Persuasion by an Imperfectly Informed Sender

Matt Schwartz*

April 2, 2025

Abstract

A sender commits to a garbling of a private signal to influence a receiver's action. I study the comparative statics of the sender's utility with respect to his information. His value function is a concave, Blackwell monotone function on the second-order belief space. If both players view a public signal, the informational environment is described by a distribution over the sender's information. I characterize an order over such distributions, which generalizes the one-agent special case of the main result from Bergemann and Morris (2016). At any fixed information, the dual to the sender's problem yields a convex price function on the belief space, which I show can be usefully interpreted as the indirect utility of a decision problem.

1 Introduction

The standard model of Bayesian persuasion, as introduced by Kamenica and Gentzkow (2011), allows the more-informed player (Sender, he) to choose any information structure for the less-informed player (Receiver, she) to view. The most notable assumption of the persuasion model is commitment power: Sender chooses an information structure before learning the state, truthfully communicates that choice to Receiver, and then truthfully reports the realized message. Many papers, such as Lipnowski, Ravid, and Shishkin (2022) and Guo and Shmaya (2021), relax various aspects of the commitment assumption. However, another key assumption of the model, which has received much less attention, is that Sender is perfectly informed at the time of communicating, in the sense that he can condition his message to Receiver on the true state. In many real-world communication scenarios, Sender does not have access to all payoff-relevant information. Consider a school designing its grading policy to maximize a student's expected job market outcomes. The school can condition its message (a transcript) on a student's performance in classes, but class performance is only partially informative about the student's ability as an employee, which is the payoff-relevant fact for

*I would like to thank my advisor, Marina Halac, for her outstanding guidance throughout the project. Elliot Lipnowski provided constant support and inspired many of my key ideas. Kai Hao Yang, Larry Samuelson, Dirk Bergemann, and Ernesto Rivera Mora gave useful feedback. All errors are my own.

employers. If Sender's information is fixed, this does not pose a problem. We can simply think of whatever Sender does know, such as class performance, as the set of states and rewrite utilities accordingly. This justifies the assumption of perfect information commonly used in the literature. However, if we want to understand how Sender's information affects outcomes, it is useful to keep the state space and utilities fixed while varying what Sender can communicate. This paper gives results on the comparative statics of Sender's utility with respect to his information.

I assume that Sender has access to a private signal, and he chooses an information structure that garbles this signal. This is the commitment analog to a cheap talk model where Sender views a partially informative signal, rather than viewing the true state, before choosing his message. The easiest comparative static for my setting is that Sender's utility improves when his information improves in the sense of Blackwell (1953). Since Sender commits to his strategy in advance, he could always ignore any extra information he gets. So, a more informative signal is just an expansion of the set of his feasible strategies. Since commitment power effectively makes the model non-strategic, Sender must do better with a larger choice set. However, it is initially unclear what we can say beyond Blackwell monotonicity.

I have two main results, Theorem 1 and Theorem 2, which respectively describe the global and local comparative statics of Sender's utility with respect to his information. To motivate Theorem 1, I imagine Sender and Receiver observing some public signal, in addition to Sender observing his private signal. Then the informational environment is fully described by a distribution over Receiver's second-order posteriors, which we can equivalently view as a distribution over Sender's information. Theorem 1 characterizes when one distribution over information is better for Sender than another distribution regardless of actions and utilities, in which case I say the first distribution dominates the second in the persuasion order. While it is necessary that Sender is more informed under the first distribution and Receiver is less informed about the state, this is not sufficient: Receiver must also know less about Sender's beliefs. It is not obvious how to capture this condition because the distribution over Sender's beliefs can change across environments. I consider which distributions over Receiver's beliefs Sender can induce in a given environment. Theorem 1 states that persuasion dominance is equivalent to this feasible set of distributions being larger.

Theorem 2 describes the local comparative statics of Sender's utility around some fixed information. If the state is binary and the set of actions is finite, then results from Dworczak and Martini (2019) imply the existence of a convex price function on the belief space, which corresponds to a supergradient of Sender's value function. I show that this convex function is the indirect utility of a decision problem which locally represents the persuasion problem. Optimal actions between the persuasion and decision problems line up, and Sender's utility is approximated by the utility in the decision problem at nearby information.

As I will show, we could rewrite Sender's problem by treating his beliefs as the states. Since Sender always knows his own belief, this rewriting gives a standard persuasion problem in the form from Kamenica and Gentzkow (2011). A change in Sender's information corresponds to a change in the prior of this rewritten problem. So, everything we know from general persuasion problems applies, including concavity of Sender's utility with respect to his information. One theme of the paper, however, is that we know more than what we would know in an arbitrary persuasion problem. For example, in addition to concavity of Sender's value function, we get Blackwell monotonicity, and there is no analogous property for arbitrary problems. Both of my main results crucially rely on the additional structure we get in the incomplete-information problem beyond what we would have in general.

After covering the related literature, in Section 2, I quickly present the model and motivate my key questions with an example. In Section 3, I define Sender's problem and obtain some easy results about his value function. Then, I look at the global comparative statics of Sender's utility in Section 4 by considering distributions over information. I define the persuasion order on the set of these distributions, and I characterize it with Theorem 1. In Section 5, I study local comparative statics by proving the existence of representative decision problems with Theorem 2. In Section 6, I graphically illustrate my results using the motivating example before concluding. Proofs of all results are in the appendix.

1.1 Related Literature

The persuasion literature began with Kamenica and Gentzkow (2011). My work relates to multiple areas of this literature. Many papers have considered modifications to players' information in the standard model. Whereas Sender commits to an information structure before his private signal realizes in my model, Perez-Richet (2014) considers what happens when Sender chooses the information structure after viewing his signal. The model becomes strategic in the sense that Sender's choice of information structure reveals something about the signal realization he saw. Tsakas and Tsakas (2021) look at the case that Sender can condition his message on the true state, but noise is added exogenously before Receiver views the message. Ball and Espín-Sánchez (2022) let Sender choose from a set of feasible experiments. My model is the special of theirs where Sender can garble any experiment in the feasible set and this feasible set has a most informative element.

Bergemann and Morris (2016) consider what an information designer can achieve by communicating with multiple agents. The one-agent special case of their model is a persuasion game where Sender is perfectly informed and Receiver has prior information. They show that all possible Senders are better off exactly when Receiver has less information. My most direct contribution to the literature is to generalize this ranking to include environments with imperfect Sender information. Then it does not only matter what Receiver knows about the state but also what she knows about Sender's belief; for example, if she always knew Sender's belief with certainty, then Sender could never persuade her of anything. As long as Receiver's prior information is public, the environment is described by a distribution over Receiver's second-order beliefs. Sender can choose any spread of this distribution by telling her about his beliefs. However, all that affects payoffs is the projection of this spread onto the corresponding distribution over first-order beliefs. This projection is non-injective, which means that the persuasion order cannot be reduced to a mean-preserving spread condition. Theorem 1 recovers the one-agent case of Bergemann and Morris (2016) as a corollary, but the result characterizes a much richer set of comparisons than that paper considers.

Two more relevant areas are the literature on mean-measurable persuasion problems and the related literature on the duality approach to persuasion. Mean-measurable problems, studied by Gentzkow and Kamenica (2016) and Dworczak and Martini (2019), have the unit interval as the state space and assume that Sender's utility only depends on the mean of Receiver's posterior belief. The special case of my model with a binary state is equivalent to a mean-measurable problem because the belief space can be identified with the unit interval. In this case, asking how Sender's utility changes with his information is equivalent to asking how Sender's utility changes with the prior of a mean-measurable problem. Dworczak and Martini (2019) study the dual to the mean-measurable problem and show that it gives a convex price function on the unit interval, which can be used to compute Sender's utility. Galperti, Levkun, and Perego (2023) consider a similar setup to mine, where Sender has access to an imperfectly informative signal, and they show the existence of a price function on the set of signal realizations. My work in Section 5 unifies these papers by identifying signal realizations with their induced beliefs, so that prices for realizations induce a convex price function on the belief space. Further, I reinterpret this price function as the indirect utility of a decision problem. I show that Sender's value function is the lower envelope of a family of decision problem utilities, which establishes a geometric relationship between persuasion and decision. I regularly use results from Dworczak and Kolotilin (2024), who study the dual persuasion problem in its most general form.

More broadly, this paper is part of the literature addressing how Sender's utility changes in communication games when communication-relevant parameters change. Green and Stokey (2007) investigate the effect of Sender's information on his utility in a certain class of communication equilibria. Lipnowski and Ravid (2020) compare Sender's utility in cheap talk to persuasion when his preferences are state independent, and Lipnowski, Ravid, and Shishkin (2022) extend this by varying a smooth parameter describing Sender's degree of commitment power. This paper asks a similar type of question, using a persuasion model and varying the signal available to Sender. Asking how Sender's persuasion utility changes with his private information is important because the persuasion utility is an upper bound on what Sender might achieve in any communication game with the same primitives. One way of rephrasing my central question is then to ask how Sender's best outcome across all communication models changes when his information changes.

2 Model

The primitives of the model are a finite state space Ω, a compact metrizable set of actions A, and continuous utilities for Sender and Receiver, $v, u : A \times \Omega \to \mathbb{R}$. Sender's private information is a measurable map $\pi : \Omega \to \Delta(S)$, where S is a compact metrizable space. $\mu_0 \in \Delta(\Omega)$ is the prior on the state space. The timing is as follows. First, Sender commits to an information structure $\sigma : S \to \Delta(A)$, and Receiver observes the choice of σ. A realization $s \in S$ is drawn according to π, and then $a \in A$ realizes according to σ given s. Receiver views the realization a and then takes an action. I assume that Receiver always takes a Sender-preferred action out of the set of her own optimal actions. As shown in Kamenica and Gentzkow (2011), it is without loss that we restrict Sender to choose A as the set of realizations for σ, in the sense that he could do no better with any richer messaging space.

In Kamenica and Gentzkow (2011), the model is the same except that Sender can choose any $\sigma : \Omega \to \Delta(A)$. So, if π is perfectly informative, the two models coincide. When π might not be perfectly informative, I argue that it is natural to view π as Sender's private information. This interpretation is justified by thinking about an analogous cheap talk model. Imagine that Sender views the realization s from π and then chooses an action recommendation. Since Sender might mix, his strategy is some $\sigma : S \to \Delta(A)$, where $\sigma(\cdot|s)$ is the distribution over his action recommendations after he views realization s. This is a cheap talk game where Sender's private information is π because π describes what Sender knows when he communicates. If we let Sender commit to σ before viewing π, then this game becomes my model.

