



Regret-minimizing Bayesian persuasion

Yakov Babichenko^a, Inbal Talgam-Cohen^a, Haifeng Xu^b, Konstantin Zabarnyi^{a,*}

^a Technion – Israel Institute of Technology, Israel

^b University of Chicago, United States of America

ARTICLE INFO

Article history:

Received 15 August 2021

Available online 19 September 2022

Keywords:

Bayesian persuasion

Regret minimization

Robust persuasion

ABSTRACT

We revisit the celebrated Bayesian persuasion setting and examine how well the Sender can perform when ignorant of the Receiver's utility. Taking an adversarial approach, we measure the signaling scheme performance via (additive) *regret* over a single persuasion instance. We focus on Receiver with two actions: *adoption* and *rejection*, assuming that Sender aims to persuade Receiver to adopt.

We show that while Sender's total ignorance of Receiver's utility is extremely harmful to her, assuming that Sender just knows Receiver's *ordinal preferences* over the states of nature upon adoption suffices to guarantee a surprisingly low regret even when the number of states tends to infinity. Moreover, we exactly pin down the minimum regret that Sender can guarantee in this case. We further show that such a positive result is impossible under the alternative performance measure of a *multiplicative approximation ratio*.

© 2022 Elsevier Inc. All rights reserved.

1. Introduction

Bayesian persuasion and applications Since the seminal paper of Kamenica and Gentzkow (2011), the model of *Bayesian persuasion* – also called *information design* – has been extensively studied in economics, computer science and operational research (for recent surveys, see Dughmi (2017); Kamenica (2019); Bergemann and Morris (2019); Candogan (2020)). The model considers an informed *Sender* who knows the *state of nature*, and an uninformed *Receiver* who does not know the state, but takes an action that affects both Receiver's and Sender's utilities. Sender has the ability to commit, before learning the state, to an information revelation policy called a *signaling scheme*. The main question of interest is: *what is the optimal Sender's utility, and what scheme should Sender choose to achieve this utility?*

One of the fundamental assumptions underlying Bayesian persuasion is that Sender perfectly knows Receiver's utility, and she can use this knowledge in her persuasion efforts. This assumption is too demanding in some circumstances. In the oft-cited words of Nobel Laureate Robert Wilson, “I foresee the progress of game theory as depending on successive reduction in the base of common knowledge required to conduct useful analyses of practical problems” Wilson (1987) (see also Hayek, 1945; Scarf, 1958).

Settings of interest We consider a setting with n states of nature, an arbitrary prior distribution over them and a binary-action Receiver whose possible actions are *adoption* and *rejection*. Sender aims to persuade Receiver to adopt – i.e., Sender's

* Corresponding author.

E-mail addresses: yakovbab@technion.ac.il (Y. Babichenko), italgam@cs.technion.ac.il (I. Talgam-Cohen), haifengxu@uchicago.edu (H. Xu), konstzab@gmail.com (K. Zabarnyi).

<https://doi.org/10.1016/j.geb.2022.09.001>

0899-8256/© 2022 Elsevier Inc. All rights reserved.

utility is 1 if Receiver adopts and 0 otherwise.¹ In particular, we study a state-independent Sender's utility. As a leading example, let us look at an interaction between a seller of a product – Sender – and a potential buyer – Receiver – and assume that the unknown state captures some characteristics of the product. In practice, the seller might be, indeed, unaware of the buyer's utility as a function of these characteristics.

The robust approach A natural way to address Sender's uncertainty is the *robust approach* – a.k.a. the *adversarial approach* or the *prior-free approach*. The goal is to design signaling schemes that perform robustly well, that is – perform well for *all* Receiver's utilities, even when an adversary is allowed to choose the *worst case* utility for Receiver. The adversarial lens has the potential to provide insights on issues such as: (a) *How harmful can it be for Sender to be unaware of Receiver's utility?* (b) *What information about Receiver's utility is sufficient to enable “reasonably good” persuasion?* (c) *How should Sender approach persuading Receiver if Receiver's utility is unknown?* We shall return to these questions in Subsection 1.2.

Main result We prove that even when the number of states n tends to infinity, it is enough for Sender to know only the ordering of the states according to Receiver's utilities upon adoption to guarantee a surprisingly low *regret* over a single persuasion instance, where the regret is the difference between Sender's optimal expected utility had she known Receiver's utility function and her actual expected utility. In the seller-buyer example, the assumption of Sender's ordinal knowledge of Receiver's preferences is natural when the state is the product's quality. Moreover, we find a simple explicit formula for the optimal regret (Theorem 3.6). Upon normalizing the regret to belong to $[0, 1]$, the optimal regret turns out to be at most $1/e$.

1.1. Related work

Bayesian persuasion has many applications besides our leading example, which may include criminal justice (Kamenica and Gentzkow, 2011), information management in organizations (Dworczak and Martini, 2019), security (Xu et al., 2015; Rabinovich et al., 2015), routing (Bhaskar et al., 2016), recommendation systems (Mansour et al., 2016), auctions (Emek et al., 2012; Bro Miltersen and Sheffet, 2012), voting (Cheng et al., 2015; Alonso and Câmara, 2016), and queuing (Lingenbrink and Iyer, 2019). In these applications and others, Sender might be unsure of Receiver's exact incentives. Common robust approaches are:

1. **Regret minimization.** It compares the performance of a signaling scheme when Sender *does not know* the Receiver's utility function to the performance of an optimal scheme *with the knowledge* of Receiver's utility by considering the *difference* between the optimal utilities in the two cases – the (additive) *regret*. Attributed to the classic work of Savage (1951) on decision theory, it is also the leading paradigm in online machine learning (Shalev-Shwartz and Ben-David, 2014). Recent examples in economic contexts are Arieli et al. (2018) (information aggregation) and Guo and Shmaya (2019a) (monopoly regulation).
2. **Adversarial approximation.** It is similar to the previous one, but the comparison is by considering the corresponding *ratio*. It is the leading paradigm in approximation algorithms (Vazirani, 2003), partially due to it being scale-free. In economic contexts, adversarial approximation underlies the well-developed research line on prior-independent auctions (Dhangwatnotai et al., 2015; Talgam-Cohen, 2020) and their sample complexity (Cole and Roughgarden, 2014) (see Roughgarden and Talgam-Cohen, 2019 for a survey). Another example is Hurwicz and Shapiro (1978) (sharecropping contracts).
3. **Minimax.** It measures the *absolute performance* of a signaling scheme with *no knowledge* of Receiver's utility function; this approach has no benchmark.² The minimax approach is attributed to the scholarship of Wald (1950). Recent applications in economics include auctions (Bandi and Bertsimas, 2014; Carroll, 2017; Gravin and Lu, 2018; Bei et al., 2019) and contracts (Carroll, 2015; Dütting et al., 2019) (see the comprehensive survey of Carroll (2019)).

In this paper, we mainly focus on the first approach of regret minimization. Our work is connected to Dworczak and Pavan (2020), who study the standard minimax approach to better capture the best policy for Sender. In Subsection 4.1, we also study adversarial approximation. We leave the minimax approach for future research.³

Our motivation is very similar to that of Castiglioni et al. (2020), whose interesting work studies Bayesian persuasion in an online learning framework, with the goal of relaxing the assumption that Sender knows Receiver's utility; our approach is complementary to theirs. In their model, Sender repeatedly faces Receiver with a non-binary action, whose type is chosen by Adversary at each round from a finite set of possibilities. Their regret notion is with respect to a best-in-hindsight single

¹ The binary-action Receiver is a fundamental case in Bayesian persuasion studied in many works – two recent examples are Kolotilin et al. (2017); Guo and Shmaya (2019b).

² Note that the name of this approach is slightly misleading, as all three robust approaches have a minimax flavor, and the difference is in what is being optimized within the minimax expression.

³ In our setting, there is not much Sender can do under the minimax approach without further assumptions (see Section 4). It goes well with Savage (1972) highlighting the “extreme pessimism” of the minimax approach compared to regret minimization (Chapters 9.8 and 13.2); see Wikipedia contributors (2020) for an example in which regret is less pessimistic than minimax.

signaling scheme. In contrast, we minimize the regret over a single persuasion instance with respect to the best scheme tailored to that instance. We consider a Receiver with a quite general class of (continuum-many) possible utilities, focusing on a binary-action case.⁴

Recently, the persuasion literature considers robustness to different ingredients of the persuasion model. Hu and Weng (2021); Dworczak and Pavan (2020) study robustness of persuasion with respect to private information that Receiver might have. Inostroza and Pavan (2018); Li et al. (2020); Mathevet et al. (2020); Morris et al. (2020); Ziegler (2020) consider persuasion with multiple receivers and study robustness with respect to the strategic behavior of the receivers after observing their signals.

1.2. Our results

We study the following three settings:

- *Arbitrary utilities*, with a completely ignorant Sender, who has no information on Receiver's utility. In our seller-buyer example, it represents a seller that is totally ignorant of customer's preferences over the characteristics – e.g., when the state is the product's color.
- *Monotonic utilities*, with Sender knowing that Receiver's utility upon adoption is monotonic as a function of the state. This assumption arises naturally when the state reflects possible qualities of a certain product.
- *Multidimensional monotonic utilities*, with the state representing the qualities of several attributes (dimensions) of a product. We assume that Receiver's utility upon adoption is monotonic in each dimension and focus on a constant number of attributes.

Results for arbitrary vs. monotonic utilities Our adversarial analysis shows that Sender cannot hope for a nontrivial bound on her regret for arbitrary utilities with a large number of states (Theorem 3.5). In contrast, even if the number of states tends to infinity, the regret remains quite low (at most $1/e$) in the monotonic utility case (Theorem 3.6). These two results answer to questions (a) and (b) from the robust approach introduction – while it might be very harmful for Sender to be unaware of Receiver's utility, knowing Receiver's ordinal preferences over the states suffices to persuade Receiver “reasonably well”. It highlights monotonicity as the distinguishing property among settings in which Sender can persuade and those in which she should seek additional information on Receiver before approaching him. Coming back to question (c) – how should Sender persuade Receiver without knowing his utility – our positive results are constructive. In particular, the idea of “pooling together” the highest-utility states – i.e., sending the same signal realization in them – is useful not only in standard persuasion, in which such threshold policies are known to be optimal (Kolotilin, 2015; Renault et al., 2017), but also for regret minimization. It makes our work fall within the “classic” theme of the robust mechanism design literature, by which well-known mechanisms have robustness properties – e.g. linear contracts (Carroll, 2015) and simultaneous ascending auctions (Milgrom, 2000). Our ignorant Sender should randomize the way she pools together the states to achieve the minimal regret, leading to a signaling scheme with continuum of signals. It stands in a sharp contrast to optimal policies in the standard persuasion model, in which binary-signal policies suffice. Several works on robust persuasion discover a similar phenomenon, including: Dworczak and Pavan (2020); Hu and Weng (2021), considering robustness to Receiver's exogenous sources of information; Perez-Richet and Skreta (2022), focusing on robustness to a state falsification; and Kosterina (2022), analyzing robustness to Receiver's prior.

Further results In Subsection 3.3, we study multidimensional states, assuming that Receiver's utility is monotonic in every dimension separately; each dimension represents an attribute of a product. Focusing on a constant number of attributes k , we provide a positive result for product priors by upper-bounding the regret in terms of k . However, as a corollary from the arbitrary utility analysis, the regret for general priors might approach 1 as the number of possible quality levels grows to infinity, even for $k = 2$.

In Subsection 4.1, we consider the alternative *adversarial approximation* approach. We establish that all our negative results immediately translate to this approach, and even for monotonic utilities – Sender can robustly achieve only a factor that is logarithmic in the inverse of the prior probability of the highest-utility state. In particular, when this prior tends to zero, the guaranteed fraction tends to zero as well.

Our techniques Robust persuasion defines a zero-sum game between Sender and an adversary (see Section 2). As both players have a rich set of strategies (signaling schemes vs. adversarial utility functions), it is unclear how to find the game value. Our main result (Theorem 3.6) uses the ordinal knowledge of states to reduce the complex game to a simpler one

⁴ Using functional analysis and convexity arguments, one can prove that for two states of nature, the optimal regret for Receiver with *arbitrarily many actions (even continuum)* when Sender knows Receiver's ordinal preferences over the states is upper-bounded by the optimal regret for a binary-action Receiver with the Sender knowing Receiver's ordinal preferences. The knowledge of the ordinal preferences over the states for two states and arbitrarily many actions means that Adversary must ensure that Sender's *indirect utility function* (Kamenica and Gentzkow, 2011) is monotonic in the posterior probability of one of the states.

in which each player chooses a threshold real number. Without the ordinal knowledge, getting a positive result requires complex analysis; in Proposition 3.4, we achieve this via geometric arguments and analytic geometry.

2. Our setting

We start by describing the classic Bayesian persuasion setting with a single Sender and a single Receiver. We impose two restrictions on the standard model: (1) Receiver has binary actions, which captures choosing whether to adopt a certain offer or not; (2) Sender's utility is state-independent, which captures caring only about whether Receiver adopts the offer, regardless of the circumstances.

Prior, posterior and signaling scheme Let $\Omega = [n] = \{1, \dots, n\}$ be the *state space*, and let $\omega \in \Omega$ be the *true state*. Denote by $\Delta(\Omega)$ the set of all the probability distributions over Ω . For every $q \in \Delta(\Omega)$ and $i \in \Omega$, let q_i be the probability the distribution q assigns to the state i . Define the *support* of q to be $\text{supp}(q) := \{i \in \Omega : q_i > 0\}$. Let $\mu \in \Delta(\Omega)$ be the *prior distribution* on Ω – or just the *prior*. Assume that μ is publicly known. Furthermore, assume w.l.o.g. that the probability of every state is strictly positive – that is, $\mu_i > 0$ for every $i \in \Omega$.⁵ Therefore, μ is an interior point of $\Delta(\Omega)$.

The *signaling scheme* is a stochastic mapping $\pi : \Omega \rightarrow S$, where S is a (finite or infinite) set of signals. Sender commits to a signaling scheme π and then observes the true state ω ; Receiver does not see ω . Upon learning ω , Sender transmits to Receiver, according to $\pi(\omega)$, a *signal realization* $s \in S$. After receiving s , Receiver updates his belief regarding the state distribution to a *posterior distribution* on Ω – or just the *posterior* – denoted by $p(s) \in \Delta(\Omega)$, where $p_i(s) = \Pr[\omega = i | s] = \frac{\Pr_{s' \sim \pi(i)}[s' = s] \mu_i}{\Pr[s]}$. Only the posterior distribution is important to the persuasion instance outcome, rather than the signal realization itself.

We slightly abuse the notation and use π also to denote the distribution over the elements of S induced by the signaling scheme π considering the prior μ . That is, for every $s_0 \in S$:

$$\Pr_{s \sim \pi}[s = s_0] = \sum_{i \in \Omega} \mu_i \cdot \Pr_{s' \sim \pi(i)}[s' = s_0].$$

Remark 2.1. It is well-known (see, e.g., Blackwell, 1953; Aumann et al., 1995) that a distribution D over Receiver's posteriors $p \in \Delta(\Omega)$ is implementable by some signaling scheme π if and only if $\mathbb{E}_{p \sim D}[p] = \mu$. We refer to this condition as *Bayes-plausibility*.

Utilities and adoption In our setting, Receiver has a binary *action space*; that is, Receiver's *action* a is selected from $\{0, 1\}$. Call action 1 *adoption* and action 0 *rejection*. Sender's utility is a function $u_s : \{0, 1\} \rightarrow [0, 1]$ of Receiver's action. We assume $u_s(a) := a$, i.e., Sender wants Receiver to take action 1 (to adopt). Thus, the expected Sender's utility equals the probability that Receiver adopts.

Receiver's utility is a function $u_r : \Omega \times \{0, 1\} \rightarrow \mathbb{R}$ of the state and Receiver's action. W.l.o.g., we normalize Receiver's utility $u_r(i, 0)$ to zero for every $i \in \Omega$. That is, *Receiver's utility for choosing to reject is always 0*, regardless of the state.

Given the posterior $p(s)$, Receiver adopts if and only if⁶:

$$\mathbb{E}_{\omega' \sim p(s)}[u_r(\omega', 1)] \geq 0. \quad (1)$$

Definition 2.2. The *adoption region* $A = A(u_r) \subseteq \Delta(\Omega)$ is the set of posteriors that lead to Receiver's adoption: $A := \{p \in \Delta(\Omega) : \mathbb{E}_{\omega' \sim p}[u_r(\omega', 1)] \geq 0\}$.

For a signaling scheme π and Receiver's utility function u_r , denote Sender's expected utility over π by:

$$u(\pi, u_r) := \Pr_{s \sim \pi}[p(s) \in A(u_r)].$$

Objective Given a Bayesian persuasion setting with (known) u_r as Receiver's utility, the standard goal of the designer (Sender) is to compute a signaling scheme $\pi^* = \pi^*(u_r)$ that maximizes Sender's expected utility: $\pi^* \in \arg \max_{\pi} \{u(\pi, u_r)\}$. In our binary-action setting, computing the optimum π^* is a well-understood problem. One general approach to tackle the problem is by first expressing u_s as a function of the posterior, then taking the concavification and evaluating it at the point representing the prior (Kamenica and Gentzkow, 2011).⁷

⁵ Otherwise, one can eliminate i from Ω .

⁶ Note that we assume here that Receiver breaks ties in Sender's favor. All our results are applicable for other tie-breaking rules.

⁷ Note that u_s equals the indicator function $\mathbf{1}_{p(s) \in A}$; since A is a half-space, the concavification of u_s can be computed and characterized.

2.1. Threshold signaling schemes

We now introduce a class of signaling schemes that can be used to optimally solve the setting with a binary-action Receiver in the standard model with Receiver's utility function known to Sender. Interestingly, this class plays a key part in our results for the adversarial model with unknown utility. We begin by describing the *knapsack method* for solving a persuasion problem; then we formulate the class of signaling schemes that arise from this approach.

The knapsack method In the special case of binary actions for Receiver, an alternative to the general concavification method is a greedy approach that reduces the problem to a fractional *knapsack* instance.⁸

The states are treated as the knapsack items, and their prior probabilities according to μ are the weights. Recall that every state induces utility $u_r(i, 1)$ to Receiver upon adoption. Sender's goal is to add to the knapsack a maximum-weight set of states (possibly – their fractions), while keeping Receiver's expected utility non-negative (the expectation is over a random state drawn with probability proportional to its weight in the knapsack). The knapsack instance can be solved greedily by sorting the states $i \in \Omega$ according to $u_r(i, 1)$ in a non-increasing order, and then continuously adding probability masses of states to the knapsack as long as the expected Receiver's utility stays non-negative. Such an approach has been adopted by Kolotilin (2015); Renault et al. (2017).

An equivalent way to present the knapsack approach is as follows. Let $i_1, \dots, i_n$ be an ordering of the states s.t. $u_r(i_n, 1) \geq u_r(i_{n-1}, 1) \geq \dots \geq u_r(i_1, 1)$. One can interpret the state space equipped with the prior distribution as drawing uniformly a real number in $[0, 1]$ – called the *real-valued state* – s.t. for every $m \in [n]$, all the realizations in the segment $(\sum_{l < m} \mu_{i_l}, \sum_{l \leq m} \mu_{i_l}]$ correspond to the state i_m .⁹ Now we define the notion of a *threshold signaling scheme*.

Definition 2.3. For every $t \in [0, 1]$, the t -threshold signaling scheme is a binary signaling scheme that reveals whether the real-valued state is below t or not.

Equivalently, for $t \in (0, 1]$, in the state space Ω , let $j \in [n]$ be s.t. $t \in (\sum_{l < j} \mu_{i_l}, \sum_{l \leq j} \mu_{i_l}]$. Then the t -threshold signaling scheme is a binary-signal scheme with a *high* and a *low* signal, s.t. the high signal is sent with probability 1 if $\omega > i_j$, with probability $\frac{\sum_{l \leq j} \mu_{i_l} - t}{\mu_{i_j}}$ if $\omega = i_j$, and with probability 0 if $\omega < i_j$. For $t = 0$, the 0-threshold signaling scheme is the scheme that reveals no information.

Given a threshold signaling scheme, we say of the states for which it sends the high (respectively, low) signal that they are *pooled together*. The knapsack approach is based on the following fact:

Fact 2.4 (See, e.g., Kolotilin, 2015; Renault et al., 2017). Every persuasion problem with a binary Receiver's action space admits an optimal x -threshold signaling scheme for some x . Moreover, for every $y \geq x$, Receiver adopts after observing the high signal in the y -threshold signaling scheme.

2.2. Robust approach

We depart from the standard model by assuming from now on that Receiver's utility function u_r is *unknown* to Sender. Besides Sender and Receiver, we introduce a third agent called *Adversary*. Given μ and π , Adversary aims to set Receiver's utility u_r in a way that makes the performance of π as bad as possible. Thus, Sender aims to design a signaling scheme with a worst-case guarantee – it should perform well for *all* possible Receiver's utility functions u_r .

Regret definition Here we formalize the (additive) regret minimization setting on which we mostly focus in this paper, as mentioned in Section 1. Fix a prior distribution μ . Consider $\pi^*(u_r)$, a signaling scheme that maximizes Sender's expected utility in the standard persuasion setting in which Receiver's utility u_r is known; denote by $u^*(u_r)$ the expected Sender's utility $\pi^*(u_r)$ yields. In our adversarial setting, given any signaling scheme π ignorant to u_r , Adversary aims to set u_r in a way that maximizes Sender's regret – that is, maximizes the difference between what Sender could have gotten had she known u_r and what she gets with π . In other words, Adversary tries to punish Sender for not choosing $\pi^*(u_r)$ as much as possible. The regret of π is, thus, defined as follows.

Definition 2.5. The *regret* of a signaling scheme π is $\text{Reg}(\pi) := \sup_{u_r} \{u^*(u_r) - u(\pi, u_r)\}$.

Sender aims to set π in a way that minimizes $\text{Reg}(\pi)$; it yields the following definition of the regret of the given persuasion setting – that is, the optimal regret for the given prior distribution and Sender's utility function.

⁸ In the *knapsack* problem, given are several items, each one having a *weight* and a *value*, and the goal is to pick a value-maximizing subset of the items subject to a weight capacity constraint; in the *fractional knapsack* problem, one is also allowed to pick fractions of items.

⁹ To demonstrate this, consider $n = 4$ states i_1, i_2, i_3, i_4 with prior $\mu(i_1, i_2, i_3, i_4) = (0.2, 0.3, 0.1, 0.4)$. Then, e.g., if the real-valued state is 0.55, the corresponding state from Ω is i_3 .

Definition 2.6. The *regret of the setting* is $\text{Reg} := \inf_{\pi} \text{Reg}(\pi)$.

The scale of regret Note that we always have $0 \leq \text{Reg} \leq 1$. Consider the following possibilities.

1. $\text{Reg} = 0$, i.e., Sender can ensure utility $u^*(u_r)$ without knowing u_r . Such a result is “too good to be true”, as indicated by our results.
2. $0 \ll \text{Reg} \ll 1$ even when the problem becomes large (the number of states grows to infinity), i.e., the additive loss from not knowing Receiver’s utility is strictly less than 1. Such a result can be viewed as positive.
3. $\text{Reg} \rightarrow 1$ as the problem becomes large; that is, in the worst case, Sender might lose her entire utility from not knowing u_r , as her expected utility might approach 0 while she could have gotten utility approaching 1 upon knowing u_r . Such a result is negative. One can also study the rate of convergence of Reg to 1 as a measure of how negative the result is.

Zero-sum game perspective Our setting can be analyzed via the following two-player zero-sum game G . The players are Adversary and Sender. Adversary’s pure strategies are mixtures over the functions $\{u_r\}$, while Sender’s pure strategies are the signaling schemes $\{\pi\}$.¹⁰ We remark that our proofs do *not* rely on the fact that G has a value. Indeed, to bound Reg from below by some bound v , we either describe explicitly a mixed strategy of Adversary ensuring $\text{Reg}(\pi) \geq v$ for every π , or show that for every choice of π by Sender, Adversary can ensure a regret of v . To bound Reg from above by some bound v , we explicitly describe a signaling scheme π s.t. $\text{Reg}(\pi) \leq v$. However, for completeness, we mention that Sion’s Minimax Theorem (Sion, 1958) holds in our setting, since the objective function is linear and lower semi-continuous as a function of π , and is linear and continuous as a function of the mixture over Receiver’s utility functions. Therefore, while we define Reg as $\inf_{\pi} \sup_{u_r} \{u^*(u_r) - u(\pi, u_r)\} = \inf_{\pi} \sup_{\Delta(u_r)} \{u^*(u_r) - u(\pi, u_r)\}$, it also equals $\sup_{\Delta(u_r)} \inf_{\pi} \{u^*(u_r) - u(\pi, u_r)\}$. That is, if Sender can ensure a regret of at most v after observing Adversary’s strategy, then she has a signaling scheme that ensures v even without knowing Adversary’s strategy.

3. Regret minimization

We consider three cases as described in Section 1: first, the case in which Sender has no information at all about Receiver’s utility (Subsection 3.1); second, the case in which Sender knows that Receiver’s utility upon adoption is monotonic as a function of the state (Subsection 3.2); finally, a case with a multidimensional state space and a Sender who knows that Receiver’s utility is monotonic in each dimension (Subsection 3.3). As mentioned in Subsection 1.2, in the first case we obtain a negative asymptotic result, while in the second case we have a positive result; in the third case, we get a positive result for product priors; however, a negative asymptotic result similar to the one from the first case holds for general priors.

3.1. Arbitrary utilities

Denote by Reg_{ARB} the regret for arbitrary utilities. That is, in Definition 2.5, the supremum is taken over all Receiver’s utilities. In the motivating example of selling a product from Section 1 – arbitrary utilities mean that the seller knows *absolutely nothing* about the buyer’s preferences among products. In this subsection, we start by characterizing the minimum regret for $n = 2$ states, as well as for $n = 3$ states with a uniform prior μ – in both cases, $\text{Reg}_{\text{ARB}} = \frac{1}{2}$. Unfortunately, such constant regret cannot be expected in general for arbitrary utilities. Indeed, our main result of this section shows that for any $n \geq 2$, the regret satisfies $1 - \frac{2}{\sqrt{n}} \leq \text{Reg}_{\text{ARB}} \leq 1 - \frac{1}{4n^2}$. In particular, $\text{Reg}_{\text{ARB}} \rightarrow_{n \rightarrow \infty} 1$. Intuitively, the more states there are, the less certainty about Receiver’s incentives Sender may have.

Observation 3.1. With arbitrary utilities, the set of all possible choices of $A(u_r)$ by Adversary coincides with the set of all the polytopes obtained by cutting $\Delta(\Omega)$ by a hyperplane.

This observation follows from Equation (1), as for a fixed u_r , the condition $\mathbb{E}_{\omega' \sim p(s)}[u_r(\omega', 1)] = 0$ specifies a hyperplane in $\Delta(\Omega)$ (see Fig. 1 for illustration).

For a binary state space we have the following result.

Proposition 3.2. For $n = 2$ and any prior μ : $\text{Reg}_{\text{ARB}} = \frac{1}{2}$.

Proof. Let us show first that $\text{Reg}_{\text{ARB}} \leq \frac{1}{2}$. For simplicity of notation, we identify a prior/posterior with the probability it specifies for $\omega = 1$. Assume w.l.o.g. that $\mu_1 \leq \frac{1}{2}$. Consider a binary-signal scheme π that induces the posteriors 0 and $2\mu_1$ with equal probability of $\frac{1}{2}$.¹¹ If $[0, 2\mu_1] \cap A \neq \emptyset$, then either $0 \in A$ or $2\mu_1 \in A$ (or both); thus, the adoption probability

¹⁰ Note that a mixture over signaling schemes is a signaling scheme.

¹¹ π is Bayes-plausible, as the expected posterior probability of $\omega = 1$ equals μ_1 ; therefore, by Remark 2.1, it indeed specifies a signaling scheme.




Fig. 1. An illustration of the adoption region for arbitrary utilities and $n = 3$.

upon π is at least $\frac{1}{2}$, and $\text{Reg}(\pi) \leq 1 - \frac{1}{2} = \frac{1}{2}$. Otherwise, A has the form $[t, 1]$ for some $t > 2\mu_1$. By Markov's inequality applied to $p(s)$, $\Pr_{s \sim \pi^*}[p(s) \in A] \leq \frac{\mu_1}{t} < \frac{1}{2}$, and again $\text{Reg}(\pi) \leq \frac{1}{2}$, as desired. Therefore, the regret is at most $\frac{1}{2}$ for every fixed u_r ; thus, it also holds for every mixture over Receiver's utilities.

Conversely, let us show that for every signaling scheme π , Adversary can pick a Receiver's utility function u_r – or, equivalently, an adoption region A – s.t. the regret would be at least $\frac{1}{2}$.

Indeed, fix π . If $\Pr_{s \sim \pi}[p(s) \in [0, \mu_1]] \leq \frac{1}{2}$, then Adversary can set $A = [0, \mu_1]$: on the one hand, $u^* = 1$ is achieved (if Sender knows A) by the no-information scheme in which $p(s) = \mu_1$ with probability 1, as $\mu_1 \in A$; on the other hand, $u(\pi, u_r) = \Pr_{s \sim \pi}[p(s) \in A] = \Pr_{s \sim \pi}[p(s) \in [0, \mu_1]] \leq \frac{1}{2}$. Hence, the regret is at least $\frac{1}{2}$, as desired.

It remains to consider the case in which $\Pr_{s \sim \pi}[p(s) \in [0, \mu_1]] > \frac{1}{2}$.

Set $\epsilon := \frac{1}{2} - \Pr_{s \sim \pi}[p(s) \in (\mu_1, 1]] \in (0, \frac{1}{2}]$ and let Adversary take $A = [\frac{\mu_1}{1-\epsilon}, 1]$; note that $0 < \mu_1 < \frac{\mu_1}{1-\epsilon} \leq \frac{\mu_1}{1-\frac{1}{2}} = 2\mu_1 \leq 1$, as $0 < \epsilon \leq \frac{1}{2}$ and $0 < \mu_1 \leq \frac{1}{2}$. Then $u^* \geq 1 - \epsilon$, as upon knowing A , Sender could have chosen a signaling scheme inducing $p(s) = 0$ with probability ϵ and $p(s) = \frac{\mu_1}{1-\epsilon}$ with probability $1 - \epsilon$ (note that this scheme is, indeed, Bayes-plausible). Since $u(\pi, u_r) = \Pr_{s \sim \pi}[p(s) \in A] \leq \Pr_{s \sim \pi}[p(s) \in (\mu_1, 1]] = \frac{1}{2} - \epsilon$, the regret is at least $(1 - \epsilon) - (\frac{1}{2} - \epsilon) = \frac{1}{2}$, as needed. $\square$

Remark 3.3. One can show that $\text{Reg}_{\text{ARB}} \geq \frac{1}{2}$ for any n and any μ using similar construction to $n = 2$.

Indeed, for a given π , Receiver should pick a hyperplane H passing through μ that does not pass through any atom in π besides, possibly, μ . Let Δ_1 be any of the two closed polytopes into which H divides $\Delta(\Omega)$. If $\Pr_{s \sim \pi}[p(s) \in \Delta_1] \leq \frac{1}{2}$, then by setting $A = \Delta_1$ Adversary ensures that Sender's expected utility is at most $\frac{1}{2}$, while the no-information scheme could have achieved utility of 1; and if $\Pr_{s \sim \pi}[p(s) \in \Delta_1] = \frac{1}{2} + \epsilon$ for some $\epsilon \in (0, \frac{1}{2}]$, then setting A to be a closed subset of $\Delta(\Omega) \setminus \Delta_1$ specified by a close enough hyperplane H' to H ensures that while $u(\pi, u_r) \leq \Pr_{s \sim \pi}[p(s) \notin \Delta_1] = \frac{1}{2} - \epsilon$, Sender could have achieved utility of $1 - \epsilon$ had she known A by an appropriate binary-signal scheme.

Another case in which we characterize the optimal persuasion is that of ternary state space and uniform prior.

Proposition 3.4. For $n = 3$ and $\mu = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$: $\text{Reg}_{\text{ARB}} = \frac{1}{2}$.

Although Proposition 3.4 deals with a specific case, its proof is not trivial at all. The explicit characterizations of the regret in case of ternary state space with non-uniform prior and in case of any larger number of states remain open problems. The formal proof of Proposition 3.4 is relegated to Appendix A. We briefly describe here the main proof ideas.

Proof idea of Proposition 3.4. Note that $\Delta(\{3\})$ is a 2-dimensional object. Each of its points is uniquely representable as a convex combination of its extreme points; the weights in this representation – a.k.a. the *Barycentric coordinates* of the point – specify the probability masses the corresponding posterior assigns to the states.

By Remark 3.3, Adversary can ensure a regret of at least $\frac{1}{2}$. To achieve such a regret, Adversary partitions the simplex $\Delta(\{3\})$ into two regions by drawing a line through the prior μ (which is the centroid of the simplex). Thereafter, Adversary sets A to be the half-space that has a weight not above $\frac{1}{2}$ according to π (for simplicity, here we ignore the possibility that π has an atom on μ).

The above argument can be applied for every line passing through μ which does not contain atoms of π . This hints at the desirable distributions of posteriors for a regret-minimizing π : it should assign an equal weight of $\frac{1}{2}$ for every half-space that contains μ on the boundary. Another useful intuition from Proposition 3.2 is that it is worthwhile for Sender to “push” the weight to the boundary. This brings us to the (somewhat educated) guess of considering the *unique* distribution that is supported on the boundary of $\Delta(\{3\})$, $\{(p_1, p_2, p_3) \in \Delta(\{3\}) : p_1 = 0 \text{ or } p_2 = 0 \text{ or } p_3 = 0\}$, which assigns probability weight of exactly $\frac{1}{2}$ to every half-space containing μ as a boundary point. The complicated part of the proof is showing that this signaling scheme indeed guarantees a regret of $\frac{1}{2}$ to Sender. Namely, to show that Adversary cannot get a regret greater than $\frac{1}{2}$ by choosing the adoption region to be a half-space that does not contain μ . This follows from geometric arguments and direct computations using analytic geometry. $\square$




Fig. 2. The regret as a function of the prior probability of the highest state.

The fact that the regret does not increase from binary to ternary state space is somewhat misleading. The following theorem provides a negative result for large state spaces and an arbitrary prior.

Theorem 3.5. For every number of states n and any prior μ : $1 - \frac{2}{\sqrt{n}} \leq \text{Reg}_{\text{ARB}} \leq 1 - \frac{1}{4n^2}$.

Theorem 3.5 indicates that not knowing Receiver's utility in large state spaces might be very costly for Sender. For every signaling scheme, there are Receiver's utilities for which adoption occurs with probability almost 0, while knowing Receiver's utility allows Sender to get adoption probability of almost 1. This negative result is not very surprising, as we focus here on the most general setting and assume nothing on Receiver's utility. Now we describe the proof idea; the full proof appears in Appendix B.

Proof idea of Theorem 3.5. We prove that $\text{Reg}_{\text{ARB}} \leq 1 - \frac{1}{4n^2}$ by describing an explicit signaling scheme ensuring such a regret bound. The more interesting part of the theorem is $\text{Reg}_{\text{ARB}} \geq 1 - \frac{2}{\sqrt{n}}$. To show this part, we consider a specific scenario in which Sender knows that most states, $\sim (n - \sqrt{n})$, are *normal* – they yield a constant negative Receiver's utility upon adoption. Among the remaining $\sim \sqrt{n}$ states, a single *good* state with very high utility is hidden, while the other states are *bad*; for clarity of exposition, assume that for bad states Receiver's utility upon adoption is $-\infty$. Had Sender known which state is the good one, she could have pooled it together with all the normal states; it would have been incentive-compatible for Receiver to adopt after receiving the signal that the state is not bad.

However, Adversary selects the state types uniformly at random; in particular, our ignorant Sender does not know which state is good. The support of any posterior distribution either with high probability contains a bad state or with high probability does not contain the good state, which would cause Receiver to reject. $\square$

3.2. Monotonic utilities

Let Reg_{MON} be the regret for monotonic utilities. That is, in Definition 2.5, the supremum is over all Receiver's utilities that are non-decreasing in the state upon adoption: $u_r(n, 1) \geq u_r(n-1, 1) \geq \dots \geq u_r(1, 1)$. In this setting, we prove a positive result that provides a full explicit characterization of the regret, for every number of states n and for every prior μ . We saw in the previous subsection that n plays a significant role for arbitrary utilities. In particular, as the number of states increases, so does the uncertainty of Sender in the arbitrary utility case. Hence, one might expect the regret to always increase with the number of states. Surprisingly, our result shows that this intuition is wrong: Sender's partial knowledge on the monotonicity of Receiver's utility turns out to be sufficient to have as good guarantees as for a binary state space.

Recall that μ_n is the prior probability of state n ; for monotonic utilities, Receiver gets the (weakly) largest utility upon adoption in this state. The following theorem gives a full characterization of Reg_{MON} in terms of μ_n . In particular, Reg_{MON} does not depend on n . The optimal regret as a function of μ_n is demonstrated in Fig. 2.

Theorem 3.6. For every number of states n and every prior $\mu = (\mu_1, \dots, \mu_n)$:

$$\text{Reg}_{\text{MON}} = \begin{cases} \frac{1}{e} & \text{if } \mu_n \leq \frac{1}{e} \\ -\mu_n \ln \mu_n & \text{if } \mu_n > \frac{1}{e}. \end{cases}$$

An intuition for why parameters other than μ_n turn out to be irrelevant appears in the proof sketch below; the relevance of μ_n is intuitively connected to the observation that Adversary must choose $u_r(n, 1) \geq 0$ to achieve a regret greater than 0 (as otherwise, Receiver would never adopt). Sender can utilize this observation in her favor: e.g., truthfully revealing that $\omega = n$ when it occurs yields her a utility of at least μ_n . As we shall see, she can utilize it in a more clever manner, which even further decreases the regret. A possible interpretation of Theorem 3.6 might be that unless the highest state is very likely (its prior probability is greater than $1/e$), the regret equals $1/e$ regardless of the number of states and the prior. Below we present the proof idea; for the full proof, see Appendix C. That appendix also contains an explicit description of the optimal Sender's and Adversary's strategies (Lemma C.1).




Fig. 3. The density function of Sender's threshold when $\mu_n = 1/4$.

Proof idea of Theorem 3.6. The proof relies on the knapsack characterization of optimal signaling schemes in the standard persuasion model. The characterization of optimal policies by Kolotilin (2015); Renault et al. (2017) states that every persuasion problem in the standard setting with binary-action Receiver has an optimal x -threshold scheme for some x (see Definition 2.3).

In our setting, Receiver's utility function is unknown to Sender; hence, the optimal threshold x is unknown. Nevertheless, we make an educated guess that threshold policies remain useful – and in fact, regret-minimizing – also in our setting. More concretely, we consider signaling schemes in which the threshold y is drawn *at random*, and thereafter the y -threshold scheme is implemented. After restricting Sender's strategies to the class of (mixed) threshold schemes, we can also view Adversary's choice of utility u_r simply as a choice of an optimal threshold $x \in [0, 1]$ (see Fact 2.4). Namely, instead of choosing a utility for Receiver, Adversary chooses the threshold x that is optimal in the standard persuasion model. This reduces the complicated original zero-sum game G (see Subsection 2.2), in which Sender chooses a signaling scheme and adversary chooses a distribution over possible functions u_r , to a much simpler zero-sum game G_{μ_n} , with both players simply choosing thresholds $x, y \in [0, 1]$.

The utility in G_{μ_n} is given by:

$$g(x, y) := (1 - x) - (1 - y)\mathbf{1}_{y \geq x},$$

where $1 - x$ is u^* (i.e., the optimal utility in the standard persuasion model) and $(1 - y)\mathbf{1}_{y \geq x}$ is the utility of our ignorant Sender: if $y \geq x$, then Receiver adopts after observing the high signal and Sender gets utility of $1 - y$; if $y < x$, then Receiver rejects after observing the high signal (recall that y is Receiver's indifference point) and Sender gets a utility of 0; anyway, Receiver rejects when observing the low signal.

In fact, not all thresholds $x \in [0, 1]$ might be optimal for some utility, but only $x \in [0, 1 - \mu_n]$. Indeed, the highest-utility state should be included entirely in the knapsack, as otherwise adoption would never occur and the regret would be zero. As Sender's utility is decreasing in y for $y \geq x$, and we know that the optimal threshold is in $[0, 1 - \mu_n]$, we can restrict ourselves to considering Sender's and Adversary's thresholds x, y from $[0, 1 - \mu_n]$.

The analysis of G_{μ_n} (see Lemma C.1 in Appendix C) leads to the game value that appears in the theorem, and also provides the optimal strategies for Sender and Adversary.

After having the (allegedly) optimal strategies of both players in hand, we still have to verify that Sender cannot ensure a smaller regret than the value of G_{μ_n} . To this end, we show that for a simple, yet optimal strategy of Adversary, Sender cannot gain more than the value of G_{μ_n} even by choosing arbitrary (not necessarily threshold) schemes. $\square$

It is interesting to note that the regret-minimizing signaling scheme randomizes the threshold over a continuum-sized support $[0, \min\{1 - \mu_n, 1 - 1/e\}]$, and, in particular, uses a continuum of signals; this stands in a sharp contrast to standard persuasion in which binary signals suffice to get the optimal utility. Inside the continuum-sized support, the density function is given by $f(y) := \frac{1}{1-y}$ and, in particular, lower thresholds (closer to no-information) are chosen with lower probability.¹² Fig. 3 demonstrates the density function of Sender's thresholds for $\mu_n = 1/4$. Note that our optimal robust signaling scheme is constructive and has a closed-form description, and thus is polynomial-time computable. For example, for $\mu_n \leq 1/e$, one can implement the optimal robust scheme – characterized in Lemma C.1 – by sampling $z \sim U[0, 1]$ and adopting the $F^{-1}(z)$ -threshold scheme, where $F(y) = -\ln(1 - y)$ for $y \in [0, 1 - 1/e]$.

3.3. Multidimensional monotonic utilities

In the setting studied in this subsection, we denote the regret by $\text{Reg}_{\text{MON-MD}}$. We assume that $\Omega = \times_{j \in [k]} \Omega_j$, where $\Omega_j = [n_j]$. We require Receiver's utility upon adoption to be monotonic in every dimension $j \in [k]$. Namely, $\forall 1 \leq j \leq$

¹² To be precise, the threshold distribution also has an atom of weight $1 + \ln \mu_n$ on $1 - \mu_n$ (i.e., Sender reveals whether ω is the highest-utility state n or not) if $\mu_n > 1/e$.

$k \forall \omega'_j \leq \omega''_j \forall \omega'_{-j} : u_r(\omega'_j \omega'_{-j}, 1) \leq u_r(\omega''_j \omega'_{-j}, 1)$. We refer to the dimension k as a constant. Indeed, for the interpretation of the dimension as the number of product attributes, it is natural to assume that k is not too large. The parameters that determine the size of the problem are $n_1, \dots, n_k$ (i.e., the numbers of quality levels of each attribute/dimension). We show that the prior plays a significant role in determining whether a constant regret (i.e., bounded away from one) can be guaranteed: on the one hand, for arbitrary priors, the minimal regret might not be bounded away from 1 even for $k = 2$; on the other hand, for *product* priors, a constant regret can be guaranteed.

For general priors, we have the following corollary from the arbitrary utility case.

Corollary 3.7. *For $k = 2$ and $n_1 = n_2 = m$ (for any $m \geq 1$), there exists a prior $\mu_m \in \Delta([m]^2)$ for which $\text{Reg}_{\text{MON-MD}}(m) = 1 - O\left(\frac{1}{\sqrt{m}}\right)$.*

Corollary 3.7 follows from Theorem 3.5 and the fact that one can assign arbitrarily small probability weights in the prior to all the states except for the m states $\{(i, j) \in \Delta([m]^2) : i + j = m + 1\}$. On these m states, the monotonicity in each dimension imposes no restriction, while the remaining states have tiny weights; thus, their effect on the regret is negligible.

Therefore, to obtain positive results, one must restrict attention to particular classes of priors. One natural class is *product* priors: $\mu = \times_{j \in [k]} \mu_j$ for some $\mu_j \in \Delta([n_j])$. For this class of priors, we have the following result, which might be viewed as positive for small values of k .

Proposition 3.8. *For every k , every sequence of attribute quality level amounts $n_1, \dots, n_k$ and every product prior $\mu = \times_{j \in [k]} \mu_j$: $\text{Reg}_{\text{MON-MD}} \leq 1 - 2^{-k}$.*

The full proof appears in Appendix D; here we only sketch the main ideas. We conjecture that a similar positive result can be obtained not only for product priors, but also for positively-correlated priors.

Proof idea of Proposition 3.8. We present a simple signaling scheme that achieves the desired regret: Sender reveals whether the product quality in all the attributes is above median, where the median is calculated according to the prior μ_j . Depending on u_r , for such a scheme one of the following holds:

1. Whenever all the attribute qualities are above median, adoption occurs – yielding expected Sender's utility of at least 2^{-k} .
2. Whenever all the attribute qualities are below median, adoption does not occur – implying that even had Sender known u_r , her expected utility could not have exceeded $1 - 2^{-k}$. $\square$

4. Discussion and future work

So far, we have adopted the adversarial regret minimization approach to Bayesian persuasion. In Section 3, we prove that while in the most general case the regret approaches 1 as the number of states grows large, Sender's knowledge of Receiver's ordinal preferences upon adoption ensures that the regret is at most $1/e$. We further provide an explicit formula for the regret, and describe the optimal Sender's and Adversary's strategies. We also study multidimensional monotonic utilities and upper bound the regret in this setting for product priors, while showing that for general priors the regret might be close to 1. In the next subsection, we discuss how our main results change when considering the adversarial approximation ratio instead of regret.

4.1. Adversarial approximation

In this subsection, the performance of a signaling scheme π is measured by $\frac{u(\pi, u_r)}{u^*(u_r)}$, i.e., the *ratio* between the utility Sender achieves with the scheme π and the optimal utility Sender can get upon knowing u_r . If $u^*(u_r) = 0$ (i.e., even the knowledgeable Sender cannot achieve any positive utility), we define the ratio to be 1.

Analogously to Definitions 2.5 and 2.6 of the additive regret, define the adversarial approximation guarantee via:

$$\text{Apr}(\pi) := \inf_{u_r} \left\{ \frac{u(\pi, u_r)}{u^*(u_r)} \right\} \text{ and } \text{Apr} := \sup_{\pi} \text{Apr}(\pi).$$

Similarly to the regret, we extend the definition of Apr to Apr_{ARB} , Apr_{MON} and $\text{Apr}_{\text{MON-MD}}$ according to the class of considered Receiver's utility functions.

Note that $0 \leq \text{Apr} \leq 1$, where $\text{Apr} \approx 0$ means that only a negligible fraction of the potential Sender's utility (i.e., with the knowledge of u_r) can be guaranteed. Our results in this section are largely negative, motivating our focus on regret minimization.

We start with the general Proposition 4.1 that connects the notions of regret (Reg) and adversarial approximation guarantee (Apr).

Proposition 4.1. $\text{Apr} \leq \frac{1}{\text{Reg}} - 1$.

This fact holds in any adversarial setting with utilities in $[0, 1]$ and not only in our context of persuasion.

Proof of Proposition 4.1. By Definitions 2.5 and 2.6, for every signaling scheme π and every $\epsilon > 0$, there exists u_r^ϵ s.t. $u^*(u_r^\epsilon) - u(\pi, u_r^\epsilon) \geq \text{Reg} - \epsilon$. In particular, $u^*(u_r^\epsilon) \geq \text{Reg} - \epsilon$ and $u(\pi, u_r^\epsilon) \leq 1 - \text{Reg} + \epsilon$. Furthermore, $\text{Apr}(\pi) \leq \frac{u(\pi, u_r^\epsilon)}{u^*(u_r^\epsilon)} \leq \frac{1 - \text{Reg} + \epsilon}{\text{Reg} - \epsilon}$. Taking $\epsilon \rightarrow 0$ yields that $\text{Apr}(\pi) \leq \frac{1 - \text{Reg}}{\text{Reg}} = \frac{1}{\text{Reg}} - 1$. As it holds for every π , the proposition follows. $\square$

Proposition 4.1 allows to translate all the negative results on the regret into negative results on adversarial approximation: if $\text{Reg} \rightarrow 1$, then $\text{Apr} \rightarrow 0$, with at least the same convergence rate. Concretely, the negative results for arbitrary utilities (Theorem 3.5) and multidimensional monotonic utilities (Corollary 3.7) have analogues in the adversarial approximation approach. It suggests that the adversarial ratio approach might be more pessimistic than regret minimization.

A natural question is whether our main positive result (Theorem 3.6) remains valid for the adversarial approximation setting. Namely, can a constant fraction of the optimal Sender's utility be guaranteed without knowing the cardinal Receiver's preferences, but only the ordinal ones (i.e., monotonicity)? Our next result shows that the answer is negative. Furthermore, we provide an exact characterization for the value in this case, for every number of states n and every prior μ .

Theorem 4.2. For every number of states n and every prior $\mu = (\mu_1, \dots, \mu_n)$: $\text{Apr}_{\text{MON}} = \frac{1}{1 + \ln \frac{1}{\mu_n}}$.

This result indicates that for large state spaces (under the natural assumption that $\mu_n \rightarrow_{n \rightarrow \infty} 0$), a constant fraction of the optimal utility *cannot* be guaranteed. Namely, unlike the regret approach in which Sender can ensure a regret bounded away from 1, Sender cannot guarantee a constant approximation ratio. The proof idea is similar to Theorem 3.6, but the simpler zero-sum game in which both players choose thresholds (real numbers in $[0, 1 - \mu_n]$) has a different value. For the proof, see Appendix E.

4.2. Future work

Some problems remain open for the setting that we study. In particular, the exact regret value in the arbitrary utility case for more than two states is unknown – the only exception is the case of uniform prior with ternary state space. The precise asymptotic convergence rate of Reg_{ARB} to 1 is also unknown to us – we only bound it between $\theta\left(\frac{1}{\sqrt{n}}\right)$ and $\theta\left(\frac{1}{n^2}\right)$. Furthermore, our setting has several natural extensions – e.g., to more than one Receiver.

Moreover, the minimax robust approach has not been discussed in this paper. Recall that in the minimax approach, Sender aims to maximize her expected utility over the signaling scheme and the unknown Adversary's strategy – that is, there is no benchmark for the performance of the signaling scheme. Unfortunately, in the absence of any Sender's knowledge about Receiver's utility, the minimax approach is meaningless: Adversary can choose u_r s.t. Receiver never adopts, yielding the minimax value of 0. An interesting question we leave for future research is to find a well-motivated Sender's *partial knowledge* about u_r , which would make the minimax approach meaningful. Similarly, it is interesting to understand whether extra assumptions on Receiver's utility, rather than just monotonicity, may yield a constant adversarial approximation ratio. One such setting might consider Sender ignorant of Receiver's private exogenous information (Hurwicz and Shapiro, 1978; Emek et al., 2012; Li et al., 2020; Dworczak and Pavan, 2020; Hu and Weng, 2021). Extending the results to adversarial choice of the prior distribution (Kosterina, 2022) is another possible research direction.

Acknowledgments

Yakov Babichenko is supported by the Binational Science Foundation BSF grant no. 2018397 and by the German-Israeli Foundation for Scientific Research and Development GIF grant no. I-2526-407.6/2019. Inbal Talgam-Cohen is a Taub Fellow supported by the Taub Family Foundation. This research is supported by the Israel Science Foundation grant no. 336/18. Yakov and Inbal are supported by the Binational Science Foundation BSF grant no. 2021680. Haifeng Xu is supported by a Google Faculty Research Award and an NSF grant no. CCF-2132506. Konstantin Zabarnyi is supported by a PBC scholarship for Ph.D. students in data science. The authors thank the anonymous reviewers and the editor for their helpful suggestions on improving the paper.




Fig. 4. Parameterization of the points.

Appendix A. Proof of Proposition 3.4

Proof of Proposition 3.4. By Remark 3.3, it is enough to show that $\text{Reg}_{\text{ARB}}(\pi) \leq \frac{1}{2}$. Consider π supported on the boundary of $\Delta(\Omega)$ that assigns to every line segment on the boundary probability proportional to the angle upon which this segment is seen from μ . We shall prove, using analytic geometry, that for this π , $\text{Reg}_{\text{ARB}}(\pi) = \frac{1}{2}$. Let $\text{at} : (-\infty, \infty] \rightarrow [0, \pi)$ be defined as $\text{at}(x) := \arctan(x) + \pi$ for $x \in (-\infty, 0)$, $\text{at}(x) := \arctan(x)$ for $x \in [0, \infty)$ and $\text{at}(\infty) := \frac{\pi}{2}$.

Fix A , and let l be the line (single-dimensional hyperplane) separating A from $\Delta(\Omega) \setminus A$. Note that $\Delta(\Omega)$ is contained in a two-dimensional plane. Represent the elements of $\Delta(\Omega)$ by Cartesian coordinates (rather than the probabilities they assign to elements of Ω). Parameterize its extreme points by $\omega_1 = (0, 0)$, $\omega_2 = (1, 0)$, $\omega_3 = (\frac{1}{2}, \frac{\sqrt{3}}{2})$.¹³ Then $\mu = (\frac{1}{2}, \frac{\sqrt{3}}{6})$. If the separating line l is disjoint to $\Delta(\Omega)$ – the regret is 0. Therefore, assume that it cuts at least two line segments from the boundary of $\Delta(\Omega)$ (possibly at the same point). Assume w.l.o.g. that l cuts $\omega_1\omega_2$ (the line connecting ω_1 and ω_2) and $\omega_1\omega_3$. Then l contains the points $D := (d, \sqrt{3}d)$ and $E := (e, 0)$ for some $0 \leq d \leq \frac{1}{2}$ and $0 \leq e \leq 1$. Assume that $\mu \notin A$ – otherwise, the adoption probability assured by π is at least $\frac{1}{2}$ and the regret is at most $\frac{1}{2}$.

Consider now π^* , an optimal signaling scheme for a Sender who knows A . By pooling together all the probability mass in A and all the probability mass in $\Delta(\Omega) \setminus A$ (separately), assume w.l.o.g. that $|\text{supp}(\pi^*)| = 2$.¹⁴

Let s be the line segment connecting the two support elements of π^* . By Bayes-plausibility, it must pass through μ . As π^* is optimal, one can assume w.l.o.g. that one support element is on l , while the other is on the boundary of $\Delta(\Omega)$ (outside A); furthermore, μ should divide s in the maximal possible ratio. A straightforward (and a well-known) trigonometric exercise shows that the ratio in which a variable line segment s passing through a fixed point, with endpoints on two fixed line segments s_1, s_2 , is divided, is maximized when one of the endpoints of s coincides with one of the endpoints of s_1, s_2 . Therefore, one of the endpoints of s is either an intersection point of l with the boundary of $\Delta(\Omega)$ or a vertex of $\Delta(\Omega)$.

Assume first that exactly one extreme point of $\Delta(\Omega)$ is in A , and assume w.l.o.g. that $\omega_1 \in A$ (that is, A is set by l as in Fig. 4). Qualitatively, we have 2 different cases to check: either $E \in \text{supp}(\pi^*)$ or $\omega_2 \in \text{supp}(\pi^*)$. Denote the second endpoint of s (that is, the support element of $\text{supp}(\pi^*)$ different from E, ω_2) by F . Note that if $E \in \text{supp}(\pi^*)$, then either $F \in \omega_2\omega_3$ or $F \in \omega_1\omega_3$.¹⁵

Consider first the case $E = (e, 0) \in \text{supp}(\pi^*)$ and $F \in \omega_2\omega_3$ (see Fig. 5). That is, $F = (f, \sqrt{3}(1-f))$, for $f = \frac{3-5e}{4-6e}$; note that as $F \in \omega_2\omega_3$ and $\mu \notin A$, we have $0 \leq e \leq \frac{1}{2}$. The regret of π for our fixed A is:

$$\frac{|\mu F|}{|EF|} - \frac{\angle D\mu E}{2\pi} = \frac{1}{3-3e} - \frac{1}{2\pi} \text{at} \left(\frac{\frac{\sqrt{3}/6}{1/2-e} - \frac{\sqrt{3}/6-\sqrt{3}d}{1/2-d}}{1 + \frac{\sqrt{3}/6}{1/2-e} \cdot \frac{\sqrt{3}/6-\sqrt{3}d}{1/2-d}} \right).$$

For a fixed e , the above expression is maximal for $d = 0$.¹⁶ Therefore (upon algebraic simplifications), it is enough to check that:



$$\frac{1}{3-3e} - \frac{1}{2\pi} \text{at} \left(\sqrt{3} \frac{e}{2-3e} \right) \leq \frac{1}{2} \iff \text{at} \left(\sqrt{3} \frac{e}{2-3e} \right) - \frac{3e-1}{3-3e} \pi \geq 0.$$

¹³ We assume that the sides of the triangle $\Delta(\Omega)$ are of length 1 when in fact they are of length $\sqrt{2}$; however, homothety (stretch) on $\Delta(\Omega)$ does not affect the regret.

