

Costly Miscalibration*

Yingni Guo[†] Eran Shmaya[‡]

July 28, 2018

Abstract

We consider a platform which provides probabilistic forecasts using some algorithm. We introduce a concept of miscalibration which measures the discrepancy between the forecast and the truth. We apply this concept to sender-receiver games in which miscalibration is costly for the sender (the platform). When the sender's miscalibration cost is sufficiently high, he can achieve his commitment solution in an equilibrium. Moreover, under some assumption about the miscalibration-cost function, the only extensive-form rationalizable strategy of the sender is his strategy in the commitment solution. For a subclass of sender-receiver games, we characterize the sender's optimal equilibrium for all miscalibration-cost levels.

JEL: D81, D82, D83

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

1 Introduction

E-commerce platforms often provide information to customers about their products. For example, the fare aggregator Kayak.com provides forecasts of future prices. The real estate

*We thank Arjada Bardhi, Amanda Friedenberg, Navin Kartik, Christoph Kuzmics, Elliot Lipnowski, Kota Murayama, Ludvig Sinander, Marciano Siniscalchi, Egor Starkov, Muhamet Yildiz, Andy Zapechelnjuk, and seminar and conference participants at OSU, Northwestern, Canadian Economic Theory Conference (2018), and Cowles Economic Theory Conference (2018) for helpful comments and discussions. We thank Cassiano Alves and Samuel Goldberg for excellent research assistance. 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.

aggregator Redfin.com identifies “hot homes” that are likely to sell quickly. The platform generates information using some algorithm, and applies the algorithm to many products. This algorithm is usually a trade secret and is not published. For example, Kayak only states 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 present a model for the platform’s communication with customers, in which customers trust the platform’s forecasts in an equilibrium even though they don’t observe the algorithm.

In our model, the platform provides information in the form of probabilistic forecasts. For example, Redfin defines a hot home as one that has a 70% chance or higher of having an accepted offer within two weeks of its debut. We say that the algorithm is *miscalibrated* if there is a discrepancy between the forecast and the realized data, for example, if only 50% of the 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, called the cost intensity, which 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 probabilistic assertions.

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

We illustrate our model with a class of sender-receiver games in which Sender promotes a product and Receiver decides whether to buy it. Many applications studied in recent papers can be viewed as such promotion games. We show that there is a Sender’s optimal equilibrium in which his strategy is calibrated and has a support of two messages, one of which induces Receiver to buy with positive probability. Using this characterization, we study how cost intensities affect Sender’s optimal equilibrium. Sender uses messages that are distant for small cost intensities. As the cost intensity increases, Sender’s strategy shifts gradually to what he would use if Receiver observed Sender’s strategy (i.e., if Sender could

commit to a strategy). Sender's optimal equilibrium payoff is monotone-increasing in the cost intensity.

Moving to general sender-receiver games, we show that for sufficiently high cost intensity, Sender's optimal equilibrium payoff is the payoff he could get if Receiver observed Sender's strategy. In addition, although other equilibrium strategies of Sender exist even for high cost intensity, we show that the only extensive-form rationalizable strategy of Sender is his strategy in the commitment solution.

The equilibrium result is intuitive, as explained by 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 hemi-continuity of the equilibrium correspondence. As usual, lower hemi-continuity is not straightforward and, in our setup, requires either some assumption on the distance function we use to measure miscalibration, or a genericity assumption on the game. The difficulty 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.

To summarize, our results show how, when communication with customers on a large scale is taken together with the available data, online platforms can achieve partial commitment despite the nonobservability of their algorithms.

Our paper is related to the literature on strategic communication with a lying cost (e.g., Kartik, Ottaviani and Squintani (2007), Kartik (2009)). In these papers, the message space is usually the state space, so Sender declares a state. As the lying cost increases, the equilibrium becomes arbitrarily close to fully revealing. Our model delivers a different asymptotic result: there is an equilibrium which is arbitrarily close to the commitment solution. Thus, in our model, when lying is costly, Sender gains the commitment power not to tell a lie, but he is not forced to reveal all his information.

Our results also connect our model to the Bayesian persuasion literature which studies sender-receiver games with commitment on Sender's side (e.g., Rayo and Segal (2010), Kamenica and Gentzkow (2011), Ely (2017)). In that literature, Sender can commit to any strategy. This is as if Receiver observes Sender's strategy. In our model, however, Receiver

does not observe Sender’s strategy. Our model bridges the cheap-talk and persuasion models, and can be used to analyze the middle ground where the talk is not completely cheap nor the commitment absolute. Viewed from this aspect, our paper is related to Fréchette, Lizzeri and Perego (2017), Min (2017) and Lipnowski, Ravid and Shishkin (2018).

Our paper is also related to Nguyen and Tan (2018). Their sender commits to an experiment that is publicly observable. Then, he privately observes the message generated, and can manipulate the message at a cost. They characterize how the chance of manipulation affects the information design. Our model differs in the sense that our sender’s entire strategy is not publicly observable. Receiver only observes the message.

