Persuading a Pessimist

Afshin Nikzad*

PRELIMINARY DRAFT

Abstract

Consider a (Bayesian) persuasion problem in which a pessimistic Receiver takes action to secure a "guaranteed payoff". Through the lens of an ordinal utility model, we demonstrate the simple structure and the strong robustness properties of the "optimal" signal. In particular, we show that when simple monotonicity conditions on Sender's and Receiver's preferences hold (essentially, when the action space and state space are ordered), a *stochastically dominant* signal exists, i.e. a signal that imposes a distribution on the realized outcome which stochastically dominates the distribution imposed by any other signal, with respect to Sender's preferences. Furthermore, we show that the stochastically dominant signal has a simple structure and satisfies strong robustness properties. In particular, the signal (i) is monotone, i.e. induces higher actions at higher states, (ii) is a partitional signal that pools "adjacent" elements, and (iii) typically has a small size, i.e. partitions the state space into a small number of subsets. The signal is robust to (i) lack of commitment power from Sender, and (ii) (mis)specification of the prior. To detect which of these properties hinge on the sufficient monotonicity conditions, we also study their relaxations. In particular, we show that even when the action space is not ordered, stochastically dominant signals still exist if the state space is compact and the prior is atomless. The structural properties of the signal might not continue to hold while the robustness properties do.

*Department of Economics, Stanford University. Email: nikzad@stanford.edu

Contents

1 Introduction

Bayesian Persuasion has become a standard framework for modeling information transmission: Sender and Receiver share the same prior on 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 her posterior after receiving the signal realization.

We focus on optimal persuasion of a pessimistic Receiver who aims to take an action that secures a "guaranteed payoff". Through the lens of an ordinal utility model, we demonstrate the simple structure and the strong robustness properties of the "optimal" signal.

Receiver has a preference list 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 list over the outcomes and evaluates signals by the probability distribution that they impose on the realized outcome. How does Sender choose her favorite signal? One of the questions that we answer is about the existence of stochastically dominant signals, i.e. a signal that imposes a distribution on outcomes which stochastically dominates the distribution imposed by *any* other signal (where the stochastic dominance relation is defined with respect to Sender's preferences over the outcomes).

We discuss our findings in detail in the paragraphs to follow; briefly, we show that stochastically dominant signals (i) exist when Sender's and Receiver's preferences satisfy simple monotonicity conditions, (ii) have a simple structure, making them easy to relay, e.g., they form partitional signals that pool "adjacent" elements and typically partition the state space into a small number of subsets, and (iii) satisfy strong robustness properties, e.g. they require much weaker commitment power from Sender than in Bayesian Persuasion. To detect which of the properties hinge on the sufficient monotonicity conditions, we also study their relaxations. We will also see that, naturally, the properties of stochastically dominant signals also hold in the classic cardinal utility model where Sender's choice of signal is the one that maximizes her (cardinal) utility.

The sufficient monotonicity conditions for the existence of stochastically dominant signals are *action-monotonicity* and *state-monotonicity*. By action-monotonicity, we mean 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. By state-monotonicity, we mean 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. When both these conditions hold, stochastically dominant signals exist. (Section 2)

"In many environments that fit into the persuasion framework, the action space and state space are ordered" [Mensch 2018], and the sufficient 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, as in [Dworczak and Martini 2018]. (See [Kamenica and Gentzkow 2011, Mensch 2018], among others, for other examples.)

The existence of stochastically dominant signals, however, does not hinge on the monotonicity conditions. We show that state-monotonicity could be relaxed by allowing Receiver to have an incomplete preference relation (e.g., a lattice) over the states. Relaxing state-monotonicity makes our framework also applicable to multidimensional persuasion problems. (Section 3.1) We also show that action-monotonicity can be dismissed when the state space is compact and the prior is atomless: stochastically dominant signals would still exist and some of their properties would be preserved. This is demonstrated in the context of an example in ride-sharing platforms where action-monotonicity fails to hold. (Section 3.2)

We prove the existence of stochastically dominant signals by constructing them. The construction also reveals the simple structure of the signals: they are "convex", "monotone", and "small", as discussed next. The signal turns out to be a partition of the state space, with each signal realization being a subset of the state space which is *convex* with respect to Receiver's possibly incomplete preference relation over the states. The signal is convex in the sense that if two states belong to the same signal realization, then so does any other state that lies in between them (with respect to Receiver's preferences over the states). When Receiver's preference relation is complete, stochastically dominant signals partition the state space into "intervals". This type of signal structure has received attention in the literature also for its practical implications. For example, [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.

When the action-monotonicity and state-monotonicity conditions hold, one can define the notion of a "monotone" signal, as defined by [Mensch 2018], to be a signal that does not recommend lower actions at higher higher states. The practical implications of such signals are discussed in the same paper. Stochastically dominant signals are monotone in that sense, and also in a broader sense when the state space is partially ordered, as in Section 3.

We also analyze the size (number of signal realizations) of stochastically dominant signals in stochastic problem instances. We show that stochastically dominant signals are typically small under the assumptions of Section 4. The analysis heavily relies on our signal construction method from Section 2 and the Principle of Deferred Decisions in probabilistic analysis.

Stochastically dominant signals also satisfy strong robustness properties, namely *No commitment to the signal choice*, and *Independence from the prior*. By the No commitment to the signal choice property, Sender has no incentive to change her signal after the state of the world is realized, so long as she abides sending *correct* signal realizations, i.e. subsets of the state space that contain the true state of the world. In other words, Sender needs no commitment power to the signal choice, so long as she commits to sending correct signal realizations. This is a milder commitment assumption than the one in Bayesian Persuasion where signal realizations are sent randomly. For instance, monitoring the correctness of a signal realization (whether it contained the true state of the world) is simpler than verifying whether the randomization for sending a signal realization was done correctly; to prove the possible incorrectness of the signal realization in the former case, the only evidence needed is the signal realization itself.

By the *Independence from the prior* property, stochastically dominant signals do not change with the prior, so long as the prior has full support (i.e. does not exclude any state). The intuition for both of the above properties is provided in Section 2, after the signal construction method is discussed.

To sum up, the above properties indicate that when the state-monotonicity and action-monotonicity conditions hold, Sender's signal is robust and has a simple structure. It is robust to (i) (mis)specification of her cardinal utilities, so long as her ordinal preferences over the outcomes do not change, (ii) lack of commitment power from Sender, and (iii) (mis)specification of the prior. Sender's signal has a simple structure, in the sense that it is "convex", "monotone", and "small", as discussed above.

All of these properties continue to hold when we relax state-monotonicity in Section 3.1. In the absence of action-monotonicity, the robustness properties continue to hold, while the structural properties may not. In particular, monotonicity of the signal cannot be defined, and also, its convexity may not hold, as in the ride-sharing example of Section 3.2. Thereby, the convexity could be counted as a consequential feature of action-monotonicity.

The significance of our constructive approach is not limited to understanding these properties. Our methods can be adapted to solve constrained persuasion problems. We provide an example in Section 3.2, where we find Sender's optimal signal in the presence of an upper quota constraint on the number of signal realizations. Such constraints can be present in applications where simplicity of the signal is a concern.

The rest of the paper is organized as follows. Section 1.1 contains a few preliminary definitions. The basic model appears in Section 2. Section 2.1 highlights the connections between the ordinal utility model and the classic cardinal utility model. We relax the state-monotonicity condition in Section 3.1 and dismiss the action-monotonicity condition in Section 3.2. Section 4 analyzes the expected size of stochastically dominant signals in stochastic problem instances. Section 5 reviews some of the related literature. All the proofs and missing arguments appear in the appendix.

1.1 Preliminary definitions

Preference relations

For any transitive preference relation $\preceq$ over a set, we use $\prec, \sim$ respectively to denote the strict preference and the equivalence preference relations imposed by $\preceq$ over that set. For any two elements x, y such that $x \preceq y$, we say that y is $\preceq$-higher than x.

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. Similarly, let $\sup^{\preceq}\{X\}$ denote $x \in X$ such that $y \preceq x$ holds for all $y \in X$.

We use the words preference relation and *order* interchangeably. A set P is *convex* with respect to a partial order $\preceq$ if for any $w, w', w'' \in P$ for which $w \preceq w' \preceq w''$ holds, $\omega, \omega'' \in P$ implies $\omega' \in P$. A family of sets is called *convex* with respect to $\preceq$ if all sets in the family are convex with respect to $\preceq$.

Asymptotic notions

We say a positive function $f(x)$ is *of the order of* a positive function $g(x)$ if $f(x)/g(x)$ approaches a constant (possibly zero) as x approaches infinity.

2 The basic setup and results

There is finite set of states, Ω, and a finite set of actions, A. Any element of $A \times \Omega$ is called an *outcome*. Sender and Receiver have complete and transitive preference relations over the set of outcomes, respectively denoted by $\preceq_s, \preceq_r$.

The state of the world, ω_0, is drawn from a prior $\mu \in \text{int}(\Delta(\Omega))$,[1] henceforth, *the prior*. Knowing the prior but not the state of the world, Sender chooses a *signal*, i.e. an arbitrary partition of Ω, namely $\pi = \{P_1, \ldots, P_n\}$. A *signal realization* $P \in \pi$ is then sent to Receiver such that $\omega_0 \in P$.[2] Receiver then takes an action $a^*(\pi, P)$ that *guarantees the best worst-case realized outcome*, in the following sense: for any $a \in A$,

$$\inf{}^{\preceq_r}\{(a,\omega) : \omega \in P\} \preceq_r \inf{}^{\preceq_r}\{(a^*(\pi,P),\omega) : \omega \in P\}$$

holds where recall that $\inf^{\preceq}\{X\}$ denotes $x \in X$ such that $x \preceq y$ holds for all $y \in X$.

Taking Receiver's objective into account, any signal π chosen by Sender imposes a distribution η_π over the set of realized outcomes. Sender's objective is choosing a *stochastically dominant* signal, i.e. a signal π^* such that η_{π^*} *stochastically dominates* η_π for any other signal π, in the following sense: Let the set of outcomes be ordered as $o_1, \ldots, o_m$ from Sender's most favorite to her least favorite. Let p_i, p_i^* respectively be the probability assigned to outcome o_i in the distributions η_π, η_{π^*}. We say η_{π^*} stochastically dominates η_π iff for any positive integer $j < m$ with $o_{j+1} \prec_s o_j$, $\sum_{k=1}^{j} p_k^* \geq \sum_{k=1}^{j} p_k$.

