Algorithmic Bayesian Persuasion

Shaddin Dughmi[*]
Department of Computer Science
University of Southern California
shaddin@usc.edu

Haifeng Xu[†]
Department of Computer Science
University of Southern California
haifengx@usc.edu

April 2, 2015

Abstract

We consider the *Bayesian Persuasion* problem, as formalized by Kamenica and Gentzkow [27], from an algorithmic perspective in the presence of high dimensional and combinatorial uncertainty. Specifically, one player (the *receiver*) must take one of a number of actions with a-priori unknown payoff; another player (the *sender*) is privy to additional information regarding the payoffs of the various actions for both players. The sender can commit to revealing a noisy signal regarding the realization of the payoffs of various actions, and would like to do so as to maximize her own payoff in expectation assuming that the receiver rationally acts to maximize his own payoff. This models a number of natural strategic interactions, in domains as diverse as e-commerce, advertising, politics, law, security, finance, and others. When the payoffs of various actions follow a joint distribution (the common prior), the sender's problem is nontrivial, and its complexity depends on the representation of the prior.

Assuming a Bayesian receiver, we study the sender's problem with an algorithmic and approximation lens. We show two results for the case in which the payoffs of different actions are i.i.d and given explicitly: a polynomial-time (exact) algorithmic solution, and a "simple" $(1-1/e)$ approximation algorithm. Both results hinge on a "symmetry characterization" of the optimal signaling scheme. The former result also involves a connection to auction theory, and in particular to Border's characterization of the space of feasible reduced-form auctions. For the latter result, our algorithm decouples the signaling problem for the different actions and signals independently for each. This decoupling drives a larger, conceptual point: collaborative persuasion by multiple parties (the senders) is a parallelizable task, requiring no coordination when actions are identical and independent and only an approximate solution is sought. We leave open the intriguing question of whether either of these two results extends to independent yet not necessarily identical payoff distributions.

We then turn our attention to the general problem, and consider distributions over action payoffs given by a *sampling oracle*. Somewhat surprisingly, we show that even in this general setting the persuasion problem admits a fully polynomial-time approximation scheme (FPTAS) in a bi-criteria sense. As our main technical tool to prove this theorem, we study the algorithmics of a natural and abstract question on vector spaces, which we term *dispersion testing*: Given two distributions $\mathcal{A}$ and $\mathcal{B}$ supported on a finite-dimensional vector space, decide whether $\mathcal{A}$ is a mean-preserving spread of $\mathcal{B}$, and if so compute the inverse of the associated spreading map. We employ tools from convex optimization and optimal transport theory to design an approximate test for the dispersion testing problem, and relate the persuasion problem to dispersion testing via a randomized Turing reduction employing the Ellipsoid method.

[*]Supported in part by NSF CAREER Award CCF-1350900.
[†]Supported by NSF grant CCF-1350900.

1 Introduction

"*One quarter of the GDP is persuasion.*"

This is both the title, and the thesis, of a 1995 paper by McCloskey and Klamer [31]. Since then, persuasion as a share of economic activity appears to be growing — a more recent estimate places the figure at 30% [5]. As both papers make clear, persuasion is intrinsic in most human endeavors. When the tools of "persuasion" are tangible — say goods, services, or money — this is the domain of traditional *mechanism design*, which steers the actions of one or many self-interested agents towards a designer's objective. What [31, 5] and much of the relevant literature refer to as persuasion, however, are scenarios in which the power to persuade derives from an *informational advantage* of some party over others. This is also the sense in which we use the term. Such scenarios are increasingly common in the information economy, and it is therefore unsurprising that persuasion has been the subject of a large body of work, motivated by domains as varied as auctions [7, 20, 19, 8], advertising [4, 26, 13], voting [3], security [38, 33], multi-armed bandits [29, 30], medical research [28], and financial regulation [23, 24]. (For an empirical survey of persuasion, we refer the reader to [16]). What is surprising, however, is the lack of systematic study of persuasion from an *algorithmic* perspective; this is what we embark on in this paper.

In the large body of literature devoted to persuasion, perhaps no model is more basic and fundamental than the *Bayesian Persuasion* model of Kamenica and Gentzkow [27], generalizing an earlier model by Brocas and Carrillo [11]. Here there are two players, who we call the *sender* and the *receiver*. The receiver is faced with selecting one of a number of *actions*, each of which is associated with an a-priori unknown payoff to both players. The *state of nature*, describing the payoff to the sender and receiver from each action, is drawn from a prior distribution known to both players. However, the sender possesses an informational advantage, namely access to the *realized* state of nature prior to the receiver choosing his action. In order to persuade the receiver to take a more favorable action for her, the sender can *commit* to a policy, often known as an *information structure* or *signaling scheme*, of releasing information about the realized state of nature to the receiver before the receiver makes his choice. This policy may be simple, say by always announcing the payoffs of the various actions or always saying nothing, or it may be intricate, involving partial information and added noise. Crucially, the receiver is aware of the sender's committed policy, and moreover is rational and Bayesian. We examine the sender's algorithmic problem of implementing the optimal signaling scheme in this paper. A solution to this problem, i.e., a signaling scheme, is an algorithm which takes as input the description of a state of nature and outputs a signal, potentially utilizing some internal randomness.

Two Examples

To illustrate the intricacy of Bayesian Persuasion, [27] use a simple example in which the sender is a prosecutor, the receiver is a judge, and the state of nature is the guilt or innocence of a defendant. The receiver (judge) has two actions, conviction and acquittal, and wishes to maximize the probability of rendering the correct verdict. On the other hand, the sender (prosecutor) is interested in maximizing the probability of conviction. As they show, it is easy to construct examples in which the optimal signaling scheme for the sender releases noisy partial information regarding the guilt or innocence of the defendant. For example, if the defendant is guilty with probability $\frac{1}{3}$, the prosecutor's best strategy is to claim "guilt" whenever the defendant is guilty, and also claim "guilt" just under half the time when the defendant is innocent. As a result, the defendant will be convicted whenever the prosecutor claims "guilt" (happening with probability just under $\frac{2}{3}$), assuming that the judge is fully aware of the prosecutor's signaling scheme. We note that it is not in the prosecutor's interest to always claim "guilt", since a rational judge aware of such a policy would ascribe no meaning to such a signal, and render his verdict based solely on his prior belief — in this case, this would always lead to acquittal.[1]

When there are more than two actions, additional counterintuitive phenomena can arise. Consider the following example with three actions and two states of nature. The sender's payoff from the three actions is constant, given by the vector $\boldsymbol{s} = (2, 1, 0)$, indexed by action, in both states of nature. The receiver's payoff vector, indexed by action, equals $\boldsymbol{r} = (2, 1, 0)$ with probability $\frac{2\epsilon}{1+2\epsilon}$ (the first state of nature), and $\boldsymbol{r}' = (0, 1 - \epsilon, 1)$ with probability $\frac{1}{1+2\epsilon}$ (the second state of nature), where $0 < \epsilon < 1/2$ is arbitrary. Given that $\boldsymbol{s} = \boldsymbol{r}$ — i.e., the sender's and receiver's utilities are aligned in the first state of nature — one might be tempted to conclude that, whatever the optimal signaling scheme might be, it should at least guarantee that the receiver chooses both players' favorite action — action 1 — in the first state of nature. However, some thought reveals that this is not the case; in fact, the optimal signaling scheme "recommends" actions 1 and 2 with equal probability in the first state of nature, and recommends action 2 in the second state of nature. The rationale is as follows: some probability mass from the first state of nature is "lumped" with the second state of nature, and this is to "bump up" the posterior expected reward of action 2. Since the second state of nature is much more probable, and the choice facing the receiver there between actions 2 and 3 has much starker consequences for the sender, it is best to sacrifice a small probability of action 1 for a much larger probability of action 2.

Results and Techniques

Motivated by these intricacies, we embark on designing algorithms for optimal persuasion in the presence of multiple actions. We first observe that a linear program with a variable for each (state-of-nature, action) pair computes a description of the optimal signaling scheme. However, when action payoffs are distributed according to a joint distribution — say exhibiting some degree of independence across different actions — the number of states of nature may be exponential in the number of actions; in such settings, both the number of variables and constraints of this linear program are exponential in the number of actions. Perhaps surprisingly, we nevertheless discover that optimal or near optimal signaling schemes can be computed efficiently — in time polynomial in the number of actions — under fairly general conditions. We consider two models: one in which action payoffs are i.i.d from an explicitly-described marginal distribution, and another in which the joint distribution of action payoffs is arbitrary and given by a black-box sampling oracle.

We start with the i.i.d model, and show two results: a "simple" and polynomial-time $(1 - 1/e)$-approximate signaling scheme, and a polynomial-time implementation of the optimal scheme. Both results hinge on a "symmetry characterization" of the optimal signaling scheme in the i.i.d. setting, closely related to analogous symmetrization results from algorithmic mechanism design by [15]. Our "simple" scheme decouples the signaling problem for the different actions and signals *independently* for each. This decoupling exposes an interesting property of optimal signaling: collaborative persuasion by multiple senders — when each sender observes the payoff of one or more actions — can be parallelized, requiring no coordination when actions are i.i.d. so long as we are willing to settle for approximate optimality. Our optimal scheme involves a connection to auction theory, and in particular to Border's characterization of the space of feasible reduced-form auctions [10, 9], as well as its algorithmic properties [12, 2, 37]. We leave open the intriguing

[1] In other words, a signal is an abstract object with no intrinsic meaning, and is only imbued with meaning by virtue of how it is used. In particular, a signal has no meaning beyond the posterior distribution on states of nature it induces.

question of whether either of these two results extends to independent yet not necessarily identical payoff distributions.

In the black-box model, we show that the persuasion problem admits a fully polynomial-time (in the number of actions) approximation scheme (FPTAS), in a bi-criteria sense. As our main technical tool to prove this theorem, we study the algorithmics of a natural and abstract question on vector spaces, which we term *dispersion testing*: Given two distributions $\mathcal{A}$ and $\mathcal{B}$ supported on a finite-dimensional vector space, decide whether $\mathcal{A}$ is a mean-preserving spread of $\mathcal{B}$, and if so compute the inverse of the associated spreading map. We employ tools from convex optimization and optimal transport theory to design an approximate test for the dispersion testing problem, and relate the persuasion problem to dispersion testing via a randomized Turing reduction employing the Ellipsoid method.

Additional Discussion of Related Work

To our knowledge, Brocas and Carrillo [11] were the first to formally consider persuasion through information control. They consider a sender with the ability to costlessly acquire information regarding the payoffs of the receiver's actions, with the stipulation that acquired information is available to both players. This is technically equivalent to our (and Kamenica and Gentzkow's [27]) informed sender who commits to a signaling scheme. Brocas and Carrillo restrict attention to a particular setting with two states of nature and three actions, and characterize optimal policies for the sender and their associated payoffs.

Kamenica and Gentzkow's [27] Bayesian Persuasion model naturally generalizes [11] to finite (or infinite yet compact) states of nature and action spaces. They establish a number of properties of optimal information structures in this model; most notably, they characterize settings in which signaling strictly benefits the sender in terms of the convexity/concavity of the sender's payoff as a function of the receiver's posterior belief.

Variants and generalizations of the Bayesian persuasion model have received a lot of attention in recent years. For example, [21] studies persuasion in the presence of multiple senders, [3] studies the case of multiple receivers in a voting setting, [22] studies optimal persuasion when the sender's information comes at a cost, and [34] examine Bayesian persuasion in the presence of an outside option for the receiver. We refer the reader to references therein for additional related work on Bayesian persuasion.

More generally, the Bayesian persuasion model is a special case of the study of optimal information structures in games. This has its roots in the early works on information economics, such as those by Akerlof [1] and Spence [35]. More recently, Bergemann and Morris [6] characterize the space of information structures in general games, and relate it to the space of correlated equilibria. Despite this appreciation of the importance of information in strategic interactions, it is only recently that researchers have started viewing the information structure of a game as a mathematical object to be designed and computed, rather than merely an exogenous variable. Recent work in the CS community, including by Emek et al. [19], Miltersen and Sheffet [32], Guo and Deligkas [25], and Dughmi et al. [18], examines optimal signaling in a variety of auction settings, and presents polynomial-time algorithms and hardness results. Moreover, Dughmi [17] examines optimal signaling in zero-sum games, and exhibits hardness results assuming the planted clique conjecture.

2 Preliminaries

2.1 The Bayesian Persuasion Model

In a persuasion game, there are two players: a *sender*, and *receiver*. The receiver is faced with selecting an action from $[n] = \{1, \ldots, n\}$, with an a-priori-unknown payoff to each of the sender and receiver. We assume payoffs are a function of an unknown *state of nature* θ, drawn from an abstract set Θ of potential realizations of nature. Specifically, the sender and receiver's payoffs are functions $s, r : \Theta \times [n] \to \mathbb{R}$, respectively. We often use $\boldsymbol{r} = \boldsymbol{r}(\theta) \in \mathbb{R}^n$ to denote the receiver's payoff vector as a function of the state of nature, where $r_i(\theta)$ is the receiver's payoff if he takes action i and the state of nature is θ. Similarly $\boldsymbol{s} = \boldsymbol{s}(\theta) \in \mathbb{R}^n$ denotes the sender's payoff vector, where $s_i(\theta)$ is the sender's payoff if the receiver takes action i and the state of nature is θ. For convenience, we sometimes conflate the abstract set Θ indexing states of nature with the family of realizable payoff vector pairs $(\boldsymbol{s}, \boldsymbol{r})$, and this is without loss of generality — i.e. we will often think of Θ as a subset of $\mathbb{R}^n \times \mathbb{R}^n$. In much of this paper we assume that Θ is finite for notational convenience, though this assumption is inconsequential for our results in Sections 4 and 5.

We assume that the state of nature is a-priori unknown to the receiver, and drawn from a common-knowledge prior distribution λ supported on Θ. The sender, on the other hand, has access to the realization of θ, and can commit to a policy of partially revealing information regarding its realization before the receiver selects his action. Specifically, the sender commits to a *signaling scheme* φ, mapping (possibly randomly) states of nature Θ to a family of *signals* Σ. For $\theta \in \Theta$, we use $\varphi(\theta)$ to denote the (possibly random) signal selected when the state of nature is θ. Moreover, we use $\varphi(\theta, \sigma)$ to denote the probability of selecting the signal σ given a state of nature θ. An algorithm *implements* a signaling scheme φ if it takes as input a state of nature θ, and samples the random variable $\varphi(\theta)$.

Given a signaling scheme φ with signals Σ, each signal $\sigma \in \Sigma$ is realized with probability $\alpha_\sigma = \sum_{\theta \in \Theta} \lambda_\theta \varphi(\theta, \sigma)$. Conditioned on the signal σ, the expected payoffs to the receiver of the various actions are summarized by the vector $\boldsymbol{r}(\sigma) = \frac{1}{\alpha_\sigma} \sum_{\theta \in \Theta} \lambda_\theta \varphi(\theta, \sigma) \boldsymbol{r}(\theta)$. Similarly, the sender's payoffs as a function of the receiver's action are summarized by $\boldsymbol{s}(\sigma) = \frac{1}{\alpha_\sigma} \sum_{\theta \in \Theta} \lambda_\theta \varphi(\theta, \sigma) \boldsymbol{s}(\theta)$. On receiving a signal, a rational receiver performs a Bayesian update and selects an action maximizing his posterior expected payoff. Specifically, on signal σ, a rational receiver selects an action $i^*(\sigma) \in \operatorname{argmax}_i r_i(\sigma)$ with expected receiver utility $\max_i r_i(\sigma)$. This induces utility $s_{i^*(\sigma)}(\sigma)$ for the sender. In the event of ties when selecting $i^*(\sigma)$, we assume those ties are broken in favor of the sender.

We adopt the perspective of a sender looking to design φ in order to maximize her expected utility $\sum_\sigma \alpha_\sigma s_{i^*(\sigma)}(\sigma)$. When φ maximizes the sender's utility, we say it is *optimal*, and when it yields a sender utility within an additive [multiplicative] ϵ of the best possible, we say it is ϵ-*optimal* [ϵ-approximate] in the additive [multiplicative] sense. For additive guarantees, we naturally constrain the range of payoffs to $[-1, 1]$, and for multiplicative guarantees we restrict to nonnegative payoffs.

A simple argument shows that an optimal signaling scheme need not use more than n signals, with one *recommending* each action. Such a *direct* scheme φ has signals $\Sigma = \{\sigma_1, \ldots, \sigma_n\}$, and satisfies $r_i(\sigma_i) \geq r_j(\sigma_i)$ for all $i, j \in [n]$. To see that a direct optimal scheme exists, note that any two signals σ, σ' with $i^*(\sigma) = i^*(\sigma')$ may be merged without changing the sender's utility. We think of σ_i as a signal recommending action i, and the requirement $r_i(\sigma_i) \geq \max_j r_j(\sigma_i)$ as an *incentive-compatibility* constraint on our signaling scheme. With these observations, we can write the sender's optimization problem as the following linear program with variables $\{\varphi(\theta, \sigma_i) : \theta \in \Theta, i \in [n]\}$.

$$
\begin{array}{lll}
\text{maximize} & \sum_{\theta\in\Theta}\sum_{i=1}^{n}\lambda_\theta\varphi(\theta,\sigma_i)s_i(\theta) & \\
\text{subject to} & \sum_{i=1}^{n}\varphi(\theta,\sigma_i)=1, & \text{for } \theta\in\Theta. \\
& \sum_{\theta\in\Theta}\lambda_\theta\varphi(\theta,\sigma_i)r_i(\theta)\geq\sum_{\theta\in\Theta}\lambda_\theta\varphi(\theta,\sigma_i)r_j(\theta), & \text{for } i,j\in[n]. \\
& \varphi(\theta,\sigma_i)\geq 0, & \text{for } \theta\in\Theta, i\in[n].
\end{array}
\tag{1}
$$

For our results in Section 5, we relax our incentive constraints by assuming that the receiver follows the recommended action so long as it approximately maximizes his utility — specifically, we replace the incentive-compatibility constraint with the approximate version $r_i(\sigma_i)\geq\max_j r_j(\sigma_i)-\epsilon$, for a parameter ϵ. This replaces the incentive constraints in (1) with the constraints

$$\sum_{\theta\in\Theta}\lambda_\theta\varphi(\theta,\sigma_i)r_i(\theta)\geq\sum_{\theta\in\Theta}\lambda_\theta\varphi(\theta,\sigma_i)(r_j(\theta)-\epsilon)$$

for all $i,j\in[n]$. We call such signaling schemes ϵ-incentive compatible. We judge the approximate optimality of an approximately incentive compatible scheme relative to the best (absolutely) incentive compatible scheme; i.e., in a bi-criteria sense.

2.2 Vector Spaces, Distributions, and Transport Theory

Vector Space Concepts. Given a real vector space $\mathcal{V}$, and $X\subseteq\mathcal{V}$, we use $\mathbf{aff}(X)$ to denote the *affine hull* of X. When $\mathcal{V}=\mathbb{R}^N$ is finite dimensional Euclidean space and an affine subspace $S\subseteq\mathcal{V}$ is clear from context, we use $\mathbf{vol}(.)$ to denote the standard Lebesgue measure on S. We consider normed vector spaces $(\mathcal{V},||.||)$; in Euclidean space, we use $||.||_p$ to denote the L^p norm. Throughout the paper, we use boldface to denote vectors, and use $\boldsymbol{x}\cdot\boldsymbol{y}$ to denote the standard inner product of $\boldsymbol{x}$ and $\boldsymbol{y}$.

Distributions. We work with distributions supported on a vector space $\mathcal{V}$. To simplify some of our derivations we assume all our distributions have finite support, though our results immediately extend to continuous distributions. For a distribution $\mathcal{D}$, we use $\mathbf{supp}(\mathcal{D})$ to denote its support. As usual, we use $\mathbf{E}[\boldsymbol{v}]$ to denote the expected value of a random vector $\boldsymbol{v}\sim\mathcal{D}$, and $\mathbf{Pr}[\boldsymbol{v}]$ to denote the probability of a particular $\boldsymbol{v}\in\mathbf{supp}(\mathcal{D})$. Crucially, we work a lot with conditional expectation as a random variable; specifically, given jointly distributed random vectors $\boldsymbol{a},\boldsymbol{b}\in\mathcal{V}$, we often interpret $\mathbf{E}[\boldsymbol{a}|\boldsymbol{b}]$ as a random vector in its own right, evaluating to the conditional expected value of $\boldsymbol{a}$ given the realized value of $\boldsymbol{b}$.

Earthmover Distance. Given two distributions $\mathcal{A}$ and $\mathcal{B}$ supported on a normed vector space $(\mathcal{V},||.||)$, we measure their *earthmover distance* $\mathrm{EMD}(\mathcal{A},\mathcal{B})$, defined as the optimal value of the following optimization problem with variables $\{x(\boldsymbol{a},\boldsymbol{b}):\boldsymbol{a}\in\mathbf{supp}(\mathcal{A}),\boldsymbol{b}\in\mathbf{supp}(\mathcal{B})\}$.

$$
\begin{array}{lll}
\text{minimize} & \sum_{\boldsymbol{a}\in\mathbf{supp}(\mathcal{A}),\boldsymbol{b}\in\mathbf{supp}(\mathcal{B})}x(\boldsymbol{a},\boldsymbol{b})||\boldsymbol{b}-\boldsymbol{a}|| & \\
\text{subject to} & \sum_{\boldsymbol{b}\in\mathbf{supp}(\mathcal{B})}x(\boldsymbol{a},\boldsymbol{b})=\mathbf{Pr}_{\mathcal{A}}[\boldsymbol{a}], & \text{for } \boldsymbol{a}\in\mathbf{supp}(\mathcal{A}). \\
& \sum_{\boldsymbol{a}\in\mathbf{supp}(\mathcal{A})}x(\boldsymbol{a},\boldsymbol{b})=\mathbf{Pr}_{\mathcal{B}}[\boldsymbol{b}], & \text{for } \boldsymbol{b}\in\mathbf{supp}(\mathcal{B}). \\
& x(\boldsymbol{a},\boldsymbol{b})\geq 0, & \text{for } \boldsymbol{a}\in\mathbf{supp}(\mathcal{A}),\boldsymbol{b}\in\mathbf{supp}(\mathcal{B}).
\end{array}
\tag{2}
$$

When the norm is not clear from context, we use $\mathrm{EMD}_{||.||}$ to denote earthmover distance relative to a particular norm $||.||$. In the above optimization problem, each feasible solution x describes a *transport map*, which is a fractional matching of the probability mass of $\mathcal{A}$ with the probability mass of $\mathcal{B}$. The objective function measures the cost of moving probability mass from $\mathcal{A}$ to $\mathcal{B}$ using transport map x, as measured by the norm $||.||$. An alternative view of x is as a coupling of random variables $\boldsymbol{a} \sim \mathcal{A}$ and $\boldsymbol{b} \in \mathcal{B}$, where $x(\boldsymbol{a}, \boldsymbol{b})$ is the joint probability mass function of $\boldsymbol{a}$ and $\boldsymbol{b}$; from this vantage point, the objective measures the expected distance of $\boldsymbol{a}$ and $\boldsymbol{b}$ over draws from this joint distribution. For more on optimal transport theory, we refer the reader to [36].

2.3 Convex Optimization and Approximate Separation Oracles

Convex Optimization. We consider *convex optimization problems* Π for which each instance $\mathcal{I} \in \Pi$ is a tuple $(N, \mathcal{P}, f)$ where N is an integer, $\mathcal{P} \subseteq \mathbb{R}^N$ is a closed and convex feasible set, and $f : \mathbb{R}^N \to \mathbb{R} \cup \{\infty\}$ is a convex objective function.[2] We assume instances are expressed as binary strings in input, and use $\langle \mathcal{I} \rangle$ to denote the description length of an instance $\mathcal{I} \in \Pi$. Moreover, we assume that instances are encoded such that f and a sub-gradient of f can be evaluated at an arbitrary point in $\text{poly}(N, \langle \mathcal{I} \rangle)$ time.

Given an instance $\mathcal{I} = (N, \mathcal{P}, f)$, we define some convenient terminology and notations. We refer to points in $\mathcal{P}$ as *feasible*. We denote $OPT = \inf \{f(\boldsymbol{x}) : \boldsymbol{x} \in \mathcal{P}\}$ and $MAX = \sup \{f(\boldsymbol{x}) : \boldsymbol{x} \in \mathcal{P}\}$, and refer to points in $\left\{\boldsymbol{x} \in \mathbb{R}^N : f(\boldsymbol{x}) \leq OPT + \epsilon[MAX - OPT]\right\}$ as *ϵ-optimal* (note that an ϵ-optimal point need not be feasible). We use $\mathbf{vol}(.)$ to denote the Lebesgue measure on $\mathbf{aff}(\mathcal{P})$.

Separation Oracles with Unstructured One-Sided Error. In Section 5 we use a convex program to solve for an approximately optimal signaling scheme, using as an approximate separation oracle an algorithm — which we term a dispersion test (Section 4) — for approximately testing a particular property of a distribution supported on the vector space of payoffs. In implementing this idea, however, we confront a technical hurdle: the "error" in our separation oracle is, invariably in this case, measured in the earthmover distance between distributions rather than the more familiar 2-norm (or, more generally, p-norm) distance defined on the space of variables representing the distribution in our convex program. Even worse, the earthmover distance does not correspond to a norm at all, nor does it approximate any norm to within a finite factor. As a consequence, our separation oracle does not satisfy the oft-stated requirements of the Ellipsoid method for (approximate) convex optimization; namely, that it errs only at points which are close to the boundary of the (true) feasible set in Euclidean (2-norm) distance.

We overcome this hurdle by analyzing the Ellipsoid algorithm when the employed separation oracle suffers unstructured "one-sided error". By this we mean that the oracle accepts every feasible point, but may accept some infeasible points, even some that are very far from being feasible. We show that, under some mild technical conditions, a particular implementation of the ellipsoid method is guaranteed to output a (possibly infeasible) solution which is (a) accepted by the separation oracle, and (b) within ϵ of the optimal feasible solution in objective value. We note that our ideas in this section are not new, and in fact a similar statement to ours appears in [37, Chapter 6.2], albeit presented in dual form— specifically, they convert an approximate oracle for linear optimization to an approximate separation oracle. Nevertheless, we present a self-contained statement and proof here.

Definition 2.1. *A* one-sided-error separation oracle *(OSO) for a convex optimization problem* Π *takes as input* $\mathcal{I} = (N, \mathcal{P}, *) \in \Pi$ *and a point* $\boldsymbol{x} \in \mathbb{R}^N$*, and either accepts* $\boldsymbol{x}$ *or outputs a vector* $\boldsymbol{h} \in \mathbb{R}^N$ *such that* $\boldsymbol{h} \cdot \boldsymbol{x} > \boldsymbol{h} \cdot \boldsymbol{y}$ *for all* $\boldsymbol{y} \in \mathcal{P}$.

For an OSO $\mathcal{O}$ we use $\mathcal{P}^+_{\mathcal{O}}$ to denote the family of all points accepted by $\mathcal{O}$ — note that $\mathcal{P} \subseteq \mathcal{P}^+_{\mathcal{O}}$. When $\mathcal{O}$ is clear from context, we refer to points in $\mathcal{P}^+_{\mathcal{O}}$ as *acceptable*.

[2] When f is defined only on a proper and convex subset of $\mathbb{R}^N$ containing $\mathcal{P}$, we simply take its value to be ∞ elsewhere as is traditional.

Definition 2.2. *An algorithm efficiently solves the* weak optimization problem *for* Π *with respect to an OSO* $\mathcal{O}$ *if it takes an input instance* $\mathcal{I} = (N, \mathcal{P}, f) \in \Pi$*, an approximation parameter* $\epsilon > 0$*, an ellipsoid* $\mathcal{E}$ *such that* $\mathcal{P} \subseteq \mathcal{E} \subseteq \mathbf{aff}(\mathcal{P})$*, and numbers* r, R *such that* $r \leq \mathbf{vol}(\mathcal{P})$ *and* $R \geq \mathbf{vol}(\mathcal{E})$*, runs in time* $\text{poly}(N, \langle \mathcal{I} \rangle, \frac{1}{\epsilon}, \log \frac{R}{r})$*, and outputs an acceptable and* ϵ*-optimal solution.*

We show that a straightforward adaptation of the Ellipsoid algorithm, given oracle access to OSO $\mathcal{O}$, solves the weak optimization problem for Π with respect to $\mathcal{O}$.

Theorem 2.3. *Let* Π *be a convex optimization problem, and let* $\mathcal{O}$ *be an OSO for* Π*. Given oracle access to* $\mathcal{O}$*, there is an algorithm for efficiently solving the weak optimization problem for* Π *with respect to* $\mathcal{O}$*.*

Proof. Let $\mathcal{I} = (N, \mathcal{P}, f) \in \Pi$, and let ϵ, $\mathcal{E}$, r, and R be as in Definition 2.2. We run the ellipsoid algorithm starting with $\mathcal{E}$, and as usual use either $\mathcal{O}$ or the objective gradient $\bigtriangledown f$ to update the working ellipsoid in each iteration. We consider the implementation of the Ellipsoid method which prioritizes $\mathcal{O}$ over $\bigtriangledown f$ in updating the ellipsoid. We show that this implementation of the ellipsoid method must encounter, after at most $\text{poly}(N, \frac{1}{\epsilon}, \log \frac{R}{r})$ iterations, an acceptable and ϵ-optimal point as the center of the working ellipsoid. Stopping the ellipsoid algorithm after a suitable number of iterations, and returning the best acceptable point — in terms of objective value — encountered over all iterations of the algorithm then yields the theorem.

First, we lower-bound the volume of the family of ϵ-optimal feasible solutions — namely $\{\boldsymbol{x} \in \mathcal{P} : f(\boldsymbol{x}) \leq OPT + \epsilon[MAX - OPT]\}$. It follows easily from Jensen's inequality and a simple geometric argument that this set has volume at least $\epsilon^N \, \mathbf{vol}(\mathcal{P})$. This, in turn, is at least $\epsilon^N r$.

Second, we show that so long the ellipsoid algorithm hasn't encountered an ϵ-optimal feasible solution, all such points must be inside the working ellipsoid. To see this, consider the first iteration at which an ϵ-optimal feasible solution is excluded by the ellipsoid method, and let $\boldsymbol{c}$ be the corresponding ellipsoid center and $\boldsymbol{h}$ be the corresponding hyperplane used to update the ellipsoid. Note that $\boldsymbol{h}$ could not have been derived from $\mathcal{O}$, since the corresponding half spaces never exclude any point from the feasible set by our definition of OSO's. Therefore, $\boldsymbol{h} = \bigtriangledown f(\boldsymbol{c})$. By our priority order on $\mathcal{O}$ and $\bigtriangledown f$, it must be that $\boldsymbol{c}$ is acceptable. Since f is convex, the only points excluded by $\boldsymbol{h}$ have objective value at least $f(\boldsymbol{c})$, and therefore our assumption that we exclude some ϵ-optimal solution implies that $f(\boldsymbol{c}) \leq OPT + \epsilon[MAX - OPT]$. To summarize, $\boldsymbol{c}$ itself is an ϵ-optimal feasible solution.

Our proof is now complete by recalling that the volume of the working ellipsoid decreases by a constant factor every N iterations. $\square$

3 The Persuasion Problem with i.i.d. Marginals

In this section we consider the case where there are exponentially many states of nature, but the payoffs of different actions are independently and identically distributed (i.i.d.) according to an explicitly-described marginal distribution. Specifically each state of nature θ is a vector in $\Theta = [m]^n$ for a parameter m, where $\theta_i \in [m]$ is the *type* of action i. Associated with each type $j \in [m]$ is a pair $(\xi_j, \rho_j) \in \mathbb{R}^2$, where ξ_j $[\rho_j]$ is the payoff to the sender [receiver] when the receiver chooses an action with type j. We are given a marginal distribution over types, described by a vector $\boldsymbol{q} = (q_1, ..., q_m) \in \Delta_m$. We assume each action's type is drawn independently according to $\boldsymbol{q}$; specifically, the prior distribution λ on states of nature is given by $\lambda(\theta) = \prod_{i \in [n]} q_{\theta_i}$. For convenience, we let $\boldsymbol{\xi} = (\xi_1, