Persuading a Pessimist: Simplicity and Robustness

Afshin Nikzad*

Abstract

The signals used in persuasion mechanisms in practice typically satisfy two well-studied simple properties: (i) they partition an ordered state space into intervals, and (ii) they do not recommend lower actions at higher states. These properties have been studied—often separately—in the Bayesian persuasion literature, where conditions for the optimality of such signals are provided in various settings.

The two properties can be defined only when the action and the state space are ordered. Under the proper ordering conditions, we show that the optimal signal features both of these properties, as well as robustness properties, when Receiver is a *pessimist*. A pessimistic receiver, rather than maximizing expected payoff, takes the action that guarantees the highest level of payoff. Through the notion of *maxmin expected utility*, our findings explain that simplicity and robustness of optimal signals can emerge from the ambiguity of the prior to Receiver.

*Department of Economics, University of Southern California (USC). Email: afshin.nikzad@usc.edu

1 Introduction

Bayesian persuasion has become a standard framework for modeling information transmission in economics: Sender and Receiver share the same prior belief about the state of the world. Before the state of the world is realized, Sender constructs a *signal*, i.e., a distribution over *signal realizations* at each possible state of the world. After the state of the world is realized, a signal realization is relayed to Receiver. Receiver then takes action to maximize her expected utility, which depends on her action and the state of the world, with the expectation taken with respect to the posterior that she forms after receiving the signal realization.

Most persuasion mechanisms in practice have a simple structure. In particular, many of the signals in practice have an "interval structure", i.e., they partition an ordered state space into intervals. The interval structure of optimal signals has been studied in the Bayesian persuasion literature. For instance, Ivanov (2015), Kolotilin (2017), Dworczak and Martini (2018) provide sufficient and, in some cases, necessary conditions for the optimality of such signals in different settings. Another structural property that is observed in practice and has been studied separately in the literature is "monotonicity", which means that the optimal signal does not recommend lower actions at higher states when both the action and state spaces are ordered (Mensch 2018).

To define the above structural properties—the interval structure and monotonicity—the action and the state space should be ordered. Under the proper ordering conditions, we show the optimal signal features both properties when Receiver is a *pessimist*. This notion of pessimism, defined more precisely next, is remarkably simple compared to the existing sufficient conditions for either of these properties.[1]

We say Receiver is a pessimist when, rather than maximizing expected utility, she evaluates each action by the least preferred outcome that it can promise and takes the most promising action. This notion of pessimism can also be interpreted through the *maxmin expected utility* model where Receiver is an expected utility maximizer who does not know the prior.[2] Hence, our findings explain how the ambiguity of the prior can lead to the simplicity

[1]The conditions in prior work often take the form of technical conditions involving payoff functions, as reviewed in Section 5. There, we also discuss that prior work does not imply our findings.

[2]Such maxmin utility specifications are used in many applications to model decision-making when probabilities are ambiguous and not objectively known. For example, they have been used to model agents' participation in stock markets based on available information about risky assets (Antoniou et al. 2015), patients' decisions in choosing treatment options or participating in medical trials (Attema et al. 2018), and investment decisions in real options (Lourens 2013). (See Machina and Siniscalchi (2014) for a survey.)

of optimal signals.

We establish our findings through the lens of an ordinal utility model. Receiver has a preference order over all *outcomes*, i.e., pairs of actions and states. After receiving the signal realization, she evaluates any action by the least preferred outcome that it can promise and takes the most promising action. Sender has a preference order over the outcomes and evaluates signals by the probability distribution that they impose on the realized outcome. She chooses a *stochastically dominant* signal when it exists, i.e., a signal that imposes a distribution over outcomes which stochastically dominates the distribution imposed by *any* other signal (where the stochastic dominance relation is defined in the usual way with respect to Sender's preference order over the outcomes).

We ask whether stochastically dominant signals satisfy the structural properties that were discussed earlier—the interval structure and monotonicity of a signal. To define these properties (let alone show that they are satisfied), the action and the state space need to be ordered. Formally, we refer to the ordering conditions as *action-monotonicity* and *state-monotonicity*. Action-monotonicity means that there exists an ordering on the set of actions such that Sender prefers any higher action in that order to any lower action, at any state. State-monotonicity means that there exists an ordering on the set of states such that Receiver prefers any higher state in that order to any lower state, when taking any action.

The two structural properties are well-defined if the action- and state-monotonicity conditions hold. We find that, then, stochastically dominant signals exist and satisfy both of the structural properties (Section 3). In addition, we show the robustness of these signals, by establishing properties such as *independence from the prior* (which means the signal does not depend on the full support prior).

In "many environments that fit into the persuasion framework, the action space and state space are ordered" (Mensch 2018), and the monotonicity conditions hold. For example, consider a pharmaceutical company persuading customers to buy its product by revealing information about its efficacy, as in Kamenica and Gentzkow (2017), or a principal persuading an agent to exert effort by providing information about the project reward (Dworczak and Martini 2018). Nevertheless, in some environments, such as in multidimensional persuasion, Receiver's preference order over the states is not complete. In the online appendix we identify conditions under which our main findings hold in such environments as well.[3]

Section 4 shows that, in our framework, some of the common signals in the literature—

[3]The condition required for dismissing state-monotonicity is fairly strong, but the one required for dismissing action-monotonicity is much weaker. These conditions are discussed in Online Appendix i and ii, respectively.

namely, fully informative, non-informative, upper censorship, and lower censorship signals—can be fully characterized. This characterization also demonstrates that a *higher alignment of Sender's and Receiver's preferences* may lead to a larger stochastically dominant signal. For example, a fully informative signal (which has the largest number of signal realizations) arises when Sender's and Receiver's preferences are *fully aligned*, whereas a non-informative signal (which has the smallest number of signal realizations) arises when Sender's and Receiver's preferences are *fully misaligned*.

Finally, we remark that the notion of pessimism that we consider can be rationalized through the *maxmin expected utility* model, where Receiver is an expected utility maximizer who does not know the prior. Hence, our findings provide a foundation for the simplicity and robustness of optimal persuasion mechanisms, by showing that these properties can emerge from the ambiguity of the prior to Receiver.

2 Example: A principal-agent setup

An agent (he) must be persuaded to exert costly effort to complete a project. The project, if completed successfully, has value v to the principal (she). The agent receives $w(v)$ of the project value in case of success, where $w : \mathbb{R}_+ \to \mathbb{R}_+$.

The agent can exert effort $e \in E$ to complete the project. We assume that E is a finite set of nonnegative real numbers.[4] If the project has value v and the agent exerts effort e, he incurs a cost $c(e)$, and the probability of success of the project would be $p(e, v)$. Then, the (ex ante) payoff of the agent from exerting effort e is defined by $u_a(e, v) = p(e, v)w(v) - c(e)$, and the (ex ante) payoff of the principal is defined by $u_p(e, v) = p(e, v)(v - w(v))$.

The value v is distributed according to a *prior distribution* μ, which is common knowledge. We assume that μ has full support on a finite set V.[5] The principal discloses information about the project value using a *signal*. A signal π consists of a finite realization space S and a family of distributions $\{\pi(\cdot|v)\}_{v \in V}$ over S. When the project value v is realized, the principal draws a signal realization from $\pi(\cdot|v)$ and sends it to the agent. The agent interprets a received signal realization s *pessimistically*: he evaluates any effort level e by the

[4] This assumption is merely for expositional simplicity; Online Appendix iii repeats the same exercise for when E is a continuum.

[5] The assumption that V is finite is for expositional simplicity; Online Appendix iii repeats the same exercise for when V is an interval.

worst-case payoff that it can promise, i.e.,

$$\min_{v: s \in \operatorname{supp}(\pi(\cdot|v))} u_a(e, v), \tag{2.1}$$

where $\operatorname{supp}(\pi(\cdot|v))$ denotes the support of the distribution $\pi(\cdot|v)$. The agent chooses an effort level $\hat{e}(s)$ that maximizes his worst-case payoff, i.e.,

$$\hat{e}(s) \in \arg\max_{e \in E} \min_{v: s \in \operatorname{supp}(\pi(\cdot|v))} u_a(e, v). \tag{2.2}$$

Remark 2.1. *If instead of (2.1) the agent evaluates an effort level e by $\mathbb{E}_{v \sim \mu_s}[u_a(e, v)]$, where μ_s is the agent's posterior for the project value after receiving a signal realization s, then he would be an expected utility maximizer and the above model would essentially be the same as the Bayesian persuasion model introduced by Kamenica and Gentzkow (2011). We take a different approach and assume (2.1) for a pessimistic agent. Evaluating actions by their worst-case payoff (2.1) can also be rationalized through the* maxmin expected utility *model where the agent is an expected utility maximizer who does not know the prior: the agent evaluates an action e by $\inf_\mu \mathbb{E}_{v \sim \mu_s}[u_a(e, v)]$, where the infimum is over all possible prior distributions μ and μ_s is the associated posterior after receiving a signal realization s.*

The *value* of a signal π for the principal is $\mathbb{E}_{v \sim \pi} \mathbb{E}_{s \sim \pi(\cdot|v)}[u_p(\hat{e}(s), v)]$. A signal π is *optimal* if no other signal has a higher value. The principal's objective is finding the optimal signal. To solve the principal's problem, we make the following assumptions.

Assumption 2.2 (*State-monotonicity*)**.** *The agent's payoff is increasing in the project value; i.e., the function $u_a(e, v)$ is increasing in v.*

Assumption 2.3 (*Action-monotonicity*)**.** *The principal's payoff is increasing in the agent's exerted effort level; i.e., the function $u_p(e, v)$ is increasing in e.*

For notational simplicity throughout this example, we assume that for every $v \in V$ there is a *unique* effort level, denoted by $e^*(v)$, that maximizes the agent's payoff when the project has value v. (It is well known that in such principal-agent models e^* may be a non-monotone function.[6]) We will construct the optimal signal using the function e^*. First, we describe the construction, and then we show that the constructed signal is optimal.

[6] That said, one can characterize conditions under which $e^*(v)$ is an increasing function. For example, when $p(e, v)$ depends only on e (i.e., $p(e, v) \equiv q(e)$ for some function q) then state-monotonicity implies that $w(v)$ is increasing. This, together with $p(e, v) \equiv q(e)$, implies that $u_a(e, v)$ is supermodular. Therefore, $e^*(v)$ is increasing by the Topkis theorem.

We say that project value $v \in V$ *induces* an effort level l if $e^*(v) = l$. The signal is constructed by an iterative algorithm. The algorithm starts at iteration $i = 1$, and at every iteration it repeats the following:

i. Choose a project value $v \in V$ that induces the highest possible effort level; i.e., $v = \arg\max_{v' \in V} e^*(v')$.
ii. Let $T_i \subseteq V$ be the subset of the project values that are no smaller than v.
iii. Remove from V all of the project values belonging to T_i; i.e., $V \leftarrow V \backslash T_i$.
iv. If V is nonempty, then increase i by one and start the next iteration; otherwise, stop.

When the algorithm stops, the signal π^* is constructed by *pooling together* the subset of the project values that belong to T_i, for every i. Formally, the signal π^* has a signal realization space $S^* = \{T_1, \ldots, T_i\}$, and sends a signal realization $T \in S^*$ to the agent when the true project value v belongs to T. Before proving that π^* is indeed the optimal signal, we illustrate its construction by an example.

Figure 1 shows how the algorithm works using an example where $V = \{1, \ldots, 12\}$. At the first iteration, $v = \arg\max_{v' \in V} e^*(v') = 12$. Thus, the algorithm sets $T_1 = \{12\}$ and removes the value 12 from V. At the beginning of the second iteration, $V = \{1, \ldots, 11\}$. Hence, the project value that induces the highest effort level is $v = \arg\max_{v' \in V} e^*(v') = 8$. The algorithm then sets $T_2 = \{8, \ldots, 11\}$ and removes the elements of T_2 from V. At the beginning of the third iteration, $V = \{1, \ldots, 7\}$. The project value that induces the highest effort level is 5, and thus $T_3 = \{5, 6, 7\}$. In the next four iterations, the algorithm sets $T_4 = \{4\}$, $T_5 = \{3\}$, $T_6 = \{2\}$, and $T_7 = \{1\}$. Hence, the constructed signal fully reveals the project value if it is below 5 or above 11, and *pools* the project values over the interval $[5, 7]$, as well as the interval $[8, 11]$, by revealing to the agent only the interval containing the value.

