

Cheap Talk with Transparent Motives

Elliot Lipnowski Doron Ravid*

September 2017

Abstract

We study a model of cheap talk with one substantive assumption: the sender's preferences are state-independent. We observe that this setting is amenable to the belief-based approach familiar from models of persuasion with commitment. Using this approach, we examine the possibility of valuable communication, assess the value of commitment, and explicitly solve for sender-optimal equilibria in a large class of examples. A key result is a geometric characterization of the value of cheap talk, described by the *quasiconcave* 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 focused on the case in which the expert's preferences are private information. However, many experts have fixed and transparent motives: salespeople want to sell products with higher commissions; politicians want to get elected; and job

*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 Benjamin Brooks, Kevin Bryan, Odilon Câmara, Johannes Hörner, Navin Kartik, Laurent Mathevet, Stephen Morris, Maher Said, and Denis Shishkin for insightful discussions. We also thank audience members at Brown, Kellogg, UCSB, CETC 2017, the 9th Israeli GT Society conference, and ESEM 2017.

candidates want to get hired. 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 transparent motives. 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 technical insight is that our model is amenable to the belief-based approach adopted by the Bayesian persuasion literature (e.g., Kamenica and Gentzkow (2011), Alonso and Câmara (2016), and Ely (2017)). Thus, we summarize communication via its induced *information policy*, a distribution over receiver posterior beliefs that averages back to the prior.

Using the belief-based approach and the language of information policies, we obtain a convenient characterization of the sender’s ability to benefit from cheap talk. Say that a payoff s is persuasive 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 that a persuasive payoff s can be obtained in equilibrium if and only if s is securable. Note that the securing information policy need not itself arise in equilibrium. Still, Theorem 1 says that the existence of a policy securing s 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 reducing said information posterior-by-posterior, one can construct an equilibrium information policy that attains the securing policy’s lower bound on the sender’s payoffs.

Given the sender’s transparent motives, one may be tempted to conjecture that, in equilibrium, informative communication in general, and persuasion in particular, is impossible. Chakraborty and Harbaugh (2010) have proven this conjecture false. In

¹Other 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 existence of 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 (2017) 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.

a special case of our model, Chakraborty and Harbaugh (2010) show that the sender can always communicate information credibly and influence the receiver’s actions by trading off issues.² In the further specialization in which the sender strictly likes to influence the receiver, their result also delivers the possibility of persuasion.

Proposition 1 shows that the possibility of informative communication and persuasion extend beyond Chakraborty and Harbaugh’s (2010) specializations. In particular, the proposition traces failures of persuasion and communication to two examples. For persuasion, the relevant example is the two-action model, in which credible persuasion is clearly impossible. Using Theorem 1, we show that a converse is also true: persuasion is impossible if and only if the model is effectively a two-action model. The same is not true for informative communication, which is possible in all two-action models but for one exception: the two-action, two-state model in which the receiver is indifferent between the two actions at the prior. Using the belief-based approach, we adapt Chakraborty and Harbaugh’s (2010) argument to show that this exception is effectively the unique example without informative equilibria. Beyond this example, informative communication is always possible.

What is the sender’s benefit from cheap talk, and how does this compare to her benefit under commitment? Kamenica and Gentzkow (2011) characterize the sender’s benefit under commitment in terms of her *value function*—i.e., the highest value the sender can obtain from the receiver’s optimal behavior given his posterior beliefs. Specifically, they show that 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 *quasiconcave* 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.

To illustrate our results, 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 over others. In particular, suppose that the status quo is the think tank’s least preferred option and that the second reform is the think tank’s favorite option. Hence, let $a \in \{0, 1, 2\}$ represent

²See Battaglini (2002) and Chakraborty and Harbaugh (2007) for applications of this idea in the unknown agenda case.




Figure 1: A function with its concave (left) and quasiconcave (right) envelopes.

both the lawmaker's choice and the think tank's payoff from that choice. Assume that exactly one reform is good, and that the lawmaker wants to implement a good reform but cares more about avoiding bad ones. Such preferences arise if we normalize the lawmaker's payoff from the status quo to be 0 and have his payoffs be 1 from a good reform and -3 from a bad one. Thus, the lawmaker willingly implements a reform if and only if it is the right reform with probability at least $\frac{3}{4}$. Being uninformed, the lawmaker's prior assigns a probability of $\frac{1}{2}$ to each reform being the right one. Hence, absent any information, the lawmaker maintains the status quo.

Suppose that the think tank could reveal the state to the lawmaker: that is, the think tank recommends that the lawmaker *implement 1* for sure when the state is 1 and *implement 2* for sure when the state is 2. Since it is incentive-compatible for the lawmaker to follow these recommendations, the think tank's ex-post payoff would be 1 when sending *implement 1* and 2 when sending *implement 2*. In 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, since every information policy must sometimes reduce 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) this bound can be achieved by reducing the informativeness of the abovementioned message protocol. For (2), consider what happens when the think tank sends the *implement 2* message according

to

$$\begin{aligned}\mathbb{P}\{\text{implement } 2|\theta = 1\} &= \frac{1}{3}, \\ \mathbb{P}\{\text{implement } 2|\theta = 2\} &= 1,\end{aligned}$$

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. Thus, the lawmaker can mix between keeping the status quo and implementing 2 with equal probabilities. Since such mixing results in indifference by the think tank, we have obtained an equilibrium.

The above example highlights that credible communication may force the receiver to mix. Such mixing is unnecessary with sender commitment. Given commitment power, every attainable sender payoff can be attained with a message protocol satisfying two properties:³ (1) the sender's messages take the form of a recommended pure action; and (2) the receiver takes the recommended action with probability one. In contrast, Theorem 3 shows that, under cheap talk, one can always guarantee either one of these properties, but not necessarily both. Moreover, we show that commitment is valuable whenever (1) and (2) cannot be simultaneously guaranteed (Proposition 2). In other words, receiver mixing may be necessary, and this hurts the sender. Intuitively, mixing generates commitment power for the sender by preventing her from steering the receiver toward her preferred actions. Such mixing, however, comes at the cost of the receiver taking the sender's preferred actions with a lower probability. One can see this cost in our example: if the think tank were committed to the information policy from the above equilibrium, the lawmaker could respond to *implement 2* by choosing the second reform instead of mixing, thereby increasing the think tank's expected payoff to $\frac{5}{3}$. Thus, eliminating mixed strategies is one channel through which commitment benefits the sender.

In Section 6, we examine how cheap talk changes persuasive information. Under cheap talk, information is constrained by the sender's incentives. Despite this additional constraint, cheap talk need not change the nature of persuasive information. Consider the information policy constructed in the simple think tank example above.

³See Kamenica and Gentzkow (2011) and Bergemann and Morris (2016), for example.

We have described two ways that the lawmaker can respond to this policy: pure and mixed. With the mixed strategy, the above protocol gives the think tank its best cheap talk payoff. With the pure strategy, Kamenica and Gentzkow’s (2011) results imply that the protocol gives the think tank’s best commitment value. Notice that, under the pure strategy, the lawmaker only ever induces the think tank’s two best payoffs. This property turns out to be sufficient for the best commitment protocol to also be best under cheap talk. More precisely, Proposition 3 shows that a most persuasive information policy under commitment remains most persuasive under cheap talk whenever it induces the receiver to take two actions that lead to consecutive payoffs for the sender.

Section 7 examines a richer version of the above think tank example. In particular, we allow the lawmaker to choose from any finite number of proposals and the reforms’ values to have any symmetric distribution. We show, as in the simple example above, that the think tank would benefit from commitment and that one can use Theorem 1 to identify the think tank’s payoff from its most preferred equilibrium. Moreover, the equilibrium’s structure does not depend on the distribution of reform values.

2 Cheap Talk: The Belief-Based Approach

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) and a receiver (R). The game begins with the realization of an unknown 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. While R’s payoffs depend on θ , S’s payoffs do not.

We impose a few technical restrictions on our model. Both Θ and A are compact metrizable spaces that contain at least two elements. The state, θ , follows some full-support distribution $\mu_0 \in \Delta\Theta$, which is known to both players.⁴ Both players’ utility functions are continuous, where we take $u_S : A \rightarrow \mathbb{R}$ to be S’s utility⁵ and $u_R : A \times \Theta \rightarrow \mathbb{R}$ to be R’s. Finally, we assume that the message space, M , is some Polish space that contains A , ΔA , and $\Delta\Theta$.

⁴For a Polish 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 $\text{supp } \gamma$ denote the support of γ . For $f : Y \rightarrow \mathbb{R}$ bounded and measurable, we let $f(\gamma) := \int_Y f \, d\gamma$.

⁵In Section 8.1, we consider more general ways of modeling state-independent S preferences.

A strategy for S maps each state of the world to a distribution over messages. A strategy for R specifies a mixed action for R conditional on every message that R may observe. We are interested in studying the game's equilibria, by which we mean perfect Bayesian equilibria. Thus, an **equilibrium** consists of three measurable maps: a strategy $\sigma : \Theta \rightarrow \Delta M$ for S; a strategy $\rho : M \rightarrow \Delta A$ for R; and a belief system $\beta : M \rightarrow \Delta \Theta$ for R; such that:

1. β is obtained from μ_0 , given message policy σ , using Bayes' rule whenever possible.⁶
2. $\rho(m)$ is supported on $\arg \max_{a \in A} u_R(a, \beta(m))$ for all $m \in M$.
3. $\sigma(\theta)$ is supported on $\arg \max_{m \in M} u_S(\rho(m))$ 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,⁷ which, in turn, induces (through β and ρ , respectively) distributions over R's equilibrium beliefs and chosen mixed action.

We begin our analysis by noting that our model is amenable to the belief-based approach used in the information design 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 $\int_{\Delta \Theta} \mu \, dp(\mu) = \mu_0$. 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 .⁸ Let $\mathcal{I}(\mu_0)$ denote the set of all information policies.

