Cheap Talk with Transparent Motives

Elliot Lipnowski Doron Ravid*

December 2018

Abstract

We study a model of cheap talk with one substantive assumption: The sender's preferences are state-independent. Our main observation is that such a sender gains credibility by garbling self-serving information. Using this observation, we examine the possibility of valuable communication, assess the value of commitment, and explicitly solve for sender-optimal equilibria in several examples. A key result is a geometric characterization of the value of cheap talk, described by the *quasi*concave envelope of the sender's value function. (*JEL* D83, D82, M37, D86, D72)

KEYWORDS: cheap talk, belief-based approach, securability, quasiconcave envelope, persuasion, information transmission, information design

1 Introduction

How much can an expert benefit from strategic communication with an uninformed agent? A large literature, starting with Crawford and Sobel (1982) and Green and Stokey (2007), has studied this question, focusing on the case in which the expert's preferences depend on the state. However, many experts have state-independent preferences: Salespeople want to sell products with higher commissions; politicians

*Department of Economics, University of Chicago, lipnowski@uchicago.edu, dravid@uchicago.edu. This research was supported by a grant from the National Science Foundation (SES-1730168). We would like to thank Emir Kamenica for his invaluable comments and suggestions. We would also like to thank Ben Brooks, Kevin Bryan, Odilon Câmara, Archishman Chakraborty, David Dillenberger, Pedro Gardete, Johannes Hörner, Navin Kartik, R. Vijay Krishna, Laurent Mathevet, Stephen Morris, Andrew Postlewaite, Maher Said, Denis Shishkin, and Vasiliki Skreta for insightful discussions.

want to get elected; lawyers want favorable rulings; and so on. This paper analyzes the extent to which such experts benefit from cheap talk.[1]

We consider a general cheap-talk model with one substantive assumption: The sender has state-independent preferences. Thus, we start with a receiver facing a decision problem with incomplete information. The relevant information is available to an informed sender who cares only about the receiver's action. Wanting to influence this action, the sender communicates with the receiver via cheap talk.

Our main insight is a sender with state-independent preferences gains credibility by garbling self-serving information. To derive this insight, we take a belief-based approach, as is common in the literature on communication.[2] Thus, we summarize communication via its induced *information policy*, a distribution over receiver posterior beliefs that averages to the prior. Say that a payoff s is *sender-beneficial* if it is larger than the sender's no-information payoff, and *securable* if the sender's lowest ex-post payoff from some information policy is at least s. Theorem 1 shows a sender-beneficial payoff s can be obtained in equilibrium if and only if s is securable. Thus, although the information policy securing s need not itself arise in equilibrium, its existence is sufficient for the sender to obtain a payoff of s in some equilibrium. Intuitively, the securing policy leads to posteriors that provide too much sender-beneficial information to the receiver. By garbling said information posterior-by-posterior, one can construct an equilibrium information policy attaining the secured value.

To illustrate our main result, consider a political think tank that advises a lawmaker. The lawmaker is contemplating whether to pass one of two possible reforms, denoted by 1 and 2, or to maintain the status quo, denoted by 0. Evaluating each proposal's merits requires expertise, which the think tank possesses. Given the think tank's political leanings, it is known to prefer certain proposals to others. In particular, suppose the status quo is the think tank's least preferred option and the second

[1]Other papers have studied cheap-talk communication between a sender and a receiver when the former has state-independent preferences. Chakraborty and Harbaugh (2010) provide sufficient conditions for the existence of an influential equilibrium in a static sender-receiver game; we discuss that paper at some length below. Schnakenberg (2015) characterizes when an expert can convince voters to implement a proposal, and when this harms the voting population. Margaria and Smolin (2018) prove a folk theorem for a repeated interaction in which both a sender and a receiver are long-lived. With a long-lived sender but short-lived receivers, Best and Quigley (2017) show that only partitional information can be credibly revealed, and that well-chosen mediation protocols can restore the commitment solution for a patient sender.

[2]For example, see Aumann and Maschler (1995), Aumann and Hart (2003), Kamenica and Gentzkow (2011), Alonso and Câmara (2016), and Ely (2017).

reform is the think tank's favorite option. Hence, let $a \in \{0, 1, 2\}$ represent both the lawmaker's choice and the think tank's payoff from that choice. To choose to implement a reform, the lawmaker must be sufficiently confident that the reform is good. Suppose one reform is good and one is bad, where the state, $\theta \in \{1, 2\}$, is the identity of the good reform. The lawmaker implements the more likely good reform whenever he believes it to be good with probability strictly above $\frac{3}{4}$. At $\frac{3}{4}$, the lawmaker is indifferent between said reform and the status quo, which the lawmaker chooses when neither reform is sufficiently likely to be good. Both reforms are equally likely to be good under the prior.

Suppose the think tank could reveal the state to the lawmaker; that is, the think tank recommends that the lawmaker *implement 1* when the state is 1 and *implement 2* when the state is 2. Because following these recommendations is incentive-compatible for the lawmaker, the think tank's ex-post payoff would be 1 when sending *implement 1* and 2 when sending *implement 2*. By contrast, under no information, the think tank's payoff is 0. Thus, revealing the state secures the think tank a payoff of 1, which is higher than its payoff under the prior. Notice that 1 is then the highest payoff that the think tank can secure, because no information policy always increases the probability that the lawmaker assigns to the state being 2. One can therefore apply Theorem 1 to learn two things: (1) 1 is an upper bound on the think tank's equilibrium payoffs, and (2) we can acheive this bound via a message-by-message garbling of said protocol. For (2), consider what happens when the think tank sends the *implement 2* message according to

$$\mathbb{P}\{\text{implement } \mathit{2} | \theta = 1\} = \frac{1}{3},$$
$$\mathbb{P}\{\text{implement } \mathit{2} | \theta = 2\} = 1,$$

and sends *implement 1* with the complementary probabilities. As with perfect state revelation, choosing proposal 1 is the lawmaker's unique best response to *implement 1.* However, given *implement 2,* the lawmaker assigns a probability of $\frac{3}{4}$ to state 2, making him indifferent between implementing 2 and the status quo. Being indifferent, the lawmaker mixes between keeping the status quo and implementing 2 with equal probabilities. Such mixing results in indifference by the think tank, yielding an equilibrium.

In the general model, Theorem 1 allows us to geometrically characterize the sender's maximal benefit from cheap talk and compare this benefit with her benefit under commitment. Kamenica and Gentzkow (2011) characterize the sender's benefit under commitment in terms of her *value function*—that is, the highest value the sender can obtain from the receiver's optimal behavior given his posterior beliefs. Specifically, they show the sender's maximal commitment value is equal to the concave envelope of her value function. As we show in Theorem 2, replacing the concave envelope with the *quasi*concave envelope gives the sender's maximal value under cheap talk. Thus, the value of commitment in persuasion is the difference between the concave and quasiconcave envelopes of the sender's value function.

Figure 1 visualizes the geometric comparison between cheap talk and commitment in the aforementioned think tank example. Because the state is binary, the lawmaker's belief can be summarized by the probability it assigns to the second reform being good ($\theta = 2$). Putting this probability on the horizontal axis, the figure plots the highest value the think tank can obtain from uninformative communication, cheap talk, and commitment. That is, the figure plots the think tank's value function (left), along with its quasiconcave (center) and concave (right) envelopes. The two envelopes describe how communication benefits the think tank by allowing it to connect points on the value function's graph. In contrast to communication with commitment, in which the think tank can connect points using any affine segment, only flat segments are allowed with cheap talk. The restriction to flat segments comes from the think tank's incentive constraints: Because the think tank's preferences are state independent, all equilibrium messages must yield the same payoff. As such, it can only connect points with the same payoff coordinate; that is, only flat segments are feasible.

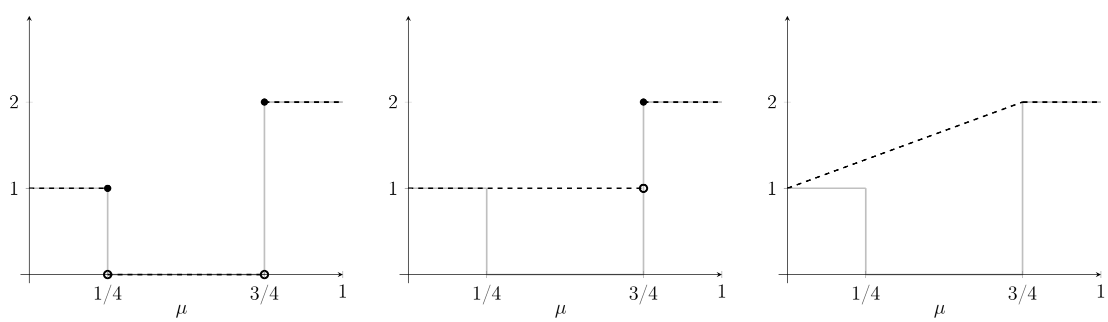

(a) Value function (b) Quasiconcave Envelope (c) Concave Envelope

Figure 1: The simple think tank example. The figure plots the highest value the think tank can obtain from no information (left), cheap talk (center), and commitment (right).

The geometric difference between cheap talk and commitment allows us to show that, in finite settings, most priors fall into one of two categories: Either the sender can get her first-best outcome with cheap talk, or she would strictly benefit from commitment. One can see this categorization holds in the simple think tank example for almost all beliefs by using Figure 1. The figure clearly shows that, unless the second reform is never good, the concave envelope lies above the quasiconcave envelope whenever the probability of the second reform being good is below $\frac{3}{4}$. Whenever the second reform is good with probability $\frac{3}{4}$ or above, the lawmaker is willing to implement the think tank's favorite reform under the prior, and so the two envelopes must coincide with the value function.

In section 5, we demonstrate how one can use our results in more specific settings. We examine three examples. In a richer version of the above think tank example, we show the think tank's best equilibrium involves giving the lawmaker noisy recommendations, where the noise is calibrated to make the lawmaker indifferent between the recommended reform and the status quo. We also study a broker-investor relationship, in which an investor consults his broker about an asset, and the broker earns a fee proportional to the investor's trades. We identify a Pareto-dominant equilibrium in which the broker tells the investor whether his holdings should be above or below a fee-independent cutoff amount. Thus, the lower the broker's fee, the better off the investor, who pays less money for the same information. Lower fees have an ambiguous effect on the broker, because they reduce her income per trade but increase equilibrium trade volume. We also conduct comparative statics in market volatility. Although higher volatility cannot hurt her, the broker strictly benefits from higher volatility only if she can effectively communicate about it to the investor. The investor's attitude toward higher volatility is ambiguous, because it changes both the investor's prior uncertainty and the usefulness of the broker's information. Our third example is a symmetric version of the multiple-goods seller example of Chakraborty and Harbaugh (2010). Specifically, we consider a seller who wants to maximize the probability of selling one of her many products to a buyer. In this setting, we show that the best the seller can do with cheap talk is to tell the buyer the identity of her best product, and that being able to benefit ex ante from providing the buyer with additional information about the best product is a necessary and sufficient condition for the seller to benefit from commitment.

That a sender with state-independent preferences can communicate effectively should come as no surprise to those familiar with Seidmann (1990) and Chakraborty and Harbaugh (2010). Seidmann (1990) shows that effective communication is possible under state independent sender preferences, and can be enabled by either the receiver holding private information or the receiver's action being multidimensional. Multidimensionality also plays a central role in Chakraborty and Harbaugh (2010), who show, in an abstract multidimensional specialization of our model, that the sender can *always* trade off dimensions in order to communicate some information credibly and influence the receiver's actions.[3] In the further specialization of Chakraborty and Harbaugh's (2010) model in which the sender strictly likes to influence the receiver, their result also shows cheap talk can benefit the sender. In contrast to these two papers, which focus on whether effective communication is possible, we focus on quantifying communication's benefits to the sender.

In section 6.1, we revisit Chakraborty and Harbaugh's (2010) argument. We point out that, stripping away the specific parametric structure of Chakraborty and Harbaugh's (2010) setting, their reasoning generates informative communication whenever there are more than two states. In general, despite being informative, this communication may not influence the receiver's actions. Using this observation, one can quickly deduce that a sender with state-independent preferences can always provide information to the receiver unless the state is binary and the sender strictly benefits from increasing the probability the receiver assigns to one state.

We conclude by discussing two ways to modify our model. The first modification generalizes the notion of state-independent preferences. As we point out, our results apply whenever the sender's preferences are state independent over all rationalizable receiver action distributions. The second modification allows for multiple rounds of communication before the receiver takes his action, that is, long cheap talk. We show that long cheap talk cannot benefit the sender but can help the receiver.

2 Cheap Talk with State-Independent Preferences

Our model is an abstract cheap-talk model with the substantive restriction that the sender has state-independent preferences. Thus, we have two players: a sender (S, she) and a receiver (R, he). The game begins with the realization of an unknown

[3] See Battaglini (2002) and Chakraborty and Harbaugh (2007) for applications of this idea in the unknown agenda case.

state, $\theta \in \Theta$, which S observes. After observing the state, S sends R a message, $m \in M$. R then observes m (but not θ) and decides which action, $a \in A$, to take. Whereas R's payoffs depend on θ, S's payoffs do not.

We impose some technical restrictions on our model.[4] Each of Θ, A, and M is a compact metrizable space containing at least two elements, and M is sufficiently rich.[5] The state, θ, follows some full-support distribution $\mu_0 \in \Delta\Theta$, which is known to both players. Both players' utility functions are continuous, where we take $u_S : A \to \mathbb{R}$ to be S's utility[6] and $u_R : A \times \Theta \to \mathbb{R}$ to be R's.

We are interested in studying the game's equilibria*,* by which we mean perfect Bayesian equilibria. An **equilibrium** consists of three measurable maps: a strategy $\sigma : \Theta \to \Delta M$ for S; a strategy $\rho : M \to \Delta A$ for R; and a belief system $\beta : M \to \Delta\Theta$ for R; such that:

1. β is obtained from μ_0, given σ, using Bayes' rule.[7]

2. $\rho(m)$ is supported on $\arg\max_{a \in A} \int u_R(a, \cdot)\ \mathrm{d}\beta(\cdot|m)$ for all $m \in M$.

3. $\sigma(\theta)$ is supported on $\arg\max_{m \in M} \int u_S(\cdot)\ \mathrm{d}\rho(\cdot|m)$ for all $\theta \in \Theta$.

Any triple $\mathcal{E} = (\sigma, \rho, \beta)$ induces a joint distribution, $\mathbb{P}_{\mathcal{E}}$, over realized states, messages, and actions,[8] which, in turn, induces (through β and ρ, respectively) distributions over R's equilibrium beliefs and chosen mixed action.

The following are a few concrete examples of our setting.

Example 1. Consider the following richer version of the think tank example from the introduction. Thus, S is a think tank that is advising a lawmaker (R) on whether to

[4]Let us describe some notational conventions we adopt throughout the paper. For a compact metrizable space Y, we let ΔY denote the set of all Borel probability measures over Y, endowed with the weak* topology. For $\gamma \in \Delta Y$, we let $\operatorname{supp}\gamma$ denote the support of γ. For a set X, a transition $g : X \to \Delta Y$, a point $\bar{x} \in X$, and a Borel subset $\hat{Y} \subseteq Y$, we let $g(\hat{Y}|\bar{x}) := g(\bar{x})(\hat{Y})$. For a set Z, a function $h : X \to Z$, and a subset $\hat{X} \subseteq X$, we let $h(\hat{X}) := \left\{h(x) :\ x \in \hat{X}\right\}$. If $X, Z \subseteq \mathbb{R}$, we write $X \ll Z$ if $x < z$ for every $x \in X$ and $z \in Z$.

[5]To simplify the statements of our results, we assume $M \supseteq A \cup \Delta A \cup \Delta\Theta$. The sender's attainable payoffs would be the same if we instead imposed that $|M| \geq \min\{|A|, |\Theta|\} < \infty$, by Proposition 3, Corollary 1, and Carathéodory's Theorem.

[6]In section 6.3, we consider more general ways of modeling state-independent S preferences.

[7]That is, $\int_{\hat{\Theta}} \sigma(\hat{M}|\cdot)\,\mathrm{d}\mu_0 = \int_{\Theta} \int_{\hat{M}} \beta(\hat{\Theta}|\cdot)\ \mathrm{d}\sigma(\cdot|\theta)\ \mathrm{d}\mu_0(\theta)$ for every Borel $\hat{\Theta} \subseteq \Theta$ and $\hat{M} \subseteq M$.

[8]Specifically, $\mathcal{E} = (\sigma, \rho, \beta)$ induces measure $\mathbb{P}_{\mathcal{E}} \in \Delta(\Theta \times M \times A)$, which assigns probability $\mathbb{P}_{\mathcal{E}}(\hat{\Theta} \times \hat{M} \times \hat{A}) = \int_{\hat{\Theta}} \int_{\hat{M}} \rho(\hat{A}|\cdot)\ \mathrm{d}\sigma(\cdot|\theta)\ \mathrm{d}\mu_0(\theta)$ for every Borel $\hat{\Theta} \subseteq \Theta, \hat{M} \subseteq M, \hat{A} \subseteq A$.

pass one of $n \in \mathbb{N}$ reforms or to pass none. A given reform $i \in \{1, \ldots, n\}$ provides uncertain benefit $\theta_i \in [0, 1]$ to the lawmaker. From the lawmaker's perspective, reforms are ex-ante identical: Their benefits are distributed according to an exchangeable prior μ_0 over $[0, 1]^n$, and each entails an implementation cost of c. The think tank prefers higher indexed reforms to lower indexed ones, and prefers some reform to no reform; that is, the think tank's payoffs are given by a strictly increasing function, $u_S : \{0, \ldots, n\} \to \mathbb{R}$, where we normalize $u_S(0) = 0$.[9] We analyze this example more fully in section 5.1.

Example 2. Suppose R is an investor consulting a broker (S) about an asset. The broker knows the investor's ideal position in the asset, $\theta \in \Theta = [0, 1]$, which is distributed according to the atomless prior, μ_0. The investor's pre-existing position is $a_0 \in [0, 1]$. After consulting his broker, the investor chooses a new position in the asset, $a \in A = [0, 1]$. The broker has state-independent preferences: Her payoff accrues from brokerage fees proportional to the net volume of trade; that is, $u_S(a) = \phi|a - a_0|$ for some $\phi > 0$. The investor wants to match the ideal holdings level, but also wishes to minimize the broker's fees: $u_R(a, \theta) = -\frac{1}{2}(a - \theta)^2 - u_S(a)$. In section 5.2, we find a Pareto-dominant equilibrium and conduct comparative statics under the assumption that the investor's existing position is correct; that is, $a_0 = \int_\Theta \theta \, \mathrm{d}\mu_0(\theta)$.

