

Perfect Bayesian Persuasion*

Elliot Lipnowski[†]
Columbia University

Doron Ravid
University of Chicago

Denis Shishkin
UC San Diego

September 12, 2023

Abstract

A sender commits to an experiment to persuade a receiver. Accounting for the sender’s experiment-choice incentives, and not presupposing a receiver tie-breaking rule when indifferent, we characterize when the sender’s equilibrium payoff is unique and so coincides with her “Bayesian persuasion” value. A sufficient condition in finite models is that every action which is receiver-optimal at some belief is uniquely optimal at some other belief—a generic property. We similarly show the equilibrium sender payoff is typically unique in ordered models. In an extension, we show uniqueness generates robustness to imperfect sender commitment.

*Lipnowski and Ravid acknowledge support from the National Science Foundation (grant SES-1730168). Lipnowski acknowledges support from the Columbia Program for Economic Research. We would like to thank Laura Doval, Christoph Kuzmics, and Stephen Morris for helpful comments. Zikai Xu provided excellent research assistance.

[†]e.lipnowski@columbia.edu, dravid@uchicago.edu, dshishkin@ucsd.edu

In recent years, many studies have explored variations on the Bayesian Persuasion model of Kamenica and Gentzkow (2011), hereafter KG (see Bergemann and Morris, 2019; Kamenica, 2019, for review). In KG’s model, a receiver (R, he) must choose an action, but a sender (S, she) controls R’s available information about a payoff-relevant state. When analyzing this model and its variations, researchers usually assume S-favorable tie-breaking. Whereas this assumption is without loss when evaluating the highest payoff S can get across all information structures, making this assumption when using the model for prediction or design is tantamount to selecting S’s favorite equilibrium. Consequently, it is unclear whether the model’s conclusions remain valid under other equilibrium selection criteria. In this paper, we address this question by asking when tie-breaking assumptions are irrelevant. More specifically, we ask: when is S’s equilibrium payoff unique?

To answer this question, we study KG’s model without imposing S favorable tie breaking. Thus, the game begins with S publicly committing to a Blackwell experiment about the payoff-relevant state $\theta \in \Theta$, whose prior distribution is $\mu_0 \in \Delta\Theta$. Formally, S chooses a measurable function $\psi: \Theta \rightarrow \Delta M$ for a given space M of messages. Then, after observing both ψ and a realized message $m \in M$, the receiver chooses an action a from a set of feasible actions A . Each player $i \in \{S, R\}$ seeks to maximize the expectation of an objective $u_i(a, \theta)$.¹ KG analyze this game using the S-favorite equilibrium concept. We diverge from this assumption by considering all equilibria of the game, and ask when all these equilibria give S the same payoff.

We defer a formal definition of equilibrium to the following section, but a brief overview of KG’s analysis is in order. They adopt a belief-based approach, casting S’s optimization problem as one of directly choosing $p \in \Delta\Delta\Theta$, the ex-ante distribution of R’s posterior belief μ concerning the state. Because R is Bayesian (and S’s experiment choice is made in ignorance of the state), it must be that p belongs to $\mathcal{I}(\mu_0)$, the set of belief distributions with barycenter equal to the prior; KG term this condition Bayes plausibility. We refer to elements of $\mathcal{I}(\mu_0)$ as **information policies**. But what payoff does S derive from a given Bayes-plausible belief distribution? By R rationality, R will choose from his best responses $A_R^*(\mu) \subseteq A$ whenever her posterior belief is μ . S’s expected payoff from such an action $a \in A_R^*(\mu)$ is then equal to $\int u_S(a, \cdot) d\mu$. Given KG’s focus on S-optimal equilibrium, they can assume without loss that R always breaks any indifferences in S’s favor. Thus, KG can

¹We maintain the technical assumptions that both A and Θ are nontrivial compact metrizable spaces; that the objectives $u_S, u_R: A \times \Theta \rightarrow \mathbb{R}$ are continuous; that M is a Polish space; and either that M is uncountable or that Θ is finite and $|M| \geq |\Theta|$. All topological spaces are viewed as measurable spaces with their Borel sigma-algebra.

summarize S’s payoff from inducing belief μ as $v(\mu) := \max_{a \in A_R^*(\mu)} \int u_S(a, \cdot) d\mu$. We call v the *value function*. Hence, S’s best equilibrium value is given by $\hat{v}(\mu_0) = \max_{p \in \mathcal{I}(\mu_0)} \int v dp$.

But what happens if R may choose best responses that differ from those S would prefer? If, in the worst case, R always chooses S’s least favorite of his best responses, then from inducing R belief μ , S can only expect a payoff of $w(\mu) := \min_{a \in A_R^*(\mu)} \int u_S(a, \cdot) d\mu$. Accordingly, S would have a profitable deviation if her payoff were ever strictly below

$$\hat{w}(\mu_0) = \sup_{p \in \mathcal{I}(\mu_0)} \int w dp.$$

As Wu (2022) has shown, this payoff lower bound is the only additional constraint imposed by S’s experiment-choice incentives. Using this result, we go on to fully characterize when S has a unique equilibrium payoff, and to provide meaningful sufficient conditions for the same.

We also apply our results to a model in which there are doubts about S’s ability to commit to her experiment. Such doubts make the equilibrium selection question salient, since they suggest S may lack the bargaining power required to dictate which equilibrium is played. Section 3 introduces a model in which S’s commitment power is limited. Specifically, we add a probability that S can covertly manipulate the experiment’s results ex-post. We characterize S’s equilibrium payoff set as this probability vanishes, and apply our uniqueness results to study robustness in that setting.

Related literature We contribute to the Bayesian persuasion literature (Aumann and Maschler, 1995; Kamenica and Gentzkow, 2011; Kamenica, 2019), which studies sender-receiver games in which a sender commits to an information-transmission strategy. Our main goal is to understand when KG’s sender payoff characterization still applies—taking her experiment-choice incentives into account, and allowing the receiver to choose actions in a manner unfavorable to the sender when he is indifferent. Our extension to the case in which the sender’s commitment is imperfect extends the analysis of Lipnowski, Ravid, and Shishkin (2022) in the same manner.²

Several papers introduce robustness considerations into information design problems (e.g., Kosterina, 2022; Dworzak and Pavan, 2022) Within this literature, most relevant are papers that study optimal information provision in multi-agent settings under planner-worst

²See also Min (2021), which develops a generalization of the limited-commitment model and shows some credibility is better than no credibility in a leading example; and Fréchette, Lizzeri, and Perego (2022), which studies communication outcomes in a laboratory experiment.

equilibrium selection (e.g. Mathevet, Perego, and Taneva, 2020; Moriya and Yamashita, 2020; Oyama and Takahashi, 2020; Ziegler, 2020; Halac, Lipnowski, and Rappoport, 2022; Morris, Oyama, and Takahashi, 2022; Li, Song, and Zhao, 2023; Inostroza and Pavan, 2023). Our work provides conditions under which such a selection rule is irrelevant for the single-agent case.

Our work is related to multiple distinct strands of the Bayesian persuasion literature that explicitly account for sender incentives in choosing an experiment. The first is the literature on competition in persuasion, in which multiple senders design flexible information either simultaneously (Gentzkow and Kamenica, 2016, 2017; Au and Kawai, 2020; Ravindran and Cui, 2022) or sequentially (Li and Norman, 2021; Wu, 2022) and must make individually rational experiment choices in equilibrium. The most related paper from this set is (Wu, 2022), which essentially proves Proposition 0. The second strand studies how experiment constraints shape chosen information (e.g., Ichihashi, 2019; Perez-Richet and Skreta, 2022), whereas the third strand studies signaling and informed principal problems for information design settings in which an experiment choice can reveal private information (e.g., Perez-Richet, 2014; Hedlund, 2017; Alonso and Câmara, 2018; Koessler and Skreta, 2022).³ Finally, beyond the Bayesian persuasion literature, work on sequential mechanism design under limited commitment (e.g., Skreta, 2006; Doval and Skreta, 2022) accounts for a principal’s incentives while her future beliefs play a prominent role in the analysis.

1. The Equilibrium Payoff Set

We now formally define an equilibrium concept for the persuasion game. Note two features of the definition. First, while R must respond optimally to his belief, we make no direct assumption concerning which best response he chooses when indifferent.⁴ Second we explicitly include an optimality condition for S at the experiment-choice stage. Note that the latter condition would have no bite under S-favorable tie-breaking, and so is rarely included in the literature.

Definition 1. *Let Ψ denote the set of all measurable functions $\tilde{\psi}: \Theta \rightarrow \Delta M$ (a.k.a. experiments), which we view as a measurable space endowed with the discrete σ -algebra.*

³In a sense, the literature on verifiable disclosure can be seen as a combination of the second and third strand, with sender incentives being a key object of study and verifiable information being limited in a type-dependent manner.

⁴Still, as we shall see, it will often be the case that *mutual* best response requires that R break indifferences in S’s favor on the path of play, just as a recipient of a zero offer accepts the offer in the unique subgame-perfect equilibrium of the ultimatum game.

A sender strategy is an experiment $\psi \in \Psi$; a receiver strategy is a measurable function $\alpha: M \times \Psi \rightarrow \Delta A$; and a receiver belief map is a measurable function $\pi: M \times \Psi \rightarrow \Delta \Theta$. A **(perfect Bayesian) equilibrium** is a triple of such maps $\langle \psi, \alpha, \pi \rangle$ such that

1. The sender's choice satisfies

$$\psi \in \operatorname{argmax}_{\tilde{\psi} \in \Psi} \int_{\Theta} \int_M \int_A u_S(a, \theta) \, d\alpha(a|\tilde{\psi}, m) \, d\tilde{\psi}(m|\theta) \, d\mu_0(\theta);$$

2. Every $\tilde{\psi} \in \Psi$ and $m \in M$ have

$$\alpha \left(\operatorname{argmax}_{a \in A} \int_{\Theta} u_R(a, \theta) \, d\pi(\theta|m, \tilde{\psi}) \mid m, \tilde{\psi} \right) = 1;$$

3. Every $\tilde{\psi} \in \Psi$, Borel $\hat{M} \subseteq M$, and Borel $\hat{\Theta} \subseteq \Theta$ have

$$\int_{\Theta} \int_{\hat{M}} \pi(\hat{\Theta}|m, \tilde{\psi}) \, d\tilde{\psi}(m|\theta) \, d\mu_0(\theta) = \int_{\hat{\Theta}} \tilde{\psi}(\hat{M}|\theta) \, d\mu_0(\theta).$$

In such a case, we say the induced **equilibrium sender payoff** is

$$\int_{\Theta} \int_M \int_A u_S(a, \theta) \, d\alpha(a|m, \psi) \, d\psi(m|\theta) \, d\mu_0(\theta).$$

The interpretation is as follows. First, S publicly chooses an experiment $\tilde{\psi} \in \Psi$.⁵ The experiment then produces a message $m \in M$ that R observes. Then, R updates his beliefs according to the message and the chosen experiment, and chooses an action $a \in A$. We require that S only choose experiments that maximize her expected payoffs, that R (having seen the realized experiment and message) only choose actions that maximize his expected payoffs with respect to his belief about the payoff state, and that R's beliefs conform to Bayesian updating.⁶

In what follows, we document the set of attainable equilibrium S payoffs, with a particular focus on understanding when it is unique.

⁵One could easily extend the model to allow S to mix over experiment choice. Doing so would entail added notational burden but would have no effect on the resulting S payoff set because the experiment choice is public, not informed by private information, and not simultaneous to any other decisions.

⁶Moreover, we assume that S cannot signal what she does not know. Indeed, our Bayesian condition implies that every $\tilde{\psi} \in \Psi$ and Borel $\hat{\Theta} \subseteq \Theta$ have $\int_{\Theta} \int_M \pi(\hat{\Theta}|m, \tilde{\psi}) \, d\tilde{\psi}(m|\theta) \, d\mu_0(\theta) = \mu_0(\hat{\Theta})$, so that the chosen experiment alone does not cause belief updating by R about the payoff state.

Remark 1. *Although our focus is on equilibrium S payoffs rather than behavior, our results have natural implications for behavior as well. In particular, when S’s equilibrium payoff is unique, our results imply that R breaks indifferences in S’s favor with probability 1 on path in every equilibrium. Hence, in this case, the results of KG (and many subsequent papers surveyed in Kamenica, 2019) are robust to allowing arbitrary tie-breaking for R and accounting for S’s experiment-choice incentives.*

1.1. Characterizing equilibrium payoffs

We begin by stating a characterization of the equilibrium S payoff set as a function of the parameters of our game. This set is a compact interval, with highest value equal to KG’s commitment solution, and lowest value equal to the supremum value S can guarantee when R breaks her indifferences adversarially. In the special case in which the state space is finite, this result is exactly Proposition 1 from Wu (2022). Although no substantive new arguments are required for the general case, we include a proof for the sake of completeness.

