Costly Miscalibration*

Yingni Guo[†] Eran Shmaya[‡]

February 28, 2019

Abstract

We consider a platform which provides probabilistic forecasts to a customer using some algorithm. We introduce a concept of miscalibration, which measures the discrepancy between the forecast and the truth. We characterize the platform's optimal equilibrium when it incurs some cost for miscalibration, and show how this equilibrium depends on the miscalibration cost: when the miscalibration cost is low, the platform uses more distant forecasts and the customer is less responsive to the platform's forecast; when the miscalibration cost is high, the platform can achieve its commitment payoff in an equilibrium, and the only extensive-form rationalizable strategy of the platform is its strategy in the commitment solution. Our results show that miscalibration cost is a proxy for the degree of the platform's commitment power.

JEL: D81, D82, D83

Keywords: Calibration, miscalibration, cheap talk, commitment, Bayesian persuasion

*Financial support from NSF Grant SES-1530608 is gratefully acknowledged.
[†]Department of Economics, Northwestern University, yingni.guo@northwestern.edu.
[‡]MEDS, Kellogg School of Management, Northwestern University, e-shmaya@kellogg.northwestern.edu.

1 Introduction

E-commerce platforms often provide information to customers about their products. For example, the fare aggregator Kayak.com provides forecasts of future prices. The real estate aggregator Redfin.com identifies "hot homes" that are likely to sell quickly. The platform generates information using some algorithm, which it applies to many products. This algorithm is usually a trade secret and is therefore not observed by outsiders. For example, Kayak states only that "our scientists develop these flight price trend forecasts using algorithms and mathematical models." Redfin states that "the hot homes algorithm automatically calculates the likelihood by analyzing more than 500 attributes of each home." In this paper, we develop a general model to analyze a platform's communication with customers, in which customers trust the platform's forecasts in an equilibrium even though they don't observe the algorithm.

Environment. 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 percent 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 percent of the homes that 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 the measure of miscalibration, which we define in the paper. The platform's cost also depends on a parameter, which we call the *cost intensity*. The cost intensity indicates how severely the platform is punished for miscalibration or how easy it is to collect data. This cost is a proxy for the reputation damage to the platform from making incorrect forecasts.

We model the interaction between the platform and a customer as a sender-receiver game: the forecast is the sender's message and the algorithm induces a strategy for the sender. As in the cheap-talk literature, the receiver observes only the message but not the sender's strategy. When the cost intensity is zero, our game is a cheap-talk game (e.g., Crawford and Sobel (1982), Green and Stokey (2007)). When the cost intensity is positive, the sender's talk is not cheap, because of the miscalibration cost. As the cost intensity increases, the sender becomes more concerned about the validity of his assertions. We examine how this miscalibration cost affects the sender's optimal equilibrium and the receiver's response to the messages.

Main results. In the sender-receiver game between the platform and the customer, the sender promotes a product and the receiver decides whether or not to buy it. We call such a game *a promotion game*. In addition to our platform-customer motivation, promotion games include many applications studied in recent papers, such as the interaction between a prosecutor and a judge as in Kamenica and Gentzkow (2011), the certification game in Henry and Ottaviani (2018) and Perez-Richet and Skreta (2018), the informational lobbying game in Minaudier (2018) and Bardhi and Guo (2018), and the media censorship game in Gehlbach and Sonin (2014) and Kolotilin et al. (2017).

Our first result is that there is a sender's optimal equilibrium in which his strategy is calibrated and has the support of two messages, one of which induces the receiver to purchase with positive probability.

Using this characterization, we first show that the sender's optimal equilibrium payoff is monotone-increasing in the cost intensity. Recall that the cost intensity models the degree of the reputation damage to the sender from making incorrect forecasts. The monotonicity result shows that the cost intensity can also be interpreted as the proxy for the degree of the sender's commitment power. The more the sender is likely to suffer from making incorrect assertions, the more he becomes credible in his communication with the receiver. A higher sender's payoff thus results. Naturally, different platforms have different cost intensities. For instance, Amazon.com has more to lose from incorrect assertions than some small stores' websites, and thus is viewed by customers as more credible in its communication. Similarly, Mayzlin, Dover and Chevalier (2014) show that small or independent hotels are more likely to engage in review manipulation than multi-unit or branded chain hotels, and Luca and Zervas (2016) show that chain restaurants are less likely to commit review fraud.

We then study how cost intensities affect the sender's optimal equilibrium. For high cost intensities, the commitment solution is an equilibrium, since even a small deviation incurs a big miscalibration cost. When the cost intensity is low, the equilibrium exhibits two features that are distinct from the commitment solution. First, the receiver purchases with only some probability after the purchase message. Second, the sender uses messages that are more distant from each other than the messages that are used in the commitment solution. Both these features reduce the sender's incentive to deviate. They show how the sender leverages his low commitment power to gain credibility.

Our second result is that our model bridges the cheap-talk and the persuasion models. For any sender-receiver game, when the cost intensity is zero, our game is (by definition) a cheap-talk game. We show that when the cost intensity is high, the sender's optimal equilibrium payoff is the payoff he could get if he could commit to a strategy. This result asserts some lower hemicontinuity of the equilibrium correspondence. As usual, lower hemicontinuity is not straightforward. In our setup, it requires either some assumption about the distance function that we use to measure miscalibration, or a generic assumption about the game. We also show that the only extensive-form rationalizable strategy of the sender is his strategy in the commitment solution.

To summarize, our contribution is threefold. We develop a general framework to analyze a platform's communication with its customer. We characterize how miscalibration cost affects the platform's optimal equilibrium and the customer's response to the messages. We show that our model bridges the cheap-talk and the persuasion models, and can be used to analyze the middle ground where the talk is not completely cheap nor the commitment absolute.

Related literature. This paper is related to the empirical literature on firms' online communication with customers. Chevalier and Mayzlin (2006) show that information on platforms has a significant impact on sales. Mayzlin, Dover and Chevalier (2014) explore the difference between a website on which faking is difficult and a website on which faking is relatively easy. They show that the cost of review manipulation determines the amount of manipulation in equilibrium and that different firms have different incentives to manipulate. We contribute to this literature by developing a general framework to analyze platforms' communication with customers. The results shed light on how firms' reputation concern affects the effectiveness of their communication.

The concept of calibration is central in the forecasting literature (e.g., Dawid (1982), Murphy and Winkler (1987), Foster and Vohra (1998), Ranjan and Gneiting (2010)). It is used for two closely related ideas. The first is the purely probabilistic sense, namely, that a forecast (a message, in our setup) matches the conditional distribution over states given the forecast. The second is the idea of calibration with the data, namely, that a forecast matches the realized distribution of states at those times in which the forecast was given. Building on these ideas, we introduce the concept of calibration to sender-receiver games, and develop a measure of miscalibration to study these games.[1]

Our paper is related to Perez-Richet and Skreta (2018) which characterizes the receiver's optimal test when the sender can falsify the state. In their baseline model, the sender's falsification is observable. Falsification improves the test results but devalues their meanings. They show that there always exists a receiver's optimal test in which the sender doesn't falsify. In our paper, the sender can deviate to any communication strategy. His strategy is not observable, but the potential miscalibration cost allows for meaningful communication. We characterize an environment in which the sender doesn't miscalibrate in his optimal

[1]There is also literature about the strategic manipulation of calibration tests. See, for instance, Foster and Vohra (1998), Lehrer (2001), and Olszewski (2015).

equilibrium, and also show that this is not always the case. Our paper is also related to Nguyen and Tan (2018). Their sender first publicly announces a test. Then, he privately observes the message generated, and can manipulate the message at a cost. The receiver's action depends on both the announced test and the final message. They characterize how the chance of manipulation affects the test design. Our model differs in that our sender chooses a test which is not observable, so the receiver observes only the message.

Our results show that our model bridges the cheap-talk and persuasion models (e.g., Rayo and Segal (2010), Kamenica and Gentzkow (2011)). From this aspect, our paper is related to Fréchette, Lizzeri and Perego (2018), Lipnowski, Ravid and Shishkin (2018) and Min (2018). In these papers, the sender first publicly announces a test. With an exogenously given probability, the message is given by this test. With the complementary probability, the sender can secretly choose a different message. The receiver in their models observes both the announced test and the final message. Our model and results propose a different measure of commitment, and generate qualitatively different predictions than the probabilistic commitment models. Lipnowski and Ravid (2018) is another recent paper that relates to the relationship between the commitment world and the cheap-talk world, and models cheap talk from a belief-based perspective.

Our paper is also related to the literature on strategic communication with a lying cost (e.g., Kartik, Ottaviani and Squintani (2007), Kartik (2009)). In the classic costly-lying model, the sender decides which message to send contingent on the state, and the lying cost is a function of the sender's state and the message. In our setup, the platform designs its algorithm before learning the product's state, and the miscalibration cost is a function of the joint distribution over states and messages. While an outsider does not observe the algorithm, he can estimate the joint distribution of states and messages when the algorithm is applied to many products, as demonstrated in the case of Redfin (see example 1). Hence, our model addresses different applications than the costly-lying model. The two models also deliver different asymptotic results, which capture different intuitions about what happens when lying becomes very costly. In the costly-lying model, as the lying cost increases, the equilibrium becomes arbitrarily close to fully revealing. In our model, as the miscalibration cost increases, the sender gains commitment power not to say something false, but he is not forced to reveal all the information.

Structure of the paper. In section 2 we use an example to illustrate our main concepts and results. Section 3 presents the model. We characterize the sender's optimal equilibrium for all promotion games and all possible cost intensities in section 4. Section 5 presents our results when the cost intensity is high. Section 6 extends our analysis to broader setups. Our results continue to hold when the sender has some information on the state but does not necessarily know it or when cheap talk is included as an option. Section 7 contains the proofs.

2 An example

Consider a platform (Sender) promoting a product to a customer (Receiver). The product's value net of its price is $s \in S = \{-2, -1, 1\}$ with the prior $p = \left(\frac{1}{8}, \frac{1}{2}, \frac{3}{8}\right)$. Receiver decides whether or not to buy the product. Receiver's payoff is s if he buys and zero otherwise. Sender gets a commission payoff of 1 whenever Receiver buys.

Sender's strategy is a mapping from the state space S to the set of distributions over the message space. The message space is $\Delta(S)$, so each message is a distribution over states. For any message m that Sender announces with positive probability, let $\tau(m)$ denote the conditional distribution over states by Bayes' rule. We call $\tau(m)$ 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 1. He sends either message m_0 or m_1. Message m_0

says that the state is -1. Message m_1 says that the distribution over states is $\left(\frac{1}{5}, \frac{1}{5}, \frac{3}{5}\right)$. Each column shows the probabilities with which Sender sends m_0 or m_1 in each state.

message \ state	-2	-1	1
$m_0 = (0, 1, 0)$	0	$\frac{3}{4}$	0
$m_1 = \left(\frac{1}{5}, \frac{1}{5}, \frac{3}{5}\right)$	1	$\frac{1}{4}$	1

Table 1: A calibrated strategy

message \ state	-2	-1	1
$m_0 = (0, 1, 0)$	0	$\frac{1}{2}$	0
$m_1 = \left(\frac{1}{5}, \frac{1}{5}, \frac{3}{5}\right)$	1	$\frac{1}{2}$	1

Table 2: A miscalibrated strategy

Under this strategy, the true conditional distribution given m_0 is $\tau(m_0) = (0, 1, 0)$. The true conditional distribution given m_1 is $\tau(m_1) = \left(\frac{1}{5}, \frac{1}{5}, \frac{3}{5}\right)$. Each message coincides with the true conditional distribution given that message. In this case, we say that the strategy is *calibrated.* 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 $\tau(m_0), \tau(m_1)$ are the true conditional distributions given the messages. For a calibrated strategy, the black circles always coincide with the gray dots.

There are many other calibrated strategies. For example, Sender can choose to reveal no information, by announcing the prior at every state. Or he can reveal all information by announcing the messages $(1, 0, 0), (0, 1, 0), (0, 0, 1)$ in states $-2, -1, 1$, respectively.

Consider another strategy in table 2. Sender uses the same messages m_0, m_1 as before. When the state is -2 or 1, he still announces m_1 for sure. However, when the state is -1, he announces m_0 and m_1 with equal probability. The true conditional distribution given m_1 is $\tau(m_1) = \left(\frac{1}{6}, \frac{1}{3}, \frac{1}{2}\right)$, which differs from m_1. Thus, this strategy is not calibrated, as illustrated in figure 2.

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 $\left(\frac{1}{5}, \frac{1}{5}, \frac{3}{5}\right)$ – is not true. If an outside group collects data and performs some statistical test, Sender might be caught.

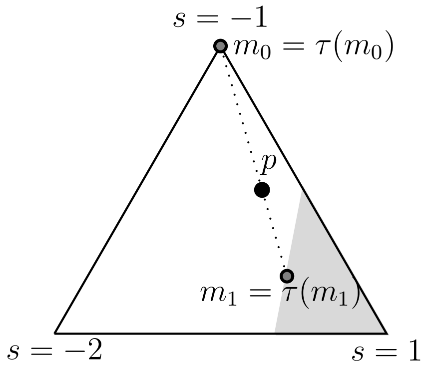

Figure 1: Strategy in table 1

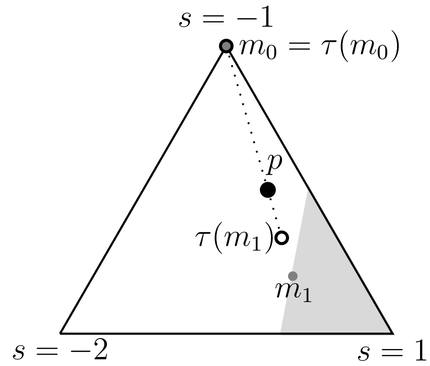

Figure 2: Strategy in table 2

Whenever Sender sends a message whose true conditional distribution differs from its asserted distribution, Sender pays some cost, which we call the *miscalibration cost*. The bigger the difference is, the higher the cost will be. We now 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 m. For the strategy in figure 2, the total-variation distance between m_1 and $\tau(m_1)$ is $\frac{2}{15}$. Since Sender announces message m_1 with probability $\frac{3}{4}$, his miscalibration cost is:

$$\lambda \cdot \frac{3}{4} \cdot \frac{1}{2}|m_1 - \tau(m_1)|_1 = \lambda \cdot \frac{3}{4} \cdot \frac{2}{15}.$$

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

Receiver sees only the message (like a customer seeing only the platform's message). He chooses the probability of buying after each message. He is willing to buy if his belief after a message is in the light-gray area in the figures. If $\lambda = 0$, Sender's only equilibrium payoff is zero. We next describe Sender's optimal equilibrium for any $\lambda > 0$, showing that (i) even for a low λ, Sender gets a positive payoff; and (ii) for a high λ, Sender gets his commitment payoff. This characterization relies on Theorem 4.1 in section 4.

For $\lambda \in \left(0, \frac{5}{4}\right]$, Sender plays the calibrated strategy in figure 1. Receiver buys with probability $\frac{4}{5}\lambda$ after m_1, and does not buy after any other message. Receiver has no incentive to deviate: given $\tau(m_1)$, both buying and not buying are his best responses; given $\tau(m_0)$, not buying is his unique best response.

Sender's equilibrium payoff is $\frac{1}{2}\lambda$. Sender could deviate by announcing m_1 more often in order to induce Receiver to buy more frequently. However, the potential miscalibration cost deters him from doing so. To illustrate, consider Sender's strategy in figure 2. His payoff from this deviation strategy is the probability that Receiver buys minus the miscalibration cost:

$$\frac{3}{4} \cdot \frac{4}{5}\lambda - \lambda \cdot \frac{3}{4} \cdot \frac{2}{15} = \frac{1}{2}\lambda,$$

which is the same as Sender's equilibrium payoff. Other deviations are not profitable either.

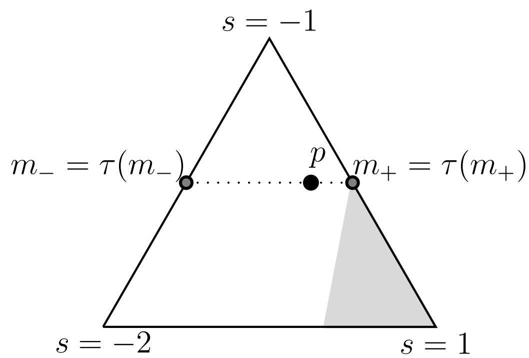

Figure 3: Sender's optimal equilibrium when $\lambda \geqslant 2$

For $\lambda \geqslant 2$, Sender's optimal equilibrium is illustrated by figure 3. He uses a calibrated strategy, splitting the prior into $m_- = \left(\frac{1}{2}, \frac{1}{2}, 0\right)$ and $m_+ = \left(0, \frac{1}{2}, \frac{1}{2}\right)$. Receiver buys for sure after m_+, and does not buy after any other message. Again, the potential miscalibration cost deters Sender from deviating and announcing m_+ more often. Sender's equilibrium payoff is $\frac{3}{4}$, which is the same as if he could commit to a strategy.

For $\lambda \in \left(\frac{5}{4}, 2\right)$, Sender's optimal equilibrium strategy shifts gradually from the strategy in figure 1 to that in figure 3. His optimal equilibrium payoff increases in the cost intensity λ, which can be interpreted as the proxy for the degree of Sender's commitment power.

Interestingly, for a low cost intensity, Sender uses messages m_0 and m_1 instead of m_-

and m_+; the latter pair is used in the commitment solution. With two distant messages like m_0 and m_1, if Sender deviates by announcing m_1 more often, he has to incur a larger amount of miscalibration, compared with the case with two nearby messages like m_- and m_+. Thus, when the cost intensity is low, Sender uses distant messages to compensate his low commitment power, in order to deter himself from deviating.

A miscalibrated optimal equilibrium: In the previous example, Sender gets his optimal equilibrium payoff in a calibrated equilibrium. This is not always true. Consider the same example but with the prior $p = \left(\frac{1}{4}, \frac{1}{4}, \frac{1}{2}\right)$ and with the Kullback-Leibler distance between $\tau(m)$ and m given by:

$$\sum_{s \in S} \tau(m)(s) \log \frac{\tau(m)(s)}{m(s)}.$$

Given this smooth distance function, for any message with full support a small amount of miscalibration has no first-order impact on the miscalibration cost.

If Sender could commit to a strategy, he would use the calibrated strategy with support $\underline{m} = (1, 0, 0)$ and $\overline{m} = \left(\frac{1}{7}, \frac{2}{7}, \frac{4}{7}\right)$. He would announce $\overline{m}$ with probability $\frac{1}{2}$ in state -2, and with probability 1 in states -1 and 1. Receiver would buy after $\overline{m}$. However, for any intensity λ, this strategy profile is not an equilibrium: Sender will deviate by announcing $\overline{m}$ more often in state -2. In particular, if Sender increases the probability of $\overline{m}$ in state -2 to $\frac{1}{2} + \varepsilon$, then the derivative of Sender's payoff w.r.t. ε at $\varepsilon = 0$ is $\frac{1}{4}$. More generally, if Receiver ever buys after a message with full support, Sender will miscalibrate to induce Receiver to buy more often, since a small amount of miscalibration has no first-order miscalibration cost.

We next construct an equilibrium in which Receiver buys the product only after the following message $\overline{m}'$:

$$\overline{m}' = \left(\frac{e^{-1/\lambda}}{7}, \frac{1}{3} - \frac{e^{-1/\lambda}}{21}, \frac{2}{3} - \frac{2e^{-1/\lambda}}{21}\right).$$

On the equilibrium path, Sender sends $\overline{m}'$ with probability one when the state is -1 or 1. When the state is -2, he sends $\underline{m}$ and $\overline{m}'$ with equal probability. The true conditional distribution given message $\overline{m}'$ is:

$$\tau(\overline{m}') = \left(\frac{1}{7}, \frac{2}{7}, \frac{4}{7}\right).$$

Figure 4 illustrates this equilibrium. Sender's strategy is not calibrated; he announces $\overline{m}'$, which is in the interior of the light-gray area, yet the true conditional distribution $\tau(\overline{m}')$ is just enough for Receiver to buy.

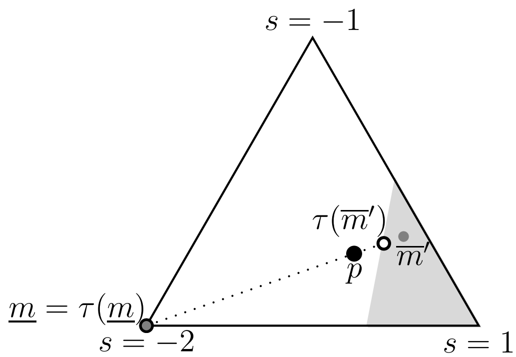

Figure 4: Sender overshoots in equilibrium to gain credibility

When a small amount of miscalibration involves minimal miscalibration cost, Sender "overshoots" in his announcement in order to gain credibility. His equilibrium payoff is $\frac{3}{4}\lambda \log\left(\frac{7e^{1/\lambda}}{6} - \frac{1}{6}\right)$, which converges to his commitment payoff as λ goes to infinity.

3 Environment

We consider a game with incomplete information between two players, Sender and Receiver. Sender sends Receiver a message that depends on a state of nature that is unobserved by Receiver. Receiver then chooses an action.

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 $\tau_\sigma : M \to \Delta(S)$ be such that $\tau_\sigma(m)$ is the conditional distribution over states given m. We sometimes use $\tau_\sigma(s|m)$ 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, and the definition of a cheap-talk equilibrium. Section 3.3 presents our definition of an equilibrium with costly miscalibration.

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..[2] Under a calibrated strategy, messages mean what they say, i.e., they can reliably be taken at face value.

We now introduce a key component of our model: a measure of miscalibration $d(\sigma)$ for a Sender's strategy σ. To define this measure, let $\kappa : \Delta(S) \times \Delta(S) \to \mathbf{R}_+$ be a continuous function, where $\kappa(q, m)$ measures the distance between a message $m \in \Delta(S)$ and a truth

[2] The term *a.s.* in our paper means "almost surely w.r.t. the probability distribution over messages induced by σ." Recall that τ_σ is defined up to a set of messages with probability zero.

$q \in \Delta(S)$. We assume that $\kappa(q, m) = 0$ if and only if $m = q$. Let

$$d(\sigma) = \sum_s p(s) \int \kappa(\tau_\sigma(m), m)\ \sigma(\mathrm{d}m|s)$$

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. A strategy σ is calibrated if and only if $d(\sigma) = 0$.

The following example illustrates the concept of miscalibration in the case of Redfin. Redfin's algorithm is not observed by Receiver or any third party. When the algorithm is applied to many products, however, the true conditional distribution given a message can be estimated.

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 collected $2,150$ hot homes and tracked their status.[3] When a home's status becomes contingent, pending, or sold, it is considered to have an accepted offer. The percentage of hot homes having an accepted offer within two weeks of its debut was 53.1. This is different from 70 percent, so Redfin's forecast was not calibrated.[4] □

3.2 Sender-receiver games

Let A be a finite set of *actions* by Receiver. 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 ρ from M to A, with the interpretation that Receiver randomizes an action from $\rho(\cdot|m)$ after message m.

[3] We examined Cook County in Illinois, which includes 165 zip codes, and collected homes that Redfin had identified as hot homes over two weeks (03/14/2018 – 03/27/2018).

[4] We thank Cassiano Alves and Samuel Goldberg for excellent research assistance.

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 (σ, ρ) are given by:

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

respectively. 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 $\lambda \geqslant 0$.

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 Equilibrium with 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:

$$V_\lambda(\sigma, \rho) = V_0(\sigma, \rho) - \lambda d(\sigma), \tag{1}$$

where $d(\sigma)$ is the measure of miscalibration, and $\lambda \geqslant 0$ is a parameter which indicates 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. We say that an equilibrium is calibrated (or miscalibrated) if Sender's strategy is calibrated (or miscalibrated).

Before proceeding, we show that any calibrated equilibrium for λ is also an equilibrium for a higher $\lambda' > \lambda$. Intuitively, if Sender has no incentive to miscalibrate for a low cost intensity, he surely has no incentive to do so when the cost is higher. However, a miscalibrated equilibrium for λ is not necessarily an equilibrium for a higher $\lambda' > \lambda$. This is because moving to a higher intensity means a higher miscalibration cost which might prompt Sender to deviate.

Claim 1. *For any sender-receiver game and any distance function, if (σ, ρ) is an equilibrium for $\lambda \geqslant 0$ and σ is calibrated, then it is an equilibrium for a higher $\lambda' > \lambda$.*

Proof. To show that (σ, ρ) is an equilibrium for λ', we first observe that Receiver has no incentive to deviate from ρ given σ. Second, Sender's payoff from deviating to any σ' is smaller than his equilibrium payoff since:

$$V_0(\sigma', \rho) - \lambda' d(\sigma') \leqslant V_0(\sigma', \rho) - \lambda d(\sigma') \leqslant V_0(\sigma, \rho).$$

The first inequality follows from the fact that $\lambda' > \lambda$. The second follows from the fact that (σ, ρ) is an equilibrium for λ and the fact that σ is calibrated. ■

4 Promotion games

Given our interest in platforms, we study a class of sender-receiver games which we call *promotion games*: Receiver decides whether or not to buy, and Sender's payoff is one if Receiver buys and zero otherwise. Section 2 provided an example of a promotion game. To avoid triviality, we assume that Receiver does not buy given the prior belief, and that he prefers to buy for some state.

Theorem 4.1 below enables us to fully characterize Sender's optimal equilibrium for any promotion game and any $\lambda > 0$. We assume that κ is a Wasserstein distance. Wasserstein distances are the natural distances over $\Delta(S)$ for a metric space S: Let $\overline{\kappa}$ be a metric on S; then the *Wasserstein distance κ over $\Delta(S)$ induced by $\overline{\kappa}$* is given by:

$$\kappa(q, q') = \inf \mathrm{E}\, \overline{\kappa}(Q, Q'),$$

where the infimum ranges over all pairs Q, Q' of S-valued random variables with marginal distributions q, q', respectively. (Such a pair Q, Q' is called a coupling of q, q'.) The total-variation distance is a Wasserstein distance when $\overline{\kappa}$ is the discrete metric so that $\overline{\kappa}(s, s') = 1$ for any $s \neq s'$. In a promotion game, it is also natural to choose $\overline{\kappa}(s, s') = |u(s, \mathrm{B}) - u(s', \mathrm{B})|$, where $u(s, \mathrm{B})$ is Receiver's payoff from buying in state s. In this case, when the state is 2, saying that the state is 1 is more costly for Sender than saying that the state is 1.[5] A Wasserstein distance κ is a convex function of q, q'. It is itself a metric, so it satisfies the triangular inequality.

Theorem 4.1. *Let $\lambda > 0$ and κ be a Wasserstein distance. In a promotion game, there is a Sender's optimal equilibrium with the following properties:*

(i) Sender's strategy is calibrated and uses two messages m_0 and m_1.

(ii) Receiver buys with a positive probability when the message is m_1 and buys with probability 0 for any other message.

(iii) Receiver is indifferent when the distribution over states is m_1, and m_0 is on the boundary of the simplex $\Delta(S)$.

Theorem 4.1 shows that Sender's optimal equilibrium payoff can be generated by a calibrated equilibrium. Therefore, according to Claim 1, Sender's optimal, calibrated equilibrium for λ is also an equilibrium for $\lambda' > \lambda$. This implies that Sender's optimal equilibrium payoff is monotone-increasing in λ. We summarize this observation in the follow corollary.

Corollary 4.2. *Let $\lambda > 0$ and κ be a Wasserstein distance. In a promotion game, Sender's optimal equilibrium payoff is monotone-increasing in λ.*

[5] We thank Bruno Strulovici for suggesting that the cost of saying s when the state is s' might depend on Receiver's payoff difference $|u(s, B) - u(s', B)|$. Wasserstein distances are rich enough to model this possibility.

Corollary 4.2 shows that the cost intensity λ can be interpreted as the proxy for the degree of Sender's commitment power. As Sender becomes more concerned about the validity of his assertions, he becomes more credible in his communication with Receiver.

Using Theorem 4.1, we can now explore how Sender's optimal equilibrium differs as the cost intensity varies. We start by describing Sender's incentive to deviate. If there is a message after which Receiver buys, Sender is tempted to announce this purchase message more frequently in order to induce Receiver to buy more often. Doing so leads to a discrepancy between the asserted and the true meanings of this purchase message. In a calibrated equilibrium, Sender must be deterred from doing so.

When λ is sufficiently high, the commitment solution is an equilibrium, since even a small deviation incurs a big miscalibration cost. When λ is low, the equilibrium exhibits two features that are distinct from the commitment solution. First, Receiver purchases with only some probability after the purchase message. This reduces Sender's gain from announcing this message more often. Second, Sender does not necessarily use the messages that are used in the commitment solution. Instead, he uses messages that are more distant from each other (like m_0, m_1 in figure 1 instead of m_-, m_+ in figure 3). The rationale is that using distant messages increases Sender's cost of miscalibration, which deters Sender from deviating. Suppose that Sender splits the prior p into m_0, m_1 such that Receiver buys after m_1. Announcing m_1 more frequently then would lead to a positive measure of miscalibration. For any fixed increase of the probability of announcing m_1, the miscalibration measure is larger when the distance between m_0 and m_1 is larger.[6]

To summarize, Sender gets a positive payoff even with a low commitment power (i.e.,

[6]We illustrate this point using the total-variation distance. Suppose that Sender splits p into m_0, m_1 in a calibrated strategy. The probabilities of m_0 and m_1 are $\frac{\kappa(m_1,p)}{\kappa(m_0,m_1)}$ and $\frac{\kappa(m_0,p)}{\kappa(m_0,m_1)}$, respectively. If Sender increases the probability of m_1 by ε, the true conditional distribution given m_1, denoted by $\tau(m_1)$, will differ from m_1:

$$\kappa(\tau(m_1), m_1) = \frac{\varepsilon\kappa(m_0, m_1)^2}{\varepsilon\kappa(m_0, m_1) + \kappa(m_0, p)}.$$

a low λ), and the optimal way to leverage the low commitment power is when Sender uses distant messages and Receiver randomizes. Thus, a less credible platform will use forecasts that are more distinct from each other than a more credible platform does, and its customers are less responsive to the purchase message in the sense that they buy only occasionally after the purchase message.

The result in Theorem 4.1 that Sender gets his optimal payoff in a calibrated equilibrium may look familiar from the cheap-talk case. Indeed, in that case (that is, when $\lambda = 0$) a simple revelation-principle argument implies that for every equilibrium (which is not necessarily Sender's optimal equilibrium), there exists an equilibrium which induces the same distribution over states and actions, and in which Sender announces the conditional distribution over states given his message (i.e., Sender uses a calibrated strategy). However, the revelation-principle argument does not apply when $\lambda > 0$. In fact, example 2 below provides a miscalibrated equilibrium which induces a distribution over states and actions that is not induced by any calibrated equilibrium. Moreover, Theorem 4.1 does not hold for all distance functions. In section 2, we showed that when κ is the Kullback-Leibler distance, Sender's optimal equilibrium is not calibrated. Theorem 4.1 relies on the triangular inequality, and also on a property of Wasserstein distances (formalized in Lemma 7.2 in the appendix) which makes them well-behaved under the splitting of probability distributions.

Example 2. Consider a promotion game with $S = \{-1, 1\}$ and the prior $p = \left(\frac{3}{4}, \frac{1}{4}\right)$. Receiver's payoff is s if he buys and zero otherwise. Let κ be the total-variation distance and let the cost intensity λ be 1. Since $M = \Delta(S)$ is one-dimensional, we use the probability of state 1 to represent the state distribution. Consider Sender's optimal equilibrium in Theorem 4.1. He uses a calibrated strategy and splits the prior $\frac{1}{4}$ into 0 and $\frac{1}{2}$. Receiver buys after

The miscalibration measure is monotone-increasing in $\kappa(m_0, m_1)$, as shown here:

$$\left(\frac{\kappa(m_0, p)}{\kappa(m_0, m_1)} + \varepsilon\right) \kappa(\tau(m_1), m_1) = \varepsilon \kappa(m_0, m_1).$$

message $\frac{1}{2}$ with probability $\frac{\lambda}{2}$, which ensures that Sender doesn't deviate. Consider now a miscalibrated equilibrium. Receiver buys after message 1 with probability λ and doesn't buy otherwise. Sender still splits the prior into 0 and $\frac{1}{2}$. When 0 is realized, he announces message 0. However, when $\frac{1}{2}$ is realized, he announces message 1. The two equilibria generate the same payoff of $\frac{1}{4}$ to Sender, but different distributions over states and actions. Conditional on state 1, Receiver buys more often under the second than the first equilibrium. In fact, no calibrated equilibrium generates the same distribution over states and actions as the second equilibrium, since otherwise Sender's equilibrium payoff would be higher than his optimal equilibrium payoff of $\frac{1}{4}$. □

5 High cost intensity

We now turn to general sender-receiver games. The three propositions in this section formalize the intuition that when the cost intensity is high, Sender gains the commitment power not to make false assertions. Propositions 5.1 and 5.2 show that Sender can achieve his commitment payoff in an equilibrium. Proposition 5.3 shows that the only extensive-form rationalizable strategy of Sender is his strategy in the commitment solution.

Formally, the commitment payoff is given by $CP = \max V_0(\sigma, \rho)$, where the maximum ranges over all profiles (σ, ρ) such that ρ is a best response to σ. We call a profile (σ, ρ) that achieves the maximum a commitment solution.

5.1 Equilibrium results

The commitment solution dispenses with the equilibrium constraints on Sender that appear in a cheap-talk equilibrium. Thus, the commitment payoff is the highest payoff that Sender can achieve if Receiver observes Sender's strategy. It is also an upper bound of Sender's equilibrium payoff in our model.

Claim 2. *Sender's payoff in any equilibrium is at most CP.*

Proof. Let (σ^*, ρ^*) be an equilibrium. Then

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

where the first inequality follows from (1), and the second from the definition of CP and the fact that ρ^* is Receiver's best response to σ^*. ■

In this subsection, we prove that when the cost intensity is high, Sender can achieve the commitment payoff in an equilibrium. On the surface, this high-cost equilibrium result can be understood from the following two observations. First, if Sender were exogenously restricted to using only calibrated strategies, then Receiver could take his message at face value. Hence, Sender's optimal commitment strategy, along with Receiver's best response to the face value of each message, would be an equilibrium. Second, as the cost intensity increases, in any equilibrium Sender uses only strategies that are close to being calibrated, in the sense that each message is close to the true posterior distribution over states given that message. Thus, our result asserts some lower hemicontinuity of the equilibrium correspondence. As usual, lower hemicontinuity is not straightforward. In our setup, it requires either some assumption about the distance function that we use to measure miscalibration, or a generic assumption about the game. The difficulty that we need to overcome is that Sender's payoff is not a continuous function of his strategy, because Receiver's strategy is typically not a continuous function of the message. Therefore, a small deviation from the optimal commitment strategy might have a big impact on Sender's payoff.

We first make the additional assumption that the distance function $\kappa(q, m)$ has a kink at $q = m$. This assumption holds for Wasserstein distances, as discussed in section 4, and in particular the total-variation distance. It also holds, for example, for the Euclidean distance $\kappa(q, m) = |q - m|_2$ and any other distance that is derived from a norm.

Proposition 5.1. *Assume that there exists some $\gamma > 0$ such that $\kappa(q, m) \geqslant \gamma |q - m|_1$. For every game there exists $\bar{\lambda}$ such that for every $\lambda > \bar{\lambda}$, there exists an equilibrium in the game with intensity λ such that (i) Sender's strategy is calibrated, and (ii) his payoff is CP.*

To prove Proposition 5.1, our key observation is Lemma 7.4 in the appendix. It asserts that for any distance function with a kink, any miscalibrated strategy of Sender is strictly dominated when λ is sufficiently high.

To prove Lemma 7.4, for any miscalibrated strategy we construct a calibrated strategy which, with high probability, announces the same messages at every state as the original strategy does. Therefore, the new strategy yields a similar payoff to Sender from his interaction with Receiver as the original strategy does. Sender may very well lose some payoff from his interaction with Receiver. Nonetheless, when we measure miscalibration using a distance function with a kink, the payoff loss is at most some constant times the miscalibration measure. Meanwhile, this new, calibrated strategy saves Sender the miscalibration cost, so it yields a higher payoff for Sender when λ is high enough.

Hence, for a high cost intensity, Sender uses only calibrated strategies and there is an equilibrium in which he obtains exactly his commitment payoff. In section 2 we demonstrated this result using the total-variation distance. The commitment solution is given by Sender's strategy in figure 3; Receiver buys the product after m_+ and does not buy after any other message. This strategy profile is an equilibrium when $\lambda \geqslant 2$.

When we now consider arbitrary distance functions, the assertion of Proposition 5.1 no longer holds. This can be seen from example 3 at the end of this section: Sender's equilibrium payoff is constantly zero for any cost intensity. However, our next proposition shows that a weaker result with a similar message still holds: in a generic set of games, as the cost intensity increases, there exists an equilibrium in which Sender's payoff approaches his commitment payoff. Here and everywhere else, when we refer to "a generic set of games," we mean that the set of payoff functions $u : S \times A \to \mathbf{R}$ for which the assertion does not hold (viewed as a subset of $\mathbf{R}^{S\times A}$) is a closed set with an empty interior and a Lebesgue measure of zero. More explicitly, our generic assumption requires that every Receiver's action be a *unique* best response to some beliefs over the states.

Proposition 5.2. *Assume that κ is convex in q. In a generic set of games, for every $\varepsilon > 0$, there exists $\bar{\lambda}$ such that for every $\lambda > \bar{\lambda}$, there exists an equilibrium in the game with intensity λ such that Sender's payoff is at least $CP - \varepsilon$.*

To illustrate the proof idea, consider the simplifying assumption that Receiver has some "punishment" action that yields a bad payoff for Sender in every state. For a generic game, any approximation for the commitment payoff can be achieved by a calibrated strategy for which Receiver's best response is unique. Then, we construct an equilibrium in which Receiver responds to messages in the support of this strategy in the same way as before, and uses the punishment action for any other message. Sender's equilibrium strategy may well be miscalibrated, but for a high cost intensity the amount of miscalibration will be small enough so that Receiver's response is still uniquely optimal. Sender thus gets an equilibrium payoff close to the commitment payoff. The proof for the case in which the simplifying assumption does not hold is more involved, but relies on a similar idea. In section 2, we gave an example that demonstrated Proposition 5.2 for the case of the Kullback-Leibler distance.

We have shown that Sender can achieve his commitment payoff (possibly up to ε) in an equilibrium with some assumption about the cost function (Proposition 5.1) or some generic assumption about the game (Proposition 5.2). In contrast, if we make neither assumption, Sender's optimal equilibrium payoff, for all cost intensities, might be bounded away from his commitment payoff, as shown in the following nongeneric game.

Example 3. Let S be $\{0, 1\}$ with the prior $p = (p_0, 1 - p_0)$. 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 in both states. 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_0. However, for any smooth distance function κ, in every equilibrium Receiver chooses only action a_0, so Sender's payoff is 0. $\square$

5.2 A rationalizablility result

Up to now, following the cheap-talk literature, we used the equilibrium as our solution concept. We argued that Sender's optimal equilibrium achieves the commitment solution for a high λ. But the game still admits other equilibria. For example, even for a high λ, there exists an uninformative equilibrium. Under this equilibrium, Sender always declares the prior, and, 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 uninformative equilibrium in our context. Consider the example with a high λ from section 2. Assume that Receiver plays according to the uninformative equilibrium, and that the message $(0, 0, 1)$ – which says that "the state is 1 with probability one" – arrives. What should Receiver do? The message is a surprise, but given the high λ, 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, and also refinements such as the perfect Bayesian equilibrium, do not capture this intuition because they allow arbitrary behaviors or beliefs off path. These behaviors do not have to be rationalizable. In order to incorporate this earlier intuition, we turn to extensive-form rationalizability (hereafter, EFR; see Pearce (1984), Battigalli (1997), and Battigalli and Siniscalchi (2002)). EFR dispenses with the assumption that players' beliefs are correct, but requires that, at every point in the game, each player form 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 in the appendix and Friedenberg (2019), for example, where similar issues arise.) However, the issue is somewhat simpler in our setup, because each player takes action only once. The important assumption behind EFR is that Receiver *strongly believes* that Sender does not use strictly dominated strategies. This means that, after observing a message, Receiver has a belief about Sender's strategies which is concentrated on those Sender's strategies that are not strictly dominated.

Proposition 5.3 says that, under the EFR solution, Receiver takes Sender's message at face value and Sender plays his optimal commitment strategies. 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:

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

Proposition 5.3. *Assume that there exists some $\gamma > 0$ such that $\kappa(q,m) \geqslant \gamma|q-m|_1$. Then, in a generic set of games, there exists $\bar{\lambda}$ such that for every $\lambda > \bar{\lambda}$, the following holds:*

1. *Receiver's extensive-form rationalizable strategies are the ones that satisfy:*

$$\textit{support}(\rho(\cdot|m)) \subseteq BR(m).$$

2. *Sender's extensive-form rationalizable strategies are his calibrated strategies σ such that $V_\lambda(\sigma,\rho) = CP$ for some best response ρ of Receiver to σ.*

To prove this result, we use the result (Lemma 7.4 in the appendix) that, for sufficiently high λ, any miscalibrated strategy of Sender is strictly dominated. Therefore, only Sender's calibrated strategies survive the first round of elimination. Given that Sender uses only calibrated strategies, Receiver takes any message at face value, and best responds to the face value. In a generic game, any approximation for the commitment payoff can be achieved by a calibrated strategy for which Receiver's best response is unique. Therefore, the only Sender's strategy that will survive is one that gives him the commitment payoff against some Receiver's best response. Note that Proposition 5.3 does not hold in the nongeneric game in example 3 even when κ has a kink. In that game, Receiver's strategy to always take a_0 is rationalizable, and therefore every calibrated strategy is rationalizable for Sender.

When we apply Proposition 5.3 to the example from section 2 with the total-variation distance, we see that, for a high λ, the only rationalizable strategy for Sender is the calibrated strategy in figure 3. Receiver's rationalizable actions after m are actions in $BR(m)$.

6 Discussion

In this section, we discuss variants of our model, and show how our analysis and results extend to these variant settings as well.

6.1 Partial information on the state

By allowing Sender to choose any strategy $\sigma : S \to \Delta(S)$, we implicitly assume that Sender knows the state. This assumption is natural since our main focus is on Sender's incentives. However, because we use terminology from the forecasting literature and because of our interest in e-commerce platforms like Redfin, it would be more realistic to assume that Sender has some information on the state but does not necessarily know it.

To model such an environment, we can add an exogenous set T of *Sender's types* and an exogenous information structure $\sigma^E : S \to \Delta(T)$, such that $\sigma^E(s)$ is the distribution over Sender's types. When Sender's type is t, his belief about the state is $\tau_{\sigma^E}(t)$. Our definitions and results will carry through *mutatis mutandis* if we assume that the message space M is the convex hull of $\{\tau_{\sigma^E}(t) : t \in T\}$, and that Sender is restricted to strategies that are less informative than σ^E (in Blackwell's sense).

6.2 Cheap talk and costly miscalibration

We next show that our results continue to hold when we include cheap-talk communication as an option. We let the message space be $M_0 \times \Delta(S)$ for some large, finite M_0. The miscalibration cost is applied to only the second argument of the message.

We first argue that the characterization in Theorem 4.1 stays the same in this hybrid model. Receiver's strategy is now $\rho : M_0 \times \Delta(S) \to \Delta(A)$. If in an equilibrium Receiver's strategy is a function of the second argument alone, then this is effectively an equilibrium in our original model. If there exists $m_a, m_b \in M_0$ and $m \in \Delta(S)$ such that $\rho(m_a, m) > \rho(m_b, m)$, then Sender will never use (m_b, m) on path since he can profitably deviate. Whenever Sender is supposed to announce (m_b, m), he can announce (m_a, m) instead. The miscalibration cost is the same as before but Sender induces a higher action by Receiver. This shows that Receiver's strategy is again effectively a function of the second argument alone. Hence, the set of equilibrium outcomes in the hybrid model is the same as that in our original model.

Second, our results for high cost intensities (Propositions 5.1 to 5.3) continue to hold as well. Because any equilibrium in our model can be transformed into an equilibrium in this hybrid model by letting both players ignore the cheap-talk message, Sender achieves his commitment payoff in an equilibrium for high intensities. Similarly, Sender won't use miscalibrated strategies for high intensities in the hybrid model, so the EFR result continues to hold here as well.

7 Proofs

7.1 Preliminaries

We extend the distance function κ to a function $\kappa : \mathbf{R}_+^S \times \Delta(S) \to \mathbf{R}_+$ given by $\kappa(\gamma q, m) = \gamma\kappa(q, m)$ 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 σ we let

$$\chi_\sigma = \sum_s p(s)\sigma(\cdot|s) \tag{3}$$

be the distribution over messages induced by σ. Lemma 7.1 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 7.1. *(i) For every strategy σ, we have*

$$\int \tau_\sigma(m)\ \chi_\sigma(\mathrm{d}m) = p. \tag{4}$$

(ii) 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 σ such that $\chi_\sigma = \chi$.

We also say that a *finite splitting* of a probability measure $p \in \Delta(S)$ is a representation $p = \sum_i t_i q_i$ such that $q_i \in \Delta(S)$, $t_i \geqslant 0$ and $\sum_i t_i = 1$. We sometimes call t_i the *weights*. The splitting lemma implies that for every such finite splitting there exists a calibrated strategy that announces the message q_i with probability t_i.

7.2 Proof of Theorem 4.1

7.2.1 Preliminaries

For a promotion game, we assume that $A = \{B, NB\}$ and that $v(s, B) = 1$ and $v(s, NB) = 0$ for each s. Also $u(s, NB) = 0$ for each s. To avoid triviality we assume that $\sum_s p(s)u(s, B) < 0$, so that Receiver does not buy absent additional information, and that $u(s, B) > 0$ for some s.

We can assume without loss that $\lambda = 1$, since all properties of the distance function κ that will be used in the proof still hold if we replace κ by $\lambda\kappa$.

We denote a Receiver's strategy by a function $\rho : \Delta(S) \to [0, 1]$, so that $\rho(m)$ is the probability that Receiver buys after message m.

For a bounded function $f : \Delta(S) \to \mathbf{R}$, let $\operatorname{Lip} f : \Delta(S) \to \mathbf{R}$ be given by

$$\operatorname{Lip} f(q) = \sup_{m \in \Delta(S)} \{f(m) - \kappa(q, m)\}.$$

Note that if κ satisfies the triangular inequality, then $\operatorname{Lip} f$ is the least 1-Lipschitz (w.r.t. κ) majorant of f.

For a bounded function $f : \Delta(S) \to \mathbf{R}$, we denote by $\operatorname{Cav} f$ the least concave majorant of f, sometimes called the concave envelope of f.

The following proposition summarizes obvious properties of the operators Lip and Cav. For functions $f, g : \Delta(S) \to \mathbf{R}$, we denote $f \leqslant g$ when $f(q) \leqslant g(q)$ for every $q \in \Delta(S)$.

Proposition 7.1. *Let $f, g : \Delta(S) \to \mathbf{R}$ be bounded. Then,*

(i) if κ satisfies the triangular inequality, then $\operatorname{Lip} \operatorname{Lip} f = \operatorname{Lip} f$*;*

(ii) $\operatorname{Cav} \operatorname{Cav} f = \operatorname{Cav} f$*;*

(iii) if $f \leqslant g$, then $\operatorname{Lip} f \leqslant \operatorname{Lip} g$ *and* $\operatorname{Cav} f \leqslant \operatorname{Cav} g$.

Claim 3. *If $\kappa(q,m)$ is convex in q, then Sender's optimal payoff against a Receiver's strategy $\rho : \Delta(S) \to [0,1]$ when the prior distribution over states is p is* $\operatorname{Cav}\operatorname{Lip}\rho(p)$.

Proof. Consider some Sender's strategy σ. Then Sender's payoff under (σ, ρ) is

$$V_1(\sigma,\rho) = \int (\rho(m) - \kappa(\tau_\sigma(m), m)) \;\; \chi_\sigma(\,\mathrm{d}m) \leqslant \int \operatorname{Lip}\rho(\tau_\sigma(m)) \;\; \chi_\sigma(\,\mathrm{d}m) \leqslant \operatorname{Cav}\operatorname{Lip}\rho(p),$$

where the first inequality follows from the definition of $\operatorname{Lip}\rho$ and $\rho(m) - \kappa(\tau_\sigma(m), m) \leqslant \operatorname{Lip}\rho(\tau_\sigma(m))$, and the second inequality follows from (4) and the definition of Cav.

For the converse, fix $\varepsilon > 0$. From the definition of Cav and Lip, it follows that there exist elements $q_i \in \Delta(S)$, weights $t_i \geqslant 0$, and messages m_i, such that $p = \sum_i t_i q_i$, $\sum_i t_i = 1$, and

$$\sum_i t_i \left(\rho(m_i) - \kappa(q_i, m_i)\right) \geqslant \operatorname{Cav}\operatorname{Lip}\rho(p) - \varepsilon. \tag{5}$$

We can assume that $m_i \neq m_j$ if $i \neq j$. Otherwise, if $m_i = m_j = \overline{m}$, we let $\overline{t} = t_i + t_j$ and $\overline{q} = (t_i q_i + t_j q_j)/\overline{t}$. Since κ is convex in q, we can replace the elements q_i and q_j with a single element $\overline{q}$, whose weight and corresponding message are $\overline{t}$ and $\overline{m}$, without violating (5).

We now consider Sender's strategy σ such that $\chi_\sigma = \sum_i t_i \delta_{m_i}$ and $\tau_\sigma : \Delta(S) \to \Delta(S)$ is a function such that $\tau_\sigma(m_i) = q_i$. Such a strategy exists by Lemma 7.1. Then (5) implies that Sender's payoff when using σ is at least $\operatorname{Cav}\operatorname{Lip}\rho(p) - \varepsilon$. ■

The following lemma shows that Wasserstein distances get along well with splitting of probability distributions. See Laraki (2004) for related results on other metrics and generalizations to infinite-dimensional spaces, including implications for Lipschitz continuity as in our Corollary 7.3.

Lemma 7.2. *Let κ be a Wasserstein distance over $\Delta(S)$. Let $q, q' \in \Delta(S)$ and let $q = t_1 q_1 + \cdots + t_n q_n$ be a splitting of q. Then there exists a splitting $q' = t_1 q'_1 + \cdots + t_n q'_n$ of q'*

with the same weights such that:

$$\kappa(q, q') = t_1\kappa(q_1, q_1') + \cdots + t_n\kappa(q_n, q_n'). \tag{6}$$

Proof. The direction $\kappa(q, q') \leqslant t_1\kappa(q_1, q_1') + \cdots + t_n\kappa(q_n, q_n')$ follows from convexity of κ for every splitting $q' = t_1 q_1' + \cdots + t_n q_n'$ of q' with the weights $t_1, \ldots, t_n$.

For the other direction, let Q, Q' be S-valued random variables with marginal distributions q, q' respectively such that $\kappa(q, q') = \mathrm{E}\overline{\kappa}(Q, Q')$. Let X be a random variable that assumes values in $\{1, \ldots, n\}$ such that $t_i = \mathrm{P}(X = i)$ and such that $q_i(s) = \mathrm{P}(Q = s|X = i)$ for every $i \in \{1, \ldots, n\}$ and $s \in S$. The existence of such a variable follows from the splitting lemma. Finally, let $q_i'(s) = \mathrm{P}(Q' = s|X = i)$. Then

$$\sum_i t_i\kappa(q_i, q_i') \leqslant \sum_i \mathrm{P}(X = i)\mathrm{E}\overline{\kappa}(Q, Q'|X = i) = \mathrm{E}\overline{\kappa}(Q, Q') = \kappa(q, q'),$$

where the inequality follows from the definition of $\kappa(q_i, q_i')$, since the marginal distributions of Q, Q' conditioned on the event $X = i$ are q_i, q_i' respectively. ∎

Corollary 7.3. *Let κ be a Wasserstein distance. Then, for every function $f : \Delta(S) \to [0, 1]$ that is 1-Lipschitz w.r.t. κ, the concave envelope* $\mathrm{Cav}\, f$ *is also 1-Lipschitz w.r.t. κ.*

Proof. Let $q, q' \in \Delta(S)$ and let $q = t_1 q_1 + \cdots + t_n q_n$ be a splitting of q such that $\mathrm{Cav}\, f(q) = t_1 f(q_1) + \cdots + t_n f(q_n)$. By Lemma 7.2 there exists a splitting $q' = t_1 q_1' + \cdots + t_n q_n'$ of q' such that (6) holds. Therefore

$$\mathrm{Cav}\, f(q') \geqslant \sum_i t_i f(q_i') \geqslant \sum_i t_i \left(f(q_i) - \kappa(q_i, q_i')\right) = \mathrm{Cav}\, f(q) - \kappa(q, q'),$$

as desired. ∎

7.2.2 Proof of Theorem 4.1

Let (σ, ρ) be an equilibrium. From the equilibrium condition for Sender and Claim 3, it follows that

$$\operatorname{Cav}\operatorname{Lip}\rho(p) = \int (\rho(m) - \kappa(\tau_\sigma(m), m)) \;\; \chi_\sigma(\,\mathrm{d}m). \tag{7}$$

since the right-hand side is Sender's payoff under the profile (σ, ρ).

Let $\mathcal{A} = \{q \in \Delta(S) : \sum_s q(s)u(s, B) \geqslant 0\}$ be the set of beliefs over states under which Receiver is willing to buy.

Let $t = \chi_\sigma(\tau_\sigma^{-1}(\mathcal{A}))$ and let $\overline{\chi}_\sigma = \chi_\sigma(\cdot \mid \tau_\sigma^{-1}(\mathcal{A}))$ and $\underline{\chi}_\sigma = \chi_\sigma(\cdot \mid \tau_\sigma^{-1}(\mathcal{A}^c))$. Let $\overline{q} = \int \tau_\sigma \;\, \mathrm{d}\overline{\chi}_\sigma$ and $\underline{q} = \int \tau_\sigma \;\, \mathrm{d}\underline{\chi}_\sigma$. Then

$$\chi_\sigma = (1-t)\underline{\chi}_\sigma + t\overline{\chi}_\sigma, \text{ and} \tag{8}$$

$$p = (1-t)\underline{q} + t\overline{q}. \tag{9}$$

From the equilibrium condition for Receiver it follows that $\rho = 0 \quad \underline{\chi}_\sigma$-almost surely. Therefore, from (7) and (8) it follows that

$$\operatorname{Cav}\operatorname{Lip}\rho(p) \leqslant t \int (\rho(m) - \kappa(\tau_\sigma(m), m)) \;\; \overline{\chi}_\sigma(\,\mathrm{d}m) \tag{10}$$

It follows from the definition of $\overline{q}$ and the convexity of $\mathcal{A}$ that $\overline{q} \in \mathcal{A}$. Since $p \notin \mathcal{A}$ there exists a belief q^* on the interval $[p, \overline{q}]$ which is on the boundary of $\mathcal{A}$, i.e., $\sum_s q^*(s)u(s, B) = 0$. From (9)

$$q^* = (1-t^*)\underline{q} + t^*\overline{q} = (1-t^*)\underline{q} + t^* \int \tau_\sigma \;\, \mathrm{d}\overline{\chi}_\sigma \tag{11}$$

for some $t^* \geqslant t$. Also,

$$p = \left(1 - \frac{t}{t^*}\right)\underline{q} + \frac{t}{t^*}q^*.$$

We now define a strategy profile as follows:

(i) Receiver's strategy ρ^* is given by $\rho^*(q^*) = \operatorname{Cav}\operatorname{Lip}\rho(q^*)$ and $\rho^*(m) = 0$ for $m \neq q^*$.

(ii) Sender's strategy σ^* is the calibrated strategy induced by the distribution over messages given by $\chi^* = (1 - \frac{t}{t^*})\delta_{\underline{q}} + \frac{t}{t^*}\delta_{q^*}$.

We claim that this is an equilibrium that gives Sender the same payoff $\operatorname{Cav}\operatorname{Lip}\rho(p)$ as the original equilibrium.

First, note that the equilibrium condition on Receiver's side is satisfied since Sender's strategy is calibrated, $\underline{q} \notin \mathcal{A}$, and Receiver is indifferent under q^*.

Second, Sender's payoff under the strategy profile (σ^*, ρ^*) satisfies

$$
\begin{aligned}
V_1(\sigma^*, \rho^*) = \frac{t}{t^*}\rho^*(q^*) = \frac{t}{t^*}\operatorname{Cav}\operatorname{Lip}\rho(q^*) \geqslant \\
t\int \operatorname{Lip}\rho(\tau_\sigma(m))\ \overline{\chi}_\sigma(\mathrm{d}m) \geqslant t\int (\rho(m) - \kappa(\tau_\sigma(m), m))\ \overline{\chi}_\sigma(\mathrm{d}m) \geqslant \operatorname{Cav}\operatorname{Lip}\rho(p),
\end{aligned}
$$

where the first inequality follows from (11), the second inequality follows from the definition of Lip, and the third inequality follows from (10).

Lastly, since $\rho^*(q) \leqslant \operatorname{Cav}\operatorname{Lip}\rho(q)$ for every $q \in \Delta(S)$, then:

$$
\operatorname{Cav}\operatorname{Lip}\rho^*(p) \leqslant \operatorname{Cav}\operatorname{Lip}\operatorname{Cav}\operatorname{Lip}\rho(p) = \operatorname{Cav}\operatorname{Cav}\operatorname{Lip}\rho(p) = \operatorname{Cav}\operatorname{Lip}\rho(p).
$$

The first inequality and the last equality follow from Proposition 7.1 (iii) and (ii), respectively. For the second equality, according to Corollary 7.3, $\operatorname{Cav}\operatorname{Lip}\rho(p)$ is 1-Lipschitz w.r.t. κ. Therefore, $\operatorname{Lip}\operatorname{Cav}\operatorname{Lip}\rho(p) = \operatorname{Cav}\operatorname{Lip}\rho(p)$.

By Claim 3 this implies that Sender's optimal payoff against ρ^* is at most $\operatorname{Cav}\operatorname{Lip}\rho(p)$, as desired.

7.3 Proof of Proposition 5.1

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 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 7.4 below, which asserts that for a sufficiently high intensity, Sender is always better off using a calibrated strategy, regardless of Receiver's strategy.

Lemma 7.4. *Assume that there exists some $\gamma > 0$ such that $\kappa(q, m) \geqslant \gamma|q - m|_1$. There exists $\bar{\lambda}$ such that for every $\lambda > \bar{\lambda}$, a Sender's strategy σ is not strictly dominated if and only if it is calibrated.*

Figure 5 illustrates the proof idea. Consider a miscalibrated strategy in figure 5: Sender splits prior p into the three true conditional distributions given by the black circles. For each true conditional distribution, he announces the message nearby given by a gray dot.

To construct a calibrated strategy, we let the messages (i.e., the gray dots) be the true conditional distributions, and let Sender announce these messages using the same probabilities as before. The barycenter of the messages is denoted by q, which may not be the same as p. We now construct a calibrated strategy. Sender first splits p into q and some q'. If q' occurs, Sender announces q'. If q occurs, Sender splits q into the gray dots and announces the true distributions. Sender doesn't pay any miscalibration cost.

Moving to this calibrated strategy, Sender may lose payoffs for two reasons. First, Receiver's response to q', $\rho(q')$, might not be favorable. But the probability that q' occurs is smaller than some constant times the measure of miscalibration of the previous strategy. Second, the true distribution given m used to be $\tau(m)$, but now it is m. We argue that the payoff loss is also smaller than some constant times the measure of miscalibration. Here, we use the property that, for any mixed action of Receiver, the difference between the expected payoffs for Sender under two different distributions (i.e., $\tau(m)$ and m) is at most some constant times the total-variation distance between these distributions. To summarize, the total payoff loss is at most some constant times the measure of miscalibration. Hence, if λ is high enough, the total payoff loss is less than the miscalibration cost that is saved.

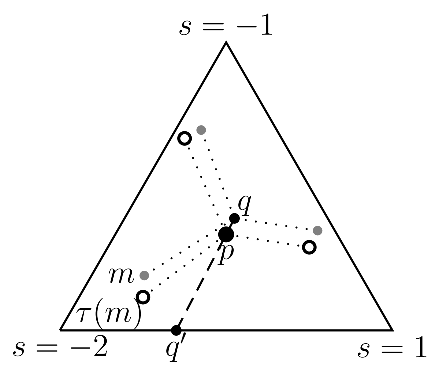

Figure 5: Proof idea for Lemma 7.4

Proof of Lemma 7.4. We first prove that any calibrated strategy is not strictly dominated. This follows from the fact that if Receiver ignores the message and plays some constant action a then all calibrated strategies give the same payoff and all miscalibrated strategies give lower payoffs.

Let $L = 1/\min\{p(s) : s \in S\}$. (Here we use the full-support assumption about 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

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

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

$$p = (1 - L|p - q|_1) \cdot q + L|p - q|_1 \cdot q'. \tag{12}$$

We now fix a Sender's strategy σ which is not strictly dominated. We need to prove that σ is calibrated. First, let σ' be any calibrated strategy of Sender and ρ be a strategy of Receiver. Then $V_\lambda(\sigma', \rho) = V_0(\sigma', \rho) \geqslant -B$ and $V_\lambda(\sigma, \rho) \leqslant B - \lambda d(\sigma)$. Since σ is not strictly dominated by σ', we can choose ρ such that $V_\lambda(\sigma', \rho) \leqslant V_\lambda(\sigma, \rho)$, which implies that:

$$d(\sigma) \leqslant 2B/\lambda. \tag{13}$$

Let χ_σ given by (3) be the distribution over messages induced by σ and let

$$q = \int m \; \chi_\sigma(\mathrm{d}m) \tag{14}$$

be the barycenter of χ_σ. It then follows from (4), (14), the convexity of the norm $|\cdot|_1$, and (16) that

$$|p - q|_1 = \left| \int (\tau_\sigma(m) - m) \;\; \chi_\sigma(\mathrm{d}m) \right|_1 \leqslant \int |\tau_\sigma(m) - m|_1 \; \chi_\sigma(\mathrm{d}m) \leqslant \frac{1}{\gamma} d(\sigma). \tag{15}$$

Let $\lambda \geqslant \frac{2BL}{\gamma}$. It then follows from (15) and (13) that $|p - q|_1 \leqslant 1/L$. Consider now the calibrated strategy σ' that is induced (via part 2 of Lemma 7.1) 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 (12) and $\delta_{q'}$ is the Dirac measure over q'. Thus, according to σ', 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 (12) and (14), it follows that the barycenter of this distribution is indeed p.)

Fix a Receiver's strategy ρ and let $\eta : \Delta(S) \times \Delta(S) \to \mathbf{R}$ 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 $|\eta(q,m)| \leqslant B$ for every q, m and satisfies the Lipschitz condition $|\eta(q,m) - \eta(q',m)| \leqslant B|q-q'|_1$ for every $q, q' \in \Delta(S)$.

The payoff $V_0(\sigma,\rho)$ and $d(\sigma)$ satisfy:

$$\begin{aligned} V_0(\sigma,\rho) &= \int \eta(\tau_\sigma(m), m)\ \chi_\sigma(\mathrm{d}m), \text{ and} \\ d(\sigma) &\geqslant \gamma \int |\tau_\sigma(m) - m|_1\ \chi_\sigma(\mathrm{d}m). \end{aligned} \tag{16}$$

And the payoff under σ' satisfy

$$\begin{aligned} &V_\lambda(\sigma',\rho) = (1 - L|p-q|_1) \int \eta(m,m)\ \chi_\sigma(\mathrm{d}m) + L|p-q|_1 \eta(q',q') \\ &\geqslant \int (\eta(\tau_\sigma(m), m) - B|\tau_\sigma(m) - m|_1)\ \chi_\sigma(\mathrm{d}m) - 2L|p-q|_1 B \\ &\geqslant V_\lambda(\sigma,\rho) + \left(\lambda - \frac{B}{\gamma}\right) d(\sigma) - 2L|p-q|_1 B \geqslant V_\lambda(\sigma,\rho) + \left(\lambda - \frac{B(1+2L)}{\gamma}\right) d(\sigma), \end{aligned} \tag{17}$$

where the first inequality follows from the Lipschitz condition and the bound on η, the second inequality from (16), and the third inequality from (15). Since σ is not strictly dominated by σ' we choose ρ such that such that $V_\lambda(\sigma',\rho) \leqslant V_\lambda(\sigma,\rho)$. If $\lambda > B(1+2L)/\gamma$, then it follows from (17) that $d(\sigma) = 0$, i.e., that σ is calibrated. ∎

Proof of Proposition 5.1. It is a standard revelation-principle argument 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 7.4 it follows that no miscalibrated strategy gives Sender more than CP against ρ^*. Therefore, σ^* is a best response against ρ^*. ■

7.4 Proof of Lemma 7.5

The first step of the proof of Proposition 5.2 is Lemma 7.5 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 given this strategy and Receiver's best response is arbitrarily close to his commitment payoff.

Lemma 7.5. *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) support$(\sigma) = \{x_a\}_{a \in A}$; and (iii) if ρ is Receiver's best response to σ, then $V_\lambda(\sigma, \rho) > CP - \varepsilon$.*

We prove Lemma 7.5 with the help of Lemma 7.6 below. The assertion in the lemma is standard. See Balkenborg, Hofbauer and Kuzmics (2015) and Brandenburger, Danieli and Friedenberg (2017) for similar arguments.

Lemma 7.6. *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 $BR(q) = \{a\}$.*

Returning to the proof of Lemma 7.5, 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 4. *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 7.6, 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 $a \neq \bar{a}$. ■

Claim 5. *There exists* $x_a \in \Delta(S), t_a > 0$ *for each* $a \in A$ *such that* $p = \sum_{a\in A} t_a x_a$*,* $\sum_{a\in A} t_a = 1$*,* $BR(x_a) = \{a\}$*, and* $\sum_{a,s} t_a x_a(s) v(s,a) > CP - \varepsilon$*.*

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 4 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 $t_a > 0$ be such that $z_a = t_a x_a$. Since $p \in \Delta(S)$, $x_a \in \Delta(S)$ for $a \in A$, and $p = \sum_{a\in A} t_a x_a$, it follows that $\sum_{a\in A} t_a = 1$. ■

By the splitting lemma there exists a calibrated strategy σ such that $\chi_\sigma = \sum_a t_a \delta_{x_a}$. From Claim 5 the strategy σ satisfies the requirements in Lemma 7.5.

7.5 Proof of Proposition 5.2

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 5; 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 6. *In the BNE of the auxiliary game, $\kappa(w^*_a/|w^*_a|, x_a) \leqslant 2B/\lambda$ 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:

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

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

Claim 7. *Sender's payoff in the Nash equilibrium of the auxiliary game is at least $CP - \varepsilon$.*

Proof. Sender can use the strategy $w_a = t_a x_a$ for every $a \in A$ and $w_\diamond = 0$. By Claim 5, Sender's payoff is at least $CP - \varepsilon$. ■

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 7. Let λ be sufficiently high such that $a \in BR(q)$ for every $q \in \Delta(S)$ such that $\kappa(q, x_a) \leqslant 2B/\lambda$. 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 6 it follows that, for every belief on-path, Receiver best responds to that belief.

7.6 Proof of Proposition 5.3

Lemma 7.4 implies that for $\lambda > \bar{\lambda}$, calibrated Sender's strategies are the strategies that are not strictly dominated.

Since any message m is on the path of some calibrated strategy of Sender, it follows that after observing a message m, Receiver's action must be in $BR(m)$.

Finally, by Lemma 7.5, for every Receiver's strategy ρ such that support$(\rho(\cdot|m)) \subseteq BR(m)$, 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 the commitment payoff against some such strategy ρ of Receiver.

Remark 1. The idea of strong belief in rationality is that if a message m is on the path of some strategy of Sender that is not strictly dominated, then Receiver's rationalizable actions for this message are the best responses against the conditional beliefs over states under such strategies of Sender. Thus, EFR requires that, for each message m and each Sender's strategy σ, we define what it means for m to be on the path of σ. The problem is that, 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. One extreme approach is that a message is on path if it appears with a strictly positive probability under the strategy. Another extreme approach is that every message is on path. In our model, both these approaches require that Receiver believe, after every message, that Sender plays a calibrated strategy. Since Receiver's best responses to all calibrated strategies are the same, any reasonable definition of EFR in our setup will give the same conclusion. □

References

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.

Bardhi, Arjada, and Yingni Guo. 2018. "Modes of persuasion toward unanimous consent." *Theoretical Economics*, 13(3): 1111–1149.

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.

Brandenburger, Adam, Alex Danieli, and Amanda Friedenberg. 2017. "How Many Levels Do Players Reason? An Observational Challenge and Solution." *Working paper.*

Chevalier, Judith A., and Dina Mayzlin. 2006. "The Effect of Word of Mouth on Sales: Online Book Reviews." *Journal of Marketing Research*, 43(3): 345–354.

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.

Foster, Dean P., and Rakesh V. Vohra. 1998. "Asymptotic Calibration." *Biometrika*, 85(2): 379–390.

Fréchette, Guillaume, Alessandro Lizzeri, and Jacopo Perego. 2018. "Rules and Commitment in Communication." *Working paper.*

Friedenberg, Amanda. 2019. "Bargaining Under Strategic Uncertainty: The Role of Second-Order Optimism." *Working paper.*

Gehlbach, Scott, and Konstantin Sonin. 2014. "Government control of the media." *Journal of Public Economics*, 118: 163 – 171.

Green, Jerry R., and Nancy L. Stokey. 2007. "A Two-person Game of Information Transmission." *Journal of Economic Theory*, 90–104.

Henry, Emeric, and Marco Ottaviani. 2018. "Research and the Approval Process: The Organization of Persuasion." *American Economic Review, forthcoming.*

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.

Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. 2017. "Persuasion of a Privately Informed Receiver." *Econometrica*, 85(6): 1949–1964.

Laraki, Rida. 2004. "On the Regularity of the Convexification Operator on a Compact Set." *Journal of Convex Analysis*, 11(1): 209–234.

Lehrer, Ehud. 2001. "Any Inspection is Manipulable." *Econometrica*, 69(5): 1333–1347.

Lipnowski, Elliot, and Doron Ravid. 2018. "Cheap Talk with Transparent Motives." *Working paper.*

Lipnowski, Elliot, Doron Ravid, and Denis Shishkin. 2018. "Persuasion via Weak Institutions." *Working paper.*

Luca, Michael, and Georgios Zervas. 2016. "Fake It Till You Make It: Reputation, Competition, and Yelp Review Fraud." *Management Science*, 62(12): 3412–3427.

Mayzlin, Dina, Yaniv Dover, and Judith Chevalier. 2014. "Promotional Reviews: An Empirical Investigation of Online Review Manipulation." *American Economic Review*, 104(8): 2421–55.

Minaudier, Clement. 2018. "The value of confidential policy information: persuasion, transparency, and influence." *Working paper.*

Min, Daehong. 2018. "Bayesian Persuasion under Partial Commitment." *Working paper.*

Murphy, Allan H., and Robert L. Winkler. 1987. "A General Framework for Forecast Verification." *Monthly Weather Review*, 115: 1330–1338.

Nguyen, Anh, and Teck Yong Tan. 2018. "Bayesian Persuasion under Costly Misrepresentation." *Working paper.*

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.

Perez-Richet, Eduardo, and Vasiliki Skreta. 2018. "Test Design under Falsification." *Working paper.*

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.

Zamir, Shmuel. 1992. "Repeated Games of Incomplete Information: Zero-sum." *Handbook of Game Theory with Economic Applications*, 1: 109–154.