Theoretical Economics 16 (2021), 477–506 1555-7561/20210477

Costly miscalibration

Yingni Guo
Department of Economics, Northwestern University

Eran Shmaya
Department of Economics, Stony Brook University

We consider a platform that 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 and, thus, provide a microfoundation for the commitment solution.

Keywords. Calibration, miscalibration, cheap talk, commitment, Bayesian persuasion, e-commerce platform.

JEL classification. D81, D82, D83.

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 only 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 its customers, in which customers trust the platform's forecasts in an equilibrium even though they do not observe the algorithm.

Yingni Guo: yingni.guo@northwestern.edu
Eran Shmaya: eran.shmaya@stonybrook.edu
Financial support from National Science Foundation Grant SES-1530608 is gratefully acknowledged.

 https://doi.org/10.3982/TE3991

Environment 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. The sender provides information in the form of probabilistic distributions. For example, Redfin defines a "hot home" as one that has a 70 percent chance or higher of having an accepted offer within 2 weeks of its debut. Given the sender's strategy, for each claim that the sender makes, we can calculate the true conditional distribution over the states. We say that the sender's strategy is *miscalibrated* if there is a discrepancy between what he claims and the true conditional distribution, for example, if only 50 percent of the homes that Redfin identifies as hot homes have an accepted offer within 2 weeks of their debut. Likewise, the sender's strategy is *calibrated* if what he claims is always the same as the true conditional distribution.

The sender incurs a cost that is a function of the miscalibration measure, which we define in the paper. The sender's cost also depends on a parameter, which we call the *cost intensity*. The cost intensity indicates how severely the sender is punished for each unit of miscalibration or how easy it is to collect data. The sender's cost is a proxy for the reputation damage from making incorrect assertions. The fact that platforms have various degrees of reputation concern is well documented in the empirical literature. Mayzlin et al. (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. There is also a large experimental literature showing that people have a preference for being seen as honest (e.g., Abeler et al. 2019).

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 to buy it. We call such a game *a promotion game*. In addition to our platform–customer application, 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 (2019) and Perez-Richet and Skreta (2021), the informational lobbying game in Bardhi and Guo (2018), Guo and Shmaya (2019), and Minaudier (2020), and the media censorship game in Gehlbach and Sonin (2014) and Kolotilin et al. (2017).

Our first result is that the sender's has an 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 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. These features show how the sender leverages his low commitment power to gain credibility.

Our result suggests that the sender makes more "extreme" statements than in the commitment solution, so as to lend himself credibility. For instance, Airbnb labels exceptional hosts as "superhosts." One criterion is that superhosts cancel less than 1 percent of the time (i.e., a maximum of 1 cancellation in 100 bookings). Given this extremely high standard, a customer who experienced a cancellation can confidently suspect miscalibration if another customer recently experienced the same.[1] Hence, Airbnb retains the credibility of its superhost label.

We also show that the sender's optimal equilibrium payoff is monotone-increasing in the cost intensity. This is because the set of calibrated equilibria expands as the cost intensity increases. Given that the sender's optimal equilibrium is calibrated, a higher cost intensity thus leads to a higher optimal equilibrium payoff for the sender. Recall that the cost intensity models the degree of reputation damage to the sender from making incorrect forecasts. The monotonicity result shows that the cost intensity can also be interpreted as a proxy for the degree of the sender's commitment power. The more the sender is likely to suffer from making incorrect assertions, the more credible he becomes in his communication with the receiver. A higher payoff to the sender thus results.

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 a generic assumption on the sender–receiver 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 an e-commerce 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 provide a microfoundation for the commitment solution; our model bridges the cheap-talk and the persuasion models, and can be used to analyze the middle ground where neither the talk is completely cheap nor the commitment absolute.

Related literature Our paper contributes to the literature on strategic communication with a lying cost (e.g., Kartik et al. 2007, Kartik 2009). The key difference is that in the

[1]For a case in which miscalibration occurred, see Patrick Collinson, "My Airbnb Superhost Stay Turned into a Super Disaster," *The Guardian* (https://www.theguardian.com/money/2019/dec/21/my-airbnb-superhost-stay-turned-into-a-super-disaster), December 21, 2019.

costly-lying model, the message space is typically the state space, so the sender's only option is to announce a state. Consequently, the only way for the sender to avoid a lying cost is to reveal everything he knows. In contrast, our sender avoids a lying cost as long as what he claims is true (i.e., his asserted distribution is the same as the true distribution).

The implication of this modeling difference is best seen when the lying cost is high. High lying costs translate into commitment power in our model, but into full disclosure in the costly-lying model. These different high cost results capture different intuitions about what happens when lying becomes very costly. In Remark 1 in Section 2, we formalize the notion that our model nests the costly-lying model.

Our notion of miscalibration differs from the notion of lying in Sobel (2020), where the message space includes some subsets of the state space. The sender has lied if the state does not belong to the subset of the state space that he has announced.

Our paper is also related to Perez-Richet and Skreta (2021), 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 a receiver's optimal test in which the sender does not falsify always exists. 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 does not miscalibrate in his optimal equilibrium, and also show that this is not always the case. Our paper is also related to Nguyen and Tan (2021). 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 that is not observable, so the receiver observes only the message (i.e., the customer sees only the platform's message, not its algorithm).

Our results show that our model bridges the cheap-talk and persuasion models (e.g., Rayo and Segal 2010, Kamenica and Gentzkow 2011). In this aspect, our paper is related to Fréchette et al. (2020), Lipnowski et al. (2019), and Min (2020). In these papers, the sender first publicly announces a test. With an exogenously given probability, the message is given by this test. With complementary probability, the sender can secretly choose a different message. The receiver in all 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 their probabilistic commitment models. Lipnowski and Ravid (2020) 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.

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 in 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 miscalibration measure to study these games.[2]

This paper is also 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 et al. (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 concerns affect the effectiveness of their communication.

Structure of the paper Section 2 presents the model. In Section 3, we characterize the sender's optimal equilibrium for promotion games. Section 4 presents our results when the cost intensity is high. Section 5 extends our analysis to setups in which the sender has some information on the state, but does not necessarily know it. Section 6 contains the proofs.

2. 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 let M be 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 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 $\operatorname{support}(\sigma) = \bigcup_s \operatorname{support}(\sigma(\cdot|s))$. 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 $m \in \operatorname{support}(\sigma)$, and is defined arbitrarily for $m \notin \operatorname{support}(\sigma)$.

For the rest of the paper, we assume that $M = \Delta(S)$. In Section 2.1, we introduce the concepts of calibrated strategies and miscalibration. These concepts are independent of Sender's interaction with Receiver. In Section 2.2, we review the model of sender–receiver games and the definition of a cheap-talk equilibrium. Section 2.3 presents our definition of an equilibrium with costly miscalibration.

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

2.1 *Sender's calibrated strategies and miscalibration*

Sender's message m is an *asserted* distribution over states. We make the assumption that $M = \Delta(S)$ for two reasons. First, it is hard to define a notion of costly miscalibration when the message space is not specified. Second, the space of beliefs is exactly the space that captures the finest grain of information needed by Receiver in a sender–receiver game.

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.[3] 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 $\kappa(\sigma)$ for Sender's strategy σ. To define this measure, let $d : \Delta(S) \times \Delta(S) \to \mathbb{R}_+$ be a continuous function, where $d(q, m)$ measures the distance between a message $m \in \Delta(S)$ and a truth $q \in \Delta(S)$. We assume that $d(q, m) = 0$ if and only if $m = q$. Let

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

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 $\kappa(\sigma) = 0$.