In general, stochastically dominant signals may not always exists. Their existence, however, is guaranteed when Sender's preferences are *monotone in actions* and Receiver's preferences are *monotone in states*. We say Sender's preferences are monotone in actions when there exists an ordering $\preceq_A$ over actions such that whenever $a \preceq_A a'$, then $(\omega, a) \preceq_s (\omega, a')$ holds for any $\omega \in \Omega$. Receiver's preferences are monotone in states when there exists an ordering $\preceq_\Omega$ over the states such that whenever $\omega \preceq_\Omega \omega'$, then $(\omega, a) \preceq_r (\omega', a)$ holds for any $a \in A$. For brevity, we sometimes refer to the former condition as *action-monotonicity* and the latter condition as *state-monotonicity*. We say Sender and Receiver have *monotone preferences* when both of these conditions hold.

We need one more definition before stating the theorem. A *monotone signal* is a signal that recommends $\preceq_A$-higher actions at $\preceq_\Omega$-higher states.

Theorem 2.1. *When Sender and Receiver have monotone preferences, there exists a stochastically dominant signal π^*. Moreover, π^* is a monotone signal, and also is convex with respect to $\preceq_\Omega$.*

[1] The assumption that the prior is an interior point is without loss of generality, as Both Sender and Receiver would ignore the states outside of prior's support.

[2] The assumptions that the signal space is the state space, signal realizations partition the state space, and are sent deterministically are without loss of generality, due to Receiver's objective.

We prove the existence of stochastically dominant signals by constructing them. Our constructive approach also reveals the following properties of the stochastically dominant signals:

(i) **No commitment to the signal choice.** Sender has no incentive to change the stochastically dominant signal after the state of the world is realized, so long as she abides sending *correct* signal realizations, i.e. subsets of the state space that contain the true state of the world. In other words, Sender needs no commitment power to the signal choice, so long as she commits to sending correct signal realizations.

(ii) **Independence from the prior.** The stochastically dominant signal does not depend on Sender's prior.

We emphasize that the *No commitment to the signal choice* property carries an implicit commitment assumption itself: that Sender commits to not sending false signal realizations. This is a milder commitment assumption than the one used in the classic Bayesian Persuasion framework where signal realizations are sent randomly. For example, verifying whether the sent signal realization has been correct (contained the true state of the world) is easier than verifying whether the randomization in sending the signal realization has been done correctly, as the only evidence required for verification in the former case is the sent signal realization itself.

To observe why the above properties hold, we first sketch the algorithm that constructs the stochastically dominant signal. This algorithm will also play a key role later in addressing the question of how large stochastically dominant signals can be (Section 4). To describe the algorithm, we need one definition: an action a is *inducible* by a subset of states if there exists a state in that subset such that Receiver's chosen action at that state is a.

The algorithm keeps track of a set of *selected actions* and a set of *covered states*, both initially empty. At each iteration, it chooses the most preferred unselected action with respect to $\preceq_A$ which is inducible by the set of uncovered states: namely, action a. It then finds the least preferred state with respect to $\preceq_\Omega$ that induces a: namely, state ω. At the end of the iteration, a is added to the set of selected states and a signal realization is added to π^* that contains all uncovered states ω' such that $\omega \preceq_\Omega \omega'$. The states in this signal realization are also added to the set of covered states. The algorithm reiterates until all states are covered. (See Section C in the appendix for a formal definition.)

One can verify that constructed signal, namely π^*, is a stochastically dominant signal through property (i) mentioned above. To see why this property holds, consider an arbitrary state ω and an arbitrary subset of the states $Q \ni \omega$. Furthermore, suppose ω belongs to the signal realization P in π^*. We will show that Sender weakly prefers the action induced by P to the action induced by Q.

Let $\delta = \inf^{\preceq_\Omega}\{Q\}$. First, verify that if δ is covered before P is constructed by the algorithm, so should be ω: since $\omega \in Q$, we have $\delta \preceq_\Omega \omega$, and therefore, ω should have been covered when δ was, by our construction of the signal realizations. This is a contradiction. Therefore, suppose that at the time of constructing P, δ is uncovered. In that case, if the action induced by Q is strictly preferred by Sender to the action induced by P, the algorithm should have chosen the action induced by Q instead of the action induced by P, since the action induced by Q is $(\Omega \backslash \Psi)$-inducible as well. This would be a contradiction. Consequently, Sender weakly prefers the action induced by P to the action induced by Q.

Property (ii), convexity, and monotonicity of the stochastically dominant signal hold by construction; the complete proof is in the appendix. We remark that convexity of the signal could be counted as a consequence of action-monotonicity: in Section 3.2, we see that the signal may be not convex when action-monotonicity is dismissed. There, we provide a different approach for constructing the signal in the absence of action-monotonicity.

2.1 Cardinal utilities and connection to Bayesian Persuasion

A natural special case of our model is a cardinal utility model. This special case fits into a slightly more general version of the Bayesian Persuasion framework of [Kamenica and Gentzkow 2011]. Since some of our examples are set up in the cardinal utility model, we present the formal model for completeness.

A signal π consists of a finite realization space S and a family of distributions $\{\pi(\cdot|\omega)\}_{\omega\in\Omega}$ over S, i.e. the usual signal structure in Bayesian Persuasion. Sender has a prior μ over Ω, which is shared by Receiver. Sender chooses a signal π. After the state of the world is realized, a signal realization s is sent to Receiver based on the true state of the world. Then, Receiver forms a posterior μ_s using the Bayes's rule and takes the action $a^*(s) = \arg\max_{a\in A} u(a, \mu_s)$, where $u(a, \eta)$ denotes her utility function when she takes action a and her posterior is η. Sender chooses the signal π that maximizes her expected utility, $\mathbb{E}_{w\sim\mu}\mathbb{E}_{s\sim\pi(\cdot|\omega)}[v(a^*(s), \omega)]$. Sender's choice in the cardinal utility model is called the *optimal* signal. (We emphasize the difference between the optimal signals in the cardinal utility model and the stochastically dominant signals in the ordinal utility model: e.g, the former always exists, but the latter not.)

With a risk-neutral Receiver, i.e. $u(a, \eta) = \mathbb{E}_{\omega\sim\eta}[u_0(a, \omega)]$ for some function $u_0 : A\times\Omega \to$

$\mathbb{R}$, the above model coincides with the Bayesian Persuasion model introduced by [Kamenica and Gentzkow 2011].[3] For a pessimistic Receiver, we suppose

$$u(a,\eta) = \inf_{\omega \in \mathrm{supp}(\eta)} u_0(a,\omega),$$

for some function $u_0 : A \times \Omega \to \mathbb{R}$.

Action-monotonicity and state-monotonicity will be defined shortly, by two orderings $\preceq_A, \preceq_\Omega$, similar to the ordinal model. When the monotonicity conditions hold, the optimal signal is convex with respect to $\preceq_\Omega$, and satisfies the robustness and the structural properties that we discussed in Section 2.

Action-monotonicity and state-monotonicity in the cardinal utility model are defined in the natural way: Sender has monotone preferences over the actions when there exists an ordering $\preceq_A$ over actions such that $v(a,\omega) < v(a',\omega)$ holds for any state ω when $a \preceq_A a'$. Receiver has monotone preferences over the states when there exists an ordering $\preceq_\Omega$ over the states such that $u(a,\omega) < u(a,\omega')$ holds for any action a when $\omega \preceq_\Omega \omega'$.

We discuss an example of the cardinal utility model in the appendix (Section A) in the context of an ad exchange platform revealing information about the customer's type to advertisers.

3 Relaxing the monotonicity conditions

Action and state-monotonicity might not hold in some applications. For example, state-monotonicity fails to hold when a principal is relaying information to an agent about both the "reward" and the "riskiness" of a project, or action-monotonicity could fail to hold when a ride-sharing platform is trying to persuade a driver to drive north or south by providing information about the surge factors in each area.

The main question of interest in this section is how relaxing or dismissing one of these conditions can affect the existence of stochastically dominant signals. In Section 3.1 we introduce the *commutativity*, a relaxation of the state-monotonicity condition, which allows Receiver to have partial preferences over the states, e.g., a lattice.

In Section 3.2, we show that in the presence of commutativity, action-monotonicity can be dismissed when the state space is compact. This is done using a new signal construction approach. The constructed signal is "almost" stochastically dominant when the probability

[3] It should be pointed out that the main results of [Kamenica and Gentzkow 2011] still hold when the risk-neutrality assumption is dismissed.

that the prior assigns to a "small" subset of the states is "small"; in particular, the signal is stochastically dominant when the state space is compact and the prior is atomless.

3.1 Relaxing state-monotonicity

Recall from Section 2 that state-monotonicity requires the existence of a complete order $\preceq_\Omega$ over the states. Here, we allow $\preceq_\Omega$ to be a partial order and relax the state-monotonicity assumption to *commutativity*, which turns out the be equivalent to state-monotonicity when $\preceq_\Omega$ is a complete order. We will provide a counterpart to Theorem 2.1 by constructing a stochastically dominant signal. The signal will be monotone, convex with respect to Receiver's partial order, and will also satisfy *No commitment to the signal choice* and *Independence from the prior* properties, which were discussed earlier in Section 2.

In the rest of this section, we use $\leq_\Omega$ to denote a partial order over Ω which is a lower semilattice.[4] For any subset $P \subseteq \Omega$, let $\bigwedge P$ denote the greatest lower bound for P with respect to $\leq_\Omega$. We say Receiver has *commutative* preferences when, for any signal realization P, Receiver's optimal choice is determined by $\bigwedge P$, in the sense that $\mathfrak{a}^*(P) = \mathfrak{a}^*(\bigwedge P)$. Note that state-monotonicity implies commutativity.

One of the simplest examples of commutative preferences is when $\Omega = \Pi_{i=1}^{|A|}\Omega_i$ provides information about Receiver's actions, with Ω_i providing information about action $i \in A$. For example, consider a principal-agent problem where Receiver is an agent whose actions correspond to working on one of the n available projects, with Ω_i providing information about project i, e.g., a "risk" or "quality" index in which Receiver's utility is monotone. Receiver's preferences then satisfy commutativity but not state-monotonicity.[5]

The choice function of a rationally bounded agent with non-commutative preferences also may satisfy commutativity. When solving the max-min optimization problem (i.e. maximizing over actions, minimizing over the signal realization P) is difficult for Receiver, choosing an action that maximizes her utility at state $\bigwedge P$ provides her with a guaranteed level of utility through a cognitively simpler task.

