

Irrationality as a mean of regularization in Bayesian Persuasion

Romain Duboscq*

Institut de Mathématiques de Toulouse ; UMR5219
Université de Toulouse ; CNRS
INSA, F-31077 Toulouse, France

Frederic de Gournay†

Institut de Mathématiques de Toulouse ; UMR5219
Université de Toulouse ; CNRS
INSA, F-31077 Toulouse, France

October 3, 2025

Abstract

We study a regularized variant of the Bayesian Persuasion problem, where the receiver’s decision process includes a divergence-based penalty that accounts for deviations from perfect rationality. This modification smooths the underlying optimization landscape and mitigates key theoretical issues, such as measurability and ill-posedness, commonly encountered in the classical formulation. It also enables the use of scalable second-order optimization methods to compute numerically the optimal signaling scheme in a setting known to be NP-hard. We present theoretical results comparing the regularized and original models, including convergence guarantees and structural properties of optimal signaling schemes. Analytical examples and numerical simulations illustrate how this framework accommodates complex environments while remaining tractable and robust. A companion Python library, *BASIL*¹, makes use of all the practical insights from this article.

*romain.duboscq@math.univ-toulouse.fr

†degourna@insa-toulouse.fr

¹<https://plmlab.math.cnrs.fr/degourna/bayesian-persuasion-by-regularization>

Contents

1	Introduction	2
2	Setting of the problem	3
3	Main results	10
4	Examples	14
5	Conclusion	23
6	Proofs	24

1 Introduction

The Bayesian Persuasion framework, introduced by Kamenica and Gentzkow [19], provides a foundational model for understanding how informed agents can strategically design information structures to influence the actions of less informed receivers. By framing persuasion as the choice of a signal structure that shapes the receiver’s posterior beliefs, the model has offered deep insights into information transmission in economics, political science, and beyond. We refer the reader to [18, 5] and the references therein for the different models and applications.

Despite its elegance, the classical Bayesian Persuasion model assumes fully rational agents and frictionless belief updating, often limiting its applicability in environments where human or institutional behavior deviates from ideal Bayesian reasoning. In many practical settings, receivers exhibit bounded rationality, behavioral biases, or computational limitations that affect how they process and respond to information. This issue has been investigated in different contexts where the agents have different priors [1, 16], make a non-Bayesian updating [14, 15, 3, 9, 25], are rationally inattentive [7, 24, 21] or have a prior bias [20].

In this work, we extend the canonical Bayesian Persuasion model by introducing a regularization term (specifically, a divergence) in the receiver’s optimization problem. This regularization induces a smoother, more tractable optimization landscape, providing both analytical clarity and computational robustness. To be more specific, this modification serves two main purposes: it models potential irrationality or sub-optimal behavior on the part of the receiver, and it addresses both theoretical and computational challenges inherent in the standard formulation. On the theoretical side, as we will show, the classical model may lead to measurability issues and difficulties in identifying the set of admissible actions. On the computational side, the persuasion problem is known to be NP-hard, and prior approaches have attempted to overcome this complexity through specialized linear programming techniques (see, e.g.,

[12, 10, 13, 4, 11, 17]). In contrast, our formulation enables the use of quasi-Newton methods, offering a more efficient and flexible computational framework.

In Section 2, we develop the mathematical framework for both the standard Bayesian Persuasion model and its regularized extension. Section 3 presents our main theoretical findings. For the classical model, we establish a lower bound on the number of messages required in an optimal signaling scheme. While this issue is typically addressed via the revelation principle, that approach does not directly apply when receivers possess private types. Regarding the regularized formulation, we demonstrate that it generally constitutes a sub-optimal approximation of the original problem. Nevertheless, under certain assumptions, we prove that it serves as a valid approximation in the sense that any sequence of minimizers converges, up to a subsequence, to a minimizer of the non-regularized problem. We also provide numerical insights and discuss the use of second-order optimization methods, which yield efficient and fast algorithms. Finally, in Section 4, we illustrate the framework and support our theoretical arguments through both analytical examples, which admit closed-form solutions, and numerical simulations, which highlight the model's flexibility and applicability to more complex environments.

2 Setting of the problem

2.1 Notations and preliminaries

For any $n \in \mathbb{N}$, we denote Δ_n the simplex of dimension n , that is

$$\Delta_n = \left\{ x \in (\mathbb{R}^+)^n \text{ such that } \sum_{k=1}^n x_k = 1 \right\}.$$

For any measurable space $(\mathcal{E}, \mathbf{E})$, where $\mathbf{E}$ is a σ -algebra of $\mathcal{E}$, we denote $\mathbf{P}(\mathcal{E})$ the set of probability measures on $(\mathcal{E}, \mathbf{E})$. When $\mathcal{E}$ is finite, we will always endow it with the σ -algebra $\mathbf{E} = \mathbf{P}(\mathcal{E})$, the set of all parts of $\mathcal{E}$. Moreover, $\mathbf{P}(\mathcal{E})$ is a compact space which identifies as $\Delta_{|\mathcal{E}|}$ where $|\mathcal{E}|$ denotes the cardinal number of $\mathcal{E}$. When $(\mathcal{E}, d)$ is a metric space, the space $(\mathcal{P}(\mathcal{E}), d_{\text{Pr}})$ is a metric space [6] where d_{Pr} is the Prokhorov metric given by

$$d_{\text{Pr}}(\mu, \nu) = \inf\{r > 0 : \mu(E) \leq \nu(E_r) + r \text{ and } \nu \leq \mu(E_r) + r, \forall E \in \mathbf{B}(\mathcal{E})\},$$

where $\emptyset_r = \emptyset$ as well as, for any $E \neq \emptyset$,

$$E_r = \{e \in \mathcal{E} : d(e, E) < r\}.$$

and $\mathbf{B}(\mathcal{E})$ is the Borel σ -algebra of $\mathcal{E}$. Furthermore, if $(\mathcal{E}, d)$ is compact, then the space $(\mathcal{P}(\mathcal{E}), d_{\text{Pr}})$ is also compact.

2.2 Agents and variables

The agents are a sender (referred to as "she") and one or more receivers (referred to as "them", oblivious to the effective number of receivers). We first describe

the variables at play, for the sake of simplicity, we suppose that each considered set is finite.

- The set of receivers is denoted $\mathcal{L}$ and is of cardinality $|\mathcal{L}| = L \geq 1$. The generic notation for a receiver is $\ell \in \mathcal{L}$ and it is common to identify $\mathcal{L}$ with $[1, \dots, L]$.
- The set of states is denoted by $\mathcal{S}$. The states are denoted s and there are $|\mathcal{S}| = S$ states in total. The states model a source of uncertainty that is common to all agents.
- Each receiver ℓ has a type $t_\ell \in \mathcal{T}_\ell$. There are $|\mathcal{T}_\ell| = T_\ell \geq 1$ types available to the receiver ℓ . The type of the receiver changes its utility and prior on $\mathcal{S}$. For each $t_\ell \in \mathcal{T}_\ell$, the probability of the ℓ -th receiver to be of type t_ℓ is given by $\eta_\ell(t_\ell)$ with $\eta_\ell \in \mathcal{P}(\mathcal{T}_\ell)$. The type of a receiver is unknown to every other agents, however every agent is aware of η_ℓ , the probability of being of a certain type. The set of types is denoted by $\mathcal{T} = \bigotimes_{\ell \in \mathcal{L}} \mathcal{T}_\ell$ and is of cardinal $T = \prod_{\ell \in \mathcal{L}} T_\ell$.
- The set of messages is denoted by $\mathcal{M}$. These will play an important role since they are at the heart of the way the information from the sender is transmitted to the receiver (see below the communication policy, or signal). The generic notation for a message is $m \in \mathcal{M}$, there are $|\mathcal{M}| = M$ messages in total.
- Each receiver ℓ has to pick an action a_ℓ in his available set of actions $\mathcal{A}_\ell$ of cardinal $|\mathcal{A}_\ell| = A_\ell$. The chosen action is known to the other agents, hence the set of available actions do not depend on the type (or else the receiver would disclose his type). The set of actions is denoted $\mathcal{A} = \bigotimes_{\ell \in \mathcal{L}} \mathcal{A}_\ell$ and is of cardinal $A = \prod_{\ell \in \mathcal{L}} A_\ell$.
- Finally, each receiver ℓ or type t_ℓ has a utility $(s, a_\ell) \rightarrow u_{t_\ell}(s, a_\ell)$ that depends on the state $s \in \mathcal{S}$ and the chosen action $a_\ell \in \mathcal{A}_\ell$. This utility drives the choice of action picked by the receiver. The sender has a utility $(s, a) \rightarrow v(s, a)$ that depends on the state $s \in \mathcal{S}$ and on each chosen action $a = (a_\ell)_{\ell \in \mathcal{L}} \in \mathcal{A}$.

2.3 The receivers' problem : choosing the action

In this section, we focus on describing the process by which the receiver ℓ of type t_ℓ chooses his action. For any prior $\nu \in \mathcal{P}(\mathcal{S})$ on the states, the receiver of type t_ℓ maximises its utility and computes

$$\Theta_{t_\ell}(\nu) = \operatorname{argmax}_{\theta \in \mathcal{P}(\mathcal{A}_\ell)} \sum_{(s, a_\ell) \in \mathcal{S} \times \mathcal{A}_\ell} u_{t_\ell}(s, a_\ell) \theta(a_\ell) \nu(s). \quad (1)$$

The elements in $\Theta_{t_\ell}(\nu)$ are the "acceptable strategies for the receiver of type t_ℓ under ν " and are *a priori* not unique. It is however well known that $\Theta_{t_\ell}(\nu)$ is

the convex hull of dirac masses located on $\mathcal{A}_{t_\ell}^*(\nu)$ the "admissible actions" (or "pure strategies") defined as

$$\mathcal{A}_{t_\ell}^*(\nu) = \operatorname{argmax}_{a \in \mathcal{A}_\ell} \sum_{s \in \mathcal{S}} u_{t_\ell}(s, a) \nu(s). \quad (2)$$

If there is only one admissible action, that is the cardinal of $\mathcal{A}_{t_\ell}^*$ is one, then the receiver picks up this action. If there are several admissible actions, the receiver picks the actions that are the most favorable to the sender. As soon as they compute their admissible strategies, they make it public to the rest of the agents and hence each agent is aware of $\Theta_t(\nu)$ given by

$$\Theta_t(\nu) = \left\{ a \in \mathcal{A} \mapsto \prod_{\ell \in \mathcal{L}} \theta_{t_\ell, \nu}(a_\ell) \text{ such that } \theta_{t_\ell, \nu} \in \Theta_{t_\ell}(\nu), \forall \ell \in \mathcal{L} \right\}.$$

Finally they choose a global strategy $\theta_{t, \nu}^* \in \Theta_t(\nu)$ a solution to

$$\max_{\theta \in \Theta_t(\nu)} \left(\sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} v(s, a) \nu(s) \theta(a) \right), \quad (3)$$

where, v is the utility of the sender and ν is her prior. Solutions to the linear programming problem (3) may not be unique, but one can always select a solution that corresponds to a vertex of the feasible polytope. Consequently, there exists at least one choice of $\theta_{t, \nu}^*$ that is a Dirac measure concentrated on a single action $a_t^*(\nu)$, such that for every type t_ℓ , we have $a_{t_\ell}^*(\nu) \in \mathcal{A}_{t_\ell}^*(\nu)$, i.e., an admissible action. By convention, we assume that receivers of type t_ℓ adopt such a (pure) strategy.

In other words, for every $t \in \mathcal{T}$ and $\nu \in \mathcal{P}(\mathcal{S})$, we define $a_t^*(\nu) \in \mathcal{A}_t^*(\nu)$, where $\mathcal{A}_t^*(\nu) = \bigotimes_{\ell \in \mathcal{L}} \mathcal{A}_{t_\ell}^*(\nu)$, as one solution to the following problem:

$$\sum_{s \in \mathcal{S}} v(s, a_t^*(\nu)) \nu(s) \geq \sum_{s \in \mathcal{S}} v(s, a) \nu(s), \quad \forall a \in \mathcal{A}_t^*(\nu).$$

Remark 1. *Throughout this section, we have implicitly assumed that all agents share the same prior $\nu \in \mathcal{P}(\mathcal{S})$. However, one can also handle the case of heterogeneous priors using the following trick: select any prior $\nu \in \mathcal{P}(\mathcal{S})$ such that each agent's prior is absolutely continuous with respect to ν , and, for each agent, scale its utility by the ratio of its own prior to ν . The agent will behave the same way under this new utility and the shared prior ν .*

2.4 The sender's problem : choosing the message

The sender is the main agent and she'll want to maximize her utility, denoted $v : \mathcal{S} \times \mathcal{A}_1 \times \dots \times \mathcal{A}_L \mapsto \mathbb{R}$, which depends on the state of the world and the actions from the receivers. We suppose that each agent share the same prior

$\mu \in \mathcal{P}(\mathcal{S})$ which is enhanced by a message conveyed by the sender. Indeed, she is able to design a communication policy (or signal) $\pi = (\pi(m|s))_{(m,s) \in \mathcal{M} \times \mathcal{S}}$ in a way that will influence the actions of receivers to her benefit. The value $\pi(m|s)$ is to be understood as the probability of receiving the message m given s , the state of the world. We observe that, for each state s , $\sum_{m \in \mathcal{M}} \pi(m|s) = 1$ and we denote $\mathcal{P}_{\mathcal{S}}(\mathcal{M})$ the set of available communication policies

$$\mathcal{P}_{\mathcal{S}}(\mathcal{M}) = \{\pi(m|s) \text{ such that } \pi(\cdot|s) \in \mathcal{P}(\mathcal{M}) \text{ for all } s\},$$

With a communication policy $\pi \in \mathcal{P}_{\mathcal{S}}(\mathcal{M})$ at hand, after receiving a message $m \in \mathcal{M}$, the receivers update their prior μ by Bayes' rule and compute their posterior

$$\nu_{m,\pi}(s) = \frac{\pi(m|s)\mu(s)}{p(m)} \quad \text{with} \quad p(m) = \sum_{\tilde{s} \in \mathcal{S}} \pi(m|\tilde{s})\mu(\tilde{s}), \quad (4)$$

where $p(m)$ is the probability of receiving the message m . With this posterior $\nu_{m,\pi}$, the receivers then compute their acceptable strategies $\Theta_{t,\nu_{m,\pi}}$ and reveal them to the sender. The sender then chooses amongst the available strategies the most favorable ones.

From here, the sender needs to design her communication policy π in order to maximize her utility. Let us notice that, when designing π , there are several random variables whose realizations she's not aware of: the state of the world, the message and the type of each receiver. The information on these variables is encoded through a distribution $\tilde{\eta} \in \mathcal{P}(\mathcal{S} \times \mathcal{M} \times \mathcal{T})$. Since the types are independent of the states and of the messages, we can factorize, for any $s \in \mathcal{S}$, $m \in \mathcal{M}$ and $t \in \mathcal{T}$,

$$\tilde{\eta}(s, m, t) = \tilde{\mu}(s, m)\eta(t) \text{ and } \eta(t) = \prod_{\ell \in \mathcal{L}} \eta_{\ell}(t_{\ell}),$$

where $\tilde{\mu}$ is computed thanks to her communication policy as well as her prior μ . That is, we have, for any $(s, m) \in \mathcal{S} \times \mathcal{M}$,

$$\tilde{\mu}(s, m) = \pi(m|s)\mu(s).$$

In the end, when designing π and by using (4), the sender faces the following maximization problem

$$\max_{\pi \in \mathcal{P}_{\mathcal{S}}(\mathcal{M})} \left(\sum_{(m,t) \in \mathcal{M} \times \mathcal{T}} \max_{\theta \in \Theta_t(\nu_{m,\pi})} \left(\sum_{(s,a) \in \mathcal{S} \times \mathcal{A}} v(s, a) \nu_{m,\pi}(s) \theta(a) \right) \eta(t) p(m) \right). \quad (5)$$

For a given t and m , the problem of maximization in θ has been discussed in Section 2.3 and the optimal solution is a Dirac located at the action $a_t^*(\nu_{m,\pi})$. Hence, the final version of the problem of the sender is

$$\max_{\pi \in \mathcal{P}_{\mathcal{S}}(\mathcal{M})} \left(\sum_{(m,t,s) \in \mathcal{M} \times \mathcal{T} \times \mathcal{S}} v(s, a_t^*(\nu_{m,\pi})) \nu_{m,\pi}(s) \eta(t) p(m) \right). \quad (6)$$

2.5 Limitations of the model

In this section, we discuss some limitations of the model.

The first caveat is somewhat technical: we have not explicitly emphasized that the mapping $(\nu, t) \mapsto a_t^*(\nu)$ must be measurable. However, this requirement is essential for the analysis to hold. Although it is theoretically possible to construct non-measurable mappings of the form $(\nu, t) \mapsto a_t^*(\nu)$, insisting on such pathological choices would be an unnecessarily adversarial stance.

A second consideration concerns the rule governing receivers when they are indifferent amongst several actions. The convention adopted here is that ‘the receiver pleases the sender.’ This choice has the advantage of making the sender’s utility upper semicontinuous with respect to the communication policy (see Lemma 2 below). An alternative rule is that, in cases of indifference, the receiver selects an action according to a predetermined distribution over the action set. This latter approach is closely related to the notion of regularization discussed in this article.

A third one is that the decision-making framework described in the previous section is not yet fully specified, as it may admit seemingly inconsistent action choices. This issue has significant implications, since the total number of admissible actions plays a critical role in estimating the optimal number of messages (see Theorem 1 below). In our setting, the space of all possible action profiles is given by

$$\mathcal{A}^T = \bigotimes_{t \in \mathcal{T}} \mathcal{A} = \{(a_t)_{t \in \mathcal{T}} \text{ such that } a_t \in \mathcal{A}, \forall t \in \mathcal{T}\},$$

which has cardinality A^T . However, allowing such a vast array of possible action profiles seems unnecessarily permissive. Indeed, suppose we have two different types t_1 and t_2 , and a group of receivers $\mathcal{R} \subset \mathcal{L}$ that share the same type, that is $(t_1)_{\mathcal{R}} = (t_2)_{\mathcal{R}}$. Further assume that the actions of all other receivers are identical across the two types, i.e., $(a_{t_1}^*(\nu))_{-\mathcal{R}} = (a_{t_2}^*(\nu))_{-\mathcal{R}}$. Then it seems reasonable?both from a modeling and intuitive standpoint?that the receivers in $\mathcal{R}$ should take the same action in both cases. Formally, we require that for all t_1, t_2 and all $\mathcal{R} \subset \mathcal{L}$,

$$\left((t_1)_{\mathcal{R}} = (t_2)_{\mathcal{R}} \text{ and } (a_{t_1}^*(\nu))_{-\mathcal{R}} = (a_{t_2}^*(\nu))_{-\mathcal{R}} \right) \Rightarrow (a_{t_1}^*(\nu))_{\mathcal{R}} = (a_{t_2}^*(\nu))_{\mathcal{R}}, \quad (7)$$

where, by $a_{-\mathcal{R}}$, we describe the actions taken by every receiver except for the ones in $\mathcal{R}$. Note, however, that it is entirely possible to fail to verify condition (7), as illustrated in Section 4.1.2. As stated in said section, a way to circumvent this issue is to suppose that the sender is never indifferent to the actions of the receivers.

2.6 The concavification of the sender's problem and the revelation principle

In this section, following the ideas of [19], we reformulate the sender's problem (6) into a linear programming problem set on the space of measures. For that purpose, for any prior $\nu \in \mathcal{P}(S)$, denote $W(\nu)$ the gain of the sender defined as

$$W(\nu) = \left(\sum_{(s,t) \in \mathcal{S} \times \mathcal{T}} v(s, a_t^*(\nu)) \nu(s) \eta(t) \right). \quad (8)$$

Then the sender's problem (6) is reformulated in

$$\max_{\pi \in \mathcal{P}_{\mathcal{S}}(\mathcal{M})} \left(\sum_{m \in \mathcal{M}} W(\nu_{m,\pi}) p(m) \right). \quad (9)$$

The first trick in the concavification of the problem relies on the introduction of τ_π which is a sum of Dirac masses at points $\nu_{m,\pi}$, that is

$$\tau_\pi = \sum_m p_m \delta_{\nu_{m,\pi}},$$

and to recast the problem into

$$\max_{\pi \in \mathcal{P}_{\mathcal{S}}(\mathcal{M})} \int_{\mathcal{P}(\mathcal{S})} W(\nu) d\tau_\pi(\nu). \quad (10)$$

The question is to describe the set of admissible τ_π when π spans $\mathcal{P}_{\mathcal{S}}(\mathcal{M})$. First remark that because $\sum p(m) = 1$, then $\tau_\pi \in \mathcal{P}(\mathcal{P}(\mathcal{S}))$. Moreover, the equation $\sum_m \nu_{m,\pi}(s) = \mu(s)$ ensures that $\int_{\mathcal{P}(\mathcal{S})} \nu d\tau_\pi(\nu) = \mu$. We are led to introduce $\mathcal{T}_{M,\mu}$, the set of probability measure on $\mathcal{P}(\mathcal{S})$ with expectation given by μ and convex combination of at most M Dirac masses.

$$\mathcal{T}_{M,\mu} = \left\{ \tau \in \mathcal{P}(\mathcal{P}(\mathcal{S})) \text{ such that } \int_{\mathcal{P}(\mathcal{S})} \nu d\tau(\nu) = \mu \text{ and } \#(\text{supp}(d\tau)) \leq M \right\}.$$

For any $\pi \in \mathcal{P}_{\mathcal{S}}(\mathcal{M})$, then τ_π belongs to $\mathcal{T}_{M,\mu}$. Reciprocally, if $\tilde{\tau} = \sum_m \rho(m) \delta_{\nu_m}$ is chosen in $\mathcal{T}_{M,\mu}$, then constructing $\pi(m, s) = \rho(m) \nu_m(s) / \mu(s)$, it is easy to check that $\pi \in \mathcal{P}_{\mathcal{S}}(\mathcal{M})$ and $\tilde{\tau} = \tau_\pi$. Hence, by a change of variable, the sender problem is equivalent to the linear programming problem

$$\max_{\tau \in \mathcal{T}_{M,\mu}} \int_{\mathcal{P}(\mathcal{S})} W(\nu) d\tau(\nu), \quad (11)$$

A companion problem to (11) is the one where the constraint on the number of messages is relaxed. We are then lead to consider the set $\mathcal{T}_\mu$, which is the set of probability measure on $\mathcal{P}(\mathcal{S})$ with expectation given by μ

$$\mathcal{T}_\mu = \left\{ \tau \in \mathcal{P}(\mathcal{P}(\mathcal{S})) \text{ such that } \int_{\mathcal{P}(\mathcal{S})} \nu d\tau(\nu) = \mu \right\}. \quad (12)$$

and the problem

$$\hat{W}(\mu) = \max_{\tau \in \mathcal{T}_\mu} \int_{\mathcal{P}(\mathcal{S})} W(\nu) d\tau(\nu). \quad (13)$$

Foreshadowing the developments in Section 3 and Theorem 1, Problem (13) is, in fact, equivalent to Problem (11) when the parameter M is sufficiently large. In this regime, a version of the revelation principle applies and each message is a prescribed action. Problem (13) turns into the following linear programming problem :

$$\max_{\pi \in \mathcal{P}_S(\mathcal{A}^\mathcal{T}), \pi \text{ verifies (15)}} \sum_{(a,s) \in \mathcal{A}^\mathcal{T} \times \mathcal{S}} \sum_{t \in \mathcal{T}} v(s, a_t) \pi(a|s) \mu(s) \eta(t), \quad (14)$$

where (15) is defined as

$$\sum_s u_{t_\ell}(s, (a_t)_\ell) \pi(a|s) \mu(s) \geq \sum_s u_\ell(s, \tilde{a}) \pi(a|s) \mu(s), \quad (15)$$

for any $t \in \mathcal{T}$, $\ell \in \mathcal{L}$, $\tilde{a} \in \mathcal{A}_\ell$ and $a \in \mathcal{A}^\mathcal{T}$.

2.7 Regularization by strict convexity

2.7.1 Problems with the standard formulation

In this section, we examine several shortcomings of the standard formulations (13) and (14), which motivate our proposal to regularize the receiver's problem. Although Problem (13) is linear in form, it is defined over a space of probability measures. Recall that for any function $f : A \rightarrow \mathbb{R}$, maximizing f is equivalent to maximizing $\int f d\mu$ over $\mu \in \mathbb{P}(A)$. Consequently, linear programming problems over $\mathbb{P}(A)$ are, in general, as challenging as directly maximizing a non-linear function over A . In our setting, since the set $\mathcal{S}$ is finite, the space $\mathbb{P}(\mathcal{S})$ can be identified with the standard simplex in dimension $S + 1$, and Problem (13) is therefore as difficult as maximizing a function over a space of dimension $S + 1$.

For formulation (14), the problem is linear in finite dimension. However, it involves $A^T \times S$ variables and approximately A^{T+1} constraints. Consequently, the sheer scale of variables and constraints often renders it unsolvable in practice. In addition, if the solution set happens to be a facet, the algorithm becomes highly sensitive to parameter changes, as even minimal variations make it switch between vertices, which destabilizes the problem-solving process.

2.7.2 Regularization

In the previous model, the receiver acts as a rational agent in choosing its action with respect to its utility. We now introduce a variant where he is biased in his choice of strategy and wishes to remain close to an arbitrary strategy (the "irrational strategy"). To do so, we rely on a divergence function on the space

$P(\mathcal{S})$ which will quantifies the dissimilarity between two elements of $P(\mathcal{S})$. The ℓ -th receiver problem (1) is reformulated as follows

$$\max_{\theta \in P(\mathcal{A}_\ell)} \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}_\ell} u_{t_\ell}(s,a) \theta(a) \nu(s) - \varepsilon \vartheta(\theta, \lambda_\ell), \quad (16)$$

for some "irrational strategy" $\lambda_\ell \in P(\mathcal{A}_\ell)$ and $\varepsilon > 0$. In practice, we choose ϑ to be the Kullback-Leibler divergence given by

$$\vartheta(\nu, \mu) = \begin{cases} \sum_{s \in \mathcal{S}} \nu(s) \log \left(\frac{\nu(s)}{\mu(s)} \right), & \text{if } \nu \ll \mu, \\ \infty, & \text{else.} \end{cases}$$

The parameter ε quantifies the degree of commitment to the irrational strategy.

Remark 2. A related model can be found in [22] where the authors introduce a costly information acquisition for the receiver. The cost is given by Shannon's entropy and connects to our model.

We assume that $\text{supp}(\lambda_\ell) = \mathcal{A}_\ell$, for any $\ell \in \{1, \dots, L\}$. The previous problem is strictly concave and admits a unique solution denoted $\theta_{t_\ell, \nu_m}^{*, \varepsilon}$. We denote, for any $(a, t) \in \mathcal{A} \times \mathcal{T}^L$ and any $\nu \in P(\mathcal{S})$,

$$\theta_{t, \nu}^{*, \varepsilon}(a) = \prod_{\ell=1}^L \theta_{t_\ell, \nu}^{*, \varepsilon}(a_\ell). \quad (17)$$

By following the same arguments as in the previous section and by denoting, for any $a \in \mathcal{A}^\mathcal{T}$ and any $\nu \in P(\mathcal{S})$,

$$\theta_\nu^{*, \varepsilon}(a) = \sum_{t \in \mathcal{T}} \theta_{t, \nu}^{*, \varepsilon}(a) \eta(t) \quad \text{and} \quad W^\varepsilon(\nu) = \sum_{(s,a) \in \mathcal{S} \times \mathcal{A}^\mathcal{T}} w(s,a) \nu(s) \theta_\nu^{*, \varepsilon}(a), \quad (18)$$

the sender's problem becomes

$$\max_{\tau \in \mathcal{T}_\mu} \left(\int_{P(\mathcal{S})} W^\varepsilon(\nu) d\tau(\nu) \right). \quad (19)$$

In contrast to Problem (13), this problem cannot be reformulated as a linear programming problem, since we have $\text{supp}(\theta_{t, \nu}^{*, \varepsilon}) = \mathcal{A}$. Nevertheless, it is possible to derive an explicit expression for $\theta_{t, \nu}^{*, \varepsilon}$, which obviously depends on the choice of the divergence ϑ . As we will see later in Section 3.3, first- or second-order methods can then be employed to solve this nonlinear optimization problem. In particular, this circumvent some of the limitations described in the previous section.

3 Main results

3.1 On the concavification

We now provide some results concerning Problem (11) and Problem (13). To begin with, their constraint spaces are compact as stated in the following lemma.

Lemma 1. *Let $M \in \mathbb{N}^*$. The spaces $\mathbb{T}_\mu$ and $\mathbb{T}_{M,\mu}$ are compact.*

To address Problem (13), the role of concavification becomes clear through the following result, which both establishes the existence of a solution and provides a method for computing the optimal value through bi-conjugate Fenchel transform of W .

Lemma 2. *The function W defined by (8) is upper semicontinuous. The function $\hat{W}$ is the smallest concave function which is greater or equal to W . It is upper semicontinuous and is then equal to the bi-conjugate Fenchel transform of W .*

A natural question when studying (13) and (11) is to determine the minimal value of M (the number of messages) such that a solution to (11) is also a solution to (13).

Remark 3. *Problem (13) is a linear programming problem with S constraints. As such, we know that it admits a solution $\tau \in \mathfrak{S}_S$ [8].*

The path to answering this question begins with the following decomposition lemma (Lemma 3) concerning optimal solutions and its corollary (Corollary 1).

Lemma 3. *Let $\tau \in \mathbb{T}_\mu$ be a solution to (13) such that there exist $r \in [0, 1]$, $\tau_1 \in \mathbb{T}_{\mu_1}$, and $\tau_2 \in \mathbb{T}_{\mu_2}$ satisfying $\tau = r\tau_1 + (1-r)\tau_2$, so that $\mu = r\mu_1 + (1-r)\mu_2$. Then, τ_1 and τ_2 are also solutions to (13) within their respective constraint sets.*

Corollary 1. *Let $\tau = \sum_{m=1}^M \rho_m \delta_{\nu_m} \in \mathbb{T}_{M,\mu}$ be a solution of (13). Then, for each m , we must have $W(\nu_m) = \hat{W}(\nu_m)$.*

Given the latest corollary, we are in position to answer to the question of the optimal number of messages with a sharp estimate

Theorem 1. *There exists a solution τ^* to (13) in $\mathbb{T}_\mu$ that is a finite sum of Dirac measures supported on points $(\nu_i)_i \in \mathcal{P}(\mathcal{S})$, such that for any two distinct indices $i \neq j$, the induced actions $a_t^*(\nu_i)$ and $a_t^*(\nu_j)$ differ for at least one type of one receiver. Since there are at most A^T such action profiles, problems (13) and (11) are equivalent whenever $M \geq A^T$. Conversely, for any given sets of receivers $\mathcal{L}$ and actions $\mathcal{A}$, if there is only one type, there exist sets of states $\mathcal{S}$ and utility functions u and v such that no solution to (13) exists with support of cardinality less than or equal to A .*

Theorem 1 is a revelation principle that states that there is a solution where each message $m = (a_t)_{t \in \mathcal{T}}$ is an element of $\mathcal{A}^{\mathcal{T}}$ of the form "If the configuration of the types is t , then perform the action $a_t \in \mathcal{A}$ ". Of course this message has to be credible, that is, once the posterior of this communication policy is computed, the action a_t has to be admissible for every receiver, that is the communication policy π has to verify, for every $a \in \mathcal{A}^{\mathcal{T}}$

$$\sum_s u_{t_\ell}(s, (a_t)_\ell) \pi(a|s) \mu(s) \geq \sum_s u_{t_\ell}(s, \tilde{a}) \pi(a|s) \mu(s), \forall t, \ell \in \mathcal{T} \times \mathcal{L}, \forall \tilde{a} \in \mathcal{A}_\ell. \quad (20)$$

We can restrict ourselves to studying such communication policies, and we obtain the following linear programming problem :

$$\max_{\pi \in \mathcal{P}_S(\mathcal{A}^\mathcal{T}), \pi \text{ verifies (20)}} \sum_{(a,s) \in \mathcal{A}^\mathcal{T} \times \mathcal{S}} \sum_{t \in \mathcal{T}} v(s, a_t) \pi(a|s) \mu(s) \eta(t). \quad (21)$$

3.2 On the regularization

We now turn to the regularized problem.

Lemma 4. *Let $(\tau^\varepsilon)_{\varepsilon>0}$ be a sequence in $\mathcal{T}_{M,\mu}$, with $M \in \mathbb{N}$, that converges to some $\tau^0 \in \mathcal{T}_{M,\mu}$. Then, we have*

$$\limsup_{\varepsilon \rightarrow 0} \int_{\mathcal{P}(\mathcal{S})} W^\varepsilon(\nu) d\tau^\varepsilon(\nu) \leq \int_{\mathcal{P}(\mathcal{S})} W(\nu) d\tau^0(\nu).$$

By using the previous lemma with a sequence of maximizers (which have a finite support of cardinality S , by Remark 3), we directly deduce the following corollary.

Corollary 2. *For any $\varepsilon > 0$, denote*

$$\bar{W}^{\varepsilon,*} = \max_{\tau \in \mathcal{T}_\mu} \left(\int_{\mathcal{P}(\mathcal{S})} W^\varepsilon(\nu) d\tau(\nu) \right) \quad \text{and} \quad \bar{W}^* = \max_{\tau \in \mathcal{T}_\mu} \left(\int_{\mathcal{P}(\mathcal{S})} W(\nu) d\tau(\nu) \right).$$

Then, we have

$$\limsup_{\varepsilon \rightarrow 0} \bar{W}^{\varepsilon,*} \leq \bar{W}^*.$$

In order to obtain the convergence of $\bar{W}^{\varepsilon,*}$ to $\bar{W}^*$ as ε goes to 0, we require the following assumption.

Assumption 1. *If $a \in \mathcal{A}^\mathcal{T}$ is a prescribed action for a certain prior, that is there exists $\nu \in \mathcal{P}(\mathcal{S})$ such that $a_t = a_t^*(\nu)$, for any $t \in \mathcal{T}$, then it is a forced action for a perhaps different prior with a support contained in $\text{supp}(\mu)$. That is, there exists $\tilde{\nu} \in \mathcal{P}(\mathcal{S})$ such that $\text{supp}(\tilde{\nu}) \subset \text{supp}(\mu)$ and $\mathcal{A}_t(\tilde{\nu}) = \{a_t\}$, for any $t \in \mathcal{T}$.*

Lemma 5. *Under Assumption 1, there exists a sequence $(\tau^\varepsilon)_{\varepsilon>0} \subset \mathcal{T}_\mu$ such that*

$$\int_{\mathcal{P}(\mathcal{S})} W^\varepsilon(\nu) d\tau^\varepsilon(\nu) \xrightarrow{\varepsilon \rightarrow 0} \bar{W}^*.$$

Corollary 3. *Under Assumption 1, we have*

$$\lim_{\varepsilon \rightarrow 0} \bar{W}^{\varepsilon,*} = \bar{W}^*.$$

Remark 4. *Assumption 1 is essential for the previous result as shown in Section 4.1.1 below.*

As a result, we have the following theorem.

Theorem 2. *Under Assumption 1, let $(\tau^\varepsilon)_{\varepsilon>0}$ be a sequence of maximizers of (19) such that, for some $M \in \mathbb{N}$ and any $\varepsilon > 0$, $\tau^\varepsilon \in \mathcal{T}_{M,\mu}$. Then, this sequence converges, up to a subsequence, to a maximizer of (13) as ε goes to 0.*

3.3 Numerical considerations

In this section, we provide a rationale for some of the numerical choices underlying our method. We begin with the now-classic Softmax function, which takes as input any vector $x \in \mathbb{R}^n$ and returns an element of the probability simplex Δ_n , defined componentwise as

$$\text{Softmax}(x)_i = \frac{e^{x_i}}{\sum_{j=1}^n e^{x_j}} \quad \text{for all } 1 \leq i \leq n.$$

This function is a staple in neural network classification tasks, typically serving as the final output layer. The reasons are straightforward: it provides a smooth, numerically stable, and easily differentiable way to turn scores into probabilities. Moreover, it plays exceptionally well with the Kullback-Leibler divergence, making it a natural fit in probabilistic modeling.