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 defines a hot home as one that has a 70 percent chance or higher of having an accepted offer within 2 weeks of its debut. Its 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.[4] 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 2 weeks of its debut was 53.1. This is different from 70 percent.

Let τ be the true probability that a hot home has an accepted offer in 2 weeks. The p-value for the null hypothesis $\tau \geq 70\%$ against the alternative hypothesis $\tau < 70\%$ is smaller than 0.0001. Redfin's forecast was not calibrated.[5] $\diamondsuit$

2.2 *Sender–receiver games*

Let A be a finite set of *actions* by Receiver. Let $v, u : S \times A \to \mathbb{R}$ be, respectively, Sender's and Receiver's payoff functions. Receiver's strategy is given by a Markov kernel ρ from

[3]The term a.s. in our paper means "almost surely with respect to the probability distribution over messages induced by σ." Recall that τ_σ is defined up to a set of messages with probability 0.

[4]We examined Cook County in Illinois, which includes 165 zip codes, and collected homes that Redfin had identified as hot homes over 2 weeks (03/14/2018–03/27/2018).

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

M to A, with the interpretation that Receiver randomizes an action from $\rho(\cdot|m)$ after message m.

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 becomes clear later when we define V_λ for every $\lambda \geq 0$.

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.

2.3 *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

$$V_\lambda(\sigma, \rho) = V_0(\sigma, \rho) - \lambda \kappa(\sigma),$$

where $\kappa(\sigma)$ is the miscalibration measure and $\lambda \geq 0$ is a parameter that 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, let (σ, ρ) be an equilibrium for $\lambda \geq 0$. If σ is calibrated so Sender pays no miscalibration cost on path, then (σ, ρ) is an equilibrium for a higher $\lambda' > \lambda$.*

Proof. Since (σ, ρ) is an equilibrium for $\lambda \geq 0$, ρ is a best response by Receiver to σ. To show that (σ, ρ) is an equilibrium for λ', we need to show that Sender's payoff from deviating to any σ' is smaller than his equilibrium payoff: $V_{\lambda'}(\sigma', \rho) \leq V_{\lambda'}(\sigma, \rho)$. This is the case, since

$$V_{\lambda'}(\sigma', \rho) = V_0(\sigma', \rho) - \lambda' \kappa(\sigma') \leq V_0(\sigma', \rho) - \lambda \kappa(\sigma') = V_\lambda(\sigma', \rho) \leq V_\lambda(\sigma, \rho) = V_{\lambda'}(\sigma, \rho).$$

The first inequality follows from the fact that $\lambda' > \lambda$ and $\kappa(\sigma') \geq 0$. The second inequality follows from the fact that (σ, ρ) is an equilibrium for λ. The last equality follows from the fact that σ is calibrated so $\kappa(\sigma) = 0$. $\square$

Remark 1. We now formalize the relationship between our model and the costly-lying model (Kartik et al. 2007, Kartik 2009). In the costly-lying model, Sender's lying measure is given by some $\hat{d}(s, s')$ if the state is s and Sender declares some $s' \in S$. For example, $\hat{d}(s, s') = (s - s')^2$. In our model, if Sender is restricted to announce deterministic messages (i.e., $m \in \{\delta_{s'} : s' \in S\}$) and for any such m the distance function $d(q, m)$ is linear in q,

$$d(q, \delta_{s'}) = \sum_s q(s)\hat{d}\big(s, s'\big),$$

then our model reduces to the costly-lying model.

3. Promotion games

Given our interest in platforms, we study a class of sender–receiver games that we call *promotion games*: Receiver decides whether to buy, and Sender's payoff is 1 if Receiver buys and 0 otherwise. Formally, we assume that $A = \{\mathrm{B}, \mathrm{NB}\}$, and that $v(s, \mathrm{B}) = 1$ and $v(s, \mathrm{NB}) = 0$ for each s. It is without loss to set $u(s, \mathrm{NB}) = 0$ for each s. It is also without loss to identify each state with Receiver's payoff from action B in that state. Thus, we set $s = u(s, \mathrm{B})$ for each s.

Promotion games are rich enough to capture games with binary actions for Receiver and state-independent payoff for Sender. To avoid triviality, we assume that $\sum_s p(s)s < 0$ and that $s > 0$ for some s. Hence, Receiver strictly prefers not to buy given the prior belief and strictly prefers to buy for some state.

It is well known that in the cheap-talk case (i.e., $\lambda = 0$), no promotion is possible in the sense that Receiver never buys. Indeed, suppose on the contrary that Receiver buys with positive probability after some message. Then Sender will announce only those messages that induce the highest buying probability. This means that Receiver is willing to buy after every message on path, contradicting the assumption that Receiver strictly prefers not to buy given the prior belief.

We now characterize 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.

Before presenting our result, we introduce a promotion-game example to illustrate the main concepts introduced in the previous section. We also use this example later to illustrate our result.

Example 2. Consider a promotion game with $S = \{-2, -1, 1\}$ and the prior $p = (1/8, 1/2, 3/8)$.

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 state distribution is $(1/5, 1/5, 3/5)$. Each column shows the probabilities with which Sender sends m_0 or m_1 in each state.

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 $\tau(m_1) = (1/5, 1/5, 3/5)$. Each message coincides with the true conditional distribution given that message, so this strategy is calibrated.

Message	State		
	-2	-1	1
$m_0 = (0, 1, 0)$	0	$\frac{3}{4}$	0
$m_1 = (\frac{1}{5}, \frac{1}{5}, \frac{3}{5})$	1	$\frac{1}{4}$	1

TABLE 1. A calibrated strategy.

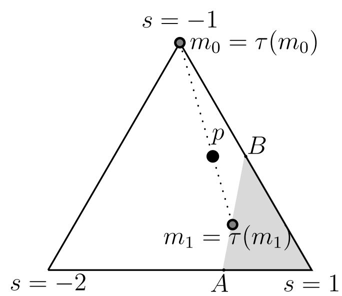

FIGURE 1. Strategy in Table 1.

Message	State		
	-2	-1	1
$m_0 = (0, 1, 0)$	0	$\frac{1}{2}$	0
$m_1 = (\frac{1}{5}, \frac{1}{5}, \frac{3}{5})$	1	$\frac{1}{2}$	1

TABLE 2. A miscalibrated strategy.

Figure 1 illustrates this strategy. Each point in the triangle represents a distribution over states. Receiver is willing to buy if his belief after a message is in the light-gray area. The black dot is the prior p. The gray dots m_0 and m_1 are the messages that Sender uses. The black circles $\tau(m_0)$ and $\tau(m_1)$ 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 $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$ in states -2, -1, and 1, respectively.

Consider another strategy, shown in Table 2 and Figure 2. Sender uses the same messages m_0 and 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_1 more often than before. The true conditional distribution given m_1 is $\tau(m_1) = (1/6, 1/3, 1/2)$, which differs from m_1. Thus, this strategy is not calibrated.