Example 3. Our setting nests the model of Chakraborty and Harbaugh (2010), who study cheap talk by an expert with state-independent preferences in a multi-dimensional environment. In their model, $\Theta = A \subseteq \mathbb{R}^N$ is a convex set with a nonempty interior where $N > 1$, the prior admits a full-support density, and u_R satisfies $\arg\max_{a \in A} \int u_R(a, \cdot) \, \mathrm{d}\mu = \left\{\int \theta \, \mathrm{d}\mu(\theta)\right\}$ for every $\mu \in \Delta\Theta$. Chakraborty and Harbaugh's (2010) main result is that this setting always admits an equilibrium in which S's messages influences R's actions. Chakraborty and Harbaugh (2010) also point out that in their model, influencing R's actions always benefits S whenever u_S is quasiconvex, and show this assumption holds in several prominent examples. In section 5.3, we take one such example and show how one can use our results to expand on their analysis.

[9] This example is related to, but formally distinct from, the respective models of Che et al. (2013) and Chung and Harbaugh (2016). The former studies a project selection model with state-dependent preferences for both players, and the latter experimentally tests a binary-state project selection model with a stochastic receiver outside option.

We analyze our model via the belief-based approach, commonly used in the communication literature. This approach uses the ex-ante distribution over R's posterior beliefs, $p \in \Delta\Delta\Theta$, as a substitute for both S's strategy and the equilibrium belief system. Clearly, every belief system and strategy for S generate some such distribution over R's posterior belief. By Bayes' rule, this posterior distribution averages to the prior, μ_0. That is, $p \in \Delta\Delta\Theta$ satisfies $\int \mu \, \mathrm{d}p\,(\mu) = \mu_0$. We refer to any p that averages back to the prior as an **information policy**. Thus, only information policies can originate from some σ and β. The fundamental result underlying the belief-based approach is that every information policy can be generated by some σ and β.[10] Let $\mathcal{I}(\mu_0)$ denote the set of all information policies.

The belief-based approach allows us to focus on the game's outcomes. Formally, an **outcome** is a pair, $(p, s) \in \Delta\Delta\Theta \times \mathbb{R}$, representing R's posterior distribution, p, and S's ex-ante payoff, s. An outcome is an **equilibrium outcome** if it corresponds to an equilibrium.[11] In contrast to equilibrium, a triple (σ, ρ, β) is a **commitment protocol** if it satisfies the first two of the three equilibrium conditions above; and (p, s) is a **commitment outcome** if it corresponds to some commitment protocol. In other words, commitment outcomes do not require S's behavior to be incentive compatible.

Using the belief-based approach, Aumann and Hart (2003) analyze, among other things, the outcomes of the cheap-talk model with general S preferences over states and actions. When S's preferences are state independent, their characterization essentially specializes to Lemma 1 below,[12] which describes the game's equilibrium outcomes using the belief-based approach. To state the lemma, let $V\,(\mu)$ be S's possible continuation values from R having μ as his posterior,[13]

$$
\begin{aligned}
V : \Delta\Theta &\rightrightarrows \mathbb{R} \\
\mu &\mapsto \text{co } u_S \left(\arg\max_{a \in A} \int u_R(a, \cdot) \, \mathrm{d}\mu \right).
\end{aligned}
$$

[10] For example, see Aumann and Maschler (1995), Benoît and Dubra (2011) or Kamenica and Gentzkow (2011).

[11] That is, an equilibrium $\mathcal{E} = (\sigma, \rho, \beta)$ exists such that $p(\hat{\Theta}) = \text{marg}_M \mathbb{P}_{\mathcal{E}} \left[\beta^{-1}(\hat{\Theta}) \right]$ for every Borel $\hat{\Theta} \subseteq \Theta$, and $s = \int_A u_S \text{ d marg}_A \mathbb{P}_{\mathcal{E}}$.

[12] Because Aumann and Hart's (2003) setting is finite, we provide a direct independent proof of said lemma for the sake of completeness.

[13] In this paper, "co" refers to the convex hull, and "$\overline{\text{co}}$" refers to the closed convex hull.

Notice V is a Kakutani correspondence,[14] and so the **value function**, $v := \max V : \Delta\Theta \to \mathbb{R}$, is upper semicontinuous.[15]

Lemma 1. *The outcome (p, s) is an equilibrium outcome if and only if:*

1. *$p \in \mathcal{I}(\mu_0)$, i.e., $\int \mu \, \mathrm{d}p\,(\mu) = \mu_0$, and*
2. *$s \in \bigcap_{\mu \in \mathrm{supp}(p)} V(\mu)$.*

The lemma's conditions reflect the requirements of perfect Bayesian equilibrium. The first condition comes from the equivalence between Bayesian updating and p being an information policy. The second condition combines both players' incentive-compatibility constraints. For S, incentive compatibility requires her continuation value to be the same from all posteriors in p's support, meaning her ex-ante value must be equal to her continuation value upon sending a message. For R, incentive compatibility requires that $V\,(\mu)$ contain S's continuation value, which equals S's ex-ante value, from any message that leaves R at posterior belief μ. Therefore, S's ex-ante value must be in $V\,(\mu)$ for all posteriors μ in p's support.

3 Securability

This section presents our main observation: S gains credibility from garbling self-serving information. Theorem 1 applies this observation to characterize S's equilibrium payoffs, and shows that, as far as S's payoffs are concerned, one can ignore S's incentive constraints by focusing on the least persuasive message in any given information policy. Thus, using the theorem, one can use non-equilibrium information policies to reason about S's possible equilibrium payoffs.

Let p be an information policy, and take s to be some possible S payoff. Say that policy p **secures** s if $p\,\{v \geq s\} = 1$,[16] and that s is **securable** if an information policy exists that secures s, that is, if $\mu_0 \in \overline{\mathrm{co}}\,\{v \geq s\}$. The main result of this section shows that securability characterizes S's equilibrium values.

[14] That is, V is a nonempty-, compact-, and convex-valued, upper hemicontinuous correspondence. This follows from Berge's Theorem, given that both objective functions are continuous.

[15] That is, $v(\mu) = \max_{s \in V(\mu)} s$ for each $\mu \in \Delta\Theta$. This maximum exists because V is nonempty- and compact-valued, and v is upper semicontinuous because V is upper hemicontinuous.

[16] Here, we use the standard notation: $\{v \geq s\} = \{\mu :\ v(\mu) \geq s\}$.

Theorem 1. *Suppose $s \geq v(\mu_0)$.*[17] *Then, an equilibrium inducing sender payoff s exists if and only if s is securable.*

The key observation behind Theorem 1 is that one can transform any policy p that secures s into an equilibrium policy by degrading information. Specifically, we replace every supported posterior μ with a different posterior μ' that lies on the line segment between μ_0 and μ. Because μ' is between μ_0 and μ, replacing μ with μ' results in a weakly less informative signal. To ensure the resulting signal is an equilibrium, we take μ' to be the closest posterior to μ_0 among the posteriors between μ_0 and μ that make s incentive compatible for R. Thus, this transformation replaces a potentially incentive-*incompatible* posterior μ with the incentive-*compatible* μ'. That μ' exists follows from two facts. First, s is between S's no-information value and her highest μ payoff, $v(\mu)$. Second, V is a Kakutani correspondence, admitting an intermediate value theorem.

The above logic also identifies a class of equilibrium information policies that span all of S's equilibrium payoffs above $v(\mu_0)$. Say that p **barely secures** s if $\{v \geq s\} \cap \text{co}\{\mu, \mu_0\} = \{\mu\}$ holds for every $\mu \in \text{supp}\, p$. Thus, barely securing policies are policies that secure a payoff higher than what S can attain at any belief between any supported posterior and the prior. The construction behind Theorem 1 transforms every securing policy into a barely securing policy that is also an equilibrium. Because all equilibrium values are securable, we thus have that any high equilibrium value can be attained in an equilibrium with a barely securing policy. Moreover, because barely securing policies are left untouched by Theorem 1's transformation, every barely securing policy must then be an equilibrium.

Using Theorem 1 yields a convenient formula for S's maximal equilibrium value, which we present in Corollary 1 below.

Corollary 1. *An S-preferred equilibrium exists, giving the sender a payoff of $v^*(\mu_0)$, where*

$$v^*(\cdot) := \max_{p \in \mathcal{I}(\cdot)} \inf\, v(\text{supp}\, p).$$

Notice that $\inf\, v(\text{supp}\, p)$ is the highest value that p secures. Thus, Corollary 1 says that maximizing S's equilibrium value is equivalent to maximizing the highest value S can secure.

[17] Given our focus on S's benefits from cheap talk, we state the theorem for high S values. For $s \leq \min\, V(\mu_0)$, one replaces the requirement that s is securable with the existence of some $p \in \mathcal{I}(\mu_0)$ such that $p\{\min\, V \leq s\} = 1$.

Corollary 1 highlights the way incentives constrain S's ability to extract value from her information. Although S can always garble self-serving information to guarantee incentives, the same cannot be done to information that is self-harming.[18] As such, S's highest value is determined by the best worst message she must send if she could commit.

The above makes it easy to see S can do no better than no information if and only if she cannot avoid sending R messages that are worse than providing no information. Said differently, the set of beliefs at which S attains a value strictly higher than no information does not contain the prior in its closed convex hull; that is, the two can be separated via a hyperplane. The following corollary summarizes this observation.

Corollary 2. *Suppose A is finite.*[19] *Then $v^*(\mu_0) = v(\mu_0)$ if and only if a continuous function $\Upsilon : \Theta \to \mathbb{R}$ exists such that $\int \Upsilon(\cdot)\, \mathrm{d}\mu_0 < 0 \leq \int \Upsilon(\cdot)\, \mathrm{d}\mu$ holds whenever $v(\mu) > v(\mu_0)$.*

One can interpret Corollary 2 as tracing S's inability to benefit from cheap talk to a two-action model in which S has a strict preference: say, $A = \{0, 1\}$ and $u_S(1) > u_S(0)$. R's incentives in such a model are captured by the function $\Upsilon(\theta) = u_R(1, \theta) - u_R(0, \theta)$, which represents R's payoff difference between his two actions at state θ. In a two-action model, μ gives S a higher payoff than the prior if and only if R is willing to take action 1 under μ but is unwilling to take it under the prior; that is, $\int \Upsilon(\cdot)\, \mathrm{d}\mu \geq 0 > \int \Upsilon(\cdot)\, \mathrm{d}\mu_0$. By the corollary, the same condition holds for all beliefs that benefit S over the prior in *all* models at which S can do no better than no information. In this sense, the two-action model is the prototypical example in which S cannot benefit from cheap talk.

4 Commitment's Value in Persuasion

The current section uses Theorem 1 to examine the value of commitment in persuasion. The main result of this section is Theorem 2, which geometrically characterizes S's maximal equilibrium value. Take $\bar{v} : \Delta\Theta \to \mathbb{R}$ and $\hat{v} : \Delta\Theta \to \mathbb{R}$ to denote

[18] The statement is true for sender-beneficial payoffs. For sender-harmful payoffs, the sender would garble self-harming information to guarantee incentives.

[19] For infinite A, one needs to adjust the statement as follows: $v^*(\mu_0) = v(\mu_0)$ holds if and only if for all $\epsilon > 0$, there exists $\Upsilon_\epsilon : \Theta \to \mathbb{R}$ such that $\int \Upsilon_\epsilon(\cdot)\, \mathrm{d}\mu_0 < 0$ and $v(\mu) \geq v(\mu_0) + \epsilon$ only if $\int \Upsilon_\epsilon(\cdot)\, \mathrm{d}\mu \geq 0$.

the **quasiconcave envelope** and **concave envelope** of v, respectively. That is, $\bar{v}$ (resp. $\hat{v}$) is the pointwise lowest quasiconcave (concave) and upper semicontinuous function that majorizes v. Because concavity implies quasiconcavity, the quasiconcave envelope lies (weakly) below the concave envelope. Figure 2 below illustrates the definitions of the concave and quasiconcave envelopes for an abstract function.

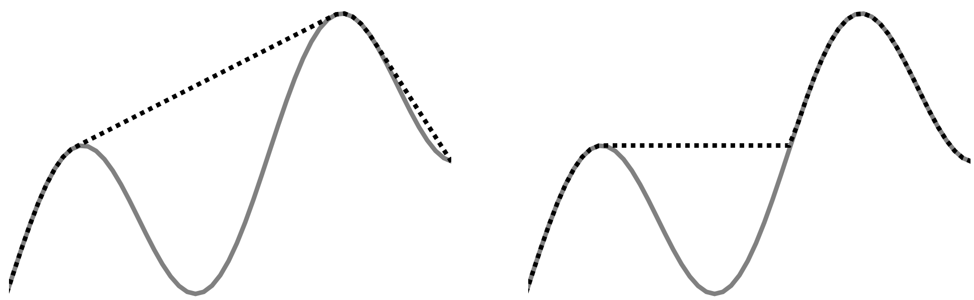

Figure 2: A function with its concave (left) and quasiconcave (right) envelopes.

As described in Aumann and Maschler (1995)[20] and Kamenica and Gentzkow (2011), $\hat{v}$ gives S's payoff from her favorite commitment outcome. Theorem 2 below shows $\bar{v}$ gives S's maximal value under cheap talk.

Theorem 2. *S's maximal equilibrium value is given by v's quasiconcave envelope; that is,* $v^* = \bar{v}$.

We now prove Theorem 2 for the case in which Θ is finite, with the proof for the general model being relegated to the appendix. The appendix also contains the proofs of two auxillary facts, which the proof below takes as given. First, the quasiconcave envelope exists. And second, for finite Θ, the quasiconcave envelope is below *every* quasiconcave function that majorizes v, not just ones that are upper semicontinuous.

Proof of Theorem 2 with finite Θ. We begin by showing the value identified by Corollary 1, v^*, majorizes $\bar{v}$. Because S can always obtain $v(\mu_0)$ by providing no information, v^* clearly majorizes v. Therefore, we can show v^* majorizes $\bar{v}$ by showing v^* is quasiconcave. For this purpose, fix μ', μ'' and $\lambda \in (0,1)$, and consider the following observations. First, if $p' \in \mathcal{I}(\mu')$, and $p'' \in \mathcal{I}(\mu'')$, then $\lambda p' + (1-\lambda) p'' \in$

[20]According to Hart (2006), the relevant results published in Aumann and Maschler (1995) were circulated in the 1960s.

$\mathcal{I}(\lambda\mu' + (1-\lambda)\mu'')$. Second, the support of the convex combination of two distributions is the union of their supports. Taken together, these observations imply the following inequality chain:

$$\begin{aligned} v^*(\lambda\mu' + (1-\lambda)\mu'') &= \max_{p \in \mathcal{I}(\lambda\mu + (1-\lambda)\mu')} \inf\, v(\operatorname{supp} p) \\ &\geq \max_{p' \in \mathcal{I}(\mu'), p'' \in \mathcal{I}(\mu'')} \inf\, v(\operatorname{supp} p' \cup \operatorname{supp} p'') \\ &= \max_{p' \in \mathcal{I}(\mu'), p'' \in \mathcal{I}(\mu'')} \min\{\inf\, v(\operatorname{supp} p'),\ \inf\, v(\operatorname{supp} p'')\} \\ &= \min\{v^*(\mu'), v^*(\mu'')\}, \end{aligned}$$

where the last equality follows from reasoning separately for p' and p''. Thus, v^* is a quasiconcave function that majorizes v; that is, $v^* \leq \bar{v}$.

We now show $\bar{v}$ majorizes v^*. To do so, fix some prior $\mu \in \Delta\Theta$ and let $p \in \mathcal{I}(\mu)$ be an information policy securing S's favorite equilibrium value, $v^*(\mu)$. Because Θ is finite, Carathéodory's Theorem delivers a finite subset $B \subseteq \operatorname{supp} p$ whose convex hull includes the prior. Combined with $\bar{v}$ being a quasiconcave function majorizing v, we have that

$$v^*(\mu) = \inf\, v(\operatorname{supp} p) \leq \min\, v(B) \leq \min\, \bar{v}(B) \leq \bar{v}(\mu).$$

Thus, we have shown both $\bar{v} \leq v^*$ and $\bar{v} \geq v^*$, meaning the two functions coincide. □

Theorem 2 provides a geometric comparison between persuasion's value under cheap talk and under commitment. With commitment, communication is only restricted by R's incentives and Bayes' rule. The value function's concave envelope describes the maximal payoff S can attain in this manner. Replacing the value function's concave envelope with its quasiconcave envelope expresses the value S loses in cheap talk due to her incentive constraints. Graphically, both envelopes allow S to extract value from connecting points on the graph of S's value function. However, although with commitment S can connect points via any affine segment, cheap talk restricts her to flat ones. One can see the associated value loss for the introduction's example in Figure 1: For beliefs in $\mu \in \left(0, \frac{3}{4}\right)$, S's highest cheap-talk value is 1, whereas with commitment, her highest payoff is given by $1 + \frac{4}{3}\mu$.

Corollary 3 below uses the geometric difference between cheap talk and commitment to show that, in a finite setting, commitment is valuable for most priors. In particular, with finite actions and states, the following is true for all priors lying outside a measure zero set: Either S attains her first-best feasible payoff, or S strictly benefits from commitment.

Corollary 3. *Suppose A and Θ are finite. Then, for Lebesgue almost all $\mu_0 \in \Delta\Theta$, either $\bar{v}(\mu_0) = \max v(\Delta\Theta)$ or $\bar{v}(\mu_0) < \hat{v}(\mu_0)$.*

The intuition for the corollary is geometric: Except at S's first-best feasible payoff, the concave envelope, $\hat{v}$, must lie above the interior of any of the quasiconcave envelope's flat surfaces. To see why, notice any prior μ_0 in the interior of such a surface can be expressed as a convex combination of another belief on the same surface and a belief yielding S's first-best feasible value. Said formally, some $\lambda \in (0,1)$, μ, and μ' exist such that $\bar{v}(\mu) = \bar{v}(\mu_0)$, $\bar{v}(\mu') = \max v(\Delta\Theta)$, and $\mu_0 = \lambda\mu + (1-\lambda)\mu'$. Because $\bar{v}$ lies below $\hat{v}$, and because $\hat{v}$ is concave, we obtain

$$\bar{v}(\mu_0) < \lambda\bar{v}(\mu) + (1-\lambda)\bar{v}(\mu') \leq \lambda\hat{v}(\mu) + (1-\lambda)\hat{v}(\mu') \leq \hat{v}(\mu_0),$$

as required.

5 Applications

5.1 The Think Tank

This section uses our results to analyze Example 1. We characterize the think tank's maximal equilibrium value and find an equilibrium in a barely securing policy that attains it. To ease notation, we assume in the main text that the probability that two reforms yield the same benefit to the lawmaker is zero.

Our analysis is based on the following claim, made possible by Theorem 1.

Claim 1. *The following are equivalent, given $k \in \{1, \ldots, n\}$.*

1. *The think tank can attain the value $u_S(k)$ in equilibrium.*

2. $\int \max_{i \in \{k,\ldots,n\}} \theta_i \, \mathrm{d}\mu_0(\theta) \geq c$.

3. *The policy, $p_k \in \mathcal{I}(\mu_0)$, that reveals the random variable*

$$\mathbf{i}_k := \arg\max_{i \in \{k,\ldots,n\}} \theta_i$$

to the lawmaker secures $u_S(k)$.