Proposition 0 (Payoff set (Wu, 2022)). *The set of equilibrium S payoffs is $[\hat{w}(\mu_0), \hat{v}(\mu_0)]$.*

Necessity is essentially immediate, and the proof of sufficiency is constructive. By degrading information from an S-optimal (under favorable tie-breaking) experiment and allowing for R to mix among optimal choices in the degraded experiment, one can find an experiment for S to choose and R best response to target any payoff in the given interval. Then, having R break indifference adversarially to S following off-path experiment choices ensures that this experiment choice is indeed optimal for S.

Remark 2. *It is apparent that Proposition 0 depends only on the value correspondence $V = [w, v]$, and moreover (as is clear from our proof) the only substantive property required of the environment is that the attainable S payoffs from R responding optimally to a given belief be convex.⁷ In addition to making Proposition 0 more tractable to apply, this feature also expands its applicability beyond the basic model we have considered. For example, the proposition can be applied to settings in which a receiver is subject to independent private payoff shocks. Additionally, the proposition applies to public persuasion of a set of agents who play a game, so long as the set of induced payoffs for the sender is convex for every public belief. The latter condition holds, for instance, if the receivers observe a rich public randomization device after the experiment choice (but before their gameplay).*

⁷Our analysis also uses the fact that V is nonempty-compact-valued and upper hemicontinuous, and that the set of optimal R choices is a weakly measurable correspondence of his belief.

2. Equilibrium Payoff Uniqueness

In this section we ask, when does S have a unique equilibrium payoff? Whenever she does, the traditional analysis that focuses on S -optimal equilibrium (and so assumes S -favorable tie-breaking by R) is essentially without loss.

As a starting observation, because uniqueness follows directly from Proposition 0 whenever $v = w$, a sufficient condition for S to have a unique equilibrium payoff is immediate.

Corollary 1 (No relevant ties). *S has a unique equilibrium payoff if, at any belief, S is indifferent between all of R 's best responses.*

Although restrictive, the above condition nevertheless captures many cases of interest. For example, if the action space is a convex subset of some linear space with R 's payoff being strictly concave in his action (e.g., Crawford and Sobel, 1982; Chakraborty and Harbaugh, 2010), then he has a unique best response to every belief, and so the corollary applies.

The following result gives an alternative sufficient condition for S to attain her Bayesian persuasion value in all equilibria. It says such uniqueness holds if information can serve as a stand-in for favorable tie-breaking at all relevant posterior beliefs. To state the result, the following definition is useful.

Definition 2. *Say a set $D \subseteq \Delta\Theta$ is **persuasion sufficient** if*

$$\hat{v}(\mu_0) = \sup \left\{ \int v \, dp : p \in \mathcal{I}(\mu_0) \text{ s.t. } p(D') = 1 \text{ for some Borel } D' \subseteq D \right\}$$

In words, a set of beliefs D is persuasion sufficient if the KG payoff can be approximated with policies supported within D under S -favorable tie-breaking.

As the following proposition says, if (on the relevant set of beliefs) information can be used to replace favorable selection, then S guarantees her KG payoff.

Proposition 1 (Information as selection). *S has a unique equilibrium payoff if $\hat{w}|_D \geq v|_D$ for some persuasion-sufficient $D \subseteq \Delta\Theta$.⁸*

To prove the proposition, we begin with an information policy supported on D that nearly attains the KG payoff under S -favorable selection. For each realized posterior $\mu \in D$, the sender then provides additional information—further splitting the belief μ into some belief

⁸For a function $f: \Delta\Theta \rightarrow \mathbb{R}$, we let $f|_D$ be the restriction of f to D ; that is, the function from D to \mathbb{R} defined via $\mu \mapsto f(\mu)$.

distribution centered on μ —to ensure a payoff of nearly $v(\mu)$ even under S-adversarial selection. That $\hat{w} \geq v$ on D ensures one can find such a splitting for each μ , and a standard measurable selection result says one can do so measurably.

Following directly from the above proposition, the next corollary exactly characterizes when equilibrium S payoff uniqueness holds independent of the prior: such uniqueness is equivalent to information always replicating favorable tie-breaking.

Corollary 2 (Global uniqueness). *S has a unique equilibrium payoff for every prior (holding other parameters fixed) if and only if $\hat{w} \geq v$.*

2.1. Direct sufficient conditions for uniqueness

The above conditions for payoff uniqueness were expressed in terms of the derived objects v , w , \hat{v} , and \hat{w} . We now develop sufficient conditions on the game's primitives that ensure S has a unique equilibrium payoff.

Let $A^U \subseteq A$ denote the set of potentially unique best response actions, that is, $a \in A$ such that some $\mu \in \Delta[\text{supp } \mu_0]$ has $A_R^*(\mu) = \{a\}$. For any belief $\bar{\mu} \in \Delta\Theta$, let $A^U(\bar{\mu})$ denote the set of potentially unique best response actions at $\bar{\mu}$, that is, $a \in A$ such that some $\mu \in \Delta\Theta$ has both $A_R^*(\mu) = \{a\}$ and $\varepsilon\mu \leq \bar{\mu}$ for some $\varepsilon > 0$.⁹

Definition 3. *Given $\mu \in \Delta\Theta$, say **the potentially unique best response (PUBR) property holds at μ** if¹⁰*

$$\max_{a \in A_R^*(\mu)} \int u_S(a, \cdot) d\mu = \sup_{a \in A_R^*(\mu) \cap A^U} \int u_S(a, \cdot) d\mu;$$

*and say **the strong PUBR property holds at μ** if*

$$\max_{a \in A_R^*(\mu)} \int u_S(a, \cdot) d\mu = \sup_{a \in A_R^*(\mu) \cap A^U(\mu)} \int u_S(a, \cdot) d\mu.$$

*Given a set $D \subseteq \Delta\Theta$, say **the (strong) PUBR property holds on D** if it holds at every $\mu \in D$.*

The above property says that S can obtain a value arbitrarily close to her payoff under favorable-selection with the receiver only using actions that are a unique best response to some belief—and the strong version further requires that the latter belief be boundedly

⁹When we refer to inequalities for measures, we interpret the inequality pointwise. For instance, $\varepsilon\mu \leq \bar{\mu}$ means $\varepsilon\mu(\hat{\Theta}) \leq \bar{\mu}(\hat{\Theta})$ for every measurable $\hat{\Theta} \subseteq \Theta$.

¹⁰Note the left-hand side is just $v(\mu)$.

absolutely continuous with respect to the relevant posterior belief.¹¹

The next result shows various versions of the PUBR property are sufficient to guarantee equilibrium selection.

Theorem 1 (The PUBR theorem¹²). *S has a unique equilibrium payoff if either of the following two conditions holds for some persuasion-sufficient $D \subseteq \Delta\Theta$:*

- (i) *The strong PUBR property holds on D.*
- (ii) *The PUBR property holds on D, and either A or Θ is finite.*

In particular, for full-support priors and finite states or finite actions, the theorem implies S has a unique equilibrium payoff if every action is the unique best response to some belief.

For intuition, consider the case in which the state and action spaces are both finite and the prior has full support, and let p be any Bayes-plausible belief distribution with finite support \tilde{D} satisfying the PUBR property. For any supported belief μ , the PUBR property implies some action a is both (i) an S-preferred R best response to belief μ and (ii) a unique R best response to some alternative belief μ'_μ . For any $\lambda \in (0, 1)$, we can then define the belief $\mu'_{\mu,\lambda} = \lambda\mu'_\mu + (1 - \lambda)\mu$, which also has a as a unique R best response by the linearity property of expected utility. Define now the belief distribution p_λ which modifies p by replacing each supported μ with $\mu'_{\mu,\lambda}$. This belief distribution is Bayes-plausible for $\mu_\lambda = (1 - \lambda)\mu_0 + \lambda \sum_\mu p(\mu)\mu'_\mu$, which converges to μ_0 as $\lambda \rightarrow 0$. Moreover, because R has a unique best response to every p_λ -supported message, it follows that $\int w dp_\lambda$ converges to $\int v dp$ as $\lambda \rightarrow 0$. Therefore, every limit point of $\hat{w}(\mu_\lambda)$ as $\lambda \rightarrow 0$ is at least $\int v dp$. Finally, the function \hat{w} is concave—for the exact same reason that the optimal value \hat{v} is—and so is continuous at the relatively interior point μ_0 of its domain. Hence, $\hat{w}(\mu_0) \geq \int v dp$, delivering the result.

There are two hurdles to generalizing the above argument to more general action and state spaces. The first hurdle is that for some beliefs, there may be no action that satisfies both (i) and (ii) above. To circumvent this issue, we replace (ii) with the requirement that the action can be approximated by actions that are unique R-best responses at some alternative beliefs. The second hurdle is that, with infinite states, $\Delta\Theta$ has an empty interior when viewed as a subset of the set of all countably additive measures. Consequently, \hat{w} may be discontinuous

¹¹Note, the PUBR property is strictly stronger than the requirement that R has no duplicate actions. The latter condition is insufficient for ensuring uniqueness, as witnessed by $A = \Theta = \{0, 1\}$, $\mu_0 = \frac{1}{2}$, $u_S(a, \theta) = -a$, and $u_R(a, \theta) = a\theta$. The condition is reminiscent of the refined best reply correspondence that Balkenborg, Hofbauer, and Kuzmics (2013) and Balkenborg, Hofbauer, and Kuzmics (2015) study in finite games of complete information.

¹²Cf. Milne (1926).

at the prior. To overcome this challenge, we employ a construction that holds the prior fixed. We refer the reader to the appendix for the exact details.

Our next result uses Theorem 1 to show unique equilibrium S payoffs are a generic feature of finite environments.¹³

Proposition 2 (Generic uniqueness). *If A and Θ are finite, then an open dense $\mathcal{U}_R \subseteq \mathbb{R}^{A \times \Theta}$ of full Lebesgue measure exists such that S has a unique equilibrium payoff (for any prior and any S preferences) as long as $u_R \in \mathcal{U}_R$.*

The proposition’s proof shows that global PUBR is generic. Intuitively, a failure of PUBR at some belief implies that, for some fixed action and fixed set of states, R’s highest possible expected payoff gain from using said action rather than another is exactly zero—a knife-edge condition.

The above proposition, which makes no structural assumptions on players’ payoffs, is completely silent on infinite persuasion models. Meanwhile, papers applying the Bayesian persuasion framework (e.g., Kolotilin et al., 2017; Kolotilin, 2018; Dworzak and Martini, 2019; Guo and Shmaya, 2019) often focus on settings in which the state space, action space, or both are infinite.¹⁴ Because many such models are in some sense one-dimensional, we next turn our attention to specifications enjoying some ordinal structure.

Definition 4. *Say the environment is **ordered** if $A, \Theta \subseteq \mathbb{R}$; the function u_R exhibits strictly increasing differences; and every $\mu \in \Delta\Theta$ has either $\int u_R(a, \cdot) d\mu$ strictly quasiconcave in $a \in A$ or $\int u_S(a, \cdot) d\mu$ weakly quasiconvex in $a \in A$.*¹⁵

The above condition says that the action space and state space are both subsets of the real line, and that preferences respect this ordered structure. The increasing-differences condition says that R would like to take higher actions in higher states, and the quasiconvexity/quasiconcavity condition (which for instance holds if the the first-order approach is valid

¹³A related result follows directly from Proposition 3 of Li and Norman (2021) in the generic finite case. That result implies that all equilibria with S-favorable tie-breaking result in the same state-contingent action distribution. Hence, *assuming* R breaks indifferences in S’s favor generically yields a behavioral uniqueness property, not just payoff uniqueness.

¹⁴Moreover, in a setting in which R has private information and takes a binary action, such as Kolotilin et al. (2017), public persuasion can be reinterpreted as a continuous-action model in which R chooses a cutoff type at which to take the high action.

¹⁵The quasiconcavity condition says that any $a_L, a_M, a_H \in A$ with $a_L < a_M < a_H$ have $\int u_R(a_M, \cdot) d\mu > \min\{\int u_R(a_L, \cdot) d\mu, \int u_R(a_H, \cdot) d\mu\}$. This condition holds vacuously with binary actions, holds if u_R is strictly concave in its first argument, and holds if u_R is the restriction of a function $A \times \text{co}\Theta \rightarrow \mathbb{R}$ that is strictly quasiconcave in its first argument and affine in its second. Similarly, the quasiconvexity condition holds if u_S is weakly convex or weakly monotone in its first argument, or if u_S is the restriction of a function $A \times \text{co}\Theta \rightarrow \mathbb{R}$ that is weakly quasiconvex in its first argument and affine in its second.

in determining R’s optimal behavior, or if S always prefers higher actions) ensures that S always prefers either R’s highest or lowest best response. Note, the assumption that the environment is ordered is silent on whether R has unique best responses, on whether preferences are affine in R’s posterior belief, and on whether the state or action space is finite or infinite.

The next result shows that, except under a specific knife-edge condition (on R’s preferences and the prior), S has a unique equilibrium payoff. In particular, the result says S has a unique equilibrium payoff in ordered models with an atomless prior.

