Costly Miscalibration in Communication

Yingni Guo[*] Eran Shmaya[†]

January 3, 2018

Abstract

An informed sender influences a receiver's decision by strategically disclosing information. The sender pays miscalibration cost whenever the asserted probability distribution of his message differs from the true probability distribution given the message. We show that, when the sender's miscalibration cost is sufficiently high, the sender can achieve his optimal commitment solution in an equilibrium. Under some assumption on the miscalibration cost function, the only rationalizable Sender's strategy is his strategy in the commitment solution.

JEL: D81, D82, D83

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

1 Introduction

The cheap talk literature (Crawford and Sobel (1982) and Green and Stokey (2007)) studies communication games in which an informed party (Sender) sends a message to another party (Receiver) whose action affects the payoffs of both. Sender's message has no impact on his payoff, and can be arbitrarily different from the real description of the world. In reality, Sender's talk is not completely cheap. He may be concerned about evidence that might show up in the future and contradict his message. He may have an intrinsic aversion to stating falsehoods. In this paper, we develop a model that incorporates Sender's preferences for being truthful.

In our model, Sender's messages are probabilistic assertions. For example, a defendant is either guilty or innocent. A prosecutor can say messages like "the defendant is guilty

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

[†]MEDS, Kellogg, Northwestern University, e-shmaya@kellogg.northwestern.edu.

with probability 50%" or "the defendant is guilty with probability 100%." Therefore, each message is an asserted distribution over some state space. Sender incurs some cost, which we call miscalibration cost, if the asserted distribution is different from the true conditional distribution given the message.

The miscalibration cost is a proxy for unmodeled things that happen outside of Sender's interaction with Receiver. Consider the defendants that the prosecutor claimed to be guilty with 50% probability. If only 30% of them are indeed guilty, then the prosecutor pays some cost that depends on the discrepancy between the asserted distribution (50%) and the true conditional distribution (30%) given this assertion. Thus, the cost enables us to incorporate Sender's consideration about the validity of his probabilistic assertions and, in particular, about their consistency with the statistical evidence that might show up later, without explicitly modeling the statistical test. This abstraction is useful in situations where Sender's assertions might be tested in the future, but it is not known how, when and by whom, what kind of statistical evidence will be available and what evidence will have actually been realized. In addition, the abstraction offers a tractable framework to analyze how costly miscalibration, i.e., Sender's caution not to be too far from the truth, affects communication.[1]

Our cost function depends on a parameter which we call the cost intensity. The cost intensity measures how severely Sender is punished for discrepancy between his messages and the true distributions. When the cost intensity is zero, the game is a cheap talk game. We show that when the cost intensity is sufficiently high, Sender's optimal payoff in equilibria is arbitrarily close to the payoff he could get if he were able to commit to a strategy. The results apply to a wide class of distance functions which we use to measure the discrepancy between Sender's messages and the true distributions.

The results connect our model to the Bayesian persuasion literature which studies sender-receiver games with commitment on Sender's side (Rayo and Segal (2010) and Kamenica and Gentzkow (2011)). In that literature, Sender is able to commit to any strategy. This is as if Sender's strategy is observed by Receiver. Alternatively, commitment means that whatever description Sender gives about the state of the world, the description must be true. In our model, as in the cheap talk literature, Receiver does not observe Sender's strategy. Our model hence provides a foundation for the commitment assumption in the Bayesian persuasion literature. From this aspect, our paper is also related to Best and Quigley (2017) who provide a foundation for the commitment assumption by analyzing repeated interactions

[1] There is also an active literature showing that people have intrinsic aversion to stating falsehoods, regardless of whether they are caught. See for instance Gneezy (2005), Fischbacher and Föllmi-Heusi (2008), Abeler, Becker and Falk (2014), Gneezy, Kajackaite and Sobel (2017), Abeler, Nosenzo and Raymond (2017).

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 and the commitment is not absolute either.

Our paper is related to the literature on strategic communication with lying cost (Kartik, Ottaviani and Squintani (2007) and Kartik (2009)). In these papers, the set of messages is the state space, so Sender always declares a state. When the lying cost grows, the equilibrium is arbitrarily close to a fully revealing equilibrium. Our model, with a larger message space which allows the Sender to make probabilistic statements, delivers a different asymptotic result: There is an equilibrium which is arbitrarily close to the commitment solution. Thus, our model delivers a different intuition than Kartik, Ottaviani and Squintani (2007) and Kartik (2009). They show that when lying is costly, Sender reveals all his information. We show that when lying is costly, Sender obtains commitment power not to say a lie, but he is not forced to reveal all his information.

Our paper is also related to the signaling literature (Spence (1973)). In that literature, Sender's payoff depends on the state, on the action by Receiver, and directly on the message chosen by himself. Because we focus on Sender's preferences for being truthful, the messages have an impact on Sender's payoff only when a message's asserted distribution differs from the true conditional distribution given the message. This payoff structure distinguishes ours from the signaling models. Moreover, unlike the signaling literature as well as Kartik (2009) and Kartik, Ottaviani and Squintani (2007), Sender's miscalibration cost depends on the entire strategy, not just the realized state and the message.

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 to buy the product or not. Receiver's payoff is $(s - 3)$ if he buys the product and zero otherwise. Sender gets commission payoff 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 says with positive probability, we let $\beta(m)$ denote the conditional distribution over states by Bayes' rule. We also refer to $\beta(m)$ as the true conditional distribution given m. (The true conditional distribution given a message depends on Sender's strategy. We omit this dependence in notation $\beta(m)$ in this section.)

Consider a Sender's strategy in Table 1. Sender says 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 saying m_0, the online platform says that the product is a lemon. By saying m_1, it says that the quality is weakly above 2. Each column shows the probabilities with which Sender says m_0 or m_1 in each state.

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

Table 1: A calibrated strategy

message \ 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 2: A miscalibrated strategy

Given the uniform prior and the strategy in Table 1, the posterior after message m_0 is $\beta(m_0) = (1, 0, 0)$. The posterior after message m_1 is $\beta(m_1) = (0, 1/2, 1/2)$. 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 choose to reveal all information by announcing the messages $(1, 0, 0), (0, 1, 0), (0, 0, 1)$ under the states $0, 2, 4$, respectively.

Consider another Sender's strategy in Table 2. Sender still uses the same messages m_0, m_1 as before. The only difference occurs when the state is 0: Sender now says m_0 and m_1 with equal probability. In this case, the posterior after message m_0 is $\beta(m_0) = (1, 0, 0)$, which still coincides with m_0. However, the posterior after message m_1 is $\beta(m_1) = (1/5, 2/5, 2/5)$, 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 $\beta(m_1)$ given m_1 differs from the asserted distribution m_1, this strategy is not calibrated. If Sender uses this strategy to make recommendations, then the distribution of the products about which Sender says m_1 is in fact $\beta(m_1)$. 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, Sender might be caught.

In our model, whenever Sender says a message whose true conditional distribution differs from its asserted distribution, Sender pays some cost which we call miscalibration cost. The bigger the difference is, the higher the cost will be. In this section, we use the total-variation

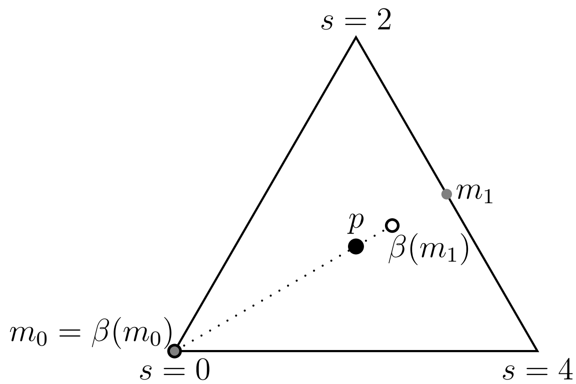

Figure 1: The miscalibrated strategy in Table 2

distance $\frac{1}{2}|m-\beta(m)|_1$ to measure how much a message m differs from the true conditional distribution $\beta(m)$ given this message. For the strategy in Table 2, the total-variation distance between m_1 and $\beta(m_1)$ is 1/5. Since Sender says message m_1 with probability 5/6, Sender pays the following miscalibration cost for using this miscalibrated strategy:

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

where $\lambda \geqslant 0$ is a parameter which measures the intensity of miscalibration cost. If Sender uses a calibrated strategy, the asserted distribution and the true conditional distribution of his message always coincide, so he pays no miscalibration cost.

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

$$\sigma(m_1)=\begin{cases}\lambda, \text{ for } 0<\lambda<1,\\ 1, \text{ for } \lambda\geqslant 1,\end{cases}$$

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

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

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

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

Since $\sigma(m_1) \leqslant \lambda$ for any $\lambda > 0$, Sender prefers to choose $\alpha = 0$. It can be shown that other deviations are not profitable either.

It is easy to verify that this is Sender's most preferred equilibrium. For sufficiently high $\lambda \geqslant 1$, Sender achieves the same payoff as if he could commit to a strategy and pay no miscalibration cost. 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 say m_1 for sure not only in state 2 and 4 but also in state 0. In other words, 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 say m_1 in state 0. This compensates the lack of commitment power on Sender's side. Sender's equilibrium payoff $2\min\{\lambda,1\}/3$ increases in λ, which can be interpreted as the proxy of 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 κ from S to M: When the state is s, Sender randomizes a message from $\kappa(\cdot|s)$. For a Sender's strategy κ, let $\beta_\kappa : M \to \Delta(S)$ be such that $\beta_\kappa(m)$ is the conditional distribution over states given m. (We suppressed the dependence of β_κ on the prior p when no confusion arises.) We sometimes use the notation $\beta_\kappa(s|m)$ for $\beta_\kappa(m)(s)$, both of which denote the conditional probability of s given m.

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

$$\beta_\kappa(s|m) = \frac{p(s)\kappa(m|s)}{\sum_{s'} p(s')\kappa(m|s')},$$

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

In Section 3.1 we consider Sender's strategies with message space $M = \Delta(S)$, and introduce the concepts of calibrated strategies and miscalibration. 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 that Sender achieves his commitment payoff in an equilibrium when the miscalibration cost is sufficiently high.

3.1 Calibrated Sender's 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 $\beta_\kappa(m)$ as the true conditional distribution over states given a message m. We say that a Sender's strategy κ is *calibrated* if $\beta_\kappa(m) = m$, a.s.. 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 (Dawid (1982), Murphy and Winkler (1987), Ranjan and Gneiting (2010)) for two related ideas: the purely probabilistic sense which we use here, and the idea of calibration with the data, that a probabilistic forecast (or a message in our terminology) matches the realized distribution of states on those times in which the forecast was given.[2] The following example illustrates the definition of "calibrated strategies" in the familiar context of recommendation letters.

Example 1. Assume that the set of states, representing possible ranks of job candidates is $S = \{1, 2, \ldots, 10\}$, with a uniform prior. Consider two senders (references) I and J who use the following strategies:

– Reference I, when writing a letter to a candidate of rank s, says that the candidate's rank is uniform in $\{1, \ldots, s\}$.

– Reference J, when writing a letter to a candidate of rank s, says that the candidate's

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

rank is uniform in $\{1, \dots, 5\}$ when $s \in \{1, \dots, 5\}$ and says that the candidate's rank is uniform in $\{6, \dots, 10\}$ when $s \in \{6, \dots, 10\}$

In our terminology, Reference I is not calibrated and Reference J is calibrated. When the candidate is of rank 5, both references deliver the message "uniform over $\{1, \dots, 5\}$." But, while the empirical distribution of ranks of candidates about whom Reference J delivers this message is indeed uniform over $\{1, \dots, 5\}$, so that the message can be taken at face value, the candidates about whom Reference I delivers this same message are all of rank 5. This example shows why, unlike Kartik (2009) and the signaling literature, Sender's miscalibration cost depends on the entire strategy, not just the realized state and the message. □

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

$$d(\kappa) = \sum_s p(s) \int \rho(\beta_\kappa(m), m) \; \kappa(\mathrm{d}m|s) \tag{1}$$

be the expected distance between an asserted distribution over states (given by the message m) and the true conditional distribution given the message under the strategy κ. Note that κ is calibrated if and only if $d(\kappa) = 0$.

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 based on the message.

Let A be a finite set of *actions*. Let $v, u : S \times A \to \mathbf{R}$ be, respectively, Sender's and Receiver's payoff functions. A Receiver's strategy is given by a Markov kernel σ from M to A, with the interpretation that Receiver randomizes an action from $\sigma(\cdot|m)$ after a message m. Let $BR(\cdot)$ be the correspondence from $\Delta(S)$ to A such that $BR(q)$ is the set of all Receiver's best actions when his belief about the state is q:

$$BR(q) = \arg\max_a \sum_s q(s) u(s, a). \tag{2}$$

A Receiver's strategy σ is *a best response to Sender's strategy* κ if [3]

$$\text{support}(\sigma(\cdot|m)) \subseteq BR(\beta_\kappa(m)) \text{ a.s..}$$

For a profile (κ, σ) of Sender's and Receiver's strategies, we denote by $\pi_{\kappa,\sigma} \in \Delta(S \times A)$ the induced distribution over states and actions when the players follow this profile. The payoff to Sender under strategy profile (κ, σ) is given by

$$V_0(\kappa, \sigma) = \int v \, \mathrm{d}\pi_{\kappa,\sigma}.$$

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 $\lambda \geqslant 0$.

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

We now turn to defining a cheap talk equilibrium. Here, the 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. In addition, following the cheap talk literature, we work with a perfect Bayesian equilibrium, which specifies Receiver's belief at the interim stage, after he observes the message m, and his action given that belief. In an equilibrium, Receiver's interim beliefs will be correct and his interim response will be optimal.

Formally, a *Receiver's interim belief* is a function $\mu : M \to \Delta(S)$. We say that Receiver's interim belief μ is *correct with respect to a Sender's strategy* κ if $\mu(m) = \beta_\kappa(m)$, a.s.. A *Receiver's interim response* is a Markov kernel α from $\Delta(S)$ to A. We say that Receiver's interim response α is *optimal* if $\text{support}(\alpha(q)) \subseteq BR(q)$ for every interim belief q. Note that every pair (μ, α) generates a strategy σ of Receiver such that $\sigma(a|m) = \alpha(a|\mu(m))$ for every message m and action a. If μ is correct with respect to some Sender's strategy κ and α is optimal, then the induced σ is a best response to κ.

Definition 1. A *perfect Bayesian equilibrium (PBE) with cheap talk* is given by $(\kappa^*, (\mu^*, \alpha^*))$ such that:

[3]The term a.s. means almost surely with respect to the probability distribution over messages induced by κ. (Recall that β_κ is defined up to a set of messages with probability zero.)

1. Sender best responds to Receiver's behavior: For every κ it holds that $V_0(\kappa^*, \sigma^*) \geqslant V_0(\kappa, \sigma^*)$, where σ^* is Receiver's strategy generated by (μ^*, α^*).
2. Receiver's interim belief μ^* is correct with respect to κ^*.
3. Receiver's interim response α^* is optimal.

When the message space is $M = \Delta(S)$, it is a standard argument, following the revelation principle, that any commitment solution and any cheap talk equilibrium can be implemented with calibrated Sender's strategies.

3.3 Costly miscalibration

We now introduce PBE in the game with costly miscalibration. The definition is the same as Definition 1 except that Sender's payoff is now given by

$$V_\lambda(\kappa, \sigma) = V_0(\kappa, \sigma) - \lambda d(\kappa). \tag{3}$$

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

Definition 2. A *perfect Bayesian equilibrium (PBE)* is given by $(\kappa^*, (\mu^*, \alpha^*))$ such that:

1. Sender best responds to Receiver's behavior: For every κ it holds that $V_\lambda(\kappa^*, \sigma^*) \geqslant V_\lambda(\kappa, \sigma^*)$ where σ^* is Receiver's strategy generated by (μ^*, α^*).
2. Receiver's interim belief μ^* is correct with respect to κ^*.
3. Receiver's interim response α^* is optimal.

Proposition 3.1. *Sender's payoff under any PBE is at most CM.*

Proof. Let $(\kappa^*, (\mu^*, \alpha^*))$ be a PBE, and let σ^* be Receiver's strategy induced by (μ^*, α^*). Then

$$V_\lambda(\kappa^*, \sigma^*) \leqslant V_0(\kappa^*, \sigma^*) \leqslant CM,$$

where the first inequality follows from (3) and the second from the definition of CM and the fact that σ^* is Receiver's best response to κ^*. □

We now turn to our main results 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 $\rho(q,m)$ has a kink at $q=m$. This assumption holds for the total-variation distance $\rho(q,m)=\frac{1}{2}|q-m|_1$, which is arguably the most natural distance function between distributions over a finite set of states with no structure, and for the Euclidean distance $\rho(q,m)=|q-m|_2$.

Theorem 3.1. *Assume that there exists some $\delta>0$ such that $\rho(q,m)\geqslant\delta|q-m|_1$. Then for every game there exists $\bar{\lambda}$ such that for every $\lambda>\bar{\lambda}$ there exists a PBE in the game with intensity λ such that Sender's strategy is calibrated and his payoff is CM.*

In Section 2 we demonstrated this result using an example with $S=\{0,2,4\}$ and the total-variation distance. In that example, the commitment solution was given by Sender's strategy in Table 1. Receiver buys the product after m_1 and does not buy otherwise. This strategy profile is a PBE when $\lambda\geqslant 1$.

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

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

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 $m_1=(1/10,3/10,6/10)$. It is easy to see that, for any intensity λ this strategy profile is not an equilibrium: If Receiver buys under m_1 then Sender has incentive to deviate a bit and announce the message m_1 with some probability when the state is 0.[4]

We next construct an equilibrium in which Receiver buys the product only after the following message m_1':

$$m_1'=\left(\frac{e^{-1/\lambda}}{10},\frac{1}{3}-\frac{e^{-1/\lambda}}{30},\frac{2}{3}-\frac{e^{-1/\lambda}}{15}\right).$$

[4]In the commitment solution, Sender says m_1 with probability 1/3 in state 0 and probability 1 in state 2 or 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 incentive to deviate for all $\lambda>0$. Note that the miscalibration cost has no first-order impact on Sender's payoff around $\alpha=0$.

On the equilibrium path, Sender says 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

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

Receiver is indeed willing to buy after m'_1. Figure 2 illustrates this equilibrium. Receiver is willing to buy if his belief is in the light-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 light-gray area, yet the true conditional distribution $\beta_\kappa(m'_1)$ is just enough for Receiver to buy. Sender's payoff is $\frac{3}{4}\lambda \log\left(10e^{1/\lambda}/9 - 1/9\right)$, which converges to Sender's commitment payoff as λ goes to infinity.

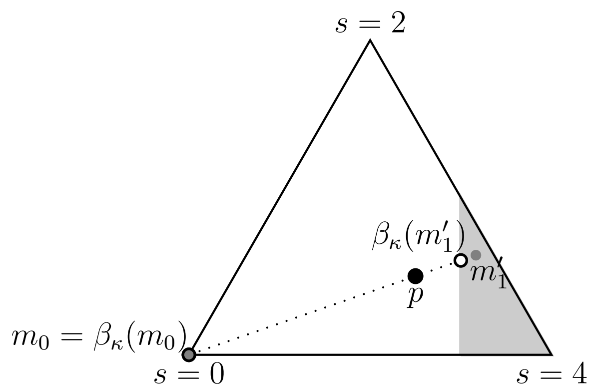

Figure 2: An equilibrium with the Kullback-Leibler distance

As the intensity of miscalibration cost increases, there exists an equilibrium in which Sender's payoff converges to his commitment payoff. The next theorem shows that this holds in general. In Theorem 3.2 and later in Lemma 4.2, a generic game means a game that has no weakly dominated actions that are not strictly dominated for Receiver. (Example 2 in Section 6.5 shows that Theorem 3.2 does not hold for nongeneric games.)

Theorem 3.2. *For a generic game it holds that for every $\varepsilon > 0$ there exists $\bar{\lambda}$ such that for every $\lambda > \bar{\lambda}$ there exists a PBE in the game with intensity λ such that Sender's payoff is at least $CM - \varepsilon$.*

4 Proof ideas

4.1 Theorem 3.1

As we already mentioned, 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 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. The difficulty is to prove that Sender will not deviate from this strategy. The key observation is the following Lemma 4.1, which asserts that for 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 $\rho(q, m) \geqslant \delta |q - m|_1$. There exists $\bar{\lambda}$ such that when $\lambda > \bar{\lambda}$, if κ is a Sender's strategy which is not calibrated and σ is a Receiver's strategy, then there exists a calibrated strategy of Sender that gives him a higher payoff than κ against σ.*

The 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 can completely change Receiver's behavior. In fact, the lemma does not hold for distance functions like the Kullback-Leibler distance.

The proof idea for Lemma 4.1 is that we can improve every miscalibrated strategy κ by changing the probability in which Sender says each message of κ 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 very property allows us to show that any payoff gain from miscalibration is overshadowed by the miscalibration cost for sufficiently high intensity.

The formal proofs are in Section 6.

4.2 Theorem 3.2

The first step of the proof is the following Lemma 4.2. It asserts that Sender can achieve payoffs arbitrarily close to his commitment payoff using a calibrated strategy under which

Receiver's best response to every message is unique.

Lemma 4.2. *For a generic game, for every $\varepsilon > 0$, there exists $x_a \in \Delta(S)$ for each $a \in A$ and a Sender's calibrated strategy κ such that (i) $BR(x_a) = \{a\}$ for each $a \in A$, (ii) support(κ) $= \{x_a\}_{a \in A}$, and (iii) if σ is Receiver's best response to κ then $V_\lambda(\kappa, \sigma) > CM - \epsilon$.*

Figure 3 illustrates Lemma 4.2 for the example in 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 of 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 CM.

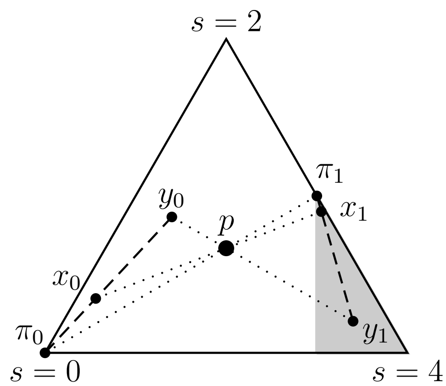

Figure 3: Decomposing p into $\{x_a\}_a$ so that $BR(x_a) = \{a\}$

Using Lemma 4.2, we consider an auxiliary game in which Sender, instead of announcing a distribution over states, can either announce an action a or stay silent. Whenever Sender says a, he pays cost relative to the asserted distribution x_a and Receiver must take action a. Whenever Sender stays silent, he pays no miscalibration cost and Receiver's equilibrium action is a best response to his equilibrium belief when Sender stays silent. Sender's payoff in this auxiliary game is at least $CM - \varepsilon$, since he can always use the strategy κ 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. Sender uses the same strategy 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 be silent, he now says the true conditional distribution given the silence message. Receiver plays a when he hears x_a and plays the strategy which was played when Sender was silent in the auxiliary game after every other message. 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 distribution conditional on message x_a is sufficiently close to x_a and a is the unique best response to belief x_a, a is also a best response to a ball of beliefs around x_a.

5 Discussion

5.1 Rationalizable communication

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

There is, however, an unsatisfactory aspect to the babbling equilibrium in our context. Consider the online platform example from Section 2 with high cost intensity λ. Assume that Receiver plays according to the babbling equilibrium and 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 is suffering 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. PBE does not capture this intuition because it allows arbitrary beliefs "off path." These beliefs do not have to be rationalizable. In order to incorporate this intuition in our analysis, we turn to the concept of extensive-form rationalizability (EFR hereafter, Pearce (1984), Battigalli (1997), Battigalli and Siniscalchi (2002)). EFR dispenses with the assumption that players' beliefs are correct, but it requires that each player, at every point in the game, 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. This means a countable set of messages in Sender-Receiver games. Extending the definition to Sender-Receiver games with a continuum message space is not obvious. (See Remark 1 below for a detailed discussion and Friedenberg (2017) in which similar issues arise.) We say that a message m is an *atom* of Sender's strategy κ if m is an atom of $\kappa(\cdot|s)$ for some state s. This means that under κ there is a strictly positive probability that Sender will say m.

Definition 3. Let K_n be sets of Sender's strategies and let Σ_n be correspondences from M to A of possible Receiver's actions defined recursively as follows:

1. K_0 is the set of all Sender's strategies.
2. $\Sigma_0(m) = A$ for every message $m \in M$.
3. For every $n \geqslant 1$, K_n is the set of all $\kappa \in K_{n-1}$ that are Sender's best responses to some Receiver's strategy σ such that $\text{support}(\sigma(\cdot|m)) \subseteq \Sigma_{n-1}(m)$.
4. For every $n \geqslant 1$ and every message m, $\Sigma_n(m)$ is the set of all actions a such that $a \in BR(\beta_\kappa(m))$ for some strategy $\kappa \in K_{n-1}$ and m is an atom of κ. If no such κ exists then $\Sigma_n(m) = \Sigma_{n-1}(m)$.

We let $K_\infty = \cap_n K_n$ and $\Sigma_\infty(m) = \cap_n \Sigma_n(m)$. The strategies in K_∞ are *Sender's rationalizable strategies* and the actions in $\Sigma_\infty(m)$ are *Receiver's rationalizable actions after message* m.

Theorem 5.1. *Assume that there exists some $\delta > 0$ such that $\rho(q, m) \geqslant \delta |q - m|_1$. There exists $\bar{\lambda}$ such that for every $\lambda > \bar{\lambda}$ Sender's rationalizable strategies are his calibrated commitment strategies, i.e., calibrated strategies κ such that $V_\lambda(\kappa, \sigma) = CM$ for some Receiver's best response σ to κ.*

Based on Theorem 5.1, in the example from Section 2, the only rationalizable strategy for Sender is the calibrated strategy in Table 1. Receiver's rationalizable actions after m are actions in $BR(m)$.

Remark 1. The idea is that if a message m is on path of some Sender's strategy 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 Sender's strategies. If the message m is not on 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 κ, defining whether m is on path of κ and, if the message is on path, the conditional belief $\beta_\kappa(m)$.

However, when the set of messages is a continuum, it is typical 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 $\beta_\kappa(m)$ are only defined 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 on a set of messages m with probability zero.) When there is only one κ on the table, as is the case of Definition 2 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 the elimination there is a set of Sender's strategies, and Receiver needs to simultaneously choose versions of the conditional distributions for all these strategies. The resulting set of rationalizable actions for Receiver will depend on the choice of the version. 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 Sender's strategy κ. For such a message m— a message that practically shouts "I am not a surprise," the conditional distribution $\beta_\kappa(m)$ is unambiguously defined. As Theorem 5.1 shows, even under this weak definition of EFR, it still narrows down the possible outcomes significantly in our environment.

6 Proofs

6.1 Preliminaries

By a generic game we mean a game that has no weakly dominated actions that are not strictly dominated for Receiver. Since strictly dominated actions are not played in any equilibrium or commitment solution, we can throw them away and consider a game with no weakly dominated strategies.

We extend the distance function ρ to a function $\rho : \mathbf{R}_+^S \times \Delta(S) \to \mathbf{R}_+$ given by $\rho(\gamma q, m) = \gamma\rho(q, m)$ for every $q \in \Delta(S)$ and $\gamma \geqslant 0$. We extend the best-response correspondence $BR(\cdot)$ to a correspondence from $\mathbf{R}_+^S$ to A given by the same formula (2).

For a Sender's strategy κ we let

$$\chi_\kappa = \sum_s p(s)\kappa(\cdot|s) \tag{4}$$

be the distribution over messages induced by κ. Lemma 6.1 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 κ it holds that*

$$\int \beta_\kappa(m)\chi_\kappa(\mathrm{d}m) = p. \tag{5}$$

2. If χ is a distribution over messages in $\Delta(S)$ such that $\int m\ \chi(\mathrm{d}m) = p$ then there exists a calibrated strategy κ such that $\chi_\kappa = \chi$.

6.2 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 κ^* and Receiver's strategy α^* such that $\alpha^*(\cdot|m)$ is concentrated on Sender's preferred actions among $BR(m)$. In particular, if we let $\mu^*(m) = m$ then $(\kappa^*, (\mu^*, \alpha^*))$ satisfies Conditions 2 and 3 of Definition 2 and gives Sender payoff CM. To complete the proof, we need to show that Sender does not gain from deviating. We know that (i) κ^* gives Sender CM against α^*, (ii) α^* is a best response against all calibrated strategies, so by the definition of CM no calibrated strategy gives more than CM against α^*, (iii) by Lemma 4.1 it follows that no miscalibrated strategy gives Sender more than CM against α^*. Therefore, κ^* is a best response against α^*.

6.3 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 it holds that

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

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

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

(Here, we implicitly assume that $1 - L|p - q|_1 > 0$. We show later in the proof that it is without loss to focus on q such that $1 - L|p - q|_1 > 0$.)

Fix a Receiver's strategy σ and let $\eta : \Delta(S) \times \Delta(S) \to \mathbf{R}$ be such that $\eta(q, m) = \sum_{s,a} q(s)\sigma(a|m)v(s, a)$ is Sender's expected payoff when the distribution over states is q, the message is m and Receiver follows σ. From the bound on v, there exists $B > 0$ such that η is bounded $|\eta(q, m)| \leqslant B$ for every q, m and satisfies the Lipschitz condition $|\eta(q, m) - \eta(q', m)| \leqslant B|q - q'|_1$ for every $q, q' \in \Delta(S)$.

For every Sender's strategy κ with induced distribution χ_κ over messages given by (4), the payoff $V_0(\kappa, \sigma)$ and the distance $d(\kappa)$ under (κ, σ) can be written as

$$\begin{aligned} V_0(\kappa, \sigma) &= \int \eta(\beta_\kappa(m), m)\ \chi_\kappa(\mathrm{d}m), \text{ and} \\ d(\kappa) &\geqslant \delta \int |\beta_\kappa(m) - m|_1\ \chi_\kappa(\mathrm{d}m). \end{aligned} \tag{7}$$

We now fix a Sender's strategy κ, with the aim of proving that if κ is miscalibrated then it cannot be a best response. Suppose not. First, if κ' is any calibrated strategy then $V_\lambda(\kappa', \sigma) = V_0(\kappa', \sigma) \geqslant -B$ and $V_\lambda(\kappa', \sigma) \leqslant V_\lambda(\kappa, \sigma) \leqslant B - \lambda d(\kappa)$. Therefore we can assume that

$$d(\kappa) \leqslant 2B/\lambda, \tag{8}$$

otherwise κ cannot be a best response.

Let χ_κ given by (4) be the distribution over messages induced by κ and let

$$q = \int m\ \chi_\kappa(\mathrm{d}m), \tag{9}$$

be the barycenter of χ_κ. Then it follows from the convexity of the norm $|\cdot|_1$, (5) and (9) that

$$|p - q|_1 = \left| \int (\beta_\kappa(m) - m)\ \ \chi_\kappa(\mathrm{d}m) \right|_1 \leqslant \int |\beta_\kappa(m) - m|_1\ \chi_\kappa(\mathrm{d}m) \leqslant \frac{1}{\delta} d(\kappa). \tag{10}$$

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

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

where q' is given by (6) and $\delta_.$ is the Dirac measure: With probability $L|p - q|_1$ Sender says q', and with the complementary probability $1 - L|p - q|_1$ Sender uses χ_κ. (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

$$
\begin{aligned}
&(1-L|p-q|_1)\int \eta(m,m)\ \chi_\kappa(\mathrm{d}m) + L|p-q|_1\eta(q',q')\\
\geqslant &\int (\eta(\beta_\kappa(m),m) - B|\beta_\kappa(m)-m|_1)\ \chi_\kappa(\mathrm{d}m) - 2L|p-q|_1B\\
\geqslant &V_\lambda(\kappa,\sigma) + \left(\lambda - \frac{B}{\delta}\right) d(\kappa) - 2L|p-q|_1B \geqslant V_\lambda(\kappa,\sigma) + \left(\lambda - \frac{B(1+2L)}{\delta}\right) d(\kappa),
\end{aligned}
$$

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 $d(\kappa) > 0$ and $\lambda > B(1+2L)/\delta$ then the calibrated strategy defined above gives Sender a strictly better payoff than κ.

6.4 Proof of Lemma 4.2

We divide the proof into several steps.

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. For every action a, since a is not weakly dominated, 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 $a \neq \bar{a}$. □

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 $\sum_{a,s} \gamma_a x_a(s) v(s,a) > CM - \varepsilon$.*

Proof. Let (κ,σ) be a commitment solution such that $CM = \sum_{s,a} v(s,a)\pi_{\kappa,\sigma}(s,a)$ where $\pi_{\kappa,\sigma}$ is the distribution induced by the profile (κ,σ) over $S \times A$. Let $z_a(s) = \varepsilon' y_a(s) + (1-\varepsilon')\pi_{\kappa,\sigma}(s,a)$ where y_a are given by Claim 1 and $\varepsilon' > 0$ is a sufficiently small number. Then $z_a \in \mathbf{R}^S_+ \setminus \{0\}$ and a is the unique best response to z_a. 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 $\sum_{a\in A} \gamma_a = 1$. □

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

6.5 Proof of Theorem 3.2

We now 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 says message $a \in A$ then Receiver must play action a and Sender pays miscalibration cost relative to x_a given in Claim 2, while if Sender says the silent message $\diamond$, Receiver can choose any action from A and there is 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 the state is s and Sender says m. Receiver's strategies are given by elements $\sigma(\cdot|\diamond) \in \Delta(A)$ (mixed actions, to be played after the silent message). The payoff to Sender under the profile $(w, \sigma(\cdot|\diamond))$ is given by

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

The payoff to Receiver under this profile is given by

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

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 ρ), and Receiver has a finite set of pure actions. Therefore a PBE exists. Let $w^* = \{w_m^* \in \mathbf{R}_+^S\}_{m \in A \cup \{\diamond\}}$ be Sender's strategy under the PBE and let $\sigma^*(\cdot|\diamond) \in \Delta(A)$ be Receiver's strategy which is a best responds to some belief $\bar{w} \in \Delta(S)$.

For Sender's equilibrium strategy w^*, we let $|w_m^*| = \sum_s w_m^*(s)$ be the probability that Sender says m. If $|w_m^*| > 0$ then the posterior distribution over states conditioned on Sender saying m is $w_m^*/|w_m^*|$. Note that from the equilibrium property it follows that $\bar{w} = w_\diamond^*/|w_\diamond^*|$ if $|w_\diamond^*| > 0$. The following claim says that in the auxiliary game Sender does not miscalibrate too much.

Claim 3. *In the PBE of the auxiliary game, it holds that $\rho(w_a^*/|w_a^*|, x_a) \leqslant B/\lambda$ for every $a \in A$ such that $|w_a^*| > 0$. Here, $B > 0$ is a constant.*

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 $w'_\diamond = w_\diamond^* + w_a^*$. So under w' Sender plays like w^* except that he is silent every time he was supposed to say a. Then

$$\tilde{V}_\lambda(w', \sigma(\cdot|\diamond)) - \tilde{V}_\lambda(w^*, \sigma(\cdot|\diamond)) = \sum_s w^*_a(s)\left(\sum_{a'\in A}\sigma(a'|\diamond)v(s,a') - v(s,a)\right) + \lambda\rho(w^*_a, x_a)$$
$$\geqslant -B|w^*_a| + \lambda\rho(w^*_a, x_a),$$

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

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

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

Recall that $\bar{w}$ is Receiver's belief in the auxiliary game when Sender says $\diamond$. We now define the PBE $(\kappa^*, (\mu^*, \alpha^*))$ in the original game as follows: κ^* has a finite support $\{x_a\}_{a\in A} \cup \{\bar{w}\}$ and satisfies

$$p(s)\kappa^*(x_a|s) = w^*_a(s), \text{ and } p(s)\kappa(\bar{w}|s) = w^*_\diamond(s).$$

Thus, Sender plays the same as in the auxiliary game except that when he wants to send the message a in the auxiliary game he now says x_a and when he wants to be silent he now says $\bar{w}$. Receiver's belief is given by $\mu^*(x_a) = w^*_a/|w^*_a|$ for every $a \in A$ such that $|w^*_a| > 0$ and $\mu(m) = \bar{w}$ otherwise. Receiver's response is given by $\alpha^*(w^*_a/|w^*_a|) = a$ and $\alpha^*(\bar{w}) = \sigma^*(\cdot|\diamond)$. Thus, Receiver's strategy is to play a when Sender says x_a and to play the mixed action $\sigma^*(\cdot|\diamond)$ 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 says x_a Receiver plays a and when he says anything else Receiver plays $\sigma^*(\cdot|\diamond)$. Therefore playing the same strategy in the original game is also an equilibrium, with payoff at least $CM - \varepsilon$ by Claim 4. By the definition of μ and κ Receiver has correct beliefs. Finally, Let λ be sufficiently high such that $a \in BR(q)$ for every $q \in \Delta(S)$ such that $\rho(q, x_a) \leqslant B/\lambda$. Such λ exists by the continuity of ρ and the fact that a is the unique best response to belief x_a. Then by Claim 3 it follows that Receiver best responds to his belief for every belief.

This following example shows that Theorem 3.2 does not hold for nongeneric games.

Example 2. Let S be $\{0, 1\}$ with the prior $\{p, 1-p\}$. Let A be $\{a_0, a_1\}$. Sender's payoff is 1 if Receiver chooses a_1 and 0 if he chooses a_0. If Receiver chooses a_0, his payoff is 0. If Receiver chooses a_1, his payoff is 0 in state 0 and -1 in state 1. This game is nongeneric since action a_1 is weakly dominated but not strictly dominated. 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 ρ, in every equilibrium Receiver only chooses the safe action a_0.

6.6 Proof of Theorem 5.1

Lemma 4.1 implies that for $\lambda > \bar{\lambda}$ only calibrated Sender's strategies survive first round of elimination. Since all calibrated strategies are best responses to the Receiver's strategy that always chooses some fixed action $a \in A$ it follows that K_1 is the set of all calibrated strategies. Since any message m is an atom of some Sender's calibrated strategy, it follows that $\Sigma_2(m) = BR(m)$. Therefore, by Lemma 4.2, for every Receiver's strategy σ that survived, Sender believes that he can get at least $CM - \varepsilon$ against σ for every ϵ. Therefore a rationalizable Sender's strategy must be calibrated and give him CM against some Receiver's strategy σ such that $\text{support}(\sigma(\cdot|m)) \subseteq BR(m)$.

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.

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.

Bergemann, Dirk, and Stephen Morris. 2016*a*. "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games." *Theoretical Economics*, 11(2): 487–522.

Bergemann, Dirk, and Stephen Morris. 2016*b*. "Information Design, Bayesian Persuasion, and Bayes Correlated Equilibrium." *American Economic Review: Papers and Proceedings*, 106(5): 586–591.

Best, James, and Daniel Quigley. 2017. "Persuasion for the Long Run." *Working paper.*

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.

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. 2004. "Informal Communication." *Journal of Economic Theory*, 117(2): 180–200.

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.