The belief-based approach allows us to focus on the game's equilibrium 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.⁹ In contrast to equilibrium, a triple (σ, ρ, β) is a

⁶i.e., $\int_{\hat{\Theta}} \sigma(\hat{M}|\cdot) \, d\mu_0 = \int_{\Theta} \int_{\hat{M}} \beta(\hat{\Theta}|\cdot) \, d\sigma(\cdot|\theta) \, d\mu_0(\theta)$ for every Borel $\hat{\Theta} \subseteq \Theta$ and $\hat{M} \subseteq M$.

⁷Specifically, $\mathcal{E} = (\sigma, \rho, \beta)$ induces measure $\mathbb{P}_{\mathcal{E}} \in \Delta(\Theta \times M \times A)$, which assigns probability $\mathbb{P}_{\mathcal{E}}(\hat{\Theta} \times \hat{M} \times \hat{A}) = \int_{\hat{\Theta}} \int_{\hat{M}} \rho(\hat{A}|\cdot) \, d\sigma(\cdot|\theta) \, d\mu_0(\theta)$ for every Borel $\hat{\Theta} \subseteq \Theta$, $\hat{M} \subseteq M$, $\hat{A} \subseteq A$.

⁸For example, see Aumann et al. (1995), Benoît and Dubra (2011) or Kamenica and Gentzkow (2011).

⁹That is, if there exists an equilibrium $\mathcal{E} = (\sigma, \rho, \beta)$ such that $p = (\text{marg}_M \mathbb{P}_{\mathcal{E}}) \circ \beta^{-1}$ and $s = u_S(\text{marg}_A \mathbb{P}_{\mathcal{E}})$.

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.

Lemma 1 shows how to describe 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,¹⁰

$$V : \Delta\Theta \rightrightarrows \mathbb{R}$$

$$\mu \mapsto \text{co } u_S \left(\arg \max_{a \in A} u_R(a, \mu) \right).$$

Notice that V is a Kakutani correspondence.¹¹ Define the value function, $v := \max V : \Delta\Theta \rightarrow \mathbb{R}$, an upper semicontinuous function.

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

1. $p \in \mathcal{I}(\mu_0)$, i.e., $\int_{\Delta\Theta} \mu \, dp(\mu) = \mu_0$; and
2. $s \in \bigcap_{\mu \in \text{supp}(p)} V(\mu)$.

The above 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. R’s incentive-compatibility constraint implies that $V(\mu)$ must contain S’s continuation value from any message that leaves R at posterior belief μ . S’s incentive-compatibility requires S’s continuation value to be the same from all posteriors in p ’s support, which implies that her ex-post value is equal to her ex-ante value. Notice that this condition requires the entire correspondence V and cannot be expressed in terms of v alone. By contrast, in the literature on optimal information with commitment (e.g. Kamenica and Gentzkow (2011)), it is without loss to assume that R breaks indifferences in S’s favor.

If Lemma 1 characterizes equilibrium outcomes, Theorem 1 below characterizes S’s equilibrium payoffs. This characterization shows that one can relax S’s incentive constraint by focusing on the least persuasive message in any given information policy.

¹⁰In this paper, “co” refers to the convex hull, and “ $\overline{\text{co}}$ ” refers to the closed convex hull.

¹¹That is, a nonempty-, compact-, and convex-valued, upper hemicontinuous correspondence.

For concreteness, let p be an information policy, and take s to be some possible S payoff. Say that policy p **secures** s if $p\{v \geq s\} = 1$; and that s is **securable** if there exists an information policy that secures s —i.e., if $\mu_0 \in \overline{\text{co}}\{\mu : v(\mu) \geq s\}$. The following theorem shows that securability characterizes S’s equilibrium values.

Theorem 1. *Suppose that $s \geq v(\mu_0)$.¹² Then, there exists an equilibrium inducing sender payoff s 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, whenever s is strictly below $V(\mu)$, one can find a posterior μ' between μ_0 and μ such that s is in $V(\mu')$. This is possible because (1) s is between S’s no-information value and her highest μ payoff; and (2) V is a Kakutani correspondence, admitting a form of the intermediate value theorem. This transformation, therefore, replaces a potentially incentive-*incompatible* posterior μ with the incentive-*compatible* μ' . Since μ' is in the line segment between μ_0 and μ , this results in a weakly less informative signal.

3 Effective Communication

This section asks two questions. First, when is communication via cheap talk possible under transparent motives? And second, when can S benefit from such communication? Formally, each of these questions corresponds to the existence of a different kind of equilibrium: the first corresponds to an informative equilibrium, the second corresponds to a strictly persuasive one. An equilibrium \mathcal{E} is **informative** if R’s beliefs change on path, i.e., the belief distribution is not equal to δ_{μ_0} . For $\epsilon > 0$, an equilibrium is **ϵ -persuasive** if S’s payoffs are at least ϵ higher than in any uninformative equilibrium. Proposition 1 below characterizes the existence of both kinds of equilibrium.

Proposition 1 traces failures in strictly persuasive and informative equilibria, respectively, to two simple examples. For persuasion, the relevant example is a model with two actions over 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 $U_R(\theta) =$

¹²Given 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 $p\{\min V \leq s\} = 1$.

$u_R(1, \theta) - u_R(0, \theta)$, representing R's payoff difference between his two actions, given θ . Note that, in the two-action model, S is better off inducing a belief μ over the prior only if R views action 1 more favorably under μ than under the prior. Given this condition, p can be part of an ϵ -persuasive equilibrium only if it is supported on the set $\{U_R > U_R(\mu_0)\}$. But this is impossible, since the hyperplane $\{U_R = U_R(\mu_0)\}$ separates this set from μ_0 by construction. Therefore, persuasion is impossible.

To rule out informative communication, one needs an even more specialized example: two actions, two states, and R indifferent between his two actions at the prior.¹³ Since there are only two states, any non-degenerate information policy must sometimes increase and sometimes decrease U_R 's value. However, R's indifference at the prior means that S strictly prefers to increase U_R 's value—i.e., condition 2 below must hold—leaving no information as the unique equilibrium information policy.

In what follows, let \mathcal{U}_R be the set of all nonconstant continuous functions from Θ to \mathbb{R} .

Proposition 1. *Fix $\epsilon > 0$. There is no ϵ -persuasive equilibrium if and only if there exists $U_R \in \mathcal{U}_R$ such that $U_R(\mu_0) \leq 0$ and, for all $\mu \in \Delta\Theta$,*¹⁴

$$U_R(\mu) \leq U_R(\mu_0) \implies v(\mu) < v(\mu_0) + \epsilon. \quad (1)$$

There is no informative equilibrium if and only if $|\Theta| = 2$ and there exists $U_R \in \mathcal{U}_R$ such that $U_R(\mu_0) = 0$ and, for all $\mu \in \Delta\Theta$,

$$U_R(\mu) < U_R(\mu_0) < U_R(\mu') \implies V(\mu) \ll V(\mu'). \quad (2)$$

To see why condition 1 is necessary, suppose no ϵ -persuasive equilibrium exists for some $\epsilon > 0$. By Theorem 1, this is equivalent to μ_0 being outside $\overline{c\bar{0}}\{\mu : v(\mu) \geq v(\mu_0) + \epsilon\}$. Condition 1 then follows from a separating hyperplane theorem.

To prove necessity of condition 2, we begin by adapting an ingenious argument of Chakraborty and Harbaugh (2010) to obtain an informative equilibrium whenever there are three or more states.¹⁵ Chakraborty and Harbaugh (2010) study the spe-

¹³Note that an informative equilibria need not be *influential*. In other words, R's behavior may be independent of S's message.

¹⁴Recall our convention from footnote 4: here, $U_R(\mu) := \int_{\Theta} U_R d\mu$.

¹⁵To be precise, Chakraborty and Harbaugh (2010) give conditions for existence of an influential equilibrium, which is a weaker notion than persuasive equilibrium and, in their setting, stronger than informative equilibrium.

cialization of our 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. They show that, in their setting, S can credibly convey information to R by trading off issues.

Chakraborty and Harbaugh’s (2010) proof applies a fixed-point argument, relying on two important features of their setting: (1) the state is a multidimensional real vector; and (2) 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. 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 that condition 2 must be satisfied. Indeed, whenever the condition is violated, one can construct an informative equilibrium by finding two payoff-equivalent messages for S, each of which increases the probability of a different state. Necessity follows.

4 The Value of Commitment

The current section examines the value of commitment in persuasion. Let a **most persuasive equilibrium** be an ex-ante S-preferred equilibrium, and take a **most persuasive commitment outcome** to be an ex-ante S-preferred commitment outcome. In this section, we characterize S’s value from a most persuasive equilibrium and compare it with her value from a most persuasive commitment outcome.

We begin with Theorem 2, which geometrically characterizes S’s value from a most persuasive equilibrium. Take $\bar{v} : \Delta\Theta \rightarrow \mathbb{R}$ and $\hat{v} : \Delta\Theta \rightarrow \mathbb{R}$ to denote 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 .¹⁶ Since concavity implies quasiconcavity, the quasiconcave envelope lies (weakly) below the concave envelope. As described in Kamenica and Gentzkow (2011), \hat{v} gives S’s payoff from a most persuasive commitment outcome. Theorem 2 below shows that, under cheap talk, S’s optimal value reduces to \bar{v} .

¹⁶In the appendix, we show that \bar{v} is well-defined in Lemma 6. In the case that Θ is finite, the qualifier “upper semicontinuous” can be omitted from the definition without change.