To be more precise, if ϑ is chosen as the Kullback-Leibler divergence then the explicit solution of the optimization problem (16) is given by

$$\theta_{t_\ell, \nu}^{\star, \varepsilon}(\cdot) = \lambda_\ell(\cdot) \text{Softmax} \left(\frac{\sum_s u_{t_\ell}(s, \cdot) \nu(s)}{\varepsilon} \right) \quad (22)$$

Choice of optimization variable. In our implementation, we invoke the Softmax function not once, but twice. Instead of directly optimizing over π under the usual probabilistic constraints, we adopt a relaxed formulation: we assume $\pi = \text{Softmax}(x)$ for some unconstrained variable x , and we perform second-order optimization directly on x . This sidesteps the need for constrained optimization, which can often be more cumbersome. In summary, the optimization pipeline unfolds as follows:

- Given a function $x : \mathcal{M} \times \mathcal{S} \rightarrow \mathbb{R}$, compute $\pi \in \mathcal{P}_{\mathcal{S}}(\mathcal{M})$ given by

$$\pi(\cdot \mid s) = \text{Softmax}(x(\cdot, s)), \quad \forall s \in \mathcal{S}.$$

- For each message $m \in \mathcal{M}$, compute $\nu_{m, \pi}$ and $p(m)$ by (4) and then for each receiver $\ell \in \mathcal{L}$ of type $t_\ell \in \mathcal{T}_\ell$, compute $\theta_{t_\ell, \nu_{m, \pi}}^{\star, \varepsilon}$ by (22) then $\theta_{\nu_{m, \pi}}^{\star, \varepsilon}$ by (18).
- Finally, evaluate the objective

$$\mathbf{v}(x) = \sum_{m \in \mathcal{M}} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}^T} v(s, a) \nu_{m, \pi}(s) \theta_{\nu_{m, \pi}}^{\star, \varepsilon}(a) p(m).$$

To maximize the objective $x \mapsto \mathbf{v}(x)$, we rely on a standard BFGS algorithm with Wolfe line search and tested tools for unconstrained smooth optimization. For gradient computations, we turn to PyTorch [23], which offers automatic differentiation and a highly optimized native implementation of the Softmax function. This makes it an ideal companion for the task at hand: fast, reliable, and doing most of the hard work under the hood.

Choice of λ_ℓ and the regularization parameter ε . The weighting function λ_ℓ is quite flexible: any choice is admissible as long as $\lambda_\ell(a) > 0$ for all $a \in \mathcal{A}_\ell$. For simplicity and to avoid introducing additional tuning parameters, we adopt the uniform distribution, i.e.,

$$\lambda_\ell(a) = \frac{1}{A_\ell} \quad \text{for all } a \in \mathcal{A}_\ell.$$

We now turn to the choice of the regularization parameter ε . A natural impulse is to take ε as small as possible in order to better approximate the unregularized problem. However, there is a catch: when ε becomes too small, the Softmax function begins to closely approximate the arg max operator. This is problematic, as the gradient of arg max is either zero or undefined, both of which are highly undesirable in a gradient-based optimization routine. This effect is illustrated below in Section 4.2.3. In practice, if utility values are roughly of order one, a value of $\varepsilon = 10^{-2}$ provides a good compromise: it preserves enough smoothness for reliable optimization while remaining close enough to the original (non-regularized) objective to yield meaningful results. If small values of ε are crucial, we propose in Section 4.2.3 to iteratively reduce ε using a strategy inspired by the interior-point method.

4 Examples

4.1 Analytical examples

4.1.1 Sharpness of the assumption of Theorem 2

This example aims to demonstrate a sequence of maximizers of the regularized problem whose utilities fail to converge to that of the original problem. In this setting, Assumption 1 is not satisfied, and the conclusion of Corollary 3 does not hold. We consider $L = 1$, $\mathcal{S} = \{s_1, s_2\}$, $\mathcal{A}_1 = \{a_1, a_2\}$ and $\mathcal{T}_1 = \{t_1\}$. We set the divergence to be the Kullback-Leibler divergence with $\lambda = (1/2, 1/2)$ and $\mu = (1/2, 1/2)$. The utilities are

$$u = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad \text{and} \quad v = \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix}.$$

In this case, we observe that, for any $\nu \in \mathcal{P}(\mathcal{S})$,

$$\sum_{s \in \mathcal{S}} u(s, a_1) \nu(s) = \nu(s_1) \quad \text{and} \quad \sum_{s \in \mathcal{S}} u(s, a_2) \nu(s) = 0.$$

Thus, we have $\mathcal{A}_{1,t_1}(\nu) = \{a_1\}$, for all $\nu \in \mathcal{P}(\mathcal{S})$ such that $\nu \neq \delta_{s_2}$, and $\mathcal{A}_{1,t_1}(\delta_{s_2}) = \{a_1, a_2\}$. Thus, Assumption 1 does not hold. Indeed, there exists no $\nu \in \mathcal{P}(\mathcal{S})$ such that $\mathcal{A}_{1,t_1}(\nu) = \{a_2\}$. We have, for any $\nu \in \mathcal{P}(\mathcal{S})$,

$$\theta_\nu^{*,\varepsilon} = \frac{e^{\nu(s_1)/\varepsilon}}{1 + e^{\nu(s_1)/\varepsilon}} \delta_{a_1} + \frac{1}{1 + e^{\nu(s_1)/\varepsilon}} \delta_{a_2},$$

and it follows that

$$\lim_{\varepsilon \rightarrow 0} \theta_{\nu}^{\star, \varepsilon} = \begin{cases} \delta_{a_1}, & \text{if } \nu \neq \delta_{s_2}, \\ (\delta_{a_1} + \delta_{a_2})/2, & \text{if } \nu = \delta_{s_2}. \end{cases}$$

We can see that

$$W^{\varepsilon}(\nu) = \frac{1}{1 + e^{\nu(s_1)/\varepsilon}},$$

so that a solution of (19) is

$$\tau = \frac{1}{2}\delta_{\delta_{s_1}} + \frac{1}{2}\delta_{\delta_{s_2}},$$

and it follows that $\bar{W}^{\star, \varepsilon} = \frac{1}{2} \left(\frac{1}{1+e^{1/\varepsilon}} + \frac{1}{2} \right) \rightarrow \frac{1}{4}$ as $\varepsilon \rightarrow 0$. However, we can see that

$$\theta_{\nu}^{\star} = \mathbb{1}_{\{\nu(s_1) > 0\}} \delta_{a_1} + \mathbb{1}_{\{\nu(s_1) = 0\}} \delta_{a_2},$$

which yields

$$W(\nu) = \mathbb{1}_{\{\nu(s_1) = 0\}} \quad \text{and} \quad \int_{P(S)} W(\nu) d\tau(\nu) = 1/2 = \bar{W}^{\star}.$$

In the end, we have $\lim_{\varepsilon \rightarrow 0} \bar{W}^{\star, \varepsilon} = 1/4 < \bar{W}^{\star} = 1/2$.

4.1.2 Inconsistency in admissible action profiles

In this section, we demonstrate that the model may be inconsistent, in the sense that condition (7) need not hold. We consider the case of two states, two receivers with the second one having two different types. That is, $L = 2$, $\mathcal{S} = \{s_1, s_2\}$, $\mathcal{A}_1 = \{a_{11}, a_{12}\}$, $\mathcal{A}_2 = \{a_{21}, a_{22}\}$, $\mathcal{T}_1 = \{t_1\}$ and $\mathcal{T}_2 = \{t_{21}, t_{22}\}$. The utilities are the following

$$\begin{aligned} v(\cdot, a_{11}, \cdot) &= v(\cdot, a_{12}, \cdot) = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix}, \quad u_{t_1} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, \\ u_{t_{21}} &= \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} \quad \text{and} \quad u_{t_{22}} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}. \end{aligned}$$

For any state $s \in \mathcal{S}$, we observe that both the sender and the first receiver are indifferent to the choice of actions in $\mathcal{A}_1$. Furthermore, the sender strictly prefer action a_{21} . The second receiver of the first type also strictly prefer action a_{21} while the second type is completely indifferent.

Thus, for any $\nu \in P(\mathcal{S})$, the admissible actions are

$$\mathcal{A}_{t_1}^{\star}(\nu) = \mathcal{A}_1, \quad \mathcal{A}_{t_{21}}^{\star}(\nu) = \{a_{21}\} \quad \text{and} \quad \mathcal{A}_{t_{22}}^{\star}(\nu) = \mathcal{A}_2.$$

It follows that there are several possible optimal actions in each case, but assume that the sender chooses

$$a_{t_1, t_{21}}^{\star}(\nu) = (a_{11}, a_{21}) \quad \text{and} \quad a_{t_1, t_{22}}^{\star}(\nu) = (a_{12}, a_{21}),$$

which clearly violates condition (7).

Now, consider again how the receivers arrive at their choices. As previously assumed, each receiver discloses their set of admissible actions, and together they select one that is most favorable to the sender. In this case, receiver a realizes that he is facing two distinct situations (because the admissible actions of receiver b differ), and so perceives no inconsistency.

However, if the receivers privately communicate their admissible sets of actions to the sender, after which the sender selects and reveals a final action, receiver a experiences an inconsistency because he observes b 's chosen action but not the set of actions available to him.

4.1.3 Regularization of the judge and prosecutor example

In this example, we carry out a step-by-step analysis of the well-known prosecutor-judge example from [19], applying a regularization based on the Kullback-Leibler divergence. We have a single receiver, one type, $\mathcal{S} = \{s_1, s_2\}$ and $\mathcal{A} = \{a_1, a_2\}$. The prior is $\mu = (0.7, 0.3)$ and the utilities are given by

$$u = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad v = \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix}.$$

We recall that, in this setting, the optimal solution to (13) is given by

$$\tau^0 = (1 - p_0)\delta_{s_1} + p_0\delta_{\nu_0},$$

where $p_0 = 2\mu(s_2)$ and $\nu_0 = (1/2, 1/2)$. Furthermore, the functions W and $\hat{W}$ are illustrated in Figure 1a.

For any $\nu \in \mathcal{P}(\mathcal{S})$, we compute

$$\sum_{s \in \mathcal{S}} u(s, a_1)\nu(s) = \nu(s_1) = 1 - \nu(s_2) \quad \text{and} \quad \sum_{s \in \mathcal{S}} u(s, a_2)\nu(s) = \nu(s_2),$$

so that

$$\begin{aligned} \theta_\nu^{*,\varepsilon} &= \frac{\lambda(a_1)}{\lambda(a_1) + \lambda(a_2)e^{(2\nu(s_2)-1)/\varepsilon}}\delta_{a_1} + \frac{\lambda(a_2)}{\lambda(a_2) + \lambda(a_1)e^{-(2\nu(s_2)-1)/\varepsilon}}\delta_{a_2} \\ &= (1 - \sigma_{\lambda,\varepsilon}(\nu(s_2)))\delta_{a_1} + \sigma_{\lambda,\varepsilon}(\nu(s_2))\delta_{a_2}, \end{aligned}$$

where

$$\sigma_{\lambda,\varepsilon}(x) = \frac{\lambda(a_2)}{\lambda(a_2) + \lambda(a_1)e^{-(2x-1)/\varepsilon}}.$$

Thus, we obtain

$$W^\varepsilon(\nu) = \sigma_{\lambda,\varepsilon}(\nu(s_2)).$$

The concavification of W^ε is such that

$$\hat{W}^\varepsilon(\nu) = \begin{cases} \sigma_{\lambda,\varepsilon}(r_\varepsilon) + \sigma'_{\lambda,\varepsilon}(r_\varepsilon)(r_\varepsilon - \nu(s_2)), & \text{if } \nu(s_2) \leq r_\varepsilon, \\ \sigma_{\lambda,\varepsilon}(\nu(s_2)), & \text{if } \nu(s_2) > r_\varepsilon, \end{cases}$$


where $r_\varepsilon \in [0.5, 1]$ is the solution of

$$\sigma_{\lambda,\varepsilon}(0) = \sigma_{\lambda,\varepsilon}(r_\varepsilon) + \sigma'_{\lambda,\varepsilon}(r_\varepsilon)r_\varepsilon.$$


In particular, a solution of (19) is

$$\tau^\varepsilon = (1 - p_\varepsilon)\delta_{s_1} + p_\varepsilon\delta_{\nu_\varepsilon},$$


where $p_\varepsilon = \mu(s_2)/r_\varepsilon$ and $\nu_\varepsilon = (1 - r_\varepsilon, r_\varepsilon)$. The function W^ε and $\hat{W}^\varepsilon$, as well as ν_ε , are depicted in Figures 1b-1c-1d for different irrational strategies λ . We observe that $\hat{W}^\varepsilon(\mu) \leq \hat{W}(\mu)$, for any choice of λ . Furthermore, $\hat{W}^\varepsilon(\mu)$ increases with $\lambda(a_2)$, which is naturally expected in this model.




(a) The unregularized problem



(b) Regularization with a uniform irrational strategy $\lambda = (1/2, 1/2)$



(c) Regularization with a gullible irrational strategy $\lambda = (1/4, 3/4)$



(d) Regularization with a stubborn irrational strategy $\lambda = (3/4, 1/4)$

Figure 1: The effects of regularization on the standard example of the judge and prosecutor for $\varepsilon = 0.1$ and different irrational strategies.

4.2 Numerical examples

The numerical examples below are all done with *BASIL* (Bayesian Signaling Library), a library in Python publicly available².