Theorem 3.1. *When Receiver's preferences are commutative and Sender's preferences are monotone in actions, there exists a stochastically dominant signal, which is also a monotone signal and convex with respect to $\leq_\Omega$.*

[4] 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, any lattice is a lower semilattice.

[5] Observe that for any action i and any subset of states S, $\min_{\omega\in S} u(i,\omega) = u\left(i, \min_{w\in S}(\omega)\right)$, where the min operator on the right-hand side is the component-wise minimum. This implies commutativity.

The stochastically dominant signal still satisfies *No commitment to the signal choice* and *Independence from the prior* properties from Section 2. The intuition remains the same as before. In Section 4, we also analyze the expected size of the signal for a special class of commutative preferences, *product orders*, which also contain the above principal-agent example, and show that the signal is small under the assumptions of that section.

The proof of Theorem 3.1 follows a signal construction approach similar to the one we discussed for Theorem 2.1: At each step, the most preferred unselected action inducible by the uncovered states is selected, and then a signal realization is constructed that contains all the states weakly preferred to that state by $\leq_\Omega$. The states in the signal realization are then added to the set of covered states and the process is repeated until all states are covered.

3.2 Dismissing action-monotonicity

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 that the prior assigns to any "small" subset of the states is "small". In particular, our construction turns out to be optimal when the state space is compact and the prior is atomless. We then provide some insight on this new approach by looking at an example in the context of ride-sharing platforms, where action-monotonicity fails to hold.

3.2.1 The greedy signal

We provide a simple *greedy* approach that constructs a signal for Sender when Receiver's preferences are commutative, as defined in Section 3.1. 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 below.

Algorithm: Greedy construction of the signal

Initialize π, R to the empty set.
For any action a, let Ω_a be the subset of states that induce a.
For any action a with $\Omega_a \neq \varnothing$, define $\omega_a = \bigwedge \Omega_a$, add ω_a to R, and let $P_a = \{\omega_a\}$.
For any state $\omega \notin R$, define a_ω by letting $(a_\omega, \omega) = \sup^{\preceq_s}\{(a, \omega) : a \in A,\ \omega_a \leq_\Omega \omega\}$.
For any action a with $\Omega_a \neq \varnothing$, add $\{\omega : a = a_\omega\}$ to P_a.
Define the greedy signal as $\pi_G = \bigcup_{a \in A : P_a \neq \varnothing} \{P_a\}$.

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. It will construct 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 $\omega_{a'} \leq_\Omega \omega$. Finally, each representative state together with the states assigned to it are defined as a signal realization in the greedy signal, π_G.

By construction, at any non-representative state, Sender takes her most preferred action that could be taken at that state in any signal. This holds essentially because any ω that satisfies the constraint $\omega_a \leq_\Omega \omega$ in the maximization problem in line 4 of the algorithm cannot be in any signal realization that induces an action b for which $\omega_b \leq_\Omega \omega$ does not hold. (A consequence of commutativity.) This is not necessarily the case for representative states, and therefore, the algorithm does not always construct a stochastically dominant signal. However, if the probability that the prior assigns to the representative states is "negligible", then π_G should be "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.

3.2.2 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. a member of the partition being a fat contour set). To rule out some of the 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 η_π-measurable. 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 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,

$$\int_{p \preceq_s q} 1 \, \mathrm{d}\, \eta_\pi(q) \geq \int_{p \preceq_s q} 1 \, \mathrm{d}\, \eta_{\pi'}(q).$$

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

Proposition 3.2. *When the prior is atomless, if the greedy signal is proper, it is also stochastically dominant.*

By construction, the stochastically dominant signal satisfies *No commitment to the signal choice* and *Independence from the prior* properties that we saw in Section 2. Unlike there, the signal is not necessary convex, as we will see in the ride-sharing example of Section 3.2.3. Therefore, convexity of the stochastically dominant signal can be seen as a consequential feature of action-monotonicity.

In a cardinal utility model, our definition of properness boils down to a simple one: integrability of Sender's utility function. In general, Sender's (expected) utility under an arbitrary signal may not be "well-defined" when the state space is compact, i.e. the integral over the state space with respect to the prior may not exist. In a cardinal utility model, our definition of "properness" only requires integrability of Sender's utility function with respect to the prior. The details are discussed in the appendix, Section D.3. There, we also discuss some of the mild sufficient conditions that ensure the properness of the greedy signal, i.e. the integrability of Sender's utility function under the greedy signal. The example in Section 3.2.3 provides some insight.

Our approach 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 solving a series of subproblems: for each subset of the representative states of the allowed size, namely $S \subseteq R$, consider a signal with $|S|$ signal realizations, each containing one of the states in S. Each other state is then added to one of the $|S|$ signal realizations, specifically the one preferred the most by Sender. It can be shown that if the signal produced in one of these subproblems stochastically dominates the other produced signals, it is also a stochastically dominant signal. In a cardinal utility model, Sender's optimal signal, which always exists, is her most preferred signal produced in the subproblems.

It is worth pointing out that the problem of finding the optimal signal with a size limit can be reduced to an optimization problem known as the *Generalized maximum coverage* problem in the computer science literature [Wiki-MCP 2017], where computationally "fast" greedy-based methods are developed for finding approximately optimal (that is, within a constant factor of the optimal) solutions.

Finally, we remark that in the presence of action-monotonicity, it can be shown that the greedy algorithm constructs the stochastically dominant signal whether the state space is discrete or compact. The greedy construction, however, reveals less about the structural properties of stochastically dominant signals (such as their convexity or their size) than the construction approach that we used in Section 2.[6]

3.2.3 The ride-sharing example

Ride-sharing platforms relay information to drivers about the 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) [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 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 the platform for commission fee, and $\omega = (\omega_0, \omega_1)$ is the state of the world.

Platform's payoff function is simply $v(a, \omega) = (1 - \beta)\omega_a$. 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 its prior μ over the states?

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 1. There are two signal realizations: one corresponding to the shaded area, and the other one 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

[6] The construction approach of Section 2 is also used in Section 4 to analyze the expected size of stochastically dominant signals.

surge multiplier in area 0 is above a predetermined fixed threshold, and also larger than the multiplier in area 1. We explain why this signal is optimal next.

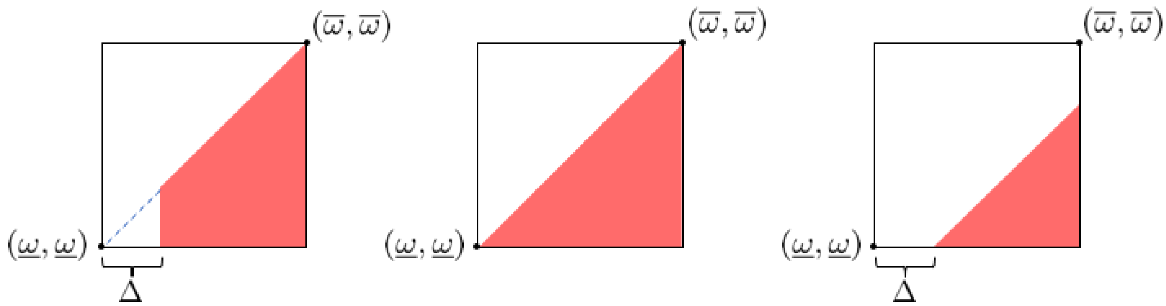

Figure 1: The horizontal and vertical axes respectively correspond to Ω_0 and Ω_1. From left to right: Receiver's action at each state under Sender's optimal signal, the optimal action for Sender at each state, and the optimal action for Receiver at each state. The shaded area corresponds to action 0.

One can verify that Receiver's preferences are commutative with the partial order $\leq_\Omega$ defined by $(x, y) \leq_\Omega (x', y')$ if $x \leq x'$ and $y \leq y'$. Therefore, the greedy algorithm can be used to construct a signal. The optimality of the greedy signal is guaranteed by Proposition D.2 in the appendix, which is a counterpart to Proposition 3.2 but for the cardinal utility 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 $\omega_a \leq_\Omega \omega$. This partitions the set of states to two subsets, producing the optimal signal for Sender, as illustrated in Figure 1. Observe that the convexity of the optimal signal (with respect to $\leq_\Omega$) does not hold here, unlike in Theorems 2.1 and 3.1 where action-monotonicity holds.

4 Size of the signal

In this section we analyze the size of the stochastically dominant signal in stochastic instances. First, this is done in the presence of both action-monotonicity and state-monotonicity. Then, we will see that similar results hold when state-monotonicity is replaced with commutativity. The main finding is that the *alignment between Receiver's choice of actions and Sender's preferences over actions* plays an important role in determining the size: even small levels of misalignment could lead to signals with significantly smaller size. First, we consider an extreme point and give the necessary and sufficient condition for the stochastically dominant signal to have the largest possible size, $|\Omega|$. A signal of this size is also called the *fully revealing signal*.

For ease of exposition, in this section we suppose that Receiver has strict preferences over the outcomes. The qualitative insights from the next proposition do not hinge on this assumption, and the rest of the propositions hold identically without this assumption.

Proposition 4.1. *The fully revealing signal is the unique stochastically dominant signal if and only if* $\mathfrak{a}^*(\omega) \prec_A \mathfrak{a}^*(\omega')$ *when* $\omega \prec_r \omega'$.

Proof. The proof is by contradiction. Suppose any stochastically dominant signal, namely π^*, has size $|\Omega|$, and suppose there exists $\omega \prec_r \omega'$ such that $\mathfrak{a}^*(\omega') \preceq_A \mathfrak{a}^*(\omega)$. Define the signal π' as

$$\pi' = \pi \backslash \{\{\omega\}, \{\omega'\}\} \bigcup \{\{\omega, \omega'\}\},$$

and observe that either $\eta_{\pi'} = \eta_\pi$ (which contradicts the uniqueness) or $\eta_{\pi'}$ stochastically dominates η_π. Contradiction. □

The sufficient and necessary condition in Proposition 4.1 is that Sender's preference list over actions is the same as the list of Receiver's optimal actions at each state when the states are ordered with respect to $\preceq_r$. For brevity, we will refer to this condition as *Receiver's choices and Sender's preferences over actions are aligned*.

On one extreme, for the stochastically dominant signal to have the largest possible size, Receiver's choices and Sender's preferences over actions should be aligned. The other extreme is when they are fully misaligned (i.e. $\mathfrak{a}^*(\omega') \prec_A \mathfrak{a}^*(\omega)$ when $\omega \prec_r \omega'$); then, the unique stochastically dominant signal is the uninformative singleton signal.