2 An example

Consider a platform (Sender) promoting a product to a customer (Receiver). The product’s quality is $s \in S = \{0, 1, 3\}$ with the prior $p = (\frac{1}{8}, \frac{1}{2}, \frac{3}{8})$. Receiver decides whether or not to buy the product. Receiver’s payoff is $(s - 2)$ 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 $\Delta(S)$, the set of distributions 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 refer to $\tau(m)$ as the true conditional distribution given m . (The true conditional distribution given a message depends on Sender’s strategy. We omit this dependence in our notation $\tau(m)$ in this section.)

Consider Sender’s strategy in table 1a. He sends either message m_0 or m_1 . Each message is a distribution over states. Message m_0 says that the state is 1. Message m_1 says that the distribution over states is $(\frac{1}{5}, \frac{1}{5}, \frac{3}{5})$. Each column shows the probabilities with which Sender sends m_0 or m_1 in each state.

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

Table 1a: A calibrated strategy

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

Table 1b: 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 $\tau(m_1) = (\frac{1}{5}, \frac{1}{5}, \frac{3}{5})$. Note that 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 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 $0, 1, 3$, respectively.

Consider another strategy in table 1b. Sender uses the same messages m_0, m_1 as before. When the state is 0 or 3 , he still announces m_1 for sure. When the state is 1 , he announces m_0 and m_1 with equal probability. The true conditional distribution given m_0 is $\tau(m_0) = (0, 1, 0)$, which still coincides with m_0 . The true conditional distribution given m_1 is $\tau(m_1) = (\frac{1}{6}, \frac{1}{3}, \frac{1}{2})$, which differs from m_1 . This strategy is not calibrated, as illustrated in figure 2.

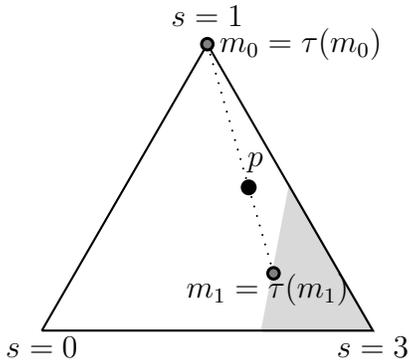


Figure 1: Strategy in table 1a

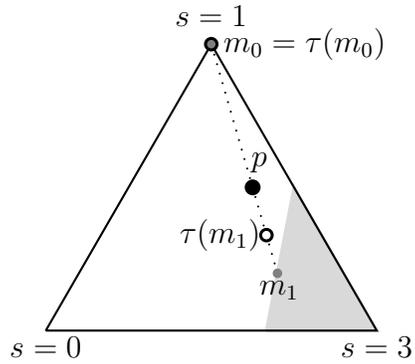


Figure 2: Strategy in table 1b

If Sender uses this strategy to make recommendations, then the probability distribution of the products about which Sender announces m_1 is in fact $\tau(m_1)$. Therefore, what Sender claims about these products—that they are $(\frac{1}{5}, \frac{1}{5}, \frac{3}{5})$ —is not true. If an outside group collects data and performs some statistical test, Sender might be caught.

Whenever Sender sends a message whose true conditional distribution differs from its asserted distribution, Sender pays some cost, which we call a *miscalibration cost*. The bigger the difference is, the higher the cost will be. In this section, we use the total-variation (TV) distance, $\frac{1}{2}|m - \tau(m)|_1$, to measure how much a message m differs from the true conditional distribution $\tau(m)$ given this message. For the strategy in figure 2, the TV 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:

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

Here, $\lambda \geq 0$ 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. 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 $\lambda > 0$, showing that (i) even for small λ , Sender gets a positive payoff; and (ii) for high λ , Sender gets his commitment payoff. The derivation relies on Proposition 3.1 in section 3.3.

For $\lambda \in (0, \frac{5}{4}]$, Sender plays the 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 $\frac{1}{2}\lambda$. He also has no incentive to deviate. To illustrate, consider Sender's strategy in figure 2, which induces Receiver to buy more frequently by announcing m_1 more often. Sender's payoff from this deviation strategy is the probability that Receiver buys minus the miscalibration cost:

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

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

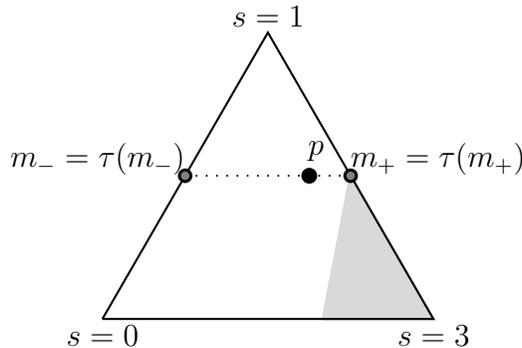


Figure 3: Sender's optimal equilibrium when $\lambda \geq 2$

For $\lambda \geq 2$, Sender's optimal equilibrium is illustrated by figure 3. Sender uses a calibrated strategy, splitting the prior to messages $m_- = (\frac{1}{2}, \frac{1}{2}, 0)$ and $m_+ = (0, \frac{1}{2}, \frac{1}{2})$. Receiver buys

for sure after m_+ , and does not buy otherwise. Using similar arguments, one can show that neither player benefits from deviating. Sender’s equilibrium payoff is $\frac{3}{4}$, which is the same as if he could commit to a strategy.

For $\lambda \in (\frac{5}{4}, 2)$, Sender’s optimal equilibrium strategy shifts gradually from the strategy in figure 1 to that in figure 3. Sender’s optimal equilibrium payoff increases in λ for $\lambda > 0$, which can be interpreted as the proxy for Sender’s commitment power.