¹⁴ $|\text{supp}(\pi^*)| > 1$ since we assume that $\mu \notin A$.

¹⁵ F cannot be on $\omega_1\omega_2$, as s passes through μ .

¹⁶ One can obtain this step geometrically: setting $D = \omega_1$ does not affect the first term and does not increase the subtracted term.

Fig. 5. An illustration of the case $E \in \text{supp}(\pi^*)$, $F \in \omega_2\omega_3$.Fig. 6. An illustration of the case $E \in \text{supp}(\pi^*)$, $F \in \omega_1\omega_3$.

Note that the above expression is continuously differentiable for $e \in [0, \frac{1}{2}]$, and its derivative is¹⁷:

$$\sqrt{3} \frac{1 \cdot (2-3e) - e \cdot (-3)}{(2-3e)^2} \cdot \frac{1}{1 + \left(\sqrt{3} \frac{e}{2-3e}\right)^2} - \frac{3(3-3e) - (3e-1) \cdot (-3)}{(3-3e)^2} \pi =$$

$$\frac{2\sqrt{3}}{3e^2 + (2-3e)^2} - \frac{6\pi}{(3-3e)^2} < 0,$$

where the last transition follows from: $2\sqrt{3}(3-3e)^2 - 6\pi(3e^2 + (2-3e)^2) = (18\sqrt{3} - 72\pi)e^2 + e(72\pi - 36\sqrt{3}) + (18\sqrt{3} - 24\pi) = (18\sqrt{3} - 24\pi)(1-e)^2 + 24\pi e(1-2e) \leq_{e_{\text{OPT}} = \frac{2\pi - \sqrt{3}}{4\pi - \sqrt{3}}} \frac{4\pi^2(18\sqrt{3} - 24\pi)}{(4\pi - \sqrt{3})^2} + 24\pi \frac{2\pi\sqrt{3} - 3}{(4\pi - \sqrt{3})^2} < 0$.

Thus:

$$\text{at}\left(\sqrt{3} \frac{e}{2-3e}\right) - \frac{3e-1}{3-3e} \pi \geq_{e_{\text{OPT}} = \frac{1}{2}} \text{at}\left(\sqrt{3} \frac{0.5}{2-1.5}\right) - \frac{1.5-1}{3-1.5} \pi = \text{at}\left(\sqrt{3}\right) - \frac{\pi}{3} = 0,$$

as needed.

Now, let us still assume that $E = (e, 0) \in \text{supp}(\pi^*)$, but $F \in \omega_1\omega_3$ (see Fig. 6). That is, $F = (f, \sqrt{3} \cdot f)$ for $f = \frac{e}{6e-2}$. Note that $\frac{1}{2} \leq e \leq 1$ (otherwise, F is outside the line segment $\omega_1\omega_3$, a contradiction). Just as in the previous case, we obtain that for a fixed e , the regret is maximized when $d = 0$. The regret equals for $d = 0$ (here we assume that $e \neq 1$; for $e = 1 \implies E = \omega_2$, the adoption probability according to π^* is $\frac{1}{3}$; thus, the regret is smaller than $\frac{1}{2}$):

$$\frac{|\mu F|}{|EF|} - \frac{\angle D\mu E}{2\pi} = \frac{1}{3e} - \frac{1}{2\pi} \text{at}\left(\frac{\frac{\sqrt{3}/6}{1/2-e} - \frac{\sqrt{3}/6 - \sqrt{3}d}{1/2-d}}{1 + \frac{\sqrt{3}/6}{1/2-e} \cdot \frac{\sqrt{3}/6 - \sqrt{3}d}{1/2-d}}\right) \stackrel{d=0}{=} \frac{1}{3e} - \frac{1}{2\pi} \text{at}\left(\sqrt{3} \frac{e}{2-3e}\right).$$

To prove that this expression is at most $\frac{1}{2}$, it is enough to show that:

$$\text{at}\left(\sqrt{3} \frac{e}{2-3e}\right) - \frac{2-3e}{3e} \pi \geq 0.$$

¹⁷ For $e \in [0, \frac{1}{2}]$, the angle $\angle D\mu E = \angle \omega_1\mu E$ is acute; thus, $\text{at}\left(\sqrt{3} \frac{e}{2-3e}\right) = \arctan\left(\sqrt{3} \frac{e}{2-3e}\right)$.




Fig. 7. An illustration of the case $E \notin \text{supp}(\pi^*)$.

The last expression is increasing in e . Therefore, the minimum is attained for $e = \frac{1}{2}$, for which the condition holds:

$$\text{at}\left(\sqrt{3} \frac{0.5}{2-1.5}\right) - \frac{2-1.5}{1.5}\pi = \text{at}(\sqrt{3}) - \frac{\pi}{3} = 0.$$

Now we should consider the case in which $E \notin \text{supp}(\pi^*)$ (see Fig. 7). Then $\omega_2 = (1, 0) \in \text{supp}(\pi^*)$ and $E \neq \omega_2$ (i.e., $e \neq 1$). As $E \in l$, $\omega_1 \in A$ and l separates A from $\Delta(\Omega) \setminus A$, we must have $\omega_2 \notin A$. Since the second endpoint of s , F , must belong to A – it must be on l . That is, $F = \left(f, \frac{1-f}{\sqrt{3}}\right)$ for $f = \frac{e-d-3de}{e-4d}$. Note that one must have $\frac{1}{4} \leq d \leq \frac{1}{2}$ (otherwise, F cannot be on the segment DE) and $0 \leq e \leq \frac{2d}{6d-1}$ (as $f \leq \frac{1}{2}$, since $\mu \notin A$). The regret is:

$$\frac{|\mu\omega_2|}{|F\omega_2|} - \frac{\angle D\mu E}{2\pi}.$$

Fix F and let D, E vary so that DE passes through F , while keeping $\frac{1}{4} \leq d \leq \frac{1}{2}$ and $0 \leq e \leq \frac{2d}{6d-1}$. Consider the regret as a function of d . Note that $e = \frac{d(4f-1)}{f+3d-1}$. The regret equals:

$$\begin{aligned} \frac{|\mu\omega_2|}{|F\omega_2|} - \frac{\angle D\mu E}{2\pi} &= \frac{1}{2(1-f)} - \frac{1}{2\pi} \text{at}\left(\frac{\frac{\sqrt{3}/6}{1/2-e} - \frac{\sqrt{3}/6 - \sqrt{3}d}{1/2-d}}{1 + \frac{\sqrt{3}/6}{1/2-e} \cdot \frac{\sqrt{3}/6 - \sqrt{3}d}{1/2-d}}\right) = \\ &= \frac{1}{2(1-f)} - \frac{1}{2\pi} \text{at}\left(\sqrt{3} \cdot \frac{d(e-1/3) - e/6}{d(1-e) + e/2 - 1/3}\right) = \\ &= \frac{1}{2(1-f)} - \frac{1}{2\pi} \text{at}\left[\sqrt{3} \cdot \frac{d(d(4f-1) - \frac{f+3d-1}{3}) - \frac{d(4f-1)}{6}}{d(f+3d-1-d(4f-1)) + \frac{d(4f-1)}{2} - \frac{f+3d-1}{3}}\right] = \\ &= \frac{1}{2(1-f)} - \frac{1}{2\pi} \text{at}\left[2\sqrt{3}(2f-1) \frac{d^2 - \frac{d}{4}}{d^2(4(1-f)) + d(3f - \frac{5}{2}) + \frac{1-f}{3}}\right]. \end{aligned}$$


The regret is continuous as a function of d ; its derivative w.r.t. d has the same sign as (apart from finitely many values of d for which the derivative of the regret is not defined):

$$\begin{aligned} & - \left(2d - \frac{1}{4}\right) \left[d^2(4(1-f)) + d\left(3f - \frac{5}{2}\right) + \frac{1-f}{3} \right] + \left(d^2 - \frac{d}{4}\right) \left[8d(1-f) + \left(3f - \frac{5}{2}\right) \right] = \\ & d^2 \left(\frac{3}{2} - 2f \right) - d \left(\frac{2}{3}(1-f) \right) + \frac{1-f}{12}. \end{aligned}$$

Recall that $\frac{1}{4} \leq d \leq \frac{1}{2}$ and $\frac{1}{4} \leq f \leq \frac{1}{2}$ (because $\omega_2 F$ contains μ and F is inside $\Delta(\Omega)$). The above function is decreasing for $\frac{1}{4} \leq d \leq \frac{2(1-f)}{3(3-4f)}$ and is increasing for $\frac{2(1-f)}{3(3-4f)} \leq d \leq \frac{1}{2}$.¹⁸ Furthermore, its value at $d = \frac{1}{4}$ is:

$$\frac{1}{16} \left(\frac{3}{2} - 2f \right) - \frac{1}{4} \cdot \frac{2}{3} (1-f) + \frac{1-f}{12} = \frac{1-4f}{96} \leq 0.$$

¹⁸ Note that $\frac{1}{4} \leq \frac{2(1-f)}{3(3-4f)} \leq \frac{1}{2}$ for $\frac{1}{4} \leq f \leq \frac{1}{2}$.

Fig. 8. An illustration of the case $\omega_2, \omega_3 \in A$.

Therefore, for a fixed f , the regret is maximized for a boundary value of d : either for $d = \frac{1}{4}$ or for $d = \frac{1}{2}$. If $d = \frac{1}{4}$, then necessarily $F = D$ and the regret is:

$$\frac{|\mu\omega_2|}{|D\omega_2|} - \frac{\angle D\mu E}{2\pi} = \frac{2}{3} - \frac{\angle D\mu E}{2\pi} \leq \frac{2}{3} - \frac{\pi/3}{2\pi} = \frac{1}{2},$$

as needed. If $d = \frac{1}{2}$, then the value of the regret (after algebraic simplifications) is:

$$\frac{1}{2(1-f)} - \frac{1}{2\pi} \text{at} \left(\frac{3\sqrt{3}(2f-1)}{2f+1} \right)$$

The derivative of this expression is (note that it is well-defined for every $\frac{1}{4} \leq f \leq \frac{1}{2}$):

$$\begin{aligned} & \frac{1}{2(1-f)^2} - \frac{3\sqrt{3}}{2\pi} \cdot \frac{1}{1 + \left(\frac{3\sqrt{3}(2f-1)}{2f+1} \right)^2} \cdot \frac{2(2f+1) - 2(2f-1)}{(2f+1)^2} = \\ & \frac{1}{2(1-f)^2} - \frac{6\sqrt{3}}{\pi((2f+1)^2 + 27(2f-1)^2)} = \frac{1}{2(1-f)^2} - \frac{3\sqrt{3}}{2\pi(28f^2 - 26f + 7)} = \\ & \frac{(28\pi - 3\sqrt{3})f^2 - (26\pi - 6\sqrt{3})f + (7\pi - 3\sqrt{3})}{2\pi(1-f)^2(28f^2 - 26f + 7)} > 0, \end{aligned}$$

where the last transition holds as both the numerator and the denominator are positive for every $f \in \mathbb{R}$. Thus, given that $d = \frac{1}{2}$, the regret is maximal for $f = \frac{1}{2}$, for which the regret is exactly $\frac{1}{2}$, as desired.

To finish the proof, it remains to consider the case in which A contains two different extreme points of $\Delta(\Omega)$.¹⁹ If s intersects l at a boundary point of $\Delta(\Omega)$, then Adversary can increase the regret by rotating l around that point in a way which decreases the area of A ; this way u^* cannot decrease, while $u(\pi, u_r)$ decreases. Therefore, it is enough to consider the case in which s contains the unique extreme point of $\Delta(\Omega)$ that is not in A . Assume w.l.o.g. that $\omega_2, \omega_3 \in A$, while $\omega_1 \notin A$ is an endpoint of s . Then the second element of $\text{supp}(\pi)$, which belongs to l , is of the form $F = (f, \frac{f}{\sqrt{3}})$ for $f = \frac{3de}{2d+e}$ (see Fig. 8). We must have $\frac{1}{2} \leq f \leq \frac{3}{4}$, as the segment $\omega_1 F$ contains μ and $F \in \Delta(\Omega)$. The regret is:

$$\frac{|\mu\omega_1|}{|F\omega_1|} - \frac{\angle D\mu E}{2\pi} = \frac{2d+e}{6de} - \frac{1}{2\pi} \text{at} \left(\frac{\frac{\sqrt{3}/6}{1/2-e} - \frac{\sqrt{3}/6 - \sqrt{3}d}{1/2-d}}{1 + \frac{\sqrt{3}/6}{1/2-e} \cdot \frac{\sqrt{3}/6 - \sqrt{3}d}{1/2-d}} \right) = \frac{2d+e}{6de} - \frac{1}{2\pi} \text{at} \left(\sqrt{3} \cdot \frac{d(e-1/3) - e/6}{d(1-e) + e/2 - 1/3} \right).$$

Set $x := de$ and $y := 2d + e$. Then the above expression equals (note that $y = \frac{3x}{f}$):

$$\begin{aligned} & \frac{y}{6x} - \frac{1}{2\pi} \text{at} \left(\sqrt{3} \cdot \frac{x - y/6}{y/2 - x - 1/3} \right) = \frac{1}{2f} - \frac{1}{2\pi} \text{at} \left(\sqrt{3} \cdot \frac{x - x/(2f)}{(3x)/(2f) - x - 1/3} \right) = \\ & \frac{1}{2f} - \frac{1}{2\pi} \text{at} \left(\sqrt{3} \cdot \frac{x(2f-1)}{x(3-2f) - 2f/3} \right). \end{aligned}$$

¹⁹ If A contains no extreme points or all the extreme points of $\Delta(\Omega)$ – the regret is 0.