Remark 2.4. *One interpretation of the algorithm is that it constructs the* smallest *increasing function that is point-wise larger than* e^**; this function is the one that assigns the value* $\max_{v' \in T_i} e^*(v')$ *to every element in* T_i*, for every* i*. Appendix C discusses the details.*

The following definition is helpful in analyzing the algorithm and showing the optimality of π^*. For every $P \subseteq V$ we define $e^*(P) = e^*(\min_{v \in P} v)$ with a slight abuse of notation. By the agent's pessimism (2.2), $e^*(P)$ is the effort level exerted by the agent when he knows (only) that the project value belongs to the subset P. To prove the optimality of π^*, we first note the following properties:

Property A: For every $T_i \in S^*$ and $t \in T_i$, $e^*(T_i) \geq e^*(t)$.

Property B: For every $T_j, T_k \in S^*$ with $j \leq k$, $e^*(T_j) \geq e^*(T_k)$.

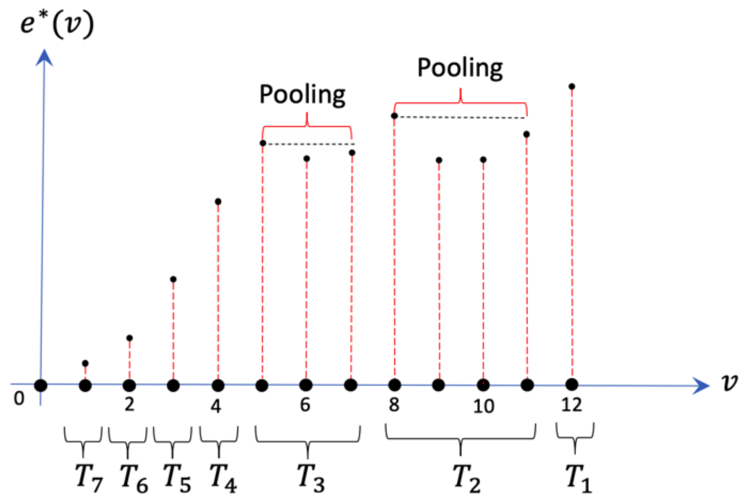

Figure 1: The set $V = \{1, \ldots, 12\}$ is illustrated by equidistant dots on the horizontal axis. The point $(v, e^*(v))$ is marked on the graph for every $v \in V$.

To see why Property A holds, we observe that in step ii of the algorithm, T_i is set to include a value v and all of the higher project values in V. Step i of the algorithm chose $v \in V$ to be the value that induces the highest possible effort. Thus, the smallest element of T_i (i.e., v) induces the highest effort level among all of its elements. Property B holds for a similar reason: When the next iteration (i.e., iteration $i+1$) starts, every remaining element in V induces a lower effort level than $e^*(T_i)$. Thus, $e^*(T_i) \geq e^*(T_{i+1})$.

We next use these properties to prove the optimality of π^*.

Proposition 2.5. *The signal π^* constructed by the algorithm is optimal.*

Proof. Suppose that the project value is v and that $v \in T_j$ for some j. If the principal uses the signal π^*, the agent exerts an effort level $e^*(T_j)$. We will show that the highest possible effort level that the agent may exert under *any* signal is $e^*(T_j)$ when the project value is v. This would prove the optimality of π^*, since the latter claim holds for arbitrary $v \in V$.

Consider an arbitrary signal ρ with a signal realization space S. For an arbitrary signal realization $s \in S$, let $V_s \subseteq V$ denote the set of values $v' \in V$ such that $\rho(s|v') > 0$ (i.e., ρ may send s with a positive probability if the project has value v'). Suppose that the principal uses the signal ρ and that the agent receives a signal realization s. Then,

$$e^*(V_s) = e^*(\min\{v'|v' \in V_s\}) \leq \max\{e^*(v'')|v'' \in V, v'' \leq v\}), \tag{2.3}$$

where the equality holds by the agent's pessimism and the inequality holds since $v \in V_s$.

On the other hand, from Properties A,B it follows that

$$e^*(T_j) = \max\{e^*(v')|v' \in V, v' \leq v\}.$$

The above equation together with (2.3) implies that $e^*(V_s) \leq e^*(T_j)$, which is the promised claim: the effort that the agent exerts under the (arbitrary) signal ρ is no larger than the effort he exerts under the signal π^*. □

We next briefly discuss some of the structural and robustness properties of π^*. The above constructive approach reveals that π^*

1. **pools adjacent project values:** if $v', v'' \in V$ belong to the same signal realization in π^*, then any $v''' \in V$ with $v' < v''' < v''$ also belongs to that signal realization,

2. **never induces a lower effort level at a higher project value:** when the principal uses π^*, then the agent always exerts a weakly higher effort level at a higher project value, and

3. **is independent of the prior:** the signal π^* is constructed independently of the prior μ, and hence π^* remains optimal if μ is replaced with any other full support prior.

Property 1 above holds because, at every iteration i of the algorithm, the signal realization T_i contains a value v and all of the higher values that belong to V in that iteration. Property 2 follows from the fact discussed earlier that for every two signal realizations T_j, T_k with $j \leq k$, $e^*(T_j) \geq e^*(T_k)$. Property 3 holds simply because the algorithm does not use the prior to construct the optimal signal.

The principal-agent setup that we defined here is based on a *cardinal* preference model where the agent's and the principal's preferences are represented by utility functions. In the next section we present our main setup, which extends the above model to an ordinal preference model (where *Receiver* and *Sender* have preference orders instead of utility functions). The algorithm that we use there to construct the optimal signal will be quite similar to the algorithm in this section. Briefly put, the idea behind this extension is that the algorithm in this section can construct the optimal signal by solely performing (ordinal) comparisons. Formalizing this idea requires extending the notion of signal optimality to an ordinal setting. After doing that, we will present natural counterparts to the above findings in the ordinal setup.

The analysis in the next section formally shows that optimal persuasion does not require knowledge about *preference intensities* (cardinal utilities) in our setup, unlike the existing models in the persuasion literature that often rely on preference intensities to construct optimal signals.

3 The general setup

We first define the following notions regarding preference orders. For a transitive preference order $\preceq$ over a set X, we use $\prec$ and $\sim$ respectively to denote the strict and the equivalence preference orders imposed by $\preceq$ over X. For any two elements x, y such that $x \preceq y$, we say that y is $\preceq$-higher than x, and that x is $\preceq$-lower than y. Let $\inf^{\preceq}\{X\}$ denote $x \in X$ such that $x \preceq y$ holds for all $y \in X$; if there is more than one such x, choose one arbitrarily. We say a subset $P \subseteq X$ is *convex* with respect to $\preceq$ if for any $\omega, \omega', \omega'' \in X$ such that $\omega \preceq \omega' \preceq \omega''$, having $\omega, \omega'' \in P$ implies that $\omega' \in P$.

There exists a finite set of states, Ω, and a finite set of actions, A. Any element of $A \times \Omega$ is called an *outcome*. Sender (she) and Receiver (he) have complete and transitive preference orders over the set of outcomes, respectively denoted by $\preceq_s, \preceq_r$.[7]

The state of the world, ω_0, is drawn from a prior $\mu \in \text{int}(\Delta(\Omega))$,[8] henceforth *the prior*. Sender discloses information about the state of the world to Receiver using a *signal*. A signal π consists of a finite realization space S and a family of distributions over S, $\{\pi(\cdot|\omega)\}_{\omega \in \Omega}$. When the state of the world is realized, Sender draws a signal realization s from $\pi(\cdot|\omega)$ and sends it to Receiver. Receiver interprets the signal realization s *pessimistically*: he evaluates any action by the *worst-case outcome that it can promise*, i.e.,

$$\inf^{\preceq_r}\{(a, \omega) : s \in \text{supp}(\pi(\cdot|\omega))\}, \tag{3.1}$$

where $\text{supp}(\pi(\cdot|\omega))$ denotes the support of $\pi(\cdot|\omega)$. We denote this outcome by $\underline{o}(a, s)$. For brevity, we also call $\underline{o}(a, s)$ the *promised outcome* of action a when Receiver has received a signal realization s.

Conditional on receiving a signal realization s, Receiver chooses an action with the *most preferred* promised outcome, i.e., an action $\hat{a}(s)$ such that

$$\underline{o}(a, s) \preceq_r \underline{o}(\hat{a}(s), s) \tag{3.2}$$

[7] In the principal-agent example of Section 2, the set of states (project values) is denoted by V, the set of actions (effort levels) is denoted by E, and $\preceq_s$ and $\preceq_r$ respectively correspond to the functions u_p and u_a.

[8] The assumption that the prior is an interior point of the probability simplex $\Delta(\Omega)$ is without loss of generality, as both Sender and Receiver would ignore the states outside of the prior's support.

holds for every action $a \in A$.

Given that Receiver takes action as above, any signal π chosen by Sender *imposes* a distribution D_π over the set of outcomes: $D_\pi(x)$ is the probability mass that outcome x is realized when the chosen signal is π. We say a signal π *stochastically dominates* another signal ρ if, for every outcome x,

$$\sum_{x \preceq_s y} D_\pi(y) \geq \sum_{x \preceq_s y} D_\rho(y). \tag{3.3}$$

A signal is *stochastically dominant* if it stochastically dominates every other signal. Sender's objective is choosing a signal π^* that is stochastically dominant.

We say Receiver's preferences are *monotone in states* when there exists a total order $\preceq_\Omega$ over the states such that whenever $\omega \preceq_\Omega \omega'$, then $(\omega, a) \preceq_r (\omega', a)$ holds for any $a \in A$. We say Sender's preferences are *monotone in actions* when there exists a total order $\preceq_A$ over actions such that whenever $a \preceq_A a'$, then $(\omega, a) \preceq_s (\omega, a')$ holds for any $\omega \in \Omega$. For brevity, we sometimes refer to the former condition as *state-monotonicity* and the latter condition as *action-monotonicity*. The two conditions are respectively counterparts to Assumptions 2.2 and 2.3 in the principal-agent model of Section 2. We say that Sender and Receiver have *monotone preferences* when both of these conditions hold.

A signal is *monotone* if it recommends $\preceq_A$-higher actions at $\preceq_\Omega$-higher states.

Theorem 3.1. *When Sender and Receiver have monotone preferences, there exists a stochastically dominant signal π^*. Moreover, π^* satisfies the following properties:*

i. The signal realization space in π^* is a family of subsets of Ω; furthermore, these subsets partition Ω. When the realized state of the world is $\omega \in \Omega$, π^* sends to Receiver the signal realization that contains ω.

ii. π^* is monotone. Also, every signal realization in π^* is convex with respect to $\preceq_\Omega$.

iii. π^* is independent of the prior μ, i.e., π^* remains stochastically dominant if μ is replaced with any other full-support prior.

The proof, presented in Appendix A.1, is similar to the analysis of the principal-agent setup in Section 2. In particular, we construct the signal π^* using a similar algorithm. Recall that, in the proof of Proposition 2.5, we established the optimality of the signal π^* at *each* realized state of the world. It is essentially this property that implies π^* is stochastically dominant here. The structural and robustness properties of π^* follow from the construction.

This analysis shows that optimal persuasion in our setup does not require knowledge about preference intensities, unlike most of the existing models in the persuasion literature.

We complement the above theorem by showing that the general existence of stochastically dominant signals depends on both the state- and the action-monotonicity assumptions.

Proposition 3.2. *Stochastically dominant signals do not generally exist if either the state- or action-monotonicity assumption does not hold.*

The proof provides one example for when state-monotonicity does not hold and one example for when action-monotonicity does not hold (Appendix A.2). In both examples we represent Receiver's and Sender's preferences by utility functions (as in the principal-agent setup of Section 2) and then we observe that optimal signals do not exist.

4 Common signals and the signal size

We next use our framework to characterize four of the common signals in the literature, namely, *fully informative*, *non-informative*, *upper censorship*, and *lower censorship* signals. A fully informative signal always reveals the state of the world to Receiver. A non-informative signal always sends the same signal realization regardless of the state of the world. An upper censorship signal picks a *threshold state* $\omega^* \in \Omega$, fully reveals ω_0 (the state of the world) if $\omega_0 \preceq_\Omega \omega^*$, and otherwise reveals only that $\omega^* \preceq_\Omega \omega_0$ holds. A lower censorship signal picks a *threshold state* $\omega^* \in \Omega$, fully reveals the state of the world if $\omega^* \preceq_\Omega \omega_0$, and otherwise reveals only that $\omega_0 \preceq_\Omega \omega^*$ holds.