Interestingly, for small λ , 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 small, Sender uses distant messages to complement his small “commitment power,” in order to deter himself from deviating.

3 Environment and main results

Let S be a finite set of *states* equipped with a prior distribution p with full support, and M a Borel space of *messages*. A *Sender’s strategy with message space M* is given by a Markov kernel σ from S to M : when the state is s , Sender randomizes a message from $\sigma(\cdot|s)$. For a Sender’s strategy σ , let $\tau_\sigma : M \rightarrow \Delta(S)$ be such that $\tau_\sigma(m)$ is the conditional distribution over states given m . We sometimes use $\tau_\sigma(s|m)$ 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. Sections 3.3 to 3.5 present our definition of an equilibrium with costly miscalibration and our equilibrium results. Section 3.6 discusses rationalizable communication.

3.1 Sender’s calibrated strategies and miscalibration

From now on we assume that $M = \Delta(S)$, so that a message m is an asserted distribution over states. We thus refer to $\tau_\sigma(m)$ as the true conditional distribution over states given a message m . We say that a strategy σ is *calibrated* if $\tau_\sigma(m) = m$, a.s..¹ Under a calibrated strategy, messages mean what they say, i.e., they can reliably be taken at face value.

The term *calibrated* is used in the forecasting literature (e.g., Dawid (1982), Murphy and Winkler (1987), Foster and Vohra (1998), Ranjan and Gneiting (2010)) for two closely related ideas. The first is the purely probabilistic sense, as in our definition of a calibrated strategy of Sender; the second is the idea of calibration with the data, namely, that a probabilistic forecast (or a message, in our terminology) matches the realized distribution of states at those times in which the forecast was given.²

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

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

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

Example 1. The Redfin example:

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

Remark 1. By allowing Sender to choose any strategy $\sigma : S \rightarrow \Delta(S)$, we implicitly assume that Sender knows the state. This assumption is natural since our main focus is on Sender’s

¹The 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.)

²There is also literature about the strategic manipulation of calibration tests. See for instance Lehrer (2001) and Olszewski (2015) for a survey.

³We examine Cook County in Illinois, which includes 165 zip codes, and collect homes that Redfin identifies as hot homes over two weeks (03/14/2018—03/27/2018).

incentives. However, because we use terminology from the forecasting literature and because of our interests 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 $\sigma^E : S \rightarrow \Delta(T)$, 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).

3.2 Sender-receiver games

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

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

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:

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

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 $\lambda \geq 0$.

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 .

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:

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

where $d(\sigma)$ is the measure of miscalibration, and $\lambda \geq 0$ 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 .

Before presenting the results for general sender-receiver games, we focus on a subclass of 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 provides 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. The following proposition enables us to characterize a Sender’s optimal equilibrium for any promotion game when κ is the TV distance.⁴

Proposition 3.1. *Let $\lambda > 0$ and κ be the TV 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’s strategy is such that 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 $\Delta(S)$.*

In the cheap-talk case (i.e., $\lambda = 0$), it is well known that for every 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., he uses a calibrated strategy). When $\lambda > 0$, this argument no longer holds. In fact, we show in section 3.5 that when κ is the Kullback-Leibler (KL) distance, Sender’s optimal equilibrium is not calibrated. Proposition 3.1 says that when κ is the TV distance, Sender can achieve his optimal payoff in a calibrated equilibrium. We emphasize that there are other Sender’s optimal equilibria in which Sender uses noncalibrated strategies. These equilibria typically generate different distributions over states and actions than the one in Proposition 3.1. (See Example 2 below.) This observation, as well as the example with the KL distance, shows that the fact that there is a calibrated Sender’s optimal equilibrium is not a simple revelation-principle argument as in the cheap-talk case.

⁴The proposition holds for other distance functions as well. In the proof in the appendix we describe the property of the distance function that is needed.

Since Sender's optimal payoff can be generated by a calibrated equilibrium, his optimal payoff is monotone-increasing in λ .⁵

Example 2. Consider a promotion game. Let S be $\{-1, 1\}$ with the prior $p = (\frac{3}{4}, \frac{1}{4})$, and $u(s, B) = s$. Assume that the cost intensity is $\lambda = 1$. 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 Proposition 3.1. He uses a calibrated strategy and splits the prior $\frac{1}{4}$ into 0 and $\frac{1}{2}$. Receiver buys after message $\frac{1}{2}$ with probability $\frac{\lambda}{2}$, which ensures that Sender doesn't deviate. Consider another equilibrium. Receiver buys after message 1 with probability λ and doesn't buy otherwise. Sender splits the prior into 0 and $\frac{1}{2}$. When $\frac{1}{2}$ realizes, Sender announces message 1. When 0 realizes, Sender announces message 0. The two equilibria generate the same payoff, $\frac{1}{4}$, to Sender but different distributions over states and actions. Conditional on state 1, Receivers buys more often under the second than the first equilibrium. In fact, it is easy to verify that no equilibrium with a Sender's calibrated strategy generates the same distribution over states and actions as the second equilibrium.

3.4 High cost intensity

We now turn to general sender-receiver games. Our goal is to show that when the cost intensity is high, Sender can achieve his commitment payoff in an equilibrium.