Theorem 2. *A most persuasive equilibrium exists, providing value $\bar{v}(\mu_0)$ to the sender.*

Theorem 2 provides a geometric comparison between persuasion’s value under cheap talk and under commitment. Figure 2 below graphs v , \bar{v} and \hat{v} as a function of μ for the simple think tank example from the introduction. Since the state is binary, we can identify each posterior μ with the probability that it assigns to θ being 2. All figures have the think tank’s value correspondence, V , in light gray. The dashed black line represents v in the left figure, v ’s quasiconcave envelope in the middle figure, and v ’s concave envelope in the right figure. These correspond to the highest value that S can obtain from uninformative communication, cheap talk, and commitment, respectively.




Figure 2: The simple think tank example

As Figure 2 illustrates, the theorem allows us to compare the value of cheap talk persuasion to two natural benchmarks: uninformative communication and full commitment. Say that **cheap talk is costly** if no equilibrium gives S a most persuasive commitment outcome, and that **cheap talk is valuable** if a most persuasive equilibrium is strictly persuasive.¹⁷ Then, cheap talk is costly if and only if $\hat{v}(\mu_0) > \bar{v}(\mu_0)$ and is valuable if and only if $\bar{v}(\mu_0) > v(\mu_0)$.

Combining Theorem 2 with Theorem 1 gives the following convenient formula for \bar{v} :

Corollary 1. *S’s highest equilibrium value is $\bar{v}(\cdot) = \max_{p \in \mathcal{I}(\cdot)} \inf v(\text{supp } p)$.*

¹⁷If, rather than choosing what to say to the receiver, S were choosing an experiment, then the appropriate limited commitment model would not be one of cheap talk. Rather, it would be an informed principal problem, as in Perez-Richet (2014).

The above formula gives an alternative perspective on the value of commitment. Specifically, given a most persuasive equilibrium policy \bar{p} and a most persuasive commitment outcome \hat{p} , the value of commitment, $(\hat{v} - \bar{v})(\mu_0)$, can be expressed as $\int_{\Delta\Theta} v \, d\hat{p} - \inf v(\text{supp}(\bar{p}))$. This implies that

$$\int_{\Delta\Theta} v \, d\bar{p} - \inf v(\text{supp} \bar{p}) \leq (\hat{v} - \bar{v})(\mu_0) \leq \int_{\Delta\Theta} v \, d\hat{p} - \inf v(\text{supp} \hat{p}). \quad (3)$$

Notice that each of the above bounds on commitment's value can be derived by solving only one of the cheap talk and commitment models. With only a most persuasive equilibrium (resp. commitment) information policy in hand, one can deduce a lower (upper) bound on commitment's value.

5 Mixed Messages

What kind of messages does S send to R in equilibrium? Kamenica and Gentzkow (2011) show that any commitment payoff for S can be generated by a commitment protocol in which S makes a pure action recommendation to R, and R always complies. The following theorem shows that a similar result holds with cheap talk, with one important caveat: R must be allowed to mix. To state the theorem, for any S strategy σ , define \mathcal{M}_σ as the set of messages in σ 's support.¹⁸

Theorem 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 theorem suggests two ways in which one can implement a payoff of s using mixed action 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.

Proving that 1 implies 2 follows an analogous line to similar results in mechanism design (Myerson (1979)), mediated games (Myerson (1986)), persuasion with commitment (Kamenica and Gentzkow (2011)), and other information design settings

¹⁸i.e., let $\mathcal{M}_\sigma = \cup_{\theta \in \Theta} \text{supp} \sigma(\cdot|\theta)$.

(Bergemann and Morris (2016)). One can take any equilibrium and pool all messages leading to the same mixed action without changing R’s or S’s incentives. However, unlike the aforementioned similar results, S’s incentives do change if one tries to use the induced distribution over pure action recommendations to replace a mixed action. In particular, if a mixed action includes two actions that give S two different payoffs, recommending the low-payoff action will never be incentive-compatible. As such, some equilibrium payoffs may be impossible to implement using pure action recommendations that are always followed.¹⁹

To prove that 1 implies 3, we rely on the construction underlying Theorem 1. Start with a minimally informative information policy that secures s .²⁰ The transformation underlying Theorem 1 implies that there is a weakly less informative policy, p' , such that (p', s) is an equilibrium outcome. Since p is minimally informative, it must be that p is equal to p' . Thus, it is sufficient to show that one can implement p as in part 3. Let \mathcal{E} be the equilibrium constructed for part 2, 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, as pooling any posteriors that induce the same $\mathbf{a}(\cdot)$ value would yield an even less informative policy that secures s . Thus, \mathbf{a} takes distinct beliefs to distinct (on-path) actions. 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 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 that one can restrict attention to direct mechanisms in which the agent reports truthfully with positive probability. While 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.

The necessity of mixing by R can have payoff consequences for S. To see why, consider the argument in the above paragraph, and suppose that S could commit. Then, one could increase S’s payoffs by replacing any mixed action recommendation

¹⁹This issue would be sidestepped if the players had access to a disinterested mediator, as in Lugo (2016), where one can apply Myerson’s (1986) revelation principle.

²⁰Here, we use Blackwell’s information ordering. See Subsection 8.3 for a formal definition.

with S's most preferred action in the mixture's support. If doing this is impossible, then S's incentive constraint must be binding, which means that S must be losing value from limited commitment. We formalize this result in Proposition 2 below.

Proposition 2. *Suppose that (σ, ρ, β) is a most persuasive equilibrium, and that $(\sigma, \tilde{\rho}, \beta)$ is not a most persuasive equilibrium for any receiver strategy $\tilde{\rho}$ without mixing. Then cheap talk is costly, i.e., $\hat{v}(\mu_0) > \bar{v}(\mu_0)$.*

At this stage, the reader may wonder whether the proposition's hypothesis can be replaced by a hypothesis on primitives, rather than on a particular equilibrium. That is, are there cases in which we can directly show that optimal cheap talk persuasion requires R to mix? Consider the simple think tank example presented in the introduction. In this example, the lawmaker can take one of three actions, each of which leads to a different payoff for the think tank. In other words, the think tank's payoffs are a one-to-one function of the investor's actions. This property turns out to be sufficient for cheap talk persuasion to require mixing. To see why, consider an equilibrium in which R never mixes. As S is never indifferent between distinct actions, R must be choosing the same action after every message. These messages can then be pooled: R would rationally choose the same action at the prior. Therefore, any strictly persuasive equilibrium entails some mixing by R. Since such mixing entails a value loss (Proposition 2), we obtain the corollary below.

Corollary 2. *Suppose that u_S is one-to-one. Then, one of the following holds:*

1. *Persuasion is impossible:* $\hat{v}(\mu_0) = \bar{v}(\mu_0) = v(\mu_0)$.
2. *Cheap talk is costly:* $\hat{v}(\mu_0) > \bar{v}(\mu_0)$.

There are two distinct sources of value loss to a sender who lacks commitment. Mixing is a specific instance of one: R breaks indifference against S's interests to provide S with appropriate incentives. The other is a change to the information revealed in equilibrium. Let (\bar{p}, s) be a most persuasive equilibrium outcome, and \hat{p} a most persuasive commitment policy. The value of commitment can be decomposed as:

$$\hat{v}(\mu_0) - \bar{v}(\mu_0) = \underbrace{\int_{\Delta\Theta} v \, d(\hat{p} - \bar{p})}_{\text{information distortion}} + \underbrace{\int_{\Delta\Theta} (v - s) \, d\bar{p}}_{\text{incentive cost}}. \quad (4)$$

The first expression is the difference between S's payoff when committing to \hat{p} and when committing to \bar{p} . This loss comes from S providing R with less persuasive information. The second expression is the difference (posterior by posterior) between S's highest possible payoff given R's belief and S's equilibrium payoff from the same. This represents the loss due to costly incentive provision. Together, these comprise the loss caused by S's lack of commitment power.

6 Persuasive Information

Unlike persuasion with full commitment, cheap talk limits the parties to using only incentive-compatible communication. While incentives clearly constrain what S *can* say, they do not always constrain what S *wants* to say. To illustrate, we consider the policy identified in the simple example from the introduction. This information policy is both a most persuasive cheap talk policy *and* an optimal policy for S under commitment. A property of this information policy is that it induces two consecutive payoffs for S. More precisely, say that payoffs $s > s'$ are **consecutive** if $u_S(A) \cap [s', s] = \{s, s'\}$ and that $p \in \mathcal{I}(\mu_0)$ is **supported on two consecutive payoffs** if there are two consecutive s and s' such that $p\{v = s \text{ or } s'\} = 1$. Proposition 3 below shows that a most persuasive commitment policy supported on two consecutive payoffs must also be a most persuasive cheap talk policy.

Proposition 3. *Let \hat{p} be a most persuasive commitment information policy. If \hat{p} is supported on two consecutive payoffs, then \hat{p} is a most persuasive equilibrium information policy.*

To understand Proposition 3, consider a most persuasive commitment policy, \hat{p} , which is supported on two consecutive payoffs, $s > s'$. The proof involves two steps. The first step is to show that (\hat{p}, s') is an equilibrium outcome. Since \hat{p} is optimal, whenever \hat{p} induces R to take an s action, R must also be willing to take another action that gives S a lower payoff. Otherwise, one could increase the weight that \hat{p} places on s by replacing s posteriors at which R is not indifferent with other posteriors that are closer to the prior. But R's indifference then means that R can replace the s action with a mixture that yields S a payoff of s' , thereby making \hat{p} incentive-compatible for S.

The second step is to show that \hat{p} must be a most persuasive equilibrium policy.

The key observation here is that, s' and s being consecutive, any policy that secures S a payoff strictly greater than s' would secure s , which cannot happen given the optimality of \hat{p} . Theorem 1 then tells us that \hat{p} is a most persuasive equilibrium policy.