Theorem 2 (Uniqueness in the ordered model). *Suppose the environment is ordered; and for each $\bar{\theta} \in \{\min \text{supp } \mu_0, \max \text{supp } \mu_0\}$, either $\mu_0(\bar{\theta}) = 0$ or R’s best response at $\delta_{\bar{\theta}}$ is unique. Then, S has a unique equilibrium payoff.*

For a proof sketch, it is useful to consider the case in which A is finite. In this special case, one can derive payoff uniqueness as a consequence of Theorem 1. We need to argue that the PUBR property holds at any belief μ supported by an information policy. Given our generic assumption on the prior, we only need to worry about the case of beliefs at which R is neither certain of the highest state nor certain of the lowest state. The quasiconcavity/quasiconvexity condition of ordered models implies that S prefers either S’s highest or lowest best response at μ —say the highest. But then the degenerate belief on the highest state is a valid witness to the PUBR property: slightly increasing the probability on the high state breaks any R indifference in S’s preferred direction, because of the increasing differences hypothesis. To prove Theorem 2 in the general case, a qualitatively similar argument applies (though one cannot appeal to Theorem 1). This argument relies on the model’s continuity hypotheses. Instead of perturbing beliefs in the direction of a degenerate belief, we do so in the direction of a *conditional* belief of the prior conditional on the state exceeding a high cutoff; and instead of ensuring an action is selected as a unique best response, we ensure that all R best responses are very near to the targeted action.

3. Limited Commitment

One way of arguing in favor of sender-favorite equilibrium selection is to assert that a principal with full commitment power should also have the ability to steer others towards her favorite equilibrium. This argument, however, may appear less compelling in settings with a weaker sender, such as when the sender’s ability to commit is compromised. In this section, we show how to apply our results to such settings by extending the model to one in which S has only limited commitment power, as modeled in Lipnowski, Ravid, and Shishkin

(2022)—hereafter LRS1. To do so, we first specialize the existing parameters of our environment for tractability:

Assumption 1. *The spaces A , Θ , and M are all finite with $|M| \geq 2|\Theta|$, and u_S is state independent (i.e., constant in its second argument).*

In mild abuse of notation, write $u_S : A \rightarrow \mathbb{R}$. In addition to the parameters of our baseline model, we parameterize our limited-commitment model by $\chi \in [0, 1]$, which denotes the sender’s *credibility*.

The augmented game begins with S provisionally choosing an experiment, $\psi : \Theta \rightarrow \Delta M$, an “official” report. The state $\theta \in \Theta$ then realizes and, independent of the state, one of two possibilities occurs. With probability χ , the message $m \in M$ is sent in accordance with the report, and so is distributed according to $\psi(\cdot|\theta)$. With complementary probability $1 - \chi$, reporting is “influenced” and so S observes θ and can freely choose $m \in M$. Crucially, R observes the message m but observes neither θ nor whether reporting was influenced.

In what follows, we formalize an appropriate solution concept for the game with compromised reporting (again taking experiment-choice incentives into account), provide a result comparing the S equilibrium payoff set to that under perfect commitment, and apply our uniqueness results to address robustness to imperfect credibility.

3.1. The equilibrium payoff set

We formalize the solution concept for the augmented game as follows.

Definition 5. *A sender strategy is a pair (ψ, σ) consisting of an experiment $\psi \in \Psi$ and a measurable function $\sigma : \Theta \times \Psi \rightarrow \Delta M$; a receiver strategy is a measurable function $\alpha : M \times \Psi \rightarrow \Delta A$; and a receiver belief map is a measurable function $\pi : M \times \Psi \rightarrow \Delta \Theta$. A **perfect Bayesian χ -equilibrium** (χ -PBE) is a quadruple of such maps $\langle \psi, \sigma, \alpha, \pi \rangle$ such that*

1. *The sender’s experiment choice satisfies*

$$\psi \in \operatorname{argmax}_{\tilde{\psi} \in \Psi} \int_{\Theta} \int_M \int_A u_S \, d\alpha(\cdot|m, \tilde{\psi}) \left[\chi \, d\tilde{\psi}(m|\theta) + (1 - \chi) \, d\sigma(m|\theta, \tilde{\psi}) \right] d\mu_0(\theta);$$

2. *Every $\tilde{\psi} \in \Psi$ and $m \in M$ have*

$$\alpha \left(\operatorname{argmax}_{a \in A} \int_{\Theta} u_R(a, \theta) \, d\pi(\theta|m, \tilde{\psi}) \mid m, \tilde{\psi} \right) = 1;$$

3. Every $\tilde{\psi} \in \Psi$, Borel $\hat{M} \subseteq M$, and Borel $\hat{\Theta} \subseteq \Theta$ have

$$\int_{\Theta} \int_{\hat{M}} \pi(\hat{\Theta}|\cdot, \tilde{\psi}) d[\chi\tilde{\psi}(\cdot|\theta) + (1-\chi)\sigma(\cdot|\theta, \tilde{\psi})] d\mu_0(\theta) = \int_{\hat{\Theta}} [\chi\tilde{\psi}(\hat{M}|\cdot) + (1-\chi)\sigma(\hat{M}|\cdot, \tilde{\psi})] d\mu_0;$$

4. Every $\tilde{\psi} \in \Psi$ and $\theta \in \Theta$ have

$$\sigma\left(\operatorname{argmax}_{m \in M} \int_A u_S(a) d\alpha(a|\tilde{\psi}, m) \mid \theta, \tilde{\psi}\right) = 1.$$

In such a case, we say the induced χ -PBE payoff is

$$\int_{\Theta} \int_M \int_A u_S(a) d\alpha(a|m, \psi) [\chi d\psi(m|\cdot) + (1-\chi) d\sigma(m|\cdot, \psi)] d\mu_0.$$

Let $w_{\chi}^*(\mu_0)$ denote the infimum χ -PBE payoff for S.

We begin our analysis of this richer model with a partial description of the range of χ -PBE payoff for S. With full credibility, the range of S payoffs is naturally identical to the game without an influencing stage (as influence occurs with zero probability): one completes the equilibrium by incorporating some best response for an influencing S. But which payoffs can S attain when unable to perfectly commit? With S-favorable selection, and ignoring experiment-choice incentives, it is nearly immediate that higher credibility can only help S: she could always use her additional commitment power to replicate the way she would behave without it. But the effect of credibility is less clear when one considers S incentives and the full range of equilibria. First, the above replication argument is only available if S finds it optimal to mimic her hypothetical influencing behavior. Second, because influencing S is subject to incentive constraints more often under lower credibility, it is natural to wonder whether her incentive constraints would refine away some bad equilibria.

The following result shows that, in spite of these concerns, lower credibility is in fact worse for S than full commitment power in a set-valued sense.¹⁶

Proposition 3 (Payoff set under partial credibility). *The set of χ -PBE payoffs for S is nonempty and weakly below $[\hat{w}(\mu_0), \hat{v}(\mu_0)]$ in the strong set order, coinciding with it for $\chi = 1$.*

The part of the proposition for $\chi = 1$ is straightforward. We now briefly summarize the proof showing S's χ -PBE payoffs lies below $[\hat{w}(\mu_0), \hat{v}(\mu_0)]$ when $\chi < 1$. Our explanation

¹⁶Note, the proposition also establishes that a χ -PBE exists.

takes existence of a χ -PBE as given. As a starting point, note that additional incentive considerations can only constrain S, and so no χ -PBE payoff can be greater than $\hat{v}(\mu_0)$. Now, fix any $\chi < 1$, and let $v_\chi^*(\mu_0)$ be the highest χ -PBE payoff for S. If \hat{v}_χ were below $\hat{w}(\mu_0)$, the desired set ranking immediately follows. Suppose then that $v_\chi^*(\mu_0) > \hat{w}(\mu_0)$. In this case, we argue that every payoff between $\hat{w}(\mu_0)$ and $v_\chi^*(\mu_0)$ is compatible with χ -PBE. For a rough sketch, consider first S's set of attainable payoffs if we ignored her experiment-choice incentives. Since one can always pair an uninformative official report with a babbling equilibrium, this set always includes $w(\mu_0)$. Moreover, this set is an interval, by Lemma 7 in LRS1. Since $w(\mu_0) \leq \hat{w}(\mu_0)$, it follows the set of payoffs S can obtain ignoring experiment-choice incentives is an interval whose lower bound is weakly below $\hat{w}(\mu_0)$.

Given the above, to prove the proposition, it is sufficient to claim that for any experiment S could initially choose, some continuation equilibrium gives S a payoff of no more than $\hat{w}(\mu_0)$. This claim would follow immediately from the definition of \hat{w} if R were to choose adversarially to S whenever he is indifferent between multiple actions. However, it is unclear whether a continuation equilibrium with the latter feature exists, since w is not continuous (hence, not upper semicontinuous)—hence S-adversarial tie-breaking by R could mean S does not have any optimal message to send when allowed to influence the state. To resolve this issue, we search for a continuation equilibrium in which R gives S a continuation payoff in $[w(\mu), z(\mu)]$ whenever his posterior belief is μ , where z is the smallest upper semicontinuous function above w (and so $z \leq v$). Because the correspondence $[w, z]$ is upper hemicontinuous, an appropriate application of Kakutani's fixed point theorem delivers such continuation play and beliefs. This continuation play generates an S payoff of at most $\hat{z}(\mu_0)$ by KG's results, where \hat{z} is the concave envelope of z . Finally, because the concave function \hat{w} on a finite-dimensional simplex is automatically continuous on the interior, we can show that the concave envelopes \hat{z} and \hat{w} agree at the prior.

3.2. Strong robustness

LRS1 establishes a robustness result (Proposition 3 of that paper) concerning the best S payoff attainable when credibility is only slightly imperfect ($\chi \approx 1$). That result implies her highest attainable payoff converges to the Bayesian persuasion value for most payoff specifications—in particular, whenever the PUBR property holds globally. In this subsection, we demonstrate a tight link between two conceptually distinct notions of robustness. We show that S payoffs are robust to slightly imperfect credibility *and* equilibrium selection if and only if they are robust to equilibrium selection in the full-credibility case. With this

result in hand, the results of the previous sections apply directly to address this stronger form of robustness in persuasion models.

Proposition 4 (Strong robustness). *The lowest χ -PBE values satisfy $\lim_{\chi \nearrow 1} w_\chi^*(\mu_0) = \hat{w}(\mu_0)$. In particular, the Bayesian persuasion value is strongly robust to partial credibility ($\lim_{\chi \nearrow 1} w_\chi^*(\mu_0) = \hat{w}(\mu_0)$) if and only if it is robust to equilibrium selection ($\hat{w}(\mu_0) = \hat{v}(\mu_0)$).*

The proof constructs lower payoff bounds for S in any χ -PBE by computing her payoff following viable deviations, and shows that these payoff bounds can be made arbitrarily close to $\hat{w}(\mu_0)$ as credibility becomes arbitrarily close to perfect. In brief, we consider the deviation to an experiment that would be approximately optimal under full S commitment when R breaks ties against S's interests. Using lower semicontinuity of w , we show S attains a payoff approximating her full-commitment payoff from this experiment as χ approaches 1, because the belief a message induces is very nearby the full-commitment version of the same.

Finally, Propositions 2 and 4 immediately yield the following conclusion.

Corollary 3 (Generic strong robustness). *For any finite A and Θ , for all but a Lebesgue-null (and nowhere dense) set of R objectives $u_R \in \mathbb{R}^{A \times \Theta}$, and every S objective $u_S \in \mathbb{R}^A$, the full commitment value is strongly robust to partial credibility.*

Thus, for most payoff specifications, S obtains essentially her Bayesian persuasion payoff in any equilibrium, even if her ability to commit to the information she provides is slightly imperfect.

A. Proofs

Before proceeding to formal proofs, we review for convenience several key notations.

$$\begin{aligned}
A_R^*: \Delta\Theta &\rightrightarrows A \\
\mu &\mapsto \operatorname{argmax}_{a \in A} \int u_R(a, \cdot) d\mu \\
V: \Delta\Theta &\rightrightarrows \mathbb{R} \\
\mu &\mapsto \operatorname{co} \left\{ \int u_S(a, \cdot) d\mu : a \in A_R^*(\mu) \right\} \\
v: \Delta\Theta &\rightarrow \mathbb{R} \\
\mu &\mapsto \max V(\mu) \\
w: \Delta\Theta &\rightarrow \mathbb{R} \\
\mu &\mapsto \min V(\mu) \\
\mathcal{I}: \Delta\Theta &\rightrightarrows \Delta\Delta\Theta \\
\mu &\mapsto \left\{ p \in \Delta\Delta\Theta : \int \tilde{\mu} dp(\tilde{\mu}) = \mu \right\} \\
\hat{v}: \Delta\Theta &\rightarrow \mathbb{R} \\
\mu &\mapsto \max_{p \in \mathcal{I}(\mu)} \int v dp \\
\hat{w}: \Delta\Theta &\rightarrow \mathbb{R} \\
\mu &\mapsto \sup_{p \in \mathcal{I}(\mu)} \int w dp.
\end{aligned}$$

Note (given that the barycenter map is continuous and by Berge's theorem) that all three correspondences are nonempty-compact-valued and upper hemicontinuous, the functions v and \hat{v} are upper semicontinuous, and the function w is lower semicontinuous.