If Sender uses this strategy to make recommendations, then the probability distribution of the products about which Sender announces m_1 is in fact $\tau(m_1)$ in Figure 2. Therefore, Sender's claim about these products—that they are $(1/5, 1/5, 3/5)$—is not

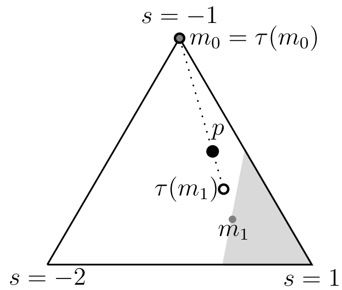

FIGURE 2. Strategy in Table 2.

true. If an outside group collects data and performs some statistical test, Sender might be caught.

Whenever Sender sends a message that has a true conditional distribution that differs from its asserted distribution, Sender pays a miscalibration cost. We now use the total-variation distance, $|m-\tau(m)|_1/2$, 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

$$\frac{1}{2}\left|\left(\frac{1}{5},\frac{1}{5},\frac{3}{5}\right)-\left(\frac{1}{6},\frac{1}{3},\frac{1}{2}\right)\right|_1=\frac{1}{2}\left(\left|\frac{1}{5}-\frac{1}{6}\right|+\left|\frac{1}{5}-\frac{1}{3}\right|+\left|\frac{3}{5}-\frac{1}{2}\right|\right)=\frac{2}{15}.$$

Since Sender announces message m_1 with probability 3/4, the miscalibration measure is $3/4\cdot 2/15=1/10$ and the miscalibration cost is $\lambda/10$. ◊

Theorem 1 below enables us to fully characterize Sender's optimal equilibrium for any promotion game and any $\lambda>0$. For this theorem, we assume that d is a Wasserstein distance. Wasserstein distances are the natural distances over $\Delta(S)$ for a metric space S. Let $\overline{d}$ be any metric on S; then the *Wasserstein distance d over $\Delta(S)$ induced by $\overline{d}$* is given by

$$d(q,q')=\inf\mathbb{E}\overline{d}(Q,Q'),$$

where the infimum ranges over all pairs Q and Q' of S-valued random variables with marginal distributions q and q', respectively. (Such a pair Q and Q' is called a coupling of q and q'.) A Wasserstein distance d is a convex function of q and q'. It is itself a metric, so it satisfies the triangular inequality.[6]

The total-variation distance is a Wasserstein distance when $\overline{d}$ is the discrete metric so that $\overline{d}(s,s')=1$ for any $s\neq s'$. In a promotion game, it is also natural to choose the metric $\overline{d}(s,s')=|s-s'|$, where s and s' are Receiver's payoffs from buying in states s and

[6]Wasserstein distance, also known as Kantorovich–Monge–Rubinstein distance and earth mover's distance, is a central concept in optimal transport theory. It is equivalently defined as the cost of the optimal transport plan for moving the mass in q to that in q'. (See Galichon 2018 for numerous applications in economics.) In our setup, q and q' correspond to a true conditional distribution and an asserted distribution. Frogner et al. (2015) provide a similar application of Wasserstein distance as a loss measure for online learning algorithms.

s', respectively. In this case, the cost of sending the message that “the state is s with probability 1” when the state is s' is increasing in Receiver’s payoff difference $|s - s'|$. For instance, when the state is -2, sending the message that “the state is 1 with probability 1” is more costly for Sender than sending the message that “the state is -1 with probability 1.” Among the frequently used distances, Kullback–Leibler and Euclidean distances are not Wasserstein distances.

Theorem 1. *Let $\lambda > 0$ and let d be a Wasserstein distance. In a promotion game, Sender’s optimal equilibrium has 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 between buying and not buying when the distribution over states is m_1 and m_0 is on the boundary of the simplex $\Delta(S)$.*

Theorem 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 following corollary.

Corollary 1. *Let $\lambda > 0$ and let d be a Wasserstein distance. In a promotion game, Sender’s optimal equilibrium payoff is monotone-increasing in λ.*

Corollary 1 shows that the cost intensity λ can be interpreted as a 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 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 so as 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 and m_1 in Figure 1 instead of m_- and m_+ in Figure 3 below). The rationale is that using messages that are distant from each other increases Sender’s cost of miscalibration, which deters Sender from deviating. Suppose that Sender splits the prior p into m_0 and m_1 such that Receiver buys after m_1. We show in Example 2 that for any fixed increase in the probability of announcing m_1, the miscalibration measure is larger when the distance between m_0 and m_1 is larger.

To summarize, Sender gets a positive payoff even with a low commitment power (i.e., 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 what a more credible platform would use, and its customers are less responsive to the purchase message in the sense that they buy only occasionally after the purchase message.

EXAMPLE 2 (Continued). According to Theorem 1, Sender splits the prior into two posteriors: one on the line between A and B in Figure 1, and the other on the boundary of the simplex $\Delta(S)$. He announces those posteriors truthfully. Figure 1 gives an example of such a strategy used by Sender. On path, Sender announces m_1 with probability 5/8 and m_0 with probability 3/8.

Receiver buys with a positive probability if m_1 is sent and does not buy after any other message. Hence, Sender is tempted to say m_1 more often. When the cost intensity λ is low, we must lower Receiver's buying probability after m_1 so as to lower Sender's temptation to say m_1 more often. This lowered temptation, combined with the miscalibration cost, prevents Sender from deviating.

To illustrate, suppose that Sender deviates by saying m_1 with probability $5/8 + \varepsilon$ (instead of 5/8) for some small $\varepsilon > 0$. The true distribution $\tau(m_1)$ given m_1 solves

$$\left(\frac{3}{8} - \varepsilon\right)m_0 + \left(\frac{5}{8} + \varepsilon\right)\tau(m_1) = p, \quad \Longrightarrow \quad \tau(m_1) = \left(\frac{1}{8\varepsilon + 5}, \frac{8\varepsilon + 1}{8\varepsilon + 5}, \frac{3}{8\varepsilon + 5}\right).$$

The miscalibration measure is the probability that m_1 is sent times the distance $d(m_1, \tau(m_1))$:

$$\left(\frac{5}{8} + \varepsilon\right)d\bigl(m_1, \tau(m_1)\bigr) = \left(\frac{5}{8} + \varepsilon\right)\frac{1}{2}\bigl|m_1 - \tau(m_1)\bigr|_1 = \frac{4}{5}\varepsilon. \tag{1}$$

The corresponding miscalibration cost is λ times this miscalibration measure, so it is $4\lambda\varepsilon/5$.

This deviation strategy increases by ε the probability that m_1 is sent, and also increases the miscalibration cost by $4\lambda\varepsilon/5$. To deter Sender from taking this deviation, Receiver's buying probability after m_1 cannot be higher than $4\lambda/5$.

For $\lambda \in (0, 5/4]$, the calibrated strategy in Figure 1, combined with a buying probability of $4\lambda/5$ after m_1, constitutes an equilibrium. Sender's equilibrium payoff is $\lambda/2$. It can be readily verified that other calibrated equilibria that are consistent with Theorem 1 give Sender a lower payoff, so Sender's optimal equilibrium payoff is $\lambda/2$.

For $\lambda \geq 2$, we argue that Sender obtains his commitment payoff 3/4 by using the strategy in Figure 3. He splits the prior into $m_- = (1/2, 1/2, 0)$ and $m_+ = (0, 1/2, 1/2)$, and announces those posteriors truthfully. Receiver buys for sure after m_+ and does not buy after any other message. On path, Sender announces m_+ with probability 3/4 and m_- with probability 1/4.