The claim says the think tank attaining a value of $u_S(k)$ in equilibrium is equivalent to two other conditions. First, always choosing the status quo is ex-ante worse for the lawmaker than always choosing the best reform from $\{k, \dots, n\}$ (Part 2). Second, telling the lawmaker nothing but the identity of the best reform from $\{k, \dots, n\}$ secures $u_S(k)$ (Part 3).

Notice that Claim 1's Part 2 provides a simple necessary and sufficient condition for $u_S(k)$ to be an equilibrium value. Using this condition, we can find S's maximal value across all equilibria: it is given by $u_S(k^*)$, where

$$k^* := \max \left\{ k \in \{1, \dots, n\} : \int \max_{i \in \{k, \dots, n\}} \theta_i \, \mathrm{d}\mu_0(\theta) \geq c \right\}.$$

That is, k^* is the highest k for which Part 2 holds. With k^* in hand, we can identify a best equilibrium for the think tank using the claim's Part 3. Part 3 tells us the think tank's favorite equilibrium value, $u_S(k^*)$, is securable by the information policy, p_{k^*}, that reveals to the lawmaker the identity of the best reform from the set $\{k^*, \dots, n\}$. Thus, to find an equilibrium, we can take p_{k^*} and garble information message by message to obtain a new policy that barely secures $u_S(k^*)$. Doing so results in a policy that has the think tank randomizing between accurately recommending the lawmaker's best reform from $\{k^*, \dots, n\}$ with probability $1 - \epsilon$, and recommending a uniformly drawn reform from $\{k^*, \dots, n\}$ with probability ϵ. By choosing ϵ appropriately, one can degrade information so as to make the lawmaker indifferent between the suggested recommendation and the status quo. One gets an equilibrium by having the lawmaker implement the suggested reform i with probability $\frac{k^*}{i}$ and maintains the status quo with complementary probability. Thus, all that remains is to calculate k^* and ϵ, which depend on the prior: For example, if $\theta_1, \dots, \theta_n$ are i.i.d. uniformly distributed on $[0, 1]$ and $c > \frac{1}{2}$,[21] then

$$k^* = \left\lfloor n - \frac{2c-1}{1-c} \right\rfloor, \text{ and } \epsilon = 2\left[(1-c) - \frac{2c-1}{n-k^*}\right].$$

The policy p_{k^*} also yields an easy lower bound on commitment's value. Specifically, the value of commitment is at least the difference between k^* and the think tank's

[21] When $c \leq \frac{1}{2}$, the think tank can obtain its first-best outcome under no information; that is, $v(\mu_0) = u_S(n)$.

value function's expectation under p_{k^*},

$$\int v(\cdot)\ \mathrm{d}p_{k^*} - u_S(k^*) = \tfrac{1}{n-k^*+1}\sum_{i=k^*}^{n} u_S(i) - u_S(k^*),$$

which simplifies to $\frac{1}{2}(n-k^*)$ in the special case of $u_S(a)=a$.

5.2 The Broker

We now revisit the setting of Example 2 under the assumption that the investor's initial holdings are correct given her information, that is, that $a_0 = \int \theta\ \mathrm{d}\mu_0(\theta)$. Even without this assumption, characterizing optimal behavior by the investor is straightforward. For any posterior belief $\mu \in \Delta\Theta$, simple calculus yields that the investor's unique best response is

$$a^*(\mu) = \begin{cases} \int \theta\ \mathrm{d}\mu(\theta) + \phi & : \int \theta\ \mathrm{d}\mu(\theta) - a_0 \le -\phi \\ a_0 & : \int \theta\ \mathrm{d}\mu(\theta) - a_0 \in [-\phi, \phi] \\ \int \theta\ \mathrm{d}\mu(\theta) - \phi & : \int \theta\ \mathrm{d}\mu(\theta) - a_0 \ge \phi. \end{cases}$$

As such, V is a single-valued correspondence, with $v(\mu) = \phi\left[\left|\int \theta\ \mathrm{d}\mu(\theta) - a_0\right| - \phi\right]_+$.[22]

The above expression demonstrates this example is a specific instance of a class of models in which $\Theta \subseteq \mathbb{R}$ and S's value function is a quasiconvex function of R's expectation of the state. The special one-dimensional structure of this class allows us to focus on *cutoff policies*. Formally, p is a θ^*-**cutoff policy** if it reports whether the state is above or below $\theta^* \in \Theta$.[23] The following proposition shows that garblings of cutoff policies are sufficient to attain any S equilibrium value in one-dimensional settings.

Claim 2. *Suppose $\Theta \subseteq \mathbb{R}$, μ_0 is atomless, and that $v(\mu) = v_M\left(\int \theta\ \mathrm{d}\mu(\theta)\right)$, for some weakly quasiconvex $v_M : \mathrm{co}\Theta \to \mathbb{R}$. Then the following are equivalent for all $s \ge v(\mu_0)$.*

1. *S can attain payoff s in equilibrium.*

2. *The payoff s is securable by a cutoff policy.*

[22] We let $[\cdot]_+ := \max\{\cdot, 0\}$.

[23] In section B.2 in the appendix, we provide a definition of cutoff policies (and prove a version of Claim 2) that applies for general priors. The two definitions coincide when the prior is atomless.

Moreover, an S-preferred equilibrium outcome (p, s) exists such that p is a cutoff policy.

We now apply the claim to our specific broker example. Notice the broker's value function is given by $v(\mu) = v_M\left(\int \theta \,\mathrm{d}\mu(\theta)\right)$, where $v_M(\theta) = \phi\left[|\theta - a_0| - \phi\right]_+$. Because v_M is a convex function, Claim 2 implies an S-preferred equilibrium exists in which S uses a cutoff policy. Consider the median-cutoff policy, where the broker tells the investor whether the state is above or below the median. Let $\theta_<$ and $\theta_>$ denote the investor's expectation of the state conditional on it being below or above the median, respectively. Because $a_0 = \int \theta \,\mathrm{d}\mu_0(\theta) = \frac{1}{2}\theta_< + \frac{1}{2}\theta_>$, one has $|\theta_> - a_0| = |\theta_< - a_0|$, meaning $v_M(\theta_<) = v_M(\theta_>)$. Thus, the median cutoff policy is an equilibrium policy. Moreover, v_M decreases on $[\theta_<, a_0]$ and increases on $[a_0, \theta_>]$, and so no alternative cutoff policy can secure a higher value. Therefore, Claim 2 tells us the median cutoff policy yields a broker-preferred equilibrium. We can therefore calculate the broker's maximal equilibrium payoff,

$$\bar{v}(\mu_0) = \phi\left[\frac{1}{2}(\theta_> - \theta_<) - \phi\right]_+ . \tag{1}$$

In the median-cutoff equilibrium, the broker's information does not depend on ϕ. This observation simplifies the task of conducting comparative statics in ϕ: The broker's maximal equilibrium payoff is single-peaked in ϕ, with the optimal ϕ being $\frac{1}{4}(\theta_> - \theta_<)$. Intuitively, increasing ϕ reduces trade but increases the broker's income per trade, with the latter effect dominating for low ϕ and the former dominating for high ϕ.

It is easy to see the broker's maximal equilibrium payoff increases with mean-preserving spreads of μ_0; that is, the more volatile the market is, the better off is the broker. However, not all volatility is made equal: Mean-preserving spreads strictly increase the broker's payoff if and only if they increase $\theta_> - \theta_<$. Thus, for the broker to strictly benefit from market volatility, she must be able to communicate about it to the investor.

How does the investor fare in the broker's preferred equilibrium? Simple algebra reveals the investor's payoff is $\frac{1}{2\phi^2}s^2 - \mathrm{Var}_{\theta\sim\mu_0}(\theta)$ in any equilibrium yielding the broker a payoff of s.[24] Two consequences are immediate. First, the investor's equilibrium payoffs increase with the broker's, meaning the broker's favorite equilibrium is Pareto

[24] The reader can find said algebra in appendix B.2.

dominant. And, second, the investor's payoffs in the Pareto-dominant equilibrium are given by

$$\frac{1}{2}\left\{\left[\frac{1}{2}\left(\theta_{>}-\theta_{<}\right)-\phi\right]_{+}\right\}^{2}-\operatorname{Var}_{\theta\sim\mu_0}(\theta).$$

Notice the investor is always better off with lower brokerage fees: Because the broker's information does not change with ϕ, a lower ϕ means the investor pays less for the same information. By contrast, the investor's attitude toward higher prior volatility (in the sense of mean-preserving spreads) is ambiguous. Intuitively, increased market volatility both increases the investor's risk and increases the usefulness of the broker's recommendations. As such, higher volatility that does not change the broker's recommendations unambiguously hurts the investor.

5.3 The Salesperson

In this section, we analyze an example presented in Chakraborty and Harbaugh (2010). A buyer (R) can take an outside option or buy one of N goods from a seller (S). The seller knows the buyer's net value from each product i, which we denote by θ_i. Thus, the seller knows the vector, $\theta=(\theta_1,\ldots,\theta_n)$, that we take without loss to be distributed according to μ_0 over the unit cube; that is, $\Theta=[0,1]^n$. We assume here that the products' values are i.i.d. and that the seller wants to maximize the probability of a sale, but does not care which product is sold. Hence, the seller receives a value of 1 if the buyer chooses to purchase product $i\in\{1,\ldots,n\}$, and 0 if the buyer chooses the outside option, which we denote by 0. Only the buyer knows her value from the outside option, ϵ, which is distributed independently from θ according to G, a continuous CDF over $[0,1]$.

In this example, Chakraborty and Harbaugh's (2010) insight delivers the existence of an influential equilibrium, that is, an equilibrium in which different messages lead to different action distributions by the buyer. Because every influential equilibrium in this setting benefits the seller, Chakraborty and Harbaugh's (2010) result implies the seller benefits from cheap talk. In this section, we find a seller-preferred equilibrium and obtain a full characterization of when commitment is valuable for the seller.

Because the buyer has private information, this example does not formally fall within our model. Our analysis, however, still applies. Given a belief $\mu\in\Delta\Theta$, the buyer purchases the good with probability $\mathbb{P}\left\{\epsilon\leq\max_i\int\theta_i\,\mathrm{d}\mu(\theta)\right\}=G\left(\max_i\int\theta_i\,\mathrm{d}\mu(\theta)\right)$.

As such, $v(\mu) := G\left(\max_i \int \theta_i \, \mathrm{d}\mu(\theta)\right)$ is the seller's continuation value from sending a message that gives the buyer a posterior of μ. Using the continuous function v as the seller's value function, we can directly apply our results to this example.

Applying Theorem 2 yields an upper bound on the seller's equilibrium values. To obtain this bound, define the continuous function $\bar{v}^*(\mu) := G\left(\int \max_{j\in\{1,\dots,n\}} \theta_j \, \mathrm{d}\mu(\theta)\right)$. Being an increasing transform of an affine function, $\bar{v}^*$ is quasiconcave. Moreover, because G is increasing, Jensen's inequality tells us

$$\bar{v}^*(\mu) = G\left(\int \max_{j\in\{1,\dots,n\}} \theta_j \, \mathrm{d}\mu(\theta)\right) \geq G\left(\max_{j\in\{1,\dots,n\}} \int \theta_j \, \mathrm{d}\mu(\theta)\right) = v(\mu).$$

In other words, $\bar{v}^*$ is a quasiconcave function that majorizes the seller's value function, and so lies above the value function's quasiconcave envelope. Theorem 2 then implies $\bar{v}^*(\mu_0)$ is above any equilibrium seller value.

We now show one can attain the upper bound $\bar{v}^*(\mu_0)$. Let p^* be the information policy in which the seller tells the buyer the identity of the most valuable product.[25] Assuming the buyer believes the seller, the seller's expected value from recommending product i is

$$G\left(\mathbb{E}_{\theta\sim\mu_0}\left[\theta_i \,\middle|\, i \in \arg\max_{j\in\{1,\dots,n\}} \theta_j\right]\right) = G\left(\int \max_{j\in\{1,\dots,n\}} \theta_j \, \mathrm{d}\mu_0(\theta)\right) = \bar{v}^*(\mu_0),$$

where the first equality follows from product values being i.i.d. Two things are worth noting. Notice all recommendations yield the seller the same value, meaning p^* is an equilibrium. Moreover, p^* attains the upper bound $\bar{v}^*(\mu_0)$ on the seller's equilibrium values. In other words, $(p^*, \bar{v}^*(\mu_0))$ is a seller-preferred equilibrium outcome.

When does the seller benefit from commitment? The answer depends on the relationship between G and its concave envelope, $\hat{G}$, evaluated at $t_0^* := \int \max_{j\in\{1,\dots,n\}} \theta_j \, \mathrm{d}\mu_0(\theta)$.

Claim 3. *The seller benefits from commitment if and only if $\hat{G}(t_0^*) > G(t_0^*)$.*

To see that commitment can benefit the seller only if $\hat{G}(t_0^*) > G(t_0^*)$, observe that $\hat{v}^*(\mu) := \hat{G}\left(\int \max_{j\in\{1,\dots,n\}} \theta_j \, \mathrm{d}\mu(\theta)\right)$ is a continuous and concave function that lies everywhere above the seller's value function. As such, the concave envelope of the seller's value function, $\hat{v}$, lies below $\hat{v}^*$. Thus, if the seller benefits from commitment

[25] That is, the seller reveals the identity of $\arg\max_{i\in\{1,\dots,n\}} \theta_i$.

(i.e., if $\hat{v}(\mu_0) > \bar{v}(\mu_0)$), then

$$\hat{G}(t_0^*) \geq \hat{v}(\mu_0) > \bar{v}(\mu_0) = G(t_0^*).$$

Conversely, suppose $\hat{G}(t_0^*) > G(t_0^*)$. Then, by reasoning analogous to Kamenica and Gentzkow's (2011) Proposition 3,[26] a seller with commitment power can strictly outperform p^* by providing additional information about the value of the best good. Thus, commitment always benefits the seller when $\hat{G}(t_0^*) > G(t_0^*)$.

Claim 3 reduces the question of whether commitment benefits the seller to comparing a one-dimensional function with its concave envelope. Such a comparison is simple when G is well behaved. Indeed, if G admits a decreasing, increasing, or single-peaked density, G itself is concave, convex, or convex-concave, respectively, and so characterizing its concave envelope is straightforward.

Claim 4. *Suppose G admits a continuous density g.*

1. *If g is weakly decreasing, the seller does not benefit from commitment.*
2. *If g is nonconstant and weakly increasing, the seller benefits from commitment.*
3. *If g is strictly quasiconcave, the seller benefits from commitment if and only if $g(t_0^*) > \frac{1}{t_0^*}\int_0^{t_0^*} g(t)\ \mathrm{d}t$.*

The claim's first part says the seller does not benefit from commitment when g is decreasing, that is, when G is concave. The second part says that when G is convex and non-affine, the seller always benefits from commitment. The third part discusses the seller's benefits from commitment when G is S-shaped. Specifically, it shows commitment is valuable in this case if and only if G's density at t_0^* is strictly larger than the average density up to t_0^*.

[26]Proposition 3 of Kamenica and Gentzkow (2011) assumes the state space is finite, and so it does not directly apply here. However, the extension to this continuous setting is straightforward given that G is continuous.

6 Discussion and Related Literature

6.1 Informative Communication

In this section, we ask when a sender with state-independent preferences can provide information credibly using cheap talk. We show the answer is: almost always. Specifically, we show the model almost always admits an **informative equilibrium,** that is, an equilibrium in which R's beliefs change on path.[27] As Proposition 1 below shows, such an equilibrium fails to exist only in the special case in which the state is binary and S strictly benefits from increasing the probability R assigns to one state.

Proposition 1. *No informative equilibrium exists if and only if* $|\Theta| = 2$*, and some* $\theta \in \Theta$ *exists such that, for all* $\mu, \mu' \in \Delta\Theta$*,*

$$\mu(\theta) < \mu_0(\theta) < \mu'(\theta) \implies V(\mu) \ll V(\mu'). \tag{2}$$

The fundamental insight behind Proposition 1 comes from Chakraborty and Harbaugh (2010), who show that a large special case of our model always admits an influential equilibrium, namely, an equilibrium in which R's action is non-constant across S's messages. Chakraborty and Harbaugh (2010) study a model in which both state and action spaces are the same convex, multidimensional Euclidean set; the prior admits a density; and R's optimal action is his expectation of the state. In this setting, they show S can credibly influence R's action by trading off dimensions.

To prove Proposition 1, we slightly generalize Chakraborty and Harbaugh's (2010) insight by showing that whereas S's ability to influence R's actions depends on their parametric setting, S's ability to convey information does not. Intuitively, Chakraborty and Harbaugh's (2010) fixed-point argument relies on two important features of their setting: First, the state is a multidimensional real vector; and second, S's payoff is a single-valued continuous function of R's beliefs. The belief-based approach allows us to mirror the structure of the first feature in *any* setting with three or more states (summarized in Lemma 7 in the appendix). For the second feature, we show that a weaker condition that applies in the general case—S's possible values form a Kakutani correspondence of R's beliefs—lets one apply similar fixed-point reasoning. Therefore, we are left with the two-state case in which we show (2) must be satisfied. Indeed, whenever (2) is violated, one can construct an informative

[27] That is, inducing an information policy that is not equal to δ_{μ_0}.

equilibrium by finding two payoff-equivalent messages for S, each of which increases the probability of a different state. Necessity follows.

6.2 The Equilibrium Payoff Set

Despite our focus on S's favorite equilibrium, our approach is useful for analyzing the entire equilibrium payoff set. To find S's payoff set, notice that because S optimality is defined via indifference, the game's equilibrium set is the same regardless of whether S's objective is u_S or $-u_S$. Just as applying Theorem 1 to the original game characterizes S's high payoffs, one can apply the theorem to the game with S objective $-u_S$ to find S's low equilibrium payoffs. Under this objective, S's value function is given by $-w$, where $w(\cdot) := \min V(\cdot)$. Theorem 1 then implies $s \leq w(\mu_0)$ is an equilibrium payoff in the *original* game if and only if some $p \in \mathcal{I}(\mu_0)$ exists such that $p\{w \leq s\} = 1$. Applying Theorem 2 then tells us S's lowest equilibrium payoff is given by the quasiconvex of envelope of w, which we denote by $\underline{w}$.[28] This gives S's entire equilibrium payoff set: s is an S equilibrium payoff if and only if $s \in [\underline{w}(\mu_0), \bar{v}(\mu_0)]$.

With S's equilibrium payoffs in hand, we can find R's possible equilibrium payoffs using two observations. First, one can implement any particular payoff profile in an equilibrium in which S recommends a pair of actions to R, and R responds by mixing only over the recommended actions. Second, if S's equilibrium payoff is s, then S's recommended action pair must consist of one action yielding S a payoff above s, and one action yielding S a payoff below s. Taking s as given, we can thus reduce the number of action pairs that S may recommend in equilibrium. We discuss these observations more formally in online appendix C.2.