A.1. Proofs for Section 1

The following lemma shows information policies can be replaced with payoff-equivalent ones of small support when the state space is finite.

Lemma 1. *If Θ is finite, then any $\bar{\mu} \in \Delta\Theta$, Borel $D \subseteq \Delta\Theta$, bounded measurable $f: \Delta\Theta \rightarrow \mathbb{R}$ and $p \in \mathcal{I}(\bar{\mu}) \cap \Delta D$ admit some $q \in \mathcal{I}(\bar{\mu}) \cap \Delta D$ with $\operatorname{supp}(q)$ affinely independent (hence of cardinality no greater than $|\Theta|$) and $\int f dq \geq \int f dp$.*

Proof. The set $\mathcal{I}(\bar{\mu})$ is convex compact metrizable, and so Choquet's theorem yields some $Q \in \Delta[\mathcal{I}(\bar{\mu})]$ with barycenter p such that Q is supported on the extreme points of $\mathcal{I}(\bar{\mu})$. But this set $\operatorname{ext}[\mathcal{I}(\bar{\mu})]$ consists exactly of those $q \in \mathcal{I}(\bar{\mu})$ with affinely independent support.

By definition of the barycenter, $\iint g dq dQ(q) = \int g dp$ for every continuous $g: \Delta\Theta \rightarrow \mathbb{R}$. However, because the barycenter of this Q is unique, it follows that $p(B) = \int q(B) dQ(q)$ for

every Borel $B \subseteq \Delta\Theta$. Hence, $Q(\Delta D) = 1$ because $p(D) = 1$, and $\iint g \, dq \, dQ(q) = \int g \, dp$ for every bounded measurable $g: \Delta\Theta \rightarrow \mathbb{R}$ —in particular for $g = f$.

Letting $\mathcal{J} := \text{ext}[\mathcal{I}(\bar{\mu})] \cap \Delta D$, we have $Q(\mathcal{J}) = 1$. Therefore,

$$0 = \iint f \, dq \, dQ(q) - \int f \, dp = \int_{\mathcal{J}} \left[\int f \, dq - \int f \, dp \right] dQ(q).$$

So the integrand is somewhere nonnegative: some $q \in \mathcal{J}$ has $\int f \, dq \geq \int f \, dp$. \square

Now, we prove the characterization of all S equilibrium payoffs.

Proof of Proposition 0. To begin, we recall some well-known facts about experiments and Bayesian updating—which collectively tell us that S choosing from Ψ and choosing from $\mathcal{I}(\mu_0)$ are equivalent formalisms. First, any experiment $\tilde{\psi} \in \Psi$ admits some compatible belief map, that is, some measurable $\tilde{\pi} = \tilde{\pi}_{\tilde{\psi}}: M \rightarrow \Delta\Theta$ such that every Borel $\hat{M} \subseteq M$ and Borel $\hat{\Theta} \subseteq \Theta$ have $\int_{\Theta} \int_{\hat{M}} \tilde{\pi}(\hat{\Theta}|m, \tilde{\psi}) \, d\tilde{\psi}(m|\theta) \, d\mu_0(\theta) = \int_{\hat{\Theta}} \tilde{\psi}(\hat{M}|\theta) \, d\mu_0(\theta)$. Second, given $\tilde{\psi} \in \Psi$ if we define the belief distribution $p_{\tilde{\psi}, \tilde{\pi}} \in \Delta\Delta\Theta$ via $p_{\tilde{\psi}, \tilde{\pi}}(D) := \int_{\Theta} \tilde{\psi}(\tilde{\pi}^{-1}(D) | \theta) \, d\mu_0(\theta)$ for each Borel $D \subseteq \Delta\Theta$, then $p_{\tilde{\psi}, \tilde{\pi}} = p_{\tilde{\psi}, \tilde{\pi}'}$ for any two such compatible $\tilde{\pi}$ and $\tilde{\pi}'$. We therefore refer to the associated belief distribution simply as $p_{\tilde{\psi}}$. Third, every $\tilde{\psi} \in \Psi$ has $p_{\tilde{\psi}} \in \mathcal{I}(\mu_0)$. Fourth, every $p \in \mathcal{I}(\mu_0)$ with $|\text{supp}(p)| \leq |M|$ admits some $\tilde{\psi}_p \in \Psi$ such that $p_{\tilde{\psi}_p} = p$.

Now we proceed to show $s \in [\hat{w}(\mu_0), \hat{v}(\mu_0)]$ is necessary and sufficient for s to be an equilibrium S payoff.

First, to see the condition is necessary, fix an arbitrary equilibrium $\langle \psi, \alpha, \pi \rangle$, and let $s \in \mathbb{R}$ be the induced S payoff; we will show $s \in [\hat{w}(\mu_0), \hat{v}(\mu_0)]$. For any $\tilde{\psi} \in \Psi$ and $m \in M$ the R optimality condition implies $\alpha(A_R^*(\pi(\cdot|m, \tilde{\psi})) | m, \tilde{\psi}) = 1$, so that $\int_A u_S(a, \theta) \, d\alpha(a|m, \tilde{\psi}) \in V(\pi(\cdot|m, \tilde{\psi}))$. Therefore, any $\tilde{\psi} \in \Psi$ has

$$\begin{aligned} \int_{\Theta} \int_M \int_A u_S(a, \theta) \, d\alpha(a|m, \tilde{\psi}) \, d\tilde{\psi}(m|\theta) \, d\mu_0(\theta) &\in \int_{\Theta} \int_M V(\pi(\cdot|m, \tilde{\psi})) \, d\tilde{\psi}(m|\theta) \, d\mu_0(\theta) \\ &= \int_{\Delta\Theta} V \, dp_{\tilde{\psi}} \\ &= \left[\int_{\Delta\Theta} w \, dp_{\tilde{\psi}}, \int_{\Delta\Theta} v \, dp_{\tilde{\psi}} \right] \\ &\subseteq \left[\int_{\Delta\Theta} w \, dp_{\tilde{\psi}}, \hat{v}(\mu_0) \right]. \end{aligned}$$

Hence, $s \leq \hat{v}(\mu_0)$. Moreover, for any $p \in \mathcal{I}(\mu_0)$ if Θ is infinite, and for any $p \in \mathcal{I}(\mu_0)$ such that $|\text{supp}(p)| \leq |\Theta|$ if Θ is finite, taking $\tilde{\psi} = \tilde{\psi}_p$ implies (by S rationality)

$$\begin{aligned} &\int_{\Theta} \int_M \int_A u_S(a, \theta) \, d\alpha(a|m, \psi) \, d\psi(m|\theta) \, d\mu_0(\theta) \\ &\geq \int_{\Theta} \int_M \int_A u_S(a, \theta) \, d\alpha(a|m, \tilde{\psi}) \, d\tilde{\psi}(m|\theta) \, d\mu_0(\theta) \\ &\geq \int_{\Delta\Theta} w \, dp. \end{aligned}$$

Applying this observation to every such p (and applying Lemma 1 if Θ is finite) implies $s \geq \hat{w}(\mu_0)$.

Conversely, take any $s \in [\hat{w}(\mu_0), \hat{v}(\mu_0)]$. Letting $p_1 \in \mathcal{I}(\mu_0)$ with $\int_{\Delta\Theta} v \, dp_1 = \hat{v}(\mu_0)$ and $|\text{supp}(p_1)| \leq |\Theta|$ if Θ is finite—which exists by Lemma 1—define $p_\lambda := \int \delta_{\lambda\mu + (1-\lambda)\mu_0} \, dp_1(\mu) \in \mathcal{I}(\mu_0)$ for each $\lambda \in [0, 1]$; observe $|\text{supp}(p_\lambda)| \leq |\text{supp}(p_1)| \leq |M|$. As $\lambda \mapsto p_\lambda$ is continuous, it follows that the $\lambda \mapsto \int V \, dp_\lambda$ is nonempty-compact-convex-valued and upper hemicontinuous because V is. Moreover,

$$\int V \, dp_0 = V(\mu_0) \ni w(\mu_0) \leq s \leq \hat{v}(\mu_0) \in \int V \, dp_1.$$

The intermediate value theorem for correspondences (e.g., Lemma 2 from de Clippel, 2008) therefore delivers some $\lambda \in [0, 1]$ such that $s \in \int V \, dp_\lambda$. Some measurable $\zeta: \Delta\Theta \rightarrow [0, 1]$ then exists such that $s = \int [(1 - \zeta)w + \zeta v] \, dp_\lambda$. By the measurable maximum theorem (Theorem 18.19 from Aliprantis and Border, 2006), a pair of measurable functions $\alpha_w, \alpha_v: \Delta\Theta \rightarrow \Delta A$ exist such that, each $\mu \in \Delta\Theta$ has $\int_{A \times \Theta} u_S \, d[\alpha_w(\cdot|\mu) \otimes \mu] = w(\mu)$, $\int_{A \times \Theta} u_S \, d[\alpha_v(\cdot|\mu) \otimes \mu] = v(\mu)$, and $\alpha_w(A_R^*(\mu)|\mu) = \alpha_v(A_R^*(\mu)|\mu) = 1$. With these objects in hand, we can define our candidate $\psi: \Theta \rightarrow \Delta M$, $\alpha: M \times \Psi \rightarrow \Delta A$, and $\pi: M \times \Psi \rightarrow \Delta\Theta$ via

$$\begin{aligned} \psi(\cdot|\theta) &:= \tilde{\psi}_{p_\lambda} \\ \pi(\cdot|m, \tilde{\psi}) &:= \tilde{\pi}_{\tilde{\psi}}(\cdot|m) \\ \alpha(\cdot|m, \tilde{\psi}) &:= \begin{cases} (1 - \zeta)\alpha_w(\cdot|\pi(\cdot|m, \tilde{\psi})) + \zeta\alpha_v(\cdot|\pi(\cdot|m, \tilde{\psi})) & : \tilde{\psi} = \tilde{\psi}_{p_\lambda} \\ \alpha_w(\cdot|\pi(\cdot|m, \tilde{\psi})) & : \tilde{\psi} \neq \tilde{\psi}_{p_\lambda} \end{cases} \end{aligned}$$

It is immediate from the construction that all three maps are measurable and that R rationality and the Bayesian property are both satisfied. Moreover, direct computation shows that choosing experiment $\tilde{\psi} \in \Psi$ gives S a continuation payoff of

$$\int_{\Theta} \int_M \int_A u_S(a, \theta) \, d\alpha(a|m, \tilde{\psi}) \, d\tilde{\psi}(m|\theta) \, d\mu_0(\theta) = \begin{cases} s & : \tilde{\psi} = \tilde{\psi}_{p_\lambda} \\ \int_{\Delta\Theta} w \, dp_{\tilde{\psi}} & : \tilde{\psi} \neq \tilde{\psi}_{p_\lambda} \end{cases}$$

Therefore, S gets payoff s if the triple is an equilibrium. Finally, the triple is indeed an equilibrium: In particular, S rationality is confirmed because $s \geq \hat{w}(\mu_0) \geq \int w \, dp_{\tilde{\psi}}$ for every alternative $\tilde{\psi} \in \Psi$. \square

A.2. Proofs for Section 2

Lemma 2. *The function $\hat{w}: \Delta\Theta \rightarrow \mathbb{R}$ is lower semicontinuous (hence measurable) and concave, and every $\bar{p} \in \mathcal{I}(\mu_0)$ has $\hat{w}(\mu_0) \geq \int \hat{w} \, d\bar{p}$.*

Proof. First, the function $\Delta\Delta\Theta \rightarrow \mathbb{R}$ given by $p \mapsto \int w \, dp$ is lower semicontinuous because w is. Moreover, the barycenter map $\Delta\Delta\Theta \rightarrow \Delta\Theta$ is open (O'Brien, 1976, Corollary 1), meaning \mathcal{I} is lower hemicontinuous. It follows from (Aliprantis and Border, 2006, Lemma

17.29) that \hat{w} is lower semicontinuous. In the next paragraph we establish the last assertion in the lemma's statement. Specializing that assertion to the case where \bar{p} has finite support (and varying μ_0) delivers that \hat{w} is concave. The lemma follows.