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*. The commitment solution dispenses with Sender's equilibrium constraints that appear in a cheap-talk equilibrium. Thus, CP is the highest payoff that Sender could achieve if Receiver could observe Sender's strategy.

Proposition 3.2. *Sender's payoff under any BNE is at most CP.*

Proof. Let (σ^*, ρ^*) be a BNE. Then

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

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

We first make the additional assumption that the distance function $\kappa(q, m)$ has a kink at $q = m$. This assumption holds for the TV distance $\kappa(q, m) = \frac{1}{2}|q - m|_1$, which is arguably

⁵If (σ, ρ) is an equilibrium for $\lambda > 0$ and σ is calibrated, then it is an equilibrium for a higher $\lambda' > \lambda$.

the most natural distance function between distributions over a finite set of states with no structure. It also holds, for example, for the Euclidean distance $\kappa(q, m) = |q - m|_2$ and the Wasserstein distance (see appendix 4.2.2).

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

To prove Theorem 3.1, we first show that for any distance function with a kink, when λ exceeds some constant, Sender's best response to any receiver's strategy is calibrated. Hence, 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 TV distance. The commitment solution was given by Sender's strategy in figure 3; Receiver buys the product after m_+ and does not buy otherwise. This strategy profile is a BNE when $\lambda \geq 2$.

When we now consider arbitrary distance functions, the assertion of Theorem 3.1 does not hold. However, our next theorem 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.

Theorem 3.2. *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 a BNE in the game with intensity λ such that Sender's payoff is at least $CP - \varepsilon$.*

In section 3.5, we give an example that demonstrates Theorem 3.2 for the case of KL distance. Example 3 in the appendix is a nongeneric game in which Sender's optimal equilibrium payoff, for all cost intensities, is bounded away from his commitment payoff.

3.5 An example with KL distance

We now consider the same example from section 2 but with the prior $p = (\frac{1}{4}, \frac{1}{4}, \frac{1}{2})$ and the KL distance given by

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

Our goal is to show that Sender's optimal equilibrium is not necessarily calibrated and that Sender's optimal equilibrium payoff approaches his commitment payoff as λ goes to infinity.

In the commitment solution, Sender uses a calibrated strategy with support $m_0 = (1, 0, 0)$ and $m_1 = (\frac{1}{7}, \frac{2}{7}, \frac{4}{7})$, and Receiver buys after m_1 . However, for any intensity λ , this strategy

profile is not an equilibrium: Sender has an incentive to deviate by announcing m_1 more often when the state is zero, because when a message has full support a small amount of miscalibration has no first-order impact on Sender's payoff.⁶

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

$$m'_1 = \left(\frac{e^{-1/\lambda}}{7}, \frac{1}{3} - \frac{e^{-1/\lambda}}{21}, \frac{2}{3} - \frac{2e^{-1/\lambda}}{21} \right).$$

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

$$\tau_\sigma(m'_1) = \left(\frac{1}{7}, \frac{2}{7}, \frac{4}{7} \right).$$

Receiver is indeed willing to buy after m'_1 .

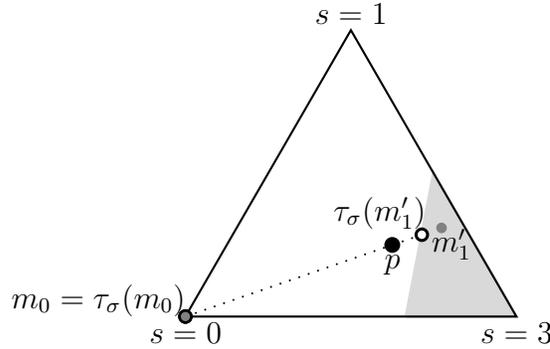


Figure 4: With KL distance, Sender overshoots in equilibrium

Figure 4 illustrates this equilibrium. Sender's strategy is not calibrated; he "overshoots" by announcing m'_1 , which is in the interior of the light-gray area, yet the true conditional distribution $\tau_\sigma(m'_1)$ is just enough for Receiver to buy. Sender's payoff is $\frac{3}{4}\lambda \log \left(\frac{7e^{1/\lambda}}{6} - \frac{1}{6} \right)$, which converges to his commitment payoff as λ goes to infinity.

⁶In the commitment solution, Sender announces m_1 with probability $\frac{1}{2}$ in state 0, and with probability 1 in states 1 and 3. If Sender increases the probability of m_1 in state 0 to $\frac{1}{2} + \alpha$, then his payoff is

$$\frac{7 + 2\alpha}{8} \left(1 - \lambda \frac{6 \log \left(\frac{7}{2\alpha+7} \right) + (2\alpha + 1) \log \left(\frac{7(2\alpha+1)}{2\alpha+7} \right)}{2\alpha + 7} \right).$$

The derivative of this payoff w.r.t. α at $\alpha = 0$ is $\frac{1}{4}$. Hence, Sender has an incentive to deviate for all $\lambda > 0$.

3.6 Rationalizable communication

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 , the “babbling equilibrium” exists. Under this equilibrium, Sender always declares the prior. After every possible Sender’s message, Receiver believes that the state is distributed according to the prior, and best responds to the prior.