6.3 State-Independent Preferences

Whereas S's objective depends only on R's action in our model, one can write more general models of state-independent preferences. Fix a continuous S objective, $\tilde{u}_S : A \times \Theta \to \mathbb{R}$. We say $\tilde{u}_S$ exhibits **ordinally (cardinally) state-independent preferences** if, for any $\theta, \theta' \in \Theta$ and $a, a' \in A$ ($\alpha, \alpha' \in \Delta A$): $\tilde{u}_S(a, \theta) \geq \tilde{u}_S(a', \theta)$ implies $\tilde{u}_S(a, \theta') \geq \tilde{u}_S(a', \theta')$ ($\int \tilde{u}_S(\cdot, \theta)\, \mathrm{d}\alpha \geq \int \tilde{u}_S(\cdot, \theta)\, \mathrm{d}\alpha'$ implies $\int \tilde{u}_S(\cdot, \theta')\, \mathrm{d}\alpha \geq$

[28]More precisely, $\underline{w}$ is the highest quasiconvex and lower semicontinuous function that is everywhere below w.

$\int \tilde{u}_S(\cdot, \theta')\ \mathrm{d}\alpha')$. Proposition 2 below shows our analysis often generalizes to these notions of state independence.

Proposition 2. *Fix the action space, state space, receiver preferences, and prior. Then the game with sender objective $\tilde{u}_S$ and the game with sender objective $u_S(\cdot) := \int_\Theta \tilde{u}_S(\cdot, \theta)\ \mathrm{d}\mu_0(\theta)$ have the same equilibria, and generate the same equilibrium outcomes, if either of the following conditions holds:*

1. *$\tilde{u}_S$ exhibits cardinally state-independent preferences.*
2. *$\tilde{u}_S$ exhibits ordinally state-independent preferences, and R has a unique best response to every posterior belief.*

The proof's intuition is as follows. Each part assumes S's implied preference relation over all action distributions that can be induced at some belief is state independent. Hence, the incentive constraints implied by u_S and $\tilde{u}_S$ for S are the same for all states. This observation yields two consequences. First, the equilibria under both utility functions must be the same. Second, S must be indifferent across all equilibrium messages regardless of the state. As such, any correlation between R's mixed action and S's message must be payoff-irrelevant for S, implying identical ex-ante S payoff under both $\tilde{u}_S$ and u_S.

A remark is in order regarding the second part of Proposition 2. Although many of our results generalize to a model in which R faces preference shocks that are independent of θ,[29] the second part of the above proposition does not. The reason is that such a generalization may still require S to compare payoffs from non-degenerate action distributions despite R having a unique best response.

6.4 Long and Transparent Cheap Talk

In a pair of influential papers, Aumann and Hart (2003) and Krishna and Morgan (2004) showed that allowing for multiple rounds of bilateral communication – that is, long cheap talk – can meaningfully expand the set of feasible equilibrium outcomes. Whereas Krishna and Morgan (2004) demonstrate this fact in the setting of Crawford

[29] Such a model has a compact metrizable space Z of payoff parameters such that the joint distribution of $(\theta, z) \in \Theta \times Z$ is $\mu_0 \otimes \zeta_0$ for some $\zeta_0 \in \Delta Z$. R has payoffs given by $u_R : A \times Z \times \Theta \to \mathbb{R}$. In this extended model, $V : \Delta\Theta \rightrightarrows \mathbb{R}$ should take the form $V(\mu) = \int_Z \mathrm{co}\, u_S\left(\arg\max_{a \in A} u_R(a, z, \mu)\right)\ \mathrm{d}\zeta_0(z)$, a Kakutani correspondence.

and Sobel (1982), Aumann and Hart (2003) characterize the long cheap-talk outcome set in terms of dimartingales and separation by diconvex functions for general R and S preferences.[30] When S's preferences are state independent, one can obtain such a separating function for S's payoffs using Theorem 1. Using this separating function, one can show that every S payoff attainable in a Nash equilibrium with long cheap talk is also attainable in PBE of the one-shot cheap-talk game. The same, however, is not true for R, who can benefit from long cheap talk. We refer the reader to appendix C.4 for the formal details.

7 Conclusion

In this paper, we study cheap talk under the assumption that the sender has state-independent preferences. Our main observation is that such a sender gains credibility from garbling self-serving information. This observation leads to our main result: Any securable sender value above what the sender can get with no information is attainable in equilibrium. This result allows us to describe the sender-preferred equilibrium value geometrically, and to contrast this value with that of persuasion with commitment. Using our results, we shed light on a rich class of examples and speak to the broader literature on cheap-talk communication.

References

Aliprantis, Charalambos D and Kim Border, *Infinite Dimensional Analysis: A Hitchhiker's Guide*, Springer Science & Business Media, 2006.

Alonso, Ricardo and Odilon Câmara, "Persuading voters," *American Economic Review*, 2016, *106* (11), 3590–3605.

Aumann, Robert J and Michael Maschler, *Repeated games with incomplete information*, MIT press, 1995.

_ **and Sergiu Hart**, "Bi-convexity and bi-martingales," *Israel Journal of Mathematics*, 1986, *54* (2), 159–180.

[30]Although the formal results therein are limited to finite settings, the Aumann and Hart (2003) setting is conceptually richer than ours, featuring a sender who may also make payoff-relevant decisions after communication concludes.

_ **and** _ , "Long cheap talk," *Econometrica*, 2003, *71* (6), 1619–1660.

Battaglini, Marco, "Multiple Referrals and Multidimensional Cheap Talk," *Econometrica*, 2002, *70* (4), 1379–1401.

Benoît, J.P. and Juan Dubra, "Apparent Overconfidence," *Econometrica*, 2011.

Bergemann, Dirk and Stephen Morris, "Bayes correlated equilibrium and the comparison of information structures in games," *Theoretical Economics*, 2016, *11* (2), 487–522.

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

Bester, Helmut and Roland Strausz, "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case," *Econometrica*, 2001, *69* (4), 1077–1098.

Chakraborty, Archishman and Rick Harbaugh, "Comparative cheap talk," *Journal of Economic Theory*, 2007, *132* (1), 70–94.

_ **and** _ , "Persuasion by cheap talk," *American Economic Review*, 2010, *100* (5), 2361–2382.

Che, Yeon Koo, Wouter Dessein, and Navin Kartik, "Pandering to persuade," *American Economic Review*, 2013, *103* (1), 47–79.

Chung, Wonsuk and Rick Harbaugh, "Biased Recommendations from Biased and Unbiased Experts," 2016.

Crawford, Vincent P and Joel Sobel, "Strategic information transmission," *Econometrica*, 1982, *50* (6), 1431–1451.

de Clippel, Geoffroy, "An axiomatization of the inner core using appropriate reduced games," *Journal of Mathematical Economics*, 2008, *44* (3), 316–323.

Dworczak, Piotr and Giorgio Martini, "The simple economics of optimal persuasion," *Journal of Political Economy*, 2018, *Forthcoming.*

Ely, Jeffrey C, "Beeps," *American Economic Review*, 2017, *107* (1), 31–53.

Gentzkow, Matthew and Emir Kamenica, "A Rothschild-Stiglitz approach to Bayesian persuasion," *American Economic Review*, 2016, *106* (5), 597–601.

Green, Jerry R and Nancy L Stokey, "A two-person game of information transmission," *Journal of Economic Theory*, 2007, *135* (1), 90–104.

Hart, Sergiu, "Robert aumann's game and economic theory," *Scandinavian Journal of Economics*, 2006, *108* (2), 185–211.

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

Krishna, Vijay and John Morgan, "The art of conversation: eliciting information from experts through multi-stage communication," *Journal of Economic theory*, 2004, *117* (2), 147–179.

Margaria, Chiara and Alex Smolin, "Dynamic communication with biased senders," *Games and Economic Behavior*, 2018, *110*, 330–339.

Myerson, Roger B, "Incentive compatibility and the bargaining problem," *Econometrica*, 1979, pp. 61–73.

_ , "Multistage games with communication," *Econometrica*, 1986, pp. 323–358.

Phelps, Robert R, *Lectures on Choquet's theorem*, Springer Science & Business Media, 2001.

Schnakenberg, Keith E, "Expert advice to a voting body," *Journal of Economic Theory*, 2015, *160*, 102–113.

Seidmann, Daniel J, "Effective cheap talk with conflicting interests," *Journal of Economic Theory*, 1990, *50* (2), 445–458.

A Omitted Proofs: Main Results

A.1 Preliminaries and additional notation

We begin by noting an abuse of notation that we use throughout the appendix. For a compact metrizable space Y, a Borel measure over it $\gamma \in \Delta Y$, and a γ-integrable function $f : Y \to \mathbb{R}$, we let $f(\gamma) = \int_Y f \, \mathrm{d}\gamma$.

We now document the (standard) notion of information ranking used throughout the paper.

Definition 1. Given $p, p' \in \Delta\Delta\Theta$, say p **is more (Blackwell) informative than** $p^{'}$ if p is a mean-preserving spread of $p^{'}$, that is, if there exists a measurable selector r of $\mathcal{I} : \Delta\Theta \rightrightarrows \Delta\Delta\Theta$ such that $p(D) = \int_{\Delta\Theta} r(D|\cdot) \, \mathrm{d}p^{'}$ for all Borel $D \subseteq \Delta\Theta$.

Now, we record a useful measurable selection result.

Lemma 2. *If $D \subseteq \Delta\Theta$ is Borel and $f : D \to \mathbb{R}$ is any measurable selector of $V|_D$, then a measurable function $\alpha_f : D \to \Delta A$ exists such that, for all $\mu \in D$, the measure $\hat{\alpha} = \alpha_f(\cdot|\mu)$ satisfies:*

1. $u_S(\hat{\alpha}) = f(\mu)$;
2. $\hat{\alpha} \in \arg\max_{\alpha \in \Delta A} u_R(\alpha, \mu)$;
3. $|\mathrm{supp}(\hat{\alpha})| \leq 2$.

Proof. Recall the sender's value correspondence V is upper hemicontinuous. Below, we use several results from Aliprantis and Border (2006) concerning measurability of correspondences. Define

$$\begin{aligned} A^* : D &\rightrightarrows A \\ \mu &\mapsto \arg\max_{a \in A} u_R(a, \mu), \end{aligned}$$

and define $A_+ := A^* \cap [u_S^{-1} \circ \max V]$ and $A_- := A^* \cap [u_S^{-1} \circ \min V]$. As semicontinuous functions, $\max V$, $\min V$ are both Borel measurable and (viewed as correspondences) compact-valued. Also, the correspondence $u_S^{-1} : u_S(A) \rightrightarrows A$ is upper hemicontinuous and compact-valued. Therefore, by Theorem 18.10, the compositions $u_S^{-1} \circ \max V$, $u_S^{-1} \circ \min V$ are weakly measurable. Lemma 18.4(3) then tells us that A_+, A_- are weakly measurable too, and the Kuratowski & Ryll-Nardzewski selection theorem (Theorem 18.13) provides measurable selectors $a_+ \in A_+$, $a_- \in A_-$. Now, define the measurable map:

$$\begin{aligned} \alpha_f : D &\to \Delta A \\ \mu &\mapsto \begin{cases} \frac{v(\mu) - f(\mu)}{v(\mu) - \min V(\mu)} \delta_{a_-(\mu)} + \frac{f(\mu) - \min V(\mu)}{v(\mu) - \min V(\mu)} \delta_{a_+(\mu)} & : \ f(\mu) \in V(\mu) \setminus \{\min V(\mu)\} \\ \delta_{a_-(\mu)} & : \ \min V(\mu) = f(\mu). \end{cases} \end{aligned}$$

By construction, α_f is as desired. $\square$

We now prove a variant of the intermediate value theorem, which is useful for our setting. This result is essentially proven in Lemma 2 of de Clippel (2008). Because the statement of that lemma is slightly weaker than we need, however, we provide a proof here for the sake of completeness.

Lemma 3. *If $F : [0,1] \rightrightarrows \mathbb{R}$ is a Kakutani correspondence with $\min F(0) \leq 0 \leq \max F(1)$, and $\bar{x} = \inf\{x \in [0,1] : \max F(x) \geq 0\}$, then $0 \in F(\bar{x})$.*

Proof. By definition of $\bar{x}$, some weakly decreasing $\{x_n^+\}_{n=1}^{\infty} \subseteq [\bar{x}, 1]$ exists that converges to $\bar{x}$ such that $\max F(x_n^+) \geq 0$ for every $n \in \mathbb{N}$. Define the sequence $\{x_n^-\}_{n=1}^{\infty} \subseteq [0, \bar{x}]$ to be the constant 0 sequence if $\bar{x} = 0$ and to be any strictly increasing sequence that converges to $\bar{x}$ otherwise. By definition of $\bar{x}$ (and, in the case of $\bar{x} = 0$, because $\min F(0) \leq 0$), it must be that $\min F(x_n^-) \leq F(\bar{x}) \leq \max F(x_n^+)$.

Passing to a subsequence if necessary, we may assume (as a Kakutani correspondence has compact range) that $\{\max F(x_n^+)\}_{n=1}^{\infty}$ converges to some $y \in \mathbb{R}$, which would necessarily be nonnegative. Upper hemicontinuity of F then implies $\max F(\bar{x}) \geq 0$. An analogous argument shows $\min F(\bar{x}) \leq 0$. Because F is convex-valued, it follows that $0 \in F(\bar{x})$. $\square$

A.2 Proof for Section 2

Below is the proof of Lemma 1, which initializes our belief-based approach. For finite states, the result can be easily proven from results in Aumann and Hart (2003). Although their ideas easily generalize to infinite state spaces such as ours, we include a direct proof here for completeness.

Proof. First take any equilibrium (σ, ρ, β), and let (p, s) be the induced outcome. That $p \in \mathcal{I}(\mu_0)$ follows directly from the Bayesian property.

Define the interim payoff, $\hat{s} : M \to \mathbb{R}$ via $\hat{s}(m) := u_S(\rho(m))$. S incentive compatibility tells us that some $M^* \subseteq M$ exists such that $\int_\Theta \beta(M^*|\cdot)\, \mathrm{d}\mu_0 = 1$, and for every $m \in M^*$ and $m' \in M$, we have $\hat{s}(m) \geq \hat{s}(m')$. In particular, $\hat{s}(m) = \hat{s}(m')$ for every $m, m' \in M^*$; that is, some $\hat{s}^* \in \mathbb{R}$ exists such that $\hat{s}|_{M^*} = \hat{s}^*$. But then

$$s = \int_\Theta \int_{M^*} u_S(\rho(m))\, \mathrm{d}\sigma(m|\theta)\, \mathrm{d}\mu_0(\theta) = \int_\Theta \int_{M^*} \hat{s}^*\, \mathrm{d}\mu_0(\theta) = \hat{s}^*,$$

so that, by receiver incentive-compatibility, $s \in V(\beta(\cdot|m))$ for every $m \in M^*$. By definition of p, then, $s \in V(\mu)$ for p-almost every $\mu \in \Delta\Theta$. Because V is upper hemicontinuous, it follows that $s \in \bigcap_{\mu \in \text{supp}(p)} V(\mu)$.

Now suppose (p, s) satisfies the three conditions. Define the compact set $D := \text{supp}(p)$. It is well known (see Benoît and Dubra (2011) or Kamenica and Gentzkow (2011)) every that $p \in \mathcal{I}(\mu_0)$ exhibits some S strategy σ and Bayes-consistent belief map $\beta : M \to \Delta\Theta$ that induce distribution p over posterior beliefs.[31] Without disrupting the Bayesian property, we may without loss assume $\beta(m) \in D$ for all $m \in M$. Now let $\alpha = \alpha_s : D \to \Delta A$ be as given by Lemma 2. We can then define the receiver strategy $\sigma := \alpha \circ \beta$, which is incentive compatible for R by definition of α. Finally, by construction, $\int_A u_S \, \mathrm{d}\rho(\cdot|m) = s$ for every $m \in M$, so that every S strategy is incentive-compatible. Therefore, (σ, ρ, β) is an equilibrium that generates outcome (p, s). □

A.3 Proofs for Section 3

A.3.1 Proof of Theorem 1

Below, we prove a lemma that is at the heart of Theorem 1. It constructs an equilibrium (a barely securing policy, which we then show to be compatible with equilibrium) of S value s from an arbitrary information policy securing s. The constructed equilibrium policy is less informative than the original policy and requires fewer messages to implement.

Lemma 4. *Let $p \in \mathcal{I}(\mu_0)$ and $s \in \mathbb{R}$.*

1. *If p secures s and $s \geq v(\mu_0)$, then some $p^* \in \mathcal{I}(\mu_0)$ exists such that p^* barely secures s, p^* is weakly less Blackwell-informative than p, and $|\text{supp}(p^*)| \leq |\text{supp}(p)|$.*

2. *If p barely secures s, then (p, s) is an equilibrium outcome.*

Proof. If $s = v(\mu_0)$, both results are trivial: In this case, the uninformative policy is the unique one that barely secures s. From this point, we focus on the case of $s > v(\mu_0)$.

[31] In particular, such (σ, β) exist with $\sigma(\Delta\Theta|\theta) = 1$ for all $\theta \in \Theta$ and $\beta(\cdot|\mu) = \mu$ for all $\mu \in D$.

Toward the first point, let $p \in \mathcal{I}(\mu_0)$ secure s, and $D := \text{supp}\,(p)$. Notice that $v(\mu) \geq s$ for every $\mu \in D$ because v is upper semicontinuous. Define the measurable function,

$$\begin{aligned}\lambda = \lambda_{p,s} : D &\to [0,1] \\ \mu &\mapsto \inf\left\{\hat{\lambda} \in [0,1] : \; v\left((1-\hat{\lambda})\mu_0 + \hat{\lambda}\mu\right) \geq s\right\}.\end{aligned}$$

By Lemma 3, it must be that $s \in V([1-\lambda(\mu)]\mu_0 + \lambda(\mu)\mu)$ for every $\mu \in S$.

Notice that some number $\epsilon > 0$ exists such that $\lambda \geq \epsilon$ uniformly. If no such ϵ existed, then a sequence $\{\mu_n\}_n \subseteq D$ would exist such that $\lambda(\mu_n)$ converges to zero. But then the sequence $\{([1-\lambda(\mu_n)]\mu_0 + \lambda(\mu_n)\mu_n, \; s)\}_n$ from the graph of V would converge to (μ_0, s). Because V is upper hemicontinuous, such convergence would contradict $s > v(\mu_0)$. Therefore, such an ϵ exists, and $\frac{1}{\lambda}$ is a bounded function.

Now, define $p^* = p^*_s \in \Delta\Delta\Theta$ by letting

$$p^*(\hat{D}) := \left(\int_{\Delta\Theta} \tfrac{1}{\lambda} \; \mathrm{d}p\right)^{-1} \cdot \int_{\Delta\Theta} \tfrac{1}{\lambda(\mu)} \mathbf{1}_{[1-\lambda(\mu)]\mu_0 + \lambda(\mu)\mu \in \hat{D}} \; \mathrm{d}p(\mu)$$

for every Borel $\hat{D} \subseteq \Delta\Theta$. Direct computation shows $p^* \in \mathcal{I}(\mu_0)$, and p^* barely secures s by construction.

