbol{v}}\}$, where co(·) denotes the convex hull. An *outcome path* in period t consists of the states and the sender's recommendations from period 0 to period $t-1$, with $O(h^t) \equiv \big((\omega_0, r_0), ..., (\omega_{t-1}, r_{t-1})\big)$. For every on-path history h^t, $o(h^t)$ consists only of (g, g), (b, b), and (b, g). A *reduced outcome path* $o(h^t)$ is derived from $O(h^t)$ by ignoring periods where the state is g. Hence, $o(h^t)$ can be summarized by a sequence of recommendations.

Let $o(h^t)$ be the reduced outcome path at h^t, which consists of a sequence of g and b with the nth element of $o(h^t)$ the sender's message when state b is realized for the nth time. For every $h^t \succ h^s$, let $o(h^t \backslash h^s)$ be the reduced outcome path between h^s and h^t. Since for every pair of histories h^s and h^t with $h^t \succ h^s$, suppose all histories between h^s and h^t belong to the active learning phase, then $\eta(h^t)$ depends only on $\eta(h^s)$ and the reduced outcomes from h^s to h^t, denoted by $o(h^t \backslash h^s)$.

- For every $r \in \{h, l\}$, let $N_r(o(h^t \backslash h^s))$ be the number of action recommendation r in reduced outcome path $o(h^t \backslash h^s)$.
- Let $|o(h^t \backslash h^s)|$ be the number of elements in $o(h^t \backslash h^s)$, i.e., the number of times with which state b is realized between history h^s and history h^t.

For every $h^t \succ h^s$, suppose all histories from h^s to h^t belong to Class 1, then according to (D.9) and (D.10), the probability of type-c_1 sender in the receiver's belief at h^t can be computed via

$$\eta(h^t) = \eta^* + (\eta(h^s) - \eta^*)\Big(1 - \lambda(1-\rho^*)\Big)^{N_g(o(h^t \backslash h^s))}\Big(1 + \lambda\rho^*\Big)^{N_b(o(h^t \backslash h^s))}. \tag{D.20}$$

Moreover, the convex weights of $\boldsymbol{v}^N$, $\boldsymbol{v}^H$, and $\boldsymbol{v}^L$ in the sender's continuation value are given by:

$$p^N(h^t) = \frac{p^N(h^s)}{\widehat{\delta}^{|o(h^t \backslash h^s)|}}, \quad p^H(h^t) = \frac{p^H(h^s) - (1-\widehat{\delta})\sum_{\tau=1}^{|o(h^t \backslash h^s)|}\widehat{\delta}^{\tau}\mathbf{1}\{r_\tau = l\}}{\widehat{\delta}^{|o(h^t \backslash h^s)|}},$$

and

$$p^L(h^t) = \frac{p^L(h^s) - (1-\widehat{\delta})\sum_{\tau=1}^{|o(h^t \backslash h^s)|}\widehat{\delta}^{\tau}\mathbf{1}\{r_\tau = h\}}{\widehat{\delta}^{|o(h^t \backslash h^s)|}}.$$

I show that for every sequence of Class 1 histories, as long as play remains in the active learning phase by the end of this sequence, the discounted average frequency of action recommendation b divided by the discounted average frequency of action recommendation g in the reduced outcome path (i.e., only counting periods where the state is bad) is below some cutoff. This provides an upper bound on the sender's continuation value.

Lemma D.0. *Suppose the sender's strategy in Class 1 histories is given by (D.9) and (D.10), with (η^*, λ) satisfying (D.4) and (D.7). For every $\underline{\eta} \in (\eta^*, 1)$, there exist $T \in \mathbb{N}$ and $\underline{\delta} \in (0,1)$, such that when $\eta(h^s) \geq \underline{\eta}$ and $\widehat{\delta} > \underline{\delta}$, if $h^t \succ h^s$ and all histories between h^s and h^t belong to Class 1, then:*

$$(1-\widehat{\delta})\sum_{\tau=1}^{|o(h^t \backslash h^s)|}\widehat{\delta}^{\tau-1}\mathbf{1}\{r_\tau = b\} \leq (1-\widehat{\delta}^T) + (1-\widehat{\delta})\sum_{\tau=1}^{|o(h^t \backslash h^s)|}\widehat{\delta}^{\tau-1}\mathbf{1}\{r_\tau = g\}\cdot\frac{1-\widetilde{\rho}}{\widetilde{\rho}}. \tag{D.21}$$