There is, however, an unsatisfactory aspect to the babbling equilibrium in our context. Consider the example from section 2 with a high λ . Assume that Receiver plays according to the babbling equilibrium, and that the message $(0, 0, 1)$ —which says that “the state is 3 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 does not capture this intuition because it allows arbitrary behaviors off path. These behaviors do not have to be rationalizable. In order to incorporate this intuition, we turn to the concept of 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 2 for a detailed discussion, as well as Friedenber (2017), where similar issues arise.)

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

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

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

- (i) Σ_0 is the set of all Sender’s strategies.

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

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

Theorem 3.3. *Assume that there exists some $\gamma > 0$ such that $\kappa(q, m) \geq \gamma|q - m|_1$. Then, in a generic set of games, there exists $\bar{\lambda}$ such that for every $\lambda > \bar{\lambda}$, Sender's rationalizable strategies are his calibrated commitment strategies, i.e., calibrated strategies σ such that $V_\lambda(\sigma, \rho) = CP$ for some best response ρ of Receiver to σ .*

When we apply Theorem 3.3 to the example from section 2, 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 $BR(m)$.

Remark 2. The idea is that if a message m is on the path of some strategy of Sender that has not yet been eliminated, then Receiver's rationalizable actions for this message are the best responses against the conditional beliefs over states under such strategies of Sender. If the message m is not on the path of any of Sender's strategies that have not yet been eliminated, then the message is a surprise and all Receiver's actions which have not yet been eliminated remain. Thus, EFR requires, for each message m and each Sender's strategy σ , (i) defining whether m is on the path of σ , and (ii) defining the conditional belief $\tau_\sigma(m)$ if the message is on path.

However, when the set of messages is a continuum, it is possible that every message m is produced with probability zero. In this case, it is not obvious what it means for a message to be on path, and the conditional distributions $\tau_\sigma(m)$ are defined only almost-surely in m . More formally, there is an equivalence class of functions, called *versions* of the conditional distributions, and these versions differ from one another in a set of messages m with probability zero. When there is only one σ on the table, as is the case in the definition of a BNE, this is not a big issue since the definition does not depend on the version of the conditional distribution. But in the definition of rationalizability, at each stage of elimination

there is a set of Sender's strategies, and Receiver needs to choose versions of the conditional distributions simultaneously for all these strategies. The resulting set of rationalizable actions for Receiver will depend on which version is chosen. Our approach is to require Receiver to form a belief only when he receives a message that is produced with a strictly positive probability for some remaining strategy σ of Sender. For such a message m —a message that practically shouts “I am not a surprise”—the conditional distribution $\tau_\sigma(m)$ is unambiguously defined. As Theorem 3.3 shows, even this weak definition of EFR narrows down significantly the possible outcomes in our environment.

4 Appendix

4.1 Preliminaries

We extend the distance function κ to a function $\kappa : \mathbf{R}_+^S \times \Delta(S) \rightarrow \mathbf{R}_+$ given by $\kappa(\gamma q, m) = \gamma \kappa(q, m)$ for every $q \in \Delta(S)$ and $\gamma \geq 0$. We extend the best-response correspondence $BR(\cdot)$ to a correspondence from \mathbf{R}_+^S to A given by the same formula (3).

For a Sender's strategy σ we let

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

be the distribution over messages induced by σ . Lemma 4.1 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 4.1. (i) *For every strategy σ , we have*

$$\int \tau_\sigma(m) \chi_\sigma(dm) = p. \tag{5}$$

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

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 a finite splitting there exists a calibrated strategy that announces the message q_i with probability t_i .

4.2 Proof of Proposition 3.1

4.2.1 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 $\rho : \Delta(S) \rightarrow [0, 1]$, so that $\rho(m)$ is the probability that Receiver buys after message m .

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

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

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

For a bounded function $f : \Delta(S) \rightarrow \mathbf{R}$, we denote by $\text{Cav } f$ 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 $f, g : \Delta(S) \rightarrow \mathbf{R}$, we denote $f \leq g$ when $f(q) \leq g(q)$ for every $q \in \Delta(S)$.

Proposition 4.1. *Let $f, g : \Delta(S) \rightarrow \mathbf{R}$ be bounded. Then,*

- (i) *if κ satisfies the triangular inequality, then $\text{Lip } \text{Lip } f = \text{Lip } f$;*
- (ii) *$\text{Cav } \text{Cav } f = \text{Cav } f$;*
- (iii) *if $f \leq g$, then $\text{Lip } f \leq \text{Lip } g$ and $\text{Cav } f \leq \text{Cav } g$.*

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

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

$$V_1(\sigma, \rho) = \int (\rho(m) - \kappa(\tau_\sigma(m), m)) \chi_\sigma(dm) \leq \int \text{Lip } \rho(\tau_\sigma(m)) \chi_\sigma(dm) \leq \text{Cav } \text{Lip } \rho(p),$$

where the first inequality follows from the definition of $\text{Lip } \rho$ and $\rho(m) - \kappa(\tau_\sigma(m), m) \leq \text{Lip } \rho(\tau_\sigma(m))$, and the second inequality follows from (5) 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 $t_i \geq 0$, and messages m_i , such that $p = \sum_i t_i q_i$, $\sum_i t_i = 1$, and

$$\sum_i t_i (\rho(m_i) - \kappa(q_i, m_i)) \geq \text{Cav } \text{Lip } \rho(p) - \varepsilon. \quad (6)$$

We can assume that $m_i \neq m_j$ if $i \neq j$. Otherwise, if $m_i = m_j = \bar{m}$, we let $\bar{t} = t_i + t_j$ and $\bar{q} = (t_i q_i + t_j q_j) / \bar{t}$. Since κ is convex in q , we can replace the elements q_i and q_j with a single element \bar{q} , whose weight and corresponding message are \bar{t} and \bar{m} , without violating (6).

We now consider Sender's strategy σ such that $\chi_\sigma = \sum_i t_i \delta_{m_i}$ and $\tau_\sigma : \Delta(S) \rightarrow \Delta(S)$ is a function such that $\tau_\sigma(m_i) = q_i$. Such a strategy exists by Lemma 4.1. Then (6) implies that Sender's payoff when using σ is at least $\text{Cav Lip } \rho(p) - \varepsilon$. \blacksquare

4.2.2 Wasserstein distance

We will prove Proposition 3.1 under a more general assumption that κ is a Wasserstein distance. Let $\bar{\kappa}$ be a metric on S . Then the Wasserstein distance κ over $\Delta(S)$ induced by $\bar{\kappa}$ is such that

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

where the infimum ranges over all pairs Q, Q' of S -valued random variables with marginal distributions q, q' respectively. (Such a pair is called a *coupling* of q, q' .) The TV distance is a special case when $\bar{\kappa}$ is the discrete metric so that $\bar{\kappa}(s, s') = 1$ when $s \neq s'$. For example, in promotion game, it is natural to choose $\bar{\kappa}(s, s') = |u(s, B) - u(s', B)|$, in which case saying that the state is 3 when it is 0 is more costly than saying that the state is 1.

The Wasserstein distance κ is a convex function of q, q' . It is a metric, so satisfies the triangular inequality. The following lemma shows that Wasserstein distance gets along well with splitting of probability distributions. For related results on other metrics and generalizations to infinite-dimensional spaces, including implications for Lipschitz continuity as in our Corollary 4.3 see Laraki (2004).

Lemma 4.2. *Let κ be a Wasserstein distance over $\Delta(S)$. Let $q, q' \in \Delta(S)$ and let $q = t_1 q_1 + \dots + t_k q_k$ be a splitting of q . Then there exists a splitting $q' = t_1 q'_1 + \dots + t_k q'_k$ of q' with the same weights such that:*

$$\kappa(q, q') = t_1 \kappa(q_1, q'_1) + \dots + t_k \kappa(q_k, q'_k). \quad (7)$$

Proof. The direction $\kappa(q, q') \leq t_1 \kappa(q_1, q'_1) + \dots + t_k \kappa(q_k, q'_k)$ follows from convexity of κ for every splitting $q' = t_1 q'_1 + \dots + t_k q'_k$ of q' with the weights t_1, \dots, t_k .

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

for every $i \in \{1, \dots, k\}$ and $s \in S$. The existence of such a variable follows from the splitting lemma. Finally, let $q'_i(s) = \mathbb{P}(Q' = s | X = i)$. Then

$$\sum_i t_i \kappa(q_i, q'_i) \leq \sum_i \mathbb{P}(X = i) \mathbb{E} \bar{\kappa}(Q, Q' | X = i) = \mathbb{E} \bar{\kappa}(Q, Q') = \kappa(q, q'),$$

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

Proof. Let $q, q' \in \Delta(S)$ and let $q = t_1 q_1 + \dots + t_k q_k$ be a splitting of q such that $\text{Cav } f(q) = t_1 f(q_1) + \dots + t_k f(q_k)$. By Lemma 4.2 there exists a splitting $q' = t_1 q'_1 + \dots + t_k q'_k$ of q' such that (7) holds. Therefore

$$\text{Cav } f(q') \geq \sum_i t_i f(q'_i) \geq \sum_i t_i (f(q_i) - \kappa(q_i, q'_i)) = \text{Cav } f(q) - \kappa(q, q'),$$

as desired. ■

4.2.3 Proof of Proposition 3.1

Let (σ, ρ) be an equilibrium. From the equilibrium condition for Sender and Claim 1, it follows that

$$\text{Cav Lip } \rho(p) = \int (\rho(m) - \kappa(\tau_\sigma(m), m)) \chi_\sigma(dm). \quad (8)$$

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

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

$$\chi_\sigma = (1 - t)\underline{\chi}_\sigma + t\bar{\chi}_\sigma, \quad \text{and} \quad (9)$$

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

From the equilibrium condition for Receiver it follows that $\rho = 0$ $\underline{\chi}_\sigma$ -almost surely. There-

fore, from (8) and (9) it follows that

$$\text{Cav Lip } \rho(p) \leq t \int (\rho(m) - \kappa(\tau_\sigma(m), m)) \bar{\chi}_\sigma(dm) \quad (11)$$

It follows from the definition of \bar{q} and the convexity of \mathcal{A} that $\bar{q} \in \mathcal{A}$. Since $p \notin \mathcal{A}$ there exists a belief q^* on the interval $[p, \bar{q}]$ which is on the boundary of \mathcal{A} , i.e., $\sum_s q^*(s)u(s, B) = 0$. From (10)

$$q^* = (1 - t^*)\underline{q} + t^*\bar{q} = (1 - t^*)\underline{q} + t^* \int \tau_\sigma d\bar{\chi}_\sigma \quad (12)$$