Lastly, we note p^* has the other required properties. The map $\mu \mapsto [1-\lambda(\mu)]\mu_0 + \lambda(\mu)\mu$ is a surjection from $\text{supp}(p^*)$ to $\text{supp}(p)$, so that $|\text{supp}(p^*)| \leq |\text{supp}(p)|$. Also by construction, p^* is weakly less informative than $\left(1 - \int_{\Delta\Theta} \lambda \; \mathrm{d}p\right) \delta_{\mu_0} + \left(\int_{\Delta\Theta} \lambda \; \mathrm{d}p\right) p$, which in turn is less informative than p. This proves (1).

Toward (2), suppose p barely secures s. Consider any $\mu \in \text{supp}(p)$. Some subsequence of $\{v((1-2^{-n})\mu + 2^{-n}\mu_0)\}_{n=1}^{\infty} \subseteq [\min u_S(A), \; s]$ converges, leading to (as V is upper hemicontinuous) some element of $V(\mu)$ that is weakly less than s. Because $v(\mu) \geq s$ by hypothesis, and V being convex-valued, it follows that $s \in \text{supp}(p)$. Because $\mu \in \text{supp}(p)$ was arbitrary, Lemma 1 delivers an equilibrium that generates S value s and information policy p. □

We now prove the Securability Theorem (Theorem 1).

Proof. The "only if" direction follows directly from Lemma 1: For any equilibrium outcome (p, s), information policy p secures payoff s. The "if" direction is a direct consequence of (both parts of) Lemma 4. □

A.3.2 Convexity of the equilibrium payoff set, and Corollary 1

For convenience, we record an easy consequence of the Theorem 1.

Corollary 4. *The set of sender equilibrium payoffs is a compact interval.*

Proof. Let Π^* be the set of equilibrium S payoffs, $\Pi_+ := \{s \in \Pi^* : \ s \geq \max V(\mu_0)\}$, $\Pi_- := \{s \in \Pi^* : \ s \leq \min V(\mu_0)\}$, and $\Pi_0 := \{s \in \Pi^* : \ \min V(\mu_0) \leq s \leq \max V(\mu_0)\}$.

As V is convex-valued, $\Pi_0 = \Pi^* \cap V(\mu_0)$. By considering uninformative equilibria, we see that $\Pi_0 = V(\mu_0) = [\min V(\mu_0), \max V(\mu_0)]$.

It follows immediately from Theorem 1 that Π_+ is convex. Letting $s_+ := \sup(\Pi_+) \geq v(\mu_0)$, a sequence $\{s_n\}_{n=1}^{\infty} \subseteq [v(\mu_0),\ s_+]$ exists that converges to s_+. Dropping to a subsequence, if necessary, we may assume some $\{p_n\}_{n=1}^{\infty} \subseteq \mathcal{I}(\mu_0)$ exists such that p_n secures s_n for each n, and $\{p_n\}_n$ converges to some $p_+ \in \mathcal{I}(\mu_0)$. But then p_+ secures s_+ because v is upper semicontinuous, so that (by Theorem 1) $s_+ \in \Pi_+$. It follows that $\Pi_+ = [v(\mu_0),\ s_+]$, a compact interval. By an identical argument, Π_- is a compact interval, say, $[s_-, \min V(\mu_0)]$ as well.[32]

Therefore, $\Pi^* = [s_-, \min V(\mu_0)] \cup [\min V(\mu_0), \max V(\mu_0)] \cup [\max V(\mu_0), s_+] = [s_-, s_+]$. □

Notice, Corollary 1 follows directly from Corollary 4 (which establishes that a best equilibrium for S exists) and Theorem 1.

A.3.3 Corollary 2: Sender-beneficial communication

Below, we prove that, for any $\epsilon > 0$: $v^*(\mu_0) < v(\mu_0) + \epsilon$ holds if and only if some continuous $\Upsilon_\epsilon : \Theta \to \mathbb{R}$ exists such that $\int_\Theta \Upsilon_\epsilon \, \mathrm{d}\mu_0 < 0$, and $v(\mu) \geq v(\mu_0) + \epsilon$ only if $\int_\Theta \Upsilon_\epsilon \, \mathrm{d}\mu \geq 0$.

Proof. Fix any $\epsilon > 0$, and let $s := v(\mu_0) + \epsilon$. By Theorem 1, $v^*(\mu_0) < s$ if and only if $\mu_0 \notin \overline{\mathrm{co}}\, \{v \geq s\}$. The Hahn-Banach Theorem says this holds if and only if some affine continuous $\psi : \mathrm{ca}(\Theta) \to \mathbb{R}$ exists such that $\psi(\mu_0) < 0 \leq \max \psi\, \{v \geq s\}$. But then the Riesz Representation Theorem says this holds if and only if a continuous $\Upsilon_\epsilon : \Theta \to \mathbb{R}$ exists such that $\int_\Theta \Upsilon_\epsilon \, \mathrm{d}\mu_0 < 0$ and $\int_\Theta \Upsilon_\epsilon \, \mathrm{d}\mu \geq 0$ for every $\mu \in \{v \geq s\}$. □

[32]Notice the only property of V used in the proofs—that it is a Kakutani correspondence — is also true of $-V$.

An immediate consequence of the above is a general characterization of when a sender-beneficial equilibrium exists. Indeed, $v^*(\mu_0) = v(\mu_0)$ if and only if $v^*(\mu_0) < v(\mu_0) + \epsilon$ for every $\epsilon > 0$, which holds if and only if every $\epsilon > 0$ admits a Υ_ϵ of the given form.

Another consequence, specializing to the case that A is finite, is Corollary 2. Indeed, then v has a finite range, and (by Corollary 1) v^*'s range is a subset of the same. The corollary then follows from considering any $\epsilon > 0$ small enough that no two elements of $v(\Delta\Theta)$ are within ϵ of each other.

A.4 Proofs for Section 4

A.4.1 The quasiconcave envelope is well-defined

Lemma 5. *There exists a pointwise lowest quasiconcave upper semicontinuous function which majorizes v, and there exists a pointwise lowest quasiconcave function which majorizes v. In particular, the function $\bar{v}$ is well defined.*

Proof. Let $\mathcal{F} := \{f : \Delta\Theta \to \mathbb{R} :\ f \geq v,\ f \text{ is quasiconcave and upper semicontinuous}\}$ or $\{f : \Delta\Theta \to \mathbb{R} :\ f \geq v,\ f \text{ is quasiconcave}\}$, and let $\tilde{v}$ be the pointwise infimum of $\mathcal{F}$. By construction, $\tilde{v}$ lies above v and below any $f \in \mathcal{F}$. Moreover, any infimum of upper semicontinuous functions is itself upper semicontinuous, because a union of open sets is open. Thus, all that remains is to show $\tilde{v}$ as defined is quasiconcave. Given a function $f : \Delta\Theta \to \mathbb{R}$,

$$
\begin{aligned}
f \text{ is quasiconcave} &\iff \{\mu \in \Delta\Theta : f(\mu) \geq f(\tilde{\mu})\} \text{ is convex } \forall \tilde{\mu} \in \Delta\Theta \\
&\iff \{\mu \in \Delta\Theta : f(\mu) \geq s\} \text{ is convex } \forall s \in f(\Delta\Theta) \\
&\iff \{\mu \in \Delta\Theta : f(\mu) \geq s\} \text{ is convex } \forall s \in \mathbb{R},
\end{aligned}
$$

where the last equivalence holds because a union of a chain of convex sets is convex. Therefore,

$$
\begin{aligned}
\tilde{v} \text{ is quasiconcave} &\iff \{\mu \in \Delta\Theta :\ \tilde{v}(\mu) \geq s\} \text{ is convex } \forall s \in \mathbb{R} \\
&\iff \bigcap_{f \in \mathcal{F}} \{\mu \in \Delta\Theta :\ f(\mu) \geq s\} \text{ is convex } \forall s \in \mathbb{R}.
\end{aligned}
$$

The result follows because an intersection of convex sets is convex. $\square$

In the subsection below, we prove $v^* = \bar{v}$, the pointwise lowest quasiconcave upper semicontinuous majorant of v. Given our finite-state proof, in the main text, that v^* is the pointwise lowest quasiconcave majorant of v, we learn the two definitions coincide (i.e., the qualifier "upper semicontinuous" is redundant) when Θ is finite.

A.4.2 Theorem 2: Quasiconcavification

Now, we prove in general our geometric characterization of the sender's best equilibrium value.

Proof. Let $\tilde{v}^* : \Delta\Delta\Theta \to \mathbb{R}$ be given by $\tilde{v}^*(p) := \inf v(\operatorname{supp} p)$, so that $v^*(\mu) := \max_{p \in \mathcal{I}(\mu)} \tilde{v}^*(p)$ for every $\mu \in \Delta\Theta$. Given Corollary 1, our goal is to show $v^* = \bar{v}$.

First, $v^* \geq v$ because $\delta_\mu \in \mathcal{I}(\mu)$ for every $\mu \in \Delta\Theta$. Next, to verify quasiconcavity, any $\mu, \mu' \in \Delta\Theta$, and $\lambda \in (0,1)$ have

$$\begin{aligned}
v^*((1-\lambda)\mu' + \lambda\mu) &= \max_{p^* \in \mathcal{I}((1-\lambda)\mu' + \lambda\mu)} \inf v(\operatorname{supp}(p^*)) \\
&\geq \max_{p^* \in (1-\lambda)\mathcal{I}(\mu') + \lambda\mathcal{I}(\mu)} \inf v(\operatorname{supp}(p^*)) \\
&= \max_{p' \in \mathcal{I}(\mu'),\ p \in \mathcal{I}(\mu)} \min\left\{\inf v(\operatorname{supp}(p')),\ \inf v(\operatorname{supp}(p))\right\} \\
&= \min\left\{v^*(\mu'),\ v^*(\mu)\right\}.
\end{aligned}$$

Now, the correspondence $\operatorname{supp} : \Delta\Delta\Theta \rightrightarrows \Delta\Theta$ is lower hemicontinuous (Theorem 17.14, Aliprantis and Border, 2006). Because v is upper semicontinuous, it follows (Lemma 17.29, Aliprantis and Border, 2006) that $\tilde{v}^*$ is upper semicontinuous. Next, the correspondence $\mathcal{I} : \Delta\Theta \rightrightarrows \Delta\Delta\Theta$ is upper hemicontinuous because the barycentre map $(p \mapsto \int_{\Delta\Theta} \mu \, \mathrm{d}p(\mu))$ is continuous (Proposition 1.1, Phelps, 2001), yielding upper semicontinuity of v^* (Lemma 17.30, Aliprantis and Border, 2006). So v^* is quasiconcave, upper semicontinuous, and above v; therefore, $v^* \geq \bar{v}$.

Finally, because $\tilde{v}^*$ is upper semicontinuous and $\mathcal{I}(\mu_0)$ compact, some $p \in \mathcal{I}(\mu_0)$ exists such that $\tilde{v}^*(p) = v^*(\mu_0)$. So $D := \operatorname{supp} p$ has $\inf v(D) = v^*(\mu_0)$. Because $\mu_0 \in \overline{\operatorname{co}}D$, and $\bar{v}$ is upper semicontinuous, quasiconcave, and above v:

$$\bar{v}(\mu_0) \geq \inf \bar{v}(\overline{\operatorname{co}}D) = \inf \bar{v}(\operatorname{co}D) = \inf \bar{v}(D) \geq \inf v(D) = v^*(\mu_0).$$

So $v^*(\mu) = \bar{v}(\mu)$. $\square$

A.4.3 Corollary 3: Commitment is usually valuable

We now establish Corollary 3, which says that almost every prior μ_0 has either $\bar{v}(\mu_0) = s^{\mathrm{FB}} := \max v(\Delta\Theta)$ or $\hat{v}(\mu_0) > \bar{v}(\mu_0)$.

Proof. First, because v has a (finite) range contained in $u_S(A)$, so does its quasiconcave envelope $\bar{v}$. Next, that $\bar{v}$ is quasiconcave and implies $\{\bar{v} \geq s\}$ is convex for every $s \in u_S(A)$. Let

$$D := (\Delta\Theta)^\circ \setminus \bigcup_{s \in u_S(A)} \partial \{\bar{v} \geq s\},$$

the set of full-support beliefs that are not on the boundary of any $\bar{v}$-upper contour set. Being the boundary of a bounded convex set in a $(|\Theta|-1)$-dimensional space, the set $\partial \{\bar{v} \geq s\}$ is a manifold of dimension strictly lower than $|\Theta|-1$ for each $s \in u_S(A)$, as is the set $\partial\Delta\Theta$. Therefore, the finite union $\Delta\Theta \setminus D$ is low dimensional as well. Being low dimensional, $\Delta\Theta \setminus D$ is Lebesgue-null and nowhere dense.

Suppose $\mu_0 \in D$, and fix some belief $\mu \in \Delta\Theta$ such that some action in $u_S^{-1}(s^{\mathrm{FB}})$ is a best response for R to belief μ. By definition of D, sufficiently small $\epsilon \in (0,1]$ will have $\epsilon\mu \leq \mu_0$ and $\bar{v}\left(\frac{\mu_0 - \epsilon\mu}{1-\epsilon}\right) \geq \bar{v}(\mu_0)$. But then, $\hat{v}$ being concave and lying above $\bar{v}$,

$$\hat{v}(\mu_0) \geq (1-\epsilon)\bar{v}\left(\tfrac{\mu_0 - \epsilon\mu}{1-\epsilon}\right) + \epsilon\bar{v}(\mu) \geq (1-\epsilon)\bar{v}(\mu_0) + \epsilon s^{\mathrm{FB}}.$$

So $\bar{v}(\mu_0) < \hat{v}(\mu_0)$ if $\bar{v}(\mu_0) < s^{\mathrm{FB}}$. □

B Omitted Proofs: Applications

B.1 Proofs for Section 5.1: The Think Tank

In this example: $A = \{0, \ldots, n\}$, $\Theta = [0,1]^n$, μ_0 is exchangeable, u_S is increasing with $u_S(0) = 0$, and

$$u_R(a, \theta) = \begin{cases} \theta_i - c & : \ a = i \in \{1, \ldots, n\} \\ 0 & : \ a = 0. \end{cases}$$

We now invest in some notation. For $\theta \in \Theta$ and $k \in \{1, \ldots, n\}$, let $\theta_{k,n}^{(1)} := \max_{i \in \{k, \ldots, n\}} \theta_i$ be the first-order statistic among reforms better (for S) than k. For finite $\hat{M} \subseteq M$, let $\mathcal{U}(\hat{M}) \in \Delta(\hat{M}) \subseteq \Delta M$ be the uniform measure over $\hat{M}$. For

$k \in \{1, \ldots, n\}$, let

$$\begin{aligned} \sigma_k : \Theta &\to \Delta\{k, \ldots, n\} \subseteq \Delta M \\ \theta &\mapsto \mathcal{U}\left(\arg\max_{i \in \{k, \ldots, n\}} \theta_i\right) \end{aligned}$$

be the S strategy that reports the best reform from among those the think tank prefers to k; let $\beta_k : M \to \Delta\Theta$ be some belief map such that σ_k and β_k are together Bayes consistent; and let $p_k \in \mathcal{I}(\mu_0)$ be the associated information policy. For any measurable $f : \Theta \to [0, 1]$, let $\mathbb{E}_0 f(\theta) := \int f \, \mathrm{d}\mu_0$; and for $k \in \{1, \ldots, n\}$ and $i \in \{k, \ldots, n\}$, let $\mathbb{E}_i^k f(\theta) := \int f \, \mathrm{d}\beta_k(\cdot|i)$. Finally, for any $k \in \{1, \ldots, n\}$, let $\hat{\theta}^k := \mathbb{E}_0 \theta_{k,n}^{(1)}$.

B.1.1 Claim 1: Ranking the best reforms

Toward the proof of Claim 1, we first show the following.

Claim. Fix $k \in \{1, \ldots, n\}$ and $i \in \{k, \ldots, n\}$. Then, $i \in \arg\max_{a \in A} u_R(a, \beta_k(i))$ if and only if $\hat{\theta}^k \geq c$.

Proof. For a given $i \in \{k, \ldots, n\}$, exchangeability of μ_0 implies the following four facts:

(1) $\mathbb{E}_0\theta_i = \mathbb{E}_0\theta_j = \mathbb{E}_i^k\theta_j$ for $j \in \{1, \ldots, k-1\}$.
(2) $\mathbb{E}_0\theta_i \in \text{co}\left\{\mathbb{E}_i^k\theta_i, \ \mathbb{E}_i^k\theta_j\right\}$ for $j \in \{k, \ldots, n\} \setminus \{i\}$.
(3) $\mathbb{E}_i^k\theta_i \geq \mathbb{E}_0\theta_i$.
(4) $\mathbb{E}_i^k\theta_i = \hat{\theta}^k$.

The first three facts collectively tell us $\mathbb{E}_i^k\theta_i \geq \mathbb{E}_i^k\theta_j$ for $j \in \{1, \ldots, n\} \setminus \{i\}$. As an implication, $i \in \arg\max_{a \in A} u_R(a, \beta_k(i))$ if and only if $\mathbb{E}_i^k\theta_i \geq c$. The fourth fact completes the proof of the claim. □

Proof of Claim 1. Now, we prove the three-way equivalence of Claim 1. First, that Part 2 implies Part 3 follows from the above claim. Next, that Part 3 implies Part 1 follows directly from Theorem 1. Now, to show Part 1 implies Part 2, consider any equilibrium yielding S value $u_S(k)$. In this equilibrium, every on-path message yields value $u_S(k)$ to S, implying some reform from $\{k, \ldots, n\}$ is incentive compatible for R. That is, R has an optimal strategy in which his gross benefit is one of $\{\theta_i\}_{i=k}^n$ almost surely. But then R's ex-ante payoff is no greater than the prior expectation of $\max_{i \in \{k, \ldots, n\}} \theta_i - c$. This is then nonnegative by R's incentives: He does not want to deviate to the status quo ex ante. Thus, Part 1 implies Part 2, completing the proof of Claim 1. $\square$

B.1.2 Additional calculations

Finally, Corollary 1 tells us the sender's best equilibrium value lies in $\{0, \ldots, n\}$, so that the S-optimal equilibrium payoff is $u_S(k^*)$, where

$$k^* = \begin{cases} \max\left\{k \in \{1, \ldots, n\} : \hat{\theta}^k \geq c\right\} & : \hat{\theta}^1 \geq c \\ 0 & : \hat{\theta}^1 < c. \end{cases}$$

As described in section 5.1, we can use the constructive proof of Theorem 1 to explicitly derive the modification of p_{k^*} that supports payoff $u_S(k^*)$ as an equilibrium payoff when $k^* > 0$. Let $\epsilon := \frac{\hat{\theta}^{k^*} - c}{\hat{\theta}^{k^*} - \hat{\theta}^n}$, and consider the truth-or-noise signal $\sigma^* := (1-\epsilon)\sigma_{k^*} + \epsilon\mathcal{U}\{k^*, \ldots, n\}$. That is, among the proposals that the think tank weakly prefers to k^*, it either reports the best (with probability $1-\epsilon$, independent of the state) or a random one. Following a recommendation $i \in \{k, \ldots, n\}$, the lawmaker is indifferent between reform i and no reform at all. He responds with $\rho(i|i) = \frac{u_S(k^*)}{u_S(i)}$ and $\rho(0|i) = 1 - \rho(i|i)$. The proof of Lemma 4 shows such play is in fact equilibrium play.