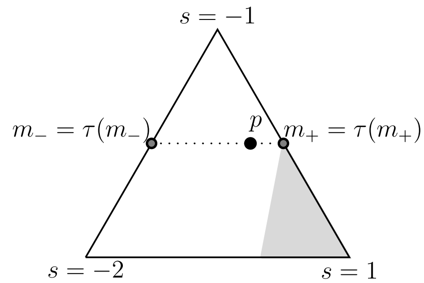

FIGURE 3. Sender's optimal equilibrium when $\lambda \geq 2$.

Suppose that Sender deviates by saying m_+ with probability $3/4+\varepsilon$ for some small $\varepsilon > 0$.[7] The true distribution $\tau(m_+)$ given m_+ solves

$$\left(\frac{1}{4}-\varepsilon\right)m_- + \left(\frac{3}{4}+\varepsilon\right)\tau(m_+) = p, \quad \Longrightarrow \quad \tau(m_+) = \left(\frac{2\varepsilon}{4\varepsilon+3}, \frac{1}{2}, \frac{3}{8\varepsilon+6}\right).$$

The miscalibration measure is the probability that m_+ is sent times the distance $d(m_+, \tau(m_+))$:

$$\left(\frac{3}{4}+\varepsilon\right)d\big(m_+, \tau(m_+)\big) = \left(\frac{3}{4}+\varepsilon\right)\frac{1}{2}\big|m_+ - \tau(m_+)\big|_1 = \frac{1}{2}\varepsilon. \tag{2}$$

The corresponding miscalibration cost is λ times this miscalibration measure, so it is $\lambda\varepsilon/2$. For any $\lambda \geq 2$, this miscalibration cost suffices to deter Sender from deviating, even if Receiver buys for sure after m_+.

For $\lambda \in (5/4, 2)$, 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 intensity λ.

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 and in Sender's optimal equilibrium when $\lambda \geq 2$. With two distant messages like m_0 and m_1, if Sender increases the probability of saying m_1 by ε, he has to incur a larger miscalibration measure, compared with the case having two nearby messages like m_- and m_+. This can be seen by observing that the miscalibration measure in (1) is larger than that in (2).

Figure 4 uses a seesaw analogue to illustrate the idea that more distant messages imply a larger miscalibration measure. In a calibrated strategy, Sender splits the prior p into posteriors m_0 and m_1. The two squares, m_0 and m_1, are balanced around a pivot at the prior. The size of each square corresponds to the probability of saying that message. If Sender increases the probability of saying m_1 by ε, this causes the posterior to change to $\tau(m_1)$. The more distant m_0 and m_1 are from each other (i.e., the longer the seesaw is), the larger the miscalibration measure is from this move of ε probability. Thus, when the cost intensity is low, Sender uses distant messages to compensate for his low commitment power to deter himself from deviating. $\diamondsuit$

[7]The probabilities with which Sender sends m_+ in states -2, -1, and 1 are 4ε, $3/4+\varepsilon$, and 1, respectively.

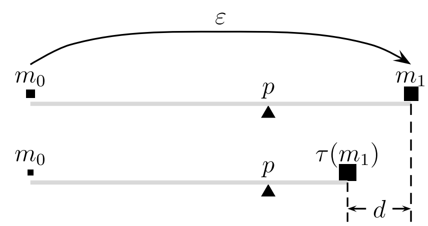

FIGURE 4. More distant messages imply larger miscalibration measure.

The result in Theorem 1, namely, that Sender gets his optimal payoff in a calibrated equilibrium, may look familiar from the cheap-talk case. Indeed, in the cheap-talk case (i.e., $\lambda = 0$), a simple revelation-principle argument implies that for every equilibrium (which is not necessarily Sender's optimal equilibrium), there exists a calibrated equilibrium that induces the same distribution over states and actions. However, this revelation-principle argument does not apply when $\lambda > 0$. Example 3 below provides a miscalibrated equilibrium that induces a distribution over states and actions that cannot be induced by any calibrated equilibrium. Moreover, Theorem 1 does not hold for all distance functions. In Section 4.1, we show that when d is the Kullback–Leibler distance, Sender's optimal equilibrium is not calibrated.

EXAMPLE 3. Consider a promotion game with $S = \{-1, 1\}$ and the prior $p = (3/4, 1/4)$. Let d 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 1. He splits the prior $1/4$ into 0 and $1/2$, and announces the posterior beliefs truthfully, so his strategy is calibrated. Receiver buys after message $1/2$ with probability $\lambda/2 = 1/2$ and does not buy otherwise. Sender's equilibrium payoff is $1/4$.

A sender-optimal, miscalibrated equilibrium also exists. Sender still splits the prior into 0 and $1/2$. When 0 is realized, he announces message 0. However, when $1/2$ is realized, he announces message 1. Receiver buys after message 1 with probability $\lambda = 1$ and does not buy otherwise. Sender's equilibrium payoff is the total probability that Receiver buys, namely, $1/2$, minus the miscalibration cost $1/4$.

Both equilibria generate the same payoff of $1/4$ to Sender, but different distributions over states and actions. Conditional on state 1, Receiver buys more often under the second than under the first equilibrium. In fact, no calibrated equilibrium generates the same distribution over states and actions that the second equilibrium does, since otherwise Sender's equilibrium payoff would be $1/2$, which is higher than his optimal equilibrium payoff of $1/4$. ◊

3.1 *Intuition for the proof of Theorem 1*

Theorem 1 relies on the triangular inequality and also on the following property of Wasserstein distances, which makes them well behaved under the splitting of probability distributions.

Lemma 1. *Let d be a Wasserstein distance over $\Delta(S)$ and let $q, q' \in \Delta(S)$. Let $q = t_1 q_1 + \cdots + t_n q_n$ be a splitting of q, where $t_i \geq 0$, $q_i \in \Delta(S)$ for all i, and $\sum_{i=1}^n t_i = 1$. Then there exists a splitting $q' = t_1 q'_1 + \cdots + t_n q'_n$ of q' with the same weights such that*

$$d(q, q') = t_1 d(q_1, q'_1) + \cdots + t_n d(q_n, q'_n). \tag{3}$$

To prove Theorem 1, we begin with an arbitrary equilibrium and construct a new equilibrium that has the properties in Theorem 1 and generates the same payoff for Sender. Roughly speaking, the proof can be divided into two steps. In the first, we make the equilibrium calibrated. In the second, we combine into a single message all the messages under which Receiver buys. Both steps use the same logic. We first construct a new strategy profile that has the desired property and generates the same payoff for Sender. We then show that for every deviation of Sender from the new profile he has a more profitable deviation from the original profile. Since the original profile was an equilibrium, the new profile must also be an equilibrium.

The first step uses the fact that Wasserstein distances satisfy the triangular inequality. Consider, for example, an equilibrium under which (i) with some probability, Sender announces a message m, after which Receiver buys with probability 1, and (ii) the true conditional distribution given m is $q = \tau_\sigma(m)$. So Sender's payoff when he announces m is $1 - \lambda d(q, m)$. Now we change Receiver's strategy so that he buys with probability $1 - \lambda d(q, m)$ after message q, and we change Sender's strategy so that he announces the true distribution q instead of m. This new strategy profile gives Sender the same payoff and his message is now calibrated. We now argue that Sender has no profitable deviation. Suppose that Sender considers deviating to some miscalibrated strategy, so that the true conditional distribution when he announces q is in fact q'. Sender then suffers a miscalibration cost of $\lambda d(q', q)$, so his payoff when he sends q is

$$1 - \lambda d(q, m) - \lambda d(q', q) \leq 1 - \lambda d(q', m),$$

where the inequality follows from the triangular inequality. The right-hand side is the payoff to Sender under the original profile if Sender deviates by announcing message m when the truth is q', and he suffers a miscalibration cost of $\lambda d(q', m)$. Since the original profile was an equilibrium, this implies that the new profile must also be an equilibrium. We note that this argument is not completely accurate, because if Receiver strictly prefers to buy when his belief about the states is q, then he must buy with probability 1 after message q in a calibrated equilibrium. Because of this nuisance, the proof has an additional modification to the equilibrium profile: it replaces q with a belief that makes Receiver indifferent.