Throughout this section, without loss of generality, we assume that $\Omega = \{\omega_1, \ldots, \omega_n\}$ such that $\omega_1 \preceq_\Omega \ldots \preceq_\Omega \omega_n$. For expositional simplicity, we assume that Sender and Receiver have strict preference orders over the outcomes.[9] For every $\omega \in \Omega$, let $a^*(\omega)$ denote the action chosen by Receiver when (he knows that) the state of the world is ω. We say that a^* is *increasing over* $\Omega' \subseteq \Omega$ when, for every $\omega, \omega' \in \Omega'$, $a^*(\omega) \preceq_A a^*(\omega')$ if $\omega \preceq_\Omega \omega'$. We say that a^* is *maximized over* $\Omega' \subseteq \Omega$ at $\omega' \in \Omega'$ when, for every $\omega \in \Omega'$, $a^*(\omega) \preceq_A a^*(\omega')$.

Proposition 4.1. *When Sender and Receiver have monotone preferences, there exists a stochastically dominant signal that is*

i. fully informative if and only if a^* is increasing over Ω;

ii. non-informative if and only if a^* is maximized over Ω at ω_1;

[9] This assumption can be dismissed without significant change in the proof.

iii. upper censorship if and only if there exists $\omega_t \in \Omega$ such that a^* is maximized over Ω at ω_t and is increasing over $\{\omega_1, \ldots, \omega_t\}$;

iv. lower censorship if and only if there exists $\omega_t \in \Omega$ such that a^* is maximized over $\{\omega_1, \ldots, \omega_t\}$ at ω_1 and and is increasing over $\{\omega_t, \ldots, \omega_n\}$.

For example, consider the principal-agent example of Section 2, where the state space is the set of project values V. Then, the sufficient and necessary condition for the optimality of upper censorship is the existence of a v^* such that $e^*(v)$ is maximized over V at $v = v^*$ and $e^*(v) \leq e^*(v')$ holds for every $v, v' \in V$ with $v < v' \leq v^*$.

Proposition 4.1 also characterizes two extreme points concerning the size (i.e., number of signal realizations) of stochastically dominant signals. At one extreme, we have stochastically dominant signals with size $|\Omega|$; this happens when Sender's and Receiver's preferences are *fully aligned*, in the sense that a^* is increasing over Ω. At the other extreme, we have stochastically dominant signals with size 1. This happens when Sender's and Receiver's preferences are *fully misaligned*, in the sense that a^* is maximized over Ω at a_1.

5 Related literature

In terms of structural properties of optimal signals, Ivanov (2015), Kolotilin (2017), Dworczak and Martini (2018), Mensch (2018) are perhaps among the closest to our work from the large literature on Bayesian persuasion. They provide sufficient, and in some cases necessary, conditions for the sender's optimal signal to have an "interval structure" in different settings. (Mensch (2018) focuses on providing sufficient conditions for monotonicity of a signal, a necessary condition for which turns out to be the "interval structure".) We discuss these works, among others, below. These findings do not imply ours, as we will see. One common difference is that we do not consider Receiver an expected utility maximizer.

Dworczak and Martini (2018) study the case where the sender's ex post utility depends only on the mean of posterior beliefs. They provide a sufficient and, in a sense, necessary condition for the optimal signal to have an interval structure. Their condition is in terms of the shape of the sender's utility function, u: the function u needs to be "affine-closed", which, loosely speaking, means that $u + q$ has at most one local interior maximum for any affine function q.[10]

[10]This is a consequence of their general analysis of the case where the sender's utility depends only on the mean of posterior beliefs.

Ivanov (2015) studies a setting where the action of an expected utility maximizing receiver depends on the mean of posterior beliefs and the order of this mean in the sequence of possible means. He provides a sufficient condition for the optimal signal to have an interval structure. The condition requires the sender's interim payoff to be linear in posterior means and actions.

Kolotilin (2017) considers a model where an expected utility maximizing receiver and sender have private types, and the receiver takes a binary decision to "act" or not to. To study the interval structure, it is assumed that the sender's utility depends only on the receiver's type, and the receiver's utility is the difference between the sender's and the receiver's types. Taking a linear programming approach, he provides a sufficient condition for the optimality of "interval revelation schemes", a type of revelation scheme that pools all the "low" states together, fully reveals all the "medium" states, and pools all the "high" states together. Loosely speaking, the condition requires the receiver's interim utility function to be convex in certain intervals and requires a first-oder approximation of the receiver's interim utility function to bound the utility function from above in another interval.

Mensch (2018) considers an expected utility maximizing receiver and defines the notion of monotone signals as signals that do not recommend lower actions at higher states. He shows that such signals should partition the state space into intervals. Building on the literature on monotone comparative statics, he provides sufficient conditions for a signal to be monotone. In particular, he shows that supermodularity of the sender's and the receiver's utility functions is sufficient when the state space is binary. He also considers the case of a continuum of actions. For that case, the sufficient condition is "quasi-supermodularity" of a certain function: the Gateaux derivative of the sender's payoff in the direction of the change of the conditional distribution.

The interval structure of messages also appears in other models of communication, e.g., in the cheap talk game of Crawford and Sobel (1982).[11]

Rappoport (2020) studies a communication game played between a sender (she) and an expected utility maximizing receiver (he), where the sender has a type known only to herself. The sender can send a message to the receiver about her type: in this message, she can pretend to be of any type weakly lower[12] than her actual type. After receiving the message, the receiver takes an action. His payoff depends on both his own action and the

[11] Interval structures appear in other literature as well. For example, Dubey and Geanakoplos (2017) consider a firm whose objective is incentivizing workers to work by designing a *reward schedule*. By definition, a reward schedule partitions a worker's output space into consecutive intervals and assigns a wage and a title to each interval, such that the wages and titles cannot go down as the intervals go up. They characterize cost-minimizing reward schedules under various conditions.

[12] In fact, *lower* can be defined generally according to a fixed partial order.

sender's type, whereas the sender's payoff depends only on the receiver's action. In a Perfect Bayesian Equilibrium of this game, a sender of type t pretends to be of a weakly lower type t' that gives her the highest payoff, taking the receiver's strategy as given. Thus all types between t' and t would be pooled together. This leads to the sender's possible types being partitioned into intervals in any equilibrium of the game. In our model, however, the interval structure is generically unique and independent of the prior, and it arises due to the receiver's exogenous pessimism.[13] Among the possibly multiple equilibria, Rappoport (2020) focuses on the *receiver-optimal* equilibrium and characterizes when a change in the prior distribution of *evidence* induces more *skepticism* in this equilibrium.

Milgrom (2008) studies a persuasion game between a seller who provides information about the quality of a product and an expected utility maximizing buyer. He shows that, at equilibrium, *skepticism* on the part of the buyer arises endogenously. That is, when the actual quality of the good is q, all equilibrium reports by the seller lead to the same outcome, in the sense that any equilibrium report specifies that the product quality is exactly q or that it belongs to a subset of possible qualities with their minimum being q; also, the buyer believes that the actual quality of the good is equal to the minimum quality in the reported subset. This is due to an *unraveling* argument: the highest-quality sellers always make reports of quality that distinguish their products from all others, and then the remaining sellers face a similar game, resulting in the product quality being (effectively) fully revealed at equilibrium. In our model, however, the receiver is exogenously pessimistic, and the optimal signal may be not fully revealing.[14,15,16]

Unlike the above settings and relevant to ours, Beauchêne et al. (2019) study ambiguity in persuasion games. In their setting, players are ambiguity-averse with maxmin expected

[13] Due to the receiver's pessimism, from the support of any signal realization, only the lowest state matters to him. Thus, if the sender pools a state ω with higher states, she would not change the receiver's action compared to when the state ω is fully revealed to him. Hence, the sender would pool a state ω with higher states if she prefers the receiver's optimal action at ω to his optimal action at each higher state. This leads to a generically unique interval structure in the optimal signal.

[14] A notable difference between the unraveling argument and our analysis is that our analysis orders the outcomes according to the receiver's preferences (e.g., project values in Figure 1), whereas the unraveling argument orders the sender's types and works from the highest-type to the lowest-type sender.

[15] There are other aspects that differentiate between the two models. For example, uncertainty about the state of the world can mute skepticism in the persuasion game model, because then the seller may sometimes benefit from pretending to be unaware of the true state, whereas in our model the optimal signal does not change under such uncertainty, due to the independence from the prior property.

[16] One can alter the persuasion game of Milgrom (2008) by considering a buyer who has *maxmin expected utility* preferences. The resulting game has multiple equilibria. Full information revelation does not generally correspond to an equilibrium of the game, whereas no information revelation and the signal π^* do (Appendix D).

utility. Even though there is no prior ambiguity, the sender may still choose to use ambiguous communication devices. They characterize the value of optimal ambiguous persuasion, and find it to be often higher than what is feasible under Bayesian persuasion. Their analysis hence provides some justification for how ambiguity may emerge endogenously in persuasion games.

There has also been some work on the robustness of the optimal signals in persuasion problems (e.g., robustness to lack of commitment power or information). Best and Quigley (2017) address the problem of the lack of commitment power in a cheap talk game with a long-lived sender and short-lived receivers. They show that optimal persuasion can be attained by altering the game. Hu and Weng (2018) consider an ambiguity-averse sender with limited knowledge about the receiver's private information and a max-min expected utility function. They show that when the sender faces full ambiguity, full disclosure is optimal, and when she faces vanishing ambiguity, she can do almost as well as when receiver has no private information.

References

Constantinos Antoniou, Richard D.F. Harris, and Ruogu Zhang. Ambiguity aversion and stock market participation: An empirical analysis. *Journal of Banking & Finance*, 58:57 – 70, 2015. ISSN 0378-4266. 1

Arthur E. Attema, Han Bleichrodt, and Olivier L'Haridon. Ambiguity preferences for health. *Health Economics*, 2018. 1

Dorian Beauchêne, Jian Li, and Ming Li. Ambiguous persuasion. *Journal of Economic Theory*, 179:312 – 365, 2019. ISSN 0022-0531. doi: https://doi.org/10.1016/j.jet.2018.10.008. URL http://www.sciencedirect.com/science/article/pii/S0022053118306653. 13

James Best and Daniel Quigley. Persuasion for the long run. *Working paper*, 2017. 14

Harry Campbell. What it's like to be part of lyft's new prime time experiment, 2018. URL https://therideshareguy.com/what-its-like-to-be-part-of-lyfts-new-prime-time-experiment. 36

Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 1982. 12

Pradeep K Dubey and John Geanakoplos. Money and status: How best to incentivize work. 2017. 12

Piotr Dworczak and Giorgio Martini. The simple economics of optimal persuasion. *Journal of Political Economy, forthcoming*, 2018. 1, 2, 11

Ju Hu and Xi Weng. Robust persuasion of a privately informed receiver. *Working paper*, 2018. 14

Maxim Ivanov. Optimal signals in bayesian persuasion mechanisms. *Working paper*, 2015. 1, 11

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, October 2011. doi: 10.1257/aer.101.6.2590. URL http://www.aeaweb.org/articles?id=10.1257/aer.101.6.2590. 4

Emir Kamenica and Matthew Gentzkow. Competition in persuasion. *Review of Economic Studies*, 2017. 2

Anton Kolotilin. Optimal information disclosure: A linear programming approach. *Theoretical Economics*, 2017. 1, 11, 12

Johan Lourens. Real options and ambiguity aversion: a theoretical approach. *ERASMUS UNIVERSITY Thesis*, 2013. 1

Mark Machina and Marciano Siniscalchi. Ambiguity and ambiguity aversion. *Handbook of the Economics of Risk and Uncertainty*, 2014. 1

Jeffrey Mensch. Monotone persuasion. *working paper*, 2018. 1, 2, 11, 12

Paul Milgrom. What the seller won't tell you: Persuasion and disclosure in markets. *Journal of Economic Perspectives*, 22(2):115–131, June 2008. doi: 10.1257/jep.22.2.115. URL http://www.aeaweb.org/articles?id=10.1257/jep.22.2.115. 13, 25, 27

Boris Mityagin. The zero set of a real analytic function. 2015. URL https://arxiv.org/abs/1512.07276. 36

Daniel Rappoport. Evidence and skepticism in verifiable disclosure games. *Available at SSRN 2978288*, 2020. 12, 13

A Proofs for Section 3

A.1 Proof of Theorem 3.1

The proof is quite similar to the proof of optimality that we provided in the principal-agent example of Section 2. Before presenting the formal proof, we define a few notions.

For every $\omega \in \Omega$, let $a^*(\omega)$ denote the action chosen by Receiver when he knows that the state of the world is ω. We say that a state ω *induces* an action a if $a^*(\omega) = a$. For every subset $P \subseteq \Omega$, with a slight abuse of notation let $a^*(P)$ denote the action taken by Receiver when he knows (only) that the state of the world belongs to P.[17] For every $P \subseteq \Omega$, we say an action a is P-inducible if there exists a state $\psi \in P$ that induces a.

The proof first constructs π^* using an algorithm similar to the one in Section 2, and then shows that π^* satisfies the promised properties.

The algorithm starts at iteration $i = 1$, and at every iteration, it repeats the following:

i. From the set of all Ω-inducible actions, find the most preferred action with respect to $\preceq_A$; call it action a.
ii. Let ω be the $\preceq_\Omega$-lowest state in Ω that induces action a; i.e., $\omega \leftarrow \inf^{\preceq_\Omega}\{\omega' \in \Omega : \omega' \text{ induces } a\}$.
iii. Let $T_i \subseteq \Omega$ contain every state in Ω that is $\preceq_\Omega$-higher than ω; i.e., $T_i \leftarrow \{w' \in \Omega : w \preceq_\Omega w'\}$.
iv. Remove from Ω all of the states belonging to T_i; i.e., $\Omega \leftarrow \Omega \backslash T_i$.
v. If Ω is nonempty, then increase i by one and start the next iteration; otherwise, stop.

When the algorithm stops, the signal π^* is constructed as follows: The signal π^* has a signal realization space $S^* = \{T_1, \dots, T_i\}$, and sends a signal realization $T \in S^*$ to Receiver if and only if the realized state of the world belongs to T.

Claim A.1. *For every $T_i \in S^*$ and $t \in T_i$, $a^*(t) \preceq_A a^*(T_i)$. In addition, for every $T_j, T_k \in S^*$ with $k \geq j$, $a^*(T_k) \preceq_A a^*(T_j)$.*

Proof. To prove the first claim, we observe that step iii of the above algorithm chooses T_i so that it includes a state ω and all of the $\preceq_\Omega$-higher states in Ω. The state $\omega \in \Omega$ was chosen in the previous step (step ii) to be the state that induces the most preferred action with respect to $\preceq_A$. Thus, among all of the elements of T_i, its least preferred element with respect to $\preceq_\Omega$ (i.e., ω) induces the most preferred action with respect to $\preceq_A$. This proves the first claim.

[17] Recall that Receiver chooses an action with the most preferred promised outcome, as defined in Section 3.

To prove the second claim, observe that when iteration $i+1$ starts, every remaining element in Ω induces an action that is $\preceq_A$-lower than $a^*(T_i)$. Thus, $a^*(T_{i+1}) \preceq_A a^*(T_i)$. This implies that $a^*(T_k) \preceq_A a^*(T_j)$ for every $k \geq j$. □

Claim A.2. *For every $\omega \in \Omega$ and an arbitrary subset of the states $Q \ni \omega$, the following holds: If ω belongs to a signal realization T in π^*, then $(a^*(Q), \omega) \preceq_s (a^*(T), \omega)$.*

Proof. We observe that

$$a^*(Q) = a^*(\inf^{\preceq_\Omega}\{\omega' | \omega' \in Q\}) \preceq_A \sup^{\preceq_A}\{a^*(\omega'') | \omega'' \in \Omega, \omega'' \preceq_\Omega \omega\}, \tag{A.1}$$

where the equality holds by the agent's pessimism and the inequality holds since $\omega \in Q$.

On the other hand, from Claim A.1 it follows that

$$a^*(T) = \sup^{\preceq_A}\{a^*(\omega'') | \omega'' \in \Omega, \omega'' \preceq_\Omega \omega\}.$$

This equation, together with (A.1), implies that $a^*(Q) \preceq_A a^*(T)$. This, together with the action-monotonicity assumption, proves the claim. □

For every signal π, let the random variable $\mathfrak{a}_\omega^\pi$ denote the action taken by Receiver when Sender uses signal π and the state of the world is ω. We emphasize that $\mathfrak{a}_\omega^\pi$ is a random variable by the definition of a signal. Also, we note that $\mathfrak{a}_\omega^{\pi^*}$ is a constant random variable. This holds because, when Sender uses the signal π^*, Receiver always receives the same signal realization $T \ni \omega$ when the state of the world is ω, due to the construction of π^*.

Claim A.3. *Consider an arbitrary $\omega \in \Omega$ and an arbitrary signal π. Then, $\mathfrak{a}_\omega^\pi \preceq_A \mathfrak{a}_\omega^{\pi^*}$ holds for every realization of $\mathfrak{a}_\omega^\pi$.*

Proof. Let S be the signal realization space of π. For every $s \in S$, define

$$Q_s = \{\omega' : s \in \operatorname{supp} \pi(\cdot | \omega')\}.$$

We note that $a^*(Q_s)$ is the action taken by Receiver if she receives the signal realization s.

For every s that Receiver may receive when the state of the world is ω, it holds that $\omega \in Q_s$. Therefore, Claim A.2 applies, which implies that $a^*(Q_s) \preceq_A a^*(T)$. This proves the claim. □

Claim A.4. *The signal π^* is stochastically dominant.*

Proof. Consider an arbitrary signal π. By Claim A.3, for every $\omega \in \Omega$, $(\mathfrak{a}^{\pi}_{\omega}, \omega) \preceq_s (\mathfrak{a}^{\pi^*}_{\omega}, \omega)$ holds for every realization of $\mathfrak{a}^{\pi}_{\omega}$. Therefore, the distribution imposed by the signal π^* over the outcomes (i.e., D_{π^*}) stochastically dominates the distribution imposed by the signal π (i.e., D_{π}), in the sense of (3.3). $\square$

We have proved that the constructed signal π^* is stochastically dominant. To complete the proof of the theorem, it remains to show that properties (i), (ii), and (iii) in its statement hold. Property (i) holds by the construction of the signal. Monotonicity of π^* in property (ii) follows immediately from Claim A.1. The signal realizations in π^* are convex with respect to $\preceq_{\Omega}$ because, in step iii of the algorithm, when a state ω is included in a signal realization T_i, every other state that belongs to Ω at that time and is $\preceq_{\Omega}$-higher than ω is also added to T_i, and then T_i is removed from Ω. Thus, the convexity property holds. Property (iii), i.e., independence from the prior, holds by construction: the algorithm does not use the prior μ in constructing π^*.

A.2 Proof of Proposition 3.2

The proof approach is as follows. We will prove a stronger claim: we will show that under the cardinal preference model (where Sender's and Receiver's preferences are determined by utility functions as in Section 2), an optimal signal does not always exist if either of the state- or action-monotonicity assumptions is dismissed. This is done in Step 1. Then, in Step 2, we use this fact to prove the proposition.

Step 1. We next show that under the cardinal preference model, an optimal signal does not always exist if the state-monotonicity assumption is dismissed. After that, we prove the same but for action-monotonicity instead of state-monotonicity.

State-monotonicity does not hold. Let $\Omega = \{\omega_1, \omega_2, \omega_3\}$ and $A = \{a_1, a_2, a_3\}$. Also, let μ be the uniform distribution over Ω. Receiver's and Sender's utilities at every outcome are defined in Figure 2.

For every subset $P \subseteq \Omega$, with a slight abuse of notation let

$$\hat{a}(P) = \arg\max_{a \in A} \inf_{\omega \in P} u(a, \omega).$$

Actions \ States	ω_1	ω_2	ω_3
a_1	9	6	2
a_2	3	8	4
a_3	1	5	7

Actions \ States	ω_1	ω_2	ω_3
a_1	7	8	9
a_2	4	5	6
a_3	1	2	3

Figure 2: The left and right tables respectively define the utility functions of Receiver and Sender, $u, v : A \times \Omega \to \mathbb{R}_+$.

Observe that $\hat{a}(\{\omega_i\}) = a_i$ for every i, and that

$$\left.\begin{aligned} \hat{a}(\{\omega_1, \omega_2\}) &= a_1, \\ \hat{a}(\{\omega_1, \omega_3\}) &= a_2, \\ \hat{a}(\{\omega_2, \omega_3\}) &= a_3, \\ \hat{a}(\{\omega_1, \omega_2, \omega_3\}) &= a_2. \end{aligned}\right\} \tag{A.2}$$

This holds by the definition of Receiver's choice of action (2.2). (In particular, the definition implies that Receiver uses a pure strategy and does not randomize over different actions.)

Consider the signal π_ϵ with $S_\epsilon = \{s, s'\}$, where $s = \{\omega_1, \omega_2\}$ and $s' = \{\omega_1, \omega_3\}$. Also, $\pi_\epsilon(\cdot \mid \cdot)$ is defined as follows: $\pi_\epsilon(\cdot \mid \omega_1)$ assigns probability $1 - \epsilon$ to s and probability ϵ to s', and $\pi_\epsilon(\cdot \mid \omega_2)$ assigns probability 1 to s, and $\pi_\epsilon(\cdot \mid \omega_3)$ assigns probability 1 to s'. Hence, given the definition of $\hat{a}(\cdot)$, the value of π^ϵ equals

$$\frac{1}{3}\left((1-\epsilon)v(a_1, \omega_1) + \epsilon v(a_2, \omega_1) + v(a_1, \omega_2) + v(a_2, \omega_3)\right).$$

Therefore, to show that an optimal signal does not exist, it suffices to show that every signal has a value less than

$$\frac{1}{3}\left(v(a_1, \omega_1) + v(a_1, \omega_2) + v(a_2, \omega_3)\right). \tag{A.3}$$

The proof is by contradiction. Consider a signal π with the above value. Suppose that Receiver has received a signal realization s, and has a posterior belief μ_s, and that $\omega_0 = \omega_3$. Since $\omega_3 \in \text{supp}(\mu_s)$, then $\hat{a}(s) \neq a_1$ holds by (A.2). Hence, when $\omega_0 = \omega_3$, Receiver does not take action a_1. Therefore, for a signal to have value (A.3), Receiver must always take action a_1 when $\omega_0 \neq \omega_3$, and take action a_2 otherwise. This is, however, impossible, because $\hat{a}(\omega_3) = a_3$. Contradiction.

Action-monotonicity does not hold. Let $\Omega = \{\omega_1, \omega_2, \omega_3\}$ and $A = \{a_1, a_2, a_3\}$. Also, let μ be the uniform distribution over Ω. Receiver's and Sender's utilities at every outcome are defined in Figure 3.

States / Actions	ω_1	ω_2	ω_3
a_1	3	4	7
a_2	2	6	8
a_3	1	5	9

States / Actions	ω_1	ω_2	ω_3
a_1	100	101	1
a_2	2	3	102
a_3	4	5	6

Figure 3: The left and right tables respectively define the utility functions of Receiver and Sender, $u, v : A \times \Omega \to \mathbb{R}_+$.

Observe that $\hat{a}(\{\omega_i\}) = a_i$ for every i, and that

$$\left.\begin{aligned} \hat{a}(\{\omega_1, \omega_2\}) &= a_1, \\ \hat{a}(\{\omega_1, \omega_3\}) &= a_1, \\ \hat{a}(\{\omega_2, \omega_3\}) &= a_2, \\ \hat{a}(\{\omega_1, \omega_2, \omega_3\}) &= a_1. \end{aligned}\right\} \tag{A.4}$$

Consider the signal π_ϵ with $S_\epsilon = \{s, s'\}$, where $s = \{\omega_1, \omega_2\}$ and $s' = \{\omega_2, \omega_3\}$. Also, $\pi_\epsilon(\cdot \mid \cdot)$ is defined as follows: $\pi_\epsilon(\cdot \mid \omega_1)$ assigns probability 1 to s, $\pi_\epsilon(\cdot \mid \omega_2)$ assigns probability $1 - \epsilon$ to s and probability ϵ to s', and $\pi_\epsilon(\cdot \mid \omega_3)$ assigns probability 1 to s'. Hence, given the definition of $\hat{a}(\cdot)$, the value of π^ϵ equals

$$\frac{1}{3}\left(v(a_1, \omega_1) + (1 - \epsilon)v(a_1, \omega_2) + \epsilon v(a_2, \omega_2) + v(a_2, \omega_3)\right).$$

Therefore, to show that an optimal signal does not exist, it suffices to show that every signal has a value less than

$$\frac{1}{3}\left(v(a_1, \omega_1) + v(a_1, \omega_2) + v(a_2, \omega_3)\right). \tag{A.5}$$

The proof is by contradiction. Consider a signal π with the above value. By the definition of Sender's utilities (Figure 3), for π to have the above value, Receiver must always take action a_2 when $\omega_0 = \omega_3$ and take action a_1 otherwise. This, however, is not possible because $\hat{a}(w_3) = a_3$. Contradiction.

Step 2. We call each of the two examples provided above a *problem in the cardinal preference model*. To complete the proof of the proposition, we will *transform* a problem in the cardinal preference model into a problem in the ordinal preference model of Section 3. We will see that under this transformation, (i) if either of the action- or state-monotonicity conditions does not hold in the original problem, then that condition also does not hold in the transformed problem, and (ii) if an optimal signal does not exist in the original problem, then a stochastically dominant signal does not exist in the transformed problem. These two facts together would conclude the proof.

To define the transformation, consider a problem in the cardinal preference model. The function $v : A \times \Omega \to \mathbb{R}$ induces a preference order over the outcomes; let this preference order be denoted by $\preceq_s$. Similarly, the function $u_0 : A \times \Omega \to \mathbb{R}$ induces a preference order over the outcomes; let this preference order be denoted by $\preceq_r$. We define the *transformed problem* to be the problem in which (i) Receiver and Sender respectively have preference orders $\preceq_r$ and $\preceq_s$ over the outcomes, and (ii) the prior distribution over the states is the same as in the original problem.

Next, we show that, under every signal π, Receiver's action after receiving a signal realization s in the original problem is the same as his action in the transformed problem. Consider the scenario in the original problem where the state of the world is ω, Sender uses an arbitrary signal π, and Receiver receives a signal realization $s \sim \pi(\cdot|\omega)$. Recall from Remark 2.1 that, for a pessimistic Receiver in the cardinal preference model,

$$u(a, \mu_s) = \inf_{\omega \in \text{supp}(\mu_s)} u_0(a, \omega),$$

where μ_s denotes the posterior formed by Receiver using Bayes's rule after observing s. The action taken by Receiver, $\hat{a}(s)$, belongs to the set $\arg\max_{a \in A} u(a, \mu_s)$. That is,

$$\hat{a}(s) \in \arg\max_{a \in A} \inf_{\omega \in \text{supp}(\mu_s)} u_0(a, \omega). \tag{A.6}$$

Define

$$\underline{o}(a, s) = \inf\nolimits^{\preceq_r} \{(a, \omega) : s \in \text{supp}(\pi(\cdot|\omega))\}.$$

By (A.6), the action $\hat{a}(s)$ satisfies

$$\underline{o}(a, s) \preceq_r \underline{o}(\hat{a}(s), s)$$

for every $a \in A$. This coincides with Receiver's choice of action in the transformed problem, as defined by (3.2).

We have shown that, under every signal π, Receiver's action after receiving a signal realization s in the original problem is the same as his action in the transformed problem. Therefore, the distribution D_π imposed over the outcomes is the same in the original problem and in the transformed problem. Hence, if there exists a stochastically dominant signal π^* in the transformed problem, then π^* is an optimal signal in the original problem. Since there is no optimal signal in the original problem (by Step 1), there is no stochastically dominant signal in the transformed problem. This completes the proof.

B Proofs from Section 4

Proof of Proposition 4.1. **Part i.** The proof for the "if" direction of part (i) is straightforward: we observe that, if a^* is increasing over Ω, then Theorem 3.1 constructs a signal with size $|\Omega|$.

We next show that if there is a stochastically dominant signal, namely π, with size $|\Omega|$, then a^* is increasing over Ω. The proof is by contradiction. Suppose there exist distinct $\omega_i, \omega_j \in \Omega$ such that $\omega_i \prec_r \omega_j$ and $a^*(\omega_j) \prec_A a^*(\omega_i)$. Define the signal π' as the signal that reveals to Receiver the state of the world if it belongs to $\Omega \setminus \{\omega_i, \omega_j\}$, and otherwise it reveals to Receiver only that the state of the world belongs to $\{\omega_i, \omega_j\}$. Observe that π' stochastically dominates π, which is a contradiction.

Part ii. To prove the "if" direction of part (ii), we observe that, under the non-informative signal, Receiver always takes the action $a^*(\omega_1)$. When a^* is maximized over Ω at ω_1, then $a^*(\omega_1)$ is the Ω-inducible action that Sender most prefers (with respect to $\preceq_A$). Therefore, under the non-informative signal, at every state of the world Receiver takes the Ω-inducible action that Sender most prefers. Hence, the non-informative signal is stochastically dominant.

The proof for the "only if" direction is by contradiction. Suppose that there exists a non-informative stochastically dominant signal when $a^*(\omega_1) \prec_A a^*(\omega_i)$ for some $\omega_i \in \Omega$. Then, consider the signal that reveals the state of the world if it is ω_i and otherwise reveals only that the state of the world is not ω_i. Observing that this signal stochastically dominates the non-informative signal gives a contradiction.

Part iii. To prove the "if" direction, define $t^* = \min\{i : a^*(\omega_i) = a^*(\omega_t)\}$. Then, observe that the algorithm of Section A.1 constructs a signal with a realization space

$$S = \{\{\omega_1\}, \ldots, \{\omega_{t^*}\}, \{\omega_{t^*+1}, \ldots, \omega_n\}\},$$

where a signal realization $P \in S$ is sent to Receiver if and only if it contains the realized state of the world. This, by definition, is an upper censorship signal.

To prove the converse, suppose that π is a stochastically dominant signal and an upper censorship signal. Hence, without loss of generality, we can assume that there exists t such that the realization space of π is

$$S = \{\{\omega_1\}, \ldots, \{\omega_{t'}\}, \{\omega_{t'+1}, \ldots, \omega_n\}\},$$

where a signal realization $P \in S$ is sent to Receiver if and only if it contains the realized state of the world. Then, by the same argument as in part (i), a^* must be increasing over $\{\omega_1, \ldots, \omega_t'\}$. Also, by the same argument as in part (ii), a^* must be maximized over $\{\omega_{t'+1}, \ldots, \omega_n\}$ at $\omega_{t'+1}$. If $a^*(\omega_{t'+1}) \preceq_A a^*(\omega_{t'})$, then define $t = t'$. Otherwise, then define $t = t' + 1$. Observing that such t satisfies the condition stated in part (iii) proves the claim.

Part iv. To prove the "if" direction, we define t' to be the smallest positive integer such that a^* is increasing over $\{\omega_{t'}, \ldots, \omega_n\}$ and $a^*(\omega_1) \prec_A a^*(\omega_t')$. If such an integer does not exist, then let $t' = n + 1$. We observe that the algorithm of Section A.1 constructs a signal that fully reveals the state of the world to Receiver if it belongs to the set $\{\omega_{t'}, \ldots, \omega_n\}$ (which would be the empty set when $t' = n + 1$), and otherwise it reveals to Receiver only that the state of the world belongs to the set $\{\omega_1, \ldots, \omega_{t'-1}\}$ (which would be the empty set when $t' = 1$). Such a signal, by definition, is a lower censorship signal.

To prove the converse, suppose that π is a stochastically dominant signal and a lower censorship signal. If π is a non-informative signal, then part (ii) concludes the proof. Hence, suppose that π is not a non-informative signal. Therefore, there exists a positive integer $t' \leq n$ such that the signal π fully reveals the state of the world to Receiver if it belongs to the set $\{\omega_{t'}, \ldots, \omega_n\}$, and otherwise it reveals to Receiver only that the state of the world belongs to the set $\{\omega_1, \ldots, \omega_{t'-1}\}$ (which would be the empty set when $t' = 1$). By part (ii), a^* should be maximized over $\{\omega_1, \ldots, \omega_{t'-1}\}$ at ω_1. Also, by part (i), a^* should be increasing over $\{\omega_{t'}, \ldots, \omega_n\}$. Let j be the smallest integer larger than $t' - 1$ such that $a^*(\omega_1) \preceq_A a^*(\omega_j)$. Such j must exist since π is not a non-informative signal. We observe that setting $t = j$ satisfies the condition stated in part (iv). This completes the proof. □

C Alternative construction of the optimal signal

Below we describe an alternative approach to prove the main theorem. For expositional simplicity, we limit the discussion in this section to the principal-agent setup from Section 2.

For functions $f, g : V \to \mathbb{R}$, we say f is *point-wise larger* than g if $f(v) \geq g(v)$ for all $v \in V$. We say g is *point-wise smaller* than f if f is point-wise larger than g. The *monotone orbit* of f is the set of all increasing functions from V to $\mathbb{R}$ that are point-wise larger than f. The *monotone envelope* of f is a function $\widehat{f}$ in the monotone orbit of f that is point-wise smaller than every other function in the monotone orbit of f.

Proposition C.1. *Every function $f : V \to \mathbb{R}$ has a monotone envelope.*

Proof. Let $\mathcal{F}$ denote the monotone orbit of f. Define the function $g : V \to \mathbb{R}$ for every $v \in V$ by

$$g(v) = \inf_{h \in \mathcal{F}} \{h(v)\}.$$

For every $v_1, v_2 \in V$ with $v_1 \leq v_2$, it holds that $g(v_1) \leq g(v_2)$. The reason is that otherwise $g(v_1) > g(v_2)$ implies there exists a function in $\mathcal{F}$ that is not increasing, which would be a contradiction. Thus, g is increasing and point-wise smaller than every other function in $\mathcal{F}$. This means that g is the monotone envelope of f.

□

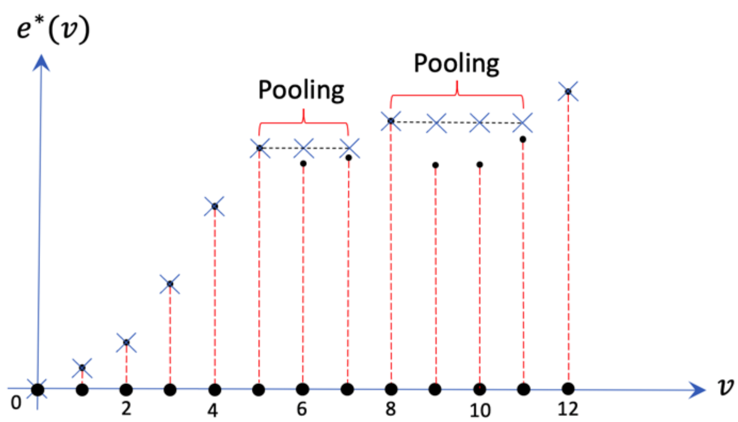

Figure 4: Points $(v, e^*(v))$ are marked by solid dots for every $v \in V$, and the points $(v, \widehat{e^*}(v))$ are marked by x's.

The intuitive description of the optimal signal is that this signal *pools* a subset of states $V' \subseteq V$ if $\widehat{e^*}$ assigns the same value to every state belonging to V'. We next elaborate this using the principal-agent model of Section 2. We consider the example of Figure 1 from that section with $V = \{1, \ldots, 12\}$, and plot the monotone envelope of e^* in Figure 4. The values of e^* and $\widehat{e^*}$ are marked in the figure. Let $T_0 = \varnothing$, and for every $i \geq 0$, let

$$T_{i+1} = \{v : \widehat{e^*}(v) = \max_{v' \in V \setminus (T_1 \cup \ldots \cup T_i)} \widehat{e^*}(v')\}.$$

Consider the signal which discloses to the agent (only) the subset T_i that contains the realized project value. We note that this signal is identical to the signal constructed in the example of Figure 1, and thus is optimal.

D Connection to communication games

Here we consider a communication game akin to Milgrom (2008), with the difference that the agent has *maxmin expected utility* preferences. We will study the equilibria of this game, and show that one of them corresponds to the optimal signal in our setup. For expositional simplicity, we limit the discussion in this section to the principal-agent setup in Section 2.

Suppose that, after observing the project value v, the principal sends a *message* to the agent, relaying that the project value is *at least* x, for some $x \in V$. (For brevity, we say that the principal *reports* x.) We assume that the principal does not make a false report (e.g., due to a legal/institutional structure that prevents the principal from doing so, as in Milgrom (2008)). The agent, after receiving the message, chooses an effort level $e \in E$. The agent's and principal's payoffs are then $u_a(e, v)$ and $u_p(e, v)$, respectively, as in Section 2. Also as in that section, we assume that the functions u_a and u_p satisfy the state-monotonicity and action-monotonicity conditions, respectively.

An equilibrium of this game consists of two types of objects: *strategies* and a *belief function*. The principal's strategy is a report function $\tilde{R} : V \to V$, indicating that the principal reports $\tilde{R}(v)$ when the project value is v. Let $\tilde{B} : V \to \Delta(V)$ be such that $B(x)$ denotes the agent's belief about the project value after hearing a report $x \in V$. Finally, let $\tilde{f} : V \to E$ be such that $f(x)$ denotes the equilibrium decision by the agent when he has received a report x. The triple $(\tilde{R}, \tilde{B}, \tilde{f})$ is called an equilibrium when, for any project value $v \in V$, the following conditions are satisfied:

i. Making the report $\tilde{R}(v)$ maximizes the principal's payoff over all truthful reports that she can make. (The principal computes her payoff by taking $\tilde{B}$ and $\tilde{f}$ as given.)
ii. After observing the report $\tilde{R}(v)$, the agent computes the *worst-case* payoff of an effort level $e \in E$ given his belief $\tilde{B}(\tilde{R}(v))$ as $\min_{v' \in \text{supp} \tilde{B}(\tilde{R}(v))} u_a(e, v')$. The agent's chosen

effort level $\tilde{f}(v)$ must have the highest worst-case payoff over all effort levels in E.

iii. The agent's belief after receiving the report is determined by Bayes' rule whenever it applies.

We will see that there is an equilibrium of the game corresponding to the signal π^* constructed in Section 2. The game has other equilibria, as discussed below. The multiplicity arises essentially because the agent's beliefs off the equilibrium path can be chosen arbitrarily. We will discuss this further later.

The principal-optimal equilibrium. Recall the algorithm of Section 2 that constructs the signal π^*, and the sets $T_1, T_2, \ldots$ from that construction. Let t_i denote the smallest element of T_i, for every i. Let the report function be such that $\tilde{R}(v) = t_i$ if $v \in T_i$. That is, for every project value belonging to T_i, the principal reports that the project value is at least t_i. Also, for $v \in T_i$, let $\tilde{B}(v)$ be the marginal distribution induced by μ on T_i and define $\tilde{f}(v) = e^*(t_i)$.

By this construction, the agent receives a message t_i if and only if the project value v belongs to T_i. Then, by Bayes' rule, the agent's belief about the project value would be $\tilde{B}(v)$. Note that $\tilde{B}(v)$ has full support on T_i since μ (which induces $\tilde{B}(v)$) has full support on V. Hence, the worst-case payoff of the agent at every effort level $e \in E$ is $u_a(e, t_i)$, by the state-monotonicity assumption. Therefore, the effort level $e^*(t_i) = \tilde{f}(v)$ has the highest worst-case payoff when the agent has received a report t_i. This confirms that the triple $(\tilde{R}, \tilde{B}, \tilde{f})$ defined here is an equilibrium.

Finally, we note that for any project value, the agent's chosen effort in the above equilibrium coincides with his choice when the signal π^* is used in the setup from Section 2. Therefore, the principal's expected payoff in this equilibrium equals her expected payoff under signal π^*. On the other hand, the principal's expected payoff at any equilibrium is at most her expected payoff under signal π^*. (The reason is that the principal has commitment power when using a signal, but not in the equilibrium notion defined in this section.) Thus, there is no equilibrium with a higher expected payoff for the principal than the equilibrium $(\tilde{R}, \tilde{B}, \tilde{f})$ defined here.

An equilibrium with no information revelation. Let v_0 denote the smallest element of V. Define $\tilde{R}(v) = v_0$, $\tilde{B}(v) = \mu$, and $\tilde{f}(v) = e^*(v_0)$ for every $v \in V$. This corresponds to an equilibrium where, regardless of the true project value, the principal always reports to the agent that the project value is at least v_0. (Hence, the principal reveals no information to the agent.) The agent's beliefs, on and off the equilibrium path, coincide with the prior

μ. Thus, on the equilibrium path, the agent always takes action while assuming that the project has the lowest possible value v_0, and the principal has no incentive to change her strategy because of the agent's beliefs off the equilibrium path.

Full information revelation is not always an equilibrium. For every $v \in V$, let $\tilde{R}(v) = v$ and $\tilde{B}(v) = \mathbb{1}_v$ where the right-hand side denotes the degenerate distribution on v. Also, let $\tilde{f}(v) = e^*(v)$. Thus, if the triple $(\tilde{R}, \tilde{B}, \tilde{f})$ is an equilibrium, it would be a fully revealing one. However, this triple would be an equilibrium if and only if $e^*(v_1) \leq e^*(v_2)$ for every to $v_1, v_2 \in V$ with $v_1 \leq v_2$. To illustrate why, we consider the example in Figure 1 where this condition does not hold. Suppose that the project value is $v = 10$. Then, the agent exerts effort $e^*(8)$ if the principal reports that the project value is at least 8. (The reason is that $\tilde{B}(8) = \mathbb{1}_8$.) Hence, by making such a report, the principal can secure a higher payoff for herself than when she reports that the project value is at least 10. Therefore, the triple $(\tilde{R}, \tilde{B}, \tilde{f})$ that we constructed here is not an equilibrium.

On the other hand, when $e^*(\cdot)$ *is* an increasing function, the triple $(\tilde{R}, \tilde{B}, \tilde{f})$ that we constructed here is an equilibrium. (Because then, if the principal makes a report that is lower than the true project value, the agent would exert only lower effort.) Milgrom (2008) assumes a similar condition, which leads to the equilibrium being fully revealing.[18] His condition implies that "if the quality level is higher, then the optimal choice of the purchased quantity will be higher—that is, the buyer will be willing to buy a higher quantity." This condition ensures that a higher-quality seller benefits from distinguishing themself from lower-quality sellers, and thus derives the so-called unravelling argument, which leads to the equilibrium being fully revealing in (Milgrom 2008).

Online Appendix

i Partially ordered state space for Receiver

In our main setup (Section 3), we assumed that Receiver's preference order over the states, $\preceq_\Omega$, is a total order. This may not hold in some applications, such as in multidimensional persuasion; e.g., as in the case of a principal informing an agent about several projects. To address such cases, we allow $\preceq_\Omega$ to be a partial order over the states, and provide a

[18] The exact condition that he imposes is that the marginal value of an increase in the quantity purchased by the buyer to him is increasing in the quality of the seller's good.

counterpart to Theorem 3.1 in this setting.

What differs here from our main setup is that, in the definition of the state-monotonicity condition, we allow $\preceq_\Omega$ to be a lower semilattice.[19] For every subset $P \subseteq \Omega$, let $\bigwedge P$ denote the greatest lower bound for P with respect to $\preceq_\Omega$.[20] Also, let $a^*(P)$ denote the action chosen by Receiver when he knows only that the state of the world belongs to P. (We recall that, Receiver chooses actions as defined in Section 3, by choosing an action that has the most preferred promised outcome.) We say Receiver has *commutative* preferences when, for every $P \subseteq \Omega$

$$a^*(P) = a^*\left(\bigwedge P\right).$$

Example i.1 (A simple example of commutative preferences)**.** *Consider a principal (she) persuading an agent (he) to work on one of the n projects, namely, $1, \ldots, n$. The set of agent's actions is $A = \{1, \ldots, n\}$; i.e., the agent may choose one project to work on. He attains utility ω_i from working on a project i, where it is common knowledge that ω_i belongs to a set Ω_i. The principal knows ω_i for every i, while the agent does not. The principal attains a utility ν_i if the agent works on project i (where $\nu_i \neq \omega_i$ may possibly hold).*

Let $\Omega = \prod_{i=1}^n \Omega_i$. For every $\omega \in \Omega$, ω_i denotes the i-th coordinate of ω. Also, we let $\preceq_\Omega$ be a lower semilattice over Ω with its meet operator $\bigwedge$ defined as

$$\bigwedge P = \left(\min_{\omega \in P} \omega_1, \ldots, \min_{\omega \in P} \omega_n\right)$$

for every $P \subseteq \Omega$.

We next verify that the agent, i.e., Receiver, has commutative preferences. Recall from Section 3 the definition of the promised outcome *of an action, and that that Receiver chooses an action with the most preferred promised outcome, as stated in (3.2). Suppose that Receiver has received a signal realization indicating (only) that the state of the world belongs to a subset $P \subseteq \Omega$. Then, the promised outcome of an action $i \in A$ is $\min_{\omega \in P} \omega_i$. Hence,*

$$a^*(P) = \arg\max_{i \in A} \min_{\omega \in P} \omega_i.$$

The right-hand side equals $a^*\left(\bigwedge P\right)$, which means that the agent has commutative preferences.

[19] Recall that a lower semilattice is a partial ordering in which any non-empty finite subset of its elements have a greatest lowerbound. In particular, every lattice is a lower semilattice.

[20] We recall that, for every subset of the elements of a lower semilattice, there is a unique element of the lower semilattice that is the greatest lower bound for that subset.

We remark that when $\preceq_\Omega$ is a total order, commutativity is implied by state-monotonicity.[21] Thus, the setup in here contains our main setup in Section 3.

We assume that Receiver's preferences are commutative. All else remains the same as in our main setup. In particular, we say that Sender's preferences are monotone in actions if there exists a total order $\preceq_A$ over the actions such that whenever $a \preceq_A a'$, then $(\omega, a) \preceq_s (\omega, a')$ holds for every $\omega \in \Omega$. A signal is *monotone* if it recommends $\preceq_A$-higher actions at $\preceq_\Omega$-higher states. A subset $P \subseteq \Omega$ is *convex* with respect to $\preceq_\Omega$ if for every $\omega, \omega', \omega'' \in \Omega$ for which $\omega \preceq_\Omega \omega' \preceq_\Omega \omega''$ and $\omega, \omega'' \in P$ hold, $\omega' \in P$ holds as well.

When Receiver's preferences are commutative and Sender's preferences are monotone in actions, there exists a stochastically dominant signal π^*. This can be shown by an argument similar to the proof of Theorem 3.1. In particular, the signal π^* can be constructed by a similar algorithm. In addition, π^* satisfies the same properties as in Theorem 3.1.

ii Unordered action space for Sender

We start this section with a new signal construction approach which, in the absence of action-monotonicity, produces an "almost stochastically dominant" signal when the probability mass that the prior assigns to every single state is "small". In particular, our construction turns out to be stochastically dominant when the state space is compact and the prior is atomless. We also adapt our approach here to solve persuasion problems in the presence of upper quota constraints on the signal size, which may be imposed to ensure simplicity. Finally, we provide some insight into this new approach by looking at an example in the context of ride-sharing platforms, where action-monotonicity fails to hold.

Setup. Throughout this section, we suppose that Receiver has commutative preferences. (Recall the definition of commutativity from Section i.) Thus, Receiver's preference order over the states (i.e., $\preceq_\Omega$) would be a lower semilattice. Let $\bigwedge$ denote the meet operator associated with $\preceq_\Omega$. We do not make any assumptions about Sender's preference order over the outcomes; in particular, Sender's preferences may not be monotone in actions.

[21] To see why, suppose that state-monotonicity holds. Let $P \subseteq \Omega$, and define $\omega = \inf^{\preceq_\Omega}\{P\}$. By state-monotonicity, $a^*(P) = a^*(w)$. Since $\omega = \bigwedge P$, then Receiver has commutative preferences.

ii.1 The greedy signal

We provide a simple *greedy* approach that constructs a signal for Sender when Receiver's preferences are commutative. This signal, henceforth the *greedy* signal, is constructed by assigning each state of the world to one of the potential signal realizations. The formal construction is given in Algorithm 1, henceforth the *greedy* algorithm.

For a preference order $\preceq$ defined over a set X, let $\sup^{\preceq}\{X\}$ denote $x \in X$ such that $y \preceq x$ holds for all $y \in X$. If there are multiple such x's available, $\sup^{\preceq}\{X\}$ picks one arbitrarily. We denote the signal by π_G and its realization space by S_G.

Algorithm 1: The greedy algorithm

1 $R \leftarrow \varnothing$.
2 For every action a, let Ω_a be the subset of states that induce a.
3 For every action a with $\Omega_a \neq \varnothing$, define $\omega_a = \bigwedge \Omega_a$, add ω_a to R, and let $P_a = \{\omega_a\}$.
4 For any state $\omega \notin R$, define a_ω by letting $(a_\omega, \omega) = \sup^{\preceq_s}\{(a, \omega) : a \in A,\, \omega_a \preceq_\Omega \omega\}$.
5 For any action a with $\Omega_a \neq \varnothing$, add $\{\omega : a = a_\omega\}$ to P_a.
6 Let $S_G = \bigcup_{a \in A : P_a \neq \varnothing} \{P_a\}$.
7 Define π_G to be the signal that, when the state of the world is ω, sends a signal realization $Q \in S_G$ such that $\omega \in Q$.

Henceforth, we call Algorithm 1 the *greedy algorithm*. The algorithm keeps track of a set of representative states, R. At most one representative state is added to R for each action a, namely ω_a. The algorithm constructs a signal that contains a signal realization per representative state, as follows. Each state $\omega \notin R$ is greedily *assigned* to one of the representative states: to the state ω_a corresponding to the action a, where a is Sender's most preferred action at state ω among all actions a' that satisfy $\omega_{a'} \preceq_\Omega \omega$. (Line 4 of the algorithm.) Finally, each representative state together with the states assigned to it are defined as a signal realization in the greedy signal, π_G.

The key point is that, under the greedy signal, at any nonrepresentative state, Sender takes her most preferred action that could be taken at that state under *any* signal. This point is a consequence of commutativity, as elaborated next. Consider the scenario where Sender uses an arbitrary signal π, the state of the world is a nonrepresentative state ω, and Receiver receives a signal realization $s \sim \pi(\cdot|\omega)$. Define $Q_s = \{\omega' : s \in \operatorname{supp} \pi(\cdot|\omega')\}$. Let $a^*(Q_s) = b$. (We recall the definition of $a^*(\cdot)$ from the proof of Theorem 3.1) Then, $\omega_b \preceq_\Omega \omega$ must hold, because $\omega \in Q_s$. This means that (b, ω) appears in the argument of $\sup^{\preceq_s}$ in line 4 of the algorithm, which proves the claim.

This point does not necessarily hold for representative states, which is why the algorithm does not always construct a stochastically dominant signal. However, if the probability measure that the prior assigns to the set of representative states is "negligible", then π_G is "almost" stochastically dominant. In particular, when the state space is compact and the prior is atomless, π_G turns out to be stochastically dominant, as we will formalize next.[22]

ii.2 The greedy signal in compact state spaces

When the state space is compact, there are pathological partitions of the state space that are unlikely to be of practical interest (e.g., if a member of the partition is a fat contour set). To rule out some pathological cases, we need a definition.

A signal π is *proper* if the upper contour set (with respect to the order $\preceq_s$) of any outcome is D_π-measurable.[23] Properness ensures that, given π and any outcome o, the probability that Sender prefers the realized outcome under π to the outcome o is well-defined.

Properness allows the stochastic dominance relation between signals to be defined in the natural way when the state space is compact: a proper signal π stochastically dominates a proper signal π' if, for any outcome p,

$$\int_{p \preceq_s q} \mathrm{d}D_\pi(q) \geq \int_{p \preceq_s q} \mathrm{d}D_{\pi'}(q).$$

A proper signal is then called a stochastically dominant if it stochastically dominates any other proper signal.

Proposition ii.1. *When the prior is atomless, if the greedy signal is proper, it is also stochastically dominant.*

Proof. Let π_G denote the greedy signal. For a state ω, we say Sender ω-prefers signal π_G to π if, when the state of the world is ω, she always weakly prefers the outcome realized under the signal π_G to any outcome that may be realized under signal π'.

Claim ii.2. *For any proper signal π, Sender ω-prefers π_G to π for all but a μ-measure zero of the states $\omega \in \Omega$.*

[22] We remark that in the presence of action-monotonicity, the greedy algorithm constructs a stochastically dominant signal even when the state space is discrete. The greedy construction, however, reveals less about the structural properties of stochastically dominant signals (such as their convexity) than the construction approach that we used in Section 3.

[23] Recall the defnition of D_π from Section 3.

If the above claim holds, the claim of the proposition must hold as well. Suppose not; then there exist a signal π and an outcome p such that

$$\int_{p \preceq_s q} \mathrm{d}D_{\pi_G}(q) < \int_{p \preceq_s q} \mathrm{d}D_{\pi}(q).$$

But that means there exists a positive μ-measure of states such as ω such that Sender ω-prefers π to π_G, which would contradict Claim ii.2. Therefore, it remains to prove Claim ii.2.

Proof of Claim ii.2. We will prove that Sender ω-prefers π_G to π for any nonrepresentative state ω (i.e., any $w \notin R$). Since there are only a finite number of representative states, this will prove the claim.

Consider the scenario where the state of the world is $\omega \notin R$ and Receiver receives a signal realization s under π. Define $Q_s = \{\omega : s \in \operatorname{supp}(\pi(\cdot|w))\}$. Note that $\omega \in Q_s$. To prove the claim, it suffices to show that $(a_\omega, \omega) \prec_s (a^*(Q_s), \omega)$ does not hold. For contradiction, suppose that it does. Let $b = a^*(Q_s)$. By commutativity, Ω_b contains $\bigwedge Q_s$, and therefore it is nonempty. Define $\omega_b = \bigwedge \Omega_b$. Note that $\omega_b \preceq_\Omega \omega$ holds, since Ω_b contains $\bigwedge Q_s$. Therefore, from line 4, we should have

$$(b, \omega) \preceq_s \sup\nolimits^{\preceq_s}\{(a, \omega) : a \in A,\, \omega_a \preceq_\Omega \omega\},$$

which implies that

$$(a_\omega, \omega) \prec_s (b, \omega) \preceq_s (a_\omega, w).$$

Contradiction.

□

□

By construction, the stochastically dominant signal satisfies the *independence from the prior* property defined in Section 3. Unlike there, the signal realizations are not necessarily convex, as we will see in the ride-sharing example in Section ii.4.

In the cardinal preference model, our definition of the properness of a signal boils down to a simple one: integrability of Sender's utility function with respect to the prior, under that signal. The details are discussed in Section ii.3. There, we also discuss some of the (mild) sufficient conditions that ensure the properness of the greedy signal. The example of ride-sharing platforms in Section ii.4 provides some insight.

ii.2.1 Exogenous constraints on the signal structure.

Our approach here can be adapted to solve persuasion problems in the presence of exogenous constraints. For example, suppose there is an upper quota constraint on the size of Sender's signal, i.e., the number of signal realizations. Such constraints could be present to ensure simplicity. The stochastically dominant signal, when it exists, can be found by constructing a series of signals: for each subset of the representative states of the allowed size, namely $S \subseteq R$, consider a signal with $|S|$ signal realizations. Initially, each signal realization contains precisely one of the states in S. Each other state is then added to one of the $|S|$ signal realizations; specifically, to the one that Sender prefers the most.[24] It can be shown that if one of the produced signals stochastically dominates the other produced signals, then that signal is also a stochastically dominant signal. In the cardinal preference model, Sender's optimal signal, which always exists, is the produced signal she most prefers.

ii.3 Optimality of the greedy signal in the cardinal preference model

We provide a counterpart for Proposition ii.1 for when Sender's preferences over the outcomes are cardinal. The properness condition will be replaced with *validity*, a weaker condition.

For any signal ρ, let $\rho(\omega)$ denote Sender's (expected) utility conditioned on the state of the world being equal to ω. A signal ρ is *valid* if $\int_{\omega\in\Omega} \rho(\omega)\mathrm{d}\mu(\omega)$ is well-defined, i.e., the Lebesgue integral exists.

Proposition ii.3. *Suppose Ω is compact and μ is atomless. Then, when the signal constructed by the greedy algorithm is valid, it is also optimal.*

Proof. For any valid signal ρ, let $\rho(\omega)$ denote Sender's (expected) payoff conditioned on the state of the world being equal to ω. (The expectation applies in case Receiver uses randomization.) We will show that, for any signal ρ,

$$\pi_G(\omega) \geq \rho(\omega), \quad \forall \omega \notin R \tag{ii.1}$$

holds. Given this inequality, the proof would be complete: Sender's payoff under the signal π_G is just equal to $\int_{\omega\in\Omega} \pi_G(\omega)\mathrm{d}\mu(\omega)$. Since there is only a finite number of representative

[24] Formally, "Sender prefers the most" means the following. Fix the $|S|$ signal realizations and let $\mathcal{P}$ denote the set containing all of them. Consider a state ω which is supposed to be added to one of the elements of $\mathcal{P}$. Each element of $\mathcal{P}$, namely P, induces an action known to Sender, namely a_P. The most preferred signal realization (to which ω will be added) for Sender is then Q, where $(a_Q, \omega) = \sup^{\preceq_s}\{(a_P, \omega) : P \in \mathcal{P}\}$.

states (with measure 0), and since μ is atomless, then by (ii.1)

$$\int_{\omega\in\Omega} \pi_G(\omega)\mathrm{d}\mu(\omega) \geq \int_{\omega\in\Omega} \rho(\omega)\mathrm{d}\mu(\omega)$$

holds for any signal ρ. Therefore, π_G is optimal.

It remains to prove (ii.1). Recall from the greedy algorithm that, for any action a with $\Omega_a \neq \varnothing$, we define $\omega_a = \bigwedge \Omega_a$. Also, for any state $\omega \notin R$, we define

$$a_\omega = \arg\max_{a\in A:\, \omega_a = \bigwedge\{\omega_a,\omega\}} v(a,\omega).$$

To prove (ii.1), consider the scenario where the state of the world is an arbitrary state $\omega \notin R$, Sender uses an arbitrary signal ρ, and Receiver receives a signal realization s. Define $Q_s = \{\omega : s \in \mathrm{supp}(\rho(\cdot|w))\}$. Note that $\omega \in Q_s$. Let $b = a^*\left(\bigwedge Q_s\right)$. Because $\omega \in Q_s$, $\bigwedge Q_s \in \Omega_b$, and $\omega_b = \bigwedge \Omega_b$, then $\omega_b = \bigwedge\{\omega_b, \omega\}$. Therefore, $v(a_\omega, \omega) \geq v(b, \omega)$, due to the definition of a_ω in line 4 of the greedy algorithm. The latter inequality holds for every signal realization received by Receiver when the state of the world is ω; hence, (ii.1) holds. □

The signals constructed by the greedy algorithm are generally valid for applications of practical interest: Producing signals that are not valid means producing signals that induce non-integrable payoff functions, which is unlikely in practical applications. Next, we discuss some of the possible sufficient conditions that ensure the validity of the constructed signal.

ii.3.1 Validity of greedy signals

Recall that when the state space is compact, we say a signal ρ valid when Sender's expected payoff, $\int_{\omega\in\Omega} v(\rho(\omega),\omega)\,\mathrm{d}\,\mu(\omega)$, is well-defined, i.e., the latter Lebesgue integral exists. When the points of discontinuity of the integrand has μ-measure zero, the existence of the integral is guaranteed. Given that there are only a finite set of actions, this turns out to be satisfied under mild conditions, as discussed next.

We introduce two technical conditions which ensure that the lower semilattice and the payoff functions are "well-behaving". That, in turn, would guarantee the validity of the constructed signal. After briefly discussing these conditions below, we discuss them more extensively.

Condition i: on the lower semilattice. For any $\omega \in \Omega$, its upper contour set with respect to $\preceq_\Omega$, i.e., $\{\omega' \in \Omega : \omega \preceq_\Omega \omega'\}$, is a closed set.[25]

Condition ii: on the payoff functions. Let a *subproblem* be defined as a closed subset of the states $\Psi \subseteq \Omega$ and a subset of the actions $B \subseteq A$. The condition is that the payoff functions of Sender and Receiver are Lebesgue-integrable under the fully revealing signal in any subproblem. That is, the integrals $\int_{\omega \in \Psi} u(\rho(\omega), \omega) \, \mathrm{d}\, \mu(\omega)$ and $\int_{\omega \in \Psi} v(\rho(\omega), \omega) \, \mathrm{d}\, \mu(\omega)$ exist for any closed $\Psi \subseteq \Omega$ where $\rho(\omega)$ is Receiver's optimal action at state ω when the set of available actions is limited to B.

For example, when Ω is a compact subset of $\mathbb{R}^n$, the above condition is always satisfied if the functions $u(\cdot, a), v(\cdot, a)$ are real analytic functions for any action $a \in A$.

Proposition ii.4. *If Conditions i and ii (defined above) hold, then the greedy signal is valid.*

Proof. For any representative state $\omega_a \in R$ corresponding to action a, let U_a denote the upper contour set of ω_a with respect to $\preceq_\Omega$. For any subset of actions $B \subseteq A$ define

$$\xi(B) = \left(\bigcap_{x \in B} U_x \right) - \left(\bigcup_{y \notin B} U_y \right).$$

Observe that for any $\omega \in \Omega$ that could be induced by an action, there exists a unique subset $B \subseteq A$ that contains ω. Therefore, to prove that the constructed signal is valid, it suffices to show that the integrals $\int_{\omega \in \xi(B)} u(\rho_B(\omega), \omega) \, \mathrm{d}\, \mu(\omega)$ and $\int_{\omega \in \xi(B)} v(\rho_B(\omega), \omega) \, \mathrm{d}\, \mu(\omega)$ exist where $\rho_B(\omega)$ is Receiver's optimal action at state ω when the set of available actions is limited to B. We give the proof for the existence of the former integral, $\int_{\omega \in \xi(B)} u(\rho_B(\omega), \omega) \, \mathrm{d}\, \mu(\omega)$. The proof for the existence of the latter integral follows similarly. To prove the existence of the former integral, it suffices to prove that the integrals

$$\int_{\omega \in \left(\bigcap_{x \in B} U_x \right)} u(\rho_B(\omega), \omega) \mathrm{d}\mu(\omega)$$

and

$$\int_{\omega \in \left(\bigcup_{y \notin B} U_y \right)} u(\rho_B(\omega), \omega) \mathrm{d}\mu(\omega)$$

exist. Since both $\left(\bigcap_{x \in B} U_x \right)$ and $\left(\bigcup_{y \notin B} U_y \right)$ are closed sets by Condition i, the existence of the integrals is guaranteed by Condition ii.

[25] An alternative condition, e.g., is that the boundary of the upper contour set has μ-measure zero.

We will show that the function ρ_B is discontinuous only over a μ-measure zero set of points in $\xi(B)$. To this end, let f be the correspondence

$$f(x) = \{a : u(a, x) = \max_{b \in A} u(b, x)\},$$

defined for any $x \in B$. We say a point $x \in \xi(B)$ is *stable* if there exists an open ball around it such that for any point y in that ball $f(x) = f(y)$. To prove the claim, it suffices to show that the set of unstable points in $\xi(B)$ has a μ-measure zero. First, we prove this assuming $|B| = 2$. The proof for $|B| > 2$ follows by induction, as we will see later.

The case of $|B| = 2$. Suppose $B = \{a, b\}$. Define $g(x) = u(a, x) - u(b, x)$ for all $x \in \xi(B)$. Observe that any point x with $g(x) \neq 0$ is stable, by the continuity of $u(a, \cdot)$ and $u(b, \cdot)$. Therefore, to prove the claim, it suffices to show that the set of roots of g has μ-measure zero. This is readily inferred from the fact that g is an analytic function itself, and therefore the set of its roots has Lebesgue-measure zero; see, e.g., Mityagin (2015).

The case of $|B| > 2$. For any two actions a, b, define the function $g_{a,b}(x) = u(a, x) - u(b, x)$. Let $Z_{a,b}$ denote the set of roots of $g_{a,b}$. By an argument similar to the case of $|B| = 2$, $Z_{a,b}$ has μ-measure zero. Let $Z = \cup_{\{a,b : a \neq b, a,b \in A\}} Z_{a,b}$. Since A is finite, Z also has μ-measure zero. The set of stable points is a subset of Z, and therefore has μ-measure zero as well. □

ii.4 The ride-sharing example

Ride-sharing platforms relay information to drivers about excess demand, or the "surge multipler", in order to control the flow of drivers to different areas. This has been practiced by using illustrative *heat maps* (that may or may not vary across drivers); see, e.g., Campbell (2018). The platform's preferences might not satisfy action-monotonicity: e.g., the platform may want to persuade drivers to go to North only when the surge factor is higher than in South, and vice versa.

Formally, suppose there are two areas, indexed by $0, 1$, located at the two extremes of the unit interval. Let $\Omega = \Omega_0 \times \Omega_1$, with $\Omega_i = [\underline{\omega}, \overline{\omega}]$ for $i \in \{0, 1\}$, where $\omega_i \in \Omega_i$ denotes the *surge multipler* at area i. For simplicity, we suppose that Driver (i.e., Receiver) will surely find a ride if she drives to either of the areas, and that her income from the ride is a dollar times the surge multiplier in that area.

Driver is located at some point on the unit interval, and she has to take an action

$a \in \{0, 1\}$, corresponding to driving to area 0 or 1, respectively. Let $c(a)$ denote the cost of taking action a. Driver's payoff from action a is then defined by

$$u(a, \omega) = \beta\omega_a - c(a),$$

where $1 - \beta \in (0, 1)$ is the fraction cut by Platform (i.e., Sender) for its commission fee, and $\omega = (\omega_0, \omega_1)$ is the state of the world.

Platform's payoff function is $v(a, \omega) = (1-\beta)\omega_a$. Therefore, Platform always prefers that Driver drives to the area with the higher surge factor to pick up a ride. Driver's decision, however, also depends on the cost of each action. What is Platform's optimal signal, given that its prior over Ω is μ? As we will see, when Driver is a pessimist, the optimal signal does not depend on the prior, so long as the prior has full support.

To keep both areas relevant in the solution, assume that $|c(0) - c(1)| \leq \beta(\overline{\omega} - \underline{\omega})$. (Otherwise, one of the actions will never be taken by Driver, regardless of the state of the world.) Without loss of generality, suppose $c(0) > c(1)$, and define $\Delta = (c(0) - c(1))/\beta$. Sender's optimal signal is illustrated in Figure 5. There are two signal realizations: one corresponding to the shaded area, and the other corresponding to the rest of the state space. When the state of the world falls in the shaded area, Driver receives a message asserting that the surge multiplier in area 0 is above a predetermined fixed threshold, and also larger than the multiplier in area 1. We explain next why this signal is optimal.

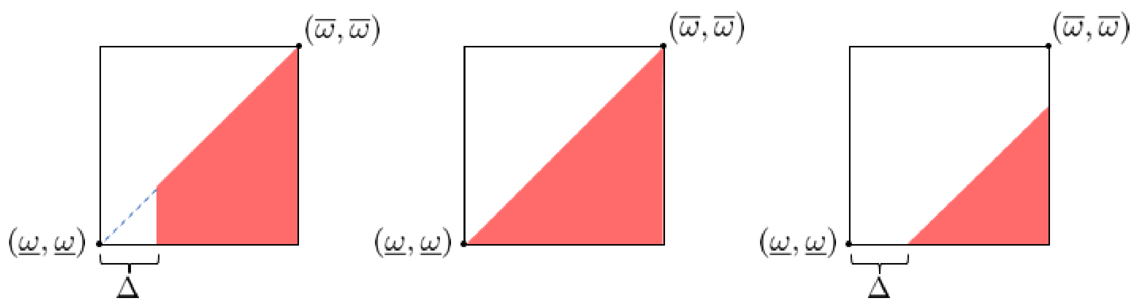

Figure 5: The horizontal and vertical axes respectively correspond to Ω_0 and Ω_1. From left to right: Receiver's action at each state under Sender's optimal signal, the optimal action for Sender at each state, and the optimal action for Receiver at each state. The shaded area corresponds to action 0.

One can verify that Receiver has commutative preferences, since her partial order $\preceq_\Omega$ over Ω is defined by $(x, y) \preceq_\Omega (x', y')$ if and only if $x \leq x'$ and $y \leq y'$. Therefore, the greedy algorithm can construct a signal. The optimality of the greedy signal is guaranteed by Proposition ii.3, which is a counterpart to Proposition ii.1 but for the cardinal preference model. When the algorithm is run on this example, it first finds two representative states, namely $\omega_0 = (\Delta, 0)$ and $\omega_1 = (0, 0)$, corresponding to actions $0, 1$, respectively. Then, each other state ω is assigned to one of the representative states: to the state ω_a corresponding to the action a that maximizes $v(a, \omega)$, subject to the constraint that $\omega_a \preceq_\Omega \omega$. This partitions the set of states into two subsets, producing the optimal signal for Sender, as illustrated in Figure 5. Observe that the convexity of the optimal signal (with respect to $\preceq_\Omega$) does not hold here, unlike in Theorem 3.1 where action-monotonicity holds.

iii Example: Continuous state space and action space

We extend the example from Section 2 by allowing the state space (i.e., the set of project values) and the action space (i.e., the set of effort levels) to be continuous. In particular, we let the project value belong to the set $V = [\underline{v}, \overline{v}]$. Also, we let the agent's effort level belong to the set $E = \mathbb{R}_+$.

The state of the world (i.e., the project value) is drawn from an atomless distribution μ with support V. A signal π is a measurable mapping $\pi : V \to \Delta(S)$, for some signal realization space S. When the agent receives a signal realization s, he chooses an effort level $\hat{e}(s)$, as defined by (2.2). The *value* of a signal π for the principal is defined, as in Section 2, by $\mathbb{E}_{v\sim\pi}\mathbb{E}_{s\sim\pi(\cdot|v)}\left[u_p(\hat{e}(s), v)\right]$. A signal π^* is *optimal* if no other signal has a higher value. The principal's objective is finding the optimal signal.

We can find the optimal signal under the same assumptions made in Section 2, namely, state- and action-monotonicity (i.e., Assumptions 2.2 and 2.3). Under these assumptions, the greedy algorithm from Section ii.1 finds the optimal signal, by Proposition ii.3. We use the greedy algorithm, instead of the algorithm used in Section 2, to handle the technicalities involved when the state space is continuous. Intuitively, however, these two algorithms are similar. This is demonstrated in Figure 6, as discussed next.

Let $e^*(v)$ denote the effort level that maximizes the agent's utility when the project has value v. (It is commonly known that, in such principal-agent models, the function e^* may be non-monotone.) Figure 6 demonstrates the signal constructed by the greedy algorithm for a particular e^*. The optimal signal fully reveals the project value to the agent if the value belongs to any of the intervals $[\underline{v}, v_1]$, $[v_2, v_3]$, or $[v_4, \overline{v}]$. Otherwise, the value belongs to the interval $[v_1, v_2]$ or $[v_3, v_4]$; in that case, only the interval that contains the project value is revealed to the agent.

The formal proof of optimality is by Proposition ii.3. To see the intuition, suppose that the interval $[\underline{v}, \overline{v}]$ is discretized using an arbitrary finite set of points $\tilde{V}$ that contains v_1, v_2, v_3, v_4. Then, assuming that the project value belongs to $\tilde{V}$, run the algorithm that constructs the optimal signal in Section 2. The constructed signal would be similar to the above signal: it fully reveals the project value to the agent if the value belongs to either of the intervals $[\underline{v}, v_1]$, $[v_2, v_3]$, or $[v_4, \overline{v}]$. Otherwise, the value belongs to the interval $[v_1, v_2]$ or $[v_3, v_4]$; in that case, only the interval that contains the project value is revealed to the agent.

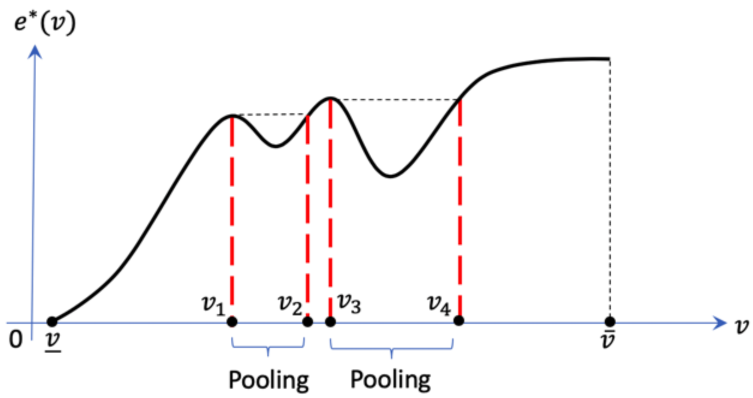

Figure 6: The optimal signal pools the project values belonging to $[v_1, v_2]$ together, and those belonging to $[v_3, v_4]$ together.