Before moving on, I will make some notational and technical notes. I repeatedly refer to the sets $\Delta(\Omega), \Delta\Delta(\Omega)$, and $\Delta\Delta\Delta(\Omega)$, so I will simply write $\Delta^1(\Omega), \Delta^2(\Omega)$, and $\Delta^3(\Omega)$, respectively. Since Ω is finite, give $\Delta^1(\Omega)$ the Euclidean metric. In line with Dworczak and Kolotilin (2024), take $\Delta^2(\Omega)$ with the Kantorovich-Rubinstein metric $\rho_{KR}(\tau, \tau') = \sup \left\{ \int_{\Delta^1(\Omega)} pd(\tau - \tau') : p \in L_1(\Delta^1(\Omega)), p\,(\mu) = 0 \right\}$ for some arbitrary $\mu \in \Delta^1(\Omega)$, where $L_1(\Delta^1(\Omega))$ is the set of 1-Lipschitz functions on $\Delta^1(\Omega)$. Then ρ_{KR} metrizes the weak* topology on $\Delta^2(\Omega)$. Similarly, take $\Delta^3(\Omega)$ with the weak* topology.

2.1 Motivating Example

Consider the classic prosecutor-judge setup from Kamenica and Gentzkow (2011): a prosecutor (Sender) tries to convince a judge (Receiver) to convict a defendant. Say the defendant is guilty with prior probability $\mu_0 = \frac{1}{2}$, and the judge will convict exactly when her belief is at least $\frac{2}{3}$. The prosecutor designs an experiment to influence the judge's action, trying to maximize the chance she convicts. However, the twist on the standard problem is that with probability q, there is evidence revealing whether the defendant is guilty or innocent, and with probability $1 - q$, no evidence exists at all. The prosecutor's private information is a signal revealing the defendant's guilt or innocence if evidence exists and revealing nothing if no evidence exists. Before viewing this signal, he commits to an information structure, which specifies the probability he recommends each action after viewing each signal realization. The parameter q determines the prosecutor's information. My results characterize how his utility depends on q, or more generally, how his utility depends on any possible private information he might have.

One approach to the problem would be to treat it as a three-state persuasion problem, where the three states are that evidence exists and the defendant is innocent, evidence exists and the defendant is guilty, or no evidence exists. Then different values of q just correspond to different priors in this three-state problem. My paper, however, treats the problem as a two-state problem with a fixed prior $\frac{1}{2}$, where q just determines the prosecutor's choice set. To see the value of this, consider that without even specifying either player's utility, we know that the prosecutor's utility must increase in q: higher q corresponds to a higher likelihood that the prosecutor has access to evidence. However, we could find utilities in arbitrary three-state problems where priors corresponding to higher q are worse for the sender.

The main results in my paper correspond to two questions in this example. First, consider an alternative scenario where the judge has exogenous information revealing whether evidence exists. This does not give the judge any prior information about the state, but it might make the prosecutor worse off because he can no longer pretend to have evidence when he does not. So, how do we formalize and generalize this insight? This requires thinking about the global properties of Sender's value function, which I do in Section 4. Second, we know that if $q = 1$, then the setting is exactly the one from Kamenica and Gentzkow (2011). We can find the prosecutor's utility by concavifying his indirect utility and evaluating at the prior. However, can we generalize this to $q \leq 1$? In Section 5, I show that we can do this by looking for the worst decision problem out of a family of decision problems and evaluating utility at Sender's information.

3 Sender's Value Function

Before defining Sender's problem, I introduce some standard objects. Given any utility $w : A \times \Omega \to \mathbb{R}$, I write $w_a(\mu) = \int_\Omega w(a, \omega) d\mu(\omega)$, so that w_a is an affine function on $\Delta^1(\Omega)$ for each $a \in A$. When Receiver's belief is $\mu \in \Delta^1(\Omega)$ at the time she makes her decision, her set of optimal actions is $a^*(\mu) = \arg\max_{a \in A} u_a(\mu)$. Assuming that Receiver chooses a Sender-preferred optimal action, Sender's expected utility when Receiver's belief is μ is

$$V(\mu) = \max_{a \in a^*(\mu)} v_a(\mu).$$

This is Sender's indirect utility function $V : \Delta^1(\Omega) \to \mathbb{R}$. V describes everything relevant to Sender's utility, though in Section 5 I will return to a description of the problem using A, v, and u.

3.1 Sender's Problem

Recall that Sender's private information is $\pi : \Omega \to \Delta(S)$. If Sender chooses the information structure $\sigma : S \to \Delta(A)$, Receiver's information is the garbling of π by σ. Since A is a sufficiently rich messaging space, Sender's problem is equivalent to choosing *any* garbling of

π. That is, he can choose any $\pi' \precsim_B \pi$, where $\precsim_B$ denotes the Blackwell order as defined by Blackwell (1953). If we let $\tau_S \in \Delta^2(\Omega)$ denote the distribution over beliefs induced by π under prior μ_0, then a choice of $\pi' \precsim_B \pi$ is equivalent to a choice of $\tau \preceq_{MPS} \tau_S$, where $\preceq_{MPS}$ is the mean-preserving spread order. So, Sender chooses any distribution over Receiver's beliefs which is less dispersed than the distribution over his own beliefs.

As shown by Kamenica and Gentzkow (2011), when the distribution over Receiver's posterior beliefs is $\tau \in \Delta^2(\Omega)$, Sender's expected utility is $\int_{\Delta^1(\Omega)} V d\tau$. So, when π induces distribution over beliefs τ_S for Sender, his maximal utility is

$$V^*(\tau_S) := \sup_{\tau \preceq_{MPS} \tau_S} \int_{\Delta^1(\Omega)} V d\tau \tag{P_B}$$

The problem (P_B) is the beliefs-based formulation of the incomplete-information persuasion problem. Its value function $V^* : \Delta^2(\Omega) \to \mathbb{R}$ is the object of study in this paper. Since Sender might face any $\tau_S \in \Delta^2(\Omega)$ given the right choice of μ_0 and π, understanding this function is equivalent to understanding Sender's utility across all priors and private signals.[1] Note that this formulation of Sender's problem is similar to the one in the mean-measurable literature, such as Gentzkow and Kamenica (2016) and Dworczak and Martini (2019). Those papers take the state space as [0,1], and Sender's utility is pinned down by the mean of Receiver's posterior. The prior is a cdf F on [0,1]. Sender's problem is to choose a mean-preserving contraction of F as the distribution over posterior means. In the case that Ω is binary, so that we can identify $\Delta^1(\Omega)$ with [0,1], my problem is then equivalent to a mean-measurable problem.

(P_B) has a clear relationship with the problem from Kamenica and Gentzkow (2011). The objective is the same, but their Bayes-plausibility constraint, which is that $\int_{\Delta^1\Omega} \mu d\tau(\mu) = \mu_0$, is replaced by the constraint that τ is a mean-preserving contraction of τ_S. Note that if τ_S only supports degenerate beliefs, that is, if π is perfectly informative, then these constraints are equivalent. Otherwise, the mean-preserving contraction constraint is more restrictive. Say that a function $f : \Delta^2(\Omega) \to \mathbb{R}$ is *Blackwell monotone* if for every $\tau \succeq_{MPS} \tau'$, $f(\tau) \geq f(\tau')$. Then we can generalize the insight that the mean-preserving contraction constraint is more restrictive than the Bayes-plausibility constraint with the following proposition, which states that Sender benefits from more information.

Proposition 1. *V^* is Blackwell monotone.*

Instead of modifying the Bayes-plausibility constraint, an alternative approach is to rewrite the problem in the standard persuasion form. Specifically, treat $\Delta^1(\Omega)$ as the set of states, where $\mu \in \Delta^1(\Omega)$ represents Sender's belief after viewing his signal. The prior on the new state space $\Delta^1(\Omega)$ is just τ_S, the distribution over Sender's beliefs. Define $T : \Delta^2(\Omega) \to \Delta^1(\Omega)$ by $T(\tau) = \int_{\Delta^1(\Omega)} \mu d\tau(\mu)$, so that T is the natural projection from the

[1] In many ways, it is most interesting to think about how Sender's utility changes when just his signal changes with the prior fixed. So, we could restrict the domain of V^* to a subset of distributions which average to the same prior. However, this is an unnecessary restriction.

second-order belief space to the first-order belief space. When Receiver has posterior belief τ about Sender's belief, he has belief $T(\tau)$ about the state, so that Sender's utility is $V(T(\tau))$. Sender can choose any distribution $\psi \in \Delta^3(\Omega)$ over Receiver's posterior τ which averages to the prior τ_S; this is Bayes plausibility. This rewriting of the problem gives the intuition for the following result.

Proposition 2. *For every $\tau_S \in \Delta^2(\Omega)$,*

$$V^*(\tau_S) = \sup_{\psi \in \Delta^3(\Omega)} \int_{\Delta^2(\Omega)} V \circ T d\psi$$
$$s.t. \int_{\Delta^2(\Omega)} \tau d\psi(\tau) = \tau_S.$$

Therefore, $V^* = \text{cav}(V \circ T)$, where cav *is the concavification operator. Further, there exists a solution ψ to the above problem, and for such ψ, the push-forward $T_\star\psi \in \Delta^2(\Omega)$ solves problem* (P_B)*.*

The key to the proof is to note that for a distribution ψ which averages to τ_S, $T_\star\psi$ is a mean-preserving contraction of τ_S, and any mean-preserving contraction can be written in this way. This provides a mapping between the two ways of writing the problem. The result shows us that changing Sender's information corresponds to changing the prior in a persuasion problem, just a problem with a higher-order state space. This is mathematically important because it allows us to use what we know about standard persuasion problems. For example, Dworczak and Kolotilin (2024) prove that these problems always have an optimal solution, and this implies existence of a solution to the mean-preserving contraction formulation of the problem, as is stated in the above proposition. Two more important properties of V^* are below.

Proposition 3. *V^* is concave and upper semicontinuous.*

The proof leverages the rewriting of the problem from Proposition 2: the persuasion utility, written as a function of the prior, is always concave and upper semicontinuous. Propositions 1 and 3 establish the key properties of V^* which I investigate in Sections 4 and 5. I analyze the global implications of these properties by thinking about distributions over information. Distributions which place more weight on more informative information give higher expected utilities because of Blackwell monotonicity, and less dispersed distributions give higher expected utility because of concavity. Further, Propositions 1 and 3 imply that V^* is the infimum of some family of Blackwell monotone, affine functions on $\Delta^2(\Omega)$. In Section 5, I show that these affine functions naturally correspond to decision problems.