The proof follows from that of Lemma A.2 in Pei (2021). The only difference is caused by the i.i.d. state ω_t. This is being taken care by replacing the discount factor δ with the effective discount factor $\widehat{\delta}$, since the sender's continuation value does not change when the realized state is h.

I provide some intuition on the motivation for showing this lemma. Recall that starting from a Class 1 history of the active learning phase, play reaches to absorbing phase either because (1) the sender lies about the bad state too frequently, in which case $p^L(h^t)$ goes to 0 after which learning stops and the continuation value is a convex combination of $\boldsymbol{v}^H$ and $\boldsymbol{v}^N$, or (2) the sender truthfully reveals the bad state too frequently, in which case his reputation $\eta(h^t)$ reaches 1 after which the continuation value is a convex combination of $\boldsymbol{v}^H$ and $\boldsymbol{v}^N$ such that:

1. Type-c_1 sender is indifferent between obtaining a perfect reputation by revealing the bad state truthfully (after which his continuation value is the above convex combination of $\boldsymbol{v}^H$ and $\boldsymbol{v}^N$) and milking his reputation by lying about the bad state,
2. All other types strictly prefer to milk his reputation by lying about the bad state.

In the active learning phase, $p^L(h^t)$ increases when the sender reveals the bad state truthfully except at histories where his reputation will reach 1 after doing so. Hence, there is a concern that p^L will become too large if the sender reveals the bad state too frequently, in which case his continuation value will escape V^*. Lemma D.0 alleviates this concern by showing that when the sender is patient, his reputation will reach 1 before his continuation value escapes V^*.

Then I state four lemmas, which imply that Class 2 histories and rebounding phase histories have negligible impact on the sender's continuation value as $\delta \to 1$. The proofs are in subsequent sections of this appendix. These lemmas together with Lemma D.0 imply that the sender's continuation value belongs to V^* at every history. Lemma D.1 establishes a lower bound on the receiver's posterior belief after observing message b in state b at any history that belongs to Class 2.

Lemma D.1. *For any history h^t belonging to Class 2,*

- *If $\underline{c}(h^t) \leq c_3$, then $\eta(h^t,(b,b)) \geq \eta(h^0)$ and $\eta(h^t,(b,g)) = 0$.*
- *If $\underline{c}(h^t) = c_2$, then $\eta(h^t,(b,b)) = \min\{1, \frac{\eta(h^t)}{1-\rho^*}\}$ and $\eta(h^t,(b,g)) = 0$.*

The next lemma establishes a uniform upper bound on the number histories that belong to Class 2 and the realized state in the previous period is b.

Lemma D.2. *There exist $\underline{\delta} \in (0,1)$ and $M \in \mathbb{N}$, such that when $\delta > \underline{\delta}$ and along every on-path play, the number of histories that belong to Class 2 while the state in the period before being b is no more than M.*

Lemma D.2 implies that for every on-path history h^t, the number of periods that belong to Class 2 in the reduced outcome $o(h^t)$ is no more than M. This upper bound does not depend on the sender's discount factor δ. Lemma D.2 implies that Class 2 histories have negligible impact on the sender's continuation value when δ is close enough to 1. Lemma D.3 establishes a uniform lower bound on $p^H(h^t)$ for all histories belonging to the active learning phase.

Lemma D.3. *There exist $\underline{\delta} \in (0,1)$ and $\underline{Q} > 0$, such that when $\delta > \underline{\delta}$, we have $p^H(h^t) \geq \underline{Q}$ for all h^t belonging to the active learning phase.*

Lemma D.3 leads to a lower bound on $p^H(h^t)$ if h^t is the *first history* that reaches the absorbing phase, i.e., h^t is such that $p^L(h^t) = 0$ and $p^L(h^s) > 0$ for all $h^s \prec h^t$.