Next, we consider a middle ground where Receiver's choices and Sender's preferences over actions are, in a sense, *independent*. We show that in this case as well, the size of the stochastically dominant signal is quite small. Then, we will show that even "small levels of misalignment" shrink the size of the stochastically dominant signal significantly.

First, we need a definition. Recall that when Senders' preferences are monotone in actions, $\preceq_s$ imposes an order $\preceq_A$ on the set of actions. When the preference relation of Sender, $\preceq_s$, is a random variable, then so is $\preceq_A$.

Definition 4.2. *Suppose $\preceq_s$ is drawn from a distribution $\mathcal{D}$ (which may depend on $\preceq_r$). We say Sender has* random independent preferences over actions *under $\mathcal{D}$ when $\preceq_A$ is a random variable independent of $\preceq_r$ and is distributed uniformly at random over the set of all permutations over A.*

When $\mathcal{D}$ is clearly known from the context, we briefly say Sender has *random independent preferences over actions*. As an example, consider a cardinal utility model where Sender's utility only depends on Receiver's action (and not to the state of the world). If Sender's payoff from each action is i.i.d. across actions, then Sender has random independent preferences over actions. When Sender has random independent preferences, the stochastically dominant signal, π^*, is small.

Proposition 4.3. *Suppose Sender has random independent preferences over actions under distribution $\mathcal{D}$. Then, $\mathbb{E}_{\mathcal{D}}[|\pi^*|] \leq 1 + \log_2 |\Omega|$, where π^* is the (smallest) stochastically dominant signal.*

The proof builds on our construction of stochastically dominant signals in Section 2. To compute the expected size of the signal, suppose, by the Principle of Deferred Decisions, that Sender's preferences are generated (independently) while the algorithm is run. Therefore, in the first iteration of the algorithm, one can suppose that Sender chooses her most favorite action among the set of inducible actions uniformly at random, which would just be the first action selected by the algorithm (to be induced by a signal realization). Let this action be called a.

It can then be verified that by choosing a uniformly at random, about half of the states will be covered (that is, added to the first signal realization) by the algorithm, in expectation: each state $\omega \in \Omega$ induces some action, and the probability that an action is chosen as a is proportional to the number of states inducing it. The constructed signal realization will have the smallest expected size when each state induces a distinct action. In that case, choosing a uniformly at random from the set of inducible actions is equivalent to choosing a state ω uniformly at random from Ω. By our construction, all the states higher than ω will be added to the first signal realization. The expected number of such states is at least $|\Omega|/2$. Similarly, we can show that each time an action is selected by the algorithm, in expectation half of the states are covered, which leads to the claimed bound.

In addition to showing that the stochastically dominant signal is small in expectation, we can also show that it is small with high probability, i.e. its size is of the same order of magnitude as its mean, with high probability. The proof for the latter is omitted.

Proposition 4.3 requires random independent preferences for Sender. The uniformity assumption on $\preceq_A$ is a mild assumption; e.g., it can be satisfied by relabeling the actions uniformly at random. However, together with independence of $\preceq_A$ from $\preceq_r$, they form a strong assumption. Next, we demonstrate that when the independence assumption is replaced with (even large levels of) positive (or negative[7]) correlation, the stochastically dominant signal would still be small. This demonstration will show that even small misalignments between Receiver's choices and Sender's preferences over actions can shrink the size of the stochastically dominant signal significantly.

We model such misalignments via two different approaches in two examples, one with ordinal and the other with cardinal utilities. Both examples are parameterized by a single parameter. At one extreme value of the parameter, the examples corresponds to the setting of Proposition 4.1 where Receiver's choices and Sender's preferences over actions are aligned. At the other extreme, they correspond to the setting of Proposition 4.3 with random independent preferences. The examples take different approaches in parametrizing Sender's preferences: the first example through a discrete choice model and the second example through a noisy utility function.

Discrete choice model. Suppose $\Omega = \{\omega_1, \dots, \omega_n\}$, $A = \{a_1, \dots, a_n\}$, and that Receiver's preferences satisfy $\mathfrak{a}^*(\omega_i) = a_i$ for all $\omega_i \in \Omega$. Sender's preferences over outcomes depend only on actions. Let Sender's preferences over actions, $\preceq_A$, be a random variable drawn from a discrete choice model with weight β_a for any action $a \in A$,[8] and let $\mathcal{D}$ denote its distribution. For any action a_i, let $\beta_{a_i} = \theta \log(i)$ for a positive constant θ. (Hence, the probability that action a_i is Sender's most favorite action is proportional to i^θ, which indicates that the misalignments in Receiver's and Sender's preferences over actions vanish with a fast rate as θ increases.) The case of $\theta = 0$, therefore, corresponds to Sender having random independent preferences over actions, and higher values of θ correspond to higher alignment between Receiver's and Sender's preferences over actions. For any fixed $\theta \geq 0$, the expected size of the stochastically dominant signal, π^*, is of the order of $\ln |\Omega|$.

Proposition 4.4. *For any non-negative θ, $\mathbb{E}_{\mathcal{D}}\left[|\pi^*|\right] \leq c + \log_{\frac{3+\theta}{2+\theta}} |\Omega|$, where c is a constant depending only on θ.*

The larger θ, the larger the above bound on size of the signal would be, but no matter how large Ω, the size of the signal is of the order of $\ln |\Omega|$ for any fixed θ.

[7] When there is negative correlation, the size of the stochastically dominant signal tends to be smaller than when there is positive correlation. The intuition is simple: at the opposite extreme to the setting of Proposition 4.1, the action taken by Receiver at her least favorite state is Sender's most favorite action, and therefore the stochastically dominant signal is the uninformative singleton signal.

[8] Therefore, the probability that action a is Sender's most favorite action in a subset $B \subseteq A$ of actions containing a is $\frac{e^{\beta_a}}{\sum_{b \in B} e^{\beta_b}}$.

Noisy utility function. Similar to the previous example, let $\Omega = \{\omega_1, \ldots, \omega_n\}$, $A = \{a_1, \ldots, a_n\}$, and Receiver's preferences satisfy $\mathfrak{a}^*(\omega_i) = a_i$ for all $\omega_i \in \Omega$. In this example, we model the misalignments in Sender's and Receiver's preferences by adding noise to Sender's utility function. This can be done in many ways; we choose an analytically simple one.

Let $v(a_i, \omega_j) = i + \epsilon_{i,\theta}$, where $\epsilon_{i,\theta}$ is drawn i.i.d. from a uniform distribution with support $[-\theta, \theta]$. Let the distribution imposed on $\preceq_s$ be denoted by $\mathcal{D}$. The noise variables $(\epsilon_{i,\theta})$ change Sender's preferences in a simple way: they change the position of an action in Sender's preference list over actions by at most 2θ positions. At $\theta = 0$, Receiver's best choice at ω_i is a_i. Therefore, Sender's optimal signal is the fully revealing one. When $\theta > 1$, Receiver's choices and Sender's preferences over actions are not aligned, and the size of the optimal signal, π^*, shrinks with a shrinkage factor proportional to $\sqrt{\theta}$.

Proposition 4.5. *For any $\theta \geq 1$, $\mathbb{E}_{\mathcal{D}}[|\pi^*|]$ is of the order of $|\Omega|/\sqrt{\theta}$.*

When θ does not grow with $|\Omega|$, we are in the regime that Receiver's choices and Sender's preferences over actions are well-aligned and the sizes of the optimal signal and the state space are of the same order of magnitude. In the regime where θ grows with $|\Omega|$, size of the optimal signal is always an order of magnitude smaller than size of the state space.

In the last part of this section, we show that the small size of the dominating signal does not hinge on the state-monotonicity assumption. We provide a counterpart for Proposition 4.3 for when state-monotonicity is replaced with commutativity. The counterpart focuses on a particular family of commutative preferences.

Definition 4.6. *A preference relation $\preceq$ over $\Omega = \Pi_{i=1}^{k}\Omega_i$ is a* product order *when there exists a family of complete preference relations $\{\preceq^i\}_{i=1}^{k}$ such that $(\omega_1, \ldots, \omega_k) \preceq (\omega'_1, \ldots, \omega'_k)$ holds iff $\omega_i \preceq^i \omega'_i$ holds for all positive integers $i \leq k$.*

In Section 3.1, when commutativity was introduced, we saw a simple example of a product order preference relation, where $\Omega = \Pi_{i=1}^{|A|}\Omega_i$ and Ω_i provides information about action $i \in A$. When Receiver has commutative preferences with respect to a product order and Sender has random independent preferences over actions, the (smallest) stochastically dominant signal, π^*, is small in the following sense.

Proposition 4.7. *Let $\Omega = \Pi_{i=1}^{k}\Omega_i$, $\preceq^i$ be a complete and transitive preference relation over Ω_i, and let $\leq_\Omega$ denote the product order defined over Ω by $\{\preceq^i\}_{i=1}^{k}$. Suppose Receiver has commutative preferences with respect to $\leq_\Omega$, and Sender has random independent preferences over actions drawn from a distribution $\mathcal{D}$. Then, $\mathbb{E}_{\mathcal{D}}[|\pi^*|] \leq 1 + \log_{\frac{2^k}{2^k-1}} |\Omega|$.*

The larger k, the larger the above bound on size of the optimal signal would be (suggesting that more complex partial preferences require more signal realizations), but no matter how large Ω, the size of the optimal signal is of the order of $\ln|\Omega|$ for any fixed k.

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, all of which provide sufficient (and in some cases necessary) conditions for Sender's optimal signal to have an "interval structure" in different settings.

[Dworczak and Martini 2018] provide a sufficient and necessary condition for the optimal signal to have an interval structure (that pools only adjacent elements) for when Sender's preferences depend only on the mean of posterior beliefs. The sufficient and necessary condition is in terms of the shape of Sender's utility function. [Ivanov 2015] studies a setting where Sender's preferences depend also on the order of posterior means, and provides a sufficient condition for the optimal signal to have an interval structure. [Kolotilin 2017], taking a linear programming approach to solve persuasion problems, provides a sufficient condition for the optimality of "interval revelation schemes", a type of revelation scheme that also has an interval structure. The interval structure of messages also appear in other models of communication, e.g. in the cheap talk game of [Crawford and Sobel 1982].

[Mensch 2018] defines the notion of monotone signals as signals that do not recommend lower actions at higher states, and therefore 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 Sender's and Receiver's preferences is sufficient when the state space is binary.