3.2 Sender's Value for More Information

The incomplete-information persuasion problem derives properties from the underlying persuasion problem, whether we think of the underlying problem as (Ω, A, v, u) or (Ω, V). Proposition 2 tells us that V^* is just the value function in a standard persuasion problem with state space $\Delta^1(\Omega)$. However, there is more structure to this problem than an arbitrary persuasion problem on state space $\Delta^1(\Omega)$. An arbitrary problem would have a concave value function but not a Blackwell monotone one in general. Sender's indirect utility in an arbitrary problem might be any upper semicontinuous function over $\Delta^2(\Omega)$, whereas in this problem the indirect utility must be constant among distributions with the same mean. This is expressed in Proposition 2 by composing the underlying utility V with an affine map, T, whose fibers are such collections of distributions. The following result gives one implication of this fact.

Proposition 4. *If V is convex, then $\tau = \tau_S$ solves problem (P_B), and $V^*(\tau_S) = \int_{\Delta^1(\Omega)} V d\tau_S$. If V is concave, then $\tau = \delta_{\mu_0}$ solves (P_B), and $V^*(\tau_S) = V(\mu_0)$.*

The most instructive way to see this is that when V is convex (concave), then $V \circ T$ is also convex (concave). So, when V is convex, Sender will give Receiver as much information as he can, and when V is concave, he will give no information. This is a special case of a general point, which is that since T is affine and surjective, $V \circ T$ retains all of the structure of V. This is why we can use the underlying problem to understand the imperfect-information problem.

Note that when V is convex and continuous, we know it represents the indirect utility resulting from some decision problem, so asking how Sender's utility changes with τ_S is the same as asking how a decision maker's utility changes with respect to information. This question is relatively straightforward; in particular, V^* is affine in this case. Theorem 2 shows that this insight generalizes in a local sense. Of course, V^* will not always be affine, so that we cannot globally represent Sender's problem as a decision problem. However, it turns out that around any τ_S, V^* is locally approximated by a decision problem.

When V is concave, Proposition 4 tells us that Sender does not value information because the prior alone pins down his utility. A natural question is when this happens more generally. Fix τ_S with $T(\tau_S) = \mu_0$. Say that *Sender benefits from more information at τ_S* if there exists some $\tau \in \Delta^2(\Omega)$ such that $T(\tau) = \mu_0$ and $V^*(\tau) > V^*(\tau_S)$. Say that Sender *locally benefits from more information at τ_S* if for every $\varepsilon > 0$, there exists some $\tau \in B_\varepsilon(\tau_S)$ such that $T(\tau) = \mu_0$ and $V^*(\tau) > V^*(\tau_S)$, where $B_\varepsilon(\tau_S)$ is the open ball of radius ε with respect to ρ_{KR}. Then we have the following.

Proposition 5. *For any τ_S, the following are equivalent:*

(i) Sender does not locally benefit from more information at τ_S

(ii) Sender does not benefit from more information at τ_S

(iii) $V^*(\tau_S) = (\mathrm{cav}(V))(\mu_0)$, the concavification of V at $T(\tau_S) = \mu_0$.

Of course, if Sender could locally benefit from information, then he does not reach his full-information utility, which Kamenica and Gentzkow (2011) show is $(\mathrm{cav}(V))(\mu_0)$. The converse is more interesting and comes from concavity and Blackwell monotonicity of V^*.

4 Distributions over Information

4.1 The Persuasion Order

Now that I have established some basic properties of V^*, I will analyze its global properties by considering distributions over Sender's information. The motivation is to consider that instead of the original setup, the informational environment is described by a signal $\pi : \Omega \to \Delta(S \times R)$. Only Sender observes the private realization $s \in S$, but both Sender and Receiver observe the public realization $r \in R$. Sender's problem is to choose an information structure $\sigma : S \times R \to \Delta(A)$. Since r is publicly observed, this is equivalent to both players observing r and updating their beliefs before Sender chooses an information structure $\sigma : S \to \Delta(A)$.

When r realizes, both players update their beliefs about the full realization (s, r). This full realization determines Sender's belief. So, the realization of r induces a belief $\tau_r \in \Delta^2(\Omega)$ about Sender's belief. At this point, Sender's problem is exactly the one considered in Sections 3 and 4, so his expected utility is $V^*(\tau_r)$. The marginal on R then yields some distribution $\psi \in \Delta^3(\Omega)$ over possible τ_r. Then Sender's utility is

$$\int_{\Delta^2(\Omega)} V^* d\psi.$$

In the context of the public information setting, ψ describes the distribution over Receiver's second-order beliefs. However, given the above expression, we could also interpret it as the distribution over Sender's information. Given two signals π and π', a natural question is when π is objectively better for Sender than π'. To formalize this, we want to understand when the above quantity is higher when integrating with respect to ψ than some other ψ', regardless of the utilities and actions in the persuasion problem. Define

$$\mathcal{V} := \left\{ V : \Delta^1(\Omega) \to \mathbb{R} | V \text{ is bounded and upper semicontinuous} \right\}.$$

As explained by Dworczak and Kolotilin (2024), $\mathcal{V}$ is the set of indirect utilities resulting from some persuasion problem (Ω, A, v, u), where we range across all compact metrizable spaces A and continuous utilities v and u. Given $V \in \mathcal{V}$, write V^* for the corresponding value function. This leads to the following definition.

Definition 1. *For any $\psi, \psi' \in \Delta^3(\Omega)$, define*

$$\psi \succsim_P \psi' \text{ if } \int_{\Delta^2(\Omega)} V^* d\psi \geq \int_{\Delta^2(\Omega)} V^* d\psi' \text{ for every } V \in \mathcal{V}.$$

In this case, say that ψ dominates ψ' in the persuasion order.

The persuasion order is analogous to the Blackwell order. The Blackwell order asks when one distribution over beliefs τ is preferable to τ' for every decision maker. When a decision maker's utility is w, write $\widehat{w}$ as his indirect utility. Then τ Blackwell dominates τ' when $\int_{\Delta^1(\Omega)} \widehat{w} d\tau \geq \int_{\Delta^1(\Omega)} \widehat{w} d\tau'$ for every w. On the other hand, the persuasion order asks when one distribution over information ψ is preferable to ψ' for every Sender. This happens when $\int_{\Delta^2(\Omega)} V^* d\psi \geq \int_{\Delta^2(\Omega)} V^* d\psi'$ for every V. While convexity is the key property of $\widehat{w}$ used to characterize the Blackwell order, concavity and Blackwell monotonicity are the key properties of V^*. Sender would like ψ to place more weight on more informative distributions because of Blackwell monotonicity, and he would like ψ to be more spread because of concavity. In this way, the persuasion order reflects global properties of Sender's value function V^*.

4.2 Characterizing Persuasion Dominance

As always, Sender's expected utility is pinned down by the ex-ante distribution over Receiver's posterior beliefs about the state. Specifically, if Sender's indirect utility is V and the distribution over Receiver's posteriors is τ, Sender's expected utility is $\int_{\Delta^1(\Omega)} V d\tau$. Thus, a natural approach to characterizing the persuasion order is to figure out which distributions over beliefs Sender could induce in different informational environments. Sender's ex-ante problem is to choose how much information to give Receiver about Sender's beliefs on top of Receiver's prior information ψ. Then Sender can choose any $\phi \succeq_{MPS} \psi$ as the distribution over Receiver's second-order posteriors. When Receiver has second-order posterior $\tau \in \Delta^2(\Omega)$, her belief about the state is $T(\tau) \in \Delta^1(\Omega)$. Therefore, if the distribution over her second-order posteriors is ϕ, then the distribution over her first-order posteriors is the push-forward $T_\star\phi \in \Delta^2(\Omega)$. Then we can rewrite Sender's problem as choosing any $T_\star\phi$ for $\phi \succeq_{MPS} \psi$.

Lemma 1. *For any $V \in \mathcal{V}$ and $\psi \in \Delta^3(\Omega)$,*

$$\int_{\Delta^2(\Omega)} V^* d\psi = \max_{\phi \succeq_{MPS} \psi} \int_{\Delta^1(\Omega)} V d(T_\star\phi).$$

The proof effectively shows that the problem of choosing a mean-preserving contraction of τ for each $\tau \in \Delta^2(\Omega)$ is equivalent to choosing a mean-preserving spread ϕ of ψ and then taking $T_\star\phi$. With this rewriting of Sender's problem, we can think of $\{T_\star\phi | \phi \succeq_{MPS} \psi\}$ as Sender's choice set. Of course, if one ψ gives him a larger choice set than another ψ', then he is better off under ψ than ψ'. However, the converse is also true: if all possible Senders prefer ψ to ψ', then the choice set under ψ must be larger.

Theorem 1. *For any $\psi, \psi' \in \Delta^3(\Omega)$, $\psi \succsim_P \psi'$ if and only if $\{T_\star\phi | \phi \succeq_{MPS} \psi\} \supseteq \{T_\star\phi | \phi \succeq_{MPS} \psi'\}$.*

The proof of the theorem relies heavily on Lemma 1. The backward implication is straightforward because the set containment implies an expansion of Sender's choice set. The other direction is less obvious. If there is some distribution over beliefs τ which is feasible under ψ' but not ψ, then we can use the Hahn-Banach separation theorem to find an affine function greater at τ than it is anywhere in the choice set under ψ. By the Riesz representation theorem, this affine function corresponds to a continuous $V : \Delta(\Omega) \to \mathbb{R}$, and then a Sender with utility V prefers ψ' to ψ.

One interesting implication of the proof is that we could restrict to continuous V without changing the persuasion order. Even further, we could look only at continuous V induced by a finite set of actions because we could uniformly approximate any continuous V by piecewise-affine, continuous functions. Continuous V induced by finite actions have nice properties, including that they are Lipschitz continuous, implying that the corresponding V^* is superdifferentiable by Dworczak and Martini (2024). I do not exploit that fact in this paper, but it might be useful for future work. Additionally, the ability to restrict to finite actions shows that the result above is truly a generalization of the one-receiver case of Bergemann and Morris (2016) because that paper assumes finite actions.

4.3 Implications of Theorem 1

Theorem 1 is much stronger than it would be if we considered arbitrary persuasion problems on state space $\Delta^1(\Omega)$. In that case, the setting would be similar to Bergemann and Morris (2016), just with a higher-order state space, and the condition for dominance would be $\psi \preceq_{MPS} \psi'$. This is more demanding than the condition in Theorem 1. When we consider imperfect-information persuasion problems, ψ and ψ' do not need to be ranked in the mean-preserving spread order because they do not need to have the same mean. For example, we could shift the support of ψ' in Blackwell-improving directions and obtain a persuasion-dominant ψ.

The structure of imperfect-information problems is reflected in Theorem 1: all that matters for Sender's utility is Receiver's *first-order* posterior, so we can rewrite Sender's problem as a choice of distribution over these first-order posteriors. This shows up when we apply the map $T_\star(\cdot)$, which sends distributions over second-order posteriors to distributions over first-order posteriors. Since this mapping is non-injective, the set containment is a weaker condition than $\psi \preceq_{MPS} \psi'$, leading to a much richer set of comparisons. One challenge of characterizing the persuasion order is to rank ψ and ψ' which might not have the same mean. We need to capture the intuition that Sender has more information and Receiver has less information about Sender's beliefs. However, the latter notion is not well-defined in general since the prior over Sender's beliefs can change between the environments. The theorem gives a precise notion of these two conditions.