B.2 Proofs for Section 5.2: The Broker

B.2.1 The one-dimensional model

In this section, we look at a one-dimensional version of our model, which generalizes Example 2, analyzed in section 5.2. Our task is to prove a generalization of Claim 2 for general priors.

Suppose $\Theta \subseteq \mathbb{R}$ and that some $v_M : \text{co}\Theta \to \mathbb{R}$ exists such that $v = v_M \circ E$, where $E : \Delta\Theta \to \text{co}\Theta$ maps each belief to its associated expectation of the state. This setting, which we call the **one-dimensional model**, was studied in Gentzkow and Kamenica (2016) and Dworczak and Martini (2018) under sender commitment power.

An important concept to simplify analysis of the one-dimensional model is the notion of a cutoff policy. Given $q \in [0, 1]$, the q-**quantile-cutoff policy** is the (necessarily unique) information policy $p^q \in \mathcal{I}(\mu_0)$ of the form $p^q = q\delta_{\mu_-^q} + (1-q)\delta_{\mu_+^q}$, for $\mu_-^q, \mu_+^q \in \Delta\Theta$ with $\max \text{supp}(\mu_-^q) \leq \min \text{supp}(\mu_+^q)$; and let $\theta_-^q := E\mu_-^q$ and $\theta_+^q := E\mu_+^q$. Say $p \in \mathcal{I}(\mu_0)$ is a **cutoff policy** if it is the q-quantile-cutoff policy for some $q \in [0,1]$. The following alternative characterization of cutoff policies, which is immediate, is useful for analyzing the one-dimensional model.

Fact 1. *For $q \in [0,1]$, the belief μ_-^q (μ_+^q) is the unique solution to the program $\min_{\mu \in \Delta\Theta:\ q\mu \leq \mu_0} E\mu$ ($\max_{\mu \in \Delta\Theta:\ (1-q)\mu \leq \mu_0} E\mu$).*

The q-cutoff policy reports whether the state is in the bottom q quantiles or the top $1-q$ quantiles, as measured according to the prior. More concretely, the sender simply reports whether the state is above or below some well-calibrated cutoff.[33] As the following claim shows, securability enables us to use cutoff policies to analyze many one-dimensional applications of interest, including the broker example.

Claim 5. *Suppose $\Theta \subseteq \mathbb{R}$ and $v_M : \text{co}\Theta \to \mathbb{R}$ is a weakly quasiconvex function such that $v = v_M \circ E$. Then the following are equivalent for all $s \geq v(\mu_0)$.*

1. *S can attain a payoff s in equilibrium.*

2. *The payoff s is securable by a cutoff policy.*

Moreover, an S-preferred equilibrium outcome (p, s) exists such that p is a cutoff policy.

Proof. Without loss, say $\text{co}\Theta = [0,1]$ and let $\theta_0 = E\mu_0$ be the prior mean. Because v_M is quasiconvex, v_M is either nonincreasing on $[0, \theta_0]$ or nondecreasing on $[\theta_0, 1]$. Suppose the latter holds without loss. Because uninformative communication is a cutoff policy with cutoff quantile 0 or 1, the result is immediate if $s = v(\mu_0)$, so we may assume $s > v(\mu_0)$.

That 2 implies 1 follows directly from Theorem 1. Now we suppose 1 holds and show 2 does as well. The nonempty (because s is securable) compact sets $\Theta_L := \{\theta \in [0, \theta_0] : \ v(\theta) \geq s\}$ and $\Theta_R := \{\theta \in [\theta_0, 1] : \ v(\theta) \geq s\}$ both exclude θ_0 because $s > v(\mu_0)$. Let $\theta_L := \max \Theta_L$ and $\theta_R := \min \Theta_R$. By Theorem 1 and Lemma 4, a Bayes-plausible information policy p exists that barely secures s, which then implies $p \circ E^{-1} \{\theta_L, \theta_R\} = 1$. That is, some $\hat{q} \in (0,1)$, $p_L \in \Delta E^{-1}(\theta_L)$, $p_R \in \Delta E^{-1}(\theta_R)$ exist

[33] This is correct as stated in the case in which μ_0 is atomless; if the cutoff is itself a state with positive prior probability, then S's message may need to be random conditional on the cutoff state itself occuring.

such that $p = \hat{q}p_L + (1-\hat{q})p_R$. But then Fact 1 implies $\theta_-^{\hat{q}} \le \theta_L$ and $\theta_+^{\hat{q}} \ge \theta_R$. Because $\theta_-^{\hat{q}} \le \theta_L < \theta_0 = \theta_-^1$, the intermediate value theorem (and Berge's theorem, which tells us from Fact 1 that $\theta_-^{(\cdot)}$ is continuous) delivers some $q_2 \in [\hat{q}, 1)$ such that $\theta_-^{q_2} = \theta_L$. Similarly, some $q_1 \in (0, \hat{q}]$ exists such that $\theta_+^{q_1} = \theta_R$. Now, because $\theta_-^{q_2} = \theta_L$, $\theta_+^{q_2} \ge \theta_R$, and $v_M|_{[\theta_0,1]}$ is nondecreasing, it follows that p^{q_2} secures s.

To prove the "moreover" part, we specialize to the case in which $s = \bar{v}(\mu_0)$. Let $Q := [q_1, q_2]$, $Q_+ := \{q \in Q : \; v_M(\theta_+^q) = s\}$, and $Q_- := \{q \in Q : \; v_M(\theta_-^q) = s\}$. That no value strictly above s is securable implies $Q = Q_+ \cup Q_-$. Therefore, the union of the closures has the same property: $\bar{Q}_+ \cup \bar{Q}_- = Q$. Because Q is connected (because v_M is monotone on each side of $\hat{q}$), some $q \in \bar{Q}_+ \cap \bar{Q}_-$ must then exist. That V is upper hemicontinuous (together with Lemma 1) then implies the q-cutoff policy, paired with payoff s, is an equilibrium outcome. $\square$

Although not directly relevant to the broker example, we briefly note the previous proposition lends tractability to the one-dimensional model, even in cases where giving more information to R does not benefit S.

Corollary 5. *Suppose $\Theta \subseteq \mathbb{R}$ and $V_M : \text{co}\Theta \rightrightarrows \mathbb{R}$ is such that $V = V_M \circ E$. Then, for any equilibrium sender payoff s, an equilibrium outcome of the form (p, s) exists, such that p is a garbling of a cutoff policy (with at most two supported posterior beliefs).*

Proof. We have nothing to show for $s \in V(\mu_0)$. We now focus on the case of $s > v(\mu_0)$, the alternative case being symmetric. Normalize $\text{co}\Theta = [0, 1]$, and let $\theta_0 := E\mu_0$.

Define the correspondence $\tilde{V}_M : [0,1] \rightrightarrows \mathbb{R}$ by letting $\tilde{V}_M(\theta) := V(\text{co}\,\{\theta, \theta_0\})$ for every $\theta \in [0, 1]$. Appealing to Lemma 3, $\tilde{V}_M$ is a Kakutani correspondence, so that $\tilde{V} := \tilde{V}_M \circ E : \Delta\Theta \rightrightarrows \mathbb{R}$ is as well. We can therefore apply the mathematical results of Claim 5, letting $\tilde{v}_M := \max \tilde{V}_M$ (which is quasiconvex and minimized at θ_0) replace v_M to find a cutoff $q \in [0, 1]$ such that $\tilde{v}_M(\theta_-^q), \tilde{v}_M(\theta_+^q) \ge s$. But then, by definition of $\tilde{V}_M$, some two-message garbling p' of p^q exists which secures s in the original game, that is, has $p'\{v \ge s\} = 1$. Finally, Lemma 4 delivers a further two-message garbling p of p' such that (p, s) is an equilibrium outcome. $\square$

B.2.2 The investor's payoff (equation (1))

Suppose (p, s) is an equilibrium outcome of the broker example (Example 2), and let r be R's associated payoff. Then

$$
\begin{aligned}
& r + \mathbb{V}_{\theta\sim\mu_0}(\theta) \\
&= \int \left\{ \tfrac{1}{2}\int (a_0^2 - 2a_0\theta + \theta^2)\ \mathrm{d}\mu_0 - \tfrac{1}{2}\int \left[a^*(\mu)^2 - 2a^*(\mu)\theta + \theta^2\right]\ \mathrm{d}\mu - u_S(a^*(\mu)) \right\}\ \mathrm{d}p(\mu) \\
&= \int \left\{ \int \left(\tfrac{1}{2}\theta^2 - a_0\theta\right) [\mathrm{d}\mu_0(\theta) - \mathrm{d}\mu(\theta)] + [a^*(\mu) - a_0]\, E\mu + \tfrac{1}{2}\left[a_0^2 - a^*(\mu)^2\right] - s \right\}\ \mathrm{d}p(\mu) \\
&= 0 + \int [a^*(\mu) - a_0] \left\{ E\mu - \tfrac{1}{2}[a_0 + a^*(\mu)] \right\}\ \mathrm{d}p(\mu) - s \\
&= \int [a^*(\mu) - a_0] \left\{ [E\mu - a^*(\mu)] + \tfrac{1}{2}[a^*(\mu) - a_0] \right\}\ \mathrm{d}p(\mu) - s \\
&= \int \left(\tfrac{s}{\phi}\right) \left[\phi + \tfrac{1}{2}\left(\tfrac{s}{\phi}\right)\right]\ \mathrm{d}p(\mu) - s \\
&= \tfrac{1}{2\phi^2} s^2.
\end{aligned}
$$

B.3 Proofs for Section 5.3: The Salesperson

First, we complete the proof of Claim 3.

Proof. In the main text, we demonstrated that $\hat{G}(t_0^*) > G(t_0^*)$ is necessary for commitment to strictly benefit the seller. Now, we show it is sufficient, by adapting Proposition 3 of Kamenica and Gentzkow (2011) to the present atomless case. Suppose the strict inequality holds.

That $\hat{G}(t_0^*) > G(t_0^*)$ implies some $t_L \in [0, t_0^*)$, $t_R \in (t_0^*, 1]$, and $\lambda \in (0,1)$ exist such that $t_0^* = \lambda t_L + (1-\lambda)t_R$ and $G(t_0^*) < \lambda G(t_L) + (1-\lambda)G(t_R)$. Replacing t_L and t_R with $(1-\epsilon)t_k + \epsilon t_0^*$ for both $k \in \{L, R\}$ for sufficiently small $\epsilon \in (0,1)$, we may (appealing to continuity of G) assume $t_L, t_R \in (0,1)$. Now, given that G is continuous and strictly increasing, for sufficiently small $\eta \in (0,1)$, some $t_{L-}, t_{L+}, t_{R-}, t_{R+}$ exist such that

$$t_{L-} < t_L < t_{L+} < t_0^* < t_{R-} < t_R < t_{R+},$$

$G(t_{L+}) - G(t_{L-}) = \lambda\eta$, $G(t_{R+}) - G(t_{R-}) = (1-\lambda)\eta$, $\frac{1}{\lambda\eta}\int_{t_{L-}}^{t_{L+}} t\ \mathrm{d}G(t) = t_L$, and $\frac{1}{(1-\lambda)\eta}\int_{t_{R-}}^{t_{R+}} t\ \mathrm{d}G(t) = t_R$. Direct computation and an appeal to continuity of G then shows the seller gets a value strictly higher than $\bar{v}(\mu_0) = G(t_0^*)$ by telling the buyer which product is best, and by further revealing whether that project has value in (t_{L-}, t_{L+}), in (t_{R-}, t_{R+}), or in neither. $\square$

Now, en route to Claim 3, we prove the following slightly more general result about when a CDF G admitting a single-peaked continuous density g is concave at the prior mean. For this purpose, let

$$\begin{aligned} \varphi_G : [0,1] &\to \mathbb{R} \\ t &\mapsto G(t) - G(0) - tg(t) = -\int_0^t \tilde{t}\ \mathrm{d}g(\tilde{t}). \end{aligned}$$

Lemma 6. *Suppose G admits a continuous, weakly quasiconcave density g. Let $t_M :=$ $\min\left[\arg\max_{t\in[0,1]} g(t)\right]$. Then $\hat{G}(t_0^*) = G(t_0^*)$ if and only if $t_0^* \geq t_M$ and $\varphi_G(t_0^*) \geq 0$.*

Proof. First, we show $\varphi_G(t_0^*) \geq 0$ is necessary for there to be no commitment gap. To that end, suppose $\varphi_G(t_0^*) < 0$. Recall that full support of μ_0 implies $t_0^* \in (0,1)$. Then, letting $\epsilon \in (0, 1-t_0^*]$, we have

$$\begin{aligned} \tfrac{t_0^*+\epsilon}{\epsilon}\left[\hat{G}(t_0^*) - G(t_0^*)\right] &\geq \tfrac{t_0^*+\epsilon}{\epsilon}\left[\tfrac{t_0^*}{t_0^*+\epsilon}G(t_0^*+\epsilon) + \tfrac{\epsilon}{t_0^*+\epsilon}G(0) - G({t^*}_0)\right] \\ &= t_0^*\frac{G(t_0^*+\epsilon) - G(t_0^*)}{\epsilon} - \left[G(t_0^*) - G(0)\right], \end{aligned}$$

which tends to $-\varphi_G(t_0^*) > 0$, as $\epsilon \to 0$. Therefore, $\frac{t_0^*+\epsilon}{\epsilon}\left[\hat{G}(t_0^*) - G(t_0^*)\right] > 0$ when $\epsilon > 0$ is sufficiently small, so that $\hat{G}({t^*}_0) > G({t^*}_0)$.

Now, we verify that $t_0^* \geq t_M$ is necessary for there to be no commitment gap. Suppose $t_0^* < t_M$. Then $g|_{[0,t_M]}$ is continuous, weakly increasing, and nonconstant. Therefore, $G|_{[0,t_M]}$ is weakly convex and not affine, implying

$$\hat{G}(t_0^*) \geq \tfrac{t_M - t_0^*}{t_M}G(0) + \tfrac{t_0^*}{t_M}G(t_M) > G(t_0^*).$$

Conversely, suppose $t_0^* \geq t_M$ and $\varphi_G(t_0^*) \geq 0$. Observing that $\varphi_G(0) = 0$ and g is weakly inceasing on $[0, t_M]$, it follows that $\varphi_G(t_M) \leq 0$. But φ_G is continuous, and so the intermediate value theorem delivers some $t_* \in [t_M, t_0^*]$ with $\varphi_G(t_0^*) = 0$. Now define

$$\begin{aligned} G^* : [0,1] &\to \mathbb{R} \\ t &\to \begin{cases} G(0) + tg(t_0^*) & : \ t \le t_0^* \\ G(t) & : \ t \ge t_0^*. \end{cases} \end{aligned}$$

By definition of t_0^*, the function G^* is well defined and differentiable at t_0^*. But it is also affine on $[0, t_0^*]$; and G^* is continuous and weakly concave on $[t_*, 1]$ because (because $t_0^* \ge t_M$) g is decreasing on $[t_*, 1]$. Therefore, G^* is concave and continuous. We next show $G^* \ge G$ on $[0,1]$. First, it is on $[t_0^*, 1]$ by construction. Next, for $t \in [t_M, t_0^*]$, that $g|_{[t_M, t_0^*]}$ is weakly decreasing tells us

$$G^*(t) - G(t) = [G(t_0^*) - G(t)] - [G^*(t_0^*) - G^*(t)] = \int_t^{t_*} [g(\tilde{t}) - g(t^*)] \, \mathrm{d}\tilde{t} \ge 0.$$

Finally, because G^* is affine on $[0, t_M]$, G is convex on $[0, t_M]$, and $G^* \ge G$ on $\{0, t_M\}$, collectively implying $G^* \ge G$ on $[0, t_M]$. Therefore, $G^* \ge \hat{G} \ge G$. But $G^*(t_0^*) = G(t_0^*)$, so that $\hat{G}(t_0^*) = G(t_0^*)$. □

From this, we can prove Claim 3 easily.

Proof of Claim 3. First, suppose g is weakly decreasing. Then $t_0^* \ge 0 = t_M$ and $G(t_0^*) - G(0) = \int_0^{t_0^*} g(t) \, \mathrm{d}t \ge t_0^* g(t_0^*)$, and Lemma 6 applies.

Second, suppose g is nonconstant and weakly increasing. If $t_0^* < t_M$, then $\hat{G}(t_0^*) > G(t^*)$ by Lemma 6. If $t_0^* \ge t_M$, then $g(t_0^*) \ge g(t_M) > g(0)$, implying $g|_{[0,t_0^*]}$ is continuous, nonconstant, and weakly increasing. So this function is everywhere weakly below $g(t_0^*)$ and strictly below it for some nondegenerate interval. Therefore, $t_0^* g(t_0^*) > \int_0^{t_0^*} g(t) \, \mathrm{d}t = G(t_0^*) - G(0)$, and Lemma 6 applies.

Third, suppose g is strictly quasiconcave. For any $\tilde{t} \in (0, t_M]$, the function g is continuous and strictly increasing—in particular, continuous, weakly increasing, and nonconstant—on $[0, \tilde{t}]$. This tells us φ_G is nonconstant and weakly decreasing on $[0, \tilde{t}]$, implying $\varphi_G(\tilde{t}) < \varphi_G(0) = 0$. Therefore, if $\varphi_G(t_0^*) \ge 0$, then $t_0^* \ne \tilde{t}$. Because $\tilde{t} \in (0, t_M]$ was arbitrary, we now know that if $\varphi_G(t_0^*) \ge 0$, then $t_0^* \ge t_0^*$. The claim then follows from the Lemma 6. □

C Online Appendix

In this online appendix, we elaborate on the results mentioned in section 6 and discuss some additional relevant results.

C.1 Proof of Proposition 1: Informative Communication

We use the following abstract lemma to show that informative communication is possible in equilibrium whenever there are more than two states. The lemma formally generalizes the main theorem of Chakraborty and Harbaugh (2010), showing that whenever the state is multidimensional, some equilibrium exists in which different on-path messages lead R to different posterior expectations of the state.

Lemma 7. *Suppose $\Theta \subseteq \mathcal{X}$ for some locally convex space $\mathcal{X}$, and let $E : \Delta\Theta \to \overline{\text{co}}\Theta$ take every belief to its associated expectation of the state. If Θ is noncollinear, an equilibrium outcome (p, s) exists such that $p \circ E^{-1}$ is nondegenerate.*

The proof is based on Chakraborty and Harbaugh's (2010) insight of using Borsuk-Ulam to show the possibility of effective communication. We begin by representing the prior as an average of three posterior beliefs, μ_1, μ_2, and μ_3, such that the three induced posterior expectations of the state are noncollinear; one can always find such beliefs when Θ is itself multidimensional. Next, we find a circle of beliefs around the prior within the convex hull of $\{\mu_1, \mu_2, \mu_2\}$. By construction, each belief on said circle yields a different posterior expectation of the state. We then generalize the one-dimensional case of the Borsuk-Ulam theorem from continuous functions to Kakutani correspondences, yielding an antipodal pair of beliefs μ and μ' on the circle such that $V(\mu) \cap V(\mu')$ is nonempty. Therefore, we can split the prior across μ and μ' to obtain an equilibrium information policy.

Proof. Let $\theta_0 := E\mu_0$, the prior mean. Because Θ is noncollinear, distinct $\theta_1, \theta_2 \in \Theta$ exist such that $\theta_0 \notin \text{co}\{\theta_1, \theta_2\}$. Next, because μ_0 has full support, both $\mu_0(N_1), \mu_0(N_2) > 0$ for any open neighborhoods N_1 of θ_1 and N_2 of θ_2. We can then define the conditional distribution $\mu_i(\cdot) := \frac{\mu_0(N_i \cap (\cdot))}{\mu_0(N_i)}$ for $i \in \{1, 2\}$. Letting N_1, N_2 be sufficiently small neighborhoods, we may assume $N_1 \cap N_2 = \emptyset$ and $\theta_0 \notin \text{co}\{E\mu_1, E\mu_2\}$. Therefore, letting $\mu_3(\cdot) := \frac{\mu_0((\cdot) \setminus (N_1 \cup N_2))}{1 - \mu_0(N_1 \cup N_2)}$, we know that $\mu_0 \in \text{co}\{\mu_1, \mu_2, \mu_3\}$, that μ_0 is not in the convex hull any two of $\{\mu_1, \mu_2, \mu_3\}$, and that the three points $\{E\mu_1, E\mu_2, E\mu_3\}$

are affinely independent in $\mathcal{X}$. So $\mu_0 = \sum_{i=1}^{3} \lambda_i \mu_i$ for some $\mu_1, \mu_2, \mu_3 \in \Delta\Theta$ and $\lambda_1, \lambda_2, \lambda_3 > 0$. Therefore, for small enough $\epsilon > 0$, we can embed the circle

$$\mathbb{S} = \left\{(x, y) \in \mathbb{R}^2 : \ x^2 + y^2 = 1\right\}$$

into the space of beliefs via the continuous map

$$\begin{aligned} &\varphi : \mathbb{S} \to \Delta\Theta \\ &(x, y) \mapsto (\lambda_1 + \epsilon x)\mu_1 + (\lambda_2 + \epsilon y)\mu_2 + [\lambda_3 - \epsilon(x + y)]\mu_3, \end{aligned}$$

and affine independence of $E\mu_1, E\mu_2, E\mu_3$ then implies $E \circ \varphi$ is injective. Next, define the function

$$\begin{aligned} f \ &: \mathbb{S} \to \mathbb{R} \\ &z \mapsto \max V(\varphi(z)) - \min V(\varphi(-z)). \end{aligned}$$

Two properties of f are immediate. First, f is upper semicontinuous because V is upper hemicontinuous. Second, any $z \in \mathbb{S}$ satisfies $f(z) + f(-z) \geq 0$ because $\max V \geq \min V$.

We use the above two properties to show some $z \in \mathbb{S}$ exists for which $f(z)$ and $f(-z)$ are both nonnegative. Assume otherwise, for a contradiction. Then every $z \in \mathbb{S}$ has the property that one of $f(z), f(-z)$ is strictly negative and (as $f(z) + f(-z) \geq 0$) the other is strictly positive; in particular, f is globally nonzero. Because $\mathbb{S}$ is connected, some $z \in \mathbb{S}$ exists that is a limit point of both $f^{-1}((0, \infty))$ and $f^{-1}((-\infty, 0))$. But then $-z$ shares this same property, because f takes opposite points to opposite signs. Finally, upper semicontinuity of f tells us both $f(z)$ and $f(-z)$ are nonnegative, a contradiction.

Now we have $z \in \mathbb{S}$ with $f(z), f(-z) \geq 0$. That is, $\max V(\varphi(z)) \geq \min V(\varphi(-z))$ and $\max V(\varphi(-z)) \geq \min V(\varphi(z))$. Said differently (recall that V is convex-valued), $V(\varphi(z)) \cap V(\varphi(-z)) \neq \emptyset$. Lemma 1 then guarantees the existence of an equilibrium that generates information policy $p = \frac{1}{2}\delta_{\varphi(z)} + \frac{1}{2}\delta_{\varphi(-z)}$. In particular, $p \circ E^{-1}$ is nondegenerate. $\square$

Armed with this lemma, we now prove Proposition 1.

Proof. First, suppose $|\Theta| > 2$. Then, letting $\mathcal{X} := \mathrm{ca}(\Theta)$, we can view Θ as a subset of $\mathcal{X}$ under the inclusion $\theta \mapsto \delta_\theta$. That $|\Theta| > 2$ then implies Θ is noncollinear, and Lemma 7 delivers an informative equilibrium.

Now focus on the case of $|\Theta| = 2$. To conserve notation, assume without loss that $\Theta = \{0, 1\}$, and identify $\Delta\{0, 1\}$ with $[0, 1]$ in the obvious way.

Suppose now that neither $\theta \in \Theta$ satisfies the property described in the second part of the theorem. Then $(\mu, \mu') \in [0, \mu_0) \times (\mu_0, 1]$ are neither guaranteed to satisfy $V(\mu) \ll V(\mu')$ nor guaranteed to satisfy $V(\mu) \gg V(\mu')$. Said differently, some $\underline{\mu}, \underline{\nu} \in [0, \mu_0)$ and $\bar{\mu}, \bar{\nu} \in (\mu_0, 1]$ exist such that $\max V(\underline{\mu}) \geq \min V(\bar{\mu})$ and $\min V(\underline{\nu}) \leq \max V(\bar{\nu})$. Define $\varphi : [0, 1] \to [0, 1]^2$ via $\varphi(\lambda) := (1 - \lambda)(\underline{\mu}, \bar{\mu}) + \lambda(\underline{\nu}, \bar{\nu})$, and define $G : [0, 1] \rightrightarrows \mathbb{R}$ via $G(\lambda) := V(\varphi_2(\lambda)) - V(\varphi_1(\lambda))$. Notice G is a Kakutani correspondence because V is. By assumption, $G(0)$ contains a nonpositive number and $G(1)$ contains a nonnegative number. By Lemma 3, then, $G(\lambda) \ni 0$ for some $\lambda \in [0, 1]$. Therefore, by Lemma 1, an equilibrium exists that has support $\{\varphi_1(\lambda), \varphi_2(\lambda)\}$ and is then informative.

Conversely, suppose $\theta \in \Theta$ is as described in the second part of the theorem; without loss, say $\theta = 1$. Any $p \in \mathcal{I}(\mu_0) \setminus \{\delta_{\mu_0}\}$ has some $\underline{\mu} \in [0, \mu_0)$ and some $\bar{\mu} \in (\mu_0, 1]$ in its support. By hypothesis, $V(\underline{\mu}) \cap V(\bar{\mu}) = \emptyset$. Lemma 1 then tells us it cannot be an equilibrium policy. This completes the proof. $\square$

C.2 The Equilibrium Payoff Set

In this subsection, we briefly comment on how our tools, and the belief-based approach more broadly, can generate a more complete picture of the world of cheap talk with state-independent S preferences. As will be clear, the results outlined herein are all very straightforward to derive given earlier results in the paper.

C.2.1 Other sender payoffs

Following the recent literature on communication with S commitment, our focus has largely been on high equilibrium S values, that is, those providing payoffs at least as high as those attainable under uninformative communication. Without a specific story of how S might successfully coordinate R toward her preferred equilibrium, this focus may appear misplaced. Luckily, the tools developed in our paper work equally well to characterize bad sender payoffs. Indeed, the proof of Lemma 1 used no special features of V other than it being a Kakutani correspondence, which $-V$ is as well.

Therefore, our game has the same equilibria as the game with S objective $-u_S$. To deliver the mirror-image versions of our main results, define the value function from S-adversarial tiebreaking, $w := \min V : \Delta\Theta \to \mathbb{R}$.

Theorem 1 implies a sender payoff $s \leq w(\mu_0)$ is an equilibrium payoff if and only if some $p \in \mathcal{I}(\mu_0)$ exists such that $p\{w \leq s\} = 1$. Combining this observation with the original statement of the securability theorem tells us $s \in \mathbb{R}$ is an equilibrium S payoff if and only if $p_+, p_- \in \mathcal{I}(\mu_0)$ exist such that $p_+\{v \geq s\} = p_-\{w \leq s\} = 1$. An easy consequence is that the equilibrium S payoff set is convex, which we document in Corollary 4. Corollary 1 has a mirror image as well, telling us the set of S equilibrium payoffs is exactly

$$\left[\min_{p\in\mathcal{I}(\mu_0)} \sup w\left(\operatorname{supp} p\right),\ \max_{p\in\mathcal{I}(\mu_0)} \inf v\left(\operatorname{supp} p\right)\right].$$

Note that convexity of the set of attainable S payoffs is special to the case in which S's payoffs are state independent; indeed, the leading example of Crawford and Sobel (1982) does not share this feature.

The mirrored counterpart of our geometric Theorem 2 is that the lowest S payoff attainable in equilibrium is $\underline{w}(\mu_0)$, where $\underline{w}$ is the quasiconvex envelope of w, that is, the pointwise highest quasiconvex and lower semicontinuous function that minorizes w. Therefore, if we do not want to take a stance on equilibrium selection in the cheap-talk model, we can geometrically characterize the value of commitment as lying anywhere between $\hat{v}(\mu_0) - \bar{v}(\mu_0)$ and $\hat{v}(\mu_0) - \underline{w}(\mu_0)$.

C.2.2 Receiver payoffs

Our most powerful tools (the securability theorem and its descendants) pertain to S payoffs. However, the belief-based approach (i.e., Lemma 1) can be used to describe R payoffs as well. Indeed, let $v_R : \Delta\Theta \to \mathbb{R}$ be R's value function, given by $v_R(\mu) := \max_{a\in A} \int_\Theta u_R(a,\cdot)\ \mathrm{d}\mu$. It follows from R's interim rationality that any equilibrium that generates outcome (p, s) will deliver a payoff of $r = \int_{\Delta\Theta} v_R\ \mathrm{d}p$ to R.

Given equilibrium S payoff s, we can then more explicitly derive the set of equilibrium R payoffs compatible with an in which S gets payoff s. Let

$$B_s := \{w \leq s \leq v\} = \left\{\mu \in \Delta\Theta :\ \exists a_+, a_- \in \arg\max_{a\in A} \int_\Theta u_R(a,\cdot)\ \mathrm{d}\mu \text{ s.t. } u_S(a_-) \leq s \leq u_S(a_+)\right\}.$$

Then (s, r) is an equilibrium payoff profile if and only if $r = \int_{\Delta\Theta} v_R \, \mathrm{d}p$ for some $p \in \mathcal{I}(\mu_0) \cap \Delta(B_s)$. The best such R payoff (given s) is given by $\widehat{v^s_R}(\mu_0)$, where $v^s_R : B_s \to \mathbb{R}$ is the restriction of v_R and $\widehat{v^s_R} : \overline{\text{co}} B_s \to \mathbb{R}$ is the concave envelope of v^s_R.

C.2.3 Implementing equilibrium payoffs

In addition to their role in proving Theorem 1, barely securing policies generate a straightforward way of implementing any equilibrium payoff.[34] If S could commit, we could apply the revelation principle[35] to implement any S commitment payoff with a commitment protocol in which S makes a pure action recommendation to R, and R always complies. Using barely securing policies, we can show a similar result holds with cheap talk, with one important caveat: R must be allowed to mix. To state this result, for any S strategy σ, define $\mathcal{M}_\sigma$ as the set of messages in σ's support.[36]

Proposition 3. *Fix some S payoff s. Then the following are equivalent:*

1. *s is generated by an equilibrium.*
2. *s is generated by an equilibrium with $\mathcal{M}_\sigma \subseteq \Delta A$ and $\rho(\alpha) = \alpha \ \forall \alpha \in \mathcal{M}_\sigma$.*
3. *s is generated by an equilibrium with $\mathcal{M}_\sigma \subseteq A$ and $\rho(a|a) > 0 \ \forall a \in \mathcal{M}_\sigma$.*

The proposition suggests two ways in which one can implement a payoff of s via incentive-compatible recommendations. The first way has S giving R a *mixed* action recommendation that R *always* follows. The second way has S giving R a *pure* action recommendation that R *sometimes* follows. Both ways can result in R mixing.

That 1 implies 2 follows from standard revelation principle logic. To prove 1 implies 3,[37] we start with a minimally informative information policy that secures s. Because p is minimally informative, it must barely secure s, meaning (p, s) is an

[34] For S payoffs $s \leq \min V(\mu_0)$, we use the mirror image of barely securing policies, that is, information policies p such that $\{\min V(\cdot) \leq s\} \cap \text{co}\{\mu, \mu_0\} = \{\mu\}$ holds for every $\mu \in \text{supp}\, p$.

[35] See, for example, Myerson (1979), Myerson (1986), Kamenica and Gentzkow (2011), and Bergemann and Morris (2016).

[36] That is, let $\mathcal{M}_\sigma = \cup_{\theta \in \Theta} \text{supp}\, \sigma(\cdot|\theta)$.

[37] The equivalence between 1 and 3 echoes an important result of Bester and Strausz (2001), who study a mechanism design setting with one agent, finitely many types, and partial commitment by the principal. Applying a graph-theoretic argument, they show one can restrict attention to direct mechanisms in which the agent reports truthfully with positive probability. Although the proof techniques are quite different, a common lesson emerges. Agent mixing helps circumvent limited commitment by the principal: in Bester and Strausz's (2001) setting by limiting the principal's information, and in ours by limiting her control.

equilibrium. Let $\mathcal{E}$ be part 2's implementation of (p, s), and take $\mathbf{a}(\mu)$ to be some S-preferred action among all those that R plays in $\mathcal{E}$ at belief μ. By minimality of p, $\mathbf{a}(\cdot)$ must be p-essentially one-to-one, because pooling any posteriors that induce the same $\mathbf{a}(\cdot)$ value would yield an even less informative policy that secures s. Thus, $\mathbf{a}(\cdot)$ takes distinct beliefs to distinct (on-path) actions: R can infer μ from $\mathbf{a}(\mu)$. One can then conclude the proof by having S recommend $\mathbf{a}(\mu)$ and R respond to $\mathbf{a}(\mu)$ as he would have responded to μ under $\mathcal{E}$.

The formal proof is below.

Proof. Because (2) and (3) each immediately imply (1), we show the converses.

Suppose s is an equilibrium S payoff. Now take some $p \in \mathcal{I}(\mu_0)$ Blackwell-minimal among all policies securing payoff s, and let $D := \text{supp}(p) \subseteq \Delta\Theta$.[38] Lemma 4 then guarantees (p, s) is an equilibrium outcome, say, witnessed by equilibrium $\mathcal{E}_1 = (\sigma_1, \rho_1, \beta_1)$. Letting $\alpha = \alpha_s : D \to \Delta A$ be as delivered by Lemma 2, we may assume $\rho_1(\cdot|m) = \alpha(\cdot|\beta(m))$. In particular, ρ_1 specifies finite-support play for every message.

Let $\mathbb{M} := \text{marg}_M \mathbb{P}_{\mathcal{E}_1}$ and $X := \text{supp}\left[\mathbb{M} \circ \hat{\rho}^{-1}\right] \subseteq \Delta A$, and fix arbitrary $(\hat{\alpha}, \hat{\mu}) \in \text{supp}\left[\mathbb{M} \circ (\rho_1, \beta_1)^{-1}\right]$; in particular, $\hat{\alpha} \in X$. By continuity of u_R and receiver incentive compatibility, $\hat{\alpha} \in \arg\max_{\alpha \in \Delta A} u_R(\alpha \otimes \hat{\mu})$. Defining $\rho' : M \to \Delta A$ (resp. $\beta' : M \to \Delta\Theta$) to agree with ρ_1 (β_1) on path and take value $\hat{\alpha}$ ($\hat{\mu}$) off path, an equilibrium $\mathcal{E}' = (\sigma_1, \rho', \beta')$ exists such that $\mathbb{P}_{\mathcal{E}'} = \mathbb{P}_{\mathcal{E}_1}$ and $\rho'(\cdot|m) \in X$ for every $m \in M$.

Now define

$$\begin{aligned}
\sigma_2 : \Theta &\to \Delta X \subseteq \Delta M \\
\theta &\mapsto \sigma_1(\cdot|\theta) \circ \rho'^{-1} \\
\rho_2 : M &\to X \subseteq \Delta A \\
m &\mapsto \begin{cases} m & : m \in X \\ \hat{\alpha} & : m \notin X \end{cases} \\
\beta_2 : M &\to \Delta\Theta \\
m &\mapsto \begin{cases} \mathbb{E}_{m \sim \mathbb{M}}\left[\beta(m) \middle| \rho(m)\right] & : m \in X \\ \hat{\mu} & : m \notin X. \end{cases}
\end{aligned}$$

[38]Some policy secures s if s is an equilibrium payoff. The set of such policies is closed (and so compact) because v is upper semicontinuous. Therefore, because the Blackwell order is closed-continuous, a Blackwell-minimal such policy exists.

By construction, $(\sigma_2, \rho_2, \beta_2)$ is an equilibrium that generates outcome (p, s),[39] proving (1) implies (2).

Now define the (A- and D-valued, respectively) random variables $\mathbf{a}, \boldsymbol{\mu}$ on $\langle D, p\rangle$ by letting $\mathbf{a}(\mu) := \max \operatorname{supp}[\alpha(\mu)]$ and $\boldsymbol{\mu}(\mu) := \mu$ for $\mu \in D$. Next define the conditional expectation $\mathbf{f} := \mathbb{E}_p[\boldsymbol{\mu}|\mathbf{a}] : D \to D$, which is defined only up to a.e.-p equivalence. By construction, the distribution of $\boldsymbol{\mu}$ is a mean-preserving spread of the distribution of $\mathbf{f}$. That is, p is weakly more informative than $p \circ \mathbf{f}^{-1}$. By hypothesis, $\mathbf{a}(\mu)$ is incentive compatible for R at every $\mu \in D$. But $D = \operatorname{supp}(p \circ \mathbf{f}^{-1})$, which implies $p \circ \mathbf{f}^{-1}$ secures s. But then minimality of p implies $p \circ \mathbf{f}^{-1} = p$. So $\mathbf{f} = \mathbb{E}_p[\boldsymbol{\mu}|\mathbf{a}]$ and $\boldsymbol{\mu}$ have the same distribution, which implies $\mathbf{f} = \boldsymbol{\mu}$ a.s.-p. By definition, $\mathbf{f}$ is $\mathbf{a}$-measurable, so that Doob-Dynkin delivers some measurable $\mathbf{b} : A \to D$ such that $\mathbf{f} = \mathbf{b} \circ \mathbf{a}$.

Summing up, we have some measurable $\mathbf{b} : A \to D$ such that $\mathbf{b} \circ \mathbf{a} =_{\text{a.e.}-p} \boldsymbol{\mu}$. Now define

$$
\begin{aligned}
\sigma_3 : \Theta &\to \Delta A \subseteq \Delta M \\
\theta &\mapsto \sigma_2(\cdot|\theta) \circ (\mathbf{a} \circ \beta_2)^{-1} \\
\rho_3 : M &\to X \subseteq \Delta A \\
m &\mapsto \begin{cases} \alpha(\mathbf{b}(m)) & : \ m \in A \\ \hat{\alpha} & : \ m \notin A \end{cases} \\
\beta_3 : M &\to \Delta\Theta \\
m &\mapsto \begin{cases} \mathbf{b}(m) & : \ m \in A \\ \hat{\mu} & : \ m \notin A. \end{cases}
\end{aligned}
$$

By construction, $(\sigma_3, \rho_3, \beta_3)$ is an equilibrium that generates outcome (p, s), proving (1) implies (3). $\square$

Proposition 3 shows some forms of communication are without loss as far as S payoffs are concerned. First, any S equilibrium payoff is attainable in an equilibrium in which S recommends mixed actions that are (on path) followed exactly. This equivalence extends to equilibrium payoff *pairs*, with the same argument: Pooling messages that lead to the same R behavior relaxes incentive constraints and generates the same joint distribution over actions and states, preserving payoffs. Second, any S equilibrium payoff is attainable in an equilibrium in which S recommends pure actions

[39] It generates $(\tilde{p}, s)$ for some garbling $\tilde{p}$ of p. Minimality of p then implies $\tilde{p} = p$.

which are followed with positive probability. Whether this result holds in general for payoff pairs is an open question. It is easy to see why, at least, our argument does not go through as stated. The proof begins by considering an information policy that gives no "extraneous" information to R, subject to securing the relevant S value. But taking information away from R in this way can result in a payoff loss.

Still, we can leverage Lemma 1 to show a result of a similar spirit: To implement an equilibrium payoff profile, it is sufficient for R to only use binary mixed actions, the support of which is S's message.

Proposition 4. *Fix some payoff profile* (s, r)*. Then, the following are equivalent:*

1. (s, r) *is generated by an equilibrium.*
2. (s, r) *is generated by an equilibrium with* $\mathcal{M}_\sigma \subseteq \Delta A$ *and* $\rho(\alpha) = \alpha \ \forall \alpha \in \mathcal{M}_\sigma$*.*
3. (s, r) *is generated by an equilibrium with* $\mathcal{M}_\sigma \subseteq \left\{\frac{1}{2}\delta_a + \frac{1}{2}\delta_{a'} : \ a, \hat{a} \in A\right\}$ *and* $\text{supp}[\rho(\alpha)] = \text{supp}(\alpha) \ \forall \alpha \in \mathcal{M}_\sigma$*.*

We can interpret 3 as describing equilibria in which S tells R, "Play a or $\hat{a}$," for some pair of actions, but does not suggest mixing probabilities.

To see the equivalence between 1 and 3, Lemma 2 from the appendix can be used to show equilibrium payoff profile (s, r) can be implemented with an equilibrium in which R only ever uses pure actions or binary-support mixtures, with the latter only being used when S is not indifferent between the two supported actions. Without loss, say such equilibrium is as in 2, with S suggesting an incentive-compatible mixture to R. But then S rationality implies no two on-path recommendations can have the same support, because then S would have an incentive to deviate to the one putting a higher probability on the preferred action. Therefore, the same behavior could be induced by having every message replaced with a uniform distribution over its (at most binary) support, and the result follows.

With finitely many actions, Proposition 4 yields an a priori upper bound on the number of distinct messages required in equilibrium, similar to Proposition 3. Still, the upper bound of Proposition 3 is significantly smaller: Whereas Proposition 3 says no more than $n := |A|$ messages are required to span the set of equilibrium S values, Proposition 4 guarantees that any equilibrium payoff pair can be attained with at most $\frac{n(n-1)}{2}$ messages.

C.3 Proof of Proposition 2: State-Independent Preferences

Toward proving Proposition 2, we introduce some notation. For each continuous $f : \Delta\Theta \to \mathbb{R}$ and $X \subseteq \Delta A$, define the induced preferences $\succsim_{f,X}$ on X via $\alpha \succsim \alpha' \iff f(\alpha) \geq f(\alpha')$. The following lemma is the heart of the proposition.

Lemma 8. *Fix the action space, state space, and R preferences. Then the game with S objective $\tilde{u}_S$ and the game with S objective $u_S := \tilde{u}_S(\cdot, \mu_0)$ have the same equilibria, and generate the same equilibrium outcomes, if the preferences $\left\{\succsim_{\tilde{u}_S(\cdot,\theta),X}\right\}_{\theta\in\Theta}$ all coincide, where $X := \bigcup_{\mu\in\Delta\Theta} \arg\max_{\alpha\in\Delta A} \tilde{u}_S(\alpha, \mu)$.*

Notice the lemma applies directly to the two cases of Proposition 2. With cardinally state-independent preferences, $\succsim_{\tilde{u}_S(\cdot,\theta),\Delta A}$ is the same regardless of $\theta \in \Theta$. With ordinally state-independent preferences and unique R best response, one has that $\succsim_{\tilde{u}_S(\cdot,\theta),\{\delta_a\}_{a\in A}}$ does not depend on θ, and that $X \subseteq \{\delta_a\}_{a\in A}$. We now prove the lemma and, in doing so, complete the proof of the proposition.

Proof. First, consider any $\alpha, \alpha' \in X$. Because $\{\tilde{u}_S(\alpha, \theta) - \tilde{u}_S(\alpha', \theta)\}_{\theta\in\Theta}$ are all of the same sign and integration is monotone, $u_S(\alpha) - u_S(\alpha')$ is of the same sign. That is, $\left\{\succsim_{\tilde{u}_S(\cdot,\theta),X}\right\}_{\theta\in\Theta} = \{\succsim_{u_S,X}\}$.

Now take any (σ, ρ, β) satisfying the Bayesian condition and R incentive compatibility (the first and second part of the definition of equilibrium, respectively). R incentive compatibility implies $\rho(M) \subseteq X$. Therefore, $\left\{\succsim_{\tilde{u}_S(\cdot,\theta),\rho(M)}\right\}_{\theta\in\Theta} = \left\{\succsim_{u_S,\rho(M)}\right\}$. It follows that S incentive-compatibility (given ρ) is the same under $\tilde{u}_S$ and under u_S. Therefore, the set of equilibrium triples (σ, ρ, β) is the same. All that remains, then, is to show that an equilibrium (σ, ρ, β) generates the same S payoff under both models. To that end, fix some $\alpha^* \in \arg\max_{\alpha\in\rho(M)} u_S(\alpha)$, which is nonempty by S incentive compatibility. Then

$$\begin{aligned}
\int_\Theta \int_M \tilde{u}_S(\rho(m), \theta) \, \mathrm{d}\sigma(m|\theta) \, \mathrm{d}\mu_0(\theta) &= \int_\Theta \int_M \tilde{u}_S(\alpha^*, \theta) \, \mathrm{d}\sigma(m|\theta) \, \mathrm{d}\mu_0(\theta) \\
&= \int_\Theta \tilde{u}_S(\alpha^*, \theta) \, \mathrm{d}\mu_0(\theta) \\
&= u_S(\alpha^*) \\
&= \int_\Theta \int_M u_S(\alpha^*) \, \mathrm{d}\sigma(m|\theta) \, \mathrm{d}\mu_0(\theta) \\
&= \int_\Theta \int_M u_S(\rho(m)) \, \mathrm{d}\sigma(m|\theta) \, \mathrm{d}\mu_0(\theta),
\end{aligned}$$

where the first and last equality follow from S incentive compatibility. Therefore, S's ex-ante payoff is the same under equilibrium (σ, ρ, β) in both models. $\square$

The two parts of Proposition 2 apply directly, noting $X \subseteq \{\delta_a\}_{a \in A}$ when R has a unique best response to every belief.

C.4 Long Cheap Talk

Let us define the long cheap-talk game. In addition to the objects in our model section, R has some message space $\tilde{M}$, which we assume is compact metrizable. Let $\mathcal{H}_{<\infty} := \bigsqcup_{t=0}^{\infty}(M \times \tilde{M})^t$, $\mathcal{H}_\infty := (M \times \tilde{M})^{\mathbb{N}}$, and $\Omega := \mathcal{H}_\infty \times A \times \Theta$. In a long cheap-talk game, S first sees the state $\theta \in \Theta$. Then, at each time $t \in \mathbb{Z}_0$, players send simultaneous messages: S sends $m_t \in M$ and R sends $\tilde{m}_t \in \tilde{M}$. Finally, after seeing the sequence of messages, R chooses an action $a \in A$. Formally, a (behavior) strategy for S is a measurable function $\sigma : \Theta \times \mathcal{H}_{<\infty} \to \Delta M$, and a strategy for R is a pair of measurable functions $(\tilde{\sigma}, \rho)$, where $\tilde{\sigma} : \mathcal{H}_{<\infty} \to \Delta \tilde{M}$ and $\rho : \mathcal{H}_\infty \to \Delta A$. These maps induce (together with the prior μ_0) a unique distribution, $\mathbb{P}_{\sigma,\tilde{\sigma},\rho} \in \Delta\Omega$, which induces payoff $u_S(\mathrm{marg}_A \mathbb{P}_{\sigma,\tilde{\sigma},\rho})$ and $u_R(\mathrm{marg}_{A\times\Theta} \mathbb{P}_{\sigma,\tilde{\sigma},\rho})$ for S and R, respectively.

C.4.1 Extra rounds cannot help the sender

Below, we use our Theorem 1 to show that any S payoff attainable under long cheap talk is also attainable under one-shot communication.[40]

Proposition 5. *Every sender payoff attainable in a Nash equilibrium of the long cheap-talk game is also attainable in perfect Bayesian equilibrium of the one-shot cheap-talk game.*

Proof. Take any $s_* \in \mathbb{R}$ that is not an equilibrium payoff for prior μ_0 in the one-shot cheap-talk game. In particular, $s_* \notin V(\mu_0)$. Focus on the case of $s_* > v(\mu_0)$, the mirror image case being analogous. Because some sender-optimal equilibrium exists, some $s' < s_*$ exists that is strictly higher than every equilibrium S value. Letting B be the closed convex hull of $v^{-1}[s', \infty)$, the securability theorem tells us $\mu_0 \notin B$. Hahn-Banach then gives a continuous linear $\varphi : \mathrm{ca}(\Theta) \to \mathbb{R}$ such that $\varphi(\mu_0) > \max \varphi(B)$.

[40] To ease notational overhead, we employ Nash equilibrium as our solution concept in studying long cheap talk, and so have no need to define a belief map for the receiver. We therefore obtain a stronger result, because any perfect Bayesian equilibrium is also Bayes Nash.

Now define the function[41]

$$\begin{array}{rcl} F : \Delta\Theta \times \mathbb{R} & \to & \mathbb{R}_+ \\ (\mu, s) & \mapsto & [\varphi(\mu) - \max \varphi(B)]_+[s - s']_+. \end{array}$$

Observe that F is biconvex and continuous.[42]

Now consider any Nash equilibrium $(\sigma, (\tilde{\sigma}, \rho))$ of the long cheap-talk game. Let us define several random variables on the Borel probability space $\langle \Omega, \mathbb{P}_{\sigma,\tilde{\sigma},\rho} \rangle$. For $\omega = ((m_t, \tilde{m}_t)_{t=0}^{\infty}, a, \theta) \in \Omega$, let $\boldsymbol{\theta}(\omega) := \theta$ and $\mathbf{a}(\omega) := a$; and, for $t \in \mathbb{N}$, let $\mathbf{m}_{2t-1}(\omega) := m_t$ and $\mathbf{m}_{2t}(\omega) := \tilde{m}_t$. From these, we define a filtration $(\mathcal{F}_k)_{k \in K}$ with index set $K = \mathbb{Z}_+ \cup \{\infty\}$ by letting each $\mathcal{F}_k$ be the sigma-algebra generated by $\{\mathbf{m}_\ell\}_{\ell \in \mathbb{N},\ \ell < k}$. Finally, for each $k \in K$, define the ($\Delta\Theta$-valued and $\mathbb{R}$-valued, respectively) random variables $\boldsymbol{\mu}_k := \mathbb{E}\left[\delta_{\boldsymbol{\theta}} | \mathcal{F}_k\right]$ and $\mathbf{s}_k := \mathbb{E}\left[u_S(\mathbf{a}) | \mathcal{F}_k\right]$; and let $P_k \in \Delta(\Delta\Theta \times \mathbb{R})$ denote the distribution of $(\boldsymbol{\mu}_k, \mathbf{s}_k)$. Note that, by construction, P_0 has a distribution $\delta_{(\mu_0, s_0)}$ for some $s_0 \in \mathbb{R}$. Our task is to show $s_0 \neq s_*$.

In what follows, take any statements about the stochastic processes $(\boldsymbol{\mu}_k)_{k \in K}$ and $(\mathbf{s}_k)_{k \in K}$ to hold $\mathbb{P}_{\sigma,\tilde{\sigma},\rho}$-almost surely. By S rationality, $\mathbf{s}_{2t+1} = \mathbf{s}_{2t}$ for every $t \in \mathbb{Z}_+$. By construction, $\boldsymbol{\mu}_{2t} = \boldsymbol{\mu}_{2t-1}$ for every $t \in \mathbb{N}$, and both $(\boldsymbol{\mu}_k)_{k \in K}$ and $(\mathbf{s}_k)_{k \in K}$ are martingales. Because F is biconvex and continuous, $F(P_0) \leq F(P_1) \leq \cdots$. In particular, $F(P_k) \geq F(P_0) = F(\mu_0, s_0)$ for every $k \in \mathbb{Z}_+$. By the martingale convergence theorem, $\mathbf{s}_k$ converges to $\mathbf{s}_\infty$. By the same, every continuous $g : \Theta \to \mathbb{R}$ has $\int_\Theta g \, \mathrm{d}\boldsymbol{\mu}_k$ converging to $\int_\Theta g \, \mathrm{d}\boldsymbol{\mu}_\infty$; so $\boldsymbol{\mu}_k$ converges (weak*) to $\boldsymbol{\mu}_\infty$. But then, P_k converges (weak*) to P_∞. Therefore, $F(P_\infty) = \lim_{k\to\infty} F(P_k) \geq F(\mu_0, s_0)$. By R rationality, $\mathbf{s}_\infty \in V(\boldsymbol{\mu}_\infty)$, implying $F(\boldsymbol{\mu}_\infty, \mathbf{s}_\infty) = 0$, so that $F(P_\infty) = 0$ too. Therefore, $F(\mu_0, s_0) \leq 0 < F(\mu_0, s_*)$. So $s_0 \neq s_*$, as required. □

C.4.2 Extra rounds can help the receiver

Unlike S, R may benefit from long cheap talk when S's preferences are state-independent. To see this, consider the following example, which we describe informally. Let

[41] Recall that $[\cdot]_+ := \max\{\cdot, 0\}$.

[42] In the language of Aumann and Hart (1986), F separates (μ_0, s_*) from the graph of the S's value correspondence. As that paper shows, F is useful for witnessing a payoff as a non-equilibrium payoff. Although their formal results do not apply directly because Θ might be infinite, we closely follow the proof of Proposition A.8 (and results leading up to it) in Aumann and Hart (2003) in constructing the below martingales. Our focus on only S payoffs, together with state-independent S preferences, yields a less notationally cumbersome argument.

$\Theta = \{0,1\}$; $\mu_0(1) = \frac{1}{8}$; $A = \{\ell, b, t, r\}$; $u_S(b) = 0,\ u_S(\ell) = 1,\ u_S(t) = u_S(r) = 2$; and $u_R(a,\theta) = -(z_a - \theta)^2$, where $z_\ell = 0,\ z_r = 1$, and $z_b = z_t = \frac{1}{2}$. The associated value correspondence V and prior belief μ_0 are depicted in Figure 3 below.

Because every $\mu \in \Delta\Theta$ with $\mu(1) \leq \mu_0(1)$ has $V(\mu) = \{1\}$, Lemma 1 immediately implies every equilibrium outcome (p, s) of the one-shot cheap-talk game has $s = 1$ and $p\left\{\mu :\ \mu(1) \leq \frac{3}{4}\right\} = 1$. In particular, every equilibrium of the long cheap-talk game generates a "mean outcome" of y_0, as depicted in the figure.

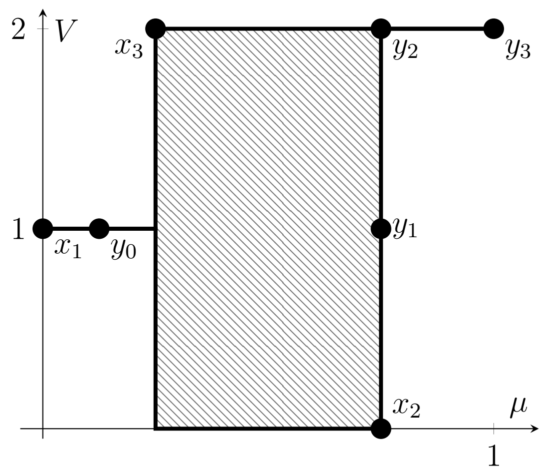

Figure 3: S's value correspondence in an example where R strictly benefits from long cheap talk.

Given the above observations, an equilibrium exists with one round of communication with R beliefs supported on $\left\{0, \frac{3}{4}\right\}$, and every other one-shot equilibrium generates less information (in a Blackwell sense) for R; we can depict this equilibrium as generating support $\{x_1, y_1\}$ in the figure. But now, with a jointly controlled lottery, this y_1 can be split in the next round to $\{x_2, y_2\}$.[43] Finally, S can provide additional information in the next round to split y_2 into $\{x_3, y_3\}$. Because action t is optimal for R at belief $\frac{3}{4}$ (i.e., that associated with y_2) but not at belief 1 (i.e., that associated with y_3), this additional information is instrumental to R. Therefore, our equilibrium is strictly better for R than any one-round equilibrium.

Thus, although additional rounds of communication do not change S's equilibrium

[43]Informally, following Aumann and Hart (2003), each player could toss a fair coin (independent of the state for S) and announce its outcome. Then the players move to x_2 if the coins come up the same and y_2 otherwise. Such jointly controlled randomization could be done simultaneously with the information that S initially conveys, so that our three-round example can be converted into a slightly more complicated two-round example.

payoff set, the static and long cheap-talk models are economically distinct, even under state-independent S preferences.