Lemma D.4. *There exist $\underline{\delta} \in (0,1)$ and $K \in \mathbb{N}$, such that when $\delta > \underline{\delta}$ and along every on-path play, the number of histories that belong to the rebounding phase is at most K.*

Lemma D.4 implies that rebounding histories have negligible impact on the sender's continuation value as δ goes to 1. These lemmas imply that when the sender's continuation value is $\boldsymbol{u}(\rho)$ in period 0 and his continuation value at every subsequent history satisfies his incentive constraint and promise-keeping constraint, then his continuation value belongs to V^* at every history.

D.1. Proof of Lemma D.1

Case 1: Consider the case in which $\underline{c}(h^t) \leq c_3$. First, suppose $\eta(h^t) \geq \eta(h^0)$, then the conclusion of Lemma D.1 follows since $\eta(h^t,(b,b)) > \eta(h^t) \geq \eta(h^0)$. Second, suppose $\eta(h^t) < \eta(h^0)$, then the posterior probability with which the sender's lying cost is c_1 is bounded from below by:

$$\frac{\eta(h^t)}{\eta(h^t)+(1-\eta(h^t))\dfrac{\pi_2+...+\pi_{j-1}+\frac{k_j-l(h^t)-1}{k_j}\pi_j}{\pi_2+...+\pi_{j-1}+\frac{k_j-l(h^t)}{k_j}\pi_j}}$$

$$\geq \frac{\eta(h^t)}{\eta(h^t)+(1-\eta(h^t))\dfrac{\pi_2+...+\pi_{j-1}+\frac{k_j-1}{k_j}\pi_j}{\pi_2+...+\pi_{j-1}+\pi_j}}$$

Let

$$X \equiv 1-\frac{\pi_2+...+\pi_{j-1}+\frac{k_j-1}{k_j}\pi_j}{\pi_2+...+\pi_{j-1}+\pi_j}=\frac{\pi_j}{k_j(\pi_2+...+\pi_{j-1}+\pi_j)}.$$

The lower bound on posterior belief $\frac{\eta(h^t)}{\eta(h^t)+(1-\eta(h^t))(1-X)}$ is greater than π_1 if and only if:

$$X \geq 1-\frac{(1-\pi_1)\eta(h^t)}{\pi_1(1-\eta(h^t))}=\frac{\pi_1-\eta(h^t)}{\pi_1(1-\eta(h^t))}.$$

The above inequality is implied by (D.4) since $\eta(h^t) \geq \eta^*$ at every h^t that belongs to Class 2.

Case 2: Consider the case in which $\underline{c}(h^t)=c_2$. If $\eta(h^t) \geq 1-\rho^*$, then type-c_2 sender plays $\mathbf{a}^L$ with probability $\min\{1,\frac{\rho^*}{1-\eta(h^t)}\}=1$, which implies that $\eta(h^t,(b,b))=1$. If $\eta(h^t) < 1-\rho^*$, then type-c_2 sender lies in the bad state with probability $\min\{1,\frac{\rho^*}{1-\eta(h^t)}\}=\frac{\rho^*}{1-\eta(h^t)}$, which implies that $\eta(h^t,(b,b))=\eta(h^t)/(1-\rho^*) \geq (1-\rho^*)\eta(h^0)/(1-\rho^*)=\eta(h^0)$.

D.2. Proof of Lemma D.2

Step 1: If h^t belongs to Class 2 and $\underline{c}(h^t)=c_j \leq c_3$, then type-$\underline{c}(h^t)$ sender sends message g in state b with probability 1 when $l(h^t)=k_j-1$, after which play reaches the absorbing phase. Therefore, along every path of play, there are at most k_j Class 2 histories satisfying $\underline{c}(h^t)=c_j$ and the state in the previous period (i.e. period $t-1$) is b. This further implies that there are at most $K \equiv k_3+...+k_n$. Class 2 histories that has $\underline{c}(h^t) \leq c_3$ and the previous period state being b.

