Bayesian Persuasion under Ex Ante and Ex Post Constraints*

Yakov Babichenko, Inbal Talgam-Cohen, Konstantin Zabarnyi
Technion – Israel Institute of Technology
yakovbab@technion.ac.il, italgam@cs.technion.ac.il, konstzab@gmail.com

Abstract

Bayesian persuasion, as introduced by Kamenica and Gentzkow in 2011, is the study of information sharing policies among strategic agents. A prime example is signaling in online ad auctions: what information should a platform signal to an advertiser regarding a user when selling the opportunity to advertise to her? Practical considerations such as preventing discrimination, protecting privacy or acknowledging limited attention of the information receiver impose constraints on information sharing. We propose a simple way to mathematically model such constraints as restrictions on Receiver's admissible posterior beliefs. We consider two families of constraints – ex ante and ex post; the latter limits each instance of Sender-Receiver communication, while the former more general family can also pose restrictions in expectation. For the ex ante family, a result of Doval and Skreta (2018) establishes the existence of an optimal signaling scheme with a small number of signals – at most the number of constraints plus the number of states of nature – and we show this result is tight. For the ex post family, we tighten the previous bound of Vølund (2018), showing that the required number of signals is at most the number of states of nature, as in the original Kamenica-Gentzkow setting. As our main algorithmic result, we provide an additive bi-criteria FPTAS for an optimal constrained signaling scheme assuming a constant number of states of nature; we improve the approximation to single-criteria under a Slater-like regularity condition. The FPTAS holds under standard assumptions, and more relaxed assumptions yield a PTAS. We then establish a bound on the ratio between Sender's optimal utility under convex ex ante constraints and the corresponding ex post constraints. We demonstrate how this result can be applied to find an approximately welfare-maximizing constrained signaling scheme in ad auctions.

1 Introduction

In many real-life situations, one entity relies on information revealed by another entity to decide which action to take. Call the former and the latter entities *Receiver* and *Sender*, respectively. Sender has the power to commit to a revelation policy, a.k.a. a *signaling scheme*. Sender would like to strategically design such a scheme to *persuade* Receiver to act in Sender's interest. Mathematically, a signaling scheme transforms Receiver's prior belief about how some unknown *state of nature* is distributed into a posterior belief, which determines Receiver's action.

Since strategic communication of information is intrinsic to most human endeavours, persuasion is of high importance in practice, and is becoming even more so in today's digital economy. Indeed, persuasion has been estimated to account for at least 30% of the total US economy (McCloskey and Klamer 1995; Antioch 2013). Persuasion has also attracted significant research interest in recent years, initiated by the celebrated Bayesian persuasion model of Kamenica and Gentzkow.

We study a theoretical model for constrained Bayesian persuasion under general families of *ex ante* and *ex post* constraints. Ex ante constraints are statistical limitations on the amount of information Receiver may learn when the Sender-Receiver communication is repeated over time; ex post constraints are a strong particular case restricting the information passage on *every* instance of the communication. These constraint families have various applications. In particular, Tsakas and Tsakas (2019) model signaling via noisy channels by ex ante-constrained persuasion. Doval and Skreta (2018) further show that optimal signaling via a capacity-constrained channel is equivalent to a constrained persuasion setting with a single entropy ex ante constraint. Vølund (2018), based on research in cognitive science, suggests ex post constraints as a possible model for human behaviour upon receiving an unwelcome signal. One of our main motivating examples in this work is online ad auctions in which ex ante constraints reduce discrimination and ex post constraints protect user privacy.

Our contribution and paper organization. Let m and k be the numbers of constraints and states of nature, respectively. Section 2 formally defines the model and describes our main motivations. Section 3 proves tightness of the $k+m$ bound of Doval and Skreta (2018) on the support size of an optimal ex ante-constrained signaling scheme; it further extends the lower bound result of Le Treust and Tomala (2019) beyond a single constraint. For ex post constraints, Section 3 shows a tight bound of k on the support size, which is the same as in the original setting of Kamenica and Gentzkow. The support size of a signaling scheme is a standard mea-

*For omitted proofs, please see our full paper version, which is available at arXiv:2012.03272.

sure of its complexity, similar to menu-size complexity in auctions (Hart and Nisan 2013, 2017). Section 4 provides an additive bi-criteria FPTAS for an optimal signaling scheme when k is constant and improves it to single-criteria under a Slater-like regularity condition. Section 5 shows that for a constant m, convex constraints and a wide family of natural objective functions, ex ante constraints outperform ex post constraints by a constant multiplicative factor. We conclude by applications to ad auctions with an exponentially large state of nature space, using the setting of Badanidiyuru, Bhawalkar, and Xu (2018).

Ex ante constraints raise technical challenges not usually encountered in the literature on persuasion. In our model, we cannot restrict attention to *straightforward policies* (Kamenica and Gentzkow 2011) in which Sender recommends an action to Receiver in an incentive-compatible way. These policies are a very central tool in persuasion problems and are widely applied across the literature (see, e.g., Dughmi 2017), but they are not descriptive enough for determining whether a given ex ante constraint is satisfied. In particular, an optimal signaling scheme in our model cannot be described by a finite linear program (LP). Note that we do not assume Receiver's action space is finite, but even such a simplifying assumption would not have resolved these issues.

Related work. The seminal work of Kamenica and Gentzkow (2011) introduces Bayesian persuasion and characterizes Sender's optimal signaling scheme using the *concavification* approach. Among the works on algorithmic aspects of persuasion we mention a negative result of Dughmi and Xu (2017), which is relevant to hardness of approximating Sender's optimal utility; see (Dughmi 2017) for a comprehensive survey of computational results.

In the context of auctions, an early work on signaling information is the classic paper of Milgrom and Weber (1982). Emek et al. (2014); Miltersen and Sheffet (2012) apply a computational approach to signaling in auctions; Fu et al. (2012) study signaling in the revenue-maximizing Myerson auction (Myerson 1981); Badanidiyuru, Bhawalkar, and Xu (2018) study this in the welfare-maximizing second price auction with exponentially many states of nature; and Daskalakis, Papadimitriou, and Tzamos (2016) design the signaling and auction mechanisms simultaneously.

The most closely related works to our own are the following: (a) Our algorithmic approach in Section 4 is related to that of Cheng et al. (2015), as both use discretization and linear programming to achieve an additive FPTAS. (b) The papers of Dughmi, Immorlica, and Roth (2014); Dughmi et al. (2015) study constrained persuasion, but their constraints are on the complexity of the Sender-Receiver communication (as measured by message length or number of signaled features), and so are fundamentally different from ours.[1] Ichihashi (2019) considers persuasion by Sender who is constrained in the information she can acquire (and therefore, send) and characterizes the set of possible equilibrium outcomes. Our Theorem 3.3 is related to this literature in that it indicates that ex post constraints on persuasion do not cause a blowup in the number of signals needed to persuade optimally. (c) Inspired by (Le Treust and Tomala 2019), Doval and Skreta (2018) prove an upper bound on the required number of signals in an ex ante-constrained optimal scheme; we prove that this bound is tight and provide an analogous tight bound for ex post constraints in Section 3. (d) Vølund (2018) studies a model of persuasion on compact subsets, which is equivalent to our ex post constraints; there is no parallel in that work to ex ante constraints, and the results on ex post in the two works do not overlap.

In Section 2, we discuss motivating applications of ex ante and ex post constraints, including limited attention, as well as privacy protection in online ad auctions. Bloedel and Segal (2018); Lipnowski, Mathevet, and Wei (2020) study persuasion with limited attention – see Section 2 for details. Eilat, Eliaz, and Mu (2019) study ex ante and ex post privacy constraints in the design of auctions rather than persuasion schemes. Ichihashi (2020) studies the economic implications of online consumer privacy; in his model, the consumer rather than the seller plays the role of Sender. It is important to note that the differential privacy paradigm (see Dwork and Roth 2014) does *not* apply to privacy protection in online ad auctions: the state of nature about which information is revealed represents characteristics of an individual rather than statistics of a large population, and it is inherent to ad personalization that these characteristics influence the outcome in a non-negligible way.

2 Our Model

Standard preliminaries. We consider Bayesian persuasion with a single *Sender* and a single *Receiver*, as introduced by Kamenica and Gentzkow (2011). Fix a space of k *states of nature* Ω and a commonly-known *prior distribution* p on them. Take a compact non-empty set A to be Receiver's *action space*. Introduce two random variables ω and x, representing the *state of nature* and Receiver's *action*, respectively. Fix a *Sender's utility function* $\tilde{u}_s : A \times \Omega \to \mathbb{R}_{\geq 0}$ and a *Receiver's utility function* $u_r : A \times \Omega \to \mathbb{R}_{\geq 0}$. The Sender-Receiver communication is specified by a *signaling scheme* Σ, a.k.a. a *signaling policy*, which is a random function – that is, a stochastic map – from Ω to some set of signals (this notion will be formalized soon). As usual, Sender must commit to Σ before learning ω.

Denote by $\Delta(\Omega)$ the set of probability distributions over Ω. Consider it to be a subset of $[0, 1]^k$, with i-th coordinate being the probability assigned to the i-th element of Ω.

Let σ be the actual *signal realization*. Note that σ induces an updated distribution on Ω in Receiver's view, called the *posterior distribution* or the *posterior*. Let $p_\sigma \in \Delta(\Omega)$ be the posterior induced by σ. The *support* of Σ, $\mathrm{supp}(\Sigma)$, is the intersection of all the closed sets $S \subseteq \Delta(\Omega)$ s.t. $\Pr_\Sigma[p_\sigma \in S] = 1$. If Σ uses only countably many signals, then $\mathrm{supp}(\Sigma)$ is the set of all the posteriors induced by signal realizations of Σ with a positive probability.

Formally, Σ is an unconditional on the state of nature distribution over the elements of $\Delta(\Omega)$ that belong to $\mathrm{supp}(\Sigma)$. For any $\omega_0 \in \Omega$, assuming $\omega = \omega_0$, Σ induces a conditional

[1] They also study a version called "bipartite signaling", which has a combinatorial flavour different than ours, in an auction setting with the strong assumption that bidder values are known.

distribution over the elements of $\mathrm{supp}(\Sigma)$ that specifies how Sender chooses the signal realization when $\omega = \omega_0$. Denote this distribution by $\Sigma(\omega_0)$. Note that given p and Σ, it can be computed by Bayes' law.

For simplicity, we introduce the following notation for the expectation of a function of the posterior over the elements of $\mathrm{supp}(\Sigma)$ according to Σ:

Notation 2.1. *For a function $f : \Delta(\Omega) \to \mathbb{R}$:*

$$E[\Sigma, f] := \mathbb{E}_{p_\sigma \sim \Sigma}[f(p_\sigma)] = \mathbb{E}_{\omega \sim p, p_\sigma \sim \Sigma(\omega)}[f(p_\sigma)].$$

By (Blackwell 1953; Aumann, Maschler, and Stearns 1995), a distribution Σ represents a signaling scheme if and only if Σ is *Bayes-plausible*. That is:

$$\forall \omega_0 \in \Omega : \quad p[\omega_0] = E[\Sigma, p_\sigma[\omega_0]].$$

The persuasion process runs as follows: (1) Sender commits to a signaling policy Σ. (2) Sender discovers the state of nature ω. (3) Sender transmits a signal realization σ to Receiver, according to $\Sigma(\omega)$. (4) Receiver chooses an action $x \in A$ s.t. $x \in \mathrm{argmax}_{x' \in A}(\mathbb{E}_{\omega' \sim p_\sigma}[u_r(x', \omega')])$; assume, as is standard, that ties are broken in Sender's favour. (5) Sender gets utility of $\tilde{u}_s(x, \omega)$, while Receiver gets utility of $u_r(x, \omega)$.

Since x depends only on p_σ, there exists $\bar{u}_s : \Delta(\Omega) \times \Omega \to \mathbb{R}_{\geq 0}$ s.t. $\tilde{u}_s(x, \omega) \equiv \bar{u}_s(p_\sigma, \omega)$. Define $u_s : \Delta(\Omega) \to \mathbb{R}_{\geq 0}$ by:

$$u_s(p_\sigma) := \mathbb{E}_{\omega' \sim p_\sigma}[\bar{u}_s(p_\sigma, \omega')] =$$
$$\sum_{\omega' \in \Omega} p_\sigma[\omega = \omega'] \cdot \tilde{u}_s(\mathrm{argmax}_{x' \in A}(\mathbb{E}_{\omega'' \sim p_\sigma}[u_r(x', \omega'')]), \omega').$$

Remark 2.2. *From now on we shall consider u_s instead of $\tilde{u}_s$ or $\bar{u}_s$, thus assuming that Sender's utility is state of nature-independent. This is w.l.o.g. for our results from Sections 3-4, since the passage from $\bar{u}_s$ to u_s preserves the conditions required there (upper semi-continuity or efficient approximability by a piecewise constant function).*[2] *While one cannot apply the results of Section 5 to the state-dependent case without strengthening Assumption 5.2, the natural applications to ad auctions discussed there have state-independent Sender's utility.*

Throughout we make the following assumption, which is a relaxation of the standard continuity assumption on u_s. In particular, u_s can be any continuous or threshold function.

Assumption 2.3. *The function u_s is upper semi-continuous.*

Ex ante and ex post constraints. So far we have described the setting of Kamenica and Gentzkow (2011). However, in our model we do not allow Sender to choose among all Bayes-plausible signaling schemes, but only among schemes that satisfy certain restrictions (see the next subsection for motivation). We define two general families of constraints: *ex ante* and *ex post*. A constraint of the latter type restricts the admissible values of a certain function of p_σ for every possible p_σ, while a constraint of the former type restricts only the expectation of such a function.

[2]In the state-dependent setting, $\bar{u}_s(\cdot, \omega_0)$ has to satisfy the theorem requirements from u_s for every $\omega_0 \in \Omega$.

Definition 2.4 (Ex ante constraints)**.** *An* ex ante constraint *on a signaling scheme Σ is a constraint of the form:*

$$E[\Sigma, f] \leq c$$

for continuous $f : \Delta(\Omega) \to \mathbb{R}$ and a constant $c \in \mathbb{R}$.

Definition 2.5 (Ex post constraints)**.** *An* ex post constraint *on a signaling scheme Σ is a constraint of the form:*

$$\forall p_\sigma \in \mathrm{supp}(\Sigma) : f(p_\sigma) \leq c$$

for continuous $f : \Delta(\Omega) \to \mathbb{R}$ and a constant $c \in \mathbb{R}$.

Definition 2.6. *A constraint defined as in either of the previous two definitions is said to be* specified *by the function f and the constant c.*

Definition 2.7. *A constraint specified by a convex f and some c is called* convex.

Observation 2.8. *Ex post constraints are a special case of ex ante constraints.*

Indeed, an ex post constraint specified by some f and c is equivalent to the ex ante constraint specified by $\max\{f, c\}$ and c. Note that if f is convex, then so is $\max\{f, c\}$.

Every ex ante constraint can be transformed into a (stronger) ex post constraint by "erasing the expectation" and vice versa. Formally:

Definition 2.9. *An ex post and an ex ante constraint* correspond *to each other if they are specified by the same function and the same constant.*

Definition 2.10. *Given a set of constraints, a signaling scheme satisfying all of them is called* valid.

Definition 2.11. *A set of constraints satisfied by every signaling scheme is called* trivial.

Motivation for constrained persuasion. In many applications of Kamenica and Gentzkow's model, Sender cannot reveal as much information as would theoretically be optimal due to constraints. Such constraints can originate from sources including law, professional integrity, political agreements, public opinion and limited attention.

As a first motivating example, consider online ad auctions. The auctioneer – an advertising platform – is Sender, while the set of bidders – which are advertisers – is Receiver.[3] The profile of the web user who will view the ad is the state of nature. This profile is known to the auctioneer, but not to the bidders; every signal reveals information about it. Such information revelation should be restricted by both privacy and fairness considerations.

The constraint families we introduce fit for protecting privacy – following Eilat, Eliaz, and Mu (2019), privacy protection can be modeled as imposing a threshold on the Kullback–Leibler (KL) divergence from the prior to the posterior. The KL divergence quantifies how much more informative the posterior is compared to the prior due to extra information about the user provided by the signal realization. An ex post constraint on the KL divergence provides a relatively

[3]We treat the bidders as a single Receiver since they all get the same signal; private signaling poses additional challenges (Arieli and Babichenko 2019) and is left for future work.

robust protection of individual privacy by ruling out sending a very informative signal even with only a small probability. In contrast, the corresponding ex ante constraint protects privacy on the group level – e.g., it limits Receiver's ability to learn the shopping habits of certain population groups, since the posterior is close, on average, to the prior.

Another important restriction on signaling in ad auctions is fairness, or anti-discrimination – e.g., ensuring that enough women compared to men are shown an ad for a high-paying job (Celis et al. 2019; Celis, Mehrotra, and Vishnoi 2019). Consider, for simplicity, a uniform prior over population groups. A simple constraint specified by $(-\min_{\omega' \in \Omega}\{p_\sigma[\omega']\})$ lower-bounds the frequency of a population group in the posterior, thus ensuring its proportional inclusion.[4] An ex ante constraint of this form ensures that on average, the advertiser does not get enough information to discriminate against particular groups.

A second motivating example involves constraints arising from Receiver's limited attention span. As Simon (1996) noted, "a wealth of information creates a poverty of attention". Our model enables limiting the signaled information so that it "fits" within Receiver's limited attention.[5] Following the *rational inattention* literature (Sims 2003), define the attention required from Receiver to process Sender's signal σ as the entropy of the posterior p_σ.[6] By constraining the entropy – either in expectation (i.e., ex ante) or of every posterior (i.e., ex post) – we enable Receiver to process the signal despite her limited attention (where the limit is either in expectation or per signal, respectively). A concrete application from Bloedel and Segal (2018) includes a busy executive as Receiver, one of her advisors as Sender and constraints on the signaled information enforced by keeping meetings and briefings short (on average or per meeting).

3 Existence Results

Doval and Skreta (2018) prove that for every set of m ex ante constraints, there exists an optimal valid signaling scheme with support size of at most $k + m$:

Fact 3.1 (Doval and Skreta (2018) – existence of an optimal valid signaling scheme under ex ante constraints with a linear-sized support)**.** *Fix m ex ante constraints. Then either there exists an optimal valid signaling scheme with support size at most $k + m$ or the set of valid signaling schemes is empty.*

In our full paper version, we show that this bound is tight:

Proposition 3.2. *The bound on the support size from Fact 3.1 is tight for every k and m.*

We further show in the full version that for any number m of ex post constraints, a stronger bound of k holds, just as in the unconstrained setting of Kamenica and Gentzkow; the same proof outline yields an alternative proof to the result of (Doval and Skreta 2018) on ex ante constraints.

Theorem 3.3 (Existence of an optimal valid signaling scheme under ex post constraints with a linear-sized support)**.** *Fix a set of ex post constraints. Then either there exists an optimal valid signaling scheme with support size at most k or the set of valid signaling schemes is empty.*

Proof Idea of Theorem 3.3. At a high level, we translate the problem into an infinite LP, with the "variables" being the distribution Σ over $\Delta(\Omega)$. We first prove that the target function of the infinite LP is upper semi-continuous. Secondly, we show, using infinite-dimensional optimization tools, that it must attain a maximum at an extreme point of the feasible set. Specifically, we use Bauer's maximum principle – see, e.g., Theorem 7.69 from (Charalambos and Border 2006). Thirdly, we argue that every extreme point has a finite support of bounded size, analyzing the effect of adding the Bayes-plausibility constraints one by one by considering the hyperplanes specifying the constraints: the maximal support size of extreme points is at most doubled upon each addition. Finally, we improve the bound on the support size of each extreme point using a finite LP. $\square$

Observation 3.4. *The bound from Theorem 3.3 is achieved by $u_s(p_\sigma) := ||p_\sigma||_\infty$ and a set of trivial ex post constraints.*

4 Computational Aspects

In this section, we provide positive computational results for a constant number of states of nature k. We focus on constant k since a hardness result of Dughmi and Xu (2017) implies that unless $P = NP$, there is no additive PTAS or constant-factor multiplicative approximation of the optimal Sender's utility in $\text{poly}(k)$-time, even for piecewise constant u_s.[7]

Call *L-Lipschitz* a function with Lipschitz constant at most L. We first present an additive bi-criteria FPTAS for ex ante-constrained persuasion problems (Theorem 4.4). Specifically, we show how to compute in polynomial time a signaling policy that achieves a utility that is additively at most ϵ-far from optimal and violates each constraint by at most ϵ; Bayes-plausibility is satisfied precisely. To achieve this, we impose regularity assumptions on both u_s and the constraints, requiring them to be efficiently approximable, respectively, by piecewise constant functions and Lipschitz functions (Assumptions 4.2-4.3). In particular, an $O(1)$-Lipschitz Sender's utility function under ex ante constraints specified by KL divergence, entropy or norms (including total variation distance of probability measures) satisfies our conditions. Then we show that under a further Slater-like regularity condition, one obtains a single-criteria additive

[4] If the prior over population groups is not uniform, then we can easily add weights to this constraint: $-\min_{\omega' \in \Omega}\{b_{\omega'} p_\sigma[\omega']\}$.

[5] An alternative model of Bloedel and Segal (2018); Lipnowski, Mathevet, and Wei (2020) allows Sender to "flood" Receiver with information, but Receiver chooses what to pay attention to. Constrained persuasion simply avoids flooding Receiver with information in expectation (the ex ante model) or always (ex post).

[6] Bloedel and Segal (2018) use mutual information of p_σ and Receiver's perception of it after paying limited attention as the measure of the attention invested by Receiver. In our model, Receiver always pays full attention, thus the mutual information coincides with the entropy of p_σ.

[7] Their result is on public persuasion with multiple Receivers, which can be replaced by a single Receiver with a large action space.

FPTAS (Theorem 4.6). By Observation 2.8, these results hold also for ex post constraints. Throughout this section, we assume that both u_s and the functions specifying the constraints are given by explicit formulae and can be evaluated at every point in constant time.

In the full paper version, we show that if one drops all the assumptions on the ex ante constraints and just assumes that u_s is either continuous or piecewise constant – with a constant number of pieces s.t. each piece covers a convex polygon in $\Delta(\Omega)$ with a constant number of vertices – then the algorithms from Theorems 4.4, 4.6 still provide an additive PTAS (rather than FPTAS).

Remark 4.1. *Our approximation algorithms output a solution of a finite LP with $k+m$ constraints; therefore, their output – which is w.l.o.g. a basic feasible solution – is a signaling scheme with support size at most $k+m$, which matches the tight theoretical bound from Fact 3.1.*

Bi-criteria FPTAS. We introduce the following assumptions on u_s and the constraints.

Assumption 4.2. *For every $\epsilon > 0$ and every $M = \mathrm{poly}\left(\frac{1}{\epsilon}\right)$, one can compute in $\mathrm{poly}\left(\frac{1}{\epsilon}\right)$-time an explicit formula for an upper semi-continuous piecewise constant $u_{\epsilon,M} : \Delta(\Omega) \to \mathbb{R}_{\geq 0}$, with a $\mathrm{poly}\left(\frac{1}{\epsilon}\right)$-number of pieces, s.t.:*

- *Every piece of $u_{\epsilon,M}$ covers a region of $\Delta(\Omega)$ that is a convex polygon with a $\mathrm{poly}\left(\frac{1}{\epsilon}\right)$-number of vertices and diameter at most $\frac{\epsilon}{M}$.*
- $\forall q \in \Delta(\Omega):\ 0 \leq u_{\epsilon,M}(q) - u_s(q) \leq \epsilon$.

Assumption 4.3. *The i-th ex ante constraint ($1 \leq i \leq m$) is specified by $f_i : \Delta(\Omega) \to \mathbb{R}$ s.t. for every $\epsilon > 0$, one can compute in $\mathrm{poly}\left(\frac{1}{\epsilon}\right)$-time an explicit formula for a $\mathrm{poly}\left(\frac{1}{\epsilon}\right)$-Lipschitz function $g_i : \Delta(\Omega) \to \mathbb{R}$ s.t. for every $q \in \Delta(\Omega)$:*

$$0 \leq f_i(q) - g_i(q) \leq \epsilon.$$

Note that Assumptions 4.2-4.3 are not restrictive. Utility functions naturally arising in applications are either $O(1)$-Lipschitz (if Receiver has a continuum of actions) or piecewise constant (if Receiver has finitely many actions), thus they satisfy Assumption 4.2 (see full paper version). The constraints are not necessarily Lipschitz, but Assumption 4.3 requires only approximability by a Lipschitz function, which is satisfied by constraints including KL divergence, entropy and norms.

Theorem 4.4 (An additive bi-criteria FPTAS for an optimal valid signaling scheme)**.** *Suppose that k is constant, u_s satisfies Assumption 4.2 and we have m ex ante constraints satisfying Assumption 4.3. Then either the set of valid signaling schemes is empty or for every $\epsilon > 0$, there exists a $\mathrm{poly}\left(m, \frac{1}{\epsilon}\right)$-time algorithm that computes an additively ϵ-optimal signaling scheme that violates each ex ante constraint at most by ϵ.*

We first strengthen Assumption 4.3 and assume that the constraints are specified by $\mathrm{poly}\left(\frac{1}{\epsilon}\right)$-Lipschitz functions. Then we restrict ourselves to the grid of the extreme points of the pieces of $u_{\epsilon,M}$, where M is the maximal Lipschitz constant of the functions specifying the constraints, and output the resultant optimal valid signaling scheme for $u_{\epsilon,M}$ rather than u_s. Finally, we bound the loss in Sender's utility and the constraint values by the approximability guarantees.

Proof of Theorem 4.4. We strengthen Assumption 4.3 to the following: the i-th ex ante constraint ($1 \leq i \leq m$) is specified by a $\mathrm{poly}\left(\frac{1}{\epsilon}\right)$-Lipschitz function $f_i : \Delta(\Omega) \to \mathbb{R}$ and some constant c_i. The original theorem follows from applying the theorem under the strengthened Assumption 4.3 with ϵ replaced by $\frac{\epsilon}{2}$ and the f_is replaced by the g_is. This is because the original Assumption 4.3 ensures that upon replacing f_i with g_i, every valid signaling scheme remains such and $E[\Sigma, f_i]$ decreases at most by ϵ.

Now we prove the theorem under the strengthened Assumption 4.3. Suppose that a valid signaling scheme exists and let OPT be Sender's expected utility under an optimal valid scheme. Fix $\epsilon > 0$ and let M be the maximal Lipschitz constant among the f_is. Compute an explicit formula for $u_{\epsilon,M}$. Denote by l the number of pieces in $u_{\epsilon,M}$. Let $q_1, ..., q_n$ be the extreme points of the regions of $\Delta(\Omega)$ covered by the pieces of $u_{\epsilon,M}$. Let us solve the following, assuming $\mathrm{supp}(\Sigma) \subseteq \{q_1, ..., q_n\}$:

$$\begin{aligned} \max \quad & E[\Sigma, u_{\epsilon,M}] \\ \text{s.t.} \quad & p[\omega_0] = E[\Sigma, p_\sigma[\omega_0]] \ \ \forall \omega_0 \in \Omega \\ & E[\Sigma, f_i] \leq c_i + \epsilon \ \ \forall 1 \leq i \leq m \end{aligned}$$

The problem defines a finite LP with n variables (representing the probability masses assigned by Σ to the q_is) and $k+m$ constraints (note that we should add a constraint for the probability masses in Σ to sum up to 1, but then we could remove one of the Bayes-plausibility constraints, as it would follow from the other constraints); this LP can be solved in time $\mathrm{poly}(n, k+m) = \mathrm{poly}\left(\frac{1}{\epsilon}, m\right)$. We return its solution Σ as the desired signaling scheme.

By the design of our LP, Σ is Bayes-plausible and violates each ex ante constraint at most by ϵ. Take now a valid optimal signaling scheme Σ_{OPT} (for Sender's utility function u_s rather than $u_{\epsilon,M}$). For every $1 \leq j \leq l$, move all the probability weight in Σ from the region covered by the j-th piece of $u_{\epsilon,M}$ to the extreme points of that region in an expectation-preserving way (so Bayes-plausibility still holds) and denote the resultant signaling scheme by Σ'_{OPT}. Since the diameter of every such region is at most $\frac{\epsilon}{M}$ and the ex ante constraints have Lipschitz constants at most M, we get that each ex ante constraint is violated at most by $\frac{\epsilon}{M} \cdot M = \epsilon$. Thus Σ'_{OPT} is a feasible solution to our LP, so $E[\Sigma'_{OPT}, u_{\epsilon,M}] \leq E[\Sigma, u_{\epsilon,M}]$.

Since $u_{\epsilon,M}$ is upper semi-continuous and piecewise constant we have: $E[\Sigma_{OPT}, u_{\epsilon,M}] \leq E[\Sigma'_{OPT}, u_{\epsilon,M}]$. Furthermore, the second bullet from Assumption 4.2 yields: $E[\Sigma, u_{\epsilon,M}] - E[\Sigma, u_s] \leq \epsilon$ and $E[\Sigma_{OPT}, u_s] \leq E[\Sigma_{OPT}, u_{\epsilon,M}]$. Combining the last four inequalities implies: $E[\Sigma, u_s] \geq E[\Sigma_{OPT}, u_s] - \epsilon = OPT - \epsilon$. $\square$

Single-criteria FPTAS. The reason for relaxing the constraints is to avoid degenerate cases. For example, finding the root of a polynomial with a single real root can be described in the language of ex ante constraints. This problem has a unique feasible distribution and if we do not relax the constraints, any algorithm missing the exact real root cannot give a satisfactory approximation. Theorem 4.4 can be improved under a Slater-like regularity condition disallowing such degenerate cases.

Assumption 4.5 (Slater-like regularity condition)**.** *There exists a signaling scheme satisfying all the given ex ante constraints with strict inequality.*

Theorem 4.6 (An additive FPTAS for an optimal valid signaling scheme)**.** *Suppose that k is constant, u_s satisfies Assumption 4.2 and we have m ex ante constraints satisfying Assumptions 4.3 and 4.5. Then for every $\epsilon > 0$, there exists a* $\text{poly}\left(m, \frac{1}{\epsilon}\right)$*-algorithm computing an additively ϵ-optimal valid signaling scheme.*

The algorithm applies Theorem 4.4 to a persuasion problem with strengthened ex ante constraints. The analysis compares the output to a convex combination of two outputs of Theorem 4.4 – one might violate the ex ante constraints, while the other satisfies them with strict inequality. We use the proof of Theorem 4.4 to bound the utility loss.

Proof of Theorem 4.6. From Assumption 4.2, u_s is bounded from above by some constant C. Assume w.l.o.g. that $C > 2$. Let OPT be Sender's optimal utility for a valid scheme. Restrict ourselves to small enough values of $0 < \epsilon < \frac{2}{C}$ s.t. strengthening each ex ante constraint by ϵ leaves the set of valid signaling schemes non-empty (it is possible by Assumption 4.5).[8] We return the signaling scheme Σ outputted by the algorithm from Theorem 4.4 on 0.5ϵ and the problem obtained by strengthening each ex ante constraint by 0.5ϵ. Then Σ satisfies the original constraints; it remains to bound its utility loss compared to OPT.

Let Σ' be the output of Theorem 4.4 on $0.125\epsilon^3$ and the original problem; denote by Σ'' the output of Theorem 4.4 on 0.5ϵ and the problem obtained by strengthening each original ex ante constraint by ϵ.

Let M be the maximal Lipschitz constant among the g_is from Assumption 4.3. Then M is not affected by adding constant factors to the constraints; furthermore, note that by Assumption 4.2, $u_{0.125\epsilon^3,M}$ can also serve as $u_{0.5\epsilon,M}$ (since $0.125\epsilon^3 < 0.5\epsilon$ and $\frac{1}{0.125\epsilon^3} = \text{poly}\left(\frac{1}{0.5\epsilon}\right)$). Therefore, by the proof of Theorem 4.4, we can assume w.l.o.g. that $\Sigma, \Sigma', \Sigma''$ are all supported on the extreme points of the pieces of $u_{0.125\epsilon^3,M}$; furthermore, $\frac{1}{1+0.25\epsilon^2}\Sigma' + \frac{0.25\epsilon^2}{1+0.25\epsilon^2}\Sigma''$ satisfies each original ex ante constraint. Note that Σ is 0.5ϵ-additively-optimal among the schemes supported on the above extreme points and satisfying the original ex ante constraints, since Σ is exactly optimal among such schemes if we replace u_s with $u_{0.5\epsilon,M}$, by Theorem 4.4 proof. Thus:

$$E[\Sigma, u_s] \geq E\left[\frac{1}{1+0.25\epsilon^2}\Sigma' + \frac{0.25\epsilon^2}{1+0.25\epsilon^2}\Sigma'', u_s\right] - \frac{\epsilon}{2} =$$
$$\frac{E[\Sigma', u_s]}{1+0.25\epsilon^2} + \frac{0.25\epsilon^2 E[\Sigma'', u_s]}{1+0.25\epsilon^2} - \frac{\epsilon}{2} \geq$$
$$\frac{OPT - 0.125\epsilon^3}{1+0.25\epsilon^2} - \frac{\epsilon}{2} \geq OPT - \epsilon,$$

where the last transition follows from $\frac{\epsilon}{2} < \frac{1}{C} \leq \frac{1}{OPT}$. $\square$

[8] To be precise, we assume that an upper bound on such values of ϵ is known in advance.

5 Ex Ante vs. Ex Post Constraints

In this section, we bound the multiplicative gap in the Sender's optimal utility between ex ante constraints and the corresponding ex post constraints; then we apply our bound to signaling in ad auctions. Note that in full generality, the gap can be arbitrarily large even for $k = 2$ states of nature and $m = 1$ convex constraints:

Example 5.1. *Fix $\epsilon \in \left(0, \frac{1}{2}\right)$; take $\Omega := \{0, 1\}$ with a uniform prior; define $f(p_\sigma) := p_\sigma[\omega = 1]$ and $c := \frac{1}{2} + \epsilon$. Let $u_s(p_\sigma)$ be 0 if $p_\sigma[\omega = 1] \in \left[0, \frac{1}{2}\right]$ and $2 \cdot p_\sigma[\omega = 1] - 1$ otherwise. The ex ante constraint specified by f and c allows full revelation, which yields expected utility of $\frac{1}{2}$ for Sender.*

Convexity of u_s implies that under the corresponding ex post constraint, there exists an optimal signaling scheme for which always $p_\sigma[\omega = 1] \in \{0, c\}$; straightforward calculations show that Sender's optimal utility is $\frac{2\epsilon}{1+2\epsilon}$. Thus, the multiplicative gap tends to ∞ as ϵ tends to 0.

We identify a multiplicatively-relaxed Jensen assumption on u_s parameterized by $M \geq 1$, which, combined with the convexity of the m constraints, yields a bound of M^m on the multiplicative gap between ex ante and ex post constraints.

Assumption 5.2 (parameterized by $M \geq 1$)**.** *For every $\lambda \in [0, 1]$ and $p_{\sigma_1}, p_{\sigma_2} \in \Delta(\Omega)$:*

$$\lambda u_s(p_{\sigma_1}) + (1 - \lambda)u_s(p_{\sigma_2}) \leq M \cdot u_s(\lambda p_{\sigma_1} + (1 - \lambda)p_{\sigma_2}).$$

For example, Assumption 5.2 holds with $M = 2$ for the welfare and revenue functions in the single-item, second-price auction (see full paper version). Note that for some u_s, the assumption does not hold for any finite M: those u_s that "grow too slowly" near 0 (in particular, if u_s maps a nonzero measure of the domain to 0, as in Example 5.1).

Theorem 5.3 (A bound on the multiplicative gap between ex ante and ex post constraints)**.** *Suppose that u_s satisfies Assumption 5.2 with parameter $M \geq 1$. Fix m convex ex ante constraints and let $\Sigma_{\textit{ex ante}}$ be a valid signaling scheme. Then there exists $\Sigma_{\textit{ex post}}$, a valid signaling scheme under the corresponding m ex post constraints, s.t.:* $E[\Sigma_{\textit{ex post}}, u_s] \geq \frac{1}{M^m} \cdot E[\Sigma_{\textit{ex ante}}, u_s]$.

The proof runs Algorithm 1 for each constraint separately. The algorithm repeatedly pools a posterior violating the ex post constraint with a posterior satisfying it with a strict inequality, replacing one of them by a posterior on which the ex post constraint is tight and decreasing the probability mass of the other posterior. We assume that $|\text{supp}(\Sigma)|$ is finite, which is w.l.o.g. by Fact 3.1. The algorithm stops as each iteration decreases the number of posteriors in $\text{supp}(\Sigma)$ on which the ex post constraint is not tight. By the constraint convexity, the resultant scheme satisfies the ex post constraint; Assumption 5.2 implies that the multiplicative loss due to the pooling process (for each constraint) is at most M.

We leave as an open question the tightness of Theorem 5.3 for general m. However, the following statement holds.

Proposition 5.4. *(a) Our analysis is tight for any m and $M = 2$;*[9] *(b) the bound from Theorem 5.3 on the multiplica-*

[9] Note that we use $M = 2$ in our applications.

Algorithm 1 Ex ante to ex post

Input: A signaling scheme Σ with a finite support satisfying: $E[\Sigma, f] \leq c$.
Parameters: A continuous convex function $f : \Delta(\Omega) \to \mathbb{R}$, a constant c.
Output: An updated signaling scheme Σ with a multiplicative expected utility loss of at most M compared to the input, s.t. $\forall p_\sigma \in \mathrm{supp}(\Sigma) : \ f(p_\sigma) \leq c$.

```
S ← supp(Σ) ∩ f^{-1}((−∞, c)).
T ← supp(Σ) ∩ f^{-1}((c, ∞)).
while S, T ≠ ∅ do
    Take q_S ∈ S, q_T ∈ T.
    r_S ← Pr_{p_σ∼Σ}[p_σ = q_S], r_T ← Pr_{p_σ∼Σ}[p_σ = q_T].
    Find λ ∈ (0, 1) s.t. f(λq_S + (1 − λ)q_T) = c.
    Define q_c := λq_S + (1 − λ)q_T.
    supp(Σ) ← supp(Σ) ∪ {q_c}.
    if λr_T ≥ (1 − λ)r_S then
        supp(Σ) ← supp(Σ) \ {q_S}.
        r_S ← 0, r_T ← r_T − (1−λ)r_S/λ, r_c ← r_S/λ.
    else
        supp(Σ) ← supp(Σ) \ {q_T}.
        r_S ← r_S − λr_T/(1−λ), r_T ← 0, r_c ← r_T/(1−λ).
    end if
    Update Σ according to r_S, r_T, r_c.
end while
return Σ.
```

tive gap between ex ante and ex post is tight for $m = 1$ and any M; (c) this gap grows with m and can be at least $m+1$.

Applications. We apply Theorem 5.3 to signaling in ad auctions. We adopt the model of Badanidiyuru, Bhawalkar, and Xu (2018) and add to it constraints on the signaling scheme; here we sketch the model and findings.

Consider a single-item second-price auction with n bidders. Recall from Section 2 that the item being sold is the opportunity to show an online advertisement to a web user, whose characteristics are known to the auctioneer, but not to the bidders. Each bidder targets a certain set of users to whom showing her ad would be most valuable; the auctioneer signals information about which targeted sets the user belongs to. In the language of persuasion, Sender is the auctioneer while Receiver is the set of bidders. The state of nature is a binary vector of length n in which the i-th coordinate specifies whether the web user is in the i-th advertiser's targeted set. Importantly, the number of states of nature is exponential, thus the results from Section 4 do not apply.

The combination of signaling and auction works as follows. The bidders have privately-known, i.i.d. *types* drawn from a commonly-known distribution. The i-th bidder's value for showing her ad is a function of her type and whether the user is in her targeted set. After the auctioneer signals information, the bidders update their values using their posterior beliefs and submit them as bids. The resultant welfare (the winner's value) is Sender's utility u_s.

Proposition 5.5. *Consider the expected welfare of a single-item second-price auction with signaling as a function of p_σ, where the expectation is taken over the bidders' private types. Then Assumption 5.2 is satisfied with $M = 2$.*

This result extends to expected revenue and to sponsored search (slot) auctions. Proposition 5.5 suggests the following "recipe" for solving signaling problems in ad auctions under a constant number of convex ex ante constraints: (approximately) solve for the corresponding ex post constraints; this yields, by Theorem 5.3, a constant-factor approximation for the original problem. The next example demonstrates.

Example 5.6. *Take a single ex ante constraint specified by the function $(-\min\{b_{\omega'} p_\sigma[\omega = \omega']\}_{\omega' \in \Omega})$ with some constant weights $\{b_{\omega'}\}_{\omega' \in \Omega}$. As mentioned in Section 2, this constraint is a possible model for anti-discrimination. Finding an optimal valid scheme $\Sigma^*_{\text{ex ante}}$ is an open question. However, the corresponding ex post constraint just restricts the posteriors to an appropriate simplex; since u_s (the social welfare) is convex – the optimal scheme $\Sigma^*_{\text{ex post}}$ is supported precisely on the vertices of this simplex, and is uniquely specified by Bayes-plausibility. By Theorem 5.3 and Proposition 5.5, $\Sigma^*_{\text{ex post}}$ is a $\frac{1}{2}$-approximation to $\Sigma^*_{\text{ex ante}}$.*

6 Future Work

We study the setting of ex ante- and ex post-constrained persuasion, which has applications to areas including ad auctions and limited attention. A future research direction, especially considering Theorem 5.3, is studying (nearly) optimal signaling schemes under common ex post constraints, such as KL divergence. Another direction is constrained persuasion with *private* signaling, e.g., when Receivers have binary actions and Sender's utility is a function of the set of Receivers adopting a certain action (Arieli and Babichenko 2019).

Acknowledgments. This research has been supported by The Israel Science Foundation (grant #336/18). The first author's research has been partially supported by The U.S.–Israel Binational Science Foundation (grant #BSF 2026924) and by The German–Israeli Foundation for Scientific Research and Development (grant #GIF 2027111); the second author is a Taub Fellow (supported by The Taub Family Foundation). The authors thank Ruggiero Cavallo for helpful conversations that motivated this research and the anonymous reviewers for their helpful suggestions on improving this paper.

References

Antioch, G. 2013. Persuasion is now 30 per cent of US GDP: Revisiting McCloskey and Klamer after a quarter of a century. *Economic Roundup* 1: 1.

Arieli, I.; and Babichenko, Y. 2019. Private Bayesian persuasion. *Journal of Economic Theory* 182: 185–217.

Aumann, R. J.; Maschler, M.; and Stearns, R. E. 1995. *Repeated games with incomplete information.* MIT Press.

Badanidiyuru, A.; Bhawalkar, K.; and Xu, H. 2018. Targeting and signaling in ad auctions. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2545–2563. SIAM.

Blackwell, D. 1953. Equivalent comparisons of experiments. *The annals of mathematical statistics* 265–272.

Bloedel, A. W.; and Segal, I. R. 2018. Persuasion with rational inattention. Available at SSRN 3164033.

Celis, E. L.; Kapoor, S.; Salehi, F.; and Vishnoi, N. K. 2019. Controlling polarization in personalization: An algorithmic framework. In *Proceedings of the conference on fairness, accountability, and transparency*, 160–169.

Celis, E. L.; Mehrotra, A.; and Vishnoi, N. K. 2019. Toward controlling discrimination in online ad auctions. In *International Conference on Machine Learning*, 4456–4465.

Charalambos, D. A.; and Border, K. C. 2006. *Infinite dimensional analysis: a hitchhiker's guide*. Springer.

Cheng, Y.; Cheung, H. Y.; Dughmi, S.; Emamjomeh-Zadeh, E.; Han, L.; and Teng, S. H. 2015. Mixture selection, mechanism design, and signaling. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, 1426–1445. IEEE.

Daskalakis, C.; Papadimitriou, C. H.; and Tzamos, C. 2016. Does Information Revelation Improve Revenue? In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC*, 233–250.

Doval, L.; and Skreta, V. 2018. Constrained information design: Toolkit. *arXiv preprint arXiv:1811.03588* .

Dughmi, S. 2017. Algorithmic information structure design: a survey. *ACM SIGecom Exchanges* 15(2): 2–24.

Dughmi, S.; Immorlica, N.; O'Donnell, R.; and Tan, L. 2015. Algorithmic Signaling of Features in Auction Design. In *Algorithmic Game Theory - 8th International Symposium, SAGT*, 150–162. Springer.

Dughmi, S.; Immorlica, N.; and Roth, A. 2014. Constrained signaling in auction design. In *Proceedings of the 25th annual ACM-SIAM symposium on Discrete algorithms*, 1341–1357. SIAM.

Dughmi, S.; and Xu, H. 2017. Algorithmic persuasion with no externalities. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 351–368.

Dwork, C.; and Roth, A. 2014. The Algorithmic Foundations of Differential Privacy. *Foundations and Trends in Theoretical Computer Science* 9(3-4): 211–407.

Eilat, R.; Eliaz, K.; and Mu, X. 2019. Optimal privacy-constrained mechanisms. CEPR Discussion Paper No. DP13536.

Emek, Y.; Feldman, M.; Gamzu, I.; Leme, R. P.; and Tennenholtz, M. 2014. Signaling schemes for revenue maximization. *ACM Transactions on Economics and Computation (TEAC)* 2(2): 1–19.

Fu, H.; Jordan, P.; Mahdian, M.; Nadav, U.; Talgam-Cohen, I.; and Vassilvitskii, S. 2012. Ad auctions with data. In *International Symposium on Algorithmic Game Theory*, 168–179. Springer.

Hart, S.; and Nisan, N. 2013. *The menu-size complexity of auctions*. Center for the Study of Rationality.

Hart, S.; and Nisan, N. 2017. Approximate revenue maximization with multiple items. *Journal of Economic Theory* 172: 313–347.

Ichihashi, S. 2019. Limiting Sender's Information in Bayesian Persuasion. *Games and Economic Behavior* 117: 276–288.

Ichihashi, S. 2020. Online privacy and information disclosure by consumers. *American Economic Review* 110(2): 569–95.

Kamenica, E.; and Gentzkow, M. 2011. Bayesian persuasion. *American Economic Review* 101(6): 2590–2615.

Le Treust, M.; and Tomala, T. 2019. Persuasion with limited communication capacity. *Journal of Economic Theory* 184: 104940.

Lipnowski, E.; Mathevet, L.; and Wei, D. 2020. Attention management. *American Economic Review: Insights* 2(1): 17–32.

McCloskey, D.; and Klamer, A. 1995. One Quarter of GDP is Persuasion. *The American Economic Review* 85(2): 191–195.

Milgrom, P. R.; and Weber, R. J. 1982. A theory of auctions and competitive bidding. *Econometrica: Journal of the Econometric Society* 1089–1122.

Miltersen, P. B.; and Sheffet, O. 2012. Send mixed signals: earn more, work less. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 234–247.

Myerson, R. B. 1981. Optimal auction design. *Mathematics of Operations Research* 6(1): 58–73.

Simon, H. A. 1996. Designing organizations for an information-rich world. *International Library of Critical Writings in Economics* 70: 187–202.

Sims, C. A. 2003. Implications of rational inattention. *Journal of monetary Economics* 50(3): 665–690.

Tsakas, E.; and Tsakas, N. 2019. Noisy persuasion. *Available at SSRN 2940681* .

Vølund, R. T. 2018. Bayesian persuasion on compact subsets. *Theoretical Models in Behavioral Economics* 64–77.