Cheap Talk with Costly Verification

Benjamin Shaver*

June 2, 2025

Abstract

I study a model of cheap talk between a biased sender and receiver. The sender, who is informed about the state of the world, communicates with a receiver, after which the receiver decides whether to pay a cost, the size of which is private information, to learn whether the sender's message is true. I show that any influential equilibrium is characterized by a threshold. If the state is above the threshold, the sender tells the truth, and if the state is below the threshold, the sender lies by pretending to be a truth-teller. In response, the receiver verifies all messages above the threshold with positive probability. I then show the receiver's ability to verify has two effects: an informational effect and a deterrence effect. Moreover, the deterrence effect can increase the expected informational effect of verification.

*Ph.D. Candidate, Harris School of Public Policy and Department of Political Science, The University of Chicago, blshaver@uchicago.edu

1 Introduction

During the COVID-19 Pandemic, Vice President Mike Pence, speaking on behalf of the administration's Coronavirus Task Force, blamed the increasing numbers of COVID-19 cases on increased testing:

> "...[w]e want the American people to understand it's almost inarguable that more testing is generating more cases. To one extent or another the volume of new cases coming in is a reflection of a great success in expanding testing across the country."[1]

Pence offered no evidence to support his statement, yet wanted voters not to blame the administration for the increased number of cases. Seminal work on cheap talk shows that when a sufficiently biased sender communicates with a receiver using cheap talk, no influential communication can arise (Crawford and Sobel, 1982). Yet, politics is full of instances like Vice President Pence's COVID-19 statement where a biased sender engages in cheap talk: a politician who wants voters to have a high opinion of her explains why her policy failed, a lobbyist who wants a legislator to support her policies attempts to convince him of the policy's merits, a candidate for office seeks to convince voters of her qualifications, etc. Are there any features of these instances that allow for influential communication?

In this paper, I explore one such feature: the endogenous choice of the receiver—the audience of the sender's message—to verify its truthfulness. A voter interested in assessing the veracity of Vice President Pence's statement about COVID-19 testing could have found and read Linda Qui's—a fact-checking reporter at *The New York Times*—article in which she deemed the statement "false," writing:

> "Ramped up testing alone does not account for the uptick in cases. Rather, the virus's spread is generating more cases..."[2]

[1] Pence (2020)

[2] Qui (2020). She supported her verdict by quoting Dr. Robert R. Redfield, the director of the Centers for Disease Control and Prevention, who had recently said "Several communities are seeing increased cases driven by multiple factors, including increased testing, outbreaks and evidence of community transmission."

Relatedly, a legislator who is told by a lobbyist that studies show similar policies have been successful elsewhere can seek out the studies himself to confirm whether the lobbyist's statement is true, and a voter who a candidate tells of her past successes in office can research this himself. Across these examples, there are three unifying features: (*i.*) the receiver can verify the truthfulness of the sender's message (e.g., he can read Linda Qui's article deeming Vice President Pence's statement false), (*ii.*) verification speaks to the truthfulness of the sender's message but not the state, (e.g., the article did not reveal the responsibility of the administration for increased COVID-19 cases), and (*iii.*) the sender does not know how costly it will be for the receiver to verify her message (e.g., Vice President Pence did not know which and how many reporters would fact-check him).

In this paper, I incorporate these features into a model of cheap talk communication. At the beginning of the game, a sender (she) learns the state of the world. Then, she communicates with a receiver (he) using cheap talk. After observing the sender's message, the receiver learns how costly it is to verify the message and chooses whether to pay the cost. If he does, he learns whether the message is true but does not learn the state of the world. If he does not, he learns nothing about the message's veracity. Then, the receiver chooses an action that is payoff relevant to both players.

In my first result, I show that any influential equilibrium of the model has a straightforward structure characterized by a threshold, the sender's expected utility in equilibrium when she lies. When the state is above this threshold, the sender tells the truth, and when the state is below this threshold, she lies by pretending to be a high type. Aside from one knife-edge case, any message the sender sends truthfully, she also sends as a lie, and the receiver verifies all messages in the support of the sender's strategy with positive probability.

This equilibrium structure reveals a dependence between endogenous verification and lying. If the receiver believes a particular message in the support of the sender's strategy is only sent truthfully, he has no incentive to verify it. But, as described above, in any influential equilibrium, the messages in the support of the sender's strategy are all above a threshold, the sender's expected utility in equilibrium when she lies. So if there is a message in the support of the sender's strategy that is only sent truthfully, a lying sender has a profitable deviation. This dependence illustrates a key difference between endogenous verification, studied in this model, and exogenous verification, studied elsewhere (e.g., Dziuda and Salas, 2018), and illustrates how the usefulness of endogenous verification relies on the accurate expectation of lying.

I then solve for all influential equilibria when there are three equally likely states and the cost of verification is uniform. If the intermediate state is closer to the highest state than the lowest state or the intermediate state is closer to the lowest state and the upper bound of verification costs is sufficiently small, in any influential equilibrium, the lowest type of sender lies by mixing between the two highest messages and the higher two types of sender tell the truth. Otherwise, in any IE, the lower two types of sender lie and the highest type tells the truth.

I then show that as the upper bound of verification costs approaches infinity, the probability of verification in equilibrium approaches zero, and the receiver's belief when he does not verify approaches his prior. On the other hand, as this upper bound approaches zero, the probability of verification approaches one, but the low-type sender continues to lie in equilibrium. This is due to the dependence between endogenous verification and lying. If the low-type sender stopped lying, the receiver would never verify, but this would provide a profitable deviation. Hence, when the upper bound of verification costs approaches infinity, my model approximates traditional cheap talk. However, when it approaches zero, there is still lying in equilibrium.

Analysis of the three-state case illustrates that there are two effects of verification. The first is the direct *informational* effect, which is the difference between the variance of the receiver's belief when he does not verify and the expected variance of his belief when he does. The second is an indirect deterrence effect of verification: verification can deter the intermediate type from lying due to an endogenous cost of lying that emerges in equilibrium.

Together, these effects mean that in any influential equilibrium, the receiver's expected utility is higher than what he obtains without this verification technology.

One might conjecture that the informational and deterrence effects act like substitutes: the greater the deterrence effect, the smaller the informational effect. However, I show this intuition is not always correct. In particular, in some cases, the expected informational effect of verification is greater when verification has a deterrence effect than when it does not.

I conclude by analyzing the sender's lying strategy when she mixes between messages when lying. In particular, I show that when the upper bound of verification costs is large, increasing the intermediate state leads the sender to lie more with the intermediate message and that when the sender often lies with the highest message, increasing the upper bound of verification costs leads her to lie more with the highest message.

2 Related Literature

This paper examines a setting where influential cheap talk communication occurs between a sender and receiver despite the sender's bias. In the seminal paper on cheap talk, Crawford and Sobel (1982) show that when the sender is sufficiently biased, no information transmission occurs in equilibrium. This paper joins others on communication with detectable lying in cheap talk in showing that if the receiver can detect a lie sent by the sender—either with exogenous probability (e.g., Dziuda and Salas, 2018; Holm, 2010) or, as in this paper, through an endogenous choice by the receiver to verify the sender's message (e.g., Sadakane and Tam, 2023; Levkun, 2022; Ball and Gao, 2025)—influential communication arises despite the sender's bias.[3]

In Sadakane and Tam (2023), the paper closest to mine, the receiver chooses whether to incur a publicly known cost to inspect a message sent by a privately informed sender. In my model, less punishing off the equilibrium path beliefs mean that in any influential

[3]Less related to this model, others study exogenously detected deceit in other communication frameworks like Bayesian persuasion (e.g., Ederer and Min, 2022; Venkatesh et al., 2025) or detectable lying when the sender's bias is limited (e.g., Balbuzanov, 2019)

equilibrium, except in one knife-edge case, all messages in the support of the sender's strategy are sent truthfully and as lies and are verified with positive probability. This is not the case in Sadakane and Tam (2023), where, if the state is uniformly distributed, some messages sent in equilibrium are not verified. In both models, the receiver's ability to verify the sender's message has two effects: an informational effect, which reveals whether the sender's message is true, and a deterrence effect, which deters intermediate types of senders from lying. However, in Sadakane and Tam (2023), the cost of verification completely offsets the informational effect. In contrast, when the receiver has private information about his cost, the informational effect exceeds the ex-ante expected cost of verification. Moreover, because the cost in my model does not offset the informational effect, I show that, in some cases, the deterrence effect can increase the expected informational effect.

My model is also related to Dziuda and Salas (2018), in which a biased sender communicates with a receiver who learns whether the sender's message is true with an exogenous probability. In contrast, in my model, the receiver endogenously chooses whether to verify the sender's message. This choice produces a distinct equilibrium structure. In Dziuda and Salas (2018), some intermediate messages are only sent truthfully, whereas, in my model, all messages in the support of the sender's strategy are sent truthfully and as lies.[4] When verification is exogenous, messages that are only sent truthfully on the equilibrium path are still verified; this prevents the sender from deviating to lying with one of these messages. When verification is endogenous, messages only sent truthfully on the equilibrium path are never verified, giving the sender a profitable deviation. Hence, my model illustrates a dependence between the usefulness of verification and the expectation of lying.

This paper is also related to models of cheap talk with an exogenous cost of lying. This cost might arise due to a psychological cost associated with lying (e.g., Kartik et al., 2007; Minozzi and Woon, 2013) or the existence of receivers who take the sender's message at face value (e.g., Kartik et al., 2007).[5] In both cases, deception emerges in equilibrium, but the

[4]Except in a knife-edge case discussed below.

[5]To a lesser degree, my model is related to a broader literature of strategic communication where lying

cost of lying constrains the sender to the extent that information can be conveyed despite disagreement between the sender and receiver(s). In my model, the receiver does not pay an exogenous cost to lie. Yet, I show that in equilibrium, an endogenous cost of lying emerges: because verification does not reveal the state, an intermediate-type sender caught lying will be pooled with the lower types who also lie. This cost is the source of the deterrence effect of verification.

Finally, in its applications, this paper is related to the literatures on fact-checking and informational lobbying. The former literature is primarily empirical (e.g., Weeks and Garrett, 2014; Weeks, 2015; Nyhan and Reifler, 2010), but includes some theoretical work (e.g., Levkun, 2022). Of particular relevance to my paper are Nyhan and Reifler (2015) and Lim (2018), who document empirical evidence of a deterrence effect of fact-checking on politicians' behavior, and Gottfried et al. (2013) and Pingree et al. (2014), who document empirical evidence of an informational effect of fact-checking. The latter literature encompasses theoretical (e.g., Ellis and Groll, 2020) and empirical work (e.g, Hojnacki and Kimball, 1998). My paper is closest to Potters and Van Winden (1992) and Austen-Smith and Wright (1992), which analyze settings in which a biased lobbyist is able to convey her private information to a politician who recognizes the lobbyist is biased. However, my setting does not rely on information acquisition or signaling being costly.

3 Model

There is a sender (S,"she") and a receiver (R, "he"). At the beginning of the game, the sender privately learns the state of the world, her "type", $\theta \in \Theta$, where Θ is a finite subset of $[0, 1]$, each $\theta \in \Theta$ occurs with probability $h(\theta) > 0$, $N = |\Theta|$, and $\theta \in \Theta$ are indexed such that $\theta_1 < ... < \theta_N$. The sender communicates with the receiver by choosing a message m from the message space $M = \theta$. If $m = \theta$, the message is "true"—the sender tells the

is costly (e.g., Nguyen and Tan, 2021; Guo and Shmaya, 2019).

truth—and if $m \neq \theta$, the message is "false"—the sender lies.[6]

After observing the sender's message, the receiver decides whether to verify it ($c = 1$) or not ($c = 0$). Verification is costly; if the receiver verifies, he pays the cost $\beta\kappa$ where $\kappa \in [0, 1]$ is private information he learns at the start of the game and $\beta > 0$ is a parameter observed by both players. I assume κ has a prior distribution g and $g(\kappa) > 0$ for all $\kappa \in [0, 1]$. If the receiver verifies the sender's message, he learns whether it is true ($v = t$) or false ($v = f$) but does not learn the state. If he does not verify, he learns nothing ($v = \emptyset$). Finally, the receiver selects an action $a \in [0, 1]$.

Preferences. The receiver has a utility function $u_R(\kappa, \theta, a) = -(\theta - a)^2 - c \cdot \beta\kappa$, and the sender has a utility function $u_S(a) = a$. Note, the receiver's payoff depends on the state—in particular, he wants to choose the action that matches the state—while the sender wants the receiver to choose $a = 1$.

Equilibrium. A (mixed) strategy for the sender is a probability function $\sigma(\cdot|\theta) : \Theta \to \Delta M$.[7] A message m is "on the equilibrium path" or in "the support of" the sender's strategy if $\sigma(m|\theta) > 0$ for at least one θ, and is "off the equilibrium path" otherwise. A strategy for the receiver is a tuple (c, a) where $c : K \times M \to \{0, 1\}$ and $a : M \times \{t, f, \emptyset\} \to [0, 1]$.[8] The game is solved for perfect Bayesian equilibrium (henceforth, an "equilibrium"), where a PBE is a triple (σ, c, a) and a belief assessment such that:

1. If $m \in M$ is in the support of $\sigma(m|\theta)$, m maximizes the sender's expected utility taking (c, a) as given.

2. If (c, a) is chosen, c maximizes the receiver's expected utility given $\mathbb{E}[\theta|m]$ and a maximizes the receiver's expected utility given $\mu_2 = \mathbb{E}[\theta|(m, v)]$.

[6] Since the sender is restricted to choosing a single message rather than allowed to choose a subset of the message space (e.g., $m = \{\theta_j, \theta_k\}$), these definitions are equivalent to the definitions in Sobel (2020).

[7] $\Delta(X)$ denotes the space of lotteries over X.

[8] I restrict attention to equilibria where the receiver plays a pure strategy. This is without loss of generality as the receiver has a unique best response to any posterior belief about the sender's type due to the strict concavity of his utility function.

3. Beliefs μ_1 and μ_2 are generated using Bayes' rule on the equilibrium path.

Throughout the paper, I focus on influential equilibria (henceforth, IEs).

Definition 1. *An equilibrium is influential if $a^*(m, v)$ is not constant on the equilibrium path (Sobel, 2013).*

Assumptions. As in many cheap talk games, the set of equilibria is large. In light of this, I make the following assumptions to refine the set of equilibria.

Assumption 1. *When lying, the sender's strategy is independent of θ.*

In equilibrium, when the sender lies, she must be indifferent between any lies within the support of her strategy. Therefore, an equilibrium where the sender conditions her lying strategy on θ requires different types of senders to use different strategies to lie even though each strategy yields the same expected utility. To simplify the analysis of the game, I assume the sender's strategy is independent of θ when she lies.

Assumption 2. *If there is a message m that is not in the support of the sender's strategy, then upon seeing m, the receiver believes $\theta = m$ and chooses $a = m$.*

This assumption can be justified by assuming that if a message is off the equilibrium path, there is an infinitesimally small probability the truth is accidentally revealed by the sender through a slip of the tongue (e.g., Hart et al., 2017).

Discussion of the Model. This paper studies a model of cheap talk communication between an informed but biased sender and receiver with three additional features: (*i.*) the receiver has access to a verification technology, (*ii.*) the verification technology reveals whether the sender's message is true but not state, and (*iii.*) the sender does not know the receiver's verification cost when she chooses her message.

<u>Communication Between a Politician and a Voter</u>

This model applies naturally to communication between a politician, who wants to avoid blame for a policy failure, and a voter, who wants to accurately assign blame to the politician, when (*i.*) journalists will issue fact-checking reports of the politician's statement, which the voter can access, (*ii.*) the fact-checking reports assess whether the politicians' statement is true but not her culpability for the policy failure, and (*iii.*) the politician does not know which journalists will fact-check her statement, which determines how costly it is for the voter to verify.

In this application, β, the parameter in the receiver's cost of verification that scales the realization κ, represents features of the political or media environment that affect the cost to access fact-checked information. For instance, larger values of β might represent a situation where few media organizations fact-check politicians' statements. Or, larger values of β might represent a situation where the politician is unlikely to be fact-checked because she is a local politician.

Informational Lobbying

This model also applies to informational lobbying of a legislator by a non-allied lobbyist. Suppose a lobbyist wants to convince a legislator to support an energy policy, and the legislator wants to support the policy if and only if it will be successful. As part of the lobbyist's efforts, she informs the legislator that studies have shown similar policies to be successful when implemented elsewhere. In this example, (*i.*) the legislator can verify whether the lobbyist's message is true by finding and reading the studies, (*ii.*) this will reveal whether the lobbyist told the truth but now whether the policy will be successful, and (*iii.*) the lobbyist does not know the legislator's cost to verify because she does not know the legislator's time and resource constraints.

In this application, β might vary with the legislator's resources. For instance, given the difference in staff sizes, β might be larger for a member of the House of Representatives than a member of the Senate.

4 Analysis

4.1 Equilibrium Structure

I begin by considering the receiver's action as a function of the message he observes and any additional information learned from verification. In the final action of the game, he solves:

$$\max_{a \in [0,1]} -\mathbb{E}[(\theta - a)^2 | (m, v)].$$

This expression is uniquely maximized when $a^*(m, v) = \mathbb{E}[\theta | (m, v)]$.

Before choosing his action, the receiver decides whether to verify the sender's message. The optimality of $a^*(m, v)$ and the fact that the receiver has quadratic loss utility means he will verify the sender's message if:

$$\begin{aligned} & Var(m|c=0) - \phi(m)Var(m|c=1, v=t) - (1-\phi(m))Var(m|c=1, v=f) \\ = & \underbrace{Var(m|c=0) - (1-\phi(m))Var(m|c=1, v=f)}_{\textbf{Informational Effect of Verification: } \Lambda(m)} > \beta\kappa, \end{aligned} \tag{1}$$

where $Var(m|\cdot)$ is the variance of the receiver's belief, $\phi(m)$ is the receiver's conjecture about the probability message m is sent truthfully, and the second line follows from the fact that $Var(m|c = 1, v = t) = 0$. The left-hand side of (1) represents the informational effect of verification, which is the difference between the variance in the receiver's belief if he does not verify and his expected variance if he does. I denote this $\Lambda(m)$. Intuitively, for the receiver to verify the sender's message, the informational effect must exceed the cost. From (1), it is clear the receiver will not verify a message he believes is only sent truthfully, nor will he verify a message he believes is only sent as a lie since $\Lambda(m) = 0$ in both cases.

In any IE, some types of sender must lie. To see why, suppose this is not the case. Then, the receiver believes all of the sender's messages without verifying them. However, if this is the case, a low-type sender has a profitable deviation to sending a higher message. Hence, in any IE, at least one type of sender lies and receives an expected payoff u^{lie}, which must be independent of θ by Assumption 1.

In the following proposition, I provide a characterization result for any IE.

Proposition 1. *In any IE, the sender tells the truth if $\theta > u^{lie}$ and lies if $\theta < u^{lie}$; when the sender lies, she lies by randomizing over each message $m > u^{lie}$ with positive probability; and the receiver verifies all messages $m > u^{lie}$ with positive probability.*

Proposition 1 states that any IE has a threshold structure: the sender tells the truth when she has a high type and lies when she has a low type by randomizing over messages that are sent truthfully, and the receiver verifies all messages $m > u_{lie}$ with positive probability.[9] Moreover, except in the knife-edge case where $\theta = u^{lie}$, any message that is sent truthfully is sent as a lie and verified with positive probability.

The intuition for this structure is as follows. As discussed above, in any IE, at least one type of sender must lie. Moreover, all messages $m \neq u^{lie}$ that are sent truthfully are also sent as lies.[10] Suppose not. Then, there is a message m that is only sent truthfully. The receiver will not verify m and will choose $a(m, \emptyset) = m$. Either $m > u^{lie}$ or $m < u^{lie}$. The sender can deviate from lying to sending m in the former. In the latter, the sender can deviate from telling the truth with m to lying.

Consider the sender with type $\theta > u^{lie}$ and suppose she lies. As a result, she receives an expected payoff of u^{lie}. Since the type-θ sender lies, $m = \theta$ must be off the equilibrium path. This is because Assumption 1 implies all senders who lie use the same strategy, meaning $m = \theta$ cannot be sent as a lie. Since $m = \theta$ is off the equilibrium path, Assumption 2 implies the sender can deviate from lying to reporting $m = \theta$, which will be believed. Thus, a sender of type $\theta > u^{lie}$ tells the truth.

Consider the sender with type $\theta < u^{lie}$, and suppose she tells the truth. When the

[9]Proposition 1 provides a characterization, not existence. Without more structure on Θ, it is difficult to derive a necessary and sufficient condition for existence. However, in Section 4.2, I solve for all IE for a particular Θ. Moreover, in the Supplementary Appendix, I provide sufficient conditions for the existence of some IEs.

[10]In the Appendix, I address the knife-edge case where $\theta = u^{lie}$.

receiver verifies $m = \theta$, he chooses $a(m,t) = \theta < u^{lie}$. As discussed above, it must be that $m = \theta$ is also sent as a lie. This implies the action the receiver takes if he does not verify $m = \theta$ is also smaller than u^{lie} since the sender tells the truth if $\theta > u^{lie}$. Hence, the sender with type $\theta < u^{lie}$ who tells the truth has a profitable deviation.

Equilibrium Structure The structure of any IE in this model is similar to the structure of equilibria in other models where the fact that the receiver can detect the sender's lies through exogenous verification (e.g., Dziuda and Salas, 2018) or endogenous verification (e.g., Sadakane and Tam, 2023) allows for influential communication despite the sender's bias. Across these models, high types tell the truth, and low types lie by pretending to be high types. Yet in Dziuda and Salas (2018) and Sadakane and Tam (2023) not all messages that are sent truthfully are sent as lies. In contrast, besides in the knife-edge case where there is a θ such that $\theta = u^{lie}$, all messages sent truthfully are also sent as lies.

Relative to Dziuda and Salas (2018), this distinction emerges because of the difference between exogenous and endogenous verification. If the receiver believes a message is only sent truthfully, he will not pay the cost to verify it when verification is endogenous. However, if he never verifies a message that is in in the support of the sender's strategy, the sender can deviate to that message without fear of getting caught. Hence, in equilibrium, the receiver must verify all messages in the support of the sender's strategy that would yield a higher payoff to a lying sender if believed than u^{lie}. Yet, the receiver cannot commit to verifying every message; he will only verify if doing so is sequentially rational. This requires that each message the sender sends truthfully she also sends as a lie. In contrast, if a message is stochastically verified, it will be verified with positive probability even if it is only sent truthfully on the equilibrium path. This constrains the behavior of the sender. Thus, this comparison illustrates a sense in which successful verification requires suspicion to be effective—otherwise, the receiver will not verify the sender's message.

Relative to Sadakane and Tam (2023), the distinction emerges due to assumed off the equilibrium path beliefs. To see this, suppose that in my model, like Sadakane and Tam (2023), the sender can send a message $m = [\theta_j, \theta_k]$ with the meaning "the state is between" θ_j and θ_k. If $m \in [\theta_j, \theta_k]$ this message is truthful and if $m \notin [\theta_j, \theta_k]$ this message is a lie. Furthermore, consider the following example:

Example 1. $\Theta = \{0, \frac{1}{32}, \frac{3}{32}, 1\}$, *$h(\theta) = \frac{1}{4}$ for all θ, $\beta = 1$, $G = \mathcal{U}[0,1]$, the sender sends $m = 1$ if $\theta = 1$, $m = [\frac{1}{32}, \frac{3}{32}]$ if $\theta \in \{\frac{1}{32}, \frac{3}{32}\}$, and lies by reporting $m = 1$ if $\theta = 0$.*

When the receiver observes $m = 1$, he verifies when $\kappa < \frac{1}{4}$, and he never verifies $m = [\frac{1}{32}, \frac{3}{32}]$. Given such a strategy, the type-0 sender and the senders of types $\frac{1}{32}$ and $\frac{3}{32}$ are indifferent between sending $m = 1$ and $m = [\frac{1}{32}, \frac{3}{32}]$ because both messages yield an expected payoff of $\frac{1}{8}$. However, among other profitable deviations, the type-$\frac{3}{32}$ sender can deviate to $m = \frac{3}{32}$, which will be believed. In contrast, Sadakane and Tam (2023) assume the receiver believes a sender who deviates is the lowest type. This punishing belief can support this behavior in equilibrium.

4.2 Three-State Example

Proposition 1 characterizes any IE. In this section, I focus on the case where $\Theta = \{0, \theta_2, 1\}$ with $\theta_2 \in (0,1)$, $h(\theta) = \frac{1}{3}$ for all θ, and $G = \mathcal{U}[0,1]$.[11] This case maps well into the example of informational lobbying about a policy that may have a negative ($\theta = 0$), middling ($\theta = \theta_2$) or positive impact ($\theta = 1$) on a politician's—the receiver's—district. A lobbyist—the sender—who wants the politician to support the policy, communicates with the politician about studies that show the policy's impact when implemented elsewhere. After observing the lobbyist's message, the politician can verify it by researching whether the studies the lobbyist cited support the claim she made.

The following proposition establishes all IEs given these preliminaries.

[11] Note, in the traditional cheap talk model with these preliminaries, there is no influential equilibrium. To see this, suppose not. Then there must be m and m' such that $a(m) \neq a(m')$. But then $a(m) > a(m')$ or $a(m) < a(m')$, which implies there is a profitable deviation.

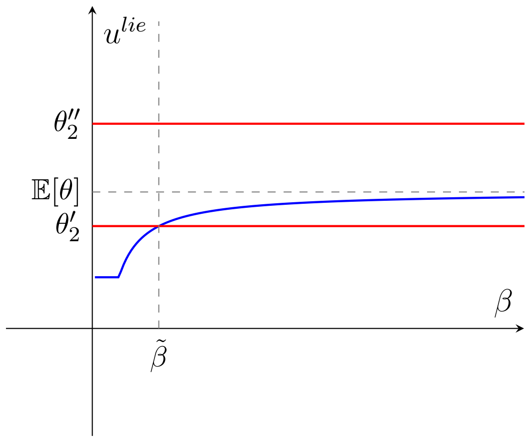

Figure 1: Expected utility from lying with $m = 1$ and from deviating to $m = \theta_2$, which is off the equilibrium path, for the sender of types $\theta \in \{0, \theta_2\}$. Assumes $\theta_2' = \frac{1}{3}$ and $\theta_2'' = \frac{2}{3}$.

Proposition 2. *Define $\tilde{\beta}$ and $\overline{\beta}$ as in (3) and (5). An IE exists in all regions of the parameter space, and any IE is of one of the following. When $\theta_2 < \frac{1}{2}$*

(a.) and $\beta > \tilde{\beta}$, an IE exists where the sender tells the truth when $\theta = 1$ and lies when $\theta \in \{0, \theta_2\}$ by sending $m = 1$;

(b.) and $\beta < \overline{\beta}$, at least one IE exists where the sender tells the truth when $\theta \in \{\theta_2, 1\}$ and lies when $\theta = 0$ by randomizing between $m = 1$ with probability σ^* and $m = \theta_2$ with probability $1 - \sigma^*$.

And when $\theta_2 \geq \frac{1}{2}$,

(c.) at least one IE exists where the sender tells the truth when $\theta \in \{\theta_2, 1\}$ and lies when $\theta = 0$ by randomizing between $m = 1$ with probability σ^* and $m = \theta_2$ with probability $1 - \sigma^*$.

Figure 1 provides an intuition for Proposition 2(a) by depicting the type-0 and type-θ_2 sender's expected utility from lying with $m = 1$ as a function of β. This expected utility is increasing in β since the probability the receiver verifies is decreasing in β, and a lower probability of verification means a higher probability of successfully pooling with the type-1 sender. For this type of IE to exist, this expected utility must be weakly larger than θ_2, which is the payoff that can be obtained by deviating to $m = \theta_2$, which is off the equilibrium path. When $\theta_2 < \frac{1}{2}$ (e.g., $\theta_2 = \theta_2'$) and β is sufficiently large—larger than $\tilde{\beta}$—this type of IE exists. Additionally, as β approaches infinity, the probability of verification approaches zero, and u^{lie} approaches the payoff the sender achieves in the traditional cheap talk model with these preliminaries, $\mathbb{E}[\theta]$. However, when $\theta_2 < \frac{1}{2}$ (e.g., $\theta_2 = \theta_2''$), this type of IE does not exist.

Figure 2 provides an intuition for Proposition 2(b) and (c). It depicts the expected utility of the type-0 sender when she lies with $m = 1$ and $m = \theta_2$ as a function of σ, the probability she lies with $m = 1$. The type-0 sender's expected utility from lying with $m = 1$ is weakly decreasing in σ since as σ increases, the receiver is more likely to verify the message and takes a lower action when he does not.[12] On the other hand, as σ increases, the receiver is less likely to verify $m = \theta_2$ and takes a higher action when he does not. So the type-0 sender's expected utility from lying with $m = \theta_2$ is weakly increasing in σ.[13]

To complete the intuition for Proposition 2(b) and (c), there are two cases to consider. First, suppose $\theta_2 < \frac{1}{2}$. If $\beta \in (\underline{\beta}, \overline{\beta})$, where $\underline{\beta}$ is defined in (7), the curves depicted in the left panel of Figure 2 have a unique intersection point, denoted σ^*. When $\beta \geq \overline{\beta}$, lying with $m = 1$ becomes too attractive to the type-0 sender, and there is no longer a σ^* such that she is indifferent between lying with $m = 1$ and $m = \theta_2$. And when $\beta \leq \underline{\beta}$, there are a continuum of σ^* such that she is indifferent between lying with $m = 1$ and $m = \theta_2$. However, for every σ^*, the receiver verifies both messages in the support of the sender's strategy with certainty. This case is depicted in the right panel of Figure 2. In particular, when $\beta \leq \underline{\beta}$, there is an interval, $[\sigma_1, \sigma_0]$, such that if $\sigma \in [\sigma_1, \sigma_0]$, the receiver verifies both messages with certainty. In all IEs in this region of the parameter space, $\sigma^* \in [\sigma_1, \sigma_0]$.

Now suppose $\theta_2 \geq \frac{1}{2}$. Then lying with $m = 1$ is never too attractive to the type-0 sender that she cannot be made indifferent between lying with $m = 1$ and $m = \theta_2$. Hence, there is

[12] In fact, as depicted in Figure 2, for β sufficiently small and σ sufficiently large ($\sigma > \sigma_1(\beta)$), the type-0 sender's expected utility from lying with $m = 1$ is zero since the receiver verifies the message with certainty.

[13] Although not depicted in Figure 2, for β sufficiently small and σ sufficiently small, the sender's expected utility from lying with $m = \frac{1}{2}$ is zero since the receiver verifies the message with certainty.

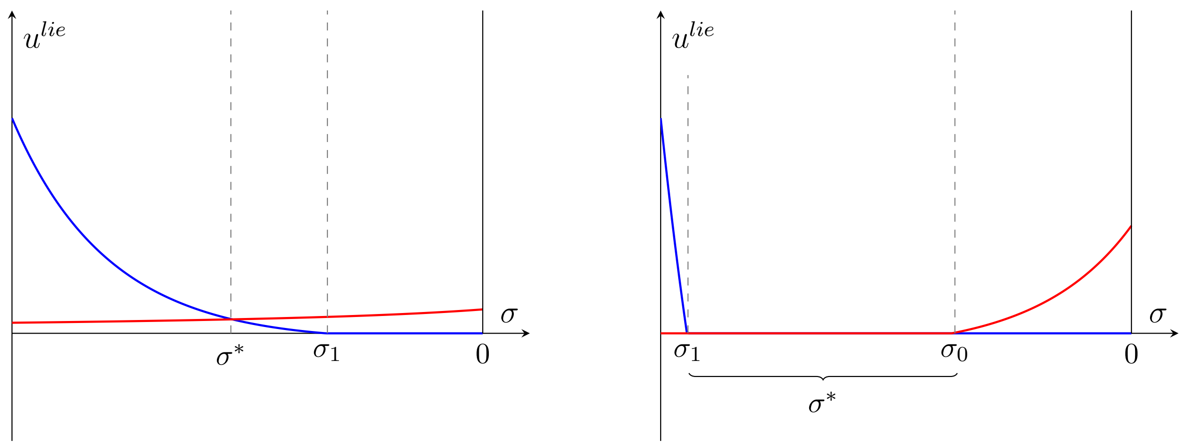

Figure 2: Left panel: Expected utility from lying with $m = 1$ and $m = \theta_2$ for a sender of type $\theta = 0$. Assumes $\beta = \frac{6}{25}$ and $\theta_2 = \frac{1}{3}$. Left panel: Expected utility from lying with $m = 1$ and $m = \theta_2$ for a sender of type $\theta = 0$. Assumes $\beta = \frac{1}{20}$ and $\theta_2 = \frac{1}{3}$.

no $\overline{\beta}$ such that this type of IE does not exist.

One might imagine that as β approaches zero, there will be an IE in which the sender always tells the truth. As discussed above, this cannot be the case case; there cannot be an IE where the sender never lies. This is due to the endogenous nature of verification. If the receiver believes the sender never lies, he has no reason to verify the sender's messages—even if the cost of doing so approaches zero. But then, the sender has the ability to lie without fear of being caught. This insight illustrates how the existence of influential communication between a biased sender and receiver requires an accurate suspicion that the sender may be lying. Without this suspicion, the receiver never verifies the sender's message, but then the sender can lie freely, destroying influential communication.

The following proposition shows what happens when β approaches infinity.

Proposition 3. *When $\beta \to \infty$, $\mathbb{E}[\theta|(m, \emptyset)] \to \mathbb{E}[\theta]$ for all m in the support of the sender's strategy.*

Proposition 3 illustrates that when the probability the receiver verifies the sender's messages approaches zero, my model approaches traditional cheap talk in which the receiver's belief on the equilibrium path is equal to the prior. The precise structure of the limit equilibrium depends on θ_2. When $\theta_2 < \frac{1}{2}$, the type-0 and type-θ_2 sender lie by sending $m = 1$.[14] As a result, when the receiver does not verify the highest message, his belief is equivalent to the prior. When $\theta_2 \geq \frac{1}{2}$, the type-0 sender lies by randomizing between $m = \theta_2$ and $m = 1$. In equilibrium, she must be indifferent between these lies. In particular, in the limit, this means she must be indifferent between the two actions the receiver takes when he does not verify.

The Effects of Verification Let Γ^{CT} denote the version of the three-state model that is identical except the receiver does not have the ability to verify. Put differently, Γ^{CT} is the traditional cheap talk model. Perhaps unsurprisingly, the ability to verify the sender's message improves the receiver's expected utility relative to her expected utility in Γ^{CT}.[15]

Proposition 4. *In any IE, the receiver's expected utility is strictly higher than his expected utility in* Γ^{CT}.

There are two reasons the receiver's expected utility is higher when he can verify the sender's message than in Γ^{CT}. The first is that verification has a direct *informational* effect, $\Lambda(m)$, which is the difference between the variance in the receiver's belief if he does not verify and his expected variance if he does. Although verification is costly, ex-ante, the informational effect outweighs the expected cost.[16]

Verification also increases the receiver's expected utility through an indirect *deterrence* effect. In equilibrium, verification generates an endogenous cost of lying for the type-θ_2 sender. If she lies and is caught, she will be pooled with the type-0 sender. In some cases, this cost of lying deters the type-θ_2 from lying.

The deterrence effect can be seen in the IE described by Proposition 2(b) and (c).[17] In

[14] In fact, this is true for any β such that this IE exists.

[15] In Γ^{CT}, no informative equilibria exists.

[16] In contrast, in Sadakane and Tam (2023), the direct informational effect is completely offset by the cost of verification because the equilibrium is constructed such that when the receiver verifies, he is exactly indifferent between verifying and not verifying.

[17] This IE exists simultaneously with the IE in which verification only provides an informational effect, described by Proposition 2(a), when $\theta_2 < \frac{1}{2}$ and $\beta \in (\tilde{\beta}, \overline{\beta})$

this IE, the type-θ_2 sender tells the truth. Deviating to lying with $m = 1$ leads to a payoff of $a(1, \emptyset) > a(\theta_2, \emptyset)$ if the receiver does not verify. But when he does, the type-θ_2 sender is punished by the receiver choosing an action that pools her with the type-0 senders who also lie. In this IE, this punishment is costly enough to deter the type-θ_2 sender from lying.

While the informational effect of verification is present in all IEs, the deterrence effect is not—in the IE described by Proposition 2(a), the type-θ_2 sender is not deterred from lying. Moreover, when $\theta_2 < \frac{1}{2}$ and $\beta \in (\tilde{\beta}, \overline{\beta})$, two IEs exist: one where verification has a deterrence effect and an informational effect and one where verification only has an informational effect. A natural conjecture is that these effects operate like substitutes: if verification deters lying, it reduces the informational value of verifying the sender's message. In the following proposition, I show this is not necessarily the case. To do so, I first introduce the concept of the expected informational effect of verification:

$$\mathbb{E}_{m\in M^*}[\Lambda(m)] = \Sigma_{m\in M^*}\rho(m)(Var(m|c=0) - (1-\phi(m))(Var(m|c=1, v=f))$$

where $\rho(m)$ is the probability m is sent on the equilibrium path.[18]

Proposition 5. *Suppose $\theta_2 \in (\overline{\theta}, \frac{1}{2})$, where $\overline{\theta}$ is defined by (12), and $\beta \in (\tilde{\beta}, \overline{\beta})$. For β sufficiently close to $\overline{\beta}$, the expected informational effect of verification is greater in the IE where there is a deterrence effect than in the IE where there is not.*

An intuition for this result is as follows. Fix θ_2. In the IE where verification has a deterrence effect increasing σ^* has two effects: it increases the variance in the receiver's belief when he observes $m = 1$ and does not verify, and it increases the probability the sender lies with $m = 1$. Together, this means the informational effect of verification is increasing σ^*. Stated differently, the deterrence effect means the sender lies with more lies in equilibrium, and the expected informational effect of verification is higher when the sender lies relatively more often by pretending to be the highest type.

[18]Recall, $\phi(m)$ is the probability m is sent truthfully.

Now consider the IE in which verification only has an informational effect. This informational effect is decreasing in θ_2. This is because as θ_2 increases, the variance in the receiver's belief when he does not verify $m = 1$ is increasing, and the variance when he does is decreasing. Put differently, without deterrence, more sender types lie, and the informational effect of verification increases the more diverse this pool of liars is.

Consider the limiting case where $\sigma^* = 1$. If $\theta_2 > \overline{\theta}$, the expected informational effect of verification is greater in the IE where there is a deterrence effect than in the IE where there is not. Furthermore, for σ^* sufficiently close to 1, this is still true. σ^* is not a parameter; it is determined in equilibrium and depends on θ_2 and β. However, for any θ, if $\beta \to \overline{\beta}$, $\sigma^* \to 1$. Since σ^* is continuous in β, which I show in the Appendix, if $\theta_2 > \overline{\theta}$, and if β is sufficiently close to $\overline{\beta}$, the expected informational effect of verification is greater in the IE where there is a deterrence effect than in the IE where there is not.

The connection between the informational effect of verification and probability of verification yields the following corollary.

Corollary 1. *Suppose $\theta_2 \in (\overline{\theta}, \frac{1}{2})$, where $\overline{\theta}$ is defined by (12), and $\beta \in (\tilde{\beta}, \overline{\beta})$. For β sufficiently close to $\overline{\beta}$, the ex-ante probability of verification is greater in the IE where there is a deterrence effect than in the IE where there is not.*

Sender's Lying Strategy When $\theta_2 \geq \frac{1}{2}$ and $\beta > \underline{\beta}$, in the unique IE the type-0 sender randomizes between lying with $m = \theta_2$ and $m = 1$. The following proposition describes how the sender's lying strategy depends on θ_2 and β.[19]

Proposition 6. *Suppose $\theta_2 \in [\frac{1}{2}, 1)$. When $\beta > \underline{\beta}$, $\sigma^*(\beta, \theta_2)$ is*

(a.) decreasing in θ_2 if β is sufficiently large,

(b.) and increasing in β if $\sigma^*(\beta, \theta_2)$ sufficiently large.

[19] While this proposition focuses on the case where $\theta_2 \geq \frac{1}{2}$, Proposition B in the Appendix obtains an identical result for $\theta_2 \in (0, 1)$

Proposition 2 implies that when $\theta_2 \in [\frac{1}{2}, 1)$ and $\beta > \underline{\beta}$, there is a unique σ^* the type-0 sender's expected utility from lying with $m = 1$ is equal to her expected utility from lying with $m = 0$. Differentiating this identity and rearranging reveals:

$$\frac{\partial \sigma^*}{\partial \theta_2} = \frac{\frac{\partial b}{\partial \theta_2}}{\frac{\partial a}{\partial \sigma} - \frac{\partial b}{\partial \sigma}} \text{ and } \frac{\partial \sigma^*}{\partial \beta} = \frac{\frac{\partial b}{\partial \beta} - \frac{\partial a}{\partial \beta}}{\frac{\partial a}{\partial \sigma} - \frac{\partial b}{\partial \sigma}},$$

where $a(\sigma, \beta)$ and $b(\sigma, \beta)$ are the type-0 sender's expected utility from lying with $m = 1$ and $m = \theta_2$ respectively. Previous analysis shows that the denominators of these expressions are negative so their sign depends on the sign of the numerator.

Consider $\frac{\partial \sigma^*}{\partial \theta_2}$, which is increasing if $\frac{\partial b}{\partial \theta_2} > 0$. As θ_2 increases, it has two effects on the type-0 sender's expected utility from lying with $m = \theta_2$. On the one hand, it increases the receiver's belief about the state when he does not verify. On the other, it increases the probability the receiver verifies the message. When β is large and the probability of verification is high, the first effect dominates and the derivative is positive. Hence, this proposition shows that when the cost of verification may be quite high, the intermediate message corresponding to a better state means the type-0 sender lies more often with the intermediate message.

Now consider $\frac{\partial \sigma^*}{\partial \beta}$, which is increasing if $\frac{\partial b}{\partial \beta} - \frac{\partial a}{\partial \beta} < 0$. These derivatives depend on σ. In particular, when σ is large, $\frac{\partial b}{\partial \beta} < \frac{\partial a}{\partial \beta}$. Hence, when the type-0 sender lies sufficiently often with $m = 1$, increasing the upper bound of verification costs increases the probability the type-0 sender lies with $m = 1$.

5 Conclusion

In this paper, I explore a common setting in politics: a biased but informed sender communicating using cheap talk with a receiver. Seminal work on cheap talk shows that the sender's bias prevents influential communication from arising in equilibrium. However, I show that when the receiver can pay a privately known cost to verify the veracity of the sender's message, there can be influential communication. This speaks to the possibility of influential communication in settings like a politician speaking to voters when the politician's message is fact-checked by a journalist or a lobbyist communicating with a legislator about a study when the legislator can verify whether the lobbyist accurately reported the study's findings.

I show that any influential equilibrium of the model has a straightforward structure characterized by threshold: high-type senders tell the truth, and low-type senders lie by pretending to be high types. Except in a knife-edge case, any message the sender sends truthfully is also sent as a lie, and all messages in the support of the sender's strategy are verified with positive probability.

I then analyze the model when there are three equally likely states and the cost to verify is drawn from a uniform distribution. In addition to showing how equilibria change depending on the upper bound of verification costs, this analysis illustrates that verification has two effects: an informational effect and a deterrence effect. I explore these effects and show that, in some cases, the presence of the deterrence effect increases the expected informational effect. I also explore how the sender's lying strategy changes with the model's parameters.

6 Appendix: Main Results

6.1 Proof of Proposition 1

Proof. Suppose an IE exists.

(i.) Claim: At least one message is sent as a lie.

Proof. Suppose not. Then all messages are sent truthfully. When the receiver observes a message, he believes it, will not verify, and will choose $a = m$. But then a sender of type $\theta < \theta_N$ has a profitable deviation to $m = \theta_N$. ■

(ii.) Claim: At least one message is sent truthfully.

Proof. Suppose not. Then all messages are sent as lies and the receiver never verifies. By Assumption 1, the sender's lying strategy cannot depend on her type, which means $\mathbb{E}[\theta|m]$ is constant for all m. But then $a^*(m, v)$ is constant on the equilibrium path. This violates the definition of an IE. ■

(iii.) Claim: The sender's expected utility from lying with message m is constant for all lies in the support of the sender's strategy.

Proof. Let M_f be the set of messages that are sent as lies on the equilibrium path. Suppose the claim is false. Then there exist lies $m \in M_f$ and $m' \in M_f$ such that the sender's expected utility from lying with m is strictly higher than lying with m'. But then a liar who was supposed to lie with m' could deviate to lying with m. Hence, this is not an IE. Denote the expected utility of lying u^{lie}. ■

(iv.) Claim: Each message m such that $m \neq u^{lie}$ and m is in the support of the sender's strategy is sent as a lie.

Proof. Let M_t be the set of messages that are sent truthfully on the equilibrium path, and let $M^* = M_f \cup M_t$ be the set of messages in the support of the sender's strategy.

Suppose the claim is false. Then there exists some $m \in M^*$ such that m is not in M_f and $m \neq u^{lie}$. When the receiver observes m, he will not verify and will choose $a = m$.

Suppose $m > u^{lie}$. Then a liar can deviate from lying according to her equilibrium strategy to lying by reporting m. The message will believed by Assumption 2, so it is a profitable deviation.

Suppose $m < u^{lie}$. Then the type-m sender can deviate from truthfully reporting m, in which case she receives a payoff of $u_S(m)$, to lying and get u^{lie}. This is a profitable deviation. ■

(v.) Claim: For any θ such that $\theta \neq u^{lie}$, the sender either lies with probability one or tells the truth with probability one.

Proof. Suppose not. Consider the type-θ where $\theta \neq u^{lie}$ who randomizes between lying and telling the truth. She must be indifferent between telling the truth and lying. By Assumption 1, all senders who lie must use the same strategy, so $m = \theta$ can only be sent truthfully. Hence, when the sender sends m, the receiver does not verify and chooses $a(m, \emptyset) = m$. It is immediate that the sender cannot be indifferent between telling the truth and lying. ■

(vi.) Claim: For any θ such that $\theta > u^{lie}$, the sender tells the truth with probability one.

Proof. Suppose not. Then there exists a θ such that $\theta > u^{lie}$ and the type-θ sender lies. Then $m = \theta$ is not in the support of the sender's strategy since the sender's strategy when lying is independent of θ and the type-θ cannot lie with $m = \theta$. Since m is off the equilibrium path, when the sender lies and receives u^{lie}, she can deviate to the message m. This message will be believed by Assumption 2, and the receiver will choose $a(m, \emptyset) = \theta$. By assumption, this is larger than u^{lie}. ■

(vii.) Claim: For any θ such that $\theta < u^{lie}$, the sender lies with probability one.

Proof. Suppose not. Then there exists θ such that $\theta < u^{lie}$ and the type-θ sender tells the truth.

The type-θ sender's expected utility from telling the truth is $\pi^*(m)\theta+(1-\pi^*(m))a(m,\emptyset)$, where $\pi^*(m)$ is the probability the receiver verifies the message m in equilibrium. Optimality implies $\pi^*(m)\theta + (1-\pi^*(m))a(m,\emptyset) \geq u^{lie}$, which requires $a(m,\emptyset) > u^{lie}$ since $\theta < u^{lie}$. However, $(vi.)$ shows that the sender tells the truth if $\theta > u^{lie}$. Hence $a(m,\emptyset) \leq \mathbb{E}[\theta|\theta \leq u^{lie}] < u^{lie}$. This is a contradiction. ■

(viii.) Claim: If there exists a θ such that $\theta = u^{lie}$, $m = \theta$ is only sent truthfully.

Proof. Suppose there exists a θ such that $\theta = u^{lie}$ and $m = \theta$ is sent as a lie.

Suppose further that $m = \theta$ is only sent as a lie. This implies the type-θ sender lies with probability one. But she cannot lie by truthfully reporting her type and Assumption 1 implies the sender's lying strategy is independent of her type.

Now suppose $m = \theta$ is sometimes sent truthfully. When $m = \theta$ is sent as a lie, it induces expected utility $\pi^*(m)a(m,f) + (1-\pi^*(m))a(m,\emptyset) = u^{lie}$. Since the sender lies if $\theta < u^{lie}$, $a(m,f) < u^{lie}$ and $a(m,\emptyset) < u^{lie}$. Hence, the required inequality cannot not hold. ■

■

6.2 Proof of Proposition 2

Let $\pi(m)$ denote the probability m is verified in equilibrium. Proposition 2 summarizes the following proposition.

Proposition A. *Define $\tilde{\beta}, \overline{\beta}$, and $\underline{\beta}$ as in (3), (5), and (7). An IE exists in all regions of the parameter space. In particular,*

(a.) if $\beta > \tilde{\beta}$, an IE exists where the sender tells the truth when $\theta = 1$, lies when $\theta \in \{0, \theta_2\}$ by sending $m = 1$, and $\pi(1) \in (0,1)$,

(b.) if $\beta \in (\underline{\beta}, \overline{\beta})$, where $\underline{\beta} < \tilde{\beta}$, an IE exists where the sender tells the truth when $\theta \in \{\theta_2, 1\}$, lies when $\theta = 0$ by randomizing between $m = 1$ with probability σ^* and $m = \theta_2$ with probability $1 - \sigma^*$ satisfying (4), and $\pi(1), \pi(\theta_2) \in (0, 1)$,

(c.) if $\beta = \overline{\beta}$, an IE exists where the sender tells the truth when $\theta \in \{\theta_2, 1\}$, lies when $\theta = 0$ by sending $m = 1$, $\pi(\theta_2) = 0$ and $\pi(1) \in (0, 1)$,

(d.) and if $\beta < \underline{\beta}$, a continuum of IEs exist where the sender tells the truth when $\theta \in \{\theta_2, 1\}$, lies when $\theta = 0$ by randomizing between $m = 1$ with probability σ^* and $m = \theta_2$ with probability $1 - \sigma^*$ satisfying (8), and $\pi(1) = \pi(\theta_2) = 1$.

Proof. **If direction:**

Proposition 1 implies that in an IE, there are three possible cases with respect to the receiver's verification: the receiver verifies all messages in the support of the sender's strategy with probability $\pi(m) \in (0, 1)$, $\pi(m) = 1$ for all messages in the support of the sender's strategy, or there exists a θ such that $\theta = u^{lie}$ and $m = \theta$ is on path but only sent truthfully.

$\pi(m) \in (0, 1)$ for m on the path:

Suppose an IE exists where the receiver verifies each message in the support of the sender's strategy probability with probability $\pi(m) \in (0, 1)$. Proposition 1 implies there are two cases to consider with respect to the sender's threshold.

First, suppose an IE exists where the sender lies when $\theta \in \{0, \theta_2\}$ by reporting $m = 1$, and tells the truth when $\theta = 1$. The type-0 sender and type-θ_2 sender's expected utility from lying with $m = 1$ is:

$$\min\left\{\frac{(2-\theta_2)^2}{18\beta}, 1\right\} u_S\left(\frac{\theta_2}{2}\right) + \max\left\{1 - \frac{(2-\theta_2)^2}{18\beta}, 0\right\} u_S\left(\frac{\theta_2+1}{3}\right). \tag{2}$$

This is a piecewise function that is constant for $\beta \in (0, \frac{(2-\theta_2)^2}{18})$ since $\min\{\frac{(2-\theta_2)^2}{18\beta}, 1\} = 1$, and is strictly increasing for $\beta > \frac{(2-\theta_2)^2}{18}$:

$$\frac{\partial(2)}{\partial\beta} = -\frac{(2-\theta_2)^2}{18\beta^2}u_S\left(\frac{\theta_2}{2}\right) + \frac{(2-\theta_2)^2}{18\beta^2}u_S\left(\frac{\theta_2+1}{3}\right) > 0$$

since u_S is linearly increasing and $\frac{\theta_2+1}{3} > \frac{\theta_2}{2}$ for $\theta_2 \in (0,1)$.

Neither the type-0 sender nor the type-θ_2 sender have have a profitable deviation from lying as long as (2) $\geq \theta_2$ since $m = \theta_2$ is off the equilibrium path and can be obtained by deviating due to Assumption 2. Since (2) is strictly increasing in β when $\pi(1) \in (0,1)$, this condition is satisfied as long as $\beta \geq \tilde{\beta}$, where

$$\tilde{\beta} \equiv \frac{(2-\theta_2)^3}{36(1-2\theta_2)}. \tag{3}$$

(3) is positive for all $\theta_2 \in (0, \frac{1}{2})$. Otherwise, $\theta_2 >$ (2).

The type-1 sender never has a profitable deviation from telling the truth since:

$$\min\left\{\frac{(2-\theta_2)^2}{18\beta}, 1\right\}u_S(1) + \max\left\{1 - \frac{(2-\theta_2)^2}{18\beta}, 0\right\}u_S\left(\frac{\theta_2+1}{3}\right) > (2)$$

Then, since $\frac{(2-\theta_2)^3}{36(1-2\theta_2)} > \frac{(2-\theta)^2}{18}$ for all $\theta_2 \in (0, \frac{1}{2})$, this IE exists for all $\theta_2 \in (0,1)$ where $\beta \geq \tilde{\beta}$, and does not exist otherwise.

Second, suppose an IE exists where the sender tells the truth when $\theta \in \{\theta_2, 1\}$, and lies when $\theta = 0$ by randomizing over $m \in \{\theta_2, 1\}$. Proposition 1 implies that in this IE, $\sigma^* = \sigma$ solves

$$\max\left\{\left(1 - \frac{\sigma}{\beta(1+\sigma)^2}\right), 0\right\}u_S\left(\frac{1}{1+\sigma}\right) = \max\left\{\left(1 - \frac{(1-\sigma)\theta_2^2}{\beta(2-\sigma)^2}\right), 0\right\}u_S\left(\frac{\theta_2}{2-\sigma}\right). \tag{4}$$

The left-hand side of (4) is a piecewise function that is strictly decreasing in σ for $\sigma \in$

$(0,1)$ when $\beta > \frac{\sigma}{(1+\sigma)^2}$:

$$\frac{\partial}{\partial \sigma} \max\left\{\left(1 - \frac{\sigma}{\beta(1+\sigma)^2}\right), 0\right\} u_s\left(\frac{1}{1+\sigma}\right) = \\ -\frac{1-\sigma}{\beta(1+\sigma)^3} u_S\left(\frac{1}{1+\sigma}\right) + \left(1 - \frac{\sigma}{\beta(1+\sigma)^2}\right) \frac{\partial}{\partial \sigma} u_S\left(\frac{1}{1+\sigma}\right) < 0,$$

since $u_S(a)$ is increasing in a. Moreover, (4) is constant for $\sigma \in (0,1)$ when $\beta < \frac{\sigma}{(1+\sigma)^2}$ since $\max\{(1 - \frac{\sigma}{\beta(1+\sigma)^2}), 0\} = 0$. Additionally, (4) $\to u_S(1)$ when $\sigma \to 0$, and (4) $\to u_S(\max\{\frac{4\beta-1}{8\beta}, 0\})$ when $\sigma \to 1$.

The right-hand side of (4) is a piecewise function that is strictly increasing in σ for $\sigma \in (0,1)$ when $\beta > \frac{(1-\sigma)\theta_2^2}{(2-\sigma)^2}$:

$$\frac{\partial}{\partial \sigma} \max\left\{\left(1 - \frac{(1-\sigma)\theta_2^2}{\beta(2-\sigma)^2}\right), 0\right\} u_S\left(\frac{\theta_2}{(2-\sigma)}\right) = \\ \frac{\sigma\theta_2^2}{\beta(2-\sigma)^3} u_S\left(\frac{\theta_2}{2-\sigma}\right) + \left(1 - \frac{(1-\sigma)\theta_2^2}{\beta(2-\sigma)^2}\right) \frac{\partial}{\partial \sigma} u_S\left(\frac{\theta_2}{2-\sigma}\right) > 0.$$

since $u_S(a)$ is increasing in a. Moreover, (4) is constant for $\sigma \in (0,1)$ when $\beta < \frac{(1-\sigma)\theta_2^2}{(2-\sigma)^2}$, which implies $\max\{(1 - \frac{(1-\sigma)\theta_2^2}{\beta(2-\sigma)^2}, 0\} = 0$. Additionally, (4) $\to u_S(\max\{\frac{(4\beta-\theta_2^2)\theta_2}{8\beta}, 0\})$ when $\sigma \to 0$, and (4) $\to u_S(\theta_2)$ when $\sigma \to 1$.

Hence, there is at least one $\sigma^* \in (0,1)$ such that(4) is satisfied as long as $max\{\frac{4\beta-1}{8\beta}, 0\} < u_S(\theta_2)$. This is always satisfied if $\theta_2 \geq \frac{1}{2}$, but is satisfied if and only if $\beta < \overline{\beta}$ when $\theta_2 < \frac{1}{2}$, where

$$\overline{\beta} \equiv \frac{1}{4(1-2\theta_2)}. \tag{5}$$

Moreover, given σ^*, the receiver verifies both messages in the support of the sender's strategy with probability $\pi(m) \in (0,1)$ unless

$$\frac{1 - \sqrt{1-4\beta} - 2\beta}{2\beta} < \frac{4\beta - \theta_2^2 + \sqrt{-4\beta\theta_2^2 + \theta_2^4}}{2\beta} \tag{6}$$

where the left-hand side of (6) is the value of σ such that $\beta = \frac{\sigma}{(1+\sigma)^2}$ and the right-hand side of (6) is the value of σ such that $\beta = \frac{(1-\sigma)\theta_2^2}{(2-\sigma)^2}$.

The left-hand side of (6) is defined for $\beta < \frac{1}{4}$, is increasing in β:

$$\frac{\partial}{\partial \beta} \frac{1 - \sqrt{1-4\beta} - 2\beta}{2\beta} = \frac{1 - \sqrt{1-4\beta} - 2\beta}{2\beta^2 \sqrt{1-4\beta}} > 0,$$

$\to 1$ when $\beta \to \frac{1}{4}$, and $\to 0$ when $\beta \to 0$.

The right-hand side of (6) is defined for $\beta < \frac{\theta_2^2}{4}$, is decreasing in β over this range:

$$\frac{\partial}{\partial \beta} \frac{\sqrt{1-36\beta} + 36\beta - 1}{18\beta} = \frac{2\beta\theta_2^2 - \theta_2^4 + \theta_2^2\sqrt{-4\beta\theta_2^2 + \theta_2^4}}{2\beta^2\sqrt{-4\beta\theta_2^2 + \theta_2^4}} < 0$$

$\to 0$ when $\beta \to \frac{\theta_2^2}{4}$, and $\to 1$ when $\beta \to 0$. Hence, for $\beta > \underline{\beta}$, where $\underline{\beta}$ is the value of β such that

$$\frac{1 - \sqrt{1-4\beta} - 2\beta}{2\beta} = \frac{4\beta - \theta_2^2 + \sqrt{-4\beta\theta_2^2 + \theta_2^4}}{2\beta}, \tag{7}$$

(6) is not satisfied and there exists a unique σ^* such that (4) is satisfied.

It is clear that a type-0 sender does not have a profitable deviation to truthfully reporting $m = 0$ as Assumption 2 implies the message will be believed, yielding her a payoff of $u_S(0)$, which is lower than the payoff she receives from lying, u^{lie}, which is positive. Additionally, neither the type-θ_2 nor the type-1 sender have a profitable deviation to lying because when the receiver verifies and learns the message was true, he chooses $a = 1 > 0$ or $a = \theta_2 > 0$.

Additionally, note that because $\underline{\beta} < \frac{\theta_2^2}{4}$ and $\tilde{\beta} = \frac{1}{4(1-2\theta_2)}$, $\underline{\beta} < \tilde{\beta}$ for all $\theta_2 \in (0, \frac{1}{2})$.

<u>$\pi(m) = 1$ for m on the path:</u>

Now, suppose an IE exists where the receiver verifies each message in the support of the sender's strategy with probability $\pi(m) = 1$. Proposition 1 implies there are two possible IEs.

First, suppose an IE exists where the sender lies when $\theta \in \{0, \theta_2\}$ by reporting $m = 1$, and tells the truth when $\theta = 1$. Since the receiver verifies $m = 1$ with probability one, the sender of type $\theta = \theta_2$ receives a payoff of $u_S(\frac{\theta_2}{2})$, which is lower than the payoff she would get from deviating to $m = \theta_2$.

Second, suppose there is an IE where the sender of type $\theta \in \{\theta_2, 1\}$ tells the truth and the sender of type $\theta = 0$ lies by randomizing between $m \in \{\theta_2, 1\}$. The type-0 sender does not have an incentive to deviate because her payoff is $u_S(0)$ on the path and $u_S(0)$ if she deviates to $m = 0$. Moreover, neither the type-θ_2 sender nor the type-1 sender have an incentive to deviate from telling the truth because doing so will lead to a payoff of $u_S(0)$ which is lower than the payoff on the equilibrium path. Given the previous analysis, this IE exists if there is a σ^* such that $\frac{\sigma^*}{(1+\sigma^*)^2} \geq \beta$ and $\frac{(1-\sigma^*)\theta_2^2}{(2-\sigma^*)^2} \geq \beta$. That is, this IE exists for σ

$$\sigma \in [\sigma_1, \sigma_0], \tag{8}$$

where $\sigma_1 = \frac{-1-\sqrt{1-16\beta}+16\beta}{8\beta}$ and $\sigma_0 = \frac{4\beta-\theta_2^2+\sqrt{-4\beta\theta_2^2+\theta_2^4}}{2\beta}$. Previously, it was shown that this interval is not empty as long as $\beta \leq \underline{\beta}$. Moreover, the previous discussion implies that $\frac{4\beta-\theta_2^2+\sqrt{-4\beta\theta_2^2+\theta_2^4}}{2\beta} < \frac{1-\sqrt{1-4\beta}-2\beta}{2\beta}$ if $\beta < \tilde{\beta}$.

<u>$\theta = u^{lie}$ and $m = \theta$ is sent truthfully</u>:

Finally, suppose an IE exists where there is a θ such that $\theta = u^{lie}$ and $m = \theta$ is on path and only sent truthfully. This implies $\pi(\theta) = 0$. Proposition 2 implies that when $\Theta = \{0, \theta_2, 1\}$, the only case where this IE might exists is if $u^{lie} = \theta_2$. In such an IE, the type-0 sender lies, and her expected utility is $\frac{4\beta-1}{8\beta}$. Hence, such an IE exists if $\theta_2 = \frac{4\beta-1}{8\beta}$. Rearranging, this IE exists if $\beta = \overline{\beta}$.

Only if direction:

Suppose $\theta_2 \in (0, 1)$ and $\beta \geq \tilde{\beta}$. It is immediate from the previous analysis that the type-0 and type-θ_2 sender prefer lying with $m = 1$ to truthfully reporting the state or deviating to an off the path message. Moreover, the type-1 sender prefers truthfully reporting the state to lying or deviating. This implies an IE exists in which the players strategies are as described in Proposition 2(a).

Suppose next that $\theta_2 \in (0,1)$, $\beta \in (\underline{\beta}, \overline{\beta})$. It is immediate from the previous analysis that there is a unique σ^* satisfying (4). Moreover, given this σ^*, the type-0 sender prefers to lie than deviating to $m = 0$ and the type-1 sender and type-θ_2 sender prefer to tell the truth to lying or deviating off path. This implies an IE exists in which the players strategies are as described in Proposition 2(b).

Suppose that $\beta = \overline{\beta}$. If the type-0 sender lies and reports $m = 1$, $u^{lie} = \theta_2$. Hence, an IE exists in which the players strategies are as described in Proposition 2(c).

Finally, suppose $\beta \leq \underline{\beta}$. It is immediate from the previous analysis that there is a continuum of σ^* such that the type-0 sender is weakly prefers to lie and the type-θ_2 and type-1 sender prefer to tell the truth than deviate off path or lie. Hence, an IE exists in which the players strategies are as described in Proposition 2(d). ∎

6.3 Proof of Proposition 3

Proof. If $\theta \in (0, \frac{1}{2})$, then as $\beta \to \infty$ the unique IE is one where the sender tells the truth when $\theta = 1$ and lies when $\theta \in \{0, \theta_2\}$. In this IE, $\mathbb{E}[\theta|(1, \emptyset)] = \frac{\theta_2+1}{3} = \mathbb{E}[\theta]$.

If $\theta \in [\frac{1}{2}, 1)$, then as $\beta \to \infty$ the unique IE is one where the sender tells the truth when $\theta \in \{\theta_2, 1\}$ and lies when $\theta = 0$ by randomizing between $m = 1$ and $m = \theta_2$. As $\beta \to \infty$, σ^* solves $\frac{1}{1+\sigma} = \frac{\theta_2}{2-\sigma} \Leftrightarrow \sigma^* = \frac{2-\theta}{1+\theta}$. Plugging σ^* in, $\mathbb{E}[\theta|(1, \emptyset)] = \mathbb{E}[\theta|(\theta_2, \emptyset)] = \frac{\theta_2+1}{3} = \mathbb{E}[\theta]$. ∎

6.4 Proof of Proposition 4

Proof. In Γ^{CT}, the receiver's expected utility is:

$$-\frac{2(1-\theta_2+\theta_2^2)}{9}. \tag{9}$$

To show the receiver's expected utility is strictly higher in any IE, I show that when he does not verify, his expected utility is weakly higher. Due to the optimality of the receiver's best response, the fact that verification reduces variance, and the fact that receiver strictly prefers to verify for some realizations of κ, this is sufficient.

In the IE described by Proposition A(a), the receiver's expected utility when he does not verify is the same as his expected utility in Γ^{CT}.

In the IE described by Proposition A(b), the receiver's expected utility when he does not verify is $-\frac{\sigma}{3(1+\sigma)} - \frac{(1-\sigma)\theta_2^2}{3(2-\sigma)}$. It can be shown through algebraic manipulation that $\overline{u}_R^b \geq$ (9) for all σ and θ_2.

In the IE described by Proposition A(c), the receiver's expected utility when he does not verify is - $\frac{1}{4}$. It can be shown through algebraic manipulation that $\overline{u}_R^b \geq$ (9) for all θ_2.

And in the IEs described by A(d), the receiver's expected utility when he does not verify is $-\frac{\sigma}{3(1+\sigma)} - \frac{(1-\sigma)\theta_2^2}{3(2-\sigma)}$. As discussed above, it can be shown through algebraic manipulation that $\overline{u}_R^b \geq$ (9) for all σ and θ_2. ■

6.5 Proof of Proposition 5

To prove Proposition 5, I begin with the following lemma.

Lemma A. *Consider the IE described by Proposition 2(b). If $\beta \in (\underline{\beta}, \overline{\beta})$, $\sigma^*(\beta)$ is continuous.*

Proof. Define an implicit function $h(\sigma, \beta) = a(\sigma, \beta) - b(\sigma, \beta)$, where $a(\sigma, \beta)$ and $b(\sigma, \beta)$ are the type-0 sender's expected utility from lying with $m = 1$ and $m = \theta_2$ respectively. Proposition 2 implies that for every $\beta \in (\underline{\beta}, \overline{\beta})$ there exists a unique $\sigma^*(\beta)$ such that $h(\sigma^*(\beta), \beta) = 0$. Proposition 2 also shows that $a(\sigma, \beta)$ and $b(\sigma, \beta)$ are continuously differentiable in β and σ. Hence, $h(\sigma, \beta)$ is too. Applying the implicit function theorem, for an arbitrary $\beta \in (\underline{\beta}, \overline{\beta})$, there exists a unique, differentiable function $\phi(\beta)$ such that in the neighborhood of β, $h(\phi(\beta), \beta) = 0$ as long as $\frac{\partial a}{\partial \sigma}(\phi(\beta), \beta) - \frac{\partial b}{\partial \sigma}(\phi(\beta), \beta) \neq 0$. In the proof of Proposition 2, I showed that $\frac{\partial a}{\partial \sigma} < 0$ and $\frac{\partial b}{\partial \sigma} > 0$, ensuring this condition is satisfied. Then, because there is a unique $\sigma^*(\beta)$ and unique $\phi(\beta)$ for all $\beta \in (\underline{\beta}, \overline{\beta})$, $\sigma^*(\beta) = \phi(\beta)$. Moreover, the local continuity of $\phi(\beta)$ on a connected space implies global continuity on the space. ■

I now prove Proposition 5.

Proof. Suppose $\theta \in (0, \frac{1}{2})$ $\beta \in (\tilde{\beta}, \overline{\beta})$. Then an IE exists where there is a deterrence benefit from verification (described by Proposition 2(a)) and an IE exists where there is not (described by Proposition 2(b)). In the latter, the expected informational effect of verification is

$$\frac{(2-\theta)^2}{18}. \tag{10}$$

In the former, the expected informational effect of verification is

$$\frac{(2-\sigma^*)\sigma^* + \theta_2^2(1-\sigma^{*2})}{3(2-\sigma^*)(1+\sigma^*)}. \tag{11}$$

(11) is increasing in σ^*:

$$\frac{\partial(11)}{\partial\sigma^*} = \frac{1}{3}\left(\frac{1}{(1+\sigma^*)^2} - \frac{\theta_2^2}{(2-\sigma^*)^2}\right) > 0$$

Moreover, for $\sigma^* = 1$, $(11) = \frac{1}{6}$. Hence, if

$$\begin{aligned} (10) &< \frac{1}{6} \\ \Leftrightarrow \overline{\theta}_2 \equiv 2 - \sqrt{3} &< \theta_2, \end{aligned} \tag{12}$$

then for σ^* sufficiently close to one, (11) > (10). Note, $\overline{\theta} < \frac{1}{2}$.

Fix $\theta_2 > \overline{\theta}_2$. Recall from the proof of Proposition 2 that for all $\beta \in (\underline{\beta}, \overline{\beta})$, there is a unique σ^* that solves (4). Taking the limit of $f(\sigma, \theta)$ and $g(\sigma, \theta)$ as $\beta \to \overline{\beta}$, it is clear that (4) is uniquely solved by $\sigma^* = 1$ for $\beta = \overline{\beta}$. Proposition 2 implies that for all $\beta \in (\tilde{\beta}, \overline{\beta})$, in the IE described by Proposition 2(b) $\sigma^* \in (0, 1)$. Moreover, Lemma A implies $\sigma^*(\beta)$ is continuous. Hence, for β sufficiently close to $\overline{\beta}$, (11) > (10). ■

6.6 Proof of Proposition 6

Proposition 6 follows from Proposition B.

Proposition B. *Suppose either* $\theta_2 \in [\frac{1}{2}, 1)$ *and* $\beta > \underline{\beta}$ *or* $\theta_2 \in (0, \frac{1}{2})$ *and* $\beta \in (\underline{\beta}, \overline{\beta})$*. In the IE where the type-*0 *sender randomizes,* $\sigma^*(\beta, \theta_2)$ *is*

(a.) decreasing in θ_2 *if* β *is sufficiently large,*

(b.) and increasing in β *if* $\sigma^*(\beta, \theta_2)$ *sufficiently large.*

Proof. Consider functions $a(\sigma, \beta)$ and $b(\sigma, \beta, \theta_2)$ where $a(\sigma, \beta)$ and $b(\sigma, \beta, \theta_2)$ are the type-0 sender's expected utility from lying with $m = 1$ and $m = \theta_2$ respectively. Proposition 2 implies that for every $\beta > \underline{\beta}$, in an IE where the type-0 sender randomizes, there exists a unique σ^* such that $a(\sigma^*, \beta) = b(\sigma^*, \beta, \theta_2)$. Differentiating both sides of the identity with respect θ_2 and rearranging, $\frac{\partial \sigma^*}{\partial \theta_2} = \frac{\frac{\partial b}{\partial \theta_2}}{\frac{\partial a}{\partial \sigma} - \frac{\partial b}{\partial \sigma}}$. The denominator is negative, since the proof of Proposition 2 shows a is decreasing in σ and b is increasing in σ. Hence, the sign depends on $\frac{\partial b}{\partial \theta_2}$:

$$\frac{1}{2-\sigma} - \frac{3\theta_2^2(1-\sigma)}{\beta(2-\sigma)^3},$$

which is positive if β is sufficiently large.

Again consider the identity $a(\sigma, \beta) = b(\sigma, \beta, \theta_2)$ but differentiate both sides with respect to β. Rearranging, $\frac{\partial \sigma^*}{\partial \theta_2} = \frac{\frac{\partial b}{\partial \beta} - \frac{\partial a}{\partial \beta}}{\frac{\partial a}{\partial \sigma} - \frac{\partial b}{\partial \sigma}}$. The denominator is negative so the sign depends on the numerator, which is positive as long as

$$\frac{\sigma}{\beta^2(1+\sigma)^3} > \frac{\theta_2^3(1-\sigma)}{\beta^2(2-\sigma)^3}. \tag{13}$$

The RHS and LHS of (13) are both increasing in σ for $\sigma < \frac{1}{2}$ and decreasing in σ for $\sigma > \frac{1}{2}$. Comparing their values when $\sigma = 0$, $\sigma = \frac{1}{2}$, and $\sigma = 1$ shows that (13) is satisfied for σ sufficiently large. ■

7 Supplementary Appendix: Additional Results

7.1 Existence of an IE

Proposition C. *Define $\check{\beta}$ and $\dot{\beta}$ as in (14) and (15). An IE exists if at least one of the following is satisfied:*

(a.) $\mathbb{E}[\theta] \geq \theta_{N-1}$ *and* $\beta \geq \check{\beta}$.

(b.) $\beta \leq \dot{\beta}$.

Proof. Consider a strategy for the sender where she lies if $\theta \in \{\theta_1, ..., \theta_{N-1}\}$ by reporting $m = \theta_N$ and tells the truth if $\theta = \theta_N$. Then when the sender lies, her expected utility is $\pi(\theta_N)\mathbb{E}[\theta|\theta \in \{\theta_1, \theta_{N-1}\}] + (1 - \pi(\theta_N))\mathbb{E}[\theta]$, where

$$\pi(\theta_N) = G\bigg(\frac{Var(\theta_N|c=0) - (1 - \phi(m)Var(\theta_n|v=f)}{\beta}\bigg).$$

The only effect of β on $\pi(\theta_N)$ is through the denominator, which means $\pi(\theta_N)$ is decreasing in β. Hence, $\mathbb{E}[\theta] > \theta_{N-1}$, there is a unique β such that

$$\pi(\theta_N)\mathbb{E}[\theta|\theta \in \{\theta_1, \theta_{N-1}\}] + (1 - \pi(\theta_N))\mathbb{E}[\theta] = \theta_{N-1}. \tag{14}$$

Denote this β as $\check{\beta}$. For all $\beta \geq \check{\beta}$, this IE exists.

Now consider a strategy for the sender where she lies if $\theta = \theta_1$ and tells the truth if $\theta > \theta_1$. In particular, when she lies, she uses a strategy where the probability she lies with $m = \theta_i \in \{\theta_2, ..., \theta_N\}$ is $\sigma_i \in (0, 1)$. A vector of lying probabilities, $\boldsymbol{\sigma}$ induces a vector of probabilities of verification, $\boldsymbol{\pi}$, where, for a particular $\pi(m) \in \boldsymbol{\pi}$, $\pi(m) = G(\frac{Var(m|c=0)}{\beta})$. Fix $\boldsymbol{\sigma}$. The only effect of β on $\pi(m)$ is through the denominator. So there exists a β such that

$$1 = \min(\pi(m) \in \boldsymbol{\pi}). \tag{15}$$

Denote this β by $\dot{\beta}$. Moreover, for all $\beta \leq \dot{\beta}$, the receiver verifies each m in the support of the sender's strategy with certainty. ■

References

Austen-Smith, D. and J. R. Wright (1992): "Competitive lobbying for a legislator's vote," *Social choice and Welfare*, 9, 229–257.

Balbuzanov, I. (2019): "Lies and consequences," *International Journal of Game Theory*, 48, 1203–1240.

Ball, I. and X. Gao (2025): "Checking cheap talk," *Working Paper*.

Crawford, V. P. and J. Sobel (1982): "Strategic information transmission," *Econometrica*, 1431–1451.

Dziuda, W. and C. Salas (2018): "Communication with detectable deceit," *Working Paper*.

Ederer, F. and W. Min (2022): "Bayesian Persuasion with Lie Detection," .

Ellis, C. J. and T. Groll (2020): "Strategic legislative subsidies: Informational lobbying and the cost of policy," *American Political Science Review*, 114, 179–205.

Gottfried, J. A., B. W. Hardy, K. M. Winneg, and K. H. Jamieson (2013): "Did fact checking matter in the 2012 presidential campaign?" *American Behavioral Scientist*, 57, 1558–1567.

Guo, Y. and E. Shmaya (2019): "The interval structure of optimal disclosure," *Econometrica*, 87, 653–675.

Hart, S., I. Kremer, and M. Perry (2017): "Evidence games: Truth and commitment," *American Economic Review*, 107, 690–713.

Hojnacki, M. and D. C. Kimball (1998): "Organized interests and the decision of whom to lobby in Congress," *American Political Science Review*, 92, 775–790.

Holm, H. J. (2010): "Truth and lie detection in bluffing," *Journal of Economic Behavior & Organization*, 76, 318–324.

Kartik, N., M. Ottaviani, and F. Squintani (2007): "Credulity, lies, and costly talk," *Journal of Economic theory*, 134, 93–116.

Levkun, A. (2022): "Communication with strategic fact-checking," Tech. rep., mimeo. University of California, San Diego.

Lim, C. (2018): "Can Fact-checking Prevent Politicians from Lying?" *Working Paper*.

Minozzi, W. and J. Woon (2013): "Lying aversion, lobbying, and context in a strategic communication experiment," *Journal of Theoretical Politics*, 25, 309–337.

Nguyen, A. and T. Y. Tan (2021): "Bayesian persuasion with costly messages," *Journal of Economic Theory*, 193, 105212.

Nyhan, B. and J. Reifler (2010): "When corrections fail: The persistence of political misperceptions," *Political Behavior*, 32, 303–330.

——— (2015): "The effect of fact-checking on elites: A field experiment on US state legislators," *American Journal of Political Science*, 59, 628–640.

Pence, M. (2020): "Mike Pence, Dr. Fauci Coronavirus Task Force Press Conference June 26," https://www.rev.com/transcripts/mike-pence-dr-fauci-coronavirus-task-force-press-conference-june-26.

Pingree, R. J., D. Brossard, and D. M. McLeod (2014): "Effects of journalistic adjudication on factual beliefs, news evaluations, information seeking, and epistemic political efficacy," *Mass Communication and Society*, 17, 615–638.

Potters, J. and F. Van Winden (1992): "Lobbying and asymmetric information," *Public choice*, 74, 269–292.

Qui, L. (2020): "As Cases Surge, Pence Misleads on Coronavirus Pandemic," *The New York Times*, https://www.nytimes.com/2020/06/26/us/politics/coronavirus-pence-fact-check.html.

Sadakane, H. and Y. C. Tam (2023): "Cheap talk and Lie detection," *Working Paper*.

Sobel, J. (2013): "Giving and receiving advice," *Advances in Economics and Econometrics*, 1, 305–341.

——— (2020): "Lying and deception in games," *Journal of Political Economy*, 128, 907–947.

Venkatesh, R. S., S. Roy, and S. Pramanik (2025): "Bayesian Persuasion and Costly Verification: A model of two-sided commitment," *Working Paper*.

Weeks, B. E. (2015): "Emotions, partisanship, and misperceptions: How anger and anxiety moderate the effect of partisan bias on susceptibility to political misinformation," *Journal of Communication*, 65, 699–719.

Weeks, B. E. and R. K. Garrett (2014): "Electoral consequences of political rumors: Motivated reasoning, candidate rumors, and vote choice during the 2008 US presidential election," *International Journal of Public Opinion Research*, 26, 401–422.