One easy case is when Sender is fully informed. Note that when ψ describes the environment, Sender's information about the state is $\int_{\Delta^2(\Omega)} \tau d\psi(\tau)$; this is just Bayes plausibility of Receiver's second-order beliefs. Receiver's prior information about the state is $T_\star\psi$ because ψ is the distribution over her second-order beliefs. Thus, for any $\psi \in \Delta^3(\Omega)$, define

$$\tau_S(\psi) := \int_{\Delta^2(\Omega)} \tau \, d\psi(\tau), \quad \tau_R(\psi) := T_\star\psi.$$

Further, say that some $\tau \in \Delta^2(\Omega)$ is perfectly informative if τ only supports degenerate beliefs.

Corollary 1. *Let $\psi, \psi' \in \Delta^3(\Omega)$ such that $\tau_S(\psi)$ and $\tau_S(\psi')$ are perfectly informative. Then $\psi \succsim_P \psi'$ if and only if $\tau_R(\psi) \preceq_{MPS} \tau_R(\psi')$.*

The above corollary recovers the result from the one-agent case of Bergemann and Morris (2016). When Sender has perfect information about the state, the problem is simpler because Receiver's beliefs about Sender's beliefs are the same as her beliefs about the state. Specifically, the important facts for the proof are that the distribution over Sender's beliefs is fixed between ψ and ψ' and that the support of this distribution is linearly independent in $\Delta^1(\Omega)$. In this case, we can effectively think of $T_\star(\cdot)$ as a bijection, and then the condition from Theorem 1 can be be rewritten as a property of $\tau_R(\psi)$ and $\tau_R(\psi')$.

An appealing thought is that we might be able to easily generalize the corollary above. In particular, it might initially feel plausible that whenever Sender is better informed in some environment and Receiver is less informed, Sender must be better off. Corollary 1 tells us that this holds when Sender is perfectly informed. Further, when Receiver has no information, Sender is better off exactly when he has more information; it is easy to see this by considering that Sender's utility might be convex. Indeed, Sender having more information and Receiver having less information are both necessary for persuasion dominance.

Corollary 2. *Let $\psi, \psi' \in \Delta^3(\Omega)$. If $\psi \succeq_P \psi'$, then $\tau_S(\psi) \succeq_{MPS} \tau_S(\psi')$ and $\tau_R(\psi) \preceq_{MPS} \tau_R(\psi')$.*

The above result is intuitive when considering that V might be convex or concave. If V is convex, Sender wants more information to give to Receiver. If V is concave, Sender wants Receiver to have less information so he can keep her less informed.

The converse to the above corollary does not hold, however. It might be that Sender is better informed under one ψ than another ψ' and Receiver is less informed, but Sender is better off under ψ'. The issue is that conditions on $\tau_S(\psi)$ and $\tau_R(\psi)$ do not capture what Receiver knows about Sender's beliefs. I demonstrate this with a counterexample. Consider the motivating example from Section 2. The state is binary, and the prior is $\mu_0 = \frac{1}{2}$. With probability $q = \frac{1}{2}$, Sender knows the state with certainty, and with probability $\frac{1}{2}$, he knows nothing and retains his prior beliefs. Say ψ describes an environment where it is public whether Sender has information, and ψ' describes an environment where nothing is public. Both Sender's and Receiver's information is the same between ψ and ψ' because the only difference is that Receiver knows whether Sender is informed under ψ, but this does not affect her belief about the state. Thus, $\tau_S(\psi) = \tau_S(\psi')$ and $\tau_R(\psi) = \tau_R(\psi')$, and in particular, $\tau_S(\psi) \succeq_{MPS} \tau_S(\psi')$ and $\tau_R(\psi) \preceq_{MPS} \tau_R(\psi')$. However, ψ does not dominate ψ' in the persuasion order. Under ψ', Sender can garble his private information to always give Receiver posterior 0.4 or 0.6. However, this is not available under ψ' because whenever Receiver knows that Sender's belief is 0.5, her posterior will be 0.5. By Theorem 1, this implies that $\psi \not\succeq_P \psi'$.

A useful future direction would be to characterize the persuasion order in terms of the signal π. I focus on ψ because I primarily use the public signal setup as motivation for considering distributions over information. However, a characterization in terms of π could give more insight into the literal interpretation of the persuasion order as ranking environments with private and public information.

5 Representative Decision Problems

Now, I turn to studying the local comparative statics of Sender's utility. As in Section 3, fix a persuasion problem (Ω, A, v, u) and some $\tau_S \in \Delta^2(\Omega)$. I begin by adding two assumptions, which I will maintain throughout the entire section.

Assumption 1. $|\Omega| = 2$*, and* A *is finite.*

These assumptions will guarantee that V^* is superdifferentiable. It might be possible to let Ω be an arbitrary finite set, but I only have the main result for a binary state. I discuss this and alternative sufficient assumptions at the end of the section.

5.1 Actions-Based Formulation of Sender's Problem

The beliefs-based approach is elegant but loses the connection to the primitives A, v, and u. So, I now rewrite Sender's problem directly in terms of these objects. If Sender's information is $\pi : \Omega \to \Delta(S)$, then identify each signal realization $s \in S$ with its induced belief in $\Delta^1(\Omega)$ under the prior μ_0.[2] Then have Sender choose a map $\sigma : \Delta^1(\Omega) \to \Delta(A)$, where $\sigma(\cdot|\mu)$ is the distribution over Sender's action recommendations when his belief is μ. From Kamenica and Gentzkow (2011), it is without loss to assume that Sender chooses an obedient signal in the sense that Receiver always takes the recommended action. So, when Sender recommends a and has belief μ, his utility is $v_a(\mu)$. He maximizes the expectation of this utility subject to the obedience constraints, which state that Receiver prefers an action a over any other a' when a is recommended. This leads to the actions-based formulation of Sender's problem (P):[3]

$$\sup_{\sigma:\Delta^1(\Omega)\to\Delta(A)} \int_{\Delta^1(\Omega)} \int_A v_a(\mu) d\sigma(a|\mu) d\tau_S(\mu) \tag{P}$$
$$\text{s.t.} \int_{\Delta^1(\Omega)} \left[u_a(\mu) - u_{a'}(\mu)\right] \sigma(a|\mu) d\tau_S(\mu) \geq 0 \quad \forall a, a' \in A.$$

To use the results in this section to understand the problem I have studied previously, we need to know that the problem above is equivalent to the beliefs-based formulation. The

[2] Assume without loss that every s is supported by $\pi(\cdot|\omega)$ for some ω so the belief induced by s is well defined.

[3] Even though A is finite, I write integrals rather than sums to suggest that we might be able to consider infinite A with additional assumptions, as I discuss in 5.4. Similarly, I will define the dual problem in a way that allows for this generalization.

following proposition establishes this. With Assumption 1, we can also guarantee that the solution is superdifferentiable, which will be crucial for the main result.

Proposition 6. *The problem* (P) *has a solution, and its value is* $V^*(\tau_S)$. *Further,* V^* *is superdifferentiable.*

For now, note that without the obedience constraint, Sender's problem (P) would just be a decision problem. Sender chooses an optimal action, or distribution over actions, at each of his beliefs to maximize the expectation of his utility v, and his information is τ_S. The dual problem will incorporate the obedience constraint into Sender's objective. I return to this point in 5.3.

5.2 Decision Problems

In this section, a decision problem refers to a utility $w : A \times \Omega \to \mathbb{R}$, where A is the set of actions from the persuasion problem. Given a decision problem $w : A \times \Omega \to \mathbb{R}$, define the indirect utility $\hat{w} : \Delta^1(\Omega) \to \mathbb{R}$ by

$$\widehat{w}(\mu) := \max_{a \in A} w_a(\mu)$$

for any $\mu \in \Delta^1(\Omega)$, so that $\widehat{w}$ is a convex function on the belief space. When a decision maker faces problem w with information τ, he chooses an optimal action at each possible realization of his belief μ, which gives utility $\widehat{w}(\mu)$. Of course, he could equivalently commit in advance to a strategy $\sigma : \Delta^1(\Omega) \to \Delta(A)$, where $\sigma(\cdot|\mu)$ is the distribution over his actions when his realized belief is μ. Then his expected utility is

$$\widetilde{w}(\tau) := \max_{\sigma:\Delta^1(\Omega)\to\Delta(A)} \int_{\Delta^1(\Omega)} \int_A w_a(\mu) d\sigma(a|\mu) d\tau(\mu) = \int_{\Delta^1(\Omega)} \widehat{w} d\tau.$$

So, $\widetilde{w} : \Delta^2(\Omega) \to \mathbb{R}$ is an affine function on the second-order belief space, and by Blackwell's theorem, it is Blackwell monotone. Call any σ which solves the above problem a solution to w at τ.

We wish to say that a decision problem w represents the persuasion problem (P) at τ_S when we could equivalently have the decision maker face w with information τ_S or have Sender face (P) with information τ_S. Then optimal actions and utilities should line up between the problems. Even stronger, we would like the utility under w to approximate the utility under (P) at information close to τ_S. These properties motivate the following.

Definition 2. *Say that a decision problem* w *represents the problem* (P) *at* τ_S *if the following both hold:*

(i) w *and* (P) *share a solution* $\sigma : \Delta^1(\Omega) \to \Delta(A)$ *at* τ_S

(ii) $\widetilde{w}$ *is a supergradient to* V^* *at* τ_S.

Now, I return to analyzing (P). Let $M(A)$ be the set of positive measures on A. Then the dual to (P) is

$$\begin{aligned}\inf_{\lambda:A\to M(A)} \sup_{\sigma:\Delta^1(\Omega)\to\Delta(A)} &\int_{\Delta^1(\Omega)} \left(\int_A v_a(\mu) d\sigma(a|\mu)\right) d\tau_S(\mu) \\ &+ \int_{A\times A} \left(\int_{\Delta^1(\Omega)} \left[u_a(\mu) - u_{a'}(\mu)\right] \sigma(a|\mu) d\tau_S(\mu)\right) d\lambda(a'|a).\end{aligned}$$

The goal is to treat the above as a decision problem for any fixed choice of λ. Specifically, given $\lambda : A \to M(A)$, define $w^\lambda : A \times \Omega \to \mathbb{R}$ by

$$w^\lambda(a,\omega) := v(a,\omega) + \int_A \left[u(a,\omega) - u(a',\omega)\right] d\lambda(a'|a).$$

Define the family of decision problems $\mathcal{W} := \{w^\lambda | \lambda : A \to M(A)\}$. Each element of $\mathcal{W}$ is some adjustment of Sender's utility so that he prefers Receiver to be better off according to a specific weighting of the actions. If the decision maker chooses action a, then a higher value of $\lambda(a'|a)$ makes him better off if Receiver preferred a to a' and worse off if Receiver preferred a' to a. Define the problem (D) as

$$\inf_{w\in\mathcal{W}} \widetilde{w}(\tau_S). \tag{D}$$

In other words, the problem (D) is to find the worst decision problem in $\mathcal{W}$ for a decision maker who has information τ_S. Now, I can state the result.

Theorem 2. *For any τ_S, there exists a solution w to* (D) *which represents* (P) *at τ_S.*

From the definition of w^λ, it is not too difficult to show that (D) is equivalent to the dual problem to (P) as previously written. This fact together with strong duality for (P) are the key steps in the proof of Theorem 2. Crucially, strong duality relies on superdifferentiability of V^*, which was established in Proposition 6 and is guaranteed by Assumption 1.

A vague intuition for the result is that each λ gives a different adjustment of Sender's utility to account for Receiver's utility. If w^λ is a representative decision problem at τ_S, then instead of having Sender persuade Receiver, he could simply choose an action to maximize his adjusted utility. If Receiver could view that action, and this were her only information, then she would take the same action if she were made to choose; this is the case because Sender's strategy satisfies the obedience constraints. The right choice of λ makes this true because Receiver knows that Sender is adjusting his decision to make her better off. The theorem tells us that such a λ exists, and that the expected adjustment to Sender's utility is 0.

The most important implication of Theorem 2 is that it gives a new interpretation of the price function studied by Dworczak and Martini (2019) in the context of my problem. Say that a convex function $p : \Delta^1(\Omega) \to \mathbb{R}$ is a price function for (P) at τ_S if the map $\tau \mapsto \int_{\Delta^1(\Omega)} p d\tau$ is a supergradient to V^* at τ_S. Then Theorem 2 implies the following.

Corollary 3. *For any τ_S, there is a representative decision problem w such that $\widehat{w}$ is a price function for* (P) *at τ_S.*

Dworczak and Martini (2019) imagine Sender endowed with his information τ_S and choosing to buy posterior beliefs μ with prices $p(\mu)$. My result gives a more direct interpretation. Consider that Sender faces the representative decision problem w instead of his persuasion problem, but he still has information τ_S. Then his expected utility is $\int_{\Delta^1(\Omega)} \widehat{w} d\tau_S = V^*(\tau_S)$, and he could take the same actions under this problem as he would recommend in the persuasion problem.

The next implication of Theorem 2 is a geometric connection between decision problems and persuasion problems.

Corollary 4. *Let $\widetilde{\mathcal{W}} = \{\widetilde{w} : w \in \mathcal{W}\}$. Then*

$$V^* = \min \widetilde{W},$$

taking the minimum pointwise.

Given a persuasion problem with utilities $v, u : A \times \Omega \to \mathbb{R}$, we can read off the family of decision problems $\mathcal{W}$ by taking w^λ for each $\lambda : A \to M(A)$. The above tells us that we could find V^* by taking the lower envelope of this family.

5.4 Relaxing Assumption 1

Assumption 1 guarantees that V^* is superdifferentiable, which is exactly what is needed to establish Theorem 2. As long as V^* is superdifferentiable at some τ_S, the dual to the actions-based formulation has a solution, and this solution gives a representative decision problem. The assumption that A is finite cannot be done away with altogether: as Dworczak and Kolotilin (2024) point out, infinite A can give a non-superdifferentiable value function. Alternatively, we could impose that V is Lipschitz continuous, which implies that $V \circ T$ is Lipschitz, and then V^* is superdifferentiable from Dworczak and Kolotilin (2024). However, this is an undesirable assumption because it is imposed on the indirect utility V rather than the primitives (Ω, A, v, u).

It is unclear whether we need the state to be binary. I use this assumption because it guarantees the existence of a dual solution at every τ_S, which is shown by Dworczak and Martini (2019). However, I find it plausible that the same holds for arbitrary finite state spaces. To prove this, we would need that V^* has bounded steepness everywhere; that is, for every τ, $\sup_{\tau'} \frac{V^*(\tau') - V^*(\tau)}{\rho_{KR}(\tau', \tau)} < \infty$. From Dworczak and Kolotilin (2024), this implies superdifferentiability of V^*. It might be that this holds for any arbitrary persuasion problem with a finite set of actions. Even if this does not hold, we could restrict to τ_S with finite support, in which case the linear program (P) is finite and thus satisfies strong duality. This is explained by Galperti, Levkun, and Perego (2023).

6 Discussion

6.1 Revisiting the Example

Before concluding, I will provide some graphical illustration of my results using the prosecutor-judge example from Section 2. Say the prosecutor gets utility 1 when the judge convicts and 0 when she acquits. Since the judge convicts exactly when her belief is at least $\frac{2}{3}$, Sender's indirect utility is $V(\mu) = \mathbb{1}\{\mu \geq \frac{2}{3}\}$.

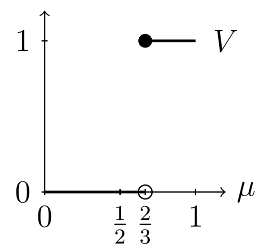

Call the judge's actions a and c for acquit and convict. The prosecutor's problem is to choose an information structure $\sigma : \{0, \frac{1}{2}, 1\} \to \Delta(\{a, c\})$, where $\sigma(\cdot|\mu)$ is the distribution over his action recommendations when his belief is μ. He wants to maximize the probability he recommends c while inducing belief $\frac{2}{3}$ from this recommendation.

Remember that q is the probability that evidence exists about the defendant's guilt or innocence. Then the distribution over the prosecutor's belief is $\tau_q = \frac{q}{2}\delta_0 + (1-q)\delta_{\frac{1}{2}} + \frac{q}{2}\delta_1$. We want to evaluate $V^*(\tau_q)$ for each $q \in [0, 1]$. From Theorem 2, we know that this is the same as finding a minimal decision problem $w \in \mathcal{W}$ at τ_q. Any w is described by two lines w_a and w_c, representing the utilities of the two actions. We must have the corresponding indirect utility $\widehat{w}$ lie above V, and out of such decision problems, we choose the one to minimize $\int_{\Delta^1(\Omega)} \widehat{w} d\tau_q$. It is easy to check that w and w' shown below solve this problem for the respective cases $q \leq \frac{1}{2}$ and $q \geq \frac{1}{2}$.

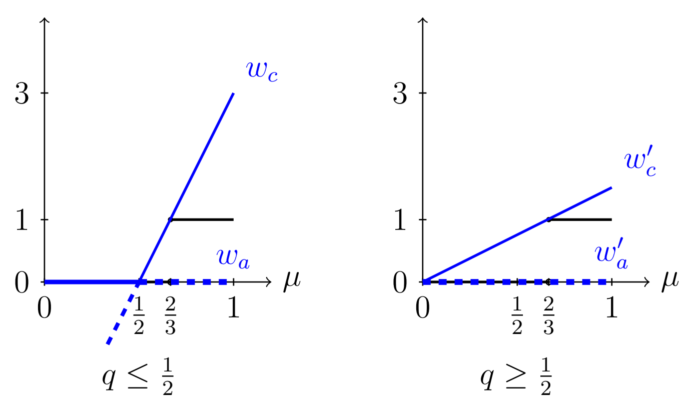

The upper envelope of w_c and w_a, which is the indirect utility $\widehat{w}$, is drawn in solid blue on the left, and similarly for $\widehat{w'}$ on the right. When $q \leq \frac{1}{2}$, the prosecutor's utility is $V^*(\tau_q) = \int_{\Delta^1(\Omega)} \hat{w} d\tau_q = \frac{q}{2}0 + (1-q)0 + \frac{q}{2}3 = \frac{3}{2}q$. When $q \geq \frac{1}{2}$, his utility is $V^*(\tau_q) = \int_{\Delta^1(\Omega)} \widehat{w'} d\tau_q = \frac{q}{2}0 + (1-q)\frac{3}{4} + \frac{q}{2}\frac{3}{2} = \frac{3}{4}$. Beyond using these decision problems to compute the value function, Theorem 2 tells us that they represent the prosecutor's optimal strategy in the persuasion problem. Regardless of q, he recommends c when his belief is 1. The decision problems show us this because $w_c(1) > w_a(1)$ and $w'_c(1) > w'_a(1)$. When $q < \frac{1}{2}$, the prosecutor mixes between his recommendations at belief $\frac{1}{2}$, which we see because $w_c(\frac{1}{2}) = w_a(\frac{1}{2})$. On the other hand, if $q \geq \frac{1}{2}$, the prosecutor always recommends c at belief $\frac{1}{2}$: when q is high, he can do this while still inducing belief $\frac{2}{3}$ from recommendation c because he has belief $\frac{1}{2}$ less often. So, when $q \geq \frac{1}{2}$, he only mixes when his belief is 0, which we see because 0 is the indifference point between w'_c and w'_a.

From the calculations for $V^*(\tau_q)$ above, we see that the prosecutor's utility is affine with respect to q at all points except $q = \frac{1}{2}$. This is because all $q < \frac{1}{2}$ have w as their representative decision problem, and all $q > \frac{1}{2}$ have w'. For a fixed decision problem, utility is affine with respect to information. We can visualize this by drawing the affine functions $\widetilde{w}(\tau_q)$ and $\widetilde{w'}(\tau_q)$, which represent the utility of the respective decision problems at information τ_q. The lower envelope of these functions is $V^*(\tau_q)$.

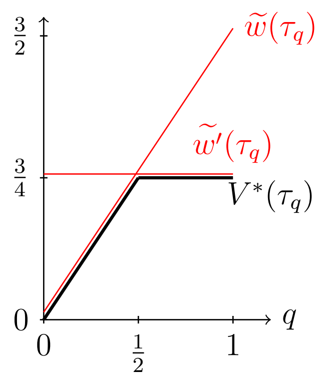

The graph above illustrates Corollary 4, which states that the value function is the lower envelope of a family of decision problem utilities. Note that in this graph, the horizontal axis corresponds to a projection of $\Delta^2(\Omega)$ onto the space of distributions τ_q for $q \in [0, 1]$. This is different from the previous two graphs, where the horizontal axis represented $\Delta^1(\Omega)$. The functions V^*, $\widetilde{w}$, and $\widetilde{w'}$ are all functions of information rather than beliefs.

Having graphed $V^*(\tau_q)$, I can now illustrate Theorem 1. If there is public information revealing whether evidence exists, then with probability q, the judge knows that the prosecutor has full information, and with probability $1-q$, the judge knows that the prosecutor has no information. In the first outcome, the prosecutor's utility is $V^*(\tau_1) = \frac{3}{4}$, and in the second, his utility is $V^*(\tau_0) = 0$. Then his expected utility is $qV^*(\tau_1) + (1-q)V^*(\tau_0)$, which is shown by the dashed line below.