Fix an arbitrary information policy $\bar{p} \in \mathcal{I}(\mu_0)$ and $\varepsilon > 0$; we want to show $\hat{w}(\mu_0) \geq \int \hat{w} d\bar{p} - \varepsilon$. To that end, define the correspondence

$$\begin{aligned} \Phi: \Delta\Theta &\rightrightarrows \Delta\Delta\Theta \\ \mu &\mapsto \left\{ p \in \mathcal{I}(\mu): \int w dp \geq \hat{w}(\mu) - \varepsilon \right\}. \end{aligned}$$

This correspondence is nonempty-valued by definition of \hat{w} , and has measurable graph because \hat{w} is measurable and the barycenter map $p \mapsto \int \mu dp(\mu)$ is continuous (hence measurable). Therefore, Corollary 18.27 from Aliprantis and Border (2006) yields a measurable function $\varphi: \Delta\Theta \rightarrow \Delta\Delta\Theta$ with $\bar{p}\{\varphi \in \Phi\} = 1$. Because $\bar{p} \in \mathcal{I}(\mu_0)$ and $\bar{p}\{\varphi \in \mathcal{I}\} = 1$, we know $p := \int \varphi d\bar{p}$ is in $\mathcal{I}(\mu_0)$. Therefore,

$$\hat{w}(\mu_0) \geq \int w dp = \iint w d\varphi(\cdot|\mu) d\bar{p}(\mu) \geq \int (\hat{w} - \varepsilon) d\bar{p} = \int \hat{w} d\bar{p} - \varepsilon$$

The lemma follows. \square

Proof of Proposition 1. Let $\mathcal{I}_D(\mu_0)$ be the set of all information policies $p \in \mathcal{I}(\mu_0)$ such that $p(\bar{D}) = 1$ for some measurable $\bar{D} \subseteq D$. Because D is persuasion sufficient and $\hat{w}|_D \geq v|_D$, and by Lemma 2, we have:

$$\hat{v}(\mu_0) = \sup_{p \in \mathcal{I}_D(\mu_0)} \int v dp \leq \sup_{p \in \mathcal{I}_D(\mu_0)} \int \hat{w} dp \leq \hat{w}(\mu_0).$$

The proposition follows. \square

Proof of Corollary 2. By Proposition 0, it suffices to show that $\hat{w} = \hat{v}$ if and only if $\hat{w} \geq v$. As it is immediate that $\hat{v} \geq \hat{w}$ and $\hat{v} \geq v$, we need only see that $\hat{w} \geq \hat{v}$ if $\hat{w} \geq v$. But this result follows directly from Proposition 1, because $\Delta\Theta$ is vacuously persuasion sufficient. \square

Lemma 3. *Suppose $D \subseteq \Delta\Theta$ is measurable, $p \in \mathcal{I}(\mu_0)$ has $p(D) = 1$, and $f: \Delta\Theta \rightarrow \mathbb{R}$ is bounded and measurable. Suppose every $\bar{\mu} \in D$ admits some $\mu \in \Delta\Theta$ and $\gamma > 0$ such that $\gamma\mu \leq \bar{\mu}$ and all sufficiently small $\lambda \in (0, 1]$ have $\hat{w}(\lambda\mu + (1-\lambda)\bar{\mu}) \geq f(\bar{\mu})$. Then $\hat{w}(\mu_0) \geq \int f dp$.*

Proof. If we establish $\hat{w}|_D \geq f|_D$, the lemma will follow from Lemma 2, because we will have $\hat{w}(\mu_0) \geq \int \hat{w} dp \geq \int f dp$. So take any $\bar{\mu} \in D$, with a view to showing $\hat{w}(\bar{\mu}) \geq f(\bar{\mu})$.

Let $\mu \in \Delta\Theta$ and $\gamma > 0$ be as guaranteed by the lemma's statement. Assuming without loss that $\gamma \in (0, 1)$, let $\mu' := \frac{\bar{\mu} - \gamma\mu}{1-\gamma} \in \Delta\Theta$ so that $\bar{\mu} = \gamma\mu + (1-\gamma)\mu'$. Then any sufficiently small $\lambda \in (0, 1]$ has

$$\begin{aligned} \bar{\mu} &= \frac{\gamma}{\gamma+\lambda(1-\gamma)} [\lambda\mu + (1-\lambda)\bar{\mu}] + \frac{\lambda(1-\gamma)}{\gamma+\lambda(1-\gamma)} \mu' \\ \implies \hat{w}(\bar{\mu}) &\geq \frac{\gamma}{\gamma+\lambda(1-\gamma)} \hat{w}(\lambda\mu + (1-\lambda)\bar{\mu}) + \frac{\lambda(1-\gamma)}{\gamma+\lambda(1-\gamma)} \hat{w}(\mu') \text{ (by Lemma 2)} \\ &\geq \frac{\gamma}{\gamma+\lambda(1-\gamma)} f(\bar{\mu}) + \frac{\lambda(1-\gamma)}{\gamma+\lambda(1-\gamma)} \min u_S(A \times \Theta) \\ &\rightarrow f(\bar{\mu}) \text{ as } \lambda \rightarrow 0. \end{aligned}$$

Thus, $\hat{w}(\bar{\mu}) \geq f(\bar{\mu})$, delivering the lemma. \square

Lemma 4. *Suppose $D \subseteq \Delta\Theta$ is measurable, $p \in \mathcal{I}(\mu_0)$ has $p(D) = 1$, and $f: \Delta\Theta \rightarrow \mathbb{R}$ is bounded and measurable. Suppose every $\bar{\mu} \in D$ admits some $\mu \in \Delta\Theta$ and $\gamma \in (0, 1)$ such that $\gamma\mu \leq \mu_0$ and all sufficiently small $\lambda \in (0, 1]$ have $\hat{w}(\lambda\mu + (1 - \lambda)\bar{\mu}) > f(\bar{\mu})$. Then, $\hat{w}(\mu_0) \geq \int f \, dp$.*

Proof. Define the correspondences

$$\begin{aligned} F: D &\rightrightarrows \Delta\Theta \\ \bar{\mu} &\mapsto \{\mu \in \Delta\Theta: \hat{w}(\mu) > f(\bar{\mu})\} \\ G: D &\rightrightarrows \Delta\Theta \\ \bar{\mu} &\mapsto \{\mu \in \Delta\Theta: \lambda\mu + (1 - \lambda)\bar{\mu} \in F(\bar{\mu}) \, \forall \text{ sufficiently small } \lambda \in (0, 1]\} \\ H: D &\rightrightarrows \Delta\Theta \times (0, 1) \\ \bar{\mu} &\mapsto \{(\mu, \zeta) \in G(\bar{\mu}) \times (0, 1): \zeta\mu \leq \mu_0\}. \end{aligned}$$

By hypothesis, H is nonempty-valued. At the end of this proof, we shall show that H has measurable graph. Before doing so, let us establish the lemma taking this fact for granted.

Extend H to $H: \Delta\Theta \rightrightarrows \Delta\Theta \times (0, 1)$ by letting $H(\bar{\mu}) := \Delta\Theta \times (0, 1)$ for every $\bar{\mu} \in (\Delta\Theta) \setminus D$. Because H has a measurable graph, the graph of the extended correspondence is also measurable, and so Corollary 18.27 from Aliprantis and Border (2006) delivers measurable functions $\varphi: \Delta\Theta \rightarrow \Delta\Theta$ and $\gamma: \Delta\Theta \rightarrow (0, 1)$ with $p\{(\varphi, \gamma) \in H\} = 1$. Let $\xi := \int (1 - \gamma) \, dp \in (0, 1)$.

For any $\lambda \in (0, 1]$, define the measurable function

$$\begin{aligned} \varphi_\lambda: D &\rightarrow \Delta\Theta \\ \mu &\mapsto \frac{1}{(1 - \lambda) + \lambda\gamma(\mu)} [(1 - \lambda)\mu + \lambda\gamma(\mu)\varphi(\mu)], \end{aligned}$$

and define the measures¹⁷

$$\mu_\lambda := \frac{1}{\lambda\xi} \left\{ \mu_0 - \int (1 - \lambda + \lambda\gamma) \varphi_\lambda \, dp \right\} \in \Delta\Theta$$

and¹⁸

$$p_\lambda := \int (1 - \lambda + \lambda\gamma) \delta_{\varphi_\lambda(\cdot)} \, dp + \lambda\xi \delta_{\mu_\lambda} \in \Delta\Delta\Theta.$$

¹⁷It is immediate from the definition of ξ that $\mu_\lambda \in \text{ca}(\Theta)$ has $\mu_\lambda(\Theta) = 1$. To see the measure is positive, hence in $\Delta\Theta$, observe that $p \in \mathcal{I}(\mu_0)$ and $p\{(\varphi, \gamma) \in H\} = 1$ implies

$$\lambda\xi\mu_\lambda = \mu_0 - \int [(1 - \lambda)\mu + \lambda\gamma(\mu)\varphi(\mu)] \, dp(\mu) = \mu_0 - \int [(1 - \lambda)\mu_0 + \lambda\gamma\varphi] \, dp = \lambda \int (\mu_0 - \gamma\varphi) \, dp \geq 0.$$

¹⁸Recall δ_μ is the degenerate probability measure on $\{\mu\}$ for any μ . Therefore, for any measurable $\tilde{D} \subseteq \Delta\Theta$, we have $p_\lambda(\tilde{D}) = \int [1 - \lambda + \lambda\gamma(\mu)] \mathbf{1}_{\varphi_\lambda(\mu) \in \tilde{D}} \, dp(\mu) + \lambda\xi \mathbf{1}_{\mu_\lambda \in \tilde{D}}$.

The definition of μ_λ ensures $p_\lambda \in \mathcal{I}(\mu_0)$ for $\lambda \in (0, 1]$. Therefore,¹⁹

$$\begin{aligned}
\hat{w}(\mu_0) &\geq \liminf_{\lambda \rightarrow 0} \int \hat{w} \, dp_\lambda && \text{(by Lemma 2)} \\
&\geq \liminf_{\lambda \rightarrow 0} \left[\int (1 - \lambda + \lambda\gamma) \hat{w} \circ \varphi_\lambda \, dp + \lambda\xi \min u_S(A \times \Theta) \right] \\
&= \liminf_{\lambda \rightarrow 0} \int (1 - \lambda + \lambda\gamma) \hat{w} \circ \varphi_\lambda \, dp \\
&\geq \int \liminf_{\lambda \rightarrow 0} [(1 - \lambda + \lambda\gamma) \hat{w} \circ \varphi_\lambda] \, dp && \text{(by Fatou's lemma)} \\
&\geq \int f \, dp && \text{(because } p\{\varphi \in G\} = 1\text{).}
\end{aligned}$$

The lemma then follows if we show H has measurable graph.

Toward measurability of the graph $\text{gr } H$, let $X := D \times \Delta\Theta$. First note that $\text{gr } F = \{(\bar{\mu}, \mu) \in X : f(\bar{\mu}) - \hat{w}(\mu) < 0\}$ is measurable because f is measurable and \hat{w} is (by Lemma 2) lower semicontinuous. Next, consider the graph of G . For any $\bar{\mu} \in D$, that F is (by Lemma 2) convex-valued and $(0, 1] = \text{co}\{\frac{1}{n} : n \in \mathbb{N}\}$ implies

$$G(\bar{\mu}) = \left\{ \mu \in \Delta\Theta : \frac{1}{n}\mu + (1 - \frac{1}{n})\bar{\mu} \in F(\bar{\mu}) \, \forall \text{ sufficiently large } n \in \mathbb{N} \right\}.$$

Moreover, for each $n \in \mathbb{N}$, the map $\psi_n : X \rightarrow X$ given by $\psi_n(\bar{\mu}, \mu) := (\bar{\mu}, \frac{1}{n}\mu + (1 - \frac{1}{n})\bar{\mu})$ is continuous, hence measurable. Therefore,

$$\begin{aligned}
\text{gr } G &= \left\{ (\bar{\mu}, \mu) \in X : \left(\bar{\mu}, \frac{1}{n}\mu + (1 - \frac{1}{n})\bar{\mu} \right) \in \text{gr } F \, \forall \text{ sufficiently large } n \in \mathbb{N} \right\} \\
&= \bigcup_{N=1}^{\infty} \bigcap_{n=N}^{\infty} \psi_n^{-1}(\text{gr } F),
\end{aligned}$$

which is measurable. Finally, the graph

$$\text{gr } H = [(\text{gr } G) \times (0, 1)] \cap \{(\bar{\mu}, \mu, \zeta) \in X \times (0, 1) : \mu_0 - \zeta\mu \geq 0\}$$

is measurable too because $(\bar{\mu}, \mu, \zeta) \mapsto \mu_0 - \zeta\mu$ is continuous and $\text{ca}_+(\Theta)$ is closed in $\text{ca}(\Theta)$.²⁰

The lemma follows. \square

Lemma 5. *Let $\bar{\mu}, \mu \in \Delta\Theta$ and $\varepsilon > 0$. If $A_R^*(\mu) = \{a\}$ for some $a \in A_R^*(\bar{\mu})$ with $\int u_S(a, \cdot) \, d\bar{\mu} > v(\bar{\mu}) - \varepsilon$, then all sufficiently small $\lambda \in (0, 1]$ have $w(\lambda\mu + (1 - \lambda)\bar{\mu}) > v(\bar{\mu}) - \varepsilon$.*

Proof. Because a is a best response to $\bar{\mu}$ and a unique best response to μ , it is (by linearity of expected utility) a unique best response to any proper convex combination μ' , yielding $w(\mu') = \int u_S(a, \cdot) \, d\mu'$, which converges to $v(\mu)$ as $\mu' \rightarrow \bar{\mu}$. \square

¹⁹Note Fatou's lemma applies because $(1 - \lambda + \lambda\gamma) \hat{w} \circ \varphi_\lambda$ is bounded and p is finite.