The second step of the proof combines all the messages under which Receiver buys into a single message. Assume that, for example, we had a calibrated equilibrium in which there are n messages $q_1, \ldots, q_n$, after which Receiver buys with probability 1. Sender announces message q_i with probability r_i. We replace this strategy profile with a new profile. In the new profile, Sender announces the message $q = t_1 q_1 + \cdots + t_n q_n$ with probability $r = r_1 + \cdots + r_n$, where $t_i = r_i / r$, and Receiver buys with probability 1 after message q. We need to show that the new profile is an equilibrium. We again construct, for every deviation that Sender has in the new profile, a deviation strategy in the original profile that is more profitable. Suppose Sender deviates to a miscalibrated strategy, under which he sends the message q with some probability r', but the true distribution when he sends q is q'. Under this deviation, Sender will suffer a miscalibration cost of $\lambda d(q', q)$ when he sends q. Lemma 1 states that Sender has a deviation strategy in the original profile such that he sends message q_i with probability $r' t_i$ and the true distribution when he sends q_i is q'_i. Moreover, the overall miscalibration cost Sender suffers, conditional on sending one of $q_1, \ldots, q_n$, is also $\lambda d(q', q)$. Hence, the aforementioned deviation to the new profile cannot be profitable.

4. High cost intensity

We now turn to general sender–receiver games, in which the set A of Receiver's actions is finite and Sender's payoff $v(s, a)$ can depend on s. The next two propositions formalize the intuition that when the cost intensity is high, Sender gains the commitment power not to make false assertions. Proposition 1 shows that Sender can achieve his commitment payoff in an equilibrium. Proposition 2 shows that if the distance function has a kink, the only extensive-form rationalizable strategy of Sender is his strategy in the commitment solution.

Formally, the commitment payoff is given by $\mathrm{CP} = \max V_0(\sigma, \rho)$, 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.

4.1 *Equilibrium result*

On the surface, Proposition 1 can be understood from the following two observations. First, if Sender were exogenously restricted to using only calibrated strategies, then Receiver could take messages 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 over states given that message. Thus, our result asserts some lower hemicontinuity of the equilibrium correspondence. As usual, lower hemicontinuity is not straightforward. The difficulty 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 make the assumption that the sender–receiver game is generic. By "a generic set of games," we mean that the set of payoff functions $u : S \times A \to \mathbb{R}$ for which the assertion does not hold (viewed as a subset of $\mathbb{R}^{S\times A}$) is a closed set with an empty interior and a Lebesgue measure of 0. Our generic assumption requires that every Receiver's action be a *unique* best response to some beliefs over the states. The following proposition requires only that d be convex in q. Wasserstein, Kullback–Leibler, and Euclidean distances all satisfy this condition.

PROPOSITION 1. *Assume that d 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* $\mathrm{CP} - \varepsilon$.

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 best responds to the face value of each message in the support of this strategy 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 the following example, Sender's optimal equilibrium must be miscalibrated, yet his optimal equilibrium payoff converges to his commitment payoff.

EXAMPLE 4 (With a different distance function). Consider the promotion game in Example 3, but with this difference: the distance between a message and a truth is the Kullback–Leibler distance

$$d\big(\tau(m), m\big) = \sum_{s \in S} \tau(m)(s) \log \frac{\tau(m)(s)}{m(s)}.$$

As in Example 3, we use the probability of state 1 to represent the state distribution.

In the commitment solution, Sender splits the prior $1/4$ into 0 and $1/2$. Receiver buys for sure after the message $1/2$. However, for any intensity $\lambda \geq 0$, this strategy profile cannot be an equilibrium: Sender will deviate by announcing $1/2$ more often. This is because, in the case of the smooth Kullback–Leibler distance, for any message with full support, a small amount of miscalibration has no first-order impact on the miscalibration cost.

Sender has to "overshoot" to gain credibility. Consider the following equilibrium. Sender splits the prior into 0 and $1/2$. When $1/2$ is realized, Sender overshoots by saying $(1 - e^{-1/\lambda}/2) > 1/2$. Receiver buys for sure after $(1 - e^{-1/\lambda}/2)$ and does not buy after other messages. Sender's equilibrium payoff is $\lambda \log(2e^{1/\lambda} - 1)/4$, which goes from 0 to the commitment payoff $1/2$ as λ goes from 0 to ∞. Our Example 1 shows that Redfin indeed overshoots. ◊

We have shown that Sender achieves his commitment payoff (up to ε) in an equilibrium with some generic assumption about the game. Example 5 is a nongeneric game in which Sender's optimal equilibrium payoff for any λ is bounded away from his commitment payoff.

Example 5. Consider a promotion game with $S = \{0, -1\}$ and the prior $p = (p_0, 1 - p_0)$. Sender's commitment payoff is p_0: Sender fully reveals the state and Receiver takes action B if the state is 0 and action NB if the state is -1.

Unlike Example 4, there exists no belief over the states such that B is uniquely optimal for Receiver, so there is no room for Sender to overshoot. Hence, for a smooth distance function d, in every equilibrium, Receiver chooses only the safe action NB, so Sender's payoff is 0. ♢

4.2 *A rationalizablility result*

Up to now, following the cheap-talk literature, we used 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 announces 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 Example 2 with a high λ. 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 1, 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 buy the product.

BNE and 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. 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. How to extend such a definition to sender–receiver games with a continuum message space is not obvious. (See Remark 2 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 that is concentrated on Sender's strategies that are not strictly dominated.

Let $\mathrm{BR}(q)$ be the set of all of Receiver's best actions when his belief about the state is q:

$$\mathrm{BR}(q) = \arg\max_a \sum_s q(s)u(s, a). \tag{4}$$

Proposition 2. *Assume that there exists some* $\gamma > 0$ *such that* $d(q, m) \geq \gamma |q - m|_1$. *Then, in a generic set of games, there exists* $\bar{\lambda}$ *such that for every* $\lambda > \bar{\lambda}$, *the following statements hold:*

(i) Receiver's extensive-form rationalizable strategies are those that satisfy

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

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

To prove this result, we use the result (Lemma 5 in the Appendix) that if the distance function d has a kink, then for sufficiently high λ, a Sender's strategy σ is not strictly dominated if and only if it is calibrated.[8] 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 2 does not hold in the nongeneric game in Example 5 even when d has a kink. In that game, Receiver's strategy to always take NB is rationalizable and, therefore, every calibrated strategy is rationalizable for Sender.

5. Discussion: 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. We argue that our analysis and results extend to this environment as well.

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 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. Proofs

6.1 *Preliminaries*

We extend the distance function d to a function $d : \mathbb{R}_+^S \times \Delta(S) \to \mathbb{R}_+$ given by $d(\gamma q, m) = \gamma d(q, m)$ for every $q \in \Delta(S)$ and $\gamma \geq 0$. We extend the best-response correspondence $\mathrm{BR}(\cdot)$ to a correspondence from $\mathbb{R}_+^S$ to A given by the same formula (4).

[8] The assumption that $d(q, m) \geq \gamma |q - m|_1$ for some $\gamma > 0$ holds for Wasserstein distances and, in particular, the total-variation distance. It also holds for the Euclidean distance $d(q, m) = |q - m|_2$ and any other distance that is derived from a norm.

For a Sender's strategy σ, we let

$$\chi_\sigma = \sum_s p(s)\sigma(\cdot|s) \tag{5}$$

be the distribution over messages induced by σ. Lemma 2 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 2. *(i) For every strategy σ, we have*

$$\int \tau_\sigma(m)\chi_\sigma(\mathrm{d}m) = p. \tag{6}$$

(ii) If χ is a distribution over messages in $\Delta(S)$ such that $\int m\chi(\mathrm{d}m) = p$, then a calibrated strategy σ such that $\chi_\sigma = \chi$ exists.

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 $q_i \in \Delta(S)$, $t_i \geq 0$ 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.

6.2 *Proof of Theorem 1*

6.2.1 *Preliminaries* We can assume without loss that $\lambda = 1$, since all properties of the distance function d that are used in the proof still hold if we replace d by λd.

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 $f : \Delta(S) \to \mathbb{R}$, let $\operatorname{Lip} f : \Delta(S) \to \mathbb{R}$ be given by

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

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

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

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

Proposition 3. *Let $f, g : \Delta(S) \to \mathbb{R}$ be bounded. Then the following statements hold:*

(i) If d satisfies the triangular inequality, then $\operatorname{Lip}\operatorname{Lip} f = \operatorname{Lip} f$.

(ii) We have $\operatorname{Cav}\operatorname{Cav} f = \operatorname{Cav} f$.

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

Claim 2. *If $d(q, m)$ 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

$$V_1(\sigma, \rho) = \int \bigl(\rho(m) - d\bigl(\tau_\sigma(m), m\bigr)\bigr)\chi_\sigma(\mathrm{d}m) \leq \int \operatorname{Lip}\rho\bigl(\tau_\sigma(m)\bigr)\chi_\sigma(\mathrm{d}m) \leq \operatorname{Cav}\operatorname{Lip}\rho(p),$$

where the first inequality follows from $\rho(m) - d(\tau_\sigma(m), m) \leq \operatorname{Lip}\rho(\tau_\sigma(m))$ by the definition of $\operatorname{Lip}\rho$, and the second follows from (6) and the definition of the concave envelope 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 $t_i \geq 0$, and messages m_i, such that $p = \sum_i t_i q_i$, $\sum_i t_i = 1$, and

$$\sum_i t_i\bigl(\rho(m_i) - d(q_i, m_i)\bigr) \geq \operatorname{Cav}\operatorname{Lip}\rho(p) - \varepsilon. \tag{7}$$

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 d 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, respectively, $\overline{t}$ and $\overline{m}$, without violating (7).

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 2. Then (7) implies that Sender's payoff when using σ is at least $\operatorname{Cav}\operatorname{Lip}\rho(p) - \varepsilon$. $\square$

Lemma 1 shows that Wasserstein distances are well behaved under the 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 2.

Proof of Lemma 1. The direction $d(q, q') \leq t_1 d(q_1, q_1') + \cdots + t_n d(q_n, q_n')$ follows from convexity of d for every splitting $q' = t_1 q_1' + \cdots + t_n q_n'$ of q' with the weights $t_1, \dots, t_n$.

For the other direction, let Q and Q' be S-valued random variables with marginal distributions q and q', respectively, such that $d(q, q') = \mathbb{E}\overline{d}(Q, Q')$. Let X be a random variable that assumes values in $\{1, \dots, n\}$ such that $t_i = \mathrm{P}(X = i)$ and such that $q_i(s) = \mathrm{P}(Q = s|X = i)$ for every $i \in \{1, \dots, n\}$ 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

$$q'(s) = \mathrm{P}\bigl(Q' = s\bigr) = \sum_i \mathrm{P}(X = i)\mathrm{P}\bigl(Q' = s|X = i\bigr) = \sum_i t_i q_i'(s),$$

$$\sum_i t_i d\bigl(q_i, q_i'\bigr) \leq \sum_i \mathrm{P}(X = i)\mathbb{E}\overline{d}\bigl(Q, Q'|X = i\bigr) = \mathbb{E}\overline{d}\bigl(Q, Q'\bigr) = d\bigl(q, q'\bigr),$$

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

Corollary 2. *Let d be a Wasserstein distance. Then, for every function $f : \Delta(S) \to [0,1]$ that is 1-Lipschitz w.r.t. d, the concave envelope* $\operatorname{Cav} f$ *is also 1-Lipschitz w.r.t. d.*

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 $\operatorname{Cav} f(q) = t_1 f(q_1) + \cdots + t_n f(q_n)$. By Lemma 1, there exists a splitting $q' = t_1 q'_1 + \cdots + t_n q'_n$ of q' such that (3) holds. Therefore,

$$\operatorname{Cav} f(q') \geq \sum_i t_i f(q'_i) \geq \sum_i t_i \big(f(q_i) - d(q_i, q'_i)\big) = \operatorname{Cav} f(q) - d(q, q'),$$

where the first inequality follows from the definition of $\operatorname{Cav} f$ and the second from the fact that f is 1-Lipschitz. □

6.2.2 *Proof of Theorem 1* Let (σ, ρ) be an equilibrium. From the equilibrium condition for Sender and Claim 2, it follows that

$$\operatorname{Cav} \operatorname{Lip} \rho(p) = \int \big(\rho(m) - d(\tau_\sigma(m), m)\big) \chi_\sigma(\mathrm{d}m), \tag{8}$$

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

Let $\mathcal{A} = \{q \in \Delta(S) : \sum_s q(s)s \geq 0\}$ 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 $\overline{\chi}_\sigma = \chi_\sigma(\cdot|\tau_\sigma^{-1}(\mathcal{A}))$ and $\underline{\chi}_\sigma = \chi_\sigma(\cdot|\tau_\sigma^{-1}(\mathcal{A}^c))$. Let $\overline{q} = \int \tau_\sigma \, \mathrm{d}\overline{\chi}_\sigma$ and $\underline{q} = \int \tau_\sigma \, \mathrm{d}\underline{\chi}_\sigma$. Then

$$\chi_\sigma = (1-t)\underline{\chi}_\sigma + t\overline{\chi}_\sigma, \tag{9}$$

$$p = (1-t)\underline{q} + t\overline{q}. \tag{10}$$

From the equilibrium condition for Receiver, it follows that $\rho = 0$, $\underline{\chi}_\sigma$-almost surely. Therefore, from (8) and (9), it follows that

$$\operatorname{Cav} \operatorname{Lip} \rho(p) \leq t \int \big(\rho(m) - d(\tau_\sigma(m), m)\big) \overline{\chi}_\sigma(\mathrm{d}m). \tag{11}$$

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 $[p, \overline{q}]$ that is on the boundary of $\mathcal{A}$, i.e., $\sum_s q^*(s)s = 0$. From (10),

$$q^* = (1-t^*)\underline{q} + t^*\overline{q} = (1-t^*)\underline{q} + t^* \int \tau_\sigma \, \mathrm{d}\overline{\chi}_\sigma \tag{12}$$

for some $t^* \geq t$. 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 $\chi^* = (1 - t/t^*)\delta_{\underline{q}} + t/t^*\delta_{q^*}$.

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

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

where the first inequality follows from (12), the second inequality follows from the definition of Lip, and the third inequality follows from (11).

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

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

The first inequality and the last equality follow from Proposition 3(iii) and (ii), respectively. For the second equality, according to Corollary 2, $\operatorname{Cav}\operatorname{Lip}\rho(p)$ is 1-Lipschitz w.r.t. d. Therefore, $\operatorname{Lip}\operatorname{Cav}\operatorname{Lip}\rho(p) = \operatorname{Cav}\operatorname{Lip}\rho(p)$. By Claim 2 this implies that Sender's optimal payoff against ρ^* is at most $\operatorname{Cav}\operatorname{Lip}\rho(p)$, as desired.

6.3 *Proving Proposition 1*

6.3.1 *Lemmas*

Lemma 3. *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) $\mathrm{BR}(x_a) = \{a\}$ for each $a \in A$, (ii) $\mathrm{support}(\sigma) = \{x_a\}_{a\in A}$, and (iii) if ρ is Receiver's best response to σ, then $V_\lambda(\sigma, \rho) > \mathrm{CP} - \varepsilon$.*

We prove Lemma 3 with the help of Lemma 4 below. The assertion in Lemma 4 is standard. See Balkenborg et al. (2015) and Brandenburger et al. (2021) for similar arguments.

Lemma 4. *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 $\mathrm{BR}(q) = \{a\}$.*

Returning to the proof of Lemma 3, 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 3. *There exists $y_a \in \mathbb{R}_+^S \setminus \{0\}$ such that $p = \sum_{a\in A} y_a$ and $\mathrm{BR}(y_a) = \{a\}$.*

Proof. By Lemma 4, 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., $\text{BR}(q_a) = \{a\}$. 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 $\bar{a} \in \text{BR}(p - \sum_a tq_a)$. Let $y_{\bar{a}} = p - \sum_a tq_a + tq_{\bar{a}}$ and $y_a = tq_a$ for every $a \neq \bar{a}$. □

Claim 4. *There exists* $x_a \in \Delta(S)$, $t_a > 0$, *for each* $a \in A$ *such that* $p = \sum_{a\in A} t_a x_a$, $\sum_{a\in A} t_a = 1$, $\text{BR}(x_a) = \{a\}$, *and* $\sum_{a,s} t_a x_a(s) v(s, a) > \text{CP} - \varepsilon$.

Proof. Let (σ, ρ) be a commitment solution such that $\text{CP} = \sum_{s,a} v(s, a)\pi_{\sigma,\rho}(s, a)$, where $\pi_{\sigma,\rho}$ is the distribution induced by the profile (σ, ρ) over $S \times A$. Let $z_a(s) = \varepsilon/(2B)y_a(s) + (1 - \varepsilon/(2B))\pi_{\sigma,\rho}(s, a)$, where y_a is given by Claim 3 and B is the bound on Sender's payoff function v. Then $z_a \in \mathbb{R}^S_+ \setminus \{0\}$, $\text{BR}(z_a) = \{a\}$, and $\sum_{s,a} v(s, a)z_a(s) > \text{CP} - \varepsilon$. 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, exists a calibrated strategy σ such that $\chi_\sigma = \sum_a t_a \delta_{x_a}$ exists. From Claim 4, the strategy σ satisfies the requirements in Lemma 3.

6.3.2 *Proof of Proposition 1* 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 4; alternatively, 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 $w = \{w_m \in \mathbb{R}^S_+\}_{m\in A\cup\{\diamond\}}$ 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

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

The payoff to Receiver under this profile is given by

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

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 d), and Receiver has a finite set of pure actions. Therefore, the auxiliary game admits a Nash equilibrium. Let $w^* = \{w^*_m \in \mathbb{R}^S_+\}_{m\in A\cup\{\diamond\}}$ 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 5. *In the BNE of the auxiliary game, $d(w^*_a/|w^*_a|, x_a) \le 2B/\lambda$ 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

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

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 6. *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 4, 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

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

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 6. Let λ be sufficiently high such that $a \in \mathrm{BR}(q)$ for every $q \in \Delta(S)$ such that $d(q, x_a) \le 2B/\lambda$. Such a λ exists by the continuity of d and the fact that a is the unique best response to belief x_a. Then, by Claim 5, it follows that for every belief on-path, Receiver best responds to that belief.

6.4 *Proof of Proposition 2*

Lemma 5. *Assume that there exists some $\gamma > 0$ such that $d(q, m) \ge \gamma|q - m|_1$. 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.*

Proof. 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

$$\big(1 - L|p-q|_1\big)q(s) \le \big(1 - L|p-q|_\infty\big)q(s) \le p(s)$$

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

$$p = \big(1 - L|p-q|_1\big) \cdot q + L|p-q|_1 \cdot q'. \tag{13}$$

We now fix a Sender's strategy σ that is not strictly dominated. We need to prove that σ is calibrated. First, let σ' be any calibrated strategy of Sender and let ρ be a strategy of Receiver. From the bound on v, there exists $B > 0$ such that $V_\lambda(\sigma', \rho) = V_0(\sigma', \rho) \ge -B$ and $V_\lambda(\sigma, \rho) \le B - \lambda\kappa(\sigma)$. Since σ is not strictly dominated by σ', we can choose ρ such that $V_\lambda(\sigma', \rho) \le V_\lambda(\sigma, \rho)$, which implies that

$$\kappa(\sigma) \le 2B/\lambda. \tag{14}$$

Let χ_σ given by (5) be the distribution over messages induced by σ and let

$$q = \int m\chi_\sigma(\mathrm{d}m) \tag{15}$$

be the barycenter of χ_σ. It then follows from (6), (15), the convexity of the norm $|\cdot|_1$, and (17) that

$$|p-q|_1 = \left|\int \big(\tau_\sigma(m) - m\big)\chi_\sigma(\mathrm{d}m)\right|_1 \le \int \big|\tau_\sigma(m) - m\big|_1 \chi_\sigma(\mathrm{d}m) \le \frac{1}{\gamma}\kappa(\sigma). \tag{16}$$

Let $\lambda \ge 2BL/\gamma$. It then follows from (16) and (14) that $|p-q|_1 \le 1/L$. Consider now the calibrated strategy σ' that is induced (via part (ii) of Lemma 2) by the distribution

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

where q' is given by (13) 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 (13) and (15), it follows that the barycenter of this distribution is indeed p.)