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 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 Receiver's private information and a max-min expected utility function. They show that when 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

J. Best and D. Quigley. Persuasion for the long run. *Working paper*, 2017. 19

H. Campbell. What it's like to be part of lyft's new prime time experiment. https://therideshareguy.com/what-its-like-to-be-part-of-lyfts-new-prime-time-experiment/, 2018. 13

V. P. Crawford and J. Sobel. Strategic information transmission. *Econometrica*, 1982. 19

P. Dworczak and G. Martini. The simple economics of optimal persuasion. *Journal of Political Economy, forthcoming*, 2018. 2, 19

J. Hu and X. Weng. Robust persuasion of a privately informed receiver. *Working paper*, 2018. 19

M. Ivanov. Optimal signals in bayesian persuasion mechanisms. *Working paper*, 2015. 2, 19

E. Kamenica and M. 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. 2, 7, 8

E. Kamenica and M. Gentzkow. Competition in persuasion. *Review of Economic Studies*, 2017. 2

A. Kolotilin. Optimal information disclosure: A linear programming approach. *Theoretical Economics*, 2017. 2, 19

J. Mensch. Monotone persuasion. *working paper*, 2018. 2, 19

B. Mityagin. The zero set of a real analytic function. 2015. URL https://arxiv.org/abs/1512.07276. 30

Wiki-MCP. Generalized maximum coverage problem. https://en.wikipedia.org/wiki/L%27H%C3%B4pital%27s_rule#Corollary, 2017. 12

A Example

First, we present the example assuming that the state space is compact. Later we will also construct the optimal signal for a discrete state space in a similar way.

An ad exchange (Sender) reveals information about a potential customer's type to an advertiser (Receiver). Type of the customer is a positive number θ belonging to the unit interval. After receiving the information, the advertiser posts a bid in an auction to win the opportunity of displaying her advertisement to the customer. In case of winning the auction, the advertiser pays her bid only if the customer clicks on the advertisement.

Suppose that the chance of the event that the advertiser wins the auction and the customer clicks on her advertisement is given by $q(b, \theta)$, where b is the advertiser's bid and θ is the customer's type.

With a customer of type θ, utility of an advertiser bidding b is given by

$$u(b, \theta) = q(b, \theta) \cdot (v(\theta) - b),$$

where $v(\theta)$ is the expected value generated for the advertiser if a customer of type θ clicks on her advertisement. We suppose that $u(b, \theta)$ is increasing in θ.

Given a customer of type θ, define $b^*(\theta)$ to be the bid that maximizes the advertiser's utility. (Note that the optimal bid, $b^*(\theta)$, may be increasing or decreasing in θ. Also, $q(b, \theta)$ may be increasing or decreasing in θ. For example, when the customer's type is "higher", there may be more competitor bids which could decrease the chance of winning the auction, holding the bid fixed.)

Advertiser's utility is given by the function $R(b, \theta; \mathbf{b}_{-})$, where b is the advertiser's bid, and the vector $\mathbf{b}_{-}$ denotes her competitors bids. For brevity, we denote this utility function by $R(b, \theta)$. Suppose that $R(b, \theta)$ is increasing in b, the advertiser's bid.

What is Sender's optimal signal, given that the advertiser's objective is defined as in Section 2.1? The solution is explained by Figure 2, which illustrates the optimal signal for a given function b^*. Optimality of the solution can be verified using the greedy algorithm of Section 3.2. A more intuitive approach is (discretizing the state space and) using the algorithm discussed in Section 2 (i.e. Algorithm 1 in the appendix), as done next.

Suppose θ takes discrete equidistance values in the unit interval. (Figure 3) In the first iteration of the algorithm, Sender's most preferred inducible action (bid) is $\sup_\theta b^*(\theta)$. The state inducing that action is $\theta = 1$. Therefore the signal realization constructed by the algorithm is a singleton containing only that state. Similarly, for any $\theta > \theta_1$ in the support of θ, there will be a signal realization containing only that state. When all such states are covered by the algorithm, the highest inducible action would be $b(\theta_1) = b(\theta_2)$. The lowest state inducing that action is θ_2. Therefore, the next signal realization would contain θ_2 and all the higher uncovered states, i.e. the states in the support of θ belonging to the interval $[\theta_2, \theta_1]$. Construction of the signal continues similarly until all states are covered.

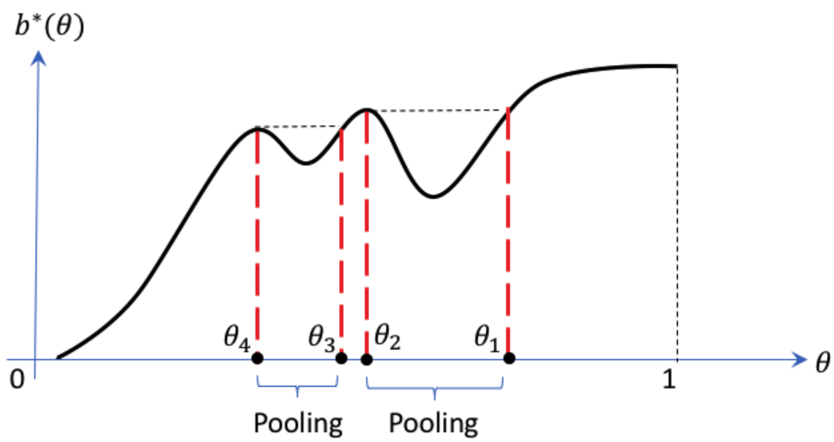

Figure 2: The optimal signal uses pooling in the intervals $[\theta_4, \theta_3]$ and $[\theta_2, \theta_1]$, and uses full revelation everywhere else.

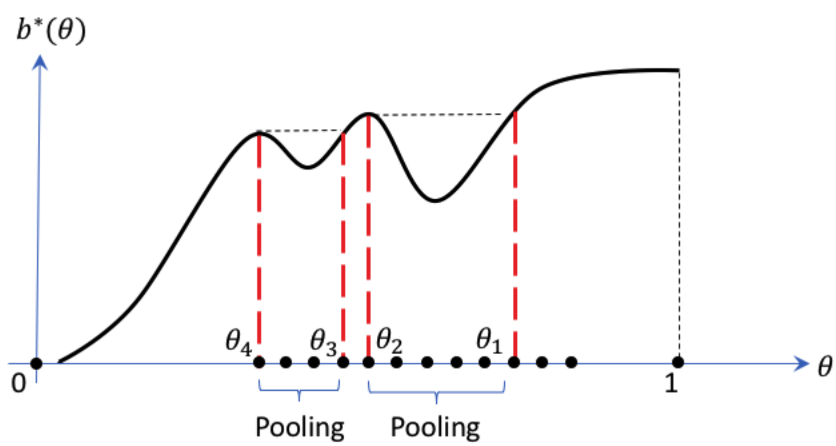

Figure 3: The same example with a discrete state space. (The θ-axis is discretized with equidistance dots.)

B Definitions

We need a few definitions for the proofs. We say a state ω *induces* action a if Receiver's chosen action is a when the signal realization is $\{\omega\}$. We say an action a is Ψ-inducible if there exists a state $\psi \in \Psi$ that induces a.

For any signal π, recall that η_π denotes the distribution induced over outcomes by π. Let $\mathfrak{a}^*(Q)$ denote the action chosen by Receiver when the signal realization is $Q \subseteq \Omega$. We call $\mathfrak{a}^*$ the choice function of Receiver. A signal realization Q *induces* an action a when $\mathfrak{a}^*(Q) = a$.

C Proofs from Section 2

Proof of Theorem 2.1. We start the proof by constructing π^*, which is done in Algorithm 1. After that, we show that π^* satisfies the promised properties.

Algorithm 1: Construction of π^*

$\Psi \leftarrow \varnothing$
$A^* \leftarrow \varnothing$
$\pi^* \leftarrow \varnothing$
while $\Psi \neq \Omega$ **do**
    Find the most preferred action with respect to $\preceq_A$ in $A \backslash A^*$ which is $(\Omega \backslash \Psi)$-inducible.
    $\omega \leftarrow \inf^{\preceq_\Omega}\{\omega' : \omega' \text{ induces } a\}$
    $A^* \leftarrow A^* \cup \{a\}$
    For any $w' \in \Omega \backslash \Psi$ such that $w \preceq_\Omega w'$, add w' to Ψ and to P.
    $\pi^* \leftarrow \pi^* \cup \{P\}$
end

Observe that the algorithm must terminate, since the size of $\Omega \backslash \Psi$ is reduced at each step. When the algorithm terminates, it has constructed a partition of the states, i.e. the signal π^*.

To prove the theorem, we need some definitions. We call A^* the set of *selected actions* and Ψ the set of *covered states*. For a state ω, we say Sender ω-prefers signal π to π' if, when the state of the world is ω, she weakly prefers the outcome induced under signal π to the outcome induced under signal π'.

To prove that π^* is a stochastically dominant signal, we will show that for any state $\omega \in \Omega$ and any signal π, Sender ω-prefers π^* to π. Suppose P is the signal realization in π that contains ω. Furthermore, suppose ω belongs to the signal realization Q in π^*.

Let a, b be the actions induced by P, Q, respectively, and suppose $b \preceq_A a$. Also, let $\delta = \inf^{\preceq_\Omega}\{Q\}$. The state δ should be covered by the algorithm at some point before constructing P: because otherwise, at the time of constructing P, δ would be uncovered and b would be a more preferred $(\Omega \backslash \Psi)$-inducible action by Sender than a, which would be a contradiction.

But then, if δ is covered before P is constructed by the algorithm, so should ω: because $\delta \preceq_\Omega \omega$ holds since $\omega \in Q$, and therefore ω should be covered when δ is, by line 8 of the algorithm. Contradiction.

The Convexity of π^* is guaranteed by construction: the only place that signal realizations are added to π^* is line 8 of the algorithm. There, we see that whenever a state is included in a signal realization (is covered), all the uncovered states that are preferred by Receiver to that state will also be added to the same signal realization. This guarantees that the convexity holds.

To see why π^* is monotone, observe that the algorithm selects the actions in a decreasing order with respect to $\preceq_A$. That is, every time that line 5 is run, the selected action is lower than the previous one. The selected action is the action induced by the signal realization constructed in line 8. (i.e. the new set covered states). The algorithm covers the states in a decreasing order with respect to $\preceq_\Omega$: every time line 8 is run, any state that is added to the signal realization P in that line is preferred less by Receiver to any state that was added in the previous run of line 8. These two facts together imply that the constructed signal is monotone.

