

Costly Miscalibration*

Yingni Guo[†] Eran Shmaya[‡]

June 4, 2018

Abstract

We consider an online platform which provides probabilistic forecasts using some algorithm. We introduce a concept of miscalibration which measures the discrepancy between the forecast and the truth. We apply this concept to sender-receiver games in which miscalibration is costly for the sender (the platform). We show that, when the sender’s miscalibration cost is sufficiently high, he can achieve his commitment solution in an equilibrium. Moreover, under some assumption about the miscalibration-cost function, the only rationalizable strategy of the sender is his strategy in the commitment solution.

JEL: D81, D82, D83

Keywords: Costly miscalibration, Cheap talk, Commitment, Bayesian persuasion

1 Introduction

E-commerce platforms often provide information to customers about their products. For example, the fair aggregator Kayak.com provides forecasts of future prices. The real estate aggregator Redfin.com identifies “hot homes” that are likely to go fast. The platform generates the information using some algorithm, which is usually a trade secret and is not published. For example, Kayak just states that “our scientists develop these flight price

*We thank Arjada Bardhi, Navin Kartik, Amanda Friedenberg, Christoph Kuzmics, Elliot Lipnowski, Kota Murayama, Ludvig Sinander, Marciano Siniscalchi, Egor Starkov, Muhamet Yildiz, Andy Zapechelnuyk for helpful comments and discussions. We thank Cassiano Alves and Samuel Goldberg for excellent research assistance. Yingni Guo gratefully acknowledges financial support from NSF Grant SES-1530608.

[†]Department of Economics, Northwestern University, yingni.guo@northwestern.edu.

[‡]MEDS, Kellogg School of Management, Northwestern University, e-shmaya@kellogg.northwestern.edu.

trend forecasts using algorithms and mathematical models.” However, even though customers don’t observe the algorithms, they still trust the platform’s forecasts. In this paper, we provide a model in which this is an equilibrium phenomenon.

In our model, the platform provides information in the form of probabilistic forecasts. For example, Redfin defines a hot home as one that has a 70% chance or higher of having an accepted offer within two weeks of its debut. We say that the algorithm is *miscalibrated* if there is a discrepancy between the forecast and the realized data, for example, if only 50% of the houses Redfin identifies as hot homes have an accepted offer within two weeks of their debut. The platform incurs a cost that is a function of a measure of miscalibration defined in the paper. The cost function depends on a parameter, called the cost intensity, which indicates how severely the platform is punished for miscalibration. This cost is a proxy for the reputation damage to the platform from making incorrect probabilistic assertions.

We model the interaction between the platform and a customer as a sender-receiver game. The algorithm corresponds to the sender’s strategy, and the forecast to his message. As in the cheap-talk literature (e.g., Crawford and Sobel (1982) and Green and Stokey (2007)), the sender’s strategy is not observed by the receiver. When the cost intensity is zero, our game is a cheap-talk game. However, when the cost intensity is positive, the sender’s talk is not completely cheap, because of the miscalibration cost. We show that when the cost intensity is sufficiently high, the sender’s optimal payoff in equilibrium is arbitrarily close to the payoff he could get if his algorithm were observed by the receiver.

Our results rely on the fact that, when asserting incorrect distributions becomes very costly, Sender’s asserted distribution in equilibrium will not be very far from the true conditional distribution over states given his assertion. However, the results are not obvious for two reasons. First, we need to show that there exists some equilibrium in which the receiver’s behavior is close to his behavior under the commitment solution. Establishing this requires some genericity assumption, without which it may happen that in any equilibrium the receiver always best responds to the prior belief (see example 2). Second, when the cost intensity is high, Sender pays a very high cost for even small discrepancies. Therefore, after showing that there is an equilibrium in which the receiver’s behavior is close to his behavior under the commitment solution, we still need to show that Sender’s payoff is not far from his commitment payoff.

The results connect our model to the Bayesian persuasion literature which studies sender-receiver games with commitment on the sender’s side (e.g., Rayo and Segal (2010) and Kamenica and Gentzkow (2011)). In that literature, the sender is able to commit to any strategy. This is as if the receiver observes the sender’s strategy. In our model, as in

the cheap-talk literature, the receiver does not observe the sender’s strategy. Our model therefore provides a foundation for the commitment assumption in the Bayesian persuasion literature. From this aspect, our paper is related to Best and Quigley (2017) which provides a foundation for the commitment assumption by analyzing interactions between a long-run sender and short-run receivers. Moreover, our model bridges the cheap-talk and Bayesian persuasion models, and can be used to analyze the middle ground where the talk is not completely cheap nor the commitment absolute.

Our paper is related to the literature on strategic communication with a lying cost (e.g., Kartik, Ottaviani and Squintani (2007) and Kartik (2009)). In these papers, the set of messages is the state space, so the sender always declares a state. As the lying cost grows, the equilibrium becomes arbitrarily close to fully revealing. Our model includes a larger message space, which allows the sender to make probabilistic statements; this is one reason why it delivers a different asymptotic result: there is an equilibrium which is arbitrarily close to the commitment solution. Thus, our model expresses a different intuition than that in Kartik, Ottaviani and Squintani (2007) and Kartik (2009). They show that when lying is costly, the sender reveals all his information. We show that when lying is costly, the sender gains the commitment power not to tell a lie, but he is not forced to reveal all his information.

2 An example

In this section, we illustrate the main idea of our model through an example. An online platform (Sender) promotes a product to a customer (Receiver). The product’s quality is $s \in S = \{0, 2, 4\}$ with a uniform prior. Receiver decides whether or not to buy the product. Receiver’s payoff is $(s - 3)$ if he buys the product and zero otherwise. Sender gets a commission payoff of 1 whenever Receiver buys the product.

Sender’s strategy is a mapping from the state space S to the message space. The message space is the set of distributions over states $\Delta(S)$. For any message m that Sender announces with positive probability, we let $\tau(m)$ denote the conditional distribution over states by Bayes’ rule. We refer to $\tau(m)$ as the true conditional distribution given m . (The true conditional distribution given a message depends on Sender’s strategy. We omit this dependence in our notation $\tau(m)$ in this section.)

Consider Sender’s strategy in table 1a. Sender sends either message m_0 or m_1 . Each message is a distribution over states. Message m_0 states that the state is 0. Message m_1 states that the state is 2 or 4 with equal probability. By sending m_0 , the online platform says that the product is a lemon. By sending m_1 , it says that the quality is weakly above 2. Each column shows the probabilities with which Sender sends m_0 or m_1 in each state.

message \ state	state		
	0	2	4
$m_0 = (1, 0, 0)$	1	0	0
$m_1 = (0, 1/2, 1/2)$	0	1	1

Table 1a: A calibrated strategy

message \ state	state		
	0	2	4
$m_0 = (1, 0, 0)$	1/2	0	0
$m_1 = (0, 1/2, 1/2)$	1/2	1	1

Table 1b: A miscalibrated strategy

Given the strategy in table 1a, the true conditional distribution given m_0 is $\tau(m_0) = (1, 0, 0)$. The true conditional distribution given m_1 is $\tau(m_1) = (0, 1/2, 1/2)$. Since each message coincides with the true conditional distribution given that message, we say that this is a calibrated strategy. There are many other calibrated strategies. For example, Sender can choose to reveal no information at all, by announcing the prior $(1/3, 1/3, 1/3)$ at every state. Or he can reveal all information by announcing the messages $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$ in states 0, 2, 4, respectively.

Consider another strategy in table 1b. Sender uses the same messages m_0, m_1 as before. The only difference occurs when the state is 0: Sender now announces m_0 and m_1 with equal probability. The true conditional distribution given m_0 is $\tau(m_0) = (1, 0, 0)$, which still coincides with m_0 . However, the true conditional distribution given m_1 is $\tau(m_1) = (1/5, 2/5, 2/5)$, which differs from m_1 . Figure 1 illustrates this strategy. Each point in the triangle represents a distribution over states. The black dot is the prior. The gray dots m_0, m_1 are the messages that Sender uses. The black circles are the true conditional distributions given the messages. Since the true conditional distribution $\tau(m_1)$ given m_1 differs from the asserted distribution m_1 , this strategy is not calibrated.




Figure 1: The miscalibrated strategy in table 1b

If Sender uses this strategy to make recommendations, then the probability distribution of the products about which Sender announces m_1 is in fact $\tau(m_1)$. Therefore, what Sender claims about these products (that they are $(0, 1/2, 1/2)$) is not true. If an outside group collects data and performs some statistical test, then Sender might be caught.

In our model, whenever Sender sends a message whose true conditional distribution differs from its asserted distribution, Sender pays some cost, which we call a miscalibration cost. The bigger the difference is, the higher the cost will be. In this section, we use the total-variation distance $\frac{1}{2}|m - \tau(m)|_1$ to measure how much a message m differs from the true conditional distribution $\tau(m)$ given this message. For the strategy in table 1b, the total-variation distance between m_1 and $\tau(m_1)$ is $1/5$. Since Sender announces message m_1 with probability $5/6$, Sender pays the following miscalibration cost:

$$\frac{5}{6} \cdot \lambda \frac{1}{2}|m_1 - \tau(m_1)|_1 = \frac{5}{6} \cdot \lambda \cdot \frac{1}{5},$$

where $\lambda \geq 0$ is a parameter which measures the intensity of the miscalibration cost. If Sender uses a calibrated strategy, then he pays no miscalibration cost.

We now claim that the game admits an equilibrium in which Sender plays the calibrated strategy in table 1a and Receiver buys the product with the following probability:

$$\rho(m_1) = \begin{cases} \lambda, & \text{for } 0 < \lambda < 1, \\ 1, & \text{for } \lambda \geq 1, \end{cases}$$

if Sender announces m_1 , and does not buy otherwise. Since $\tau(m_1) = m_1$, both buying and not buying are Receiver's best responses to m_1 . Since $\tau(m_0) = m_0$, not buying is Receiver's best response to m_0 . We next argue that Sender has no incentive to deviate. Consider a deviation strategy which induces Receiver to buy more frequently: Sender sends m_1 in states 2 and 4, and sends m_0 and m_1 with probability $1 - \alpha$ and α , respectively, in state 0. The true conditional distribution given m_1 is:

$$\tau(m_1) = \left(\frac{\alpha}{2 + \alpha}, \frac{1}{2 + \alpha}, \frac{1}{2 + \alpha} \right).$$

Sender's payoff from this deviation strategy is the probability that Receiver buys minus the cost of miscalibration:

$$\begin{aligned} & \underbrace{\frac{2 + \alpha}{3} \rho(m_1)}_{\text{prob. that Receiver buys}} - \underbrace{\frac{2 + \alpha}{3}}_{\text{prob. of } m_1} \lambda \frac{1}{2} \underbrace{\left| \left(0, \frac{1}{2}, \frac{1}{2} \right) - \left(\frac{\alpha}{2 + \alpha}, \frac{1}{2 + \alpha}, \frac{1}{2 + \alpha} \right) \right|_1}_{\text{distance between } m_1 \text{ and } \tau(m_1)} \\ & = \frac{2\rho(m_1)}{3} + \frac{\alpha}{3}(\rho(m_1) - \lambda). \end{aligned}$$

Since $\rho(m_1) \leq \lambda$ for any $\lambda > 0$, Sender prefers to choose $\alpha = 0$. It can be shown that no

other deviations are profitable.

It is easy to verify that this is Sender's most preferred equilibrium. For sufficiently high $\lambda \geq 1$, Sender achieves the same payoff as if he could commit to a strategy. This equilibrium is outcome-equivalent to Sender's commitment solution. For intermediate $\lambda \in (0, 1)$, if Receiver bought for sure after m_1 , then Sender would announce m_1 not only in states 2 and 4 but also in state 0. Put differently, the miscalibration cost is not severe enough for Sender to gain full commitment power. For intermediate $\lambda \in (0, 1)$, Receiver does not buy for sure after m_1 , so Sender has less incentive to send m_1 in state 0. This compensates for the lack of commitment power on Sender's side. Sender's equilibrium payoff of $2 \min\{\lambda, 1\}/3$ increases in λ , which can be interpreted as the proxy for Sender's commitment power.

3 Environment and main results

Let S be a finite set of *states* equipped with a prior distribution p with full support, and M a Borel space of *messages*. A *Sender's strategy with message space M* is given by a Markov kernel σ from S to M : when the state is s , Sender randomizes a message from $\sigma(\cdot|s)$. For a Sender's strategy σ , let $\tau_\sigma : M \rightarrow \Delta(S)$ be such that $\tau_\sigma(m)$ is the conditional distribution over states given m . We sometimes use the notation $\tau_\sigma(s|m)$ for $\tau_\sigma(m)(s)$, both of which denote the conditional probability of s given m .

We say that a strategy σ has finite support if $\sigma(\cdot|s)$ has finite support for every s , in which case we let $\text{support}(\sigma) = \cup_s \text{support}(\sigma(\cdot|s))$. When σ has finite support, the conditional probability $\tau_\sigma(s|m)$ is given by:

$$\tau_\sigma(s|m) = \frac{p(s)\sigma(m|s)}{\sum_{s'} p(s')\sigma(m|s')},$$

for every $m \in \text{support}(\sigma)$, and is defined arbitrarily for $m \notin \text{support}(\sigma)$.

In section 3.1 we consider Sender's strategies with message space $M = \Delta(S)$, and introduce the concepts of calibrated strategies and miscalibration. These concepts are independent of Sender's interaction with Receiver. In section 3.2 we review the model of Sender-Receiver games, the definition of a cheap-talk equilibrium, and the definition of a commitment solution. Section 3.3 presents our definition of an equilibrium with costly miscalibration and our main results: Sender achieves his commitment payoff in an equilibrium when the cost intensity is sufficiently high.

3.1 Sender's calibrated strategies and miscalibration

From now on we assume that $M = \Delta(S)$, so that a message m is an asserted distribution over states. We thus refer to $\tau_\sigma(m)$ as the true conditional distribution over states given a

message m . We say that a strategy σ is *calibrated* if $\tau_\sigma(m) = m$, a.s.¹ Under a calibrated strategy, messages mean what they say, i.e., they can reliably be taken at face value.

The term *calibrated* is used in the forecasting literature (e.g., Dawid (1982), Murphy and Winkler (1987), Ranjan and Gneiting (2010)) for two closely related ideas. The first is the purely probabilistic sense as in our definition of a calibrated strategy of Sender; the second is the idea of calibration with the data, namely, that a probabilistic forecast (or a message, in our terminology) matches the realized distribution of states at those times in which the forecast was given.²

We now introduce a key component of our model: a measure of miscalibration for a Sender’s strategy. To define this measure, let $\kappa : \Delta(S) \times \Delta(S) \rightarrow \mathbf{R}_+$ be a continuous function, where $\kappa(q, m)$ measures the distance between a message $m \in \Delta(S)$ and a truth $q \in \Delta(S)$. We assume that $\kappa(q, m) = 0$ if and only if $m = q$ and that κ is convex in q . Let

$$d(\sigma) = \sum_s p(s) \int \kappa(\tau_\sigma(m), m) \sigma(dm|s) \quad (1)$$

be the expected distance between the distribution asserted by Sender’s message and the true conditional distribution given that message when Sender’s strategy is σ . Note that σ is calibrated if and only if $d(\sigma) = 0$.

Example 1. The Redfin example:

Redfin’s hot home algorithm “identifies hot homes based on real estate conditions in each market.” This allows us to focus on a local market. We collect 2,150 hot homes and track their status.³ When a home’s status become contingent, pending or sold, it has an accepted offer. The percentage of hot homes having an accepted offer within two weeks of its debut is 53.1 (see figure 2). This is different than 70%, so Redfin’s strategy is not calibrated.

3.2 Sender-Receiver games

We consider a game with incomplete information in which Sender, who observes the realization of the state, sends a message to Receiver. Receiver then chooses an action.

Let A be a finite set of *actions*. Let $v, u : S \times A \rightarrow \mathbf{R}$ be, respectively, Sender’s and Receiver’s payoff functions. A Receiver’s strategy is given by a Markov kernel ρ from M to A , with the interpretation that Receiver randomizes an action from $\rho(\cdot|m)$ after message m .

¹The term *a.s.* in our paper means “almost surely with respect to the probability distribution over messages induced by σ .” (Recall that τ_σ is defined up to a set of messages with probability zero.)

²There is also a literature about the strategic manipulation of calibration tests. See Olszewski (2015) for a recent survey.

³We examine the area of cook county in Illinois which includes 165 zip codes, and collect homes that Redfin identifies as hot homes over two weeks (03/14/2018 – 03/27/2018).




Figure 2: Percentage of hot homes having an accepted offer

For a profile (σ, ρ) of Sender's and Receiver's strategies, we let $\pi_{\sigma, \rho} \in \Delta(S \times A)$ be the induced distribution over states and actions when the players follow this profile. The payoffs to Sender and Receiver under strategy profile (σ, ρ) are given by:

$$V_0(\sigma, \rho) = \int v \, d\pi_{\sigma, \rho}, \quad \text{and} \quad U(\sigma, \rho) = \int u \, d\pi_{\sigma, \rho}.$$

The reason we add the subscript 0 in Sender's payoff function V_0 will become clear later when we define V_λ for every $\lambda \geq 0$.

Using CP to represent the commitment payoff, we let $CP = \max V_0(\sigma, \rho)$ be Sender's maximal payoff over all profiles (σ, ρ) such that ρ is a best response to σ . (It follows from Kamenica and Gentzkow (2011) that the maximum is indeed achieved.) We call a profile (σ, ρ) that achieves the maximum a *commitment solution*. CP is the highest payoff that Sender can achieve by committing to some strategy σ . This would be the case if Receiver could observe Sender's strategy.

We now define a cheap-talk equilibrium. The contrasting assumption is that Receiver does not observe the strategy that Sender uses to generate messages. This requires adding incentive-compatibility conditions on Sender's side to the definition of a commitment solution. A *Bayesian Nash equilibrium (BNE)* for a cheap-talk game is a Nash equilibrium in the normal form game defined by the payoff functions V_0 and U .

3.3 Costly miscalibration

We now define a BNE for the game with costly miscalibration. The definition is the same as that for a cheap-talk game except that Sender's payoff is given by:

$$V_\lambda(\sigma, \rho) = V_0(\sigma, \rho) - \lambda d(\sigma), \tag{2}$$

where $d(\sigma)$ is given by (1) and $\lambda \geq 0$ is a parameter which measures the *intensity* of Sender's miscalibration cost.

A BNE is a Nash equilibrium in the normal form game defined by the payoff functions V_λ and U .

Proposition 3.1. *Sender's payoff under any BNE is at most CP.*

Proof. Let (σ^*, ρ^*) be a BNE. Then

$$V_\lambda(\sigma^*, \rho^*) \leq V_0(\sigma^*, \rho^*) \leq CP,$$

where the first inequality follows from (2), and the second from (i) the definition of CP and (ii) the fact that ρ^* is Receiver's best response to σ^* . \square

We now turn to our main result that when the cost intensity is high, Sender can achieve the commitment solution in an equilibrium. We first make the additional assumption that the cost function $\kappa(q, m)$ has a kink at $q = m$. This assumption holds, for example, for (i) the total-variation distance $\kappa(q, m) = \frac{1}{2}|q - m|_1$, which is arguably the most natural distance function between distributions over a finite set of states with no structure; and (ii) the Euclidean distance $\kappa(q, m) = |q - m|_2$.

Theorem 3.1. *Assume that there exists some $\delta > 0$ such that $\kappa(q, m) \geq \delta|q - m|_1$. For every game there exists $\bar{\lambda}$ such that for every $\lambda > \bar{\lambda}$, there exists a BNE in the game with intensity λ such that (i) Sender's strategy is calibrated and (ii) his payoff is CP.*

In section 2 we demonstrated this result using an example with the total-variation distance. In that example, the commitment solution was given by Sender's strategy in table 1a. Receiver buys the product after m_1 and does not buy otherwise. This strategy profile is a BNE when $\lambda \geq 1$.

Moving now to arbitrary cost functions, we show that the strong assertion of Theorem 3.1 does not hold. Consider the same example from section 2 but with the prior $p = (1/4, 1/4, 1/2)$ and the Kullback-Leibler distance given by

$$\kappa(q, m) := \sum_{s \in S} q(s) \log \frac{q(s)}{m(s)}.$$

The Kullback-Leibler distance is asymmetric in q, m , and is smooth. In the commitment solution, Sender uses a calibrated strategy with support $m_0 = (1, 0, 0)$ and $m_1 = (1/10, 3/10, 6/10)$, and Receiver buys after m_1 . However, for any intensity λ , this strategy profile is not an equilibrium: if Receiver buys after m_1 , then Sender has an incentive to deviate by announcing

m_1 more often when the state is 0.⁴

We next construct an equilibrium in which Receiver buys the product only after the following message m'_1 :

$$m'_1 = \left(\frac{e^{-1/\lambda}}{10}, \frac{1}{3} - \frac{e^{-1/\lambda}}{30}, \frac{2}{3} - \frac{e^{-1/\lambda}}{15} \right).$$

On the equilibrium path, Sender sends m'_1 with probability one when the state is 2 or 4, and with probability $1/3$ when the state is 0. The true conditional distribution given message m'_1 is

$$\tau_\sigma(m'_1) = \left(\frac{1}{10}, \frac{3}{10}, \frac{6}{10} \right).$$

Receiver is indeed willing to buy after m'_1 . Figure 3 illustrates this equilibrium. Receiver is willing to buy if his belief is in the gray area. In this equilibrium, Sender's strategy is not calibrated; he makes a strong statement m'_1 , which is in the interior of the gray area, yet the true conditional distribution $\tau_\sigma(m'_1)$ is just enough for Receiver to buy. Sender's payoff is $\frac{3}{4}\lambda \log(10e^{1/\lambda}/9 - 1/9)$, which converges to his commitment payoff as λ goes to infinity.




Figure 3: An equilibrium with the Kullback-Leibler distance

As the cost intensity increases, there exists an equilibrium in which Sender's payoff converges to his commitment payoff. The next theorem shows that this holds for a generic set of games.

Theorem 3.2. *In a generic set of games, for every $\varepsilon > 0$, there exists $\bar{\lambda}$ such that for every $\lambda > \bar{\lambda}$, there exists a BNE in the game with intensity λ such that Sender's payoff is at least $CP - \varepsilon$.*

⁴In the commitment solution, Sender announces m_1 with probability $1/3$ in state 0, and with probability 1 in states 2 and 4. If Sender deviates by increasing the probability of m_1 in state 0 to $1/3 + \alpha$, then his payoff is

$$\frac{1}{4} \left(\alpha + \frac{1}{3} \right) + \frac{3}{4} - \frac{1}{12} \lambda \left(9 \log \left(\frac{10}{3\alpha + 10} \right) + (3\alpha + 1) \log \left(10 - \frac{90}{3\alpha + 10} \right) \right).$$

The derivative of this payoff with respect to α at $\alpha = 0$ is $1/4$. Hence, Sender has an incentive to deviate for all $\lambda > 0$. This is because the miscalibration cost has no first-order impact on Sender's payoff around $\alpha = 0$.

We have shown that Sender can achieve his commitment payoff (possibly up to ε) in an equilibrium with some assumption on the cost function (Theorem 3.1) or some generic assumption on the game (Theorem 3.2). Without either assumption, Sender's optimal payoff in equilibrium, for all cost intensity, might be bounded by a value which is strictly below his commitment payoff, as shown in the following nongeneric game.

Example 2. Let S be $\{0, 1\}$ with the prior $\{p, 1 - p\}$. Let A be $\{a_0, a_1\}$. If Receiver chooses a_0 , both players' payoffs are 0. If Receiver chooses a_1 , his payoff is 0 in state 0 and -1 in state 1, and Sender's payoff is 1. This game is nongeneric. In the commitment solution, Sender uses a calibrated strategy and Receiver takes action a_1 only after message $(1, 0)$. Sender's commitment payoff is p . However, for any smooth cost function κ , in every equilibrium Receiver chooses only the safe action a_0 , so Sender's payoff is 0.

4 Proof ideas

4.1 Theorem 3.1

It is a standard revelation-principle argument that the commitment solution can be implemented via a calibrated strategy of Sender, with Receiver believing Sender's message and playing Sender's preferred action among all of Receiver's best responses to that message. We have to show that playing this profile is an equilibrium. Since Sender's strategy is calibrated, it is clear that Receiver best responds to Sender. What is difficult, however, is proving that Sender will not deviate from this strategy. The key observation is Lemma 4.1 below, which asserts that for a sufficiently high intensity, Sender is always better off using a calibrated strategy, regardless of Receiver's strategy.

Lemma 4.1. *Assume that there exists some $\delta > 0$ such that $\kappa(q, m) \geq \delta|q - m|_1$. There exists $\bar{\lambda}$ such that for every $\lambda > \bar{\lambda}$, if σ is a Sender's strategy which is not calibrated and ρ is a Receiver's strategy, then there exists a calibrated strategy of Sender that gives him a higher payoff than σ against ρ .*

This lemma captures the intuition that when the cost intensity is high, miscalibration simply is not worth it. However, the lemma is not obvious since Receiver's strategy is not continuous, and arbitrarily small miscalibrations on the part of Sender can completely change Receiver's behavior. In fact, the lemma does not hold for distance functions like the Kullback-Leibler distance. The lemma also relies on the assumption that the prior has full support.

The proof idea for Lemma 4.1 is that we can improve every miscalibrated strategy σ by changing the probability with which Sender announces each message in σ in a way that the

true conditional distribution given the message aligns with its asserted distribution. Thus the messages remain the same but the true conditional distributions given the messages change. In the proof, we use the property that, for any mixed action of Receiver, the difference between the expected payoffs for Sender under two different distributions is bounded by the total-variation distance between these distributions. This property allows us to show that any payoff gain from miscalibration is overshadowed by the miscalibration cost for a sufficiently high intensity.

4.2 Theorem 3.2

The first step of the proof is Lemma 4.2 below. It asserts that there exists a calibrated strategy of Sender under which (i) Receiver’s best response to every message is unique; and (ii) Sender’s payoff from using this strategy is arbitrarily close to his commitment payoff.

Let $BR(\cdot)$ be the correspondence from $\Delta(S)$ to A such that $BR(q)$ is the set of all of Receiver’s best actions when his belief about the state is q :

$$BR(q) = \arg \max_a \sum_s q(s)u(s, a). \quad (3)$$

Lemma 4.2. *In a generic set of games, for every $\varepsilon > 0$, there exists $x_a \in \Delta(S)$ for each $a \in A$ and a Sender’s calibrated strategy σ , such that (i) $BR(x_a) = \{a\}$ for each $a \in A$; (ii) $\text{support}(\sigma) = \{x_a\}_{a \in A}$; and (iii) if ρ is Receiver’s best response to σ , then $V_\lambda(\sigma, \rho) > CP - \varepsilon$.*

Figure 4 illustrates Lemma 4.2 for the example from section 2. Receiver decides whether to buy ($a = 1$) or not to buy ($a = 0$). He is willing to buy if his belief is in the gray area. At the border between the gray and white areas, Receiver is indifferent. We first show that we can decompose the prior p into the convex combination of beliefs (y_0, y_1) such that a is the *unique* best response to y_a . The commitment solution splits the prior p into the convex combination of π_0 and π_1 , with Receiver being indifferent between the two actions at π_1 . We construct (x_0, x_1) by taking the convex combination of (y_0, y_1) and (π_0, π_1) . Our construction ensures that Receiver has a unique best response at both x_0 and x_1 . As we move (x_0, x_1) arbitrarily close to (π_0, π_1) , Sender’s payoff from the calibrated strategy with support (x_0, x_1) is arbitrarily close to CP .

Using Lemma 4.2, we consider an auxiliary game in which, instead of announcing a distribution over states, Sender can either announce an action a or send a “silent” message. Whenever Sender announces a , he pays a cost relative to the asserted distribution x_a and Receiver must take action a . Whenever Sender sends “silent,” he pays no miscalibration cost and Receiver’s equilibrium action is a best response to his equilibrium belief when Sender




Figure 4: Decomposing p into $\{x_a\}_a$ so that $BR(x_a) = \{a\}$

sends “silent.” Sender’s payoff in this auxiliary game is at least $CP - \varepsilon$, since he can always use the strategy σ in Lemma 4.2 and pay no miscalibration cost. Moreover, for any equilibrium in this auxiliary game, the true conditional distribution given message a must be sufficiently close to x_a if the intensity λ is high. This is because the payoff is bounded and thus it does not pay for Sender to be far away from x_a .

For any equilibrium in this auxiliary game, we can construct an equilibrium in the original game as follows: Sender uses the same strategy that he used in the auxiliary game, but with two changes: (i) when he wants to say a in the auxiliary game, he now says x_a ; and (ii) when he wants to send “silent” in the auxiliary game, he now says the true conditional distribution given the “silent” message. Receiver plays a when he hears x_a ; after every other message, he plays the strategy that he played when Sender sent “silent” in the auxiliary game. Sender has no incentive to deviate since he is in the same situation as in the auxiliary game. Receiver has no incentive to deviate either: since the true conditional distribution given message x_a is sufficiently close to x_a and a is the unique best response to belief x_a , it follows that a is also a best response to a ball of beliefs around x_a .

5 Rationalizable communication

In this paper, as in the cheap-talk literature, we used equilibrium as our solution concept. We argued that Sender’s optimal equilibrium achieves the commitment solution. But the game still admits other equilibria. For example, even for a high intensity λ , the “babbling equilibrium” still exists. Under this equilibrium, Sender always declares the prior. After every possible Sender’s message, Receiver believes that the state is distributed according to the prior, and best responds to the prior.

There is, however, an unsatisfactory aspect to the babbling equilibrium in our context. Consider the example from section 2 with a high cost-intensity λ . Assume that Receiver plays according to the babbling equilibrium, and that the message $(0, 0, 1)$ —which says that

“the state is 4 with probability one”—arrives. What should Receiver do? The message is a surprise, but given the high intensity, Sender suffers a very high cost if the truth is far away from this message. It seems reasonable that Receiver will deduce that the truth is close enough to this message, in which case Receiver will respond by purchasing the product. BNE does not capture this intuition because it allows arbitrary behaviors off path. These behaviors do not have to be rationalizable. In order to incorporate this intuition in our analysis, we turn to the concept of extensive-form rationalizability (hereafter EFR; Pearce (1984), Battigalli (1997), Battigalli and Siniscalchi (2002)). EFR dispenses with the assumption that players’ beliefs are correct, but it requires that, at every point in the game, each player forms a belief that is, as much as possible, consistent with the opponent being rational.

EFR is usually defined in an environment with only countable information sets. In Sender-Receiver games, this means a countable set of messages. How to extend such a definition to Sender-Receiver games with a continuum message space is not obvious. (See Remark 1 below for a detailed discussion, as well as Friedenber (2017), in which similar issues arise.)

We say that a message m is an *atom* of Sender’s strategy σ if m is an atom of $\sigma(\cdot|s)$ for some state s . This means that under σ there is a strictly positive probability that Sender will announce m .

Definition 1. Let Σ_n be sets of Sender’s strategies and let R_n be correspondences from M to A of Receiver’s possible actions given the message defined recursively as follows:

1. Σ_0 is the set of all Sender’s strategies.
2. $R_0(m) = A$ for every message $m \in M$.
3. For every $n \geq 1$, Σ_n is the set of all $\sigma \in \Sigma_{n-1}$ that are Sender’s best responses to some strategy ρ of Receiver such that $\text{support}(\rho(\cdot|m)) \subseteq R_{n-1}(m)$.
4. For every $n \geq 1$ and every message m , if there exists some strategy $\sigma \in \Sigma_{n-1}$ such that m is an atom of σ , then $R_n(m)$ is the set of all actions a such that $a \in BR(\tau_\sigma(m))$ for such a strategy σ ; If no such σ exists, then $R_n(m) = R_{n-1}(m)$.

We let $\Sigma_\infty = \bigcap_n \Sigma_n$ and $R_\infty(m) = \bigcap_n R_n(m)$. The strategies in Σ_∞ are *Sender’s rationalizable strategies* and the actions in $R_\infty(m)$ are *Receiver’s rationalizable actions after message m* .

Theorem 5.1. *Assume that there exists some $\delta > 0$ such that $\kappa(q, m) \geq \delta|q - m|_1$. Then, in a generic set of games, there exists $\bar{\lambda}$ such that for every $\lambda > \bar{\lambda}$, Sender’s rationalizable strategies are his calibrated commitment strategies, i.e., calibrated strategies σ such that $V_\lambda(\sigma, \rho) = CP$ for some best response ρ of Receiver to σ .*

Based on Theorem 5.1, for sufficiently high λ the only rationalizable strategy for Sender in the example from section 2 is the calibrated strategy in table 1a. Receiver’s rationalizable actions after m are actions in $BR(m)$.

Remark 1. The idea is that if a message m is on the path of some strategy of Sender that was not yet eliminated, then Receiver’s rationalizable actions for this message are the best responses against the conditional beliefs over states under such strategies of Sender. If the message m is not on the path of any of Sender’s strategies that were not yet eliminated, then the message is a surprise and all Receiver’s actions which were not yet eliminated remain. Thus, EFR requires, for each message m and each Sender’s strategy σ , (i) defining whether m is on the path of σ , and (ii) defining the conditional belief $\tau_\sigma(m)$ if the message is on path.

However, when the set of messages is a continuum, it is possible that every message m is produced with probability zero. In this case, it is not obvious what it means for a message to be on path, and the conditional distributions $\tau_\sigma(m)$ are defined only almost-surely in m . More formally, there is an equivalence class of functions, called *versions* of the conditional distributions, and these versions differ from one another in a set of messages m with probability zero. When there is only one σ on the table, as is the case in the definition of a BNE, this is not a big issue since the definition does not depend on the version of the conditional distribution. But in the definition of rationalizability, at each stage of elimination there is a set of Sender’s strategies, and Receiver needs to choose versions of the conditional distributions simultaneously for all these strategies. The resulting set of rationalizable actions for Receiver will depend on which version is chosen. Our approach in this paper is to require Receiver to form a belief only when he receives a message that is produced with a strictly positive probability for some remaining strategy σ of Sender. For such a message m —a message that practically shouts “I am not a surprise”—the conditional distribution $\tau_\sigma(m)$ is unambiguously defined. As Theorem 5.1 shows, even this weak definition of EFR still narrows down significantly the possible outcomes in our environment.

6 Appendix

6.1 Preliminaries

We Let B be a bound on Sender's payoff function, so that $|v(s, a)| \leq B$ for every $s \in S, a \in A$. We extend the distance function κ to a function $\kappa : \mathbf{R}_+^S \times \Delta(S) \rightarrow \mathbf{R}_+$ given by $\kappa(\gamma q, m) = \gamma \kappa(q, m)$ for every $q \in \Delta(S)$ and $\gamma \geq 0$. We extend the best-response correspondence $BR(\cdot)$ to a correspondence from \mathbf{R}_+^S to A given by the same formula (3).

For a Sender's strategy σ we let

$$\chi_\sigma = \sum_s p(s) \sigma(\cdot | s) \quad (4)$$

be the distribution over messages induced by σ . Lemma 6.1 below is essentially Aumann and Maschler's splitting lemma (see, for example, Zamir (1992) Proposition 3.2). It says that a distribution over messages in $\Delta(S)$ is induced by some calibrated strategy if and only if its barycenter is the prior p .

Lemma 6.1. 1. For every strategy σ , we have

$$\int \tau_\sigma(m) \chi_\sigma(dm) = p. \quad (5)$$

2. If χ is a distribution over messages in $\Delta(S)$ such that $\int m \chi(dm) = p$, then there exists a calibrated strategy σ such that $\chi_\sigma = \chi$.

6.2 Proof of Lemma 4.1

Let $L = 1/\min\{p(s) : s \in S\}$. (Here we use the full-support assumption on p .) The choice of L is such that, for every $q \in \Delta(S)$, p can be "split" into a convex combination of q and some other belief $q' \in \Delta(S)$, where L bounds the weight on q' . More explicitly, for every belief q , we have

$$(1 - L|p - q|_1)q(s) \leq (1 - L|p - q|_\infty)q(s) \leq p(s),$$

for every state s . This implies that we can find some $q' \in \Delta(S)$ such that

$$p = (1 - L|p - q|_1) \cdot q + L|p - q|_1 \cdot q'. \quad (6)$$

Fix a Receiver's strategy ρ and let $\eta : \Delta(S) \times \Delta(S) \rightarrow \mathbf{R}$ be such that $\eta(q, m) = \sum_{s,a} q(s) \rho(a|m) v(s, a)$ is Sender's expected payoff when the distribution over states is q , the

message is m , and Receiver follows ρ . From the bound on v , there exists $B > 0$ such that η is bounded $|\eta(q, m)| \leq B$ for every q, m and satisfies the Lipschitz condition $|\eta(q, m) - \eta(q', m)| \leq B|q - q'|_1$ for every $q, q' \in \Delta(S)$.

For every Sender's strategy σ with induced distribution χ_σ over messages given by (4), the payoff $V_0(\sigma, \rho)$ and $d(\sigma)$ satisfy:

$$\begin{aligned} V_0(\sigma, \rho) &= \int \eta(\tau_\sigma(m), m) \chi_\sigma(dm), \text{ and} \\ d(\sigma) &\geq \delta \int |\tau_\sigma(m) - m|_1 \chi_\sigma(dm). \end{aligned} \tag{7}$$

We now fix a Sender's strategy σ , with the aim of proving that if σ is miscalibrated, then there exists a calibrated strategy which gives Sender a higher payoff against ρ . Suppose not (i.e., suppose that σ is miscalibrated and a best response to ρ). First, if σ' is any calibrated strategy, then $V_\lambda(\sigma', \rho) = V_0(\sigma', \rho) \geq -B$ and $V_\lambda(\sigma', \rho) \leq V_\lambda(\sigma, \rho) \leq B - \lambda d(\sigma)$. Therefore, we can assume that

$$d(\sigma) \leq 2B/\lambda; \tag{8}$$

otherwise, σ cannot be a best response.

Let χ_σ given by (4) be the distribution over messages induced by σ and let

$$q = \int m \chi_\sigma(dm) \tag{9}$$

be the barycenter of χ_σ . It then follows from the convexity of the norm $|\cdot|_1$, (5), (9), and (7) that

$$|p - q|_1 = \left| \int (\tau_\sigma(m) - m) \chi_\sigma(dm) \right|_1 \leq \int |\tau_\sigma(m) - m|_1 \chi_\sigma(dm) \leq \frac{1}{\delta} d(\sigma). \tag{10}$$

Let $\lambda \geq \frac{2BL}{\delta}$. It then follows from (10) and (8) that $|p - q|_1 \leq 1/L$. Consider now the calibrated strategy that is induced (via part 2 of Lemma 6.1) by the distribution

$$(1 - L|p - q|_1)\chi_\sigma + L|p - q|_1\delta_{q'} \in \Delta(\Delta(S)),$$

where q' is given by (6) and δ is the Dirac measure: with probability $L|p - q|_1$ Sender sends q' , and with the complementary probability $1 - L|p - q|_1$ Sender uses χ_σ . (Note that, from (6) and (9), it follows that the barycenter of this distribution is indeed p .) The payoff under

this strategy is given by

$$\begin{aligned}
& (1 - L|p - q|_1) \int \eta(m, m) \chi_\sigma(dm) + L|p - q|_1 \eta(q', q') \\
& \geq \int (\eta(\tau_\sigma(m), m) - B|\tau_\sigma(m) - m|_1) \chi_\sigma(dm) - 2L|p - q|_1 B \\
& \geq V_\lambda(\sigma, \rho) + \left(\lambda - \frac{B}{\delta} \right) d(\sigma) - 2L|p - q|_1 B \geq V_\lambda(\sigma, \rho) + \left(\lambda - \frac{B(1 + 2L)}{\delta} \right) d(\sigma),
\end{aligned}$$

where the first inequality follows from the Lipschitz condition and the bound on η , the second inequality from (7), and the third inequality from (10). Therefore, if $d(\sigma) > 0$ and $\lambda > B(1 + 2L)/\delta$, then the calibrated strategy defined above gives Sender a payoff strictly higher than σ does.

6.3 Proof of Theorem 3.1

It is a standard argument (following the revelation principle) that the commitment solution can be implemented with the message space $M = \Delta(S)$, a calibrated strategy σ^* , and Receiver's strategy ρ^* such that $\rho^*(\cdot|m)$ is concentrated on Sender's preferred actions among $BR(m)$. This strategy profile gives Sender a payoff of CP . Given that σ^* is calibrated, ρ^* is a best response to σ^* . To complete the proof, we need to show that Sender does not gain from deviating. We know that (i) σ^* gives Sender CP against ρ^* ; (ii) ρ^* is a best response against all calibrated strategies, so by the definition of CP , no calibrated strategy gives more than CP against ρ^* ; and (iii) by Lemma 4.1 it follows that no miscalibrated strategy gives Sender more than CP against ρ^* . Therefore, σ^* is a best response against ρ^* .

6.4 Proof of Lemma 4.2

We prove Lemma 4.2 with the help of Lemma 6.2 below. The assertion in the lemma is not new—similar arguments appear in Balkenborg, Hofbauer and Kuzmics (2015) and in Brandenburger, Danieli and Friedenberg (2017). To make the paper self-contained we provide an independent proof.

Lemma 6.2. *In a generic set of games, for every action a that is not strictly dominated for Receiver, there exists a belief $q \in \Delta(S)$ such that $BR(q) = \{a\}$.*

We first introduce the following lemma that characterizes actions which are a *unique* best response to some beliefs. It also follows from arguments in Lemma D.4 of Brandenburger, Friedenberg and Keisler (2008) and Theorem 2 of Balkenborg, Hofbauer and Kuzmics (2015). Both Lemma 6.3 and Lemma 6.2 are of course similar in spirit to Wald's complete class theorem (Kuzmics (2015)) and Pearce's notion of rationalizability (Pearce (1984)).

Lemma 6.3. *Let $a \in A$ be an action. The following conditions are equivalent:*

1. *There exists a belief $q \in \Delta(S)$ over states such that $BR(q) = \{a\}$.*
2. *There does not exist a mixture $\xi \in \Delta(A \setminus \{a\})$ over actions other than a , such that $u(s, a) \leq \sum_{c \in A \setminus \{a\}} \xi(c)u(s, c)$ for every state s .*

Proof of Lemma 6.2. For a subset $T \subseteq S$ of states and an action $a \in A$, let $u_a^T \in \mathbf{R}^T$ be given by $u_a^T(t) = u(t, a)$ for every $t \in T$. Consider the set of games such that, for every set $\bar{A} \subseteq A$ of actions and every set $T \subseteq S$ of states such that $|\bar{A}| \leq |T| + 1$, the vectors $\{u_a^T\}$ for $a \in \bar{A}$ are affinely independent. This is a generic set.⁵ In particular, if $\bar{A} = C \cup \{a\}$ for some $C \subseteq A$ and some action $a \in A \setminus C$, then if u_a^T is in the affine hull of $\{u_c^T\}$ for $c \in C$, then this affine hull is \mathbf{R}^T .

Consider such a game and let a be an action for which the condition in Lemma 6.3 is not satisfied. We prove that a is strictly dominated.

Let $\xi \in \Delta(A \setminus \{a\})$ be such that $u(s, a) \leq \sum_{c \in A \setminus \{a\}} \xi(c)u(s, c)$ for every state s . Let $C = \text{support}(\xi)$ and let $T = \{s : u(s, a) = \sum_{c \in C} \xi(c)u(s, c)\}$. If T is an empty set, then a is strictly dominated by ξ . Suppose that T is nonempty. Then, from the definition of T it follows that u_a^T is in the affine span of $\{u_c^T\}$ for $c \in C$. Therefore, from the generic assumption, the affine hull of $\{u_c^T\}$ for $c \in C$ is \mathbf{R}^T . In particular, there exists $\xi'(c) \in \mathbf{R}$ for $c \in C$ such that $\sum_{c \in C} \xi'(c) = 1$ and $\sum_{c \in C} \xi'(c)u(t, c) > u(t, a)$ for every $t \in T$.

We now choose $\varepsilon > 0$ sufficiently small such that $(1 - \varepsilon)\xi(c) + \varepsilon\xi'(c) > 0$ for every $c \in C$ and $u(s, a) < \sum_{c \in C} ((1 - \varepsilon)\xi(c) + \varepsilon\xi'(c))u(s, c)$ for every $s \notin T$. Then, $(1 - \varepsilon)\xi + \varepsilon\xi'$ is a mixture over actions for Receiver that strictly dominates a . \square

Returning to the proof of Lemma 4.2, since strictly dominated actions are not played in any equilibrium or commitment solution, we can discard them and consider a generic game with no strictly dominated strategies. We divide the proof into two claims.

Claim 1. *There exists $y_a \in \mathbf{R}_+^S \setminus \{0\}$ such that $p = \sum_{a \in A} y_a$ and $BR(y_a) = \{a\}$.*

Proof. By Lemma 6.2, for every action a there exists some belief $q_a \in \Delta(S)$ such that a is the unique best response to the belief q_a , i.e., $BR(q_a) = \{a\}$. Since p has full support, there exists a small $t > 0$ such that $p \gg \sum_{a \in A} tq_a$. Let \bar{a} be such that $\bar{a} \in BR(p - \sum_a tq_a)$. Let $y_{\bar{a}} = p - \sum_a tq_a + tq_{\bar{a}}$ and $y_a = tq_a$ for every $a \neq \bar{a}$. \square

⁵Consider the vector space of functions $u : S \times A \rightarrow \mathbf{R}$, viewed as $|S| \times |A|$ matrices. A violation of this condition means that the $(|T| + 1) \times |\bar{A}|$ matrix that is given by the corresponding $|T| \times |\bar{A}|$ minor of the payoff matrix u with an additional row $(1, \dots, 1)$ is not of full rank. Thus, the set of payoff functions u for which the condition is violated is a closed set with an empty interior and Lebusge's measure zero.

Claim 2. *There exists $x_a \in \Delta(S), \gamma_a > 0$ for each $a \in A$ such that $p = \sum_{a \in A} \gamma_a x_a$, $\sum_{a \in A} \gamma_a = 1$, $BR(x_a) = \{a\}$, and $\sum_{a,s} \gamma_a x_a(s) v(s, a) > CP - \varepsilon$.*

Proof. Let (σ, ρ) be a commitment solution such that $CP = \sum_{s,a} v(s, a) \pi_{\sigma, \rho}(s, a)$, where $\pi_{\sigma, \rho}$ is the distribution induced by the profile (σ, ρ) over $S \times A$. Let $z_a(s) = \frac{\varepsilon}{2B} y_a(s) + (1 - \frac{\varepsilon}{2B}) \pi_{\sigma, \rho}(s, a)$ where y_a is given by Claim 1 and B is the bound on Sender's payoff function v . Then, $z_a \in \mathbf{R}_+^S \setminus \{0\}$, $BR(z_a) = \{a\}$, and $\sum_{s,a} v(s, a) z_a(s) > CP - \varepsilon$. Let $x_a \in \Delta(S)$ and $\gamma_a > 0$ be such that $z_a = \gamma_a x_a$. Since $p \in \Delta(S)$, $x_a \in \Delta(S)$ for $a \in A$, and $p = \sum_{a \in A} \gamma_a x_a$, it follows that $\sum_{a \in A} \gamma_a = 1$. \square

By the splitting lemma there exists a calibrated strategy σ such that $\chi_\sigma = \sum_a \gamma_a \delta_{x_a}$. From Claim 2 the strategy σ satisfies the requirements in Lemma 4.2.

6.5 Proof of Theorem 3.2

We consider the following auxiliary game, in which Sender's set of messages is $A \cup \{\diamond\}$ where \diamond is a message that says "silent." In the auxiliary game, if Sender sends message $a \in A$, then Receiver must play action a and Sender pays a miscalibration cost relative to x_a given in Claim 2; on the other hand, if Sender sends the silent message \diamond , then Receiver can choose any action from A and Sender pays no miscalibration cost. Sender's strategies can be represented by $w = \{w_m \in \mathbf{R}_+^S\}_{m \in A \cup \{\diamond\}}$ such that $\sum_{m \in A \cup \{\diamond\}} w_m = p$, with the interpretation that $w_m(s)$ is the probability that (i) the state is s and (ii) Sender sends m . Receiver's strategies are given by elements $\rho(\cdot | \diamond) \in \Delta(A)$ (mixed actions, to be played after the silent message). The payoff to Sender in the auxiliary game under the profile $(w, \rho(\cdot | \diamond))$ is given by

$$\tilde{V}_\lambda(w, \rho(\cdot | \diamond)) = \sum_{a \in A, s \in S} (w_a(s) + w_\diamond(s) \rho(a | \diamond)) v(s, a) - \lambda \sum_{a \in A} \kappa(w_a, x_a).$$

The payoff to Receiver under this profile is given by

$$\tilde{U}(w, \rho(\cdot | \diamond)) = \sum_{a \in A, s \in S} (w_a(s) + w_\diamond(s) \rho(a | \diamond)) u(s, a).$$

In the auxiliary game, Sender has a convex, compact set of strategies and a concave payoff function (which follows from the convexity of the distance function κ), and Receiver has a finite set of pure actions. Therefore, the auxiliary game admits a Nash equilibrium. Let $w^* = \{w_m^* \in \mathbf{R}_+^S\}_{m \in A \cup \{\diamond\}}$ be Sender's strategy under the Nash equilibrium and let $\rho^*(\cdot | \diamond) \in \Delta(A)$ be Receiver's strategy.

For Sender's equilibrium strategy w^* , we let $|w_m^*| = \sum_s w_m^*(s)$ be the probability that Sender sends m . If $|w_m^*| > 0$, then the posterior distribution over states conditioned on

Sender announcing m is $w_m^*/|w_m^*|$. The following claim says that in the auxiliary game Sender does not miscalibrate too much.

Claim 3. *In the BNE of the auxiliary game, $\kappa(w_a^*/|w_a^*|, x_a) \leq 2B/\lambda$ for every $a \in A$ such that $|w_a^*| > 0$. Here, $B > 0$ is the bound on Sender's payoff function v .*

Proof. Fix $a \in A$ such that $|w_a^*| > 0$. Let w' be the strategy given by $w'_m = w_m^*$ for $m \in A \setminus \{a\}$, $w'_a = 0$, and $w'_\diamond = w_\diamond^* + w_a^*$. Therefore, under w' Sender plays like w^* except that he is silent every time he was supposed to announce a . It follows that:

$$\begin{aligned} \tilde{V}_\lambda(w', \rho(\cdot|\diamond)) - \tilde{V}_\lambda(w^*, \rho(\cdot|\diamond)) &= \sum_s w_a^*(s) \left(\sum_{a' \in A} \rho(a'|\diamond) v(s, a') - v(s, a) \right) + \lambda \kappa(w_a^*, x_a) \\ &\geq -2B|w_a^*| + \lambda \kappa(w_a^*, x_a), \end{aligned}$$

where the inequality follows from the bounds on v . The assertion follows from the fact that w' is not a profitable deviation for Sender. \square

Claim 4. *Sender's payoff in the Nash equilibrium of the auxiliary game is at least $CP - \varepsilon$.*

Proof. Sender can use the strategy $w_a = \gamma_a x_a$ for every $a \in A$ and $w_\diamond = 0$. By Claim 2, Sender's payoff is at least $CP - \varepsilon$. \square

We now define the BNE (σ^*, ρ^*) in the original game. Let $\bar{w} = w_\diamond^*/|w_\diamond^*|$ be the posterior distribution over states if Sender announces \diamond in the auxiliary game. If $|w_\diamond^*| = 0$, then \bar{w} is not defined. Sender's strategy σ^* has finite support $\{x_a\}_{a \in A} \cup \{\bar{w}\}$ and satisfies

$$p(s)\sigma^*(x_a|s) = w_a^*(s), \text{ and } p(s)\sigma^*(\bar{w}|s) = w_\diamond^*(s).$$

Thus, Sender plays the same as in the auxiliary game except that (i) instead of sending the message a as he did in the auxiliary game, he now sends x_a ; and (ii) instead of being silent as he was in the auxiliary game, he now announces \bar{w} . Receiver's strategy is given by $\rho^*(x_a) = \delta_a$ for every $a \in A$ such that $|w_a^*| > 0$, and $\rho^*(\cdot|m) = \rho^*(\cdot|\diamond)$ for any other m . Thus, Receiver's strategy is to play a when Sender announces x_a and to play the mixed action $\rho^*(\cdot|\diamond)$ otherwise.

We claim that this is the desired equilibrium. First, note that Sender's situation in the original game is the same as in the auxiliary game: when he announces x_a , Receiver plays a ; when he announces anything else, Receiver plays $\rho^*(\cdot|\diamond)$. Therefore, playing the same strategy in the original game is also an equilibrium, with the payoff at least $CP - \varepsilon$ by

Claim 4. Let λ be sufficiently high such that $a \in BR(q)$ for every $q \in \Delta(S)$ such that $\kappa(q, x_a) \leq 2B/\lambda$. Such a λ exists by the continuity of κ and the fact that a is the unique best response to belief x_a . Then, by Claim 3 it follows that, for every belief on-path, Receiver best responds to that belief.

6.6 Proof of Theorem 5.1

Lemma 4.1 implies that for $\lambda > \bar{\lambda}$, only calibrated strategies of Sender survive the first round of elimination. Since all calibrated strategies are best responses to Receiver's strategy that always chooses some fixed action $a \in A$, it follows that Σ_1 is the set of all calibrated strategies. Since any message m is an atom of some calibrated strategy of Sender, it follows that $R_2(m) = BR(m)$. Therefore, by Lemma 4.2, for every Receiver's strategy ρ that survived, Sender believes that he can get at least $CP - \varepsilon$ against ρ for every $\varepsilon > 0$. Therefore, a rationalizable strategy of Sender must be calibrated and give him CP against some strategy ρ of Receiver such that $\text{support}(\rho(\cdot|m)) \subseteq BR(m)$.

References

- Abeler, Johannes, Anke Becker, and Armin Falk.** 2014. “Representative evidence on lying costs.” *Journal of Public Economics*, 96–104.
- Abeler, Johannes, Daniele Nosenzo, and Collin Raymond.** 2017. “Preferences for truth-telling.” *Working paper*.
- Aumann, Robert J., and Michael Maschler.** 1995. *Repeated Games with Incomplete Information*. The MIT Press.
- Balkenborg, Dieter, Josef Hofbauer, and Christoph Kuzmics.** 2015. “The refined best-response correspondence in normal form games.” *International Journal of Game Theory*, 44(1): 165–193.
- Battigalli, Pierpaolo.** 1997. “On Rationalizability in Extensive Games.” *Journal of Economic Theory*, 74: 40–61.
- Battigalli, Pierpaolo, and Marciano Siniscalchi.** 2002. “Strong Belief and Forward Induction Reasoning.” *Journal of Economic Theory*, 106: 356–391.
- Battigalli, Pierpaolo, and Martin Dufwenberg.** 2007. “Guilt in Games.” *American Economic Review*, 97(2): 170–176.
- Best, James, and Daniel Quigley.** 2017. “Persuasion for the Long Run.” *Working paper*.
- Brandenburger, Adam, Alex Danieli, and Amanda Friedenberg.** 2017. “How Many Levels Do Players Reason? An Observational Challenge and Solution.” *Working paper*.
- Brandenburger, Adam, Amanda Friedenberg, and H. Jerome Keisler.** 2008. “Admissibility in Games.” *Econometrica*, 76(2): 307–352.
- Chen, Ying.** 2011. “Perturbed communication games with honest senders and naive receivers.” *Journal of Economic Theory*, 146(2): 401 – 424.
- Crawford, Vincent P., and Joel Sobel.** 1982. “Strategic Information Transmission.” *Econometrica*, 50(6): 1431–1451.
- Dawid, A. P.** 1982. “The Well-Calibrated Bayesian.” *Journal of the American Statistical Association*, 77(379): 612–613.

- Farrell, Joseph.** 1993. "Meaning and Credibility in Cheap-Talk Games." *Games and Economic Behavior*, 514–531.
- Fischbacher, Urs, and Franziska Föllmi-Heusi.** 2008. "Lies in disguise: An experimental study on cheating." *Journal of the European Economic Association*, 11: 525–547.
- Friedenberg, Amanda.** 2017. "Bargaining Under Strategic Uncertainty: The Role of Second-Order Optimism." *Working paper*.
- Gentzkow, Matthew, and Emir Kamenica.** 2016. "A Rothschild-Stiglitz Approach to Bayesian Persuasion." *American Economic Review*, 5(106): 597–601.
- Gneezy, Uri.** 2005. "Deception: The role of consequences." *American Economic Review*, 384–394.
- Gneezy, Uri, Agne Kajackaite, and Joel Sobel.** 2017. "Lying Aversion and the Size of the Lie." *American Economic Review*, *forthcoming*.
- Green, Jerry R., and Nancy L. Stokey.** 2007. "A Two-person Game of Information Transmission." *Journal of Economic Theory*, 90–104.
- Kajackaite, Agne, and Uri Gneezy.** 2017. "Incentives and cheating." *Games and Economic Behavior*, 433–444.
- Kamenica, Emir, and Matthew Gentzkow.** 2011. "Bayesian Persuasion." *American Economic Review*, 101(6): 2590–2615.
- Kartik, Navin.** 2009. "Strategic Communication with Lying Costs." *Review of Economic Studies*, 1359–1395.
- Kartik, Navin, Marco Ottaviani, and Francesco Squintani.** 2007. "Credulity, lies, and costly talk." *Journal of Economic Theory*, 134: 93–116.
- Kuzmics, Christoph.** 2015. "Abraham Wald's Complete Class Theorem and Knightian Uncertainty." *Browser Download This Paper*.
- Lehrer, Ehud, Dinah Rosenberg, and Eran Shmaya.** 2010. "Signaling and Mediation in Games with Common Interests." *Games and Economic Behavior*, 68(2): 670–682.
- Lehrer, Ehud, Dinah Rosenberg, and Eran Shmaya.** 2013. "Garbling of Signals and Outcome Equivalence." *Games and Economic Behavior*, 81: 179–191.

- Murphy, Allan H., and Robert L. Winkler.** 1987. "A General Framework for Forecast Verification." *Monthly Weather Review*, 115: 1330–1338.
- Olszewski, Wojciech.** 2015. "Calibration and Expert Testing." *Handbook of Game Theory with Economic Applications*, 4: 949–984.
- Pearce, David G.** 1984. "Rationalizable Strategic Behavior and the Problem of Perfection." *Econometrica*, 52(4): 1029–1050.
- Ranjan, Roopesh, and Tilmann Gneiting.** 2010. "Combining probability forecasts." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(1): 71–91.
- Rayo, Luis, and Ilya Segal.** 2010. "Optimal Information Disclosure." *Journal of Political Economy*, 118(5): 949–987.
- Spence, Michael.** 1973. "Job Market Signaling." *The Quarterly Journal of Economics*, 87(3): 355–374.
- Zamir, Shmuel.** 1992. "Repeated games of incomplete information: Zero-sum." *Handbook of Game Theory with Economic Applications*, 1: 109–154.