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

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

$$\begin{aligned} V_0(\sigma, \rho) &= \int \eta\big(\tau_\sigma(m), m\big)\chi_\sigma(\mathrm{d}m), \\ \kappa(\sigma) &\geq \gamma \int \big|\tau_\sigma(m) - m\big|_1 \chi_\sigma(\mathrm{d}m), \end{aligned} \tag{17}$$

and the payoff under σ' satisfies

$$\begin{aligned} V_\lambda\big(\sigma', \rho\big) &= \big(1 - L|p-q|_1\big) \int \eta(m, m)\chi_\sigma(\mathrm{d}m) + L|p-q|_1 \eta\big(q', q'\big) \\ &\geq \int \big(\eta\big(\tau_\sigma(m), m\big) - B\big|\tau_\sigma(m) - m\big|_1\big)\chi_\sigma(\mathrm{d}m) - 2L|p-q|_1 B \\ &\geq V_\lambda(\sigma, \rho) + \left(\lambda - \frac{B}{\gamma}\right)\kappa(\sigma) - 2L|p-q|_1 B \\ &\geq V_\lambda(\sigma, \rho) + \left(\lambda - \frac{B(1+2L)}{\gamma}\right)\kappa(\sigma), \end{aligned} \tag{18}$$

where the first inequality follows from the Lipschitz condition and the bound on η, the second inequality follows from (17), and the third inequality follows from (16). Since σ is not strictly dominated by σ', we choose ρ such that such that $V_\lambda(\sigma', \rho) \leq V_\lambda(\sigma, \rho)$. If $\lambda > B(1+2L)/\gamma$. Then it follows from (18) that $\kappa(\sigma) = 0$, i.e., that σ is calibrated. □

Lemma 5 states that for $\lambda > \bar{\lambda}$, a Sender's strategy σ is not strictly dominated if and only if it is calibrated. Hence, in the first round of elimination, calibrated Sender's strategies remain. 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 $\mathrm{BR}(m)$.

Finally, by Lemma 3, for every Receiver's strategy ρ such that $\mathrm{support}(\rho(\cdot|m)) \subseteq \mathrm{BR}(m)$, Sender believes that he can get at least $\mathrm{CP} - \varepsilon$ 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 2. 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 0. 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

Abeler, Johannes, Daniele Nosenzo, and Collin Raymond (2019), "Preferences for truth-telling." *Econometrica*, 87, 1115–1153. [478]

Balkenborg, Dieter, Josef Hofbauer, and Christoph Kuzmics (2015), "The refined best-response correspondence in normal form games." *International Journal of Game Theory*, 44, 165–193. [499]

Bardhi, Arjada and Yingni Guo (2018), "Modes of persuasion toward unanimous consent." *Theoretical Economics*, 13, 1111–1149. [478]

Battigalli, Pierpaolo (1997), "On rationalizability in extensive games." *Journal of Economic Theory*, 74, 40–61. [494]

Battigalli, Pierpaolo and Marciano Siniscalchi (2002), "Strong belief and forward induction reasoning." *Journal of Economic Theory*, 106, 356–391. [494]

Brandenburger, Adam, Alex Danieli, and Amanda Friedenberg (2021), "The implications of finite-order reasoning." Unpublished paper. Forthcoming in *Theoretical Economics*. [499]

Chevalier, Judith A. and Dina Mayzlin (2006), "The effect of word of mouth on sales: Online book reviews." *Journal of Marketing Research*, 43, 345–354. [481]

Crawford, Vincent P. and Joel Sobel (1982), "Strategic information transmission." *Econometrica*, 50, 1431–1451. [478]

Dawid, A. Philipp (1982), "The well-calibrated Bayesian." *Journal of the American Statistical Association*, 77, 605–610. [480]

Foster, Dean P. and Rakesh V. Vohra (1998), "Asymptotic calibration." *Biometrika*, 85, 379–390. [480, 481]

Fréchette, Guillaume, Alessandro Lizzeri, and Jacopo Perego (2020), "Rules and commitment in communication: An experimental analysis." Unpublished paper, Columbia University. [480]

Friedenberg, Amanda (2019), "Bargaining under strategic uncertainty: The role of second-order optimism." *Econometrica*, 87, 1835–1865. [494]

Frogner, Charlie, Chiyuan Zhang, Hossein Mobahi, Mauricio Araya, and Tomaso A. Poggio (2015), "Learning with a Wasserstein loss." In *NIPS'15: Proceedings of the 28th International Conference on Neural Information Processing Systems*, volume 2, 2053–2061. [486]

Galichon, Alfred (2018), *Optimal Transport Methods in Economics*. Princeton University Press. [486]

Gehlbach, Scott and Konstantin Sonin (2014), "Government control of the media." *Journal of Public Economics*, 118, 163–171. [478]

Green, Jerry R. and Nancy L. Stokey (2007), "A two-person game of information transmission." *Journal of Economic Theory*, 135, 90–104. [478]

Guo, Yingni and Eran Shmaya (2019), "The interval structure of optimal disclosure." *Econometrica*, 87, 653–675. [478]

Henry, Emeric and Marco Ottaviani (2019), "Research and the approval process: The organization of persuasion." *American Economic Review*, 109, 911–955. [478]

Kamenica, Emir and Matthew Gentzkow (2011), "Bayesian persuasion." *American Economic Review*, 101, 2590–2615. [478, 480]

Kartik, Navin (2009), "Strategic communication with lying costs." *Review of Economic Studies*, 76, 1359–1395. [479, 484]

Kartik, Navin, Marco Ottaviani, and Francesco Squintani (2007), "Credulity, lies, and costly talk." *Journal of Economic Theory*, 134, 93–116. [479, 484]

Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. (2017), "Persuasion of a privately informed receiver." *Econometrica*, 85(6), 1949–1964. [478]

Laraki, Rida (2004), "On the regularity of the convexification operator on a compact set." *Journal of Convex Analysis*, 11(1), 209–234. [497]

Lehrer, Ehud (2001), "Any inspection is manipulable." *Econometrica*, 69, 1333–1347. [481]

Lipnowski, Elliot, Doron Ravid, and Denis Shishkin (2019), "Persuasion via weak institutions." Unpublished paper, SSRN 3168103. [480]

Lipnowski, Elliot and Doron Ravid (2020), "Cheap talk with transparent motives." *Econometrica*, 88, 1631–1660. [480]

Luca, Michael and Georgios Zervas (2016), "Fake it till you make it: Reputation, competition, and yelp review fraud." *Management Science*, 62, 3412–3427. [478]

Mayzlin, Dina, Yaniv Dover, and Judith Chevalier (2014), "Promotional reviews: An empirical investigation of online review manipulation." *American Economic Review*, 104, 2421–2455. [478, 481]

Min, Daehong (2020), "Bayesian persuasion under partial commitment." Unpublished paper, Division of Social Science, New York University Abu Dhabi. [480]

Minaudier, Clement (2020), "The value of confidential policy information: Persuasion, transparency, and influence." Unpublished paper, University of Vienna. [478]

Murphy, Allan H. and Robert L. Winkler (1987), "A general framework for forecast verification." *Monthly Weather Review*, 115, 1330–1338. [480]

Nguyen, Anh and Teck Yong Tan (2021), "Bayesian persuasion with costly messages." Unpublished paper. Forthcoming in *Journal of Economic Theory*. [480]

Olszewski, Wojciech (2015), "Calibration and expert testing," 949–984. [481]

Pearce, David G. (1984), "Rationalizable strategic behavior and the problem of perfection." *Econometrica*, 52, 1029–1050. [494]

Perez-Richet, Eduardo and Vasiliki Skreta (2021), "Test design under falsification." Working paper, CEPR Discussion Paper 15627. [478, 480]

Ranjan, Roopesh and Tilmann Gneiting (2010), "Combining probability forecasts." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72, 71–91. [480]

Rayo, Luis and Ilya Segal (2010), "Optimal information disclosure." *Journal of Political Economy*, 118, 949–987. [480]

Sobel, Joel (2020), "Lying and deception in games." *Journal of Political Economy*, 128, 907–947. [480]

Zamir, Shmuel (1992), "Repeated games of incomplete information: Zero-sum," volume 1, 109–154. [496]

Co-editor Simon Board handled this manuscript.

Manuscript received 18 October, 2019; final version accepted 7 July, 2020; available online 10 August, 2020.