²<https://plmlab.math.cnrs.fr/degourna/bayesian-persuasion-by-regularization>

4.2.1 The voting problem

We first study the voting problem of [2] that consists in three receivers (voters A, voters B and voters C) which have to vote between two options. In this case the set of actions is the same for every receiver $\mathcal{A} = \{\text{Yes}, \text{No}\}$, there are three states $\mathcal{S} = \{A, B, C\}$ and the utility $u_i(s, a)$ of the three receivers are given by

$$u_1 = \begin{pmatrix} 1.1 & 0 \\ -1 & 0 \\ -1 & 0 \end{pmatrix}, u_2 = \begin{pmatrix} -1 & 0 \\ 1.1 & 0 \\ -1 & 0 \end{pmatrix}, u_3 = \begin{pmatrix} -1 & 0 \\ -1 & 0 \\ 1.1 & 0 \end{pmatrix},$$

The prior is given by $\mu = (1/3, 1/3, 1/3)$. The actions taken by the three receivers, as well as the optimal signal, is displayed in Figure 2.




Figure 2: The voting problem. A point in the triangle correspond to an element of $\mathcal{P}(\mathcal{S})$, the pure states are the vertices. The coloring corresponds to the different strategies of the receivers (voters) for a given state. The black diamond represents μ and the three stars represents the different ν for the three messages. Each star is located in one of the small triangle where the majority of voters approve the bill. In these triangles, the utility of the sender is 1, it is 0 elsewhere.

In this Figure, the prior is a point in the simplex Δ_3 , where the pure states are the vertices of the triangle, beginning at the lower left and proceeding in counter-clockwise order. The utility of the sender is 1 if there is a majority of "Yes" and 0 otherwise. Without persuasion, each receiver will vote "No" and the corresponding gain of the sender is 0. Consider the decomposition

$$\mu = \frac{1}{3} \sum_i \nu_i \quad \text{with } \nu_1 = \left(0, \frac{1}{2}, \frac{1}{2}\right), \nu_2 = \left(\frac{1}{2}, 0, \frac{1}{2}\right) \text{ and } \nu_3 = \left(\frac{1}{2}, \frac{1}{2}, 0\right).$$

For each ν_i two receivers will vote "Yes", hence the sender gains 1. For this particular choice of $d\tau = \sum_i \frac{1}{3} \delta_{\nu_i}$, the sender reaches the maximum of the utility. This decomposition is represented with black stars in Figure 2.

The optimal solution is obtained for $M = 3$ messages. For this problem, Figure 3 illustrates the evolution of the positions of ν_m and the corresponding probabilities p_m for different numbers of messages. The algorithm may terminate with an excessive number of messages, either because some messages have a very low probability of occurring (see Figure 3, middle or bottom) or because some messages are duplicated (see Figure 3, bottom). In *BASIL*, we employ post-processing techniques to decrease the number of messages, ultimately producing a solution that uses the minimal number of messages.

4.2.2 Numerical convergence of the minimizer with respect to ε

This example, albeit a bit artificial, is chosen to illustrate the convergence of the optimal signals when ε goes to 0. There are three states $\mathcal{S} = \{s_1, s_2, s_3\}$ and three receivers $\{R_1, R_2, R_3\}$ who need to choose among three actions $\mathcal{A} = \{a_1, a_2, a_3\}$. The utilities of the receivers are

$$u_{r_1} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad u_{r_2} = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 2 \\ 3 & 0 & 0 \end{pmatrix} \quad \text{and} \quad u_{r_3} = \begin{pmatrix} 0 & 3 & 0 \\ 2 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

The utility of the sender is, for any $s \in \mathcal{S}$,

$$v(s, a_1, a_1, a_2) = 2, \quad v(s, a_1, a_2, a_2) = 4, \quad v(s, a_2, a_1, a_1) = 2, \quad v(s, a_2, a_1, a_2) = 1, \\ v(s, a_2, a_3, a_1) = 1, \quad v(s, a_3, a_1, a_2) = 2 \quad \text{and} \quad v(s, a_3, a_1, a_3) = 4.$$

The prior is $\mu = (3/10, 4/10, 3/10)$ and the divergence is the Kullback-Leibler divergence with $\lambda = (1/3, 1/3, 1/3)$. We compute numerically the optimal signals with two messages for $\varepsilon \in \{10, 1, 10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}\}$ and depict them in Figure 4. As predicted by Theorem 2, we observe that the signals vary considerably for large values of ε before eventually converging.

4.2.3 The hunter problem

We now consider a scenario in which both the state and action spaces are relatively large. The states represent the locations of two deer within a three-mile square territory. This territory is divided into nine unit-square cells, each identified by coordinates $(i, j) \in \{1, 2, 3\}^2 = \mathcal{P}$. A state is therefore defined as a pair $s = (s_1, s_2) \in \mathcal{P}^2$, where s_1 and s_2 indicate the positions of the first and second deer, respectively.

The receiver in this setting is a hunter who must choose a cell in which to hunt each day, so that his action set is $\mathcal{A} = \mathcal{P}$. The hunter owns a cabin located at some position $p \in \mathcal{P}$, and he prefers to hunt near it. His payoff from choosing action a in state s is

$$u(s, a) = h(s, a) - d(p, a)/8,$$

where the function h captures the hunting outcome: it equals 0 if the chosen cell contains no deer, 1 if it contains exactly one deer, and 2 if both deer are present. The term $d(p, a) = |p_1 - a_1| + |p_2 - a_2|$ measures the distance between




Figure 3: The voting problem. On the left we display the evolution of ν_m through the iterations, the starting point is a diamond and the ending point is a circle. On the right, we display the value of p_m , the probability that the message is sent. When the said probability gets under 1%, the curves are dashed and the circle of the corresponding ending ν is not displayed. The number of messages is respectively 3,6,9 from the top to the bottom and the optimal number of messages is 3. On the middle, the three extra messages are not used whereas in the bottom, the algorithm ends up with more messages than needed, the gold, pink and brown messages are the same. The algorithm always end with the correct solution (dirac masses on each middle of the faces).




Figure 4: Numerical convergence of the minimizer with respect to ε . For different ε , we compute the optimal signals which are represented as pairs of stars connected by a dashed line (which passes through the prior, represented by a black diamond). Their color depends on the value of ε : the larger the ε , the darker the color.

the cabin and the hunting location, thus penalizing choices far from the cabin. The deer's exact positions are not known, but it is observed that they tend to stay close to each other and are attracted to areas with food. The availability of food in each cell is described by a function $f : \mathcal{P} \mapsto \{0, 1, 2\}$. Combining these behavioral tendencies yields a prior distribution over states:

$$\mu(s) = \frac{\nu(s)\kappa(s)}{Z},$$

where Z is a normalization constant, $\kappa(s) = 2^{\mathbb{1}_{\{s_1=s_2\}}}$ reflects the tendency of the deer to remain together and $\nu(s) = 2^{f(s_1)+f(s_2)}$ expresses the influence of food availability.

The second agent, the sender, is a forest ranger who aims to protect a designated conservation zone. This zone consists of three adjacent cells within the territory, formally represented as the set $\mathcal{J} = \{p_k\}_{1 \leq k \leq 3} \subset \mathcal{P}$. Each day, the ranger patrols the entirety of the territory and communicates a signal to the hunter in order keep him out of the sanctuary. When the hunter chooses the action a , he's utility is equal to 0 if $a \notin \mathcal{J}$ and -1 otherwise. That is, for any state s ,

$$v(s, a) = -\mathbb{1}_{\mathcal{J}}(a).$$

With this model, we now turn to numerical simulations with $\varepsilon = 10^{-4}$. For this, we choose $f(q)$ to be equal to one for $q \in \{(1, 1), (3, 1), (2, 2), (2, 3)\}$ and

two for $q = (3, 3)$. The protected area is set as $\mathcal{J} = \{(2, 2), (3, 2), (3, 3)\}$ and the hunter's cabin is located at $p = (1, 3)$. The computed prior can be seen in Figure 5.




Figure 5: Prior for the hunter problem. The number in the upper-half (resp. the lower-half) of each square is the probability (rounded up) of the presence of a single deer (resp. two deers). Special attention must be given to the upper right cell which is a protected area, close to the cabin and with a high probability of deers.

We expect the optimal number of messages to be lower than $|\mathcal{A}| = 9$. It turns out that we can find an optimal signal with 3 messages as depicted in Figure 6 where the expected utility of the ranger is (almost) 0.

However, because of the randomization of the initial conditions, the algorithm does not consistently converge to the optimal signal. Figure 7 reports the statistics, based on one hundred realizations, of the expected utility of the sender for randomized initial data under varying values of ε . As the figure illustrates, the regularization process tends to obscure information from the optimization problem, so that the optimal value is attained only when ε is sufficiently small. However, for small values of ε , the algorithm tends to yield (very) bad signals when the number of messages is small. We observe that increasing the number of messages improves the convergence of the algorithm toward signals that more closely approximate the optimal one. Our interpretation is that artificially enlarging the message space strengthens the algorithm's exploratory and selective capacity (has also seen in Section 4.2.1), thereby facilitating convergence to the optimal set of messages. Based on this observation, we iteratively solve the

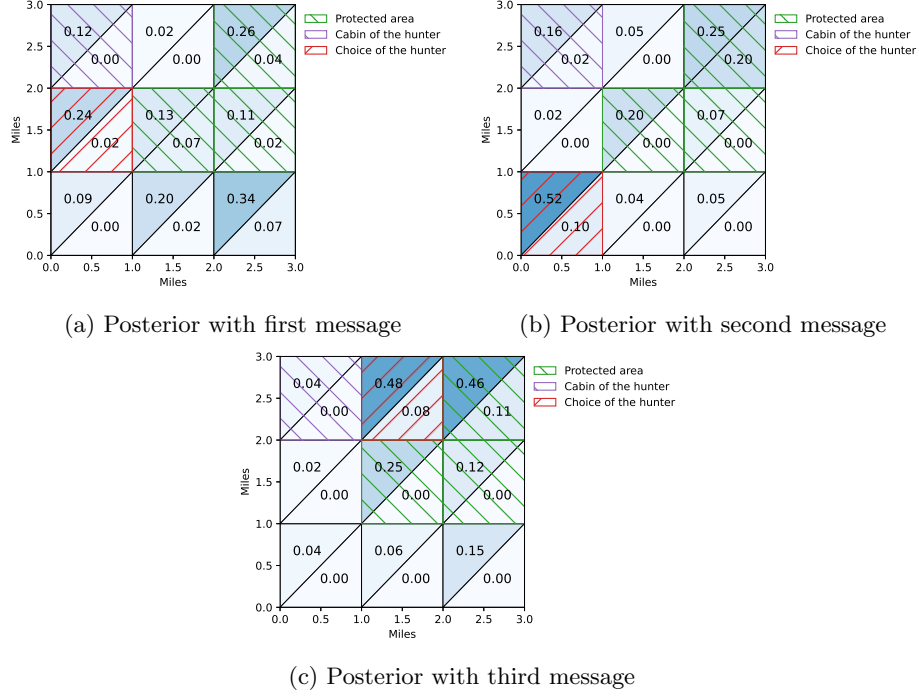


Figure 6: Depiction of an optimal signal for the hunter problem. The probability (rounded up) of receiving each message is respectively 0.366 for the first, 0.278 for the second and 0.356 for the third. Note that, for each message, the hunter is sent to an unprotected area.

optimization problem for decreasing values of ε of the form 10^{-k} , where k is linearly spaced between 0 and 4, yielding a total of ten distinct values. This gives a more robust and effective approach whose results can be seen in the last statistic, titled "Varying ε ". This approach is natively implemented in the *BASIL* library.

5 Conclusion

This article presents a new approach to solving the Bayesian persuasion problem based on regularization methods. Our method has the advantage of ensuring the receiver's solution is unique (and explicit in certain cases), which makes it possible to employ first- and second-order optimization methods. We prove that the solution of the regularized problem converges to that of the original problem as ε tends to 0. In addition, we provide a version of the revelation principle that allows one to determine the optimal number of messages for a given problem. Through various numerical examples, we examine the strengths and limitations




Figure 7: Efficiency of the algorithm for random initial data depending on the values of ε for the hunter problem.

of our method and justify our numerical choices. For these experiments, we developed a Python library, *BASIL*, which is publicly available and ensures the reproducibility of the reported results.