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 $\rho^*(q^*) = \text{Cav Lip } \rho(q^*)$ 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 $\text{Cav Lip } \rho(p)$ 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

$$\begin{aligned} V_1(\sigma^*, \rho^*) &= \frac{t}{t^*}\rho^*(q^*) = \frac{t}{t^*} \text{Cav Lip } \rho(q^*) \geq \\ &t \int \text{Lip } \rho(\tau_\sigma(m)) \bar{\chi}_\sigma(dm) \geq t \int (\rho(m) - \kappa(\tau_\sigma(m), m)) \bar{\chi}_\sigma(dm) \geq \text{Cav Lip } \rho(p), \end{aligned}$$

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

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

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

The first inequality and the last equality follow from Proposition 4.1 (iii) and (ii), respectively. For the second equality, according to Corollary 4.3, $\text{Cav Lip } \rho(p)$ is 1-Lipschitz w.r.t. κ . Therefore, $\text{Lip Cav Lip } \rho(p) = \text{Cav Lip } \rho(p)$.

By Claim 1 this implies that Sender's optimal payoff against ρ^* is at most $\text{Cav Lip } \rho(p)$, as desired.

4.3 Proof of Theorem 3.1

It is a standard revelation-principle argument that the commitment solution can be implemented via a calibrated strategy of Sender, with Receiver believing Sender's message and playing Sender's preferred action among Receiver's best responses to that message. We have to show that playing this profile is an equilibrium. Since Sender's strategy is calibrated, it is clear that Receiver best responds to Sender. What is difficult, however, is proving that Sender will not deviate from this strategy. The key observation is Lemma 4.4 below, which asserts that for a sufficiently high intensity, Sender is always better off using a calibrated strategy, regardless of Receiver's strategy.

Lemma 4.4. *Assume that there exists some $\gamma > 0$ such that $\kappa(q, m) \geq \gamma|q - m|_1$. There exists $\bar{\lambda}$ such that for every $\lambda > \bar{\lambda}$, if σ is a Sender's strategy which is not calibrated and ρ is a Receiver's strategy, then there exists a calibrated strategy of Sender that gives him a higher payoff than σ against ρ .*

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 TV distance between these distributions. To summarize, the total payoff

loss is at most some constant times the measure of miscalibration. If λ is high enough, this payoff loss is less than the miscalibration cost that is saved.

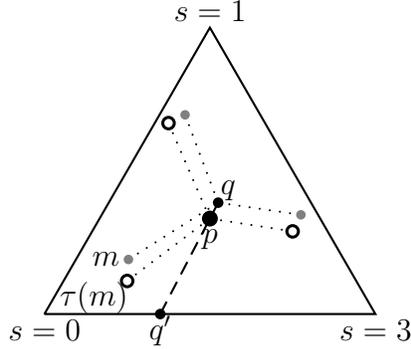


Figure 5: Proof idea for Lemma 4.4

Proof of Lemma 4.4. Let $L = 1/\min\{p(s) : s \in S\}$. (Here we use the full-support assumption on p .) The choice of L is such that, for every $q \in \Delta(S)$, p can be split into a convex combination of q and some other belief $q' \in \Delta(S)$, where L bounds the weight on q' . More explicitly, for every belief q , we have

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

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

$$p = (1 - L|p - q|_1) \cdot q + L|p - q|_1 \cdot q'. \quad (13)$$

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

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

$$\begin{aligned} V_0(\sigma, \rho) &= \int \eta(\tau_\sigma(m), m) \chi_\sigma(dm), \text{ and} \\ d(\sigma) &\geq \gamma \int |\tau_\sigma(m) - m|_1 \chi_\sigma(dm). \end{aligned} \quad (14)$$

We now fix a Sender's strategy σ , with the aim of proving that if σ is miscalibrated, then

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

$$d(\sigma) \leq 2B/\lambda; \quad (15)$$

otherwise, σ cannot be a best response.

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

$$q = \int m \chi_\sigma(dm) \quad (16)$$

be the barycenter of χ_σ . It then follows from the convexity of the norm $|\cdot|_1$, (5), (16), and (14) that

$$|p - q|_1 = \left| \int (\tau_\sigma(m) - m) \chi_\sigma(dm) \right|_1 \leq \int |\tau_\sigma(m) - m|_1 \chi_\sigma(dm) \leq \frac{1}{\gamma} d(\sigma). \quad (17)$$

Let $\lambda \geq \frac{2BL}{\gamma}$. It then follows from (17) and (15) that $|p - q|_1 \leq 1/L$. Consider now the calibrated strategy that is induced (via part 2 of Lemma 4.1) 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 (13) and δ is the Dirac measure: with probability $L|p - q|_1$ Sender sends q' , and with the complementary probability $1 - L|p - q|_1$ Sender uses χ_σ . (Note that, from (13) and (16), it follows that the barycenter of this distribution is indeed p .) The payoff under this strategy is given by

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

where the first inequality follows from the Lipschitz condition and the bound on η , the second inequality from (14), and the third inequality from (17). Therefore, if $d(\sigma) > 0$ and $\lambda > B(1 + 2L)/\gamma$, then the calibrated strategy defined above gives Sender a payoff strictly higher than σ does. ■