The value of commitment is simplified when an optimal commitment policy, \hat{p} , is also a most persuasive equilibrium policy. Let (\hat{p}, s) be a most persuasive equilibrium outcome. Then, (3) and (4) yield the following formula for the value of commitment:

$$\hat{v}(\mu_0) - \bar{v}(\mu_0) = \int_{\Delta\Theta} v \, d\hat{p} - \inf v(\text{supp } \hat{p}) = \int_{\Delta\Theta} (v - s) \, d\hat{p}.$$

Conceptually, each equality corresponds to a different perspective on commitment's value. The first equality comes from S 's most persuasive equilibrium payoff being the highest payoff that \hat{p} can secure. The second equality comes from the fact that cheap talk does not distort information. As such, (4) tells us that commitment's sole benefit is the ability to avoid costly incentive provision.

7 Think Tank Example

Consider now a richer version of the introduction's example. A think tank (S) advises a lawmaker (R) on whether to pass one of $n \in \mathbb{N}$ reforms or to pass none. A given reform $i \in \{1, \dots, n\}$ has 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 has transparent motives: its payoff is $u_S(i)$ from reform i and $u_S(0) = 0$ from the status quo, where $u_S : \{0, \dots, n\} \rightarrow \mathbb{R}$ is increasing.²¹ We now use the results of previous sections to characterize the think tank's most persuasive equilibrium value and construct a most persuasive equilibrium. In this section, we proceed informally; see the appendix for further details.

Suppose that $p \in \mathcal{I}(\mu_0)$ secures $u_S(k)$ for some $k \in \{1, \dots, n\}$. Then it must be that for any μ in p 's support, there is some $i \in \{k, \dots, n\}$ such that $i \in \arg \max_{a \in A} u_R(a, \mu)$.

²¹Che et al. (2013) study a similar application with a two important differences. First, they require the prior to be non-exchangeable (see equation (R1) in their Definition 1). Second, they assume that S has state-dependent preferences that agree with R 's on all but the default option.

In particular,

$$c \leq \max_{i \in \{k, \dots, n\}} \int_{\Theta} \theta_i \, d\mu(\theta) \leq \int_{\Theta} \max_{i \in \{k, \dots, n\}} \theta_i \, d\mu(\theta) = \int_{\Theta} \theta_{k,n}^{(1)} \, d\mu(\theta),$$

where $\theta_{k,n}^{(1)}$ is the order statistic $\theta_{k,n}^{(1)} = \max_{i \in \{k, \dots, n\}} \theta_i$. Integrating over p then tells us that $\int_{\Theta} \theta_{k,n}^{(1)} \, d\mu_0(\theta) \geq c$, so that the lawmaker is willing to adopt reform $i \in \{k, \dots, n\}$ upon learning that i is the best project among $\{k, \dots, n\}$.²² Therefore, the information policy $p_k \in \mathcal{I}(\mu_0)$, which reveals the identity of the lawmaker's most preferred reform among $\{k, \dots, n\}$, secures $u_S(k)$, as well.

The above demonstrates that, to secure a high value, the think tank could restrict attention to a small class of information policies. Picking a small set $\{k, \dots, n\}$ of reforms that it most prefers, the think tank could reveal the lawmaker's best choice among $\{k, \dots, n\}$. This secures $u_S(k)$ if and only if this information is precise enough to offset the cost of a reform. From here, Theorem 1 delivers a full characterization of the most persuasive equilibrium payoff. Let

$$K := \left\{ k \in \{1, \dots, n\} : \int_{\Theta} \theta_{k,n}^{(1)} \, d\mu_0(\theta) \geq c \right\}.$$

If K is empty, then no equilibrium is strictly persuasive: the lawmaker never implements a reform in any equilibrium. Otherwise, letting $k^* := \max K$, the think tank's best equilibrium value is $u_S(k^*)$, as secured by p_{k^*} . For instance, if $\theta_1, \dots, \theta_n$ are i.i.d. uniformly distributed on $[0, 1]$,

$$k^* = \left\lfloor n - \frac{c}{1-c} \right\rfloor.$$

The policy p_{k^*} helps us find the value of persuasion via cheap talk, but it will typically not be an equilibrium policy. The constructive proof underlying Theorem 1, however, delivers an equilibrium policy that generates a payoff $u_S(k^*)$ to the think tank. In this equilibrium, the think tank either accurately reports the lawmaker's best reform from $\{k^*, \dots, n\}$ or (with probability ϵ) suggests a uniformly random reform from $\{k^*, \dots, n\}$. Degrading information in this way makes the lawmaker indifferent between the suggested proposal and no reform. In equilibrium, the lawmaker implements the suggested reform i with probability $\frac{u_S(k^*)}{u_S(i)}$ and implements no

²²See Claim 1 in the appendix.

reform with complementary probability. Thus, the think tank enjoys some credibility endogenously, as it faces a quality-quantity tradeoff: between the *value* of having its advice followed and the *probability* of having its advice followed.

What is the value of commitment in this example? By Proposition 2, and its Corollary 2, we can immediately deduce, without needing to solve the commitment case, that either (1) cheap talk is costly, or (2) providing no information is optimal under both commitment and cheap talk.²³ Hence, except for degenerate cases, cheap talk is always costly in this application. One can bound this cost using the bound in equation (3). Specifically, the value of commitment is at least

$$\frac{1}{n-k^*+1} \sum_{i=k^*}^n u_S(i) - u_S(k^*),$$

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

8 Discussion and Related Literature

8.1 State-independent Preferences

While S's objective depends only on R's action in our model, one can write more general models of transparent motives. Fix a continuous S objective, $\tilde{u}_S : A \times \Theta \rightarrow \mathbb{R}$. We say that \tilde{u}_S exhibits **cardinally (ordinally) state-independent preferences** if, for any $\theta, \theta' \in \Theta$ and $\alpha, \alpha' \in \Delta A$ ($a, a' \in A$): $\tilde{u}_S(\alpha, \theta) \geq \tilde{u}_S(\alpha', \theta)$ implies $\tilde{u}_S(\alpha, \theta') \geq \tilde{u}_S(\alpha', \theta')$ ($\tilde{u}_S(a, \theta) \geq \tilde{u}_S(a', \theta)$ implies $\tilde{u}_S(a, \theta') \geq \tilde{u}_S(a', \theta')$). Our analysis often applies even under these more permissive notions of state-independence.

Proposition 4. *Fix the action space, state space, and receiver preferences. Then the game with sender objective \tilde{u}_S and the game with sender objective $u_S := \tilde{u}_S(\cdot, \mu_0)$ 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.

²³Case (2) holds in either of two circumstances: $k^* = n$ (so that the think tank's preferred reform is incentive-compatible under no information) or the status quo is strictly dominant for the lawmaker. Otherwise, using exchangeability, one can show that full information strictly outperforms no information. Case (1) then obtains.

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. In each part, we assume that 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 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 4. While many of our results generalize to a model in which R faces preference shocks that are independent of θ ,²⁴ 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.

8.2 Fully Revealing Equilibrium

When can S tell R everything she knows? This question has been asked in a variety of settings. For example, Baliga and Morris (2002) and Hagenbach et al. (2014) study when players can communicate what they know before a strategic interaction. Mathis (2008) provides conditions under which a sender can reveal her information when information is certifiable. Renault et al. (2013), Golosov et al. (2014), and Margarita and Smolin (2017) study, among other things, the possibility of full information revelation when the interaction between S and R is dynamic.

An equilibrium (σ, ρ, β) is **fully revealing** if, for every $m \in M$, there is some $\theta \in \Theta$ such that $\beta(\theta|m) = 1$. Lemma 1 then implies that a fully revealing equilibrium exists if and only if

$$\sup_{\theta \in \Theta} \min V(\delta_\theta) \leq \inf_{\theta \in \Theta} \max V(\delta_\theta).$$

The proof is straightforward. Note, first, that, if the above inequality does not hold,

²⁴In such a model, there is 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 $u_R : A \times Z \times \Theta \rightarrow \mathbb{R}$. In this extended model, $V : \Delta\Theta \rightrightarrows \mathbb{R}$ should take the form $V(\mu) = \int_Z \text{co } u_S(\arg \max_{a \in A} u_R(a, z, \mu)) \, d\zeta_0(z)$, a Kakutani correspondence.

there must be two states such that S necessarily strictly prefers R to be completely convinced of one rather than the other. Revealing the less beneficial state to R cannot then be incentive-compatible. Suppose, now, that the above inequality holds. Since V is a Kakutani correspondence, there is a payoff that S can obtain after revealing the state to R, regardless of which state it is. The existence of a fully revealing equilibrium then follows from Lemma 1.

8.3 Receiver Payoffs and Equilibrium Efficiency

While our focus is on credible persuasion, we can also use the belief-based approach to examine R's payoffs. There are two ways to do so. The first is to evaluate R's payoffs from any (p, s) directly via $\int_{\Delta\Theta} \max_{a \in A} u_R(a, \mu) dp(\mu)$. The second is through p 's informative content. More precisely, say that p is **more informative than** p' if and only if p is a mean-preserving spread of p' .²⁵ Then R is always weakly better off under a more informative policy (Blackwell (1953)).

Knowing that equilibrium outcomes also include information about R's payoffs allows one to examine the inefficiencies that arise under cheap talk. One inefficiency that plagues cheap talk is coordination failure. A case in point is Crawford and Sobel's (1982) leading example, which has multiple equilibria, all of which are ex-ante Pareto ranked. More precisely, given any two distinct equilibria, there is one that simultaneously has S sending more messages²⁶ and is ex-ante better for both players. The following proposition says that similar coordination failures also occur in our model.

Proposition 5. *There are equilibrium outcomes $\{(p_s^*, s)\}_{v(\mu_0) \leq s \leq \bar{v}(\mu_0)}$ such that, for all $s, s' \in [v(\mu_0), \bar{v}(\mu_0)]$ with $s > s'$:*

1. $(p_{s'}^*, s')$ is Pareto inferior to (p_s, s) .
2. $p_{s'}^*$ is less informative than p_s .
3. $|\text{supp } p_{s'}^*| \leq |\text{supp } p_s^*|$.

²⁵Say that p is a mean-preserving spread of p' if there exists a measurable map $G : \Delta\Theta \rightarrow \Delta\Delta\Theta$ such that: (1) $G(\mu) \in \mathcal{I}(\mu)$ for all μ ; and (2) $p(D) = \int_{\Delta\Theta} G(D|\cdot) dp'$ for all Borel $D \subseteq \Delta\Theta$.

²⁶More precisely, it has S inducing more beliefs with positive probability. The number of messages sent in equilibrium is indeterminate given the information policy.

Notice that this result holds for an arbitrary specification of cheap talk with state-independent S preferences.

The proposition's proof relies on the transformation behind Theorem 1. To see this, let (p, s) be an equilibrium outcome for some s strictly above $v(\mu_0)$. Then, p secures s , which implies that p also secures any value between s and $v(\mu_0)$. Letting s' be such a value, we can now apply the transformation of Theorem 1 on p to obtain a p' such that (p', s') is also an equilibrium outcome. Since p' is obtained by moving every posterior belief for p to a new belief closer to the prior, it must be that p' is both less informative than p and has smaller support. Finally, since p' is less informative than p , R is weakly worse off with p' than with p , so that (p', s') is Pareto inferior to (p, s) .

8.4 Cheap Talk and Money Burning

In Theorem 1, we show that degrading information can turn a policy that secures value s into an equilibrium generating the same value. The payoff reduction is similar to assigning a cost to each message—i.e., money burning. This is more than simply an analogy. Suppose that S could commit to having some money burned in response to a given message. S would then be solving the following problem:²⁷

$$\max_{p \in \mathcal{I}(\mu_0)} \int_{\Delta\Theta} (v - c) dp$$

$$c : \Delta\Theta \rightarrow \mathbb{R} \text{ meas.}$$

$$\text{s.t. } v(\mu) - c(\mu) \geq v(\mu') - c(\mu'), \forall \mu, \mu' \in \text{supp}(p) \quad (5)$$

$$c(\mu) \geq 0, \forall \mu \in \text{supp}(p). \quad (6)$$

The above program relaxes the program that defines S's most persuasive equilibrium value. Still, Theorem 1 implies that the two programs yield the same value. To see why, fix an information policy p and consider the problem of finding the best $c : \Delta\Theta \rightarrow \mathbb{R}_+$ for S, subject to c making p incentive-compatible. Note that making p incentive-compatible fixes the difference $c(\mu) - c(\mu')$ for all μ and μ' in p 's support.²⁸

²⁷Costs could be associated with messages rather than with their induced posterior beliefs with no essential change.

²⁸Given money burning, it is without loss for S values to have R break any indifferences in S's favor.

As such, c is completely pinned down over p 's support by its minimal value. But this minimal value must then be zero for the optimal such c , since S always prefers lower costs to higher ones. As such, the best c that makes p incentive-compatible gives S a net payoff of $\inf_{\mu \in \text{supp}(p)} v(\mu)$. Maximizing over p gives the value from Theorem 1.

As the previous paragraph shows, the ability to burn money would not benefit S. Money burning can, however, benefit R. In particular, it is always possible in our model to achieve full information revelation with money burning by taking:

$$c(\delta_\theta) = v(\delta_\theta) - \inf_{\theta \in \Theta} v(\delta_\theta)$$

to be the cost of revealing that the state is θ . Under these costs, it is clearly incentive-compatible for S to reveal the state. Since R cannot do better than under full information, money burning allows R to obtain his first-best. Austen-Smith and Banks (2000) and Kartik (2007) discuss how money burning expands the set of equilibrium outcomes in the model of Crawford and Sobel (1982).²⁹ As in our model, they show that full information revelation is an equilibrium outcome when money burning is allowed. Kartik (2007) also asks what happens if the value that S can burn is capped from above. He shows that, as the cap converges to zero, the set of achievable outcomes converges to what can be achieved via cheap talk. It is straightforward to show that the same result holds in our model.

One can further describe S's payoff from a most persuasive commitment outcome with a similar commitment technology. To do so, one requires that costs are positive ex-ante rather than ex-post. Formally, this corresponds to replacing (6) with the constraint that the expected cost is positive—i.e., $\int_{\Delta\Theta} c(\mu) dp(\mu) \geq 0$. This would be a reasonable assumption if S could place verifiable, actuarially fair bets with a third party. Note, though, that this ex-ante constraint implies that S's payoff from any information policy p , minus transfers, is bounded from above by her expected fundamental payoff, $\int_{\Delta\Theta} v dp$. Therefore, S's payoff from having such a commitment technology is bounded from above by her most persuasive commitment payoff, $\hat{v}(\mu_0)$. But clearly S can achieve this bound by taking a most persuasive commitment policy p and setting $c(\mu) = v(\mu) - \int_{\Delta\Theta} v dp$. Thus, from S's standpoint, the difference between the two assumptions—that costs are positive ex-ante versus ex-post—is exactly equal

²⁹We allow S to commit for convenience, while Austen-Smith and Banks (2000) and Kartik (2007) focus on voluntary money burning. Making money burning voluntary in our model does not alter the equivalence result.

to the value of commitment.

9 Conclusion

In this paper, we study cheap talk under the assumption that the sender has transparent motives. Adopting a belief-based approach, we derive a complete characterization of feasible equilibrium outcomes. One consequence is that any sender value that can be secured ex-post is consistent with equilibrium play. Our approach further allows us to trace persuasion and communication failures to a limited class of examples, to describe the sender-preferred equilibrium value geometrically, and to contrast this value with that of persuasion with commitment. In particular, the cost of communicating via cheap talk alone can be decomposed into two: an information distortion cost and a truthtelling incentive cost. Finally, we fully solve 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, Michael Maschler, and Richard E Stearns**, *Repeated games with incomplete information*, MIT press, 1995.
- Austen-Smith, David and Jeffrey S Banks**, “Cheap talk and burned money,” *Journal of Economic Theory*, 2000, 91 (1), 1–16.
- Baliga, Sandeep and Stephen Morris**, “Co-ordination, spillovers, and cheap talk,” *Journal of Economic Theory*, 2002, 105 (2), 450–468.
- 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.
- Blackwell, David**, “Equivalent comparisons of experiments,” *The annals of mathematical statistics*, 1953, *24* (2), 265–272.
- 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.
- 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.
- Ely, Jeffrey C**, “Beeps,” *American Economic Review*, 2017, *107* (1), 31–53.
- Golosov, Mikhail, Vasiliki Skreta, Aleh Tsyvinski, and Andrea Wilson**, “Dynamic strategic information transmission,” *Journal of Economic Theory*, 2014, *151*, 304–341.
- Green, Jerry R and Nancy L Stokey**, “A two-person game of information transmission,” *Journal of Economic Theory*, 2007, *135* (1), 90–104.
- Hagenbach, Jeanne, Frédéric Koessler, and Eduardo Perez-Richet**, “Certifiable Pre-Play Communication: Full Disclosure,” *Econometrica*, 2014, *82* (3), 1093–1131.

- Kamenica, Emir and Matthew Gentzkow**, “Bayesian persuasion,” *American Economic Review*, 2011, 101 (October), 2590–2615.
- Kartik, Navin**, “A note on cheap talk and burned money,” *Journal of Economic Theory*, 2007, 136 (1), 749–758.
- Lugo, Andrés Salamanca**, “The Value of Mediated Communication,” 2016.
- Margaria, Chiara and Alex Smolin**, “Dynamic Communication with Biased Senders,” 2017.
- Mathis, Jérôme**, “Full revelation of information in Sender–Receiver games of persuasion,” *Journal of Economic Theory*, 2008, 143 (1), 571–584.
- Myerson, Roger B**, “Incentive compatibility and the bargaining problem,” *Econometrica*, 1979, pp. 61–73.
- , “Multistage games with communication,” *Econometrica*, 1986, pp. 323–358.
- Perez-Richet, Eduardo**, “Interim bayesian persuasion: First steps,” *American Economic Review*, 2014, 104 (5), 469–474.
- Phelps, Robert R**, *Lectures on Choquet’s theorem*, Springer Science & Business Media, 2001.
- Renault, Jérôme, Eilon Solan, and Nicolas Vieille**, “Dynamic sender–receiver games,” *Journal of Economic Theory*, 2013, 148 (2), 502–534.
- Schnakenberg, Keith E**, “Expert advice to a voting body,” *Journal of Economic Theory*, 2015, 160, 102–113.

A Appendix

A.1 Technical Preliminaries

First, we record a useful measurable selection result.

Lemma 2. *If $D \subseteq \Delta\Theta$ is Borel and $f : D \rightarrow \mathbb{R}$ is any measurable selector of $V|_D$, then there exists a measurable function $\alpha_f : D \rightarrow \Delta A$ 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. $|\text{supp}(\hat{\alpha})| \leq 2$.

Proof. 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. The correspondence $u_S^{-1} : u_S(A) \rightarrow A$ 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:

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

By construction, α_f is as desired. □

We now prove a variant of the intermediate value theorem which is useful for our setting. This result is essentially proven in de Clippel (2008), Lemma 2. As the statement of that lemma is slightly weaker than we need, 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} , there exists some weakly decreasing $\{x_n^+\}_{n=1}^\infty \subseteq [\bar{x}, 1]$ which 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 which converges to \bar{x} otherwise. By definition of \bar{x} (and, in the case of $\bar{x} = 0$, since $\min F(0) \leq 0$), 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 that $\min F(\bar{x}) \leq 0$. As F is convex-valued, it follows that $0 \in F(\bar{x})$. \square

A.2 Proofs for Section 2

A.2.1 Lemma 1

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

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

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 $\mu \in \Delta\Theta$. As 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)) that $p \in \mathcal{I}(\mu_0)$ implies there is some S strategy σ and Bayes-consistent β that induce distribution p over posterior beliefs.³⁰ Without disrupting the Bayesian property, we may without loss assume that $\beta(m) \in D$ for all $m \in M$. Now, let $\alpha = \alpha_s : D \rightarrow \Delta A$ be as given by Lemma 2. We can then define the receiver strategy $\sigma := \alpha \circ \beta$, which satisfies receiver incentive-compatibility by definition of α . Finally, by construction, $\int_A u_S d\rho(\cdot|m) = s$ for every $m \in M$, so that every

³⁰In particular, there is one with $\sigma(\Delta\Theta|\theta) = 1$ for all $\theta \in \Theta$ and $\beta(\cdot|\mu) = \mu$ for all $\mu \in \Delta\Theta$.

S strategy is incentive-compatible. Therefore, (σ, ρ, β) is an equilibrium generating outcome (p, s) . \square

A.2.2 From securability to equilibrium

Below, we prove a lemma which will be the heart of Theorem 1. It constructs an 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. *Suppose $p \in \mathcal{I}(\mu_0)$ secures $s \geq v(\mu_0)$. Then there exists $p^* \in \mathcal{I}(\mu_0)$ such that:*

1. (p^*, s) is an equilibrium outcome;
2. p^* is weakly less Blackwell-informative than p , and $|\text{supp}(p^*)| \leq |\text{supp}(p)|$.

Proof. The result is trivial (focusing on S-preferred uninformative equilibrium) for $s = v(\mu_0)$, so focus on the case of $s > v(\mu_0)$. Let $p \in \mathcal{I}(\mu_0)$ secure s , and $D := \text{supp}(p)$. Notice that $v(\mu) \geq s$ for every $\mu \in D$ since v is upper semicontinuous. Define the measurable function,

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

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 there is some number $\epsilon > 0$ such that $\lambda \geq \epsilon$ uniformly. If there were no such ϵ , then there would be a sequence $\{\mu_n\}_n \subseteq D$ 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) . As V is upper hemicontinuous, this 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

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

for every Borel $\hat{D} \subseteq \Delta\Theta$. Direct computation shows that $p^* \in \mathcal{I}(\mu_0)$. Lemma 1 delivers an equilibrium generating S value s and information policy p^* .

Lastly, we note that p^* has the other desired 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 $(1 - \int_{\Delta\Theta} \lambda dp) \delta_{\mu_0} + (\int_{\Delta\Theta} \lambda dp) p$, which in turn is less informative than p . \square

A.2.3 Theorem 1

Proof. The “if” direction follows directly from Lemma 1: for any equilibrium outcome (p, s) , information policy p secures payoff s . The “only if” direction is a direct consequence of Lemma 4. \square

For convenience, we record an easy consequence of the above theorem.

Corollary 3. *The set of sender payoffs s induced by some equilibrium 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. From Theorem 2, proven below, Π^* has a maximum value s_+ . By an identical argument, Π_- is convex and has a minimum value s_- .³¹

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

A.3 Proofs for Section 3

In this subsection, we prove Proposition 1. The first part follows from Theorem 1, given a separation argument. The second part follows from a fixed-point argument.

³¹Notice that the only property of V used in the proofs—that it is a Kakutani correspondence—is also true of $-V$.

A.3.1 Proposition 1: persuasive equilibrium

Proof. Appealing to Theorem 1 tells us that:

$$\begin{aligned}
 \exists \text{ some } \epsilon\text{-persuasive equilibrium} &\iff \exists s \geq v(\mu_0) + \epsilon \text{ which is securable} \\
 &\iff v(\mu_0) + \epsilon \text{ is securable} \\
 &\iff \mu_0 \in \overline{\text{co}} \{v \geq v(\mu_0) + \epsilon\}.
 \end{aligned}$$

We must therefore show that $\mu_0 \notin \overline{\text{co}} \{v \geq v(\mu_0) + \epsilon\}$ if and only if there exists some $U_R \in \mathcal{U}_R$ as described in the first part of the theorem.

First, suppose that $\mu_0 \notin \overline{\text{co}} \{v \geq v(\mu_0) + \epsilon\}$. Then the strong separating hyperplane theorem tells us that there is some continuous linear functional $\phi : ca(\Theta) \rightarrow \mathbb{R}$ such that $\phi(\mu_0) < \min \phi(\overline{\text{co}} \{v \geq v(\mu_0) + \epsilon\})$. By the Riesz representation theorem, there is some $U_R \in \mathcal{U}_R$ such that $\phi(\mu) = \int_{\Theta} U_R d\mu$ for all $\mu \in ca(X)$. Translating, we may assume $U_R(\mu_0) \leq 0$. Clearly, U_R then takes the desired form.

Next, suppose that U_R is as described in the first part of the theorem. As integration against U_R is continuous, and a continuous function on compact domain attains a minimum, it follows that $U := \min U_R(\overline{\text{co}} \{v \geq v(\mu_0) + \epsilon\})$ exists, and is therefore strictly greater than $U_R(\mu_0)$. Integration against U_R being linear and continuous, it must be that $U_R(\mu) \geq U > U_R(\mu_0)$ for every $\mu \in \overline{\text{co}} \{v \geq v(\mu_0) + \epsilon\}$. In particular, $\mu_0 \notin \overline{\text{co}} \{v \geq v(\mu_0) + \epsilon\}$. \square

A.3.2 Proposition 1: informative equilibrium

Toward showing that there always exist informative equilibria in the case of $|\Theta| > 2$, we prove the following fixed-point result, which generalizes the one-dimensional case of the Borsuk-Ulam theorem to the case of a Kakutani correspondence. The applicability of such fixed-point reasoning to show the possibility of effective communication is an important technical insight of Chakraborty and Harbaugh (2010).

Lemma 5. *If $\mu_0 = \sum_{i=1}^3 p_i \mu_i$ for some $\mu_1, \mu_2, \mu_3 \in \Delta\Theta$ and $p_1, p_2, p_3 > 0$, then there exist distinct $q, q' \in \Delta\{1, 2, 3\}$ and an equilibrium whose generated information policy is supported on $\{\sum_{i=1}^3 q_i \mu_i, \sum_{i=1}^3 q'_i \mu_i\}$.*

Proof. For small enough $\epsilon > 0$, we can embed the circle

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

into the space of beliefs via the map

$$\begin{aligned}\phi : \mathbb{S} &\rightarrow \Delta\Theta \\ (x, y) &\mapsto (p_1 + \epsilon x)\mu_1 + (p_2 + \epsilon y)\mu_2 + [p_3 - \epsilon(x + y)]\mu_3.\end{aligned}$$

Next, define the function

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

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 that there is some $z \in \mathbb{S}$ 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 strictly positive; in particular, f is globally nonzero. Since \mathbb{S} is connected, there is some $z \in \mathbb{S}$ which is a limit point of both $f^{-1}((0, \infty))$ and $f^{-1}((-\infty, 0))$. But then $-z$ shares this same property. Finally, upper semicontinuity tells us that 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, $\max V(\phi(z)) \geq \min V(\phi(-z))$ and $\max V(\phi(-z)) \geq \min V(\phi(z))$. Said differently (recall, V is convex-valued), $V(\phi(z)) \cap V(\phi(-z)) \neq \emptyset$. Lemma 1 then guarantees existence of an equilibrium generating information policy $\frac{1}{2}\delta_{\phi(z)} + \frac{1}{2}\delta_{\phi(-z)}$. \square

Now, we can prove the second part of Proposition 1

Proof. First, suppose $|\Theta| > 2$. Then, letting $E, F \subseteq \Theta$ be sufficiently small disjoint open neighborhoods around two distinct points in the support of μ_0 , we may assume $\mu_0(E), \mu_0(F), \mu(\Theta \setminus [E \cup F]) > 0$. Lemma 5 then yields an informative equilibrium.

For the rest of the proof, assume $|\Theta| = 2$. To conserve notation, assume without loss that $\Theta = \{0, 1\}$, and identify $\Delta\{0, 1\}$ with $[0, 1]$.

Suppose now that there is no $U_R \in \mathcal{U}_R$ with 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, there

exist some $\underline{\mu}, \underline{\nu} \in [0, \mu_0)$ and $\bar{\mu}, \bar{\nu} \in (\mu_0, 1]$ such that $\max V(\underline{\mu}) \geq \min V(\bar{\mu})$ and $\min V(\underline{\nu}) \leq \max V(\bar{\nu})$. Define $\phi : [0, 1] \rightarrow [0, 1]^2$ via $\phi(\lambda) := (1 - \lambda)(\underline{\mu}, \bar{\mu}) + \lambda(\underline{\nu}, \bar{\nu})$, and define $G : [0, 1] \rightrightarrows \mathbb{R}$ via $G(\lambda) := V(\phi_2(\lambda)) - V(\phi_1(\lambda))$. 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, there is an equilibrium with support $\{\phi_1(\lambda), \phi_2(\lambda)\}$, which is then informative.

Conversely, suppose $U_R \in \mathcal{U}_R$ is as described in the second part of the theorem; without loss, say $U_R(0) < U_R(1)$. 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, $U_R(\underline{\mu}) < U_R(\mu_0) < U_R(\bar{\mu})$, so that $V(\underline{\mu}) \cap V(\bar{\mu}) = \emptyset$. Lemma 1 then tells us that it cannot be an equilibrium policy. This completes the proof. \square

A.4 Proofs for Section 4

A.4.1 The quasiconcave envelope is well-defined

Lemma 6. *The function \bar{v} is well-defined. That is, there exists a pointwise lowest quasiconcave upper semicontinuous function which majorizes v .*

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

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

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

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

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

It is worth noting that, in the finite state case (just as in the case with commitment), the qualifier “upper semicontinuous” can be omitted.

Fact 1. *If $|\Theta| < \infty$, then \bar{v} is the pointwise lowest quasiconcave function which majorizes v .*

Proof. Notice that:

$$\begin{aligned} \bar{v}(\mu) &= \max \left\{ s : p \in \mathcal{I}(\mu), s \in \bigcap_{\nu \in \text{supp}(p)} V(\nu) \right\} \text{ (by Theorem 2 and Lemma 1)} \\ &= \max \left\{ s : \mu \in \overline{\text{co}} V^{-1}(s) \right\} \text{ (by Phelps (2001) Proposition 1.2)} \\ &= \max \left\{ s : \mu \in \text{co} \overline{V^{-1}(s)} \right\} \text{ (as } |\Theta| < \infty) \\ &= \max \left\{ s : \mu \in \text{co} V^{-1}(s) \right\} \text{ (as } V \text{ is upper hemicontinuous).} \end{aligned}$$

Suppose $\tilde{v} : \Delta\Theta \rightarrow \mathbb{R}$ is a quasiconcave function which majorizes v , and take any $\mu \in \Delta\Theta$. By the above, there exists some finite-support $p \in \mathcal{I}(\mu)$ such that $\bar{v}(\mu) \in V(\nu)$ for every $\nu \in \text{supp}(p)$. Moreover, there exists no $\epsilon > 0$ such that $\bar{v}(\mu) + \epsilon \in V(\nu)$ for every $\nu \in \text{supp}(p)$. As V is a Kakutani correspondence, it must be that $\max V(\nu^*) = \bar{v}(\mu)$ for some $\nu^* \in \text{supp}(p)$. In particular, $v(\nu^*) = \bar{v}(\mu) \leq v|_{\text{supp}(p)}$. Therefore:

$$\tilde{v}(\mu) = \tilde{v} \left(\sum_{\nu \in \text{supp}(p)} p(\nu)\nu \right) \geq \min_{\nu \in \text{supp}(p)} \tilde{v}(\nu) \geq \min_{\nu \in \text{supp}(p)} v(\nu) = \bar{v}(\mu).$$

So, in the finite state world, \bar{v} is indeed the pointwise lowest quasiconcave function above v . \square

A.4.2 Theorem 2

Proof. Let $V^* : \Delta\Theta \rightrightarrows \mathbb{R}$ map any given prior to the associated set of equilibrium values for S , and let $v^* := \sup V^*$. Our goal is to show that this supremum is attained and that $v^* = \bar{v}$.

First notice that $v \in V^*$, witnessed by an uninformative equilibrium. Therefore, V^* is a nonempty-valued correspondence, and $v^* \geq v$.

Now, consider any sequence $(\mu_n, s_n)_n$ from the graph of V^* which converges to some $(\mu, s) \in (\Delta\Theta) \times \mathbb{R}$. For each n , Theorem 1 delivers some $p_n \in \mathcal{I}(\mu_n)$ such that $v(\nu) \geq s_n$ for p_n -almost every $\nu \in \Delta\Theta$. By compactness, some subsequence of $(p_n)_n$ converges to some $p \in \Delta\Delta\Theta$. The correspondence \mathcal{I} is upper hemicontinuous, so that $p \in \mathcal{I}(\mu)$. Upper semicontinuity of v then tells us that $v(\nu) \geq s$ for p -almost every $\nu \in \Delta\Theta$. Appealing to Theorem 1 again then tells us that $s \in V^*(\mu)$. Therefore, V^* has closed graph: it is an upper hemicontinuous, compact-valued correspondence. As two immediate by-products, we learn that an optimal information policy exists for every prior (as V^* is compact-valued), and that v^* is an upper semicontinuous function.

Next, consider any $\mu, \mu' \in \Delta\Theta$ and any $\lambda \in (0, 1)$, Theorem 1 implies that, for every $\nu \in \Delta\Theta$,

$$v^*(\nu) = \max \{s \in \mathbb{R} : v|_{\text{supp}(p)} \geq s \exists p \in \mathcal{I}(\nu)\} = \max_{p \in \mathcal{I}(\nu)} \inf v(\text{supp}(p)).$$

In particular,

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

That is, v^* is quasiconcave. Therefore, $v^* \geq \bar{v}$.

Finally, take any $\mu \in \Delta\Theta$. By Theorem 1, there is some $p \in \mathcal{I}(\mu)$ such that $v(\nu) \geq v^*(\mu)$ for p -almost every $\nu \in \Delta\Theta$. Therefore, for such ν , we have: $v^*(\mu) \leq v(\nu) \leq \bar{v}(\nu)$. Letting $D \subseteq \Delta\Theta$ denote the support of p , upper semicontinuity of \bar{v} implies $v^*(\mu) \leq \bar{v}|_D$. But then, as \bar{v} is quasiconcave and upper semicontinuous, $v^*(\mu) \leq \bar{v}|_{\overline{\text{co}}D}$. As p witnesses $\mu \in \overline{\text{co}}D$, we learn that $v^*(\mu) \leq \bar{v}(\mu)$. \square

A.5 Proofs for Section 5

A.5.1 Theorem 3

Proof. As (2) and (3) both 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 $D := \text{supp}(p) \subseteq \Delta\Theta$.³² Lemma 4 then guarantees that (p, s) is an equilibrium outcome, say witnessed by equilibrium $\mathcal{E}_1 = (\sigma_1, \rho_1, \beta_1)$. Letting $\alpha = \alpha_s : D \rightarrow \Delta A$ be as delivered by Lemma 2, we may assume that $\rho_1(\cdot|m) = \alpha(\cdot|\beta(m))$. In particular, ρ_1 specifies finite-support play for every message.

Let $\mathbb{M} := \text{marg}_M \mathbb{P}_{\mathcal{E}_1}$ and $X := \text{supp}[\mathbb{M} \circ \hat{\rho}^{-1}] \subseteq \Delta A$, and fix arbitrary $(\hat{\alpha}, \hat{\mu}) \in \text{supp}[\mathbb{M} \circ (\rho_1, \beta_1)^{-1}]$; 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 $\rho' : M \rightarrow \Delta A$ (resp. $\beta' : M \rightarrow \Delta\Theta$) to agree with ρ_1 (β_1) on-path and take value $\hat{\alpha}$ ($\hat{\mu}$) off-path, there is then an equilibrium $\mathcal{E}' = (\sigma_1, \rho', \beta')$ such that $\mathbb{P}_{\mathcal{E}'} = \mathbb{P}_{\mathcal{E}_1}$ and $\rho'(\cdot|m) \in X$ for every $m \in M$.

Now, define:

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

By construction, $(\sigma_2, \rho_2, \beta_2)$ is an equilibrium generating outcome (p, s) ,³³ proving that (1) implies (2).

Now, define the (A - and D -valued, respectively) random variables $\mathbf{a}, \boldsymbol{\mu}$ on $\langle D, p \rangle$ by letting $\mathbf{a}(\mu) := \max \text{supp}[\alpha(\mu)]$ and $\boldsymbol{\mu}(\mu) := \mu$ for $\mu \in D$. Next, define the conditional expectation $\mathbf{f} := \mathbb{E}_p[\boldsymbol{\mu}|\mathbf{a}] : D \rightarrow D$, 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 $D = \text{supp}(p \circ \mathbf{f}^{-1})$, which

³²Some 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, since the Blackwell order is closed-continuous, a Blackwell-minimal such policy exists.

³³It generates (\tilde{p}, s) for some garbling \tilde{p} of p . Minimality of p then implies $\tilde{p} = p$.

implies that $p \circ \mathbf{f}^{-1}$ secures s . But then, minimality of p implies that $p \circ \mathbf{f}^{-1} = p$. So $\mathbf{f} = \mathbb{E}_p[\boldsymbol{\mu}|\mathbf{a}]$ and $\boldsymbol{\mu}$ have the same distribution, which implies that $\mathbf{f} = \boldsymbol{\mu}$ a.s.- p . By definition, \mathbf{f} is \mathbf{a} -measurable, so that Doob-Dynkin tells us there is some measurable $\mathbf{b} : A \rightarrow D$ such that $\mathbf{f} = \mathbf{b} \circ \mathbf{a}$.

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

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

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

A.5.2 Proposition 2

Proof. Suppose cheap talk is costless, i.e., $\bar{v}(\mu_0) = \hat{v}(\mu_0)$. Letting (p, s) be the outcome induced by the equilibrium (σ, ρ, β) , it must be that $\int_{\Delta \Theta} v \, dp = s$. This implies that $v(\mu) = s$ for p -almost every μ , and therefore (by upper semicontinuity) for every $\mu \in D := \text{supp}(p)$. Now, define the measurable function,

$$\begin{aligned} f : D &\rightarrow \mathbb{R} \\ \mu &\mapsto \begin{cases} s & : v(\mu) = s \\ \min V(\mu) & : v(\mu) > s. \end{cases} \end{aligned}$$

Let $\alpha = \alpha_f : D \rightarrow \Delta A$ be as given by Lemma 2. By construction, $\text{supp}(\alpha(\mu)) = 1$ for every $\mu \in B$. Now, define a new receiver strategy,

$$\begin{aligned} \tilde{\rho} : M &\rightarrow \Delta\Theta \\ m &\mapsto \begin{cases} \alpha(\beta(m)) & : \beta(m) \in D, \\ \rho(m) & : \text{otherwise.} \end{cases} \end{aligned}$$

By construction, $(\sigma, \tilde{\rho}, \beta)$ is an equilibrium with no receiver mixing. Moreover, it generates value s and is therefore most persuasive. \square

A.5.3 Corollary 2

Proof. Take u_S to be one-to-one.

First, suppose that there is a most persuasive equilibrium (σ, ρ, β) with no mixing, which induces an outcome (p, s) .

Let $A^* := \{a \in A : \delta_a \in \text{supp}[\text{marg}_{\Delta A} \mathbb{P}]\}$. By S incentive-compatibility, $u_S|_{A^*}$ is constant. But then, $A^* = \{a\}$ for some action a , as u_S is injective. By receiver incentive-compatibility, $a \in \arg \max_{\hat{a} \in A} u_R(\hat{a}, \mu)$ for every p -almost every μ . Therefore, $a \in \arg \max_{\hat{a} \in A} u_R(\hat{a}, \mu_0)$ too. This implies that $s \in V(\mu_0)$, so that an uninformative equilibrium is most persuasive, i.e., cheap talk is not valuable.

Appealing to Proposition 2, if cheap talk is costless, it is not valuable. That is, if $\hat{v}(\mu_0) = \bar{v}(\mu_0)$, then $\bar{v}(\mu_0) = v(\mu_0)$. The result follows. \square

A.6 Proofs for Section 6

A.6.1 Proposition 3

Proof. Suppose \hat{p} is a most persuasive commitment policy that is supported on two consecutive actions. Let $D = \text{supp}(\hat{p})$. If $v(D) = \{s\}$, then Lemma 1 directly applies to tell us (\hat{p}, s) is an equilibrium outcome. Now, suppose that $v(D) = \{s, s'\}$ for $s > s'$. Let:

$$\begin{aligned} D_0 &= \{\mu \in B : V(\mu) = \{s\}\} \\ D_1 &= \{\mu \in B : s' \in V(\mu)\}. \end{aligned}$$

As s and s' are consecutive, $s > s'$, and V convex-valued, it must be that $D_0 \cup D_1 = D$.

Suppose now, for a contradiction, that $D_0 \neq \emptyset$. For every $\epsilon \in (0, 1)$, define

$$\begin{aligned} D_\epsilon &:= \{\mu \in D_0 : V((1 - \epsilon)\mu + \epsilon\mu_0) \subseteq (s', \infty)\} \\ &= \{\mu \in D_0 : V((1 - \epsilon)\mu + \epsilon\mu_0) \subseteq [s, \infty)\}. \end{aligned}$$

Upper hemicontinuity of V implies that $\hat{p}(D_\epsilon) > 0$ for sufficiently small $\epsilon \in (0, 1)$. Let $\lambda := 1 - \epsilon[1 - \bar{p}(D_\epsilon)] \in (0, 1)$, and define $p := \frac{1}{\lambda}[(1 - \epsilon)\hat{p}(\cdot \setminus D_\epsilon) + \hat{p}(\cdot \cap D_\epsilon)] \in \Delta\Delta\Theta$. Next, let

$$\begin{aligned} f &: \Delta\Theta \rightarrow \Delta\Theta \\ \mu &\mapsto \mu + \epsilon \mathbf{1}_{\mu \in D_\epsilon} (\mu_0 - \mu). \end{aligned}$$

For $g \in \{v, \text{id}_{\Delta\Theta}\}$, we have

$$\begin{aligned} \int_{\Delta\Theta} g \, d[p \circ f^{-1}] &= \int_{(\Delta\Theta) \setminus D_\epsilon} g \, dp + \int_{D_\epsilon} g((1 - \epsilon)\mu + \epsilon\mu_0) \, dp(\mu) \\ &= \frac{1 - \epsilon}{\lambda} \int_{(\Delta\Theta) \setminus D_\epsilon} g \, d\hat{p} + \frac{1}{\lambda} \int_{D_\epsilon} g((1 - \epsilon)\mu + \epsilon\mu_0) \, d\hat{p}(\mu) \\ \implies \frac{\lambda}{\bar{p}(D_\epsilon)} \int_{\Delta\Theta} g \, d(p \circ f^{-1} - d\hat{p}) &= \frac{1}{\bar{p}(D_\epsilon)} \int_{D_\epsilon} [g((1 - \epsilon)\mu + \epsilon\mu_0) - (1 - \epsilon)g(\mu)] \, d\hat{p}(\mu) \\ &\quad - \epsilon \int_{\Delta\Theta} g \, d\hat{p}. \end{aligned}$$

This yields two by-products. First, it tells us that $p \circ f^{-1} \in \mathcal{I}(\mu_0)$, as

$$\begin{aligned} \frac{\lambda}{\bar{p}(D_\epsilon)} \left[\int_{\Delta\Theta} \mu \, d[p \circ f^{-1}](\mu) - \mu_0 \right] &= \frac{1}{\bar{p}(D_\epsilon)} \int_{D_\epsilon} [((1 - \epsilon)\mu + \epsilon\mu_0) - (1 - \epsilon)\mu] \, d\hat{p}(\mu) \\ &\quad - \epsilon \int_{\Delta\Theta} \mu \, d\hat{p}(\mu) \\ &= \epsilon\mu_0 - \epsilon\mu_0 = 0. \end{aligned}$$

Next, it tells us that $\int_{\Delta\Theta} v d[p \circ f^{-1}] > \int_{\Delta\Theta} v d\hat{p}$, as

$$\begin{aligned} \frac{\lambda}{\hat{p}(D_\epsilon)} \int_{\Delta\Theta} v d(p \circ f^{-1} - \hat{p}) &= \frac{1}{\hat{p}(D_\epsilon)} \int_{D_\epsilon} [v((1-\epsilon)\mu + \epsilon\mu_0) - (1-\epsilon)v(\mu)] d\hat{p}(\mu) \\ &\quad - \epsilon \int_{\Delta\Theta} v d\hat{p} \\ &\geq \frac{1}{\hat{p}(D_\epsilon)} \int_{B_\epsilon} [s - (1-\epsilon)v(\mu)] d\hat{p}(\mu) - \epsilon \int_{\Delta\Theta} v d\hat{p} \\ &> \frac{1}{\hat{p}(D_\epsilon)} \int_{D_\epsilon} [s - (1-\epsilon)s] d\hat{p}(\mu) - \epsilon s = 0. \end{aligned}$$

This contradicts \hat{p} being optimal under commitment.

Therefore $p(D_1) = 1$, and so (\hat{p}, s') is an equilibrium outcome, by Lemma 1.

Now, as v is $u_S(A)$ -valued, and s' and s are consecutive, it follows that any policy which secures S a payoff strictly greater than s' must in fact secure s . But no such policy can exist, since \hat{p} is a most persuasive commitment policy which generates a payoff strictly below s . With no payoff strictly greater than s' being securable, Theorem 1 tells us that \hat{p} is a most persuasive equilibrium policy. \square

A.7 Proofs for Section 7

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

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

We now invest in some notation. For $\theta \in \Theta$ and $k \in \{1, \dots, n\}$, let $\theta_{k,n}^{(1)} := \max_{i \in \{k, \dots, n\}} \theta_i$ 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 $k \in \{1, \dots, n\}$, let:

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

be the S strategy which reports the best reform from among those preferred to k by the

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

Claim 1. Fix $k \in \{1, \dots, n\}$ and $i \in \{k, \dots, n\}$. 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 $i \in \{k, \dots, n\}$, 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 $j \in \{1, \dots, k-1\}$.
- (2) $\mathbb{E}_0 \theta_i \in \text{co} \{ \mathbb{E}_i^k \theta_i, \mathbb{E}_i^k \theta_j \}$ for $j \in \{k, \dots, n\} \setminus \{i\}$.
- (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 that $\mathbb{E}_i^k \theta_i \geq \mathbb{E}_i^k \theta_j$ for $j \in \{1, \dots, n\} \setminus \{i\}$. This implies that $i \in \arg \max_{a \in A} u_R(a, \beta_k(i))$ if and only if $\mathbb{E}_i^k \theta_i \geq c$. The fourth fact completes the proof of the claim. \square

By Theorem 1 (as described in the main text), the S-optimal equilibrium payoff is $u_S(k^*)$, where

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

In fact, more can be said from the constructive proof of Theorem 1. We can 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 $\sigma^* := (1 - \epsilon)\sigma_{k^*} + \epsilon \mathcal{U} \{k^*, \dots, n\}$. 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 $i \in \{k, \dots, n\}$, 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 that such play is in fact equilibrium play.

A.8 Proofs for Section 8

A.8.1 Proposition 4

Toward proving Proposition 4, we introduce some notation. For each continuous $f : \Delta\Theta \rightarrow \mathbb{R}$ 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 7. *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 $\{\succsim_{\tilde{u}_S(\cdot, \theta), X}\}_{\theta \in \Theta}$ all coincide, where $X := \bigcup_{\mu \in \Delta\Theta} \arg \max_{\alpha \in \Delta A} \tilde{u}_S(\alpha, \mu)$.*

Notice that the lemma applies directly to the two cases of Proposition 4. 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 $\alpha, \alpha' \in X$. As $\{\tilde{u}_S(\alpha, \theta) - \tilde{u}_S(\alpha', \theta)\}_{\theta \in \Theta}$ are all of the same sign and integration is monotone, $u_S(\alpha) - u_S(\alpha')$ is of the same sign. That is, $\{\succsim_{\tilde{u}_S(\cdot, \theta), X}\}_{\theta \in \Theta} = \{\succsim_{u_S, X}\}$.

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 that $\rho(M) \subseteq X$. Therefore, $\{\succsim_{\tilde{u}_S(\cdot, \theta), \rho(M)}\}_{\theta \in \Theta} = \{\succsim_{u_S, \rho(M)}\}$. 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:

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

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

A.8.2 Proposition 5

Proof. By Theorem 2, there is an information policy p such that $(p, \bar{v}(\mu_0))$ is an equilibrium outcome.

For any $s \in [v(\mu_0), \bar{v}(\mu_0)]$, notice that p secures s . Thus we apply Lemma 4 to arrive at an information policy p_s^* such that (p_s^*, s) is an equilibrium outcome. Let $\lambda_{p,s}$ be as described in the proof of Lemma 4. By construction, whenever $v(\mu_0) \leq s' \leq s \leq \bar{v}(\mu_0)$, we have $\lambda_{p,s'}(\cdot) = \lambda_{p,s}(\cdot)\lambda_{p_s^*,s'}(\cdot)$. The second and third point then follow directly as in Lemma 4. Finally, the first point follows from the second, as a less informative information policy generates a lower receiver value. \square