²⁰Suppose $(\eta_\beta)_\beta$ is a net in $\text{ca}_+(\Theta)$ that weak* converges to some $\eta \in \text{ca}(\Theta)$. For each β , some $\lambda_\beta \in \mathbb{R}_+$ and $\mu_\beta \in \Delta\Theta$ have $\eta_\beta = \lambda_\beta \mu_\beta$. We want to show $\eta \geq 0$. Passing to a subnet, we may assume (because $\Delta\Theta$ is compact) μ_β converges to some $\mu \in \Delta\Theta$; and because the constant function $\mathbf{1}$ is continuous, $\eta_\beta \rightarrow \eta$ implies $\lambda_\beta \rightarrow \lambda := \int \mathbf{1} \, d\eta$. Therefore, $\lambda \geq 0$ and so $\eta = \lambda\mu \geq 0$.

Lemma 6. *S has a unique equilibrium payoff if some persuasion-sufficient $D \subseteq \Delta\Theta$ exists such that every $\bar{\mu} \in D$ has*

$$\max_{a \in A_R^*(\bar{\mu})} \int u_S(a, \cdot) d\bar{\mu} = \sup_{a \in A_R^*(\bar{\mu}) \cap A^U(\mu_0)} \int u_S(a, \cdot) d\bar{\mu}.$$

Proof. Fixing $\varepsilon > 0$ and $\bar{\mu} \in D$, Lemma 4 means it suffices to find $\mu \in \Delta\Theta$ and $\gamma > 0$ such that $\gamma\mu \leq \mu_0$ and all sufficiently small $\lambda \in (0, 1]$ have $\hat{w}(\lambda\mu + (1-\lambda)\bar{\mu}) > v(\bar{\mu}) - \varepsilon$. Some $a \in A_R^*(\bar{\mu})$ has $\int u_S(a, \cdot) d\bar{\mu} > v(\bar{\mu}) - \varepsilon$. Then, by hypothesis, some $\mu \in \Delta\Theta$ and $\gamma > 0$ exist such that $\gamma\mu \leq \mu_0$ and $A_R^*(\mu) = \{a\}$. This μ is as desired by Lemma 5. \square

Proof of Theorem 1. First, we establish that the strong PUBR property implies S has a unique equilibrium payoff. Given $\varepsilon > 0$ and $\bar{\mu} \in D$, Lemma 3 means it suffices to find $\mu \in \Delta\Theta$ and $\gamma > 0$ such that $\gamma\mu \leq \bar{\mu}$ and all sufficiently small $\lambda \in (0, 1]$ have $\hat{w}(\lambda\mu + (1-\lambda)\bar{\mu}) > v(\bar{\mu}) - \varepsilon$. To do so, note some $a \in A_R^*(\bar{\mu})$ has $\int u_S(a, \cdot) d\bar{\mu} > v(\bar{\mu}) - \varepsilon$. Then, by hypothesis, some $\mu \in \Delta\Theta$ and $\gamma > 0$ exist such that $\gamma\mu \leq \bar{\mu}$ and $A_R^*(\mu) = \{a\}$. This μ is as desired by Lemma 5.

Now, we show the PUBR property implies S has a unique equilibrium payoff if $\min\{|A|, |\Theta|\} < \infty$. For both cases, we appeal to Lemma 6. In light of that lemma, it suffices to note that $A^U \subseteq A^U(\mu_0)$ if either Θ or A is finite. If Θ is finite, the result follows because every belief $\mu \in \Delta[\text{supp } \mu_0]$ has $\gamma\mu \leq \mu_0$ for $\gamma := \min_{\theta \in \text{supp } \mu} \frac{\mu_0(\theta)}{\mu(\theta)}$.²¹ If A is finite, the result follows because the set of beliefs in $\Delta[\text{supp } \mu_0]$ to which a given action is a unique best response is relatively open, and (Lemma 2, Lipnowski and Mathevet, 2018) the set of beliefs μ admitting $\gamma > 0$ with $\gamma\mu \leq \mu_0$ is dense in $\Delta[\text{supp } \mu_0]$. \square

Lemma 7. *Given finite A and Θ , define the set $\mathcal{U}_R \subseteq \mathbb{R}^{A \times \Theta}$ of R objectives u_R such that every $\mu \in \Delta\Theta$ and $a \in A_R^*(\mu)$ have some $\mu' \in \Delta[\text{supp } \mu]$ with $A_R^*(\mu') = \{a\}$. Then \mathcal{U}_R is open and dense with full Lebesgue measure.*

Proof. Toward showing these properties, let us note an algebraic characterization of \mathcal{U}_R . Define the finite index set $\mathbb{I} := \{(a, \hat{\Theta}) : a \in A, \emptyset \neq \hat{\Theta} \subseteq \Theta\}$ and, for each $i = (a, \hat{\Theta}) \in \mathbb{I}$, define

$$\begin{aligned} \varphi_i : \mathbb{R}^{A \times \Theta} &\rightarrow \mathbb{R} \\ u_R &\mapsto \max_{\mu \in \Delta\hat{\Theta}} \min_{a' \in A \setminus \{a\}} \int [u_R(a, \cdot) - u_R(a', \cdot)] d\mu. \end{aligned}$$

Then $\mathcal{U}_R = \{u_R \in \mathbb{R}^{A \times \Theta} : \varphi_i(u_R) \text{ is nonzero for every } i \in \mathbb{I}\} = \bigcap_{i \in \mathbb{I}} \varphi_i^{-1}(\mathbb{R} \setminus \{0\})$.

We can therefore show $\mathcal{U}_R \subseteq \mathbb{R}^{A \times \Theta}$ is open, dense, and of full Lebesgue measure by establishing that $\varphi_i^{-1}(\mathbb{R} \setminus \{0\})$ enjoys these properties for every $i = (a, \hat{\Theta}) \in \mathbb{I}$. First, note it is open because (by Berge's theorem) φ_i is continuous. To show it is of full measure (hence also dense), define $\bar{z} := [\mathbf{1}_{\bar{a}=a}]_{\bar{a} \in A, \bar{\theta} \in \Theta} \in \mathbb{R}^{A \times \Theta}$, and observe that $\varphi_i(u_R + \lambda\bar{z}) = \varphi_i(u_R) + \lambda$ for

²¹A more direct proof (as described in the main text) is available for the case of finite states. Because we require Lemma 6 to prove the finite-action case and Theorem 2, we find it convenient to use it here too.

any $u_R \in \mathbb{R}^{A \times \Theta}$ and any $\lambda \in \mathbb{R}$. Now, fixing some $\bar{\theta} \in \Theta$, observe that we can decompose the vector space of all R objectives as the direct sum $\mathbb{R}^{A \times \Theta} = Y \oplus Z$, where

$$Y := \{u_R \in \mathbb{R}^{A \times \Theta} : u_R(a, \bar{\theta}) = 0\} \text{ and } Z := \{\lambda \bar{z} : \lambda \in \mathbb{R}\}.$$

As $\varphi_i(y + \lambda \bar{z}) = \varphi(y) + \lambda$ for any $y \in Y$ and $\lambda \in \mathbb{R}$, and a singleton is Lebesgue-null in $\mathbb{R} \cong Z$, it follows from the law of iterated expectations that $\varphi_i^{-1}(0)$ is Lebesgue-null as well. The proposition follows. \square

Proof of Proposition 2. Lemma 7 names a set \mathcal{U}_R and shows it is open and dense in $\mathbb{R}^{A \times \Theta}$ with full Lebesgue measure. Meanwhile, by definition of \mathcal{U}_R , the PUBR property holds on $\Delta\Theta$ whenever $u_R \in \mathcal{U}_R$. The proposition then follows from Theorem 1. \square

Lemma 8. *Suppose $A, \Theta \subseteq \mathbb{R}$; the function u_R exhibits strictly increasing differences; every $\mu \in \Delta\Theta$ admits some $a \in \{\min A_R^*(\mu), \max A_R^*(\mu)\}$ such that $\int u_S(a, \cdot) d\mu = v(\mu)$; and for each $\bar{\theta} \in \{\min \text{supp } \mu_0, \max \text{supp } \mu_0\}$, either $\mu_0(\bar{\theta}) = 0$ or R 's best response at $\delta_{\bar{\theta}}$ is unique. Then, S has a unique equilibrium payoff.*

Proof. Let $\theta_L := \min \text{supp } \mu_0$ and $\theta_H := \max \text{supp } \mu_0$, and define the measurable set of beliefs

$$D := (\Delta[\text{supp } \mu_0]) \setminus \{\delta_\theta : \theta \in \{\theta_L, \theta_H\}, \mu_0(\theta) = 0\}.$$

Clearly, that $p \in \mathcal{I}(\mu_0)$ implies $p(D) = 1$, and so D is persuasion sufficient.

Fix any $\bar{\mu} \in D$ and $\varepsilon > 0$. Lemma 4 will deliver the present lemma if we can find some $\mu \in \Delta\Theta$ and $\gamma \in (0, 1)$ such that $\gamma\mu \leq \mu_0$ and $\hat{w}(\lambda\mu + (1-\lambda)\bar{\mu}) > v(\bar{\mu}) - \varepsilon$ for all sufficiently small $\lambda \in (0, 1]$. If $\bar{\mu} = \delta_{\bar{\theta}}$ for some $\bar{\theta} \in \{\theta_L, \theta_H\}$, then (since $\bar{\mu} \in D$) $\mu_0(\bar{\theta}) > 0$ and $|A_R^*(\bar{\mu})| = 1$, so that $(\mu, \gamma) = (\bar{\mu}, \mu_0(\bar{\theta}))$ is as desired. We therefore focus on the complementary case in which $\bar{\mu}$ lies strictly between δ_{θ_L} and δ_{θ_H} in the sense of first-order stochastic dominance.

By hypothesis, some $\bar{a} \in \{\min A_R^*(\bar{\mu}), \max A_R^*(\bar{\mu})\}$ has $\int u_S(\bar{a}, \cdot) d\mu = v(\bar{\mu})$. Without loss (because our hypotheses are invariant to reversing the order on both A and Θ), say $\bar{a} = \max A_R^*(\bar{\mu})$. That $\bar{\mu} \in D$ with $|A_R^*(\bar{\mu})| > 1$ implies, given the theorem's hypothesis on the prior, that $\bar{\mu} \neq \delta_{\theta_H}$.

Now, for any $\theta^* \in (\theta_L, \theta_H)$, let $\gamma_{\theta^*} := \mu_0(\theta^*, \theta_H) \in (0, 1)$ and define the conditional distribution $\mu_{\theta^*} := \frac{1}{\gamma_{\theta^*}} \mu_0[(\cdot) \cap (\theta^*, \theta_H)] \in \Delta\Theta$. By construction, we know $\gamma_{\theta^*} \mu_{\theta^*} \leq \mu_0$ for every $\theta^* \in (\theta_L, \theta_H)$. We will show that $(\mu, \gamma) = (\mu_{\theta^*}, \gamma_{\theta^*})$ are as desired when $\theta^* \in (\theta_L, \theta_H)$ is large enough.

Let $\eta > 0$ and neighborhood $B \subseteq \Delta\Theta$ of $\bar{\mu}$ be small enough that any $a \in A$ with $|a - \bar{a}| < \eta$ and any $\tilde{\mu} \in B$ have $\int u_S(a, \cdot) d\tilde{\mu} > v(\bar{\mu}) - \varepsilon$. We now claim that, for all sufficiently large $\theta^* \in (\theta_L, \theta_H)$, every $a \in A$ with $a \leq \bar{a} - \eta$ has $\int u_S(a, \cdot) d\mu_{\theta^*} < \int u_S(\bar{a}, \cdot) d\mu_{\theta^*}$. Assume otherwise, for a contradiction. Then (A being compact) some sequence (θ_n^*, a_n) from $(\theta_L, \theta_H) \times A$ converging to (θ_H, a') , for $a' \in A \cap (-\infty, \bar{a} - \eta]$, has $a_n \in A_R^*(\mu_{\theta_n^*})$ for every n . But $\int [u_S(a_n, \cdot) - u_S(\bar{a}, \cdot)] d\mu_{\theta_n^*} \geq 0$ converges to $\int [u_S(a', \cdot) - u_S(\bar{a}, \cdot)] d\delta_{\theta_H}$, which is strictly negative due to strictly increasing differences and $\bar{a} \in A^*(\bar{\mu})$ —a contradiction.

In what follows, fix $\theta^* \in (\theta_L, \theta_H)$ such that every $a \in A$ with $a \leq \bar{a} - \eta$ has $\int u_S(a, \cdot) d\mu_{\theta^*} < \int u_S(\bar{a}, \cdot) d\mu_{\theta^*}$; the previous paragraph established such θ^* exists. Letting $\mu := \mu_{\theta^*}$, and letting

$A_\lambda := A_R^*(\lambda\mu + (1 - \lambda)\bar{\mu})$ for each $\lambda \in (0, 1]$, we will next observe $A_\lambda \subseteq (\bar{a} - \eta, \bar{a} + \eta)$ for small enough $\lambda \in (0, 1]$. First, consider any $a \in A$ with $a \leq \bar{a} - \eta$. By definition of μ we have $\int u_S(a, \cdot) d\mu < \int u_S(\bar{a}, \cdot) d\mu$, and that $\bar{a} \in A_R^*(\bar{\mu})$ tells us $\int u_S(a, \cdot) d\mu \leq \int u_S(\bar{a}, \cdot) d\mu$. Therefore, $a \notin A_\lambda$. Having established $A_\lambda \subseteq (\bar{a} - \eta, \infty)$ for every $\lambda \in (0, 1]$, the desired containment for sufficiently small λ would follow if we knew that $\limsup_{\lambda \rightarrow 0} \max A_\lambda \leq \bar{a}$. But observe this inequality holds because $\max A_R^*$ is upper semicontinuous and $\bar{a} = \max A_0$.