Proof of Theorem 3.1. 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 4.4 it follows that no miscalibrated strategy gives Sender more than CP against ρ^* . Therefore, σ^* is a best response against ρ^* . ■

4.4 Proof of Lemma 4.5

The first step of the proof of Theorem 3.2 is Lemma 4.5 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 4.5. *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) $\text{support}(\sigma) = \{x_a\}_{a \in A}$; and (iii) if ρ is Receiver's best response to σ , then $V_\lambda(\sigma, \rho) > CP - \varepsilon$.*

We prove Lemma 4.5 with the help of Lemma 4.6 below. The assertion in the lemma is standard. See Balkenborg, Hofbauer and Kuzmics (2015) and Brandenburger, Danieli and Friedenberg (2017) for similar arguments.

Lemma 4.6. *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 4.5, 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 2. *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 4.6, 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 3. *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 2 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 3 the strategy σ satisfies the requirements in Lemma 4.5.

4.5 Proof of Theorem 3.2

We consider the following auxiliary game, in which Sender's set of messages is $A \cup \{\diamond\}$ where \diamond is a message that says "silent." In the auxiliary game, if Sender sends message $a \in A$, then Receiver must play action a and Sender pays a miscalibration cost relative to x_a given in Claim 3; 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 4. *In the BNE of the auxiliary game, $\kappa(w_a^*/|w_a^*|, x_a) \leq 2B/\lambda$ for every $a \in A$ such that $|w_a^*| > 0$. Here, $B > 0$ is the bound on Sender's payoff function v .*

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

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

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

Claim 5. *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 3, Sender's payoff is at least $CP - \varepsilon$. \blacksquare

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

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 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 the safe action a_0 , so Sender's payoff is 0.

4.6 Proof of Theorem 3.3

Lemma 4.4 implies that for $\lambda > \bar{\lambda}$, only calibrated strategies of Sender survive the first round of elimination. Since all calibrated strategies are best responses to Receiver's strategy that always chooses some fixed action $a \in A$, it follows that Σ_1 is the set of all calibrated strategies. Since any message m is an atom of some calibrated strategy of Sender, it follows that $R_2(m) = BR(m)$. Therefore, by Lemma 4.5, for every Receiver's strategy ρ that survived, Sender believes that he can get at least $CP - \varepsilon$ against ρ for every $\varepsilon > 0$. Therefore, a rationalizable strategy of Sender must be calibrated and give him CP against some strategy ρ of Receiver such that $\text{support}(\rho(\cdot|m)) \subseteq BR(m)$.

References

- Aumann, Robert J., and Michael Maschler.** 1995. *Repeated Games with Incomplete Information*. The MIT Press.
- Balkenborg, Dieter, Josef Hofbauer, and Christoph Kuzmics.** 2015. “The Refined Best-response Correspondence in Normal Form Games.” *International Journal of Game Theory*, 44(1): 165–193.
- Battigalli, Pierpaolo.** 1997. “On Rationalizability in Extensive Games.” *Journal of Economic Theory*, 74: 40–61.
- Battigalli, Pierpaolo, and Marciano Siniscalchi.** 2002. “Strong Belief and Forward Induction Reasoning.” *Journal of Economic Theory*, 106: 356–391.
- Brandenburger, Adam, Alex Danieli, and Amanda Friedenberg.** 2017. “How Many Levels Do Players Reason? An Observational Challenge and Solution.” *Working paper*.
- Brandenburger, Adam, Amanda Friedenberg, and H. Jerome Keisler.** 2008. “Admissibility in Games.” *Econometrica*, 76(2): 307–352.
- 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.
- Ely, Jeffrey C.** 2017. “Beeps.” *American Economic Review*, 107(1): 31–53.
- Foster, Dean P., and Rakesh V. Vohra.** 1998. “Asymptotic Calibration.” *Biometrika*, 85(2): 379–390.
- Fréchette, Guillaume, Alessandro Lizzeri, and Jacopo Perego.** 2017. “Rules and Commitment in Communication.” *Working paper*.
- Friedenberg, Amanda.** 2017. “Bargaining Under Strategic Uncertainty: The Role of Second-Order Optimism.” *Working paper*.
- Green, Jerry R., and Nancy L. Stokey.** 2007. “A Two-person Game of Information Transmission.” *Journal of Economic Theory*, 90–104.

- Kamenica, Emir, and Matthew Gentzkow.** 2011. “Bayesian Persuasion.” *American Economic Review*, 101(6): 2590–2615.
- Kartik, Navin.** 2009. “Strategic Communication with Lying Costs.” *Review of Economic Studies*, 1359–1395.
- Kartik, Navin, Marco Ottaviani, and Francesco Squintani.** 2007. “Credulity, Lies, and Costly Talk.” *Journal of Economic Theory*, 134: 93–116.
- Kuzmics, Christoph.** 2015. “Abraham Wald’s Complete Class Theorem and Knightian Uncertainty.” *Browser Download This Paper*.
- 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, Doron Ravid, and Denis Shishkin.** 2018. “Persuasion via Weak Institutions.” *Working paper*.
- Min, Daehong.** 2017. “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.
- 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.