6 Proofs

6.1 Proof of Lemma 1

We know that the space $\mathcal{P}(\mathcal{S})$ endowed with the Prokhorov metric is a compact space so that the space $\mathcal{P}(\mathcal{P}(\mathcal{S}))$ (also endowed with the Prokhorov metric) is a compact space. Let $\mu \in \mathcal{P}(\mathcal{S})$ and $(\tau_n)_{n \in \mathbb{N}}$ be a sequence in $\mathcal{T}_\mu$. We know that, up to a subsequence, $(\tau_n)_{n \in \mathbb{N}}$ converges to an element $\tau \in \mathcal{P}(\mathcal{P}(\mathcal{S}))$. Denote

$$\eta = \int_{\mathcal{P}(\mathcal{S})} \nu d\tau(\nu).$$

We wish to prove that $\eta = \mu$. For any $s \in \mathcal{S}$, we consider the continuous bounded function $f_s : \mathcal{P}(\mathcal{S}) \mapsto [0, 1]$ given by

$$f_s(\nu) = \nu(s).$$

Thus, we know that, for any $s \in \mathcal{S}$, up to a subsequence,

$$\mu(s) = \int_{\mathcal{P}(\mathcal{S})} f_s(\nu) d\tau^n(\nu) \xrightarrow{n \rightarrow \infty} \int_{\mathcal{P}(\mathcal{S})} f_s(\nu) d\tau(\nu) = \nu(s),$$

which proves that $\tau \in \mathcal{T}_\mu$ and, thus, $\mathcal{T}_\mu$ is compact.

Let $M \in \mathbb{N}^*$, $\mu \in \mathcal{P}(\mathcal{S})$ and $(\tau_n)_{n \in \mathbb{N}}$ be a sequence in $\mathcal{T}_{M,\mu}$. We have that

$$\tau_n = \sum_{m=1}^M p_{m,n} \delta_{\nu_{m,n}}.$$

Since, for any $n \in \mathbb{N}$, $(p_{m,n}, \nu_{m,n})_{1 \leq m \leq M} \in ([0, 1] \times \mathcal{P}(\mathcal{S}))^M$, where $([0, 1] \times \mathcal{P}(\mathcal{S}))^M$ is a compact space, we know that there exists $(p_m, \nu_m)_{1 \leq m \leq M} \in ([0, 1] \times \mathcal{P}(\mathcal{S}))^M$ such that, up to a subsequence,

$$\tau_n \xrightarrow{n \rightarrow \infty} \sum_{m=1}^M p_m \delta_{\nu_m} \in \mathcal{T}_{M,\mu},$$

which proves the desired result.

6.2 Proof of Lemma 2

We denote $\text{co}(W)$ the convex hull of $\text{Gr}(W) = \{(\nu, W(\nu)) \text{ for any } \nu \in \mathcal{P}(\mathcal{S})\}$. That is, for any $\nu \in \mathcal{P}(\mathcal{S})$ and $(\nu, v) \in \text{co}(W)$, there exists $\tau \in \mathcal{T}_\nu$ such that

$$v = \int_{\mathcal{P}(\mathcal{S})} W(\nu) d\tau(\nu).$$

In other words, $\text{co}(W)$ is also given by

$$\text{co}(W) = \left\{ \left(\mu, \int_{\mathcal{P}(\mathcal{S})} W(\nu) d\tau(\nu) \right) \text{ for any } \mu \in \mathcal{P}(\mathcal{S}) \text{ and } \tau \in \mathcal{T}_\mu \right\}.$$

We define, for any $\nu \in \mathcal{P}(\mathcal{S})$,

$$\tilde{W}(\nu) = \sup\{v \text{ such that } (\nu, v) \in \text{co}(W)\},$$

so that, if W is upper semicontinuous, we have $\tilde{W} = \hat{W}$. The proof is decomposed in several points.

6.2.1 W is upper semicontinuous

Let $(\nu_n)_{n \in \mathbb{N}} \subset \mathcal{P}(\mathcal{S})$ be a sequence that converges to ν as $n \rightarrow \infty$. For any $t \in \mathcal{T}$, since $(a_t^*(\nu_n))_{n \in \mathbb{N}} \subset \mathcal{A}$, we know that it converges, up to a subsequence, to a $a_t^* \in \mathcal{A}$. For any $\ell \in \{1, \dots, L\}$ and $t \in \mathcal{T}$, we have, for any $a_\ell \in \mathcal{A}_\ell$,

$$\sum_{s \in \mathcal{S}} u_{t_\ell}(s, a_\ell) \nu_n(s) \leq \sum_{s \in \mathcal{S}} u_{t_\ell}(s, (a_t^*(\nu_n))_\ell) \nu_n(s).$$

Thus, by passing to the limit, up to a subsequence, in the previous inequality yields that $a_t^* \in \mathcal{A}_t(\nu)$. It follows that, up to a subsequence,

$$\begin{aligned} W(\nu_n) &= \sum_{(s,t) \in \mathcal{S} \times \mathcal{T}} v(s, a_t^*(\nu_n)) \nu_n(s) \eta(t) \\ &\xrightarrow{n \rightarrow \infty} \sum_{(s,t) \in \mathcal{S} \times \mathcal{T}} v(s, a_t^*) \nu(s) \eta(t) \\ &\leq \sum_{(s,t) \in \mathcal{S} \times \mathcal{T}} v(s, a_t^*(\nu)) \nu(s) \eta(t) = W(\nu), \end{aligned}$$

so that $\limsup_{n \rightarrow \infty} W(\nu_n) \leq W(\nu)$, which is the desired result.

6.2.2 The function $\tilde{W}$ is concave

Indeed, for any $\nu_1, \nu_2 \in \mathcal{P}(\mathcal{S})$ and $r \in [0, 1]$, we notice that

$$(r\nu_1 + (1-r)\nu_2, rW(\nu_1) + (1-r)W(\nu_2)) \in \text{co}(W).$$

Let $(W_1^n)_{n \in \mathbb{N}}$ and $(W_2^n)_{n \in \mathbb{N}}$ such that $(\nu_1, W_1^n), (\nu_2, W_2^n) \in \text{co}(W)$ and

$$W_1^n \xrightarrow{n \rightarrow \infty} \tilde{W}(\nu_1) \quad \text{and} \quad W_2^n \xrightarrow{n \rightarrow \infty} \tilde{W}(\nu_2).$$

We have, for any $n \in \mathbb{N}$,

$$(r\nu_1 + (1-r)\nu_2, rW_1^n + (1-r)W_2^n) \in \text{co}(W),$$

so that

$$rW_1^n + (1-r)W_2^n \leq \tilde{W}(r\nu_1 + (1-r)\nu_2),$$

which yields, by passing to the limit $n \rightarrow \infty$,

$$r\tilde{W}(\nu_1) + (1-r)\tilde{W}(\nu_2) \leq \tilde{W}(r\nu_1 + (1-r)\nu_2).$$

Thus, $\tilde{W}$ is concave.

6.2.3 The function $\tilde{W}$ is the smallest concave function greater than W

Let $V : \mathcal{P}(\mathcal{S}) \mapsto \mathbb{R}$ be a concave function. On one hand, for any $(\nu, v) \in \text{co}(V)$, there exists $\tau \in \mathcal{T}_\nu$ such that

$$v = \int_{\mathcal{P}(\mathcal{S})} V(\nu) d\tau(\nu) \leq V(\nu),$$

where we used the fact that V is concave. It follows that $\tilde{V}(\nu) \leq V(\nu)$. On the other hand, since $\text{Gr}(V) \subset \text{co}(V)$, we have that $V(\nu) \leq \tilde{V}(\nu)$. We conclude that $\tilde{V} = V$ if V is concave.

Now, let $H : \mathbf{P}(\mathcal{S}) \mapsto \mathbb{R}$ be such that $W \leq H$. For any $(\nu, w) \in \text{co}(W)$, there exists $\tau \in \mathbf{T}_\nu$ such that

$$w = \int_{\mathbf{P}(\mathcal{S})} W(\nu) d\tau(\nu) \leq \int_{\mathbf{P}(\mathcal{S})} H(\nu) d\tau(\nu),$$

and, in particular, we have $\tilde{H}(\nu) \geq w$ which yields $\tilde{W}(\nu) \leq \tilde{H}(\nu)$.

It follows that, for any $V : \mathbf{P}(\mathcal{S}) \mapsto \mathbb{R}$ which is concave and such that $W \leq V$, we have $\tilde{W} \leq V$. Thus, $\tilde{W}$ is the smallest concave function greater than W .

6.3 Proof of Lemma 3

We proceed by contradiction and assume that there exists $\tilde{\tau}_1 \in \mathbf{T}_{\mu_1}$ such that

$$\int_{\mathbf{P}(\mathcal{S})} W(\nu) d\tau_1(\nu) < \int_{\mathbf{P}(\mathcal{S})} W(\nu) d\tilde{\tau}_1(\nu).$$

Then, by setting $\tilde{\tau} = r\tilde{\tau}_1 + (1-r)\tau_2$, we observe that $\tilde{\tau} \in \mathbf{T}_\mu$ and

$$\begin{aligned} \int_{\mathbf{P}(\mathcal{S})} W(\nu) d\tilde{\tau}(\nu) &= r \int_{\mathbf{P}(\mathcal{S})} W(\nu) d\tilde{\tau}_1(\nu) + (1-r) \int_{\mathbf{P}(\mathcal{S})} W(\nu) d\tau_2(\nu) \\ &> r \int_{\mathbf{P}(\mathcal{S})} W(\nu) d\tau_1(\nu) + (1-r) \int_{\mathbf{P}(\mathcal{S})} W(\nu) d\tau_2(\nu) = \int_{\mathbf{P}(\mathcal{S})} W(\nu) d\tau(\nu), \end{aligned}$$

which contradicts the fact that τ is a solution of (11). The same arguments hold for τ_2 .

6.4 Proof of Corollary 1

By Remark 3, we know that τ can be decomposed as

$$\tau = \sum_{m=1}^M p_m \delta_{\nu_m},$$

for some $M \leq |\mathcal{S}|$, $(p_m)_{1 \leq m \leq M} \in \Delta_M$ and $(\nu_m)_{1 \leq m \leq M} \subset \mathbf{P}(\mathcal{S})$ verifying $\sum_{m=1}^M p_m \nu_m = \mu$. We observe that we can decompose $\tau = r\tau_1 + (1-r)\tau_2$ with

$$\tau_1 = \delta_{\nu_1}, \quad \tau_2 = \sum_{m=2}^M p_m \delta_{\nu_m} \quad \text{and} \quad r = p_1.$$

With have $\tau_1 \in \mathbf{T}_{\nu_1}$ and, by Lemma 3, τ_1 is a solution of (13). Thus, we have

$$W(\nu_1) = \int_{\mathbf{P}(\mathcal{S})} W(\nu) d\tau_1(\nu) = \hat{W}(\nu_1).$$

By iterating this argument, we obtain that $W(\nu_m) = \hat{W}(\nu_m)$, for any $1 \leq m \leq M$.

6.5 Proof of Theorem 1

6.5.1 Upper bound on the number of messages

We now discuss the number of possible actions. We start with the first receiver ℓ and for each type t_ℓ , we choose its action in $\mathcal{A}_\ell$, there are $A_\ell^{T_\ell}$ choices and iterating over the receivers, the set of available actions is $\mathcal{A}^\# = \bigotimes_{\ell \in \mathcal{L}} \bigotimes_{t_\ell \in \mathcal{T}_\ell} \mathcal{A}_\ell$ that can be rewritten as:

$$\mathcal{A}^\# = \{a^\# = ((a_{t_\ell}^\#)_{t_\ell \in \mathcal{T}_\ell})_{\ell \in \mathcal{L}} \text{ such that } a_{t_\ell}^\# \in \mathcal{A}_\ell \quad \forall \ell \in \mathcal{L}, t_\ell \in \mathcal{T}_\ell\}$$

This set is of cardinal $\prod_{\ell \in \mathcal{L}} A_\ell^{T_\ell}$ and can be rewritten as .

We tensorize the action by types, that is we consider a new set of actions $\mathcal{A}^\# = \bigotimes_{t \in \mathcal{T}} \mathcal{A}$ of cardinal A^T which is defined as

For any $a^T \in \mathcal{A}^T$, we introduce the set $\mathcal{C}(a^T)$ which is the set of priors ν such that, for each type t the action a_t^T is admissible for the receiver ℓ of type t_ℓ , that is

$$\mathcal{C}(a^T) = \{\nu \in \mathcal{P}(\mathcal{S}) \text{ such that, for each } t \in \mathcal{T}, \ell \in \mathcal{L}, \text{ then } (a_t^T)_\ell \in \mathcal{A}_{t_\ell}^*(\nu)\}$$

The set $\mathcal{C}(a^T)$ is convex, indeed ν belongs to $\mathcal{C}(a^T)$ if and only if we have for every $t \in \mathcal{T}, \ell \in \mathcal{L}$ and for all $a \in \mathcal{A}_\ell$,

$$\sum_s u_{t_\ell}(s, a) \nu(s) \leq \sum_s u_{t_\ell}(s, (a_t^\#)_\ell) \nu(s).$$

If the above inequality is true for ν_1 and ν_2 , it is surely true for $r\nu_1 + (1-r)\nu_2$ for any $r \in [0, 1]$, and this proves that $\mathcal{C}(a^T)$ is convex.

Let $\tau \in \mathcal{T}_\mu$ be a solution of (13) and for each $a^\# \in \mathcal{A}^\#$ denote $\nu(a^\#)$, the average of the prior ν over the set of prior that admits $a^\#$ as action policies.

$$p(a^\#) = \int_{a^*(\nu)=a^\#} d\tau(\nu) \text{ and } \nu(a^\#) = \frac{\int_{a^*(\nu)=a^\#} \nu d\tau(\nu)}{p(a^\#)}.$$

Note that, thanks to the discussion in preamble of this section, we restrict our analysis to the set $\mathcal{A}^\# \subset \mathcal{A}^T$ and $\mathcal{P}(\mathcal{S})$ is partitionned into the different sets $\{\nu, a^*(\nu) = a^\#\}_{a^\# \in \mathcal{A}^\#}$. Introducing $\tilde{\tau} = \sum_{a^\# \in \mathcal{A}^\#} p(a^\#) \delta_{\nu(a^\#)}$, we have

$$\int_{\mathcal{P}(\mathcal{S})} \nu d\tilde{\tau}(\nu) = \sum_{a^\# \in \mathcal{A}^\#} p(a^\#) \nu(a^\#) = \sum_{a^\# \in \mathcal{A}^\#} \int_{a^*(\nu)=a^\#} \nu d\tau(\nu) = \int_{\mathcal{P}(\mathcal{S})} \nu d\tau(\nu) = \mu$$

Similarly, we prove that $\sum_{a^\#} p(a^\#) = 1$, so that $\tilde{\tau}$ belongs to $\mathcal{T}_\mu$ and is admissible in Problem (13). Because $a^*(\nu)$ is chosen amongst the admissible actions, then

$$a^*(\nu) = a^\# \Rightarrow \nu \in \mathcal{C}(a^\#).$$

By definition, $\nu(a^\#)$ is an average of priors ν which all belong to the convex set $\mathcal{C}(a^\#)$, hence $\nu(a^\#) \in \mathcal{C}(a^\#)$. Hence

$$W(\nu(a^\#)) = \sum_{s,t} v(s, a_t^*(\nu(a^\#))) \nu(a^\#) \eta(t) \geq \sum_{s,t} v(s, a_t^\#) \nu(a^\#) \eta(t)$$

It follows that

$$\begin{aligned}
\int_{\mathcal{P}(\mathcal{S})} W(\nu) d\tilde{\tau}(\nu) &= \sum_{a^\sharp} p(a^\sharp) W(\nu(a^\sharp)) \geq \sum_{a^\sharp} \sum_{s,t} v(s, a^\sharp) p(a^\sharp) \nu(a^\sharp) \eta(t) \\
&\geq \sum_{a^\sharp} \sum_{s,t} v(s, a^\sharp) \left(\int_{a^\star(\nu)=a^\sharp} \nu d\tau(\nu) \eta(t) \right) \\
&\geq \int_{\mathcal{P}(\mathcal{S})} \sum_{s,t} v(s, a^\star(\nu)) \nu d\tau(\nu) \eta(t) = \int_{\mathcal{P}(\mathcal{S})} W(\nu) d\tau(\nu)
\end{aligned}$$

Because $d\tau$ is optimal, so is $d\tilde{\tau}$ and every inequality becomes an equality and $a^\sharp$ is not only admissible for $\nu(a^\sharp)$ but can be defined as the action policy taken by the receivers. That is, we can suppose that $a^\sharp = a_t^\star(\nu(a^\sharp))$. Hence, as claimed, each Dirac measure that compose the optimal $d\tilde{\tau}$ is associated to a different action policy $a^\sharp$. Finally there is at most $\prod_\ell A_\ell^{\mathcal{T}_\ell}$ of them.

6.5.2 Lower bound on the number of messages

We are given $\mathcal{A}$ and we suppose that we have only one type, that is $T = 1$. We want to design utility functions u and v on a precise space state $\mathcal{S}$ so that there is no optimal solution to (13) with support of cardinality lower than or equal to A . For that purpose, we suppose that $\mathcal{S} = \mathcal{A}$. For each $s = a \in \mathcal{S}$ we define the utility of the receivers and the one of the sender to be

$$u_{t_\ell}(s, a_\ell) = \begin{cases} 1 & \text{if } a_\ell = s_\ell \\ 0 & \text{if } a_\ell \in \mathcal{A}_\ell \setminus \{s_\ell\} \end{cases} \quad \text{and} \quad v(s, a) = \begin{cases} 1 & s = a \\ 0 & \text{if not} \end{cases}.$$

Clearly, for each $s \in \mathcal{S}$, if $\nu_s = \delta_s$, then there is only one admissible action which is s for which the gain of the sender is $W(\nu_s) = 1$. Take any $\mu = \sum_{s \in \mathcal{S}} \mu(s) \nu_s$ which is not located on the vertices of $\mathcal{P}(\mathcal{S})$, that is its support is of cardinal at least 2, or equivalently $\mu(s) < 1$ for every s , recall that, for this particular μ , the receiver chose a certain action $a^\star(\mu)$ and we have, by construction of v ,

$$W(\mu) = \sum_s v(s, a^\star(\mu)) \mu(s) = \mu(a^\star(\mu))$$

Consider now $\tau = \sum_s \mu(s) \delta_{\nu_s}$, where $\nu_s = \delta_s$. It is easy to check that $\tau \in \mathcal{T}_\mu$ and then

$$\hat{W}(\mu) \geq \int_{\mathcal{P}(\mathcal{S})} W(\nu) d\tau(\nu) = \sum_s \mu(s) W(\nu_s).$$

As we said before the sender gains 1 for ν_s , so that we obtain

$$\hat{W}(\mu) \geq \sum_s \mu(s) > \mu(a^\star(\mu)) = W(\mu)$$

So that for every measure μ which is not of the form $\mu = \nu_s$, we have $\hat{W}(\mu) > W(\mu)$. Now take any μ such that $\mu(s) > 0$ for every s and take τ optimal for

the problem of finding $\hat{W}(\mu)$. If τ is of finite support with M Dirac masses, then each Dirac mass must be supported on a ν_s such that $\nu_s = \delta_s$ for some $s \in \mathcal{S}$. And then τ must be written as $\tau = \sum_s \tau_s \delta_{\nu_s}$ and the condition $\int \nu \tau(\nu) \mu$ imposes $\tau_s = \mu(s) \neq 0$ for every s and hence τ has support of cardinality $S = A$.

6.6 Proof of Lemma 4

We denote

$$\tau^\varepsilon = \sum_{m=1}^M p_m^\varepsilon \delta_{\nu_m^\varepsilon},$$

where $(p_m^\varepsilon)_{1 \leq m \leq M} \in \Delta_M$ and $(\nu_m^\varepsilon)_{1 \leq m \leq M} \subset \mathcal{P}(\mathcal{S})$ verifying $\sum_{m=1}^M p_m^\varepsilon \nu_m^\varepsilon = \mu$. We also denote, for any $\ell \in \{1, \dots, L\}$ and any $t \in \mathcal{T}_\ell$,

$$\theta_m^{\star, \varepsilon} = (\theta_{\ell, t_\ell, \nu_m^\varepsilon}^{\star, \varepsilon})_{\substack{1 \leq \ell \leq L \\ t \in \mathcal{T}^L}} \in \bigoplus_{\substack{1 \leq \ell \leq L \\ t \in \mathcal{T}^L}} \mathcal{P}(\mathcal{A}_\ell) = \mathfrak{X}.$$

Since $\mathfrak{X}^M$ is a compact space, we know that, up to a subsequence,

$$(p_m^\varepsilon, \nu_m^\varepsilon, \theta_m^{\star, \varepsilon})_{1 \leq m \leq M} \xrightarrow{\varepsilon \rightarrow 0} (p_m, \nu_m, \theta_m^\star)_{1 \leq m \leq M},$$

with

$$\tau^0 = \sum_{m=1}^M p_m \delta_{\nu_m}.$$

We denote, for any $1 \leq m \leq M$,

$$(\theta_{\ell, t_\ell, m}^\star)_{\substack{1 \leq \ell \leq L \\ t \in \mathcal{T}^L}} = \theta_m^\star.$$

Since, for any $\ell \in \{1, \dots, L\}$ and any $t_\ell \in \mathcal{T}_\ell$, $\theta_{\ell, t_\ell, \nu_m^\varepsilon}^{\star, \varepsilon}$ is a solution of (16), we have, for every $\theta \in \mathcal{P}(\mathcal{A}_\ell)$,

$$\begin{aligned} & \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}_\ell} u_{t_\ell}(s, a) \nu_m^\varepsilon(s) \theta(a) - \varepsilon \vartheta(\theta, \lambda_\ell) \\ & \leq \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}_\ell} u_{t_\ell}(s, a) \nu_m^\varepsilon(s) \theta_{\ell, t_\ell, \nu_m^\varepsilon}^{\star, \varepsilon}(a) - \varepsilon \vartheta(\theta_{\ell, t_\ell, \nu_m^\varepsilon}^{\star, \varepsilon}, \lambda_\ell) \\ & \leq \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}_\ell} u_{t_\ell}(s, a) \nu_m^\varepsilon(s) \theta_{\ell, t_\ell, \nu_m^\varepsilon}^{\star, \varepsilon}(a). \end{aligned}$$

Thus, by passing to the limit $\varepsilon \rightarrow 0$ (of the subsequence), we deduce that

$$\max_{\theta \in \mathcal{P}(\mathcal{A}_\ell)} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}_\ell} u_{t_\ell}(s, a) \nu_m(s) \theta(a) = \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}_\ell} u_{t_\ell}(s, a) \nu_m(s) \theta_{\ell, t_\ell, m}^\star(a),$$

so that $\theta_{\ell, t_\ell, m}^*$ belongs in $\Theta_{\ell, t_\ell}(\nu_m)$. In particular, by denoting $\theta_{t, m}^* = \prod_{\ell=1}^L \theta_{\ell, t_\ell, m}^*$, we observe that $\theta_{t, m}^* \in \Theta_t(\nu_m)$ and, thus,

$$\begin{aligned} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} v(s, a) \theta_{t, m}^*(a) \nu_m(s) &\leq \max_{\theta \in \Theta_t(\nu_m)} \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} v(s, a) \theta(a) \nu_m(s) \\ &\leq \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}} v(s, a) \theta_{t, \nu_m}^*(a) \nu_m(s). \end{aligned}$$

Hence, by denoting, for any $a \in \mathcal{A}^\mathcal{T}$,

$$\theta_m^*(a) = \prod_{t \in \mathcal{T}^L} \theta_{t, m}^*(a_t) \quad \text{and} \quad W_m^0 = \sum_{(s, a) \in \mathcal{S} \times \mathcal{A}^\mathcal{T}} w(s, c) \nu_m(s) \theta_m^*(a),$$

it follows that, up to a subsequence,

$$W^\varepsilon(\tau^\varepsilon) = \sum_{m=1}^M W^\varepsilon(\nu_m^\varepsilon) p_m^\varepsilon \xrightarrow{\varepsilon \rightarrow 0} \sum_{m=1}^M W_m^0 p_m \leq \sum_{m=1}^M W(\nu_m) p_m = W(\tau^0),$$

which yields the desired result.

6.7 Proof of Lemma 5

Let τ^* be a solution of (19). Theorem 1 yields the existence of a set $\mathcal{R} \subset \mathcal{A}^\mathcal{T}$ such that

$$\tau^* = \sum_{a \in \mathcal{R}} p_a^* \delta_{\nu_a^*},$$

with ν_a^* such that $a_t = a_t^*(\nu_a^*)$, for any $t \in \mathcal{T}$. Since, by Assumption 1, for any $a \in \mathcal{R}$, there exists ν_a such that $\text{supp}(\nu_a) \subset \text{supp}(\mu)$ and $\mathcal{A}_t(\nu_a) = \{a_t\}$, for any $t \in \mathcal{T}$. Furthermore, since $\mathcal{S}$ is finite, there exists $r > 0$ such that

$$\bar{\nu} = \mu + r \left(\mu - \sum_{a \in \mathcal{R}} p_a^* \nu_a \right) \in \mathcal{P}(\mathcal{S}).$$

Let $\alpha = \varepsilon^{1/2}$ as well as, for any $a \in \mathcal{R}$,

$$\nu_a^\varepsilon = \alpha \nu_a + (1 - \alpha) \nu_a^*, \quad \beta = \frac{\alpha}{\alpha + r} \quad \text{and} \quad \tau^\varepsilon = (1 - \beta) \sum_{a \in \mathcal{R}} p_a^* \delta_{\nu_a^\varepsilon} + \beta \delta_{\bar{\nu}}.$$

We have

$$(1 - \beta) \sum_{a \in \mathcal{R}} p_a + \beta = 1,$$

so that $\tau^\varepsilon \in \mathbf{P}(\mathbf{P}(\mathcal{S}))$ and, furthermore, we observe that

$$\begin{aligned}
\int_{\mathbf{P}(\mathcal{S})} \nu d\tau^\varepsilon(\nu) &= (1 - \beta) \sum_{a \in \mathcal{R}} p_a^* \nu_a^\varepsilon + \beta \bar{\nu} \\
&= (1 - \beta) \sum_{a \in \mathcal{R}} p_a^* \nu_a^\varepsilon + \beta \left(\mu + r \left(\mu - \sum_{a \in \mathcal{R}} p_a^* \nu_a \right) \right) \\
&= \beta(1 + r)\mu + \sum_{a \in \mathcal{R}} p_a^* ((1 - \beta)\nu_a^\varepsilon - \beta r \nu_a) \\
&= \beta(1 + r)\mu + (1 - \beta)(1 - \alpha) \sum_{a \in \mathcal{R}} p_a^* \nu_a^* \\
&= (\beta(1 + r) + (1 - \beta)(1 - \alpha)) \mu = \mu,
\end{aligned}$$

which yields that $\tau^\varepsilon \in \mathbf{T}_\mu$. Since $\mathcal{A}_t(\nu_a) = \{a_t\}$, we have, for any $\ell \in \{1, \dots, L\}$ and $t_\ell \in \mathcal{T}_\ell$, that $\mathcal{A}_{\ell, t_\ell}(\nu_a) = \{(a_t)_\ell\}$. In particular, for any $\bar{a} \in \mathcal{A}_\ell$ such that $\bar{a} \neq (a_t)_\ell$, there exists $\iota > 0$ such that

$$\sum_{s \in \mathcal{S}} u_{t_\ell}(s, \bar{a}) \nu_a(s) \leq \sum_{s \in \mathcal{S}} u_{t_\ell}(s, a_{t, \ell}) \nu_a(s) - \iota,$$

where we denoted $a_{t, \ell} = (a_t)_\ell$, which yields, for any $\theta \in \mathbf{P}(\mathcal{A}_\ell)$,

$$\sum_{(s, \bar{a}) \in \mathcal{S} \times \mathcal{A}_\ell} u_{t_\ell}(s, \bar{a}) \nu_a(s) \theta(\bar{a}) \leq \sum_{s \in \mathcal{S}} u_{t_\ell}(s, a_{t, \ell}) \nu_a(s) - \iota(1 - \theta(a_{t, \ell})).$$

Moreover, since $a_t \in \mathcal{A}_t(\nu_a^*)$, we also have

$$\sum_{(s, \bar{a}) \in \mathcal{S} \times \mathcal{A}_\ell} u_{t_\ell}(s, \bar{a}) \nu_a^*(s) \theta(\bar{a}) \leq \sum_{s \in \mathcal{S}} u_{t_\ell}(s, a_{t, \ell}) \nu_a^*(s).$$

Combining these two inequalities, we obtain, for any $\lambda_\ell \in \mathbf{P}(\mathcal{A}_\ell)$,

$$\begin{aligned}
&\sum_{(s, \bar{a}) \in \mathcal{S} \times \mathcal{A}_\ell} u_{t_\ell}(s, \bar{a}) \nu_a^\varepsilon(s) \theta(\bar{a}) - \varepsilon \vartheta(\theta, \lambda_\ell) \\
&\leq \sum_{s \in \mathcal{S}} u_{t_\ell}(s, a_{t, \ell}) \nu_a^\varepsilon(s) - \alpha \iota (1 - \theta(a_{t, \ell})) - \varepsilon \vartheta(\theta, \lambda_\ell).
\end{aligned}$$

It follows that

$$\begin{aligned}
&\sum_{s \in \mathcal{S}} u_{t_\ell}(s, a_{t, \ell}) \nu_a^\varepsilon(s) - \varepsilon \vartheta(\delta_{a_{t, \ell}}, \lambda_\ell) \\
&\leq \sum_{s \in \mathcal{S}} u_{t_\ell}(s, a_{t, \ell}) \nu_a^\varepsilon(s) - \alpha \iota (1 - \theta_{t_\ell, \nu_a^\varepsilon}^{*, \varepsilon}(a_{t, \ell})) - \varepsilon \vartheta(\theta_{t_\ell, \nu_a^\varepsilon}^{*, \varepsilon}, \lambda_\ell),
\end{aligned}$$

which yields

$$0 \leq -\iota(1 - \theta_{t_\ell, \nu_a^\varepsilon}^{*, \varepsilon}(a_{t, \ell})) - \varepsilon^{1/2}(\vartheta(\theta_{t_\ell, \nu_a^\varepsilon}^{*, \varepsilon}, \lambda_\ell) - \vartheta(\delta_{a_{t, \ell}}, \lambda_\ell)).$$

We know that, up to a subsequence, $(\theta_{t_\ell, \nu_a^\varepsilon}^{*, \varepsilon})_{\varepsilon > 0}$ converges in $\mathcal{P}(\mathcal{A}_\ell)$ as $\varepsilon \rightarrow 0$. Assume that the limit is not $\delta_{a_{t, \ell}}$. Then, letting $\varepsilon \rightarrow 0$ in the previous inequality leads to $0 \leq -\iota$, which is impossible. We conclude that

$$\theta_{t_\ell, \nu_a^\varepsilon}^{*, \varepsilon} \xrightarrow{\varepsilon \rightarrow 0} \delta_{a_{t, \ell}} \quad \text{and} \quad \theta_{\nu_a^\varepsilon}^* \xrightarrow{\varepsilon \rightarrow 0} \delta_a.$$

In the end, we obtain that

$$\begin{aligned} \int_{\mathcal{P}(S)} W^\varepsilon(\nu) d\tau^\varepsilon(\nu) &= (1 - \beta) \sum_{a \in \mathcal{R}} p_a^* \sum_{(s, \bar{a}) \in \mathcal{S} \times \mathcal{A}^\tau} w(s, \bar{a}) \theta_{\nu_a^\varepsilon}^{*, \varepsilon}(\bar{a}) \nu_a^\varepsilon(s) \\ &\quad + \beta \sum_{(s, \bar{a}) \in \mathcal{S} \times \mathcal{A}^\tau} w(s, \bar{a}) \theta_{\bar{\nu}}^{*, \varepsilon}(\bar{a}) \bar{\nu}(s) \\ &\xrightarrow{\varepsilon \rightarrow 0} \sum_{a \in \mathcal{R}} p_a^* w(s, a) \nu_a^*(s) = \int_{\mathcal{P}(S)} W(\nu) d\tau^*(\nu), \end{aligned}$$

since $\nu_a^\varepsilon \rightarrow \nu_a$ and $\beta \rightarrow 0$ as $\varepsilon \rightarrow 0$.

6.8 Proof of Corollary 3

By Lemma 5, we know that there exists a sequence $(\tau^\varepsilon)_{\varepsilon > 0} \subset \mathcal{T}_\mu$ such that

$$\bar{W}^* = \liminf_{\varepsilon \rightarrow 0} \int_{\mathcal{P}(S)} W^\varepsilon(\nu) d\tau^\varepsilon(\nu) \leq \liminf_{\varepsilon \rightarrow 0} \bar{W}^{*, \varepsilon} \leq \limsup_{\varepsilon \rightarrow 0} \bar{W}^{*, \varepsilon}.$$

Thus, by Corollary 2, this yields

$$\bar{W}^* = \lim_{\varepsilon \rightarrow 0} \bar{W}^{*, \varepsilon}.$$

6.9 Proof of Theorem 2

Let $(\tau^\varepsilon)_{\varepsilon > 0}$ be a sequence of maximizers of (19) belonging in $\mathcal{T}_{M, \mu}$, for some $M \in \mathbb{N}$. By Lemma 1, we know that this sequence converges, up to a subsequence, to an element $\tau \in \mathcal{T}_{M, \mu}$. Furthermore, by Corollary 3 and Lemma 4, we have

$$\bar{W}^* = \lim_{\varepsilon \rightarrow 0} \bar{W}^{*, \varepsilon} = \limsup_{\varepsilon \rightarrow 0} \int_{\mathcal{P}(S)} W^\varepsilon(\nu) d\tau^\varepsilon(\nu) \leq \int_{\mathcal{P}(S)} W(\nu) d\tau(\nu) \leq \bar{W}^*.$$

It follows that $\bar{W}^* = \int_{\mathcal{P}(S)} W(\nu) d\tau(\nu)$ and, thus, τ is a maximizer of (13).

References

- [1] Ricardo Alonso and Odilon Câmara. Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory*, 165:672–706, 2016.
- [2] Ricardo Alonso and Odilon Câmara. Persuading voters. *American Economic Review*, 106(11):3590–3605, 2016.

- [3] Ned Augenblick and Matthew Rabin. Belief movement, uncertainty reduction, and rational updating. *The Quarterly Journal of Economics*, 136(2):933–985, 2021.
- [4] Yakov Babichenko and Siddharth Barman. Algorithmic aspects of private bayesian persuasion. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*, pages 34–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2017.
- [5] Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.
- [6] Patrick Billingsley. *Convergence of probability measures*. John Wiley & Sons, 2013.
- [7] Alexander W Bloedel and Ilya R Segal. Persuasion with rational inattention. *Available at SSRN 3164033*, 2018.
- [8] Claire Boyer, Antonin Chambolle, Yohann De Castro, Vincent Duval, Frédéric De Gournay, and Pierre Weiss. On representer theorems and convex regularization. *SIAM Journal on Optimization*, 29(2):1260–1281, 2019.
- [9] Geoffroy de Clippel and Xu Zhang. Non-bayesian persuasion. *Journal of Political Economy*, 130(10):2594–2642, 2022.
- [10] Shaddin Dughmi. Algorithmic information structure design: a survey. *ACM SIGecom Exchanges*, 15(2):2–24, 2017.
- [11] Shaddin Dughmi. On the hardness of designing public signals. *Games and Economic Behavior*, 118:609–625, 2019.
- [12] Shaddin Dughmi and Haifeng Xu. Algorithmic bayesian persuasion. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 412–425, 2016.
- [13] Shaddin Dughmi and Haifeng Xu. Algorithmic persuasion with no externalities. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 351–368, 2017.
- [14] Larry G Epstein. An axiomatic model of non-bayesian updating. *The Review of Economic Studies*, 73(2):413–436, 2006.
- [15] Larry G Epstein, Jawwad Noor, and Alvaro Sandroni. Non-bayesian updating: a theoretical framework. 2008.
- [16] Simone Galperti. Persuasion: The art of changing worldviews. *American Economic Review*, 109(3):996–1031, 2019.

- [17] Ronen Gradwohl, Niklas Hahn, Martin Hoefer, and Rann Smorodinsky. Algorithms for persuasion with limited communication. *Mathematics of Operations Research*, 47(3):2520–2545, 2022.
- [18] Emir Kamenica. Bayesian persuasion and information design. *Annual Review of Economics*, 11(1):249–272, 2019.
- [19] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- [20] Yong-Ju Lee, Wooyoung Lim, and Chen Zhao. Cheap talk with prior-biased inferences. *Games and Economic Behavior*, 138:254–280, 2023.
- [21] Elliot Lipnowski, Laurent Mathevet, and Dong Wei. Attention management. *American Economic Review: Insights*, 2(1):17–32, 2020.
- [22] Ludmila Matyskova and Alfonso Montes. Bayesian persuasion with costly information acquisition. *Journal of Economic Theory*, 211:105678, 2023.
- [23] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.
- [24] Dong Wei. Persuasion under costly learning. *Journal of Mathematical Economics*, 94:102451, 2021.
- [25] Chen Zhao. Pseudo-bayesian updating. *Theoretical Economics*, 17(1):253–289, 2022.