Having seen $A_\lambda \subseteq (\bar{a} - \eta, \bar{a} + \eta)$ for sufficiently small $\lambda \in (0, 1]$, and observing that $\lambda\mu + (1 - \lambda)\bar{\mu} \in B$ for sufficiently small $\lambda \in (0, 1]$, it follows that μ is as required. \square

Proof of Theorem 2. Given Lemma 8, the theorem follows immediately if we establish that every $\mu \in \Delta\Theta$ admits $a \in \{\min A_R^*(\mu), \max A_R^*(\mu)\}$ such that $\int u_S(a, \cdot) d\mu = v(\mu)$. This property follows directly from the quasiconvexity/quasiconcavity property of the ordered model: The condition clearly holds if S's expected payoff is quasiconvex in the action, and it holds vacuously (because $|A_R^*(\mu)| \leq 2$) if R's expected payoff is strictly quasiconcave in the action. \square

A.3. Proofs for Section 3

To better understand the set of χ -PBE and the S payoffs they can generate, connecting this solution concept to the analysis of Lipnowski, Ravid, and Shishkin (2022) is useful. That paper defines a notion of a χ -**equilibrium** and characterizes the S payoffs such a solution can generate.²² Roughly, a χ -equilibrium specifies an experiment $\psi \in \Psi$, together with continuation play and continuation beliefs that satisfy the incentive and Bayesian properties in the partial-credibility game, but with no requirement that the initial experiment ψ be chosen optimally. From the definitions in the present paper and in LRS1, the following is immediate.

Fact 1. *The quadruple $\langle \psi, \sigma, \alpha, \pi \rangle$ is a χ -PBE if and only if:*

1. *For every $\tilde{\psi} \in \Psi$, the quadruple $\langle \tilde{\psi}, \sigma(\cdot, \tilde{\psi}), \alpha(\cdot, \tilde{\psi}), \pi(\cdot, \tilde{\psi}) \rangle$ is a χ -equilibrium.*
2. *We have $\psi \in \operatorname{argmax}_{\tilde{\psi} \in \Psi} s_{\tilde{\psi}}$, where $s_{\tilde{\psi}}$ is the S payoff induced by $\langle \tilde{\psi}, \sigma(\cdot, \tilde{\psi}), \alpha(\cdot, \tilde{\psi}), \pi(\cdot, \tilde{\psi}) \rangle$.*

In particular, every χ -PBE payoff is a χ -equilibrium payoff.

We will use the following notation, for an S payoff that LRS1 characterizes, throughout.²³

Notation 1. *Let $v_\chi^*(\mu_0)$ denote the highest χ -equilibrium S payoff (given prior μ_0).*

Toward constructing adversarial χ -PBE that give S an undesirable payoff in response to off-path experiment choices, we begin with a technical lemma showing any reporting protocol comprises part of a χ -equilibrium in which R always chooses from a given restricted set of best responses.

²²The model in LRS1 assumes the message space to be uncountable. However, given that $|M| \geq 2|\Theta|$, it follows readily from Carathéodory's theorem and Lemma 1 of LRS1 that the χ -equilibrium payoff set is unchanged.

²³The value's characterization (Theorem 1 of LRS1) remains valid in the present setting in light of Corollary 1 of LRS1.

Lemma 9. *If $\tilde{V} \subseteq V$ is a Kakutani correspondence and ψ is any official reporting protocol, then some χ -equilibrium $(\psi, \tilde{\sigma}, \tilde{\alpha}, \tilde{\pi})$ exists such that $u_S(\tilde{\alpha}) \in \tilde{V}(\tilde{\pi})$.*

Proof. Let $\Pi := (\Delta\Theta)^M$ be the set of all R belief mappings and define correspondences

$$\begin{aligned} \hat{S} : \Pi &\rightrightarrows \mathbb{R} \\ \tilde{\pi} &\mapsto \left[\max_{m \in M} \min \tilde{V}(\tilde{\pi}(m)), \max_{m \in M} \tilde{V}(\tilde{\pi}(m)) \right], \\ \hat{M} : \Pi &\rightrightarrows M \\ \tilde{\pi} &\mapsto \{m \in M : \tilde{V}(\tilde{\pi}(m)) \cap \hat{S}(\tilde{\pi}) \neq \emptyset\}. \end{aligned}$$

Observe that \hat{S} is Kakutani, since \tilde{V} is Kakutani and a finite maximum or minimum of upper or lower semicontinuous functions inherits the same semicontinuity. Therefore, \hat{M} is nonempty-valued with closed graph. Now let $\Sigma := (\Delta M)^\Theta$ and consider the correspondence mapping belief maps into S-IC influencing strategies (assuming R's strategy delivers S values from \tilde{V})

$$\begin{aligned} \hat{\Sigma} : \Pi &\rightrightarrows \Sigma \\ \tilde{\pi} &\mapsto \{\tilde{\sigma} \in \Sigma : \cup_{\theta \in \Theta} \text{supp}(\tilde{\sigma}(\theta)) \subseteq \hat{M}(\tilde{\pi})\}, \end{aligned}$$

and the correspondence mapping influencing strategies into consistent belief maps

$$\begin{aligned} \hat{\Pi} : \Sigma &\rightrightarrows \Pi, \\ \tilde{\sigma} &\mapsto \left\{ \tilde{\pi} \in \Pi : \tilde{\pi}(\theta|m) \int_{\Theta} \left[\chi \, d\psi(m|\cdot) + (1 - \chi)\tilde{\sigma}(m|\cdot) \right] d\mu_0 \right. \\ &\quad \left. = [\chi(\theta)\psi(m|\theta) + (1 - \chi(\theta))\tilde{\sigma}(m|\theta)] \mu_0(\theta), \forall \theta \in \Theta, m \in M \right\}. \end{aligned}$$

It then follows that $\hat{\Sigma}$ and $\hat{\Pi}$ are both Kakutani. Therefore, the Kakutani fixed point theorem delivers some $\tilde{\sigma} \in \Sigma$ and $\tilde{\pi} \in \Pi$ such that $\tilde{\sigma} \in \hat{\Sigma}(\tilde{\pi})$ and $\tilde{\pi} \in \hat{\Pi}(\tilde{\sigma})$. Now, take any $s_i \in \hat{S}(\tilde{\pi})$ and let $D := \tilde{\pi}(M)$. Note that $s_i \wedge \tilde{V}|_D$ is nonempty-valued and so admits a selector $\phi : D \rightarrow \mathbb{R}$.²⁴ Therefore, some $\hat{\alpha} : D \rightarrow \Delta A$ exists such that $u_S(\hat{\alpha}(m)) = \phi(m)$. Next, define $\tilde{\alpha} := \hat{\alpha} \circ \tilde{\pi} : M \rightarrow \Delta(A)$. It is then easy to verify that $(\psi, \tilde{\sigma}, \tilde{\alpha}, \tilde{\pi})$ is a χ -equilibrium. \square

The following lemma shows the payoff $\hat{w}(\mu_0)$ dominates some χ -equilibrium payoff.

Lemma 10. *Every official reporting protocol ψ admits some χ -equilibrium $(\psi, \tilde{\sigma}, \tilde{\alpha}, \tilde{\pi})$ with ex-ante S payoff weakly below $\hat{w}(\mu_0)$.*

Proof. Without loss, we can focus on the case that μ_0 is of full support. Indeed, if we construct a χ -equilibrium as desired (for official reporting protocol $\psi|_{\Theta_0}$) in the restricted model with state space $\Theta_0 := \text{supp} \mu_0$, then this equilibrium can be extended to a χ -equilibrium in the true model, by fixing any $\theta_0 \in \Theta_0$ and extending σ to Θ via $\sigma(\theta) := \sigma(\theta_0)$ for $\theta \in \Theta \setminus \Theta_0$.

²⁴That is, some $\phi : D \rightarrow \mathbb{R}$ has $\phi \leq s_i$ and $\phi(\mu) \in \tilde{V}(\mu)$ for every $\mu \in D$.

Note the lemma follows directly from Lemma 9 if we can find a Kakutani subcorrespondence $\tilde{V} \subseteq V$ such that the concave envelope of its upper selection satisfies $\text{cav}[\max \tilde{V}](\mu_0) \leq \hat{w}(\mu_0)$. Let us show $\tilde{V} := [w, z]$ has this property, where z is the upper semicontinuous envelope of w , given by

$$\begin{aligned} z: \Delta\Theta &\rightarrow \mathbb{R} \\ \mu &\mapsto \limsup_{\mu' \rightarrow \mu} w(\mu'). \end{aligned}$$

First, \tilde{V} is a Kakutani subcorrespondence of V since z is upper semicontinuous and lies above the lower semicontinuous function w (and hence lies below v). All that remains, then, is to show that $\hat{w}(\mu_0) \geq \text{cav}[\max \tilde{V}](\mu_0) = \hat{z}(\mu_0)$. To do so, let us establish the stronger claim that $\hat{z}|_D = \hat{w}|_D$, where $D \subseteq \Delta\Theta$ is the set of full-support beliefs.

Define

$$\begin{aligned} \tilde{z}: \Delta\Theta &\rightarrow \mathbb{R} \\ \mu &\mapsto \limsup_{\mu' \rightarrow \mu} \hat{w}(\mu'). \end{aligned}$$

It follows from concavity of \hat{w} that \tilde{z} is concave too. Hence, because $\tilde{z} \geq z$ and \tilde{z} is upper semicontinuous by construction, it follows that $\tilde{z} \geq \hat{z}$.²⁵ Moreover, Theorem 10.4 from Rockafellar (1970) implies the concave function $\hat{w}|_D$ is continuous. Hence, $\tilde{z}|_D = \hat{w}|_D$ by the definition of \tilde{z} . That $\tilde{z} \geq \hat{z} \geq \hat{w}$ then implies $\hat{z}|_D = \hat{w}|_D$. \square

Here, we provide a sufficient condition for a payoff to be compatible with χ -PBE for an arbitrary credibility level.

Lemma 11. *If $s \in [\hat{w}(\mu_0) \wedge v_\chi^*(\mu_0), v_\chi^*(\mu_0)]$, then s is a χ -PBE payoff for S .*

Proof. First, we argue a χ -equilibrium exists with ex-ante S payoff s . To that end, observe that Lemma 1 from LRS1 implies $(\delta_{\mu_0}, w(\mu_0), w(\mu_0))$ is a χ -equilibrium outcome (as witnessed by $k = \chi$ and $g = b = \delta_{\mu_0}$). But then, as Theorem 1 from LRS1 says $v_\chi^*(\mu_0)$ is the highest χ -equilibrium S payoff, it follows from Lemma 7 of LRS1 that every payoff in $[w(\mu_0), v_\chi^*(\mu_0)]$ is a χ -equilibrium S payoff. Thus, s is a χ -equilibrium S payoff because $\hat{w}(\mu_0) \wedge v_\chi^*(\mu_0) \geq w(\mu_0)$. So let $(\psi, \tilde{\sigma}, \tilde{\alpha}, \tilde{\pi})$ be some χ -equilibrium generating S payoff s .

Finally, to build a χ -PBE, construct σ , α , and π as follows. First, let $\sigma(\cdot, \psi) := \tilde{\sigma}$, $\alpha(\cdot, \psi) := \tilde{\alpha}$, and $\pi(\cdot, \psi) := \tilde{\pi}$. Second, given any $\tilde{\psi} \in \Psi \setminus \{\psi\}$, let $\langle \tilde{\psi}, \sigma(\cdot, \tilde{\psi}), \alpha(\cdot, \tilde{\psi}), \pi(\cdot, \tilde{\psi}) \rangle$ be some χ -equilibrium—which Lemma 10 exists—with ex-ante S value of at most $\hat{w}(\mu_0) \wedge v_\chi^*(\mu_0)$; this χ -equilibrium necessarily yields S payoff no greater than $v_\chi^*(\mu_0)$ by definition of the latter. Since $s \geq \hat{w}(\mu_0) \wedge v_\chi^*(\mu_0)$, the quadruple $\langle \psi, \sigma, \alpha, \pi \rangle$ is a χ -PBE as desired. \square

Next, we characterize the highest χ -PBE payoff S can attain; it coincides with her highest χ -equilibrium payoff.

²⁵In fact, one can show $\tilde{z} = \hat{z}$, but this fact is immaterial to the present argument.

Lemma 12. *The highest χ -PBE payoff for S is $v_\chi^*(\mu_0)$, her highest χ -equilibrium payoff. This payoff is weakly increasing in χ .*

Proof. By Fact 1, no χ -PBE payoff is strictly higher than $v_\chi^*(\mu_0)$. Meanwhile, Lemma 11 implies $v_\chi^*(\mu_0)$ is a χ -PBE payoff. The last statement follows from LRS1's Corollary 3. \square

Now, we characterize the set of all 1-PBE payoffs S can attain; it coincides with the 1-equilibrium payoffs.

Lemma 13. *The set of all 1-PBE payoffs for S is $[\hat{w}(\mu_0), \hat{v}(\mu_0)]$.*

Proof. Theorem 1 from LRS1 tells us $\hat{v}_1^* = \hat{v}$, and so Lemma 11 says all payoffs in $[\hat{w}(\mu_0), \hat{v}(\mu_0)]$ are 1-PBE payoffs for S. Conversely, every 1-PBE generates a 1-equilibrium by Fact 1, and so generates an equilibrium (in the sense of Definition 1) by throwing away the influencing S behavior, the set of 1-PBE payoffs is a subset of the set of equilibrium payoffs. Hence, the other containment follows from Proposition 0. \square

We can now prove the results on χ -PBE reported in the main text.

Proof of Proposition 3. Let S_χ denote the set of χ -PBE payoffs for S, and let $\bar{S} := [\hat{w}(\mu_0), \hat{v}(\mu_0)]$. First, Lemma 12 says $v^*(\mu_0) = \max S_\chi$, and $v^*(\mu_0) \leq \hat{v}(\mu_0) = \max \bar{S}$ by Theorem 1 from LRS1. Meanwhile, Lemma 11 implies $\bar{S} \cap (-\infty, \max S_\chi] \subseteq S_\chi$. Hence S_χ is weakly below \bar{S} in the strong set order. Finally, that $S_1 = \bar{S}$ is exactly Lemma 13. \square

Proof of Proposition 4. Fix any full-support prior $\mu_0 \in \Delta\Theta$, and define $\underline{s}_\chi := w_\chi^*(\mu_0)$ for each $\chi \in [0, 1]$; in particular, $\underline{s}_1 = \hat{w}(\mu_0)$ by Lemma 13. To prove the equivalence, it suffices to show that $\lim_{\chi \nearrow 1} \underline{s}_\chi = \underline{s}_1$. Further, Lemma 11 tells us $\underline{s}_\chi \leq \underline{s}_1$ for every $\chi \in [0, 1]$, so we need only show $\liminf_{\chi \nearrow 1} \underline{s}_\chi \geq \underline{s}_1$, which we do below.

Take an arbitrary $\varepsilon > 0$. By definition of \underline{s}_1 , some $p \in \mathcal{I}(\mu_0)$ exists such that $\int w \, dp > \underline{s}_1 - \varepsilon$. Moreover, by Lemma 1, we may further assume $|\text{supp}(p)| \leq |\Theta| \leq |M|$. For each $\mu \in \text{supp}(p)$, let $N(\mu) \subseteq \Delta\Theta$ be some open neighborhood of μ on which $w > w(\mu) - \varepsilon$, which exists because w is lower semicontinuous. Because $\text{supp}(p)$ is finite, which in particular implies $p(\mu) > 0$ for every $\mu \in \text{supp}(p)$, some $\underline{\chi} \in (0, 1)$ is such that

$$\frac{1}{\underline{\chi}p(\mu) + (1-\underline{\chi})} [\underline{\chi}p(\mu)\mu + (1-\underline{\chi})\Delta\Theta] \subseteq N(\mu)$$

for each $\mu \in \text{supp}(p)$, and so (since $\Delta\Theta$ is convex) the containment holds as well when we replace $\underline{\chi}$ with any $\chi \in (\underline{\chi}, 1)$.

Consider now, any $\chi \in (\underline{\chi}, 1)$, and fix some χ -PBE $\langle \psi, \sigma, \alpha, \pi \rangle$ generating S payoff $s \in \mathbb{R}$. Let $\tilde{\psi}_p \in \Psi$ be as defined in Proposition 0's proof, so that $p_{\tilde{\psi}_p} = p$. Modifying $\tilde{\psi}_p$ if necessary, we may assume without loss that any two distinct messages from $M_p := \{m \in M : \int_{\Theta} \tilde{\psi}_p(m|\cdot) \, d\mu_0 > 0\}$ would generate distinct beliefs. Hence, every belief $\mu \in \text{supp}(p)$ admits a unique $m_\mu \in M_p$ such that every $\theta \in \Theta$ has $\frac{\tilde{\psi}_p(m_\mu|\theta)\mu_0(\theta)}{\int_{\Theta} \tilde{\psi}_p(m_\mu|\cdot) \, d\mu_0} = \mu(\theta)$. If S chooses official reporting protocol $\tilde{\psi}_p$ and sends message m_μ for some $\mu \in \text{supp}(p)$, the

Bayesian property implies $\pi(m_\mu, \tilde{\psi}_p) \in \frac{1}{\chi p(\mu) + (1-\chi)} [\underline{\chi} p(\mu)\mu + (1-\chi)\Delta\Theta] \subseteq N(\mu)$, so that R's best response property implies S has continuation value exceeding $w(\mu) - \varepsilon$. But because S chooses $\psi \in \Psi$ optimally, and has the option to choose $\tilde{\psi}_p$, it must be that

$$\begin{aligned}
s &\geq \int_{\Theta} \left(\int_M \left[\int_A u_S(a) d\alpha(a|m, \tilde{\psi}_p) \right] d[\chi \tilde{\psi}_p(m|\theta) + (1-\chi) \sigma(m|\theta, \tilde{\psi}_p)] \right) d\mu_0(\theta) \\
&\geq \chi \int_{\Theta} \int_M \left[\int_A u_S(a) d\alpha(a|m, \tilde{\psi}_p) \right] d\tilde{\psi}_p(m|\theta) d\mu_0(\theta) + (1-\chi) \min w(\Delta\Theta) \\
&\geq \chi \int_{\Delta\Theta} (w - \varepsilon) dp + (1-\chi) \min w(\Delta\Theta) \\
&\geq \chi(\underline{s}_1 - \varepsilon) + (1-\chi) \min w(\Delta\Theta)
\end{aligned}$$

Because s was the payoff from an arbitrary χ -PBE, we learn that every $\chi \in (\underline{\chi}, 1)$ has $\underline{s}_\chi \geq \chi(\underline{s}_1 - \varepsilon) + (1-\chi) \min w(\Delta\Theta)$, which converges to $\underline{s}_1 - \varepsilon$ as χ converges to 1. Hence, $\liminf_{\chi \nearrow 1} \underline{s}_\chi \geq \underline{s}_1 - \varepsilon$. But ε was itself arbitrary, so that $\liminf_{\chi \nearrow 1} \underline{s}_\chi \geq \underline{s}_1$, as desired. \square

References

- Aliprantis, Charalambos D and Kim Border. 2006. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer Science & Business Media.
- Alonso, Ricardo and Odilon Câmara. 2018. "On the value of persuasion by experts." *Journal of Economic Theory* 174:103–123.
- Au, Pak Hung and Keiichi Kawai. 2020. "Competitive information disclosure by multiple senders." *Games and Economic Behavior* 119:56–78.
- Aumann, Robert J and Michael Maschler. 1995. *Repeated games with incomplete information*. MIT press.
- Balkenborg, Dieter, Josef Hofbauer, and Christoph Kuzmics. 2013. "Refined best reply correspondence and dynamics." *Theoretical Economics* 8 (1):165–192.
- . 2015. "The refined best-response correspondence in normal form games." *International Journal of Game Theory* 44:165–193.
- Bergemann, Dirk and Stephen Morris. 2019. "Information design: A unified perspective." *Journal of Economic Literature* 57 (1):44–95.
- Chakraborty, Archishman and Rick Harbaugh. 2010. "Persuasion by cheap talk." *American Economic Review* 100 (5):2361–2382.
- Crawford, Vincent P and Joel Sobel. 1982. "Strategic information transmission." *Econometrica* 50 (6):1431–1451.
- de Clippel, Geoffroy. 2008. "An axiomatization of the inner core using appropriate reduced games." *Journal of Mathematical Economics* 44 (3-4):316–323.
- Doval, Laura and Vasiliki Skreta. 2022. "Mechanism design with limited commitment." *Econometrica* 90 (4):1463–1500.
- Dworczak, Piotr and Giorgio Martini. 2019. "The simple economics of optimal persuasion." *Journal of Political Economy* 127 (5):1993–2048.
- Dworczak, Piotr and Alessandro Pavan. 2022. "Preparing for the worst but hoping for the best: Robust (Bayesian) persuasion." *Econometrica* 90 (5):2017–2051.
- Fréchette, Guillaume R, Alessandro Lizzeri, and Jacopo Perego. 2022. "Rules and commitment in communication: An experimental analysis." *Econometrica* 90 (5):2283–2318.
- Gentzkow, Matthew and Emir Kamenica. 2016. "Competition in persuasion." *The Review of Economic Studies* 84 (1):300–322.

- . 2017. “Bayesian persuasion with multiple senders and rich signal spaces.” *Games and Economic Behavior* 104:411–429.
- Guo, Yingni and Eran Shmaya. 2019. “The interval structure of optimal disclosure.” *Econometrica* 87 (2):653–675.
- Halac, Marina, Elliot Lipnowski, and Daniel Rappoport. 2022. “Addressing strategic uncertainty with incentives and information.” In *AEA Papers and Proceedings*, vol. 112. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, 431–437.
- Hedlund, Jonas. 2017. “Bayesian persuasion by a privately informed sender.” *Journal of Economic Theory* 167:229–268.
- Ichihashi, Shota. 2019. “Limiting Sender’s information in Bayesian persuasion.” *Games and Economic Behavior* 117:276–288.
- Inostroza, Nicolas and Alessandro Pavan. 2023. “Adversarial coordination and public information design.” *Available at SSRN 4531654* .
- Kamenica, Emir. 2019. “Bayesian Persuasion and Information Design.” *Annual Review of Economics* 11 (1).
- Kamenica, Emir and Matthew Gentzkow. 2011. “Bayesian persuasion.” *American Economic Review* 101 (6):2590–2615.
- Koessler, Frédéric and Vasiliki Skreta. 2022. “Informed Information Design.” .
- Kolotilin, Anton. 2018. “Optimal information disclosure: A linear programming approach.” *Theoretical Economics* 13 (2):607–635.
- Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnjuk, and Ming Li. 2017. “Persuasion of a privately informed receiver.” *Econometrica* 85 (6):1949–1964.
- Kosterina, Svetlana. 2022. “Persuasion with unknown beliefs.” *Theoretical Economics* 17 (3):1075–1107.
- Li, Fei and Peter Norman. 2021. “Sequential persuasion.” *Theoretical Economics* 16 (2):639–675.
- Li, Fei, Yangbo Song, and Mofei Zhao. 2023. “Global manipulation by local obfuscation.” *Journal of Economic Theory* 207:105575.
- Lipnowski, Elliot and Laurent Mathevet. 2018. “Disclosure to a psychological audience.” *American Economic Journal: Microeconomics* 10 (4):67–93.
- Lipnowski, Elliot, Doron Ravid, and Denis Shishkin. 2022. “Persuasion via weak institutions.” *Journal of Political Economy* 130 (10):2705–2730.

- Mathevet, Laurent, Jacopo Peregò, and Ina Taneva. 2020. “On information design in games.” *Journal of Political Economy* 128 (4):1370–1404.
- Milne, Alan Alexander. 1926. *Winnie-the-pooh*. EP Dutton.
- Min, Daehong. 2021. “Bayesian persuasion under partial commitment.” *Economic Theory* 72 (3):743–764.
- Moriya, Fumitoshi and Takuro Yamashita. 2020. “Asymmetric-information allocation to avoid coordination failure.” *Journal of Economics & Management Strategy* 29 (1):173–186.
- Morris, Stephen, Daisuke Oyama, and Satoru Takahashi. 2022. “Implementation via information design in binary-action supermodular games.” *Available at SSRN 3697335* .
- O’Brien, Richard C. 1976. “On the openness of the barycentre map.” *Mathematische Annalen* 223:207–212.
- Oyama, Daisuke and Satoru Takahashi. 2020. “Generalized Belief Operator and Robustness in Binary-Action Supermodular Games.” *Econometrica* 88 (2):693–726.
- Perez-Richet, Eduardo. 2014. “Interim bayesian persuasion: First steps.” *American Economic Review* 104 (5):469–474.
- Perez-Richet, Eduardo and Vasiliki Skreta. 2022. “Test design under falsification.” *Econometrica* 90 (3):1109–1142.
- Ravindran, Dilip and Zhihan Cui. 2022. “Competing persuaders in zero-sum games.” *arXiv preprint arXiv:2008.08517* .
- Rockafellar, R Tyrrell. 1970. *Convex Analysis*. 28. Princeton University Press.
- Skreta, Vasiliki. 2006. “Sequentially Optimal Mechanisms.” *Review of Economic Studies* 73 (4):1085–1111.
- Wu, Wenhao. 2022. “Sequential Bayesian Persuasion.” .
- Ziegler, Gabriel. 2020. “Adversarial bilateral information design.” *Working paper* .