Note that the angle $\angle D\mu E$ is obtuse for all feasible choices of x, f ; thus, for a fixed f , the above expression is continuously differentiable w.r.t. x . The derivative of the expression inside $\text{at}(\cdot)$ w.r.t. x is:

$$\sqrt{3} \frac{(2f-1) \cdot (-2f/3)}{(x(3-2f)-2f/3)^2} < 0,$$

as $f > \frac{1}{2}$. Therefore, for a fixed f , the regret is maximal for the largest possible x . As $f = \frac{3x}{y}$, maximizing x for a fixed f is equivalent to maximizing x for a fixed y . By the AM-GM inequality, $x \leq \frac{y^2}{8}$, with equality if and only if $2d = e$.²⁰ When $2d = e$, we have $f = \frac{3x}{y} = \frac{3y}{8} = \frac{3e}{4}$ and $x = \frac{e^2}{2}$; furthermore, we must have $\frac{1}{2} \leq e \leq 1$, as $\mu \notin A$. Assuming $2d = e$, the regret as a function of e is:

$$\rho(e) := \frac{2}{3e} - \frac{1}{2\pi} \text{at}\left(\sqrt{3} \cdot \frac{e^2(3e/4 - 1/2)}{e^2(3/2 - 3e/4) - e/2}\right) = \frac{2}{3e} - \frac{1}{2\pi} \text{at}\left(\sqrt{3} \cdot \frac{e(3e/4 - 1/2)}{e(3/2 - 3e/4) - 1/2}\right).$$

We have:

$$\begin{aligned} \rho'(e) &= -\frac{2}{3e^2} - \frac{\sqrt{3}}{2\pi} \frac{1}{1 + 3\left(\frac{e(3e/4 - 1/2)}{e(3/2 - 3e/4) - 1/2}\right)^2} \\ &= \frac{(3e/2 - 1/2)(e(3/2 - 3e/4) - 1/2) - e(3e/4 - 1/2)(3/2 - 3e/2)}{(e(3/2 - 3e/4) - 1/2)^2} \\ &= -\frac{2}{3e^2} - \frac{\sqrt{3}}{2\pi} \frac{3e^2 - 3e + 1}{(3e^2/2 - 3e + 1)^2 + 3(3e^2/2 - e)^2} < 0. \end{aligned}$$

Hence, the maximum of $\rho(e)$ on $[\frac{1}{2}, 1]$ is obtained for $e = \frac{1}{2}$. Since $\rho(\frac{1}{2}) = \frac{1}{2}$, the regret is always at most $\frac{1}{2}$, which completes our proof. $\square$

Appendix B. Proof of Theorem 3.5

Proof of Theorem 3.5. We start by proving the lower bound. Assume $n \geq 16$ (otherwise the result follows from Remark 3.3). Our proof constructs a difficult instance under which any algorithm suffers a regret of at least $1 - \frac{2}{\sqrt{n}}$. Fix $0 < \delta < \frac{\min_{i \in \Omega} \mu_i}{2n}$. Set $\text{supp}_\delta(p) := \{i \in \Omega : \Pr_{\omega' \sim p}[\omega' = i] > \delta\}$. Depending on what Receiver's utility a state leads to, we divide the n states in our instance into three kinds: a single *good* state, $n - \lfloor \sqrt{n} \rfloor - 1$ *normal* states and $\lfloor \sqrt{n} \rfloor$ *bad* states. Sender's uncertainty about Receiver's utilities is captured by her uncertainty on the type of the true state ω .

Suppose that Adversary chooses a uniform permutation on Ω ; then she sets the first state in the permutation to be good, the next $n - \lfloor \sqrt{n} \rfloor - 1$ states to be normal and the rest to be bad. Denote by μ_{good} the prior probability of the good state and by μ_{normal} – the sum of the prior probabilities of the normal states. Suppose that for $i \in \Omega$, $u_r(i, 1)$ equals $\frac{1}{\mu_{\text{good}}}$ if i is good, $-\frac{1}{\mu_{\text{normal}}}$ if i is normal and $-\frac{1}{\delta \cdot \mu_{\text{good}}}$ if i is bad. The utility choices ensure that:

1. The expected Receiver's utility knowing that the state is not bad is 0.
2. A posterior p with a bad state in $\text{supp}_\delta(p)$ leads to rejection.
3. A posterior p s.t. the good state is not in $\text{supp}_\delta(p)$ leads to rejection, as then p must assign a probability of at least $\frac{1-\delta}{n-1}$ to a certain not good state, resulting in Receiver's utility being smaller than:

$$\begin{aligned} \frac{\delta}{\mu_{\text{good}}} - \frac{1-\delta}{n-1} \cdot \frac{1}{\max\{\delta \cdot \mu_{\text{good}}, \mu_{\text{normal}}\}} &= \frac{\delta}{\mu_{\text{good}}} - \frac{1-\delta}{\mu_{\text{normal}}(n-1)} < \\ \frac{1}{2n\mu_{\text{normal}}} - \frac{1}{\mu_{\text{normal}}(n-1)} + \frac{1}{2\mu_{\text{normal}}(n-1)} &< 0. \end{aligned}$$

Had Sender known the type of each state, she could have used a signaling scheme revealing whether the true state ω was bad or not. This scheme would have made Receiver adopt exactly for not bad values of ω . Thus, $u^* \geq 1 - \frac{\lfloor \sqrt{n} \rfloor}{n} \geq 1 - \frac{1}{\sqrt{n}}$.

We shall prove that our ignorant Sender cannot get a utility above $\frac{1}{\sqrt{n}}$; it would imply $\text{Reg}_{\text{ARB}} \geq \left(1 - \frac{1}{\sqrt{n}}\right) - \frac{1}{\sqrt{n}} \geq 1 - \frac{2}{\sqrt{n}}$, completing the proof. In fact, we shall show a stronger result: for every Receiver's posterior $p \in \Delta(\Omega)$, the probability of adoption, under Adversary's strategy, is at most $\frac{1}{\sqrt{n}}$.

Indeed, fix some posterior $p \in \Delta(\Omega)$. We have the following cases:

²⁰ The geometric interpretation of this case is – DE is parallel to $\omega_3\omega_2$.

- If $|\text{supp}_\delta(p)| \leq \lfloor \sqrt{n} \rfloor$, then the probability of having the good state in $\text{supp}_\delta(p)$ is:

$$\frac{|\text{supp}_\delta(p)|}{n} \leq \frac{\sqrt{n}}{n} = \frac{1}{\sqrt{n}}.$$
- Otherwise ($|\text{supp}_\delta(p)| > \lfloor \sqrt{n} \rfloor$), the probability of not having a bad state – and meanwhile, having the good state – in $\text{supp}_\delta(p)$ is:

$$\begin{aligned} & \left(1 - \frac{\lfloor \sqrt{n} \rfloor}{n}\right) \left(1 - \frac{\lfloor \sqrt{n} \rfloor}{n-1}\right) \dots \left(1 - \frac{\lfloor \sqrt{n} \rfloor}{n+1-|\text{supp}_\delta(p)|}\right) \cdot \frac{|\text{supp}_\delta(p)|}{n - \lfloor \sqrt{n} \rfloor} \leq \\ & \left(1 - \frac{\lfloor \sqrt{n} \rfloor}{n}\right)^{|\text{supp}_\delta(p)|} \cdot \frac{|\text{supp}_\delta(p)|}{n - \lfloor \sqrt{n} \rfloor} \leq \left(1 - \frac{\lfloor \sqrt{n} \rfloor}{n}\right)^{\lfloor \sqrt{n} \rfloor + 1} \cdot \frac{\lfloor \sqrt{n} \rfloor + 1}{n - \lfloor \sqrt{n} \rfloor} \leq \\ & e^{-\frac{\lfloor \sqrt{n} \rfloor (\lfloor \sqrt{n} \rfloor + 1)}{n}} \cdot \frac{\lfloor \sqrt{n} \rfloor + 1}{n - \lfloor \sqrt{n} \rfloor} \leq \frac{1}{e^{1-1/\sqrt{n}}} \cdot \frac{\sqrt{n} + 1}{n - \sqrt{n}} \leq_{n \geq 16} \frac{1}{\sqrt{n}}, \end{aligned}$$

where the second inequality holds since $\left(1 - \frac{\lfloor \sqrt{n} \rfloor}{n}\right)^x \cdot x$ is non-increasing as a function of the integer variable x for $x \geq \lfloor \sqrt{n} \rfloor + 1$, because $1 - \frac{\lfloor \sqrt{n} \rfloor}{n} \leq \frac{x}{x+1}$ for $x \geq \lfloor \sqrt{n} \rfloor + 1$.

Therefore, the expected adoption probability over Adversary's strategy at every posterior is at most $\frac{1}{\sqrt{n}}$, which completes the proof of the lower bound.

Let us prove now the upper bound. Note that the full revelation scheme *fails* to provide any regret guarantee for priors that have tiny masses on some of the states. The regret guarantee of full revelation is as high as $1 - \min_{i \in [n]} \mu_i$.²¹ We shall prove that a modification of the full-revelation scheme s.t. every signal realization pools at most two states has a regret guarantee of $1 - \frac{1}{4n^2}$ for every prior.

Define $U := \{i \in \Omega : \mu_i \geq \frac{1}{2n}\}$. Note that $\sum_{i \in U} \mu_i = 1 - \sum_{i \notin U} \mu_i \geq 1 - \frac{n}{2n} = \frac{1}{2}$ (in particular, $U \neq \emptyset$).

Consider the signaling scheme π with the set of signals $S = \{s_i : i \in \Omega\} \cup \{s_{i,j} : i \in U, j \in \Omega \setminus U\}$, s.t.:

- For every $i \in U$, the signal s_i is assigned a probability mass of $(1 - \frac{1}{2n})\mu_i$ out of the prior probability of μ_i for $\omega = i$.
- For every $i \in \Omega \setminus U$, the signal s_i is assigned a probability mass of $\frac{\mu_i}{2}$ out of the prior probability for $\omega = i$.
- For every $i \in U, j \in \Omega \setminus U$, the signal $s_{i,j}$ is assigned a probability mass of $\frac{\mu_i}{2n(n-|U|)}$ out of the prior probability for $\omega = i$ and a probability mass of $\frac{\mu_j}{2|U|}$ out of the prior probability for $\omega = j$.

We claim that $\text{Reg}_{\text{ARB}}(\pi) \leq 1 - \frac{1}{4n^2}$. Indeed, fix u_r .²² Define $T := \{i \in \Omega : u_r(i, 1) \geq 0\}$. Consider the following cases:

- $T \cap U \neq \emptyset$. Since s_i leads to adoption for every $i \in T$, in this case $u(\pi, u_r) \geq \frac{1}{2n} \cdot (1 - \frac{1}{2n}) \geq \frac{1}{4n^2}$; thus, $\text{Reg}_{\text{ARB}}(\pi) = u^*(u_r) - u(\pi, u_r) \leq 1 - \frac{1}{4n^2}$.
- $T \cap U = \emptyset$, and for every $i \in U$ there exists $j_i \in \Omega \setminus U$ s.t. $\frac{\mu_i}{2n(n-|U|)} \cdot u_r(i, 1) + \frac{\mu_{j_i}}{2|U|} \cdot u_r(j_i, 1) \geq 0$. Then s_{i,j_i} yields to adoption for every $i \in U$. We saw that $\sum_{i \in U} \mu_i \geq \frac{1}{2}$; thus, adoption occurs with probability at least $\sum_{i \in U} \frac{\mu_i}{2n(n-|U|)} = \frac{\sum_{i \in U} \mu_i}{2n(n-|U|)} \geq \frac{1}{4n^2}$, and again the regret is at most $1 - \frac{1}{4n^2}$.
- $T \cap U = \emptyset$, and there exists $i \in U$ s.t. for every $j \in \Omega \setminus U$: $\frac{\mu_i}{2n(n-|U|)} \cdot u_r(i, 1) + \frac{\mu_j}{2|U|} \cdot u_r(j, 1) < 0$. Summing over $j \in \Omega \setminus U$ and multiplying by $2|U|$ yields: $\frac{|U|}{n} \cdot \mu_i \cdot u_r(i, 1) + \sum_{j \in \Omega \setminus U} \mu_j \cdot u_r(j, 1) < 0$. If $U = \Omega$, then adoption never occurs for any signaling scheme and the regret is 0. Otherwise, as $T \subseteq \Omega \setminus U$, we get that even upon the knowledge of u_r , a probability mass greater than $\left(1 - \frac{|U|}{n}\right) \cdot \mu_i \geq \frac{\mu_i}{n} \geq \frac{1}{2n^2}$ from the prior probability of μ_i for $\omega = i$ does not lead to adoption. Therefore, $u^* \leq 1 - \frac{1}{2n^2} \leq 1 - \frac{1}{4n^2}$, and $\text{Reg}_{\text{ARB}}(\pi) = u^*(u_r) - u(\pi, u_r) \leq 1 - \frac{1}{4n^2}$.

In all the cases, $\text{Reg}_{\text{ARB}}(\pi) \leq 1 - \frac{1}{4n^2}$, as desired. $\square$

Appendix C. Proof of Theorem 3.6

We start with the following lemma.

²¹ Indeed, if for some $i \in [n]$ we have $u_r(i, 1) = \frac{1}{\mu_i}$ and $u_r(j, 1) = -\frac{1}{1-\mu_i}$ for $j \neq i$, then the no-information scheme leads to adoption with probability 1, while the full-revelation scheme leads to adoption with probability μ_i ; therefore, the regret is $1 - \mu_i$.

²² The result for a mixture over possible functions u_r would follow by taking the expectation over Adversary's strategy.

Lemma C.1. Let $\alpha \in (0, 1)$ be a constant. Let G_α be a two-player zero-sum game with continuum action sets $X = Y = [0, 1 - \alpha]$ of the x - (maximizing) and the y - (minimizing) players, respectively, and utility:

$$g(x, y) := (1 - x) - (1 - y)\mathbf{1}_{y \geq x}.$$

Denote by $v = \text{Val}(G_\alpha)$ the value of G_α .

- For $\alpha \geq \frac{1}{e}$ we have $v = -\alpha \ln \alpha$. Furthermore, there exists an optimal strategy $o_x^* \in \Delta([0, 1 - \alpha])$ of the x -player that has an atom of weight α on $x = 0$ and otherwise has the density function $f_X(x) := \frac{\alpha}{(1-x)^2}$ over the entire segment $[0, 1 - \alpha]$. Moreover, there exists an optimal strategy $o_y^* \in \Delta([0, 1 - \alpha])$ of the y -player that has an atom of weight $1 + \ln \alpha$ on $y = 1 - \alpha$ and otherwise has the density function $f_Y(y) := \frac{1}{1-y}$ over the entire segment $[0, 1 - \alpha]$.
- For $\alpha < \frac{1}{e}$ we have $v = \frac{1}{e}$. Furthermore, there exists an optimal strategy $o_x^* \in \Delta([0, 1 - \alpha])$ of the x -player that has an atom of weight $\frac{1}{e}$ on $x = 0$ and otherwise has the density function $f_X(x) := \frac{1}{e(1-x)^2}$ over the segment $[0, 1 - \frac{1}{e}]$. Moreover, there exists an optimal strategy $o_y^* \in \Delta([0, 1 - \alpha])$ of the y -player that has the density function $f_Y(y) := \frac{1}{1-y}$ over the segment $[0, 1 - \frac{1}{e}]$.