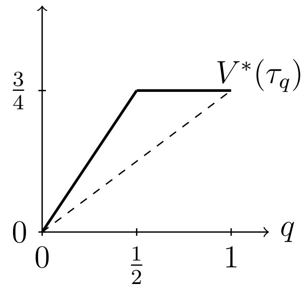

The fact that the dashed line lies below V^* is a result of concavity of V^*. This corresponds to the idea that the prosecutor wants the judge to be less informed about his belief. The fact that this line and V^* are increasing in q corresponds to Blackwell monotonicity. Theorem 1 captures and generalizes these properties. This example illustrates a key point from Section 4.3: the judge knowing whether evidence exists gives her no prior information about the state, but it tells her something about the prosecutor's beliefs. This makes the prosecutor strictly worse off whenever $q \in (0, 1)$, which is shown in the graph above.

The example highlights two economic facts about imperfect-information persuasion problems. First, in a very intuitive sense, Blackwell monotonicity and concavity mean that information has positive but decreasing marginal value to Sender. This is made precise when considering a 1-dimensional, Blackwell-ranked projection of $\Delta^2(\Omega)$ as I do above: $V^*(\tau_q)$ has a weakly decreasing derivative with respect to q where this derivative is defined. After some threshold, more information has no value. Specifically, once Sender has access to a more informative signal than his optimal full-information persuasion strategy, he does not benefit from more information. Second, in communication problems with complex informational environments, higher-order beliefs matter. The judge learning whether the prosecutor has access to evidence does not change the judge's first-order beliefs, but it does change her second-order beliefs, which hurts the prosecutor. This phenomenon never occurs when Sender has full information or Receiver has no information. Since I only consider environments where Receiver's information is public, we can still restrict attention to the distribution of her second-order beliefs. However, if Sender and Receiver each have access to private, potentially correlated information, we would have to consider their entire belief hierarchies to describe the environment. It is an interesting question whether characterizing the persuasion order over this broader class of environments is tractable.

6.2 Connection Between the Main Results

While Theorems 1 and 2 answer distinct questions, it is easy to tell a story which connects them. Theorem 2 tells us that instead of letting Sender persuade Receiver, we could equivalently have Receiver choose the worst decision problem $w \in \mathcal{W}$ for Sender. Then, Sender views his private signal, and he chooses an action to maximize his utility in w. If there is public information, then Receiver knows something about Sender's beliefs when she chooses w. If she knows more about Sender's beliefs, she can choose a worse decision problem. However, if Sender has more information about the state, he can take a better action in every decision problem. This gives an intuition for Theorem 1. One interesting question is how to write down a model which justifies Receiver's objective as choosing the worst decision problem for Sender.

Both results suggest an analogy between decision and persuasion. The indirect utility of a decision problem is convex because it is the upper envelope of a family of affine functions on $\Delta^1(\Omega)$, each corresponding to an action. The value function for a persuasion problem is concave and Blackwell monotone because it is the lower envelope of a family of Blackwell monotone, affine functions on $\Delta^2(\Omega)$, each corresponding to a decision problem. The Blackwell order on $\Delta^2(\Omega)$ ranks information for decision problems, and its properties derive from convexity of decision problem indirect utilities. The persuasion order on $\Delta^3(\Omega)$ ranks informational environments for persuasion problems, and its properties derive from concavity and Blackwell monotonicity of persuasion value functions. Formalizing this analogy might be an interesting direction.

6.3 Future Research

The most direct way to build on my research is to strengthen Theorems 1 and 2. Specifically, characterizing the persuasion order in terms of signals π would be a useful addition to Theorem 1. Relaxing Assumption 1 by letting Ω be any finite set, if possible, would make Theorem 2 more widely applicable.

Both results could yield useful economic applications. There are many communication problems where the more-informed agent has imperfect information and both players share public information. Theorem 1 gives a ranking of such environments. Theorem 2 could be useful for understanding optimal signal structures in persuasion problems, which is often a key application of the duality approach to persuasion. Since I keep the problem written in terms of primitives in Section 5, the result might let us phrase properties of optimal signals directly in terms of action recommendations. One specific question of interest is to characterize which changes in information lead to changes in the optimal signal structure. This is closely tied to the question of which changes in information lead to a different representative decision problem. For example, in 6.1, we saw that the prosecutor's signal structure changes at $q = \frac{1}{2}$ because there was a different representative decision problem for $q < \frac{1}{2}$ and $q > \frac{1}{2}$. Generalizing this observation could lead to important applications of representative decision problems.

References

[1] Ball, I. and J.-A. Espin-Sanchez (2021): "Experimental Persuasion," *Cowles Foundation Discussion Paper* No. 2298.

[2] Bergemann, D. and S. Morris (2016): "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 11(2), 487–522.

[3] Blackwell, D. (1953): "Equivalent Comparisons of Experiments," *The Annals of Mathematical Statistics*, 24, 265–272.

[4] Dworczak, P. and A. Kolotilin (2024): "The Persuasion Duality," *Theoretical Economics*, 19 (4), 1701-1755.

[5] Dworczak, P. and G. Martini (2019): "The Simple Economics of Optimal Persuasion," *Journal of Political Economy*, 127, 1993–2048.

[6] Galperti, S., A. Levkun, and J. Perego (2024): "The Value of Data Records," *Review of Economic Studies*, 91, 1007–1038.

[7] Gentzkow, M. and E. Kamenica (2011): "Bayesian Persuasion," *American Economic Review*, 101, 2590–2615.

[8] Gentzkow, M. and E. Kamenica (2016): "A Rothschild–Stiglitz Approach to Bayesian Persuasion," *American Economic Review*, 106, 597–601.

[9] Gretsky, N. E., J. M. Ostroy, and W. R. Zame (2002): "Subdifferentiability and the Duality Gap," *Positivity*, 6, 261–274.

[10] Green, J. R. and N. L. Stokey (2007): "A Two-Person Game of Information Transmission," *Journal of Economic Theory*, 135, 90–104.

[11] Guo, Y. and E. Shmaya (2021): "Costly miscalibration," *Theoretic Economics*, 16 (2), 477-506.

[12] Lipnowski, E., L. Mathevet, and D. Wei (2020): "Attention Management," *American Economic Review: Insights*, 2, 17–32.

[13] Lipnowski, E. and D. Ravid (2020): "Cheap Talk with Transparent Motives," *Econometrica*, 88, 1631–1660.

[14] Lipnowski, E., D. Ravid, and D. Shishkin (2022): "Persuasion via Weak Institutions," *Journal of Political Economy*, 130(10), 2705–2730.

[15] Perez-Richet, E. (2014): "Interim Bayesian Persuasion: First Steps," *American Economic Review: Papers & Proceedings*, 104(5), 469–474.

[16] Tsakas, E. and N. Tsakas (2021): “Noisy Persuasion,” *Games and Economic Behavior*, 130, 44–61.

Appendix: Proofs

The appendix contains restatements and proofs for all results throughout the paper.

Proposition 1. *V^* is Blackwell monotone.*

Proof. Let $\tau_S \succeq_{MPS} \tau'_S$. By transitivity of $\succeq_{MPS}$, $\{\tau : \tau \preceq_{MPS} \tau_S\} \supseteq \{\tau : \tau \preceq_{MPS} \tau'_S\}$. Thus,

$$V^*(\tau_S) = \sup_{\tau \preceq_{MPS} \tau_S} \int_{\Delta^1(\Omega)} V d\tau \geq \sup_{\tau \preceq_{MPS} \tau'_S} \int_{\Delta^1(\Omega)} V d\tau = V^*(\tau'_S).$$

□

Proposition 2. *For every $\tau_S \in \Delta^2(\Omega)$,*

$$V^*(\tau_S) = \sup_{\psi \in \Delta^3(\Omega)} \int_{\Delta^2(\Omega)} V \circ T d\psi \quad s.t. \int_{\Delta^2(\Omega)} \tau d\psi(\tau) = \tau_S.$$

Therefore, $V^* = \text{cav}(V \circ T)$, where cav is the concavification operator. Further, there exists a solution ψ to the above problem, and for such ψ, the push-forward $T_\star\psi \in \Delta^2(\Omega)$ solves problem (P_B).

Proof. For any $\tau_S \in \Delta^2(\Omega)$, we have

$$\{\tau \in \Delta^2(\Omega) : \tau \preceq_{MPS} \tau_S\} = \{T_\star\psi : \psi \in \Delta^3(\Omega), \int_{\Delta^2(\Omega)} \tau d\psi(\tau) = \tau_S\},$$

as is proven in the online appendix to Lipnowski, Mathevet, and Wei (2020). Further, for any V and ψ, we have $\int_{\Delta^2(\Omega)} V \circ T d\psi = \int_{\Delta^1(\Omega)} V d(T_\star\psi)$. These two facts give the second equality in the following:

$$V^*(\tau_S) = \sup_{\tau \preceq_{MPS} \tau_S} \int_{\Delta^1(\Omega)} V d\tau = \sup_{\psi \in \Delta^3(\Omega)} \int_{\Delta^2(\Omega)} V \circ T d\psi \quad \text{s.t.} \int_{\Delta^2(\Omega)} \tau d\psi(\tau) = \tau_S.$$

Dworczak and Kolotilin (2024) prove that persuasion problems in the form above have a solution, and the value function is the concavification of the indirect utility at the prior. Thus, the problem above has a solution ψ, and its value is $\text{cav}(V \circ T)$. The problem satisfies their assumptions because $\Delta^2(\Omega)$ is given the Kantorovich-Rubinstein metric, V is upper semicontinuous from Kamenica and Gentzkow (2011), and thus $V \circ T$ is upper semicontinuous because T is continuous. Given an optimum ψ, the mapping between the two problems implies that $T_\star\psi$ is optimal in the mean-preserving contraction formulation.

□

Proposition 3. *V^* is concave and upper semicontinuous.*

Proof. From Proposition 2, $V^*(\tau_S)$ is the utility of the persuasion problem with utility $V \circ T$ and prior τ_S. Then both properties follow from the results of Dworczak and Kolotilin (2024) for arbitrary persuasion problems. □

Proposition 4. *If V is convex, then $\tau = \tau_S$ solves problem (P_B), and $V^*(\tau_S) = \int_{\Delta^1(\Omega)} V d\tau_S$. If V is concave, then $\tau = \delta_{\mu_0}$ solves (P_B), and $V^*(\tau_S) = V(\mu_0)$.*

Proof. If V is convex, then the map $\tau \mapsto \int_{\Delta^1(\Omega)} V d\tau$ is Blackwell monotone by Blackwell's theorem. So, the problem (P_B) is solved by taking taking $\tau = \tau_S$ since this τ is a mean-preserving spread of all elements of the feasible set. This implies that $V^*(\tau_S) = \int_{\Delta^1(\Omega)} V d\tau_S$.

If V is concave, then $-V$ is convex, so that $\int_{\Delta^1(\Omega)} V d\tau$ decreases in the spread of τ. Thus, it is optimal to choose $\tau = \delta_{\mu_0}$ because this is a mean-preserving contraction of every element of the feasible set. Then $V^*(\tau_S) = \int_{\Delta^1(\Omega)} V d\delta_{\mu_0} = V(\mu_0)$. □

Proposition 5. *For any τ_S, the following are equivalent:*

(i) Sender does not locally benefit from more information at τ_S

(ii) Sender does not benefit from more information at τ_S

(iii) $V^*(\tau_S) = (\mathrm{cav}(V))(\mu_0)$, the concavification of V at $T(\tau_S) = \mu_0$.

Proof. I prove that (i) implies (ii), (ii) implies (iii), and (iii) implies (i). If (i) holds, then for some $\varepsilon > 0$ and every $\tau \in B_\varepsilon(\tau_S)$ with $T(\tau) = \mu_0$, $V^*(\tau) = V^*(\tau_S)$. Fix such ε. Let $\bar{\tau}$ denote the perfectly informative distribution with mean μ_0. Choose $\delta \in (0,1)$ sufficiently small such that $(1-\delta)\tau_S + \delta\bar{\tau} \in B_\varepsilon(\tau_S)$. Such δ must exist because ρ_{KR} is induced by the Kantorovich-Rubinstein norm on $\Delta^2(\Omega)$. Then $V^*(\tau_S) = V^*((1-\delta)\tau_S + \delta\bar{\tau}) \geq (1-\delta)V^*(\tau_S) + \delta V^*(\bar{\tau})$, where the equality is by the choice of δ and the inequality is by concavity of V^*. By $\delta \in (0,1)$, this inequality implies $V^*(\tau_S) \geq V^*(\bar{\tau})$. Since $\bar{\tau}$ is a mean-preserving spread of all τ such that $T(\tau) = \mu_0$ and V^* is Blackwell monotone, we have $V^*(\tau_S) \geq V^*(\bar{\tau}) \geq V^*(\tau)$ whenever $T(\tau) = \mu_0$. Thus, Sender does not benefit from more information at τ_S.

If (ii) holds, then $V^*(\tau_S) \geq V^*(\bar{\tau})$, so $V^*(\tau_S) = V^*(\bar{\tau})$ by Blackwell monotonicity. Note that under $\bar{\tau}$, Sender just faces the full-information persuasion problem, so $V^*(\tau_S) = V^*(\bar{\tau}) = (\mathrm{cav}(V))(\mu_0)$ from Kamenica and Gentzkow (2011).

If (iii) holds, then Sender achieves his full-information utility, which is higher than any partial-information utility, and in particular higher than any utility from information near τ_S. So, Sender does not locally benefit from more information.

□

Lemma 1. *For any $V \in \mathcal{V}$ and $\psi \in \Delta^3(\Omega)$,*

$$\int_{\Delta^2(\Omega)} V^* d\psi = \max_{\phi \succeq_{MPS} \psi} \int_{\Delta^1(\Omega)} V d(T_\star \phi).$$

Proof. Let $V \in \mathcal{V}$, so that $V^*(\tau) = \max_{\tau' \preceq_{MPS} \tau} \int_{\Delta^1(\Omega)} V d\tau'$ for every τ. First, I argue that there is a solution to the problem

$$\max_{\phi \succeq_{MPS} \psi} \int_{\Delta^1(\Omega)} V d(T_\star \phi),$$

so that the maximum is well-defined. The set of mean-preserving spreads of ψ is compact, and the map $T_\star(\cdot) : \Delta^3(\Omega) \to \Delta^2(\Omega)$ is continuous. Thus, the set $\{T_\star \phi | \phi \succeq_{MPS} \psi\}$ is compact. Further, the map $\tau \mapsto \int_{\Delta^1(\Omega)} V d\tau$ is upper semicontinuous because V is upper semicontinuous. Therefore, there is a $\tau \in \{T_\star \phi | \phi \succeq_{MPS} \psi\}$ which attains the maximum of $\int_{\Delta^1(\Omega)} V d\tau$ on this set. For some $\phi \succeq_{MPS} \psi$, $T_\star \phi = \tau$, so this ϕ solves the problem above.

Define $T' : \Delta^3(\Omega) \to \Delta^2(\Omega)$ as $T'(\phi) = \int_{\Delta^2(\Omega)} \tau d\phi(\tau)$ for all $\phi \in \Delta^3(\Omega)$. Let $\mathcal{F} = \{f : \Delta^2(\Omega) \to \Delta^3(\Omega) | T'(f(\tau)) = \tau \quad \forall \tau \in \Delta^2(\Omega)\}$. These are all maps sending second-order beliefs τ to distributions which average to τ. Then $\{\phi \in \Delta^3(\Omega) | \phi \succeq_{MPS} \psi\} = \{\int_{\Delta^2(\Omega)} f \psi | f \in \mathcal{F}\}$.

Now, we can rewrite the left-hand-side as follows:

$$\begin{aligned}
\int_{\Delta^2(\Omega)} V^* d\psi &= \int_{\Delta^2(\Omega)} \left[\max_{\tau' \preceq_{MPS} \tau} \int_{\Delta^1(\Omega)} V d\tau' \right] d\psi(\tau) \\
&= \int_{\Delta^2(\Omega)} \left[\max_{\psi', T'(\psi') = \tau} \int_{\Delta^1(\Omega)} V d(T_\star \psi') \right] d\psi(\tau) \\
&= \max_{f \in \mathcal{F}} \int_{\Delta^2(\Omega)} \int_{\Delta^1(\Omega)} V d(T_\star f(\tau)) d\psi(\tau).
\end{aligned}$$

The first equality just substitutes the definition of V^*, the second uses Proposition 2, and the third treats the pointwise choice of a Bayes-plausible ψ' for each τ as a map $f \in \mathcal{F}$. Note that for each τ, $T_\star f(\tau) \in \Delta^2(\Omega)$. As a slight abuse of notation, treat $T_\star(\cdot)$ as a map from $\Delta^3(\Omega) \to \Delta^2(\Omega)$, so that $T_\star \circ f : \Delta^2(\Omega) \to \Delta^2(\Omega)$ for any $f \in \mathcal{F}$. Note that $T_\star(\cdot)$ is linear. Then we have

$$\begin{aligned}
&\max_{f \in \mathcal{F}} \int_{\Delta^2(\Omega)} \int_{\Delta^1(\Omega)} V d(T_\star f(\tau)) d\psi(\tau) = \max_{f \in \mathcal{F}} \int_{\Delta^1(\Omega)} V d \left(\int_{\Delta^2(\Omega)} T_\star f(\tau) d\psi(\tau) \right) \\
&= \max_{f \in \mathcal{F}} \int_{\Delta^1(\Omega)} V d \left(\int_{\Delta^2(\Omega)} T_\star \circ f d\psi \right) \\
&= \max_{f \in \mathcal{F}} \int_{\Delta^1(\Omega)} V d \left(T_\star \int_{\Delta^2(\Omega)} f d\psi \right) = \max_{\phi \succeq_{MPS} \psi} \int_{\Delta^1(\Omega)} V d(T_\star \phi).
\end{aligned}$$

The first equality uses Fubini's theorem, and the last leverages the fact that a choice of $f \in \mathcal{F}$ is equivalent to a choice of $\phi \succeq_{MPS} \psi$, as previously explained. Combining the two chains of equalities gives the desired result.

□

Theorem 1. *For any $\psi, \psi' \in \Delta^3(\Omega)$, $\psi \succsim_P \psi'$ if and only if $\{T_\star\phi | \phi \succeq_{MPS} \psi\} \supseteq \{T_\star\phi | \phi \succeq_{MPS} \psi'\}$.*

Proof. First, I prove the backwards implication, which is the easier direction. If $\{T_\star\phi | \phi \succeq_{MPS} \psi\} \supseteq \{T_\star\phi | \phi \succeq_{MPS} \psi'\}$, then from Lemma 1, we have

$$\int_{\Delta^2(\Omega)} V^* d\psi = \max_{\phi \succeq_{MPS} \psi} \int_{\Delta^1(\Omega)} V d(T_\star\phi) \geq \max_{\phi \succeq_{MPS} \psi'} \int_{\Delta^1(\Omega)} V d(T_\star\phi) = \int_{\Delta^2(\Omega)} V^* d\psi'.$$

for every $V \in \mathcal{V}$. By definition, the above inequality implies that $\psi \succsim_P \psi'$.

Now, I prove the forward implication. To show the contrapositive, assume that $\{T_\star\phi | \phi \succeq_{MPS} \psi\} \not\supseteq \{T_\star\phi | \phi \succeq_{MPS} \psi'\}$. Then fix $\hat{\phi} \succeq_{MPS} \psi'$ such that $T_\star\hat{\phi} \notin \{T_\star\phi | \phi \succeq_{MPS} \psi\}$, which must exist by assumption.

Note that $\{\phi | \phi \succeq_{MPS} \psi\}$ is a compact, convex subset of $\Delta^3(\Omega)$, and the map $T_\star(\cdot) : \Delta^3(\Omega) \to \Delta^2(\Omega)$ is affine and continuous. So, $\{T_\star\phi | \phi \succeq_{MPS} \psi\}$ is a compact, convex subset of $\Delta^2(\Omega)$. By the Hahn-Banach separation theorem, there exists a continuous, affine $w : \Delta^2(\Omega) \to \mathbb{R}$ such that $w(T_\star\hat{\phi}) > w(T_\star\phi)$ for every $\phi \succeq_{MPS} \psi$. By the Riesz Representation theorem, identify this w with some $V \in C(\Delta^1(\Omega))$, so that $\int_{\Delta^1(\Omega)} V d(T_\star\hat{\phi}) > \int_{\Delta^1(\Omega)} V d(T_\star\phi)$ for every $\phi \succeq_{MPS} \psi$. Therefore

$$\max_{\phi \succeq_{MPS} \psi'} \int_{\Delta^1(\Omega)} V d(T_\star\phi) \geq \int_{\Delta^1(\Omega)} V d(T_\star\hat{\phi}) > \max_{\phi \succeq_{MPS} \psi} \int_{\Delta^1(\Omega)} V d(T_\star\phi).$$

The strict inequality uses that the maximum is achieved and that this inequality holds for each $\phi \succeq_{MPS} \psi$. Since V is continuous and defined on the compact set $\Delta^1(\Omega)$, it is upper semicontinuous and bounded, so $V \in \mathcal{V}$. Then from Lemma 1, the above inequality gives $\psi \not\succsim_P \psi'$. The contrapositive is the desired implication.

□

Corollary 1. *Let $\psi, \psi' \in \Delta^3(\Omega)$ such that $\tau_S(\psi)$ and $\tau_S(\psi')$ are perfectly informative. Then $\psi \succsim_P \psi'$ if and only if $\tau_R(\psi) \preceq_{MPS} \tau_R(\psi')$.*

Proof. We have that $\tau_S(\psi) = \int_{\Delta^2(\Omega)} \tau d\psi(\tau)$ only supports degenerate beliefs δ_ω. Thus, if τ is supported by ψ, then τ only supports degenerate beliefs. In other words, any $\tau \in \text{supp}(\psi)$ is a perfectly informative distribution, though not one with mean μ_0 in general. Define $I \subseteq \Delta^2(\Omega)$ as the set of perfectly informative distributions τ. Define

$$X = \{\hat{\psi} \in \Delta^3(\Omega) : \text{supp}(\hat{\psi}) \subseteq I\}.$$

Note that $T|_I$ is a bijection between I and $\Delta^1(\Omega)$: for each belief μ, there is exactly one perfectly informative distribution with mean μ. Then $T_\star|_X$ is a bijection between X and $\Delta^2(\Omega)$. By assumption, we have $\psi, \psi' \in X$, and thus $\phi \in X$ if $\phi \succeq_{MPS} \psi$ or $\phi \succeq_{MPS} \psi'$. Then

$$\begin{aligned} \psi \succsim_P \psi' &\iff \{T_\star\phi | \phi \succeq_{MPS} \psi\} \supseteq \{T_\star\phi | \phi \succeq_{MPS} \psi'\} \\ &\iff \{\phi | \phi \succeq_{MPS} \psi\} \supseteq \{\phi | \phi \succeq_{MPS} \psi'\} \\ &\iff \psi \preceq_{MPS} \psi', \end{aligned}$$

using Theorem 1 for the first equivalence and that $T_\star|_X$ is a bijection for the second. Since $T_\star$ is affine, $T_\star|_X$ is an isomorphism between X and $\Delta^2(\Omega)$, so it preserves the mean-preserving spread order. Thus, again using $\psi, \psi' \in X$, we have

$$\psi \preceq_{MPS} \psi' \iff T_\star\psi \preceq_{MPS} T_\star\psi' \iff \tau_R(\psi) \preceq_{MPS} \tau_R(\psi'),$$

using the definition of τ_R for the last equivalence. Combining the two chains of equivalences gives the desired result.

□

Corollary 2. *Let $\psi, \psi' \in \Delta^3(\Omega)$. If $\psi \succeq_P \psi'$, then $\tau_S(\psi) \succeq_{MPS} \tau_S(\psi')$ and $\tau_R(\psi) \preceq_{MPS} \tau_R(\psi')$.*

Proof. For any $\phi \succeq_{MPS} \psi$, we have

$$\tau_R(\psi) = T_\star\psi \preceq_{MPS} T_\star\phi \preceq_{MPS} \int_{\Delta^2(\Omega)} \tau d\psi(\tau) = \tau_S(\psi).$$

The above just states that $\tau_S(\psi)$ is the most informative distribution Sender can induce under ψ and $\tau_R(\psi)$ is the least informative. The same holds for ψ'. If $\psi \succsim_P \psi'$, then $\{T_\star\phi | \phi \succeq_{MPS} \psi\} \supseteq \{T_\star\phi | \phi \succeq_{MPS} \psi'\}$ by Theorem 1. In particular, $T_\star\psi' = \tau_R(\psi') \in \{T_\star\phi | \phi \succeq_{MPS} \psi\}$. Then $\tau_R(\psi) \preceq_{MPS} \tau_R(\psi')$ because $\tau_R(\psi)$ is a contraction of every element of the choice set under ψ as shown before. Similarly, $\tau_S(\psi') \in \{T_\star\phi | \phi \succeq_{MPS} \psi\}$, so that $\tau_S(\psi) \succeq_{MPS} \tau_S(\psi')$ because $\tau_S(\psi)$ is a spread of every element of the choice set under ψ. This gives the desired result.

□

Proposition 6. *The problem* (P) *has a solution, and its value is $V^*(\tau_S)$. Further, V^* is superdifferentiable.*

Proof. The key is to note that this problem is equivalent to (P_B), which follows from the rewriting of (P_B) in the standard persuasion form: Kamenica and Gentzkow (2011) show the equivalence for Sender of choosing a Bayes-plausible distribution over beliefs, in this case ψ, and choosing an obedient signal, in this case σ. Thus, the value of (P) at τ_S is $V^*(\tau_s)$, the value of (P_B) at τ_S. Any solution τ to (P_B), which must exist by Proposition 2, then gives a solution to (P).

For superdifferentiability, the assumption that A is finite guarantees that V satisfies the regularity conditions from Dworczak and Martini (2019). Specifically, V is piecewise affine. Since I assume that $|\Omega| = 2$, the problem (P_B) is equivalent to the problem from their paper.

They prove that the dual to this problem has a solution, and Dworczak and Kolotilin (2024) show that this implies superdifferentiability of V^*. □

Theorem 2. *For any τ_S, there exists a solution w to* (D) *which represents* (P) *at τ_S.*

Proof. As stated before, the dual to the actions-based persuasion problem (P) is

$$\inf_{\lambda:A\to M(A)} \sup_{\sigma:\Delta^1(\Omega)\to\Delta(A)} \int_{\Delta^1(\Omega)} \left(\int_A v_a(\mu)d\sigma(a|\mu)\right) d\tau_S(\mu) + \int_{A\times A} \left(\int_{\Delta^1(\Omega)} [u_a(\mu) - u_{a'}(\mu)]\,\sigma(a|\mu)d\tau_S(\mu)\right) d\lambda(a'|a).$$

Note that (P) is a linear program in that both the objective and the obedience constraints are linear in σ and τ_S. Further, the value function for (P), V^*, is superdifferentiable from Proposition 6. Gretsky, Ostroy, and Zame (2001) prove that a linear program with superdifferentiable value function has dual attainment and no duality gap. Primal attainment was established in Proposition 6. Then we have strong duality for (P), so the value of the dual problem is $V^*(\tau_S)$, and there exists a solution (λ^*, σ^*) to the dual such that σ^* solves (P).

Now, we have to work with the expression for w^{λ^*}. Remember that for each a, ω, we have

$$w^{\lambda^*}(a,\omega) := v(a,\omega) + \int_A (u(a,\omega) - u(a',\omega))\, d\lambda^*(a'|a).$$

Then for any a, $w_a^{\lambda^*}$ is the linear functional on $\Delta^1(\Omega)$ defined as

$$\begin{aligned} w_a^{\lambda^*}(\mu) &= \int_\Omega \left[v(a,\omega) + \int_A (u(a,\omega) - u(a',\omega))\, d\lambda^*(a'|a)\right] d\mu(\omega) \\ &= v_a(\mu) + \int_A [u_a(\mu) - u_{a'}(\mu)]\, d\lambda^*(a'|a). \end{aligned}$$

A decision-maker with information τ_S facing decision problem w^{λ^*} then chooses $\sigma : \Delta^1(\Omega) \to \Delta(A)$ to maximize

$$\begin{aligned} &\int_{\Delta^1(\Omega)} \left(\int_A w_a^{\lambda^*}(\mu)d\sigma(a|\mu)\right) d\tau_S(\mu) \\ &= \int_{\Delta^1(\Omega)} \left(\int_A \left[v_a(\mu) + \int_A [u_a(\mu) - u_{a'}(\mu)]\, d\lambda^*(a'|a)\right] d\sigma(a|\mu)\right) d\tau_S(\mu) \\ &= \int_{\Delta^1(\Omega)} \left(\int_A v_a(\mu)d\sigma(a|\mu)\right) d\tau_S(\mu) + \int_{\Delta^1(\Omega)} \left(\int_A \int_A [u_a(\mu) - u_{a'}(\mu)]\, d\lambda^*(a'|a)d\sigma(a|\mu)\right) d\tau_S(\mu) \\ &= \int_{\Delta^1(\Omega)} \left(\int_A v_a(\mu)d\sigma(a|\mu)\right) d\tau_S(\mu) + \int_{A\times A} \left(\int_{\Delta^1(\Omega)} [u_a(\mu) - u_{a'}(\mu)]\,\sigma(a|\mu)d\tau_S(\mu)\right) d\lambda^*(a'|a). \end{aligned}$$

This is the same expression as we saw in the dual problem, just with a fixed choice of λ^*. Then since (σ^*, λ^*) solve the dual, we know σ^* is optimal in w^{λ^*} at τ_S, and the expected utility in the decision problem is the same as value of the dual, $V^*(\tau_S)$.

Remember that $\mathcal{W} = \{w^\lambda | \lambda : A \to M(A)\}$, and the problem (D) was defined as $\inf_{w \in \mathcal{W}} \widetilde{w}(\tau_S)$. So, w^{λ^*} solves this problem. Since τ_S was arbitrary, we know that for every τ,

$$V^*(\tau) = \min_{w \in \mathcal{W}} \widetilde{w}(\tau).$$

Since $w^{\lambda^*} \in \mathcal{W}$, this implies that $\widetilde{w^{\lambda^*}}$ lies above V^* everywhere, and this function is affine and continuous because $\widehat{w^{\lambda^*}}$ is continuous. Since $\widetilde{w^{\lambda^*}}(\tau_S) = V^*(\tau_S)$, this means $\widetilde{w^{\lambda^*}}$ is a supergradient to V^* at τ_S. Thus, we have a solution w^{λ^*} to (D) which satisfies all the conditions of a representative decision problem.

□

Corollary 3. *For any τ_S, there is a representative decision problem w such that $\widehat{w}$ is a price function for* (P) *at τ_S.*

Proof. Take a representative decision problem w. Then the map $\tau \mapsto \int_{\Delta^1(\Omega)} \widehat{w} d\tau$ is just $\widetilde{w}$, so in particular it is a supergradient to V^* at τ. The function $\widehat{w}$ is convex because it is the indirect utility of a decision problem. Thus, $\widehat{w}$ is a price function for (P) at τ_S. □

Corollary 4. *Let $\widetilde{\mathcal{W}} = \{\widetilde{w} : w \in \mathcal{W}\}$. Then*

$$V^* = \min \widetilde{W},$$

taking the minimum pointwise.

Proof. Theorem 2 establishes that the value of problem (D) at τ_S is $V^*(\tau_S)$, and the problem has a solution. Then $V^*(\tau_S) = \min_{w \in W} \widetilde{w}(\tau_S) = (\min \widetilde{W})(\tau_S)$. □