$\square$

D Proofs from Section 3

D.1 Proofs from Section 3.1

Proof of Theorem 3.1. The proof takes an approach similar to the proof of Theorem 2.1. In particular, it constructs π^* by an algorithm similar to Algorithm 1. First, we recall a few definition. A state ω *induces* action a if Receiver's chosen action is a when the signal realization is $\{\omega\}$. An action a is Ψ-inducible if there exists a state $\psi \in \Psi$ that induces a. We are now ready for the algorithm that constructs π^*, Algorithm 2.

The algorithm must terminate, since the size of $\Omega \backslash \Psi$ is reduced at each step. When the algorithm terminates, it has constructed a partition of the states, i.e. the signal π^*. The proof that π^* is stochastically dominant is quite similar to the proof of Theorem 2.1. We include the full proof for completeness. First, we need a few definitions.

Algorithm 2: Construction of π^*

1 $\Psi \leftarrow \varnothing$
2 $A^* \leftarrow \varnothing$
3 $\pi^* \leftarrow \varnothing$
4 **while** $\Psi \neq \Omega$ **do**
5 Find the most preferred action with respect to $\preceq_A$ in $A \backslash A^*$ which is $(\Omega \backslash \Psi)$-inducible; call it action a.
6 $\omega \leftarrow \bigwedge\{\omega' : \omega' \text{ induces } a\}$
7 $A^* \leftarrow A^* \cup \{a\}$
8 For any $w' \in \Omega \backslash \Psi$ such that $\omega \leq_\Omega w'$, add w' to Ψ and to P.
9 $\pi^* \leftarrow \pi^* \cup \{P\}$
10 **end**

We call A^* the set of *selected actions* and Ψ the set of *covered states*. For a state ω, we say Sender ω-prefers signal π to π' if, when the state of the world is ω, she weakly prefers the outcome induced under signal π to the outcome induced under signal π'.

To prove that π^* is a stochastically dominant signal, we will show that for any state $\omega \in \Omega$ and any signal π, Sender ω-prefers π^* to π. Suppose P is the signal realization in π that contains ω. Furthermore, suppose ω belongs to the signal realization Q in π^*.

Let a, b be the actions induced by P, Q, respectively, and suppose $b \preceq_A a$. Also, let $\delta = \bigwedge Q$. The state δ should be covered by the algorithm at some point before constructing P: because otherwise, at the time of constructing P, δ would be uncovered and b would be a more preferred $(\Omega \backslash \Psi)$-inducible action by Sender than a, which would be a contradiction.

But then, if δ is covered before P is constructed by the algorithm, so should ω: because $\delta \leq_\Omega \omega$ holds since $\omega \in Q$ and $\delta = \bigwedge Q$, and therefore ω should be covered when δ is, by line 8 of the algorithm. Contradiction.

The Convexity of π^* is guaranteed by construction: the only place that signal realizations are added to π^* is line 8 of Algorithm 2. There, we see that whenever a state is included in a signal realization (is covered), all the uncovered states that are in its upper contour set with respect to $\leq_\Omega$ are also be added to the same signal realization. This guarantees that any signal realization in π^* is convex with respect to $\leq_\Omega$.

To see why π^* is monotone, observe that the algorithm selects the actions in a decreasing order with respect to $\preceq_A$. That is, every time that line 5 is run, the selected action is lower than the previous one. The selected action is the action induced by the signal realization constructed in line 8. The algorithm covers the states in a non-increasing order with respect to $\preceq_\Omega$ (the order is not necessarily decreasing, since $\leq_\Omega$ may be incomplete). Every time line 8 is run, a state of the world and the uncovered states in its upper contour set with respect to $\leq_\Omega$ are covered (and form a signal realization). Thereby, a state that is covered in later iterations of line 8 cannot be higher (with respect to $\leq_\Omega$) than a state that is covered sooner. This fact together with the fact that the algorithm selects the actions in a decreasing order completes the proof.

□

D.2 Proofs from Section 3.2

Proof of Proposition 3.2. Let π_G denote the greedy signal.

For a state ω, we say Sender ω-prefers signal π to π' if when the state of the world is ω, she weakly prefers the outcome induced under signal π to the outcome induced under signal π'.

Claim D.1. *For any proper signal π, Sender ω-prefers π_G to π for all but a μ-measure zero of the states $\omega \in \Omega$.*

If this claim holds, the proof will be complete: suppose not; then there exist a signal π and an outcome p such that

$$\int_{p \preceq_s q} 1 \,\mathrm{d}\,\eta_{\pi_G}(q) < \int_{p \preceq_s q} 1 \,\mathrm{d}\,\eta_{\pi}(q).$$

But that means there exist a positive μ-measure of states such as ω such that Sender ω-prefers π to π_G, which would be a contradiction. Therefore, it remains to prove Claim D.1.

Proof of Claim D.1. We will prove that Sender ω-prefers π_G to π for any non-representative state ω (i.e. any $w \notin R$). Since there are only a finite number of representative states, this will prove the claim.

Suppose ω is a non-representative state. Therefore in line 4 of the greedy algorithm, a_ω is defined and in line 5, ω is assigned to the signal realization represented by a_ω. To prove the claim, we will show that there is no signal realization of π, namely $P \subseteq \Omega$, such that $(a_\omega, \omega) \prec_s (\mathfrak{a}^*(P), \omega)$ and $\omega \in P$. For contradiction, suppose there is. Let $b = \mathfrak{a}^*(P)$. By commutativity, Ω_b contains $\bigwedge P$, and therefore it is not empty; so, $\omega_b = \bigwedge \Omega_b$ is well-defined. Moreover, $\omega_b {\leq_\Omega} \omega$ holds, since Ω_b contains $\bigwedge P$. Therefore, in line 5, we should have

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

which implies

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

Contradiction.

□

□

D.3 Optimality of the greedy signal in cardinal utility models

We provide a counterpart for Proposition 3.2 in a cardinal utility model. 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 $\int_{\omega\in\Omega} v(\rho(\omega), \omega)\,\mathrm{d}\,\mu(\omega)$ is well-defined, i.e. the Lebesgue integral exists.

Proposition D.2. *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,

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

holds. Given this inequality, the proof would be complete: Sender's payoff under the signal π is just equal to $\int_{\omega\in\Omega} \pi(\omega)\,\mathrm{d}\,\mu(\omega)$. Since there is only a finite number of representative states (with measure 0), and since μ is atomless, then by (D.1),

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

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

It remains to prove (D.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

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

Consider an arbitrary state $w \notin R$. By construction, $\omega \in P_{a_\omega}$. Now, consider an arbitrary signal ρ with $P \in \rho$ and $\omega \in P$. Let $b = \mathfrak{a}^* \left(\bigwedge P\right)$. If we show that $\omega_b = \bigwedge\{\omega_b, \omega\}$, then we must have $v(a_\omega, \omega) \geq v(b, \omega)$, and the proof would be complete. This holds because $\omega \in P$, $\bigwedge P \in \Omega_b$, $\omega_b = \bigwedge \Omega_b$; hence, $\omega_b = \bigwedge\{\omega_b, \omega\}$.

□

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 an unlikely case in practical applications. Next, we discuss some of the possible sufficient conditions that ensure the validity of the constructed signal.

D.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 integrand has a μ-measure zero points of discontinuity (in Ω), 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 use two technical conditions to ensure that the the lower semilattice and the payoff functions are "well-behaving", which in turn guarantees the validity of the constructed signal. After briefing 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 $\leq_\Omega$, i.e. $\{\omega' \in \Omega : \omega \leq_\Omega \omega'\}$, is a closed set.[9]

Condition ii: on the payoff functions. Let a *subproblem* be defined by 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 $a \in A$.

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

Proposition D.3. *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 $\leq_\Omega$. 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 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, therefore the existence of the integrals are guaranteed by Condition ii.

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 implied from the fact that g is an analytic function itself, and therefore the set of its roots has Lebesgue-measure zero [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 a μ-measure zero. The set of stable points is a subset of Z, and therefore has a μ-measure zero as well. $\square$

E Proofs from Section 4

Proof of Proposition 4.3. The proof uses the Principle of Deferred Decisions. Rather than fixing the composition of $\preceq_s$, we generate it while running Algorithm 1. We first need to make a few definitions. Let $\mathfrak{a}^*(Q)$ denote the action chosen by Receiver when the signal realization is $Q \subseteq \Omega$. We say a signal realization Q induces an action a when $\mathfrak{a}^*(Q) = a$. When $Q = \{\omega\}$ is singleton, with slight abuse of notation we also denote $\mathfrak{a}^*(Q)$ by $\mathfrak{a}^*(\omega)$. Let $\hat{A} = \cup_{\omega\in\Omega}\mathfrak{a}^*(\omega)$, and for any $a \in \hat{A}$, let $\mathfrak{a}^{-1}(a)$ denote the set $\{\omega : \mathfrak{a}^*(\omega) = a\}$.

When Algorithm 1 is run, it finds the most preferred action (with respect to $\preceq_A$) in $\hat{A}$ that is inducible; this is done in line 5. Consider the first time that line 5 is run. When it comes to choosing the most preferred action, we use the Principle of Deferred Decisions, and rather than defining $\preceq_A$ over $\hat{A}$ completely, we only determine $\sup^{\preceq_A}\{\hat{A}\}$. This can done simply by choosing one of the elements in $\hat{A}$ uniformly at random, since Sender has random independent preferences over actions. Suppose the chosen element is a (following the notation in Algorithm 1). The algorithm then proceeds to line 8 where the first signal realization, P, is constructed. Verify that

$$P = \{\psi : \inf^{\preceq_\Omega}\{\mathfrak{a}^{-1}(a)\} \preceq_\Omega \psi\}$$

.

We will show that the expected size of P is at least $|\Omega|/2$, and then use this fact to complete the proof by repeatedly applying the same argument.

Claim E.1. *When a is chosen uniformly at random from $\hat{A}$, then the expected size of P is at least $|\Omega|/2$.*

Proof. Define $P' \subseteq \Omega$ as follows. Choose ω' uniformly at random from Ω. Then, let $P' = \{\psi : \omega' \preceq_\Omega \psi\}$. It is straight-forward to verity that the expected size of P' is at least $|\Omega|/2$.