Proof of Lemma C.1. Assume first that $\alpha \geq \frac{1}{e}$. We shall show that for the strategy couple (o_x^*, o_y^*) , both players are indifferent between all actions in $[0, 1 - \alpha]$ and the payoff is $-\alpha \ln \alpha$; therefore, (o_x^*, o_y^*) is an equilibrium and $v = -\alpha \ln \alpha$. Indeed, every fixed $x \in [0, 1 - \alpha]$ yields expected (over y) payoff of:

$$\begin{aligned} \int_0^{1-\alpha} f_Y(y)g(x, y)dy + \Pr[y = 1 - \alpha]g(x, 1 - \alpha) &= \int_0^{1-\alpha} \frac{1}{1-y} \cdot (1-x)dy - \\ \int_x^{1-\alpha} \frac{1}{1-y} \cdot (1-y)dy + (1 + \ln \alpha) \cdot (1-x-\alpha) &= -(1-x)\ln \alpha - (1-\alpha-x) + \\ (1 + \ln \alpha)(1-x-\alpha) &= -\alpha \ln \alpha. \end{aligned}$$

Furthermore, every fixed $y \in [0, 1 - \alpha]$ yields expected (over x) payoff of:

$$\begin{aligned} \int_0^{1-\alpha} f_X(x)g(x, y)dx + \Pr[x = 0]g(0, y) &= \int_0^{1-\alpha} \frac{\alpha}{(1-x)^2} \cdot (1-x)dx - \\ \int_0^y \frac{\alpha}{(1-x)^2} \cdot (1-y)dx + \alpha \cdot y &= -\alpha \ln \alpha - (1-y) \cdot \left(\frac{\alpha}{1-y} - \alpha \right) + \alpha \cdot y = -\alpha \ln \alpha, \end{aligned}$$

as desired.

Assume now that $\alpha < \frac{1}{e}$. We shall show that for the strategy couple (o_x^*, o_y^*) , both players are indifferent between the actions in $[0, 1 - \frac{1}{e}]$ – resulting in payoff of $\frac{1}{e}$ – and worse off by taking other actions. Indeed, every fixed $x \in [0, 1 - \frac{1}{e}]$ yields expected payoff of:

$$\begin{aligned} \int_0^{1-1/e} f_Y(y)g(x, y)dy &= \int_0^{1-1/e} \frac{1}{1-y} \cdot (1-x)dy - \int_x^{1-1/e} \frac{1}{1-y} \cdot (1-y)dy = \\ (1-x) \cdot 1 - \left(1 - \frac{1}{e} - x \right) &= \frac{1}{e}, \end{aligned}$$

while for a fixed $x \in (1 - \frac{1}{e}, 1 - \alpha]$, the expected payoff is:

$$\begin{aligned} \int_{1-1/e}^{1-\alpha} f_Y(y)g(x, y)dy &= \int_{1-1/e}^{1-\alpha} \frac{1}{1-y} \cdot (1-x)dy = -(1 + \ln \alpha)(1-x) < \\ -(1 + \ln \alpha) \left(1 - \left(1 - \frac{1}{e} \right) \right) &< \frac{1}{e}, \end{aligned}$$

as desired. Moreover, every fixed $y \in [0, 1 - \frac{1}{e}]$ yields expected payoff of:

$$\int_0^{1-1/e} f_X(x)g(x, y)dx + \Pr[x=0]g(0, y) = \int_0^{1-1/e} \frac{1}{e(1-x)^2} \cdot (1-x)dx - \int_0^y \frac{1}{e(1-x)^2} \cdot (1-y)dx + \frac{1}{e} \cdot y = \frac{1}{e} \cdot 1 - \frac{1-y}{e} \cdot \left(\frac{1}{1-y} - 1 \right) + \frac{1}{e} \cdot y = \frac{1}{e},$$

while for a fixed $y \in (1 - \frac{1}{e}, 1 - \alpha]$ we have:

$$\int_0^{1-1/e} f_X(x)g(x, y)dx + \Pr[x=0]g(0, y) = \int_0^{1-1/e} \frac{1}{e(1-x)^2} \cdot ((1-x) - (1-y))dx + \frac{1}{e} \cdot y = \frac{1}{e} \cdot (1 - (1-y) \cdot (e-1)) + \frac{1}{e} \cdot y = \frac{2}{e} - 1 + y > \frac{1}{e},$$

as desired. $\square$

Proof of Theorem 3.6. Consider the two-player zero-sum game G_{μ_n} as described in Lemma C.1 and let v be its value. We shall prove that G (the zero-sum game interpretation of regret minimization; see Subsection 2.2) also has value equal to v .

First, let us show that Sender can ensure regret of at most v by using the signaling scheme π_y defined as follows: Sender picks a random $y \sim o_y^*$ (where o_y^* is defined in Lemma C.1); then she uses the y -threshold scheme (see Definition 2.3).

Indeed, fix u_r . By Fact 2.4, it defines an optimal threshold $x = x(u_r)$ (in the standard persuasion model with the knowledge of Receiver's utility). Note that not all the values $x \in [0, 1]$ might be the optimal threshold for some u_r , but only $x \in [0, 1 - \mu_n]$. Indeed, the highest-utility state n should be included entirely in the knapsack, as otherwise adoption never occurs and the regret is 0. Assume, therefore, that $x \in [0, 1 - \mu_n]$.

We have $u^* = 1 - x$, and for every $y \in [0, 1 - \mu_n]$: $u(\pi_y, u_r) = \mathbf{1}_{y \geq x}$, as Receiver adopts if and only if $y \geq x$. Therefore, the regret for π_y and u_r is (where g is as defined in Lemma C.1):

$$u^*(u_r) - u(\pi_y, u_r) = \mathbb{E}_{y \sim o_y^*} [(1-x) - (1-y)\mathbf{1}_{y \geq x}] = \mathbb{E}_{y \sim o_y^*} [g(x, y)].$$

By Lemma C.1, the last expression is at most v . Therefore, Sender can ensure that the regret is at most v for any fixed u_r ; thus, it holds also for any mixture over u_r , as needed.

It remains to prove that Adversary can ensure a regret of at least v . Consider the following strategy of Adversary: she deterministically sets $u_r(i, 1) := -\mu_n$ for all $1 \leq i \leq n-1$; then she chooses a random $t \sim o_x^*$ (with o_x^* from Lemma C.1) and sets $u_r(n, 1) := 1 - \mu_n - t$. For this Adversary's strategy, the optimal threshold x is distributed according to o_x^* . To bound the regret that Sender can guarantee against this Adversary's strategy, fix a signaling scheme π . As $u_r(i, 1)$ is the same for $1 \leq i \leq n-1$, one can refer to the states $1, 2, \dots, n-1$ as a single state that we call state 0. Concretely, π induces a signaling scheme π' in a persuasion scenario with a binary-state space $\{0, n\}$ s.t. $\mu_0 := 1 - \mu_n$; a posterior $p(s) \in \Delta(\{n\})$ is mapped to the posterior $(\sum_{1 \leq i \leq n-1} p_i, p_n)$. Therefore, computing $\text{Reg}_{\text{MON}}(\pi)$ in our original setting is reduced to computing $\text{Reg}_{\text{MON}}(\pi')$ in this binary-state setting, with Adversary's strategy being setting $u_r(0, 1) := -\mu_n$ deterministically and choosing $u_r(n, 1) = 1 - \mu_n - t$, where $t \sim o_x^*$.

From now on, we shall refer to a posterior $p \in \Delta(\{0, n\})$ as a real number $q \in [0, 1]$, where $q := p_0$. To understand the optimal Sender's utility in this binary-state persuasion problem, we compute her expected utility $u'(q)$ (when the expectation is over Adversary's mixed strategy) for each possible posterior $q \in [0, 1]$; then we evaluate the concavification of u' at the prior $1 - \mu_n$.

We shall prove that there exists Sender's best-reply signaling scheme that is a threshold signaling scheme. It would complete the proof, as the optimal threshold x is distributed according to o_x^* , which gives for a Sender's y -threshold scheme expected regret of:

$$\mathbb{E}_{x \sim o_x^*} [(1-x) - (1-y)\mathbf{1}_{y \geq x}] = \mathbb{E}_{x \sim o_x^*} [g(x, y)],$$

(with g defined in Lemma C.1) which is, by Lemma C.1, at least v for every $y \in [0, 1 - \mu_n]$.²³

Indeed, to understand Sender's best-reply we shall consider a standard Bayesian persuasion instance in which Sender's utility is the expected (over Adversary's strategy) adoption probability. Sender's expected utility, as a function of Receiver's posterior q , is:

$$u'(q) := \Pr_{t \sim o_x^*} [-\mu_n q + (1 - \mu_n - t)(1 - q) \geq 0] = \Pr_{t \sim o_x^*} \left[t \leq 1 - \frac{\mu_n}{1 - q} \right].$$

²³ We can assume w.l.o.g. that $y \in [0, 1 - \mu_n]$, as the optimal threshold x is at most $1 - \mu_n$; thus, choosing $y > 1 - \mu_n$ gives a greater regret than choosing $y = 1 - \mu_n$.




Fig. 9. The function $u'(q)$ and its concavification. The function appears in blue; its concavification appears in red.

Straightforward calculations show that for $\mu_n \geq \frac{1}{e}$:

$$u'(q) = \begin{cases} 1 - q & \text{if } 0 \leq q \leq \mu_0 \\ 0 & \text{if } \mu_0 < q \leq 1, \end{cases}$$

while for $\mu_n < \frac{1}{e}$:

$$u'(q) = \begin{cases} 1 & \text{if } 0 \leq q < 1 - \mu_n e \\ \frac{1-q}{\mu_n e} & \text{if } 1 - \mu_n e \leq q \leq \mu_0 \\ 0 & \text{if } \mu_0 < q \leq 1. \end{cases}$$

In both cases, $u'(q) = 0$ for $q > \mu_0$ and $u'(q) > 0$ otherwise. Therefore, the graph of the concavification of u' (see Fig. 9) includes a line segment connecting the point $(q = 1, u'(q) = 0)$ with a point having a value of q smaller than μ_0 . Thus, there exists an optimal Sender's signaling scheme with binary signals s.t. one of them leads to the posterior $q = 1$ (i.e., certainty that $\omega = 0$); this scheme is a threshold scheme (see Definition 2.3), as desired. $\square$

Appendix D. Proof of Proposition 3.8

Proof of Proposition 3.8. To describe the proof, we use the *fractional multidimensional knapsack approach*. Consider the k -dimensional cube $P := \times_{j \in [k]} [1, n_j]$, with integer points representing the elements of Ω . A *knapsack* K is a closed subset of P s.t. for every $p \in K$, all the points that *Pareto-dominate* p are also in K .²⁴ We define a *knapsack signaling scheme* – the multidimensional variant of a threshold scheme (see Definition 2.3) – as follows. Interpret Ω , equipped with the prior μ , as drawing uniformly $p \in P$ – called the *continuous state* – s.t. all the realizations in $\times_{j \in [k]} \left(\sum_{l_j < m_j} \Pr_{\omega'_j \sim \mu_j} [\omega'_j = l_j], \sum_{l_j \leq m_j} \Pr_{\omega'_j \sim \mu_j} [\omega'_j = l_j] \right)$ correspond to the state $(m_1, \dots, m_k)$ (when $1 \leq m_j \leq n_j$ for every $1 \leq j \leq k$). For a knapsack K , the *K-knapsack signaling scheme* is a binary-signal scheme revealing whether $p \in K$ or not.²⁵

Consider the signaling scheme π specified by the knapsack:

$$K_\pi := \left\{ (p^1, \dots, p^k) \in P : p^j \geq \frac{n_j+1}{2} \forall 1 \leq j \leq k \right\}. \text{ We claim that } \pi \text{ ensures a regret of at most } 1 - 2^{-k}.$$

Indeed, if Adversary chooses A s.t. every continuous state realization in K_π leads to adoption – Sender earns exactly 2^{-k} . Since always $u^*(u_r) \leq 1$, the regret is at most $1 - 2^{-k}$, as desired.

Otherwise, there exists some continuous state realization $p = (p^1, \dots, p^k) \in K$ leading to rejection. As $(p^1, \dots, p^k) \in K$, we have $p^j \geq \frac{n_j+1}{2}$ for every $1 \leq j \leq k$. Thus, no realization with the j -th coordinate being at most $\frac{n_j+1}{2}$ for every j leads to adoption. Similarly to the single-dimensional setting, for a given u_r , there exists an optimal knapsack signaling scheme. Had Sender known A , her optimal knapsack strategy would have been specified by a knapsack K^* disjoint to $S := \left\{ (q^1, \dots, q^k) \in P : q^j \leq \frac{n_j+1}{2} \forall 1 \leq j \leq k \right\}$. Such a strategy yields adoption probability of at most $1 - \frac{\text{Vol}(S)}{\text{Vol}(P)} = 1 - 2^{-k}$. Therefore, the regret is at most $u^* \leq 1 - 2^{-k}$, as desired. $\square$

Appendix E. Proof of Theorem 4.2

We start with the following lemma.

²⁴ A point $q = (q^1, \dots, q^k)$ *Pareto-dominates* $p = (p^1, \dots, p^k)$ if $q^j \geq p^j$ for every $1 \leq j \leq k$.

²⁵ The knapsack scheme specified by a knapsack with zero volume is the no-information signaling scheme.

Lemma E.1. Let $\alpha \in (0, 1)$ be a constant. Denote $\beta = \beta(\alpha) := \frac{1}{1+\ln \frac{1}{\alpha}}$. Let G'_α be a two-player zero-sum game with continuum action sets $X = Y = [0, 1 - \alpha]$ of the x - (minimizing) and the y - (maximizing) players, respectively, and utility:

$$h(x, y) := \frac{(1-y)\mathbf{1}_{y \geq x}}{1-x}.$$

Then the value of G'_α is β . Furthermore, there exists an optimal strategy $\sigma_x^{*'} \in \Delta([0, 1 - \alpha])$ of the x -player that has an atom of weight β on $x = 0$ and otherwise has the density function $f_X(x) := \frac{\beta}{1-x}$ over the entire segment $[0, 1 - \alpha]$. Moreover, there exists an optimal strategy $\sigma_y^{*'} \in \Delta([0, 1 - \alpha])$ of the y -player that has an atom of weight β on $y = 1 - \alpha$ and otherwise has the density function $f_Y(y) := \frac{\beta}{1-y}$ over the entire segment $[0, 1 - \alpha]$.

Proof of Lemma E.1. We shall show that for the strategy couple $(\sigma_x^{*'}, \sigma_y^{*'})$, both players are indifferent between all actions in $[0, 1 - \alpha]$ and the payoff is β ; therefore, $(\sigma_x^{*'}, \sigma_y^{*'})$ is an equilibrium and the game value is β . Indeed, every fixed $x \in [0, 1 - \alpha]$ yields expected (over y) payoff of:

$$\begin{aligned} \int_0^{1-\alpha} f_Y(y)h(x, y)dy + \Pr[y = 1 - \alpha]h(x, 1 - \alpha) &= \int_x^{1-\alpha} \frac{\beta}{1-y} \cdot \frac{1-y}{1-x} dy + \beta \cdot \frac{1-(1-\alpha)}{1-x} = \\ \frac{\beta(1-\alpha-x)}{1-x} + \frac{\alpha\beta}{1-x} &= \beta, \end{aligned}$$

as needed. Furthermore, every fixed $y \in [0, 1 - \alpha]$ yields expected (over x) payoff of:

$$\begin{aligned} \int_0^{1-\alpha} f_X(x)h(x, y)dx + \Pr[x = 0]h(0, y) &= \int_0^y \frac{\beta}{1-x} \cdot \frac{1-y}{1-x} dx + \beta \cdot \frac{1-y}{1-0} = \\ \beta(1-y) \left(\frac{1}{1-y} - 1 \right) + \beta(1-y) &= \beta, \end{aligned}$$

as desired. $\square$

Proof of Theorem 4.2. Using the notations of Lemma E.1, consider the two-player zero-sum game G'_{μ_n} ; by the lemma, it has value of $\beta := \beta(\mu_n) = \frac{1}{1+\ln \frac{1}{\mu_n}}$. We claim that the two-player zero-sum game interpretation of adversarial approximation ratio maximization, in which Sender is the maximizing player and her possible strategies are the signaling schemes, and Adversary is the minimizing player and her possible strategies are the mixtures over Receiver's utility functions, has value of β .

Using similar arguments to Theorem 3.6 proof, Sender can ensure adversarial approximation of at least β by using the signaling scheme π_y defined as follows: Sender picks a random $y \sim \sigma_y^{*'}$; then she uses the y -threshold scheme. The only difference from Theorem 3.6 proof is that in the current proof, for a fixed u_r and the corresponding optimal threshold x , the expected adversarial approximation over π_y is $\mathbb{E}_{y \sim \sigma_y^{*'}}[h(x, y)]$ (for h from Lemma E.1).

It remains to prove that Adversary can ensure an adversarial approximation of at most β . As in Theorem 3.6 proof, it is enough to prove that Adversary is guaranteed an adversarial approximation of at most β in a persuasion scenario with a binary-state space $\{0, n\}$, s.t. $\mu_0 := 1 - \mu_n$, by using the following strategy: $u_r(0, 1) := -\mu_n$ deterministically; $u_r(n, 1) := 1 - \mu_n - t$, where $t \sim \sigma_x^*$.

We shall refer to a posterior $p \in \Delta(\{0, n\})$ as a real number $q \in [0, 1]$, where $q := p_0$. As in Theorem 3.6 proof, we compute Sender's expected utility $u'(q)$ (the expectation is over Adversary's mixed strategy) for each possible posterior $q \in [0, 1]$, and then we evaluate the concavification of u' at the prior $1 - \mu_n$. Again, it is enough to prove that there exists Sender's best-reply signaling scheme that is a threshold scheme.

Indeed, to understand Sender's best-reply we can consider a standard Bayesian persuasion instance in which Sender's utility is the expected – over Adversary's strategy – ratio of the indicator whether adoption occurs at the posterior q to the optimal adoption probability upon knowing t . Sender's expected utility, as a function of the posterior q , is:

$$u'(q) := \mathbb{E}_{t \sim \sigma_x^*} \left[\frac{\mathbb{1}_{-\mu_n q + (1-\mu_n-t)(1-q) \geq 0}}{1-t} \right] = \mathbb{E}_{t \sim \sigma_x^*} \left[\frac{\mathbb{1}_{t \leq 1 - \frac{\mu_n}{1-q}}}{1-t} \right].$$

Straightforward calculations show that (see Fig. 10):

$$u'(q) = \begin{cases} \frac{\beta(1-q)}{\mu_n} & \text{if } 0 \leq q \leq \mu_0 \\ 0 & \text{if } \mu_0 < q \leq 1. \end{cases}$$




Fig. 10. The function u' (appears in blue) and its concavification (appears in red).

Therefore, there exists an optimal signaling scheme at the prior that uses the posterior $q = 1$, which is, in particular, a threshold scheme, as needed. $\square$

References

- Alonso, Ricardo, Câmara, Odilon, 2016. Persuading voters. *Am. Econ. Rev.* 106 (11), 3590–3605.
- Arieli, Itai, Babichenko, Yakov, Smorodinsky, Rann, 2018. Robust forecast aggregation. *Proc. Natl. Acad. Sci.* 115 (52), E12135–E12143.
- Aumann, Robert J., Maschler, Michael, Stearns, Richard E., 1995. *Repeated Games with Incomplete Information*. MIT Press.
- Bandi, Chaithanya, Bertsimas, Dimitris, 2014. Optimal design for multi-item auctions: a robust optimization approach. *Math. Oper. Res.* 39 (4), 1012–1038.
- Bei, Xiaohui, Gravin, Nick, Lu, Pinyan, Tang, Zhihao Gavin, 2019. Correlation-robust analysis of single item auction. In: *Proceedings of the 30th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 193–208.
- Bergemann, Dirk, Morris, Stephen, 2019. Information design: a unified perspective. *J. Econ. Lit.* 57, 44–95.
- Bhaskar, Umang, Cheng, Yu, Ko, Young Kun, Swamy, Chaitanya, 2016. Hardness results for signaling in Bayesian zero-sum and network routing games. In: *Proceedings of the 2016 ACM Conference on Economics and Computation EC*. ACM, pp. 479–496.
- Blackwell, David, 1953. Equivalent comparisons of experiments. *Ann. Math. Stat.*, 265–272.
- Bro Miltersen, Peter, Sheffet, Or, 2012. Send mixed signals: earn more, work less. In: *Proceedings of the 13th ACM Conference on Electronic Commerce*. EC '12, New York, NY, USA. ACM. ISBN 978-1-4503-1415-2, pp. 234–247.
- Candogan, Ozan, 2020. Information design in operations. In: *TutORials in Operations Research*, pp. 176–201.
- Carroll, Gabriel, 2015. Robustness and linear contracts. *Am. Econ. Rev.* 105 (2), 536–563.
- Carroll, Gabriel, 2017. Robustness and separation in multidimensional screening. *Econometrica* 85 (2), 453–488.
- Carroll, Gabriel, 2019. Robustness in mechanism design and contracting. *Annu. Rev. Econ.* 11, 139–166.
- Castiglioni, Matteo, Celli, Andrea, Marchesi, Alberto, Gatti, Nicola, 2020. Online Bayesian persuasion. *Adv. Neural Inf. Process. Syst.* 33.
- Cheng, Yu, Cheung, Ho Yee, Dughmi, Shaddin, Emamjomeh-Zadeh, Ehsan, Han, Li, Teng, Shang-Hua, 2015. Mixture selection, mechanism design, and signaling. In: *IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS 2015)*.
- Cole, Richard, Roughgarden, Tim, 2014. The sample complexity of revenue maximization. In: *Proceedings of the 46th Annual ACM Symposium on Theory of Computing STOC*, pp. 243–252.
- Dhangwatnotai, Peerapong, Roughgarden, Tim, Yan, Qiqi, 2015. Revenue maximization with a single sample. *Games Econ. Behav.* 91, 318–333.
- Dughmi, Shaddin, 2017. Algorithmic information structure design: a survey. *ACM SIGecom Exch.* 15 (2), 2–24.
- Dütting, Paul, Roughgarden, Tim, Talgam-Cohen, Inbal, 2019. Simple versus optimal contracts. In: *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*, pp. 369–387.
- Dworczak, Piotr, Martini, Giorgio, 2019. The simple economics of optimal persuasion. *J. Polit. Econ.* 127 (5), 1993–2048.
- Dworczak, Piotr, Pavan, Alessandro, 2020. Preparing for the worst but hoping for the best: Robust (Bayesian) persuasion.
- Emek, Yuval, Feldman, Michal, Gamzu, Iftah, Paes Leme, Renato, Tennenholtz, Moshe, 2012. Signaling schemes for revenue maximization. In: *Proceedings of the 13th ACM Conference on Electronic Commerce*. EC '12. ACM. ISBN 978-1-4503-1415-2, pp. 514–531.
- Gravin, Nick, Lu, Pinyan, 2018. Separation in correlation-robust monopolist problem with budget. In: *Proceedings of the 29th ACM-SIAM Symposium on Discrete Algorithms SODA*, pp. 2069–2080.
- Guo, Yingni, Shmaya, Eran, 2019a. Robust monopoly regulation. Working paper.
- Guo, Yingni, Shmaya, Eran, 2019b. The interval structure of optimal disclosure. *Econometrica* 87 (2), 653–675.
- Hayek, Friedrich A., 1945. The use of knowledge in society. *Am. Econ. Rev.* 35 (4), 519–530.
- Hu, Ju, Weng, Xi, 2021. Robust persuasion of a privately informed receiver. *Econ. Theory* 72 (3), 909–953.
- Hurwicz, Leonid, Shapiro, Leonard, 1978. Incentive structures maximizing residual gain under incomplete information. *Bell J. Econ.* 9 (1), 180–191.
- Inostroza, Nicolas, Pavan, Alessandro, 2018. Persuasion in global games with application to stress testing.
- Kamenica, Emir, 2019. Bayesian persuasion and information design. *Annu. Rev. Econ.* 11, 249–272.
- Kamenica, Emir, Gentzkow, Matthew, 2011. Bayesian persuasion. *Am. Econ. Rev.* 101 (6), 2590–2615.
- Kolotilin, Anton, 2015. Experimental design to persuade. *Games Econ. Behav.* 90, 215–226.
- Kolotilin, Anton, Mylovanov, Tymofiy, Zapechelnyuk, Andriy, Li, Ming, 2017. Persuasion of a privately informed receiver. *Econometrica* 85 (6), 1949–1964.
- Kosterina, Svetlana, 2022. Persuasion with unknown beliefs. *Theor. Econ.* 17 (3), 1075–1107.
- Li, Fei, Song, Yangbo, Zhao, Mofei, 2020. Global manipulation by local obfuscation. Working paper.
- Lingenbrink, David, Iyer, Krishnamurthy, 2019. Optimal signaling mechanisms in unobservable queues. *Oper. Res.*
- Mansour, Yishay, Slivkins, Aleksanders, Syrgkanis, Vasilis, Wu, Zhiwei Steven, 2016. Bayesian exploration: incentivizing exploration in Bayesian games. In: *Proceedings of the 2016 ACM Conference on Economics and Computation*.
- Mathevet, Laurent, Perego, Jacopo, Taneva, Ina, 2020. On information design in games. *J. Polit. Econ.* 128 (4), 1370–1404.
- Milgrom, Paul, 2000. Putting auction theory to work: the simultaneous ascending auction. *J. Polit. Econ.* 108 (2), 245–272.
- Morris, Stephen, Oyama, Daisuke, Takahashi, Satoru, 2020. Information design in binary action supermodular games. Working paper.
- Perez-Richet, Eduardo, Skreta, Vasiliki, 2022. Test design under falsification. *Econometrica* 90 (3), 1109–1142.
- Rabinovich, Zinovi, Jiang, Albert Xin, Jain, Manish, Xu, Haifeng, 2015. Information disclosure as a means to security. In: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

- Renault, Jérôme, Solan, Eilon, Vieille, Nicolas, 2017. Optimal dynamic information provision. *Games Econ. Behav.* 104, 329–349.
- Roughgarden, Tim, Talgam-Cohen, Inbal, 2019. Approximately optimal mechanism design. *Annu. Rev. Econ.* 11, 355–381.
- Savage, Leonard Jimmie, 1951. The theory of statistical decision. *J. Am. Stat. Assoc.* 46 (253), 55–67.
- Savage, Leonard Jimmie, 1972. *The Foundations of Statistics*. Dover Publications.
- Scarf, Herbert E., 1958. A min-max solution of an inventory problem. In: Arrow, Kenneth J. (Ed.), *Studies in the Mathematical Theory of Inventory and Production*, pp. 201–220.
- Shalev-Shwartz, Shai, Ben-David, Shai, 2014. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.
- Sion, Maurice, 1958. On general minimax theorems. *Pac. J. Math.* 8 (1), 171–176.
- Talgam-Cohen, Inbal, 2020. Prior-independent auctions. In: Roughgarden, Tim (Ed.), *Beyond the Worst-Case Analysis of Algorithms*. Cambridge University Press, chapter 27.
- Vazirani, Vijay V., 2003. *Approximation Algorithms*. Springer.
- Wald, Abraham, 1950. *Statistical Decision Functions*. John Wiley.
- Wikipedia contributors, 2020. Regret (decision theory) — Wikipedia, the free encyclopedia. [https://en.wikipedia.org/wiki/Regret_\(decision_theory\)](https://en.wikipedia.org/wiki/Regret_(decision_theory)). (Accessed 9 February 2021) [Online].
- Wilson, Robert B., 1987. Game-theoretic approaches to trading processes. In: Bewley, Truman F. (Ed.), *Advances in Economic Theory: Fifth World Congress*. Cambridge University Press, pp. 33–77.
- Xu, Haifeng, Rabinovich, Zinovi, Dughmi, Shaddin, Tambe, Milind, 2015. Exploring information asymmetry in two-stage security games. In: *AAAI Conference on Artificial Intelligence (AAAI)*.
- Ziegler, Gabriel, 2020. Adversarial bilateral information design. Technical report. Working paper.