Step 2: Consider the number of Class 2 histories such that (1) $\underline{c}(h^t)=c_2$, and (2) the state in the previous period is b. Let $N \equiv \lceil \frac{1}{1-\gamma} \rceil$, and recall the integer constant T in Lemma D.1. In addition to the requirements on δ specified before, I need require δ to be large enough such that $\widehat{\delta}$ satisfies

$$\widehat{\delta}^{T+1}(1+\widehat{\delta}+...+\widehat{\delta}^N) > N \text{ and } 2\widehat{\delta}^{T+N+2} > 1. \tag{D.22}$$

First, I show that after the sender sends message b when the state is b at h^t, it takes at most $T+N$ such periods for play to reach a history that belongs to either to the absorbing phase or to another Class 2 history. According to the continuation value at $(h^t,(b,b))$, we have:

$$p^L(h^t,(b,b)) = \frac{p^L(h^t)}{\widehat{\delta}} < \frac{1-\widehat{\delta}}{\widehat{\delta}}. \tag{D.23}$$

The last inequality comes from h^t belonging to Class 2, so that $p^L(h^t) < 1-\widehat{\delta}$ by definition. According to Lemma D.1, for every Class 1 history h^s such that $h^s \succ (h^t,(b,b)) \equiv h^{t+1}$ and all histories between $(h^t,(b,b))$ and h^s belong to Class 1, we have:

$$(1-\widehat{\delta})\sum_{\tau=1}^{|o(h^s\backslash h^{t+1})|}\widehat{\delta}^{\tau-1}\mathbf{1}\{r_\tau=b\} \leq (1-\widehat{\delta}^T)+(1-\widehat{\delta})\sum_{\tau=1}^{|o(h^s\backslash h^{t+1})|}\widehat{\delta}^{\tau-1}\mathbf{1}\{r_\tau=g\}\cdot\frac{1-\widetilde{\rho}}{\widetilde{\rho}}. \tag{D.24}$$

Moreover, (D.23) and the requirement that all histories between $(h^t,(b,b))$ and h^s belong to Class 1 imply that

$$(1-\widehat{\delta})\sum_{\tau=1}^{|o(h^s\backslash h^{t+1})|}\widehat{\delta}^{\tau-1}\mathbf{1}\{r_\tau=g\} < \frac{1-\widehat{\delta}}{\widehat{\delta}}. \tag{D.25}$$

Given that only $(\mathbf{a}^L,\mathbf{b}^T)$ and $(\mathbf{a}^H,\mathbf{b}^T)$ occur at active learning phase histories (Class 1 and 2):

$$\begin{aligned}1-\widehat{\delta}^{|o(h^s\backslash h^{t+1})|-1} &= (1-\widehat{\delta})\sum_{\tau=1}^{|o(h^s\backslash h^{t+1})|}\widehat{\delta}^{\tau-1}\mathbf{1}\{r_\tau=g\}+(1-\widehat{\delta})\sum_{\tau=1}^{|o(h^s\backslash h^{t+1})|}\widehat{\delta}^{\tau-1}\mathbf{1}\{r_\tau=b\}\\ &\leq (1-\widehat{\delta}^T)+\frac{1-\widehat{\delta}}{\widehat{\delta}}+\frac{1-\widehat{\delta}}{\widehat{\delta}}\frac{1-\widetilde{\rho}}{\widetilde{\rho}} \leq (1-\widehat{\delta}^T)+\frac{1-\widehat{\delta}}{\widehat{\delta}\widetilde{\rho}} \leq (1-\widehat{\delta}^T)+\frac{1-\widehat{\delta}}{\widehat{\delta}\rho}\end{aligned} \tag{D.26}$$

Next, I show that $|o(h^s\backslash h^{t+1})| \leq T+N+1$. Suppose by way of contradiction that $|o(h^s\backslash h^{t+1})| \geq T+N+2$, then

$$(1-\widehat{\delta}^T)+\frac{1-\widehat{\delta}}{\widehat{\delta}}N \geq (1-\widehat{\delta}^T)+\frac{1-\widehat{\delta}}{\widehat{\delta}\rho} \geq 1-\widehat{\delta}^{|o(h^s\backslash h^{t+1})|-1} \geq 1-\widehat{\delta}^{T+N+1},$$

which yields $\frac{1-\widehat{\delta}}{\widehat{\delta}}N \geq \widehat{\delta}^T(1-\widehat{\delta}^{N+1})$. Dividing both sides by $\frac{1-\widehat{\delta}}{\widehat{\delta}}$, we have $N \geq \widehat{\delta}^{T+1}(1+\widehat{\delta}+...+\widehat{\delta}^N)$, which contradicts the first inequality of (D.22).

Second, I focus on history h^s that has the following two features:

1. h^s belongs to Class 2, and the state in period $s-1$ is l,
2. $h^s \succeq (h^t,(b,b))$ and all histories between $(h^t,(b,b))$ and h^s, excluding h^s, belong to Class 1.

I show that there exists at most one period from $(h^t,(b,b))$ to h^s such that the stage-game outcome is such that the sender recommends action g while the state is b. Suppose by way of contradiction that there exist two or more such periods, then $(1-\widehat{\delta})\sum_{\tau=1}^{|o(h^s\backslash h^{t+1})|}\widehat{\delta}^{\tau-1}\mathbf{1}\{r_\tau=g\}\geq 2(1-\widehat{\delta})\widehat{\delta}^{T+N+1}$. The last inequality comes from the previous conclusion that $|o(h^s\backslash h^{t+1})|\leq T+N+1$. According to (D.25),

$$2(1-\widehat{\delta})\widehat{\delta}^{T+N+1}<(1-\widehat{\delta})\sum_{\tau=1}^{|o(h^s\backslash h^{t+1})|}\widehat{\delta}^{\tau-1}\mathbf{1}\{r_\tau=g\}<\frac{1-\widehat{\delta}}{\widehat{\delta}}. \tag{D.27}$$

The above inequality contradicts the second inequality of (D.22) that $2\widehat{\delta}^{T+N+2}>1$.

Let h^t be the first time play reaches a history that belongs to Class 2 with $\underline{c}(h^t)=c_2$. According to Lemma D.1, $\eta(h^t,(b,b))\geq\frac{\eta^*}{1-\rho^*}\geq\eta(h^0)=\pi_1$. Let h^s be the next history that belongs to Class 2 with $\omega_{s-1}=b$. Since we have shown that the receiver takes the bad action at most once between $(h^t,(b,b))$ and h^s, we know that $\eta(h^s,(b,b))=\min\{1,\frac{\eta(h^s)}{1-\rho^*}\}\geq\min\{1,\frac{\eta(h^t,(b,b))}{1-\rho^*}(1-\lambda(1-\rho^*))\}$ Therefore, conditional on $(h^s,(b,b))$ not being an absorbing phase history, the receiver's belief at $(h^s,(b,b))$ assigns probability at least:

$$\eta(h^s,(b,b))\geq\eta(h^t,(b,b))\frac{1-\lambda(1-\rho^*)}{1-\rho^*}\geq\eta(h^t,(b,b))\sqrt{\frac{1}{1-\rho^*}} \tag{D.28}$$

to type c_1, where the last inequality comes from $\lambda\in(0,\frac{1-\sqrt{1-\rho^*}}{1-\rho^*})$. Let $\widehat{M}\equiv\frac{\log(1/\pi_1)}{\log\sqrt{\frac{1}{1-\rho^*}}}+1$. Since $\eta(h^t,(b,b))\geq\pi_1$ for the first Class 2 history h^t satisfying $\underline{c}(h^t)=c_2$, there can be at most $\widehat{M}$ Class 2 histories with c_2 being the highest-cost type along every path of play. This is because otherwise, the receiver's posterior belief assigns probability greater than $\pi_1\Big(\frac{1}{\sqrt{1-\rho^*}}\Big)^{\widehat{M}}>1$ at the $(\widehat{M}+1)$th such history, which leads to a contradiction. Summarizing the conclusions of the two parts, we know that along every path of equilibrium play, there exist at most $M\equiv K+\widehat{M}$ histories that belong to Class 2 and the state in the previous period is b.

D.3. *Proof of Lemma D.3*

Consider any given Class 2 history h^t such that no predecessor of h^t belongs to Class 2, in another word, all predecessors of h^t belong to Class 1 or the rebounding phase. Therefore, $p^H(h^{t-1})\geq Y$, which implies that $p^H(h^t)\geq Y-(1-\widehat{\delta})$. As a result $Q(h^t)=p^H(h^t)-\frac{1-\widehat{\delta}-p^L(h^t)}{\underline{c}(h^t)}\geq Y-(1-\widehat{\delta})>0$. If play remains at the active learning phase (Class 1 or 2) after h^t, then player 1 must be sending message b when the state is b at h^t, after which

$$p^H(h^t,(b,b))\geq p^H(h^t)-(1-\widehat{\delta})\geq Y-2(1-\widehat{\delta})\text{ and }p^L(h^t,(b,b))\leq\frac{1-\widehat{\delta}}{\widehat{\delta}}.$$

According to Lemma D.1, $\eta(h^t,(b,b))\geq\eta(h^0)=\pi_1$. One can then apply Lemma D.0 again, which implies that at every Class 1 history h^s such that only one predecessor of h^s (1) belongs to Class 2 and (2) $\omega_{s-1}=b$, we have $p^H(h^s)\geq Z\equiv Y-2(1-\widehat{\delta})-\frac{1-\widehat{\delta}}{\widehat{\delta}}\frac{1-\widetilde{\rho}}{\widetilde{\rho}}-(1-\widehat{\delta}^T)$. When $\widehat{\delta}$ is large enough, $Z\geq Y/2$. Therefore, for every Class 2 history h^s such that there is only one strict predecessor history belongs to Class 2, $Q(h^s)=p^H(h^s)-\frac{1-\widehat{\delta}-p^L(h^s)}{\underline{c}(h^s)}\geq Z-(1-\widehat{\delta})>0$. Iterate this process:

1. The number of Class 2 histories with the state in the previous period being b is bounded from above by M along every path of play.
2. For every Class 2 history h^t, $p^L(h^t,(b,b))=\frac{1-\widehat{\delta}}{\widehat{\delta}}$ and $\eta(h^t,(b,b))\geq\eta(h^0)$.

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

D.4. *Proof of Lemma D.4*

I construct a constant $K\in\mathbb{N}$ that is independent of the discount factor δ such that once play enters the rebounding phase, it will go back to the active learning phase after *at most* K periods with the realized state ω being b. First, type c_1's continuation value in the rebounding phase is at least 0. Second, play goes back to the active learning phase whenever his continuation value is above $(1-\widehat{\delta})(p+(1-c_1)(1-p))$. After K periods in the rebounding phase with the realized state being b, type c_1's continuation value is at least: $\frac{1-\widehat{\delta}^K}{\widehat{\delta}^K}c(1-p)$, which is more than $(1-\widehat{\delta})(p+(1-c_1)(1-p))$ if $K\geq\left\lceil\frac{p+(1-c_1)p}{c_1(1-p)^2}\right\rceil$. Lemma D.4 is obtained by setting δ to be close enough to 1 such that the corresponding $\widehat{\delta}^K$ satisfies the requirement for δ in Lemma D.0.

Appendix E. Proof of Proposition 4

Suppose by way of contradiction that there exists a type with lying cost $c_j\in\mathcal{C}$ who plays both $\mathbf{a}_1^H$ and $\mathbf{a}_1^L$ with positive probability at every on-path history. Then playing $\mathbf{a}_1^H$ at every on-path history and playing $\mathbf{a}_1^L$ at every on-path history are both type c_j's best replies to the receiver's equilibrium strategy. Since $\mathbf{a}_1^H$ is the sender's stage-game strategy that minimizes the expected lying cost, for every $c_i>c_j$, type c_i plays $\mathbf{a}_1^H$ with probability 1 at every on-path history.

I consider two cases. First, if $j\geq2$, then type-c_1 sender plays his honest strategy $\mathbf{a}_1^H$ with probability 1 at every on-path history. Therefore, type-c_2 sender separates from type-c_1 sender the first time he recommends action g in state b, after which he becomes the highest-cost type in the support of the receivers' posterior belief. According to Proposition 1, type c_2's continuation value is no more than p. As a result, type-c_2 sender's expected payoff in period 0 is no more than $(1-\delta)+\delta p$, which is strictly less than v_2^* when δ is close to 1. This contradicts the hypothesis that type-c_2 sender's equilibrium payoff is more than $v_2^*-\varepsilon$.

Second, if $j=1$, then type-c_1 sender finds it optimal to play $\mathbf{a}_1^L$ in every period. Since the sender's equilibrium payoff is within ε of v^*, type-c_1 sender's payoff is at least $v_1^*-\varepsilon$ by playing $\mathbf{a}_1^L$ in every period, and type-c_2 sender's payoff from doing so must be no more than $v_2^*+\varepsilon$. Since $p<1/2$, the receiver's stage-game strategy of *choosing $y_t=g$ regardless of the sender's message* is strictly suboptimal, and hence, it cannot be played at any on-path history. Among the remaining three receiver stage-game strategies, the sender's stage-game payoff is $1-(1-p)c$ under $(\mathbf{a}_1^L,\mathbf{a}_2^T)$, is $-(1-p)c$ under $(\mathbf{a}_1^L,\mathbf{a}_2^N)$ and $(\mathbf{a}_1^L,\mathbf{a}_2^O)$, where

$$\mathbf{a}_2^O(r)\equiv\begin{cases} g & \text{if } r=b\\ b & \text{if } r=g.\end{cases}$$

Let Q_L be the discounted probability of the stage-game strategy profile $(\mathbf{a}_1^L,\mathbf{a}_2^T)$ when the sender plays $\mathbf{a}_1^L$ in every period and the receiver plays according to his equilibrium strategy σ_2. Type

c_1's equilibrium payoff is $Q_L\Big(1-(1-p)c_1\Big)-(1-Q_L)(1-p)c_1$. Since type-$c_1$ sender's equilibrium payoff is more than $p-\varepsilon$, we have:

$$Q_L \geq p+(1-p)c_1-\varepsilon. \tag{E.1}$$

Type-c_2 sender's payoff from playing $\mathbf{a}_1^L$ in every period is $Q_L\big(1-(1-p)c_2\big)-(1-Q_L)(1-p)c_2$, which is at least $p+(1-p)(c_1-c_2)-\varepsilon$ by inequality (E.1). Since $p<1/2$, the lower bound on type-c_2 sender's equilibrium payoff is strictly greater than $v_2^*+\varepsilon$. However, type-c_2 sender's equilibrium payoff cannot exceed $v_2^*+\varepsilon$ when δ is close enough to 1. This leads to a contradiction.

References

Aumann, Robert, Maschler, Michael, 1965. Repeated Games with Incomplete Information. Working Paper.

Benabou, Roland, Laroque, Guy, 1992. Using privileged information to manipulate markets: insiders, gurus, and credibility. Q. J. Econ. 107 (3), 921–958.

Best, James, Quigley, Daniel, 2022. Persuasion in the Long Run. Working Paper.

Chen, Ying, Kartik, Navin, Sobel, Joel, 2008. Selecting cheap talk equilibria. Econometrica 76 (1), 117–136.

Crawford, Vincent, Sobel, Joel, 1982. Strategic information transmission. Econometrica 50 (6), 1431–1451.

Cripps, Martin, Mailath, George, Sameulson, Larry, 2004. Imperfect monitoring and impermanent reputations. Econometrica 72 (2), 407–432.

Cripps, Martin, Thomas, Jonathan, 2003. Some asymptotic results in discounted repeated games of one-sided incomplete information. Math. Oper. Res. 28, 433–462.

Fudenberg, Drew, Kreps, David, Maskin, Eric, 1990. Repeated games with long-run and short-run players. Rev. Econ. Stud. 57 (4), 555–573.

Fudenberg, Drew, Levine, David, 1992. Maintaining a reputation when strategies are imperfectly observed. Rev. Econ. Stud. 59 (3), 561–579.

Fudenberg, Drew, Maskin, Eric, 1991. On the dispensability of public randomization in discounted repeated games. J. Econ. Theory 53 (2), 428–438.

Fudenberg, Drew, Tirole, Jean, 1991. Perfect Bayesian equilibrium and sequential equilibrium. J. Econ. Theory 53, 236–260.

Gneezy, Uri, 2005. Deception: the role of consequences. Am. Econ. Rev. 95 (1), 384–394.

Gossner, Olivier, 2011. Simple bounds on the value of a reputation. Econometrica 79 (5), 1627–1641.

Green, Jerry, Laffont, Jean-Jacques, 1986. Partially verifiable information and mechanism design. Rev. Econ. Stud. 53 (3), 447–456.

Guo, Yingni, Shmaya, Eran, 2021. Costly miscalibration. Theor. Econ. 16, 477–506.

Hart, Sergiu, 1985. Nonzero-sum two-person repeated games with incomplete information. Math. Oper. Res. 10 (1), 117–153.

Kamenica, Emir, Gentzkow, Matthew, 2011. Bayesian persuasion. Am. Econ. Rev. 101 (6), 2590–2615.

Kartik, Navin, 2009. Strategic communication with lying costs. Rev. Econ. Stud. 76 (4), 1359–1395.

Kartik, Navin, Ottaviani, Marco, Squintani, Francesco, 2007. Credulity, lies, and costly talk. J. Econ. Theory 134 (1), 93–116.

Kuvalekar, Aditya, Lipnowski, Elliot, Ramos, João, 2022. Goodwill in communication. J. Econ. Theory 203, 105467.

Lipnowski, Elliot, Doron, Ravid, Shishkin, Denis, 2022. Persuasion via weak institutions. J. Polit. Econ. 130 (10), 2705–2730. https://doi.org/10.1086/720462.

Margaria, Chiara, Smolin, Alex, 2018. Dynamic communication with biased senders. Games Econ. Behav. 110, 330–339.

Mathevet, Laurent, Pearce, David, Stacchetti, Ennio, 2022. Reputation and Information Design. Working Paper.

Meng, Delong, 2021. On the value of repetition for communication games. Games Econ. Behav. 127, 227–246.

Nguyen, Anh, Tan, Teck Yong, 2021. Bayesian persuasion with costly messages. J. Econ. Theory 193, 105212.

Ortner, Juan, 2015. Direct implementation with minimally honest individuals. Games Econ. Behav. 90, 1–16.

Pei, Harry, 2021. Trust and betrayals: reputational payoffs and behaviors without commitment. Theor. Econ. 16, 449–475.

Pei, Harry, 2022. Repeated communication with private lying cost. arXiv:2006.08069.

Pęski, Marcin, 2014. Repeated games with incomplete information and discounting. Theor. Econ. 9, 651–694.

Renault, Jérôme, Solan, Eilon, Vieille, Nicolas, 2013. Dynamic sender-receiver games. J. Econ. Theory 148 (2), 502–534.

Shalev, Jonathan, 1994. Nonzero-sum two-person repeated games with incomplete information and known-own payoffs. Games Econ. Behav. 7 (2), 246–259.
Sobel, Joel, 1985. A theory of credibility. Rev. Econ. Stud. 52 (4), 557–573.
Sobel, Joel, 2020. Lying and deception in games. J. Polit. Econ. 128 (3), 907–947.
Stokey, Nancy, Lucas, Robert E., Prescott, Edward, 1989. Recursive Methods in Economic Dynamics. Harvard University Press.
Titova, Maria, 2022. Persuasion with Verifiable Information. Working Paper.