To prove the claim, we will show that the expected size of P is at least equal to the expected size of P'. Note that in the construction of P, any element of Ω is a member of $\mathfrak{a}^{-1}(a)$ with probability at least $1/|\Omega|$. On the other hand, in the construction of P', each element of Ω is chosen as ω' with probability $1/|\Omega|$. This fact together with the definitions of P, P' imply that the expected size of P is at least equal to the expected size of P'. □

We are now ready to finish the proof. Suppose $x = |\Omega|$ and let $T(x)$ denote the expected size of the signal that is constructed by Algorithm 1 when the state space given its input has size x. Note that by line 8, $\Psi = P$. After the first iteration of the loop is completed, the rest of the algorithm is run essentially by removing Ψ from Ω and repeating the same loop. We therefore can write

$$T(x) = 1 + \mathbb{E}_a \left[T(x - |P|) \right], \tag{E.1}$$

where the expectation is taken over the choice of a. (Recall that a is chosen uniformly at random from $\hat{A}$.)

In the rest of the proof, we use induction to show that $T(x) \leq 1 + \log_2 x$. The base case for $x = 1$ is trivial. Suppose $x > 1$. We then can write

$$\begin{aligned} T(x) &\leq 1 + \mathbb{E}_a \left[1 + \log_2(x - |P|) \right], && \text{(E.2)} \\ &\leq 2 + \log_2 \left(\mathbb{E}_a \left[x - |P| \right] \right) && \text{(E.3)} \\ &\leq 2 + \log_2 \left(\mathbb{E}_a \left[x/2 \right] \right) = 1 + \log_2 x, && \text{(E.4)} \end{aligned}$$

where (E.2) holds by (E.1) and the induction hypothesis, (E.3) is by the Jensen inequality, and (E.4) holds by Claim E.1. The proof is complete. □

Proof of Proposition 4.7. The proof follows a similar approach as the proof of Proposition 4.3. Let $\hat{A} = \cup_{\omega \in \Omega \setminus \Psi} \mathfrak{a}^*(\omega)$, where Ψ is defined in Algorithm 2. Also, for any $a \in \hat{A}$, let $\mathfrak{a}^{-1}(a)$ denote the set $\{\omega : \mathfrak{a}^*(\omega) = a\}$. When Algorithm 2 is run, it finds the most preferred action (with respect to $\preceq_A$) in $\hat{A}$ that is inducible; this is done in line 5. When it comes to choosing the most preferred action, we use the Principle of Deferred Decisions, and rather than defining $\preceq_A$ over $\hat{A}$ completely, we only determine $\sup^{\preceq_A}\{\hat{A}\}$. This can done simply by choosing one of the elements in $\hat{A}$ uniformly at random, since Sender has random independent preferences over actions. Suppose the chosen element is a (following the notation in Algorithm 2).

The algorithm then proceeds to line 8 where the first signal realization, P, is constructed. Verify that

$$P = \{\psi : \psi \in \Omega \backslash \Psi, \bigwedge \mathfrak{a}^{-1}(a) \leq_\Omega \psi\}.$$

We will show that the expected size of P is at least $|\Omega\backslash\Psi|/2^k$. We then use this fact to complete the proof by repeatedly applying the same argument.

Claim E.2. *When a is chosen uniformly at random from $\hat{A}$, then the expected size of P is at least $|\Omega\backslash\Psi|/2^k$.*

Proof. The proof is induction. The induction basis is $k = 1$, which is proved in Claim E.1. For the induction step, suppose $k > 1$.

Define $P' \subseteq \Omega$ as follows. Choose ω' uniformly at random from $\Omega\backslash\Psi$. Then, let

$$P' = \{\psi : \psi \in \Omega\backslash\Psi, \omega' \leq_\Omega \psi\}.$$

First, we will show that the expected size of P is at least equal to the expected size of P', and then we will show that the expected size of P' is at least $|\Omega\backslash\Psi|/2^k$. For the first step, note that in the construction of P, any element of $\Omega\backslash\Psi$ is a member of $\mathfrak{a}^{-1}(a)$ with probability at least $1/|\Omega\backslash\Psi|$. On the other hand, in the construction of P', each element $\omega \in \Omega\backslash\Psi$ is chosen as ω' with probability $1/|\Omega\backslash\Psi|$. This fact together with the definitions of P, P' imply that the expected size of P is at least equal to the expected size of P'.

It remains to show that the expected size of P' is at least $|\Omega\backslash\Psi|/2^k$. To this end, define $S = \{\omega_1 : (\omega_1, \ldots, \omega_k) \in \Omega\backslash\Psi\}$. Also, denote $|S|$ by m and w.l.o.g. suppose that elements of S are $\omega_1^1 \lesssim^1 \ldots \lesssim^1 \omega_m^1$. For any positive integer $i \leq m$ define

$$T_i = \big\{(\omega_1, \ldots, \omega_k) \in \Omega\backslash\Psi : \omega_1 = \omega_i^1\big\},$$

i.e. the elements in $\Omega\backslash\Psi$ with ω_i^1 as their first component. Let $x_i = |T_i|$.

We now provide a lower bound on the expected size of P'. Recall the definition of P',

$$P' = \{\psi : \psi \in \Omega\backslash\Psi, \omega' \leq_\Omega \psi\},$$

where ω' is chosen uniformly at random from $\Omega\backslash\Psi$. Conditioned on $\omega' \in T_i$, for any $j \leq i$ the expected size of the elements from T_j which will be added to P' is at least $\frac{x_j}{2^{k-1}}$. This follows from the induction hypothesis. This fact implies that the expected size of P' is at least

$$\sum_{i=1}^{m}(i-m+1)\cdot\frac{x_i}{|\Omega\backslash\Psi|}\cdot\frac{x_i}{2^{k-1}}.$$

Denote the above quantity by $f(x_1,\ldots,x_m)$. Observe that f is a convex function of x (e.g. by verifying that it has a positive semi-definite Hessian). Therefore, the minimum value of f is attained at a point where $x_1=\ldots=x_m$, subject to the constraint that $\sum_{i=1}^{m}x_i=|\Omega\backslash\Psi|$. This implies that $f(x)$ is at least $\frac{|\Omega\backslash\Psi|}{2^k}$, which proves the claim. $\square$

We are now ready to finish the proof. Let $x=|\Omega\backslash\Psi|$ and let $T(x)$ denote the expected number of signal realizations that Algorithm 2 adds to π^* when the remaining state space, $\Omega\backslash\Psi$, has size x. We then can write

$$T(x)=1+\mathbb{E}_a\left[T(x-|P|)\right], \tag{E.5}$$

where the expectation is taken over the choice of a. (Recall that a is chosen uniformly at random from $\hat{A}$.)

In the rest of the proof, we use induction to show that $T(x)\le 1+\log_{\frac{2^k}{2^k-1}}x$. The base case for $x=1$ is trivial. Suppose $x>1$. We then can write

$$\begin{aligned}
T(x) &\le 1+\mathbb{E}_a\left[1+\log_{\frac{2^k}{2^k-1}}(x-|P|)\right], &\text{(E.6)}\\
&\le 2+\log_{\frac{2^k}{2^k-1}}\left(\mathbb{E}_a\left[x-|P|\right]\right) &\text{(E.7)}\\
&\le 2+\log_{\frac{2^k}{2^k-1}}\left(\mathbb{E}_a\left[x\cdot\frac{2^k-1}{2^k}\right]\right)=1+\log_{\frac{2^k}{2^k-1}}x, &\text{(E.8)}
\end{aligned}$$

where (E.6) holds by (E.5) and the induction hypothesis, (E.7) is by the Jensen inequality, and (E.8) holds by Claim E.2. The proof is complete.

$\square$

Proof of Proposition 4.4. We use the Principle of Deferred Decisions and rather than fixing $\preceq_A$, we construct it in the course of running Algorithm 1.

To this end, let $\hat{A}=\cup_{\omega\in\Omega}\mathfrak{a}^*(\omega)$, and for any $a\in\hat{A}$, let $\mathfrak{a}^{-1}(a)$ denote the set $\{\omega:\mathfrak{a}^*(\omega)=a\}$. Consider the first iteration of the while loop in Algorithm 1: In line 5, the most preferred action (with respect to $\preceq_A$) in $\hat{A}$ that is inducible is found. When it comes to choosing the most preferred action, we use the Principle of Deferred Decisions, and rather than defining $\preceq_A$ over $\hat{A}$ completely, we only determine $\sup^{\preceq_A}\{\hat{A}\}$. This can done simply by choosing one of the elements in $\hat{A}$ at random, with each element a being chosen with probability proportional to e^{β_a}. Let the chosen element be a (following the notation in Algorithm 1). The algorithm then proceeds to line 8 where the first signal realization, P, is constructed. Verify that

$$P = \{\psi : \inf^{\preceq_\Omega}\{\mathfrak{a}^{-1}(a)\} \preceq_\Omega \psi\}$$

.

We will show that the expected size of P is at least $|\Omega|/(\theta+3)$. We then use this fact to complete the proof by repeatedly applying the same argument.

Claim E.3. *When $|\Omega| > n_\theta$, the expected size of P is at least $|\Omega|/(\theta+3)$, where n_θ is a constant depending only on θ.*

Proof. For notational simplicity, let $n = |\Omega|$. When sampling an action from $\hat{A}$, action a_i is chosen with probability proportional to i^θ, in which case $|P| = n - i + 1$. Therefore, we can write

$$\begin{aligned}
\mathbb{E}\left[|P|\right] = \frac{\sum_{i=1}^{n}(n-i+1)i^\theta}{\sum_{i=1}^{n} i^\theta} &= n+1-\frac{\sum_{i=1}^{n} i^{\theta+1}}{\sum_{i=1}^{n} i^\theta} \geq n+1-\frac{\int_0^{n+1} i^{\theta+1}\,\mathrm{d}\,i}{\int_0^{n} i^{\theta}\,\mathrm{d}\,i} \\
&= n+1-(n+1)\cdot\frac{\theta+1}{\theta+2}\cdot(1+1/n)^{\theta+1}.
\end{aligned}$$

This implies that for all $n > n_\theta$, $\mathbb{E}\left[|P|\right] \geq \frac{n}{\theta+3}$, where

$$n_\theta = \left(\left(\frac{(\theta+2)^2}{(\theta+1)(\theta+3)}\right)^{1/(\theta+1)} - 1\right)^{-1}.$$

$\square$

We are now ready to finish the proof. Suppose $T(n)$ denote the expected size of the signal that is constructed by Algorithm 1 where $n = |\Omega|$. Note that by line 8, $\Psi = P$. After the first iteration of the loop is completed, the rest of the algorithm is run essentially by removing Ψ from Ω and repeating the same loop. We therefore can write

$$T(n) = 1 + \mathbb{E}_a\left[T(n-|P|)\right], \tag{E.9}$$

where the expectation is taken over action a.

In the rest of the proof, we use induction to show that $T(n) \leq c + \log_{\frac{3+\theta}{2+\theta}} n$, where c is a constant depending only on θ. To keep the proof simple, we define $c = n_\theta$. (E.g, $n_1 = 15$.

A tighter analysis can reduce the value of c; this is not our focus here.) The base cases for $n \le n_\theta$ are trivial. Suppose $n > n_\theta$. We then can write

$$T(n) \le 1 + \mathbb{E}_a\left[c + \log_{\frac{3+\theta}{2+\theta}}(n - |P|)\right], \tag{E.10}$$

$$\le 1 + c + \log_{\frac{3+\theta}{2+\theta}}\left(\mathbb{E}_a\left[n - |P|\right]\right) \tag{E.11}$$

$$\le 1 + c + \log_{\frac{3+\theta}{2+\theta}}\left(\mathbb{E}_a\left[n(\theta+2)/(\theta+3)\right]\right) = c + \log_{\frac{3+\theta}{2+\theta}} n, \tag{E.12}$$

where (E.10) holds by (E.9) and the induction hypothesis, (E.11) is by the Jensen inequality, and (E.12) holds by Claim E.3. The proof is complete. □

Proof of Proposition 4.5. First, we can harmlessly assume that $\epsilon_{i,\theta}$ is drawn iid from a uniform distribution with support $[-2\theta, 0]$. This is essentially just a change of variables (i.e. $\epsilon'_{i,\theta} = \epsilon_{i,\theta} - \theta$) that does not affect the size of the optimal signal. (Verify this by running Algorithm 1 on both instances)

For notational simplicity in the rest of the proof, let $\kappa = 2\theta$, $N = \{1, \ldots, n\}$, and $v_i = v(a_i, w_1)$, for all $i \in N$. Define the random variable $K = \max_{i \in N} n - v_i$, and let F_K denote its CDF. Also, define the random variable $I = n + 1 - \arg\max_{i \in N} n - v_i$, and let F_I denote its CDF. (For intuition, note that I is equal to the size of P, the signal realization being constructed by Algorithm 1) Let the random variable $(\epsilon_{1,\theta}, \epsilon_{2,\theta}, \ldots, \epsilon_{n,\theta})$ be denoted by $\boldsymbol{\epsilon}$ Observe that, for any positive integer $y > 0$,

$$1 - F_K(y) = \mathbb{P}_{\boldsymbol{\epsilon}}\left[K \ge y\right] = \begin{cases} 0, \text{ for } \kappa \le y \\ (1 - \frac{y}{\kappa})(1 - \frac{y-1}{\kappa}) \ldots (1 - \frac{y - \lfloor y \rfloor}{\kappa}), \text{ for } \kappa > y \end{cases} \tag{E.13}$$

We therefore can write

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{\epsilon}}\left[K\right] &= \int_0^{\kappa} 1 - F_K(y) \,\mathrm{d}\, y \\
&\ge \int_0^{\kappa/2} \Pi_{i=0}^{\lfloor y \rfloor} e^{-\frac{2(y-i)}{\kappa}} \,\mathrm{d}\, y && \text{(E.14)} \\
&\ge \int_0^{\kappa/2} e^{-\frac{(y+1)(y+2)}{\kappa}} \,\mathrm{d}\, y \\
&= \sqrt{\kappa} \cdot \frac{1}{2}\sqrt{\pi} e^{\frac{1}{4\kappa}} \left(\operatorname{erf}\left(\frac{\kappa+3}{2\sqrt{\kappa}}\right) - \operatorname{erf}\left(\frac{3}{2\sqrt{\kappa}}\right)\right), && \text{(E.15)}
\end{aligned}$$

where (E.14) holds because $e^{-2x} \le 1 - x$ for any positive $x \le 1/2$. Denote the second multiplicand in (E.15) by $f(\kappa)$. Let $\kappa_* = \inf_{\kappa \geq 1} f(\kappa)$. It is straight-forward to verify that $f(\kappa^*) > 0$.[10] The above inequality therefore says

$$\mathbb{E}_{\boldsymbol{\epsilon}}\left[K\right] \geq \sqrt{\kappa} \cdot f(\kappa), \tag{E.16}$$

where $f(\kappa) \geq f(\kappa^*) > 0$ for all $\kappa \geq 1$.

Claim E.4. *Suppose $n > 2\kappa$. Then, $\mathbb{E}_{\boldsymbol{\epsilon}}\left[I|K = k\right] \geq k/2$.*

Proof. For notational simplicity, let $m = n - \lfloor k \rfloor$, and let $M = \{m, \ldots, n\}$. Define the vector $\mathbf{v} = (v_m, \ldots, v_n)$, with $\mathbf{v}_i = v_i$ denoting the i-th element of $\mathbf{v}$. Let $\phi : \mathbb{R}^{n-m+1} \to \mathbb{R}_+$ be the marginal PDF induced by the random variable ϵ on $\mathbf{v}$. For any vector $\mathbf{x} \in \mathbb{R}^{n-m+1}$ let $\widehat{\mathbf{x}} = \{\mathbf{x}_1, \ldots, \mathbf{x}_{n-m+1}\}$. Also, let

$$\begin{aligned}
\mathcal{K}(\mathbf{x}) &= \max_{i \in M}(n - \mathbf{x}_{i-m+1}), \\
\mathcal{I}(\mathbf{x}) &= n + 1 - \arg\max_{i \in M}(n - \mathbf{x}_{i-m+1}).
\end{aligned}$$

Since $|\,\widehat{\mathbf{v}}\,| < n - m + 1$ is a zero probability event, we can suppose that $|\,\widehat{\mathbf{v}}\,| = n - m + 1$ in the rest of the proof. We say a vector $\mathbf{x} \in \mathbb{R}^{n-m+1}$ is *feasible* if there exist variables $\{\delta_{i,\kappa}\}_{i \in M}$ each in the interval $[-\kappa, 0]$ such that (i) $\mathbf{x}_{i-m+1} = v_i + \delta_{i,\kappa}$ for all $i \in M$, (ii) $k = \max_{i \in M} n - \mathbf{x}_i$, and (iii) all the components of $\mathbf{x}$ have distinct values.

Define $\mathcal{V}(\mathbf{x})$ to be the family of all feasible vectors that are obtained by permuting the elements of $\mathbf{x}$. Let $\mathcal{V}_i(\mathbf{x})$ denote a subset of $\mathcal{V}(\mathbf{x})$ that includes elements $\mathbf{x}'$ such that $i = \arg\min_{j \in M} \mathbf{x}'_{j-m+1}$.

We will show that for any feasible $\mathbf{x}$,

$$|\mathcal{V}_i(\mathbf{x})| \leq |\mathcal{V}_j(\mathbf{x})|, \text{ for } i \geq j. \tag{E.17}$$

If this holds, then the proof is complete: since the density function ϕ is uniform, $\phi(\mathbf{x}) = \phi(\mathbf{x}')$ for all $\mathbf{x}' \in \mathcal{V}(x)$. Therefore, (E.17) would imply that for any feasible $\mathbf{x}$,

$$\mathbb{P}\left[\mathcal{I}(\mathbf{x}) = n - i + 1 \middle| \,\widehat{\mathbf{x}}\,\right] \leq \mathbb{P}\left[\mathcal{I}(\mathbf{x}) = n - j + 1 \middle| \,\widehat{\mathbf{x}}\,\right], \text{ for } i \geq j. \tag{E.18}$$

[10]For example, this can be done by observing that f is increasing and $f(1) > 0$. In case one is interested in characterizing a quantitative upper bound for $|\pi^*|$, the fact that $f(2) \approx 0.12$ can be used right away to derive one. It is possible to derive better bounds with a sharper analysis.

The above inequality then would imply that

$$\mathbb{E}\left[I \middle| \widehat{\mathbf{x}}\right] \geq \frac{1}{\lfloor k \rfloor + 1} \cdot \sum_{i=1}^{\lfloor k \rfloor + 1} i \geq k/2.$$

To finish the proof, it remains to show that (E.17) holds. We do this by providing an injective mapping $f_{ij} : \mathcal{V}_i(\mathbf{x}) \to \mathcal{V}_j(\mathbf{x})$ for any $i > j$. For any $\mathbf{x}' \in \mathcal{V}_i(\mathbf{x})$, define $f_{ij}(\mathbf{x}') = \mathbf{x}''$, where

$$\mathbf{x}''_l = \begin{cases} \mathbf{x}'_i, & \text{for } l = j, \\ \mathbf{x}'_j, & \text{for } l = i, \\ \mathbf{x}'_l, & \text{for } l \in \{1, \ldots, n-m+1\} \backslash \{i, j\}. \end{cases}$$

Observe that (i) the mapping f_{ij} is injective, because it only swaps the i-th and j-th components of its input, (ii) $\mathbf{x}''$ is feasible because $i > j$, and (iii) $\mathbf{x}'' \in \mathcal{V}_j(\mathbf{x})$, because $\mathbf{x}''_j = \mathbf{x}'_i$. We therefore have shown that (E.17) holds. □

We are now ready to compute the expected size of π^*. Let $T(n)$ denote the expected size of the signal constructed by Algorithm (1) when $|\Omega| = n$ is the size of the given state space. Recall that I is just equal to the size of P, the signal realization being constructed by Algorithm 1. Let $n_0 = 2\kappa$. For any $n > n_0$, we can write

$$\mathbb{E}_{\boldsymbol{\epsilon}}\left[I\right] = \mathbb{E}_{\boldsymbol{\epsilon}}\left[\mathbb{E}_{\boldsymbol{\epsilon}}\left[I | K = k\right]\right] \geq \mathbb{E}_{\epsilon}\left[k/2\right] \geq \sqrt{\kappa} \cdot f(\kappa)/2, \tag{E.19}$$

which holds by (E.16) and Claim E.4. Define $g(\kappa) = \sqrt{\kappa} \cdot f(\kappa)/2$. (E.19) implies that

$$T(n) = 1 + \mathbb{E}_{\boldsymbol{\epsilon}}\left[T(n - |P|)\right] \leq 1 + \mathbb{E}_{\boldsymbol{\epsilon}}\left[T(n - g(\kappa))\right].$$

On the other hand, observe that for any $n \leq n_0$, $T(n) \leq n$ holds trivially. Therefore,

$$T(n) \leq n/g(\kappa) + n_0$$

holds for all $n \geq 1$. Noting that $g(\kappa) = \sqrt{\kappa} \cdot f(\kappa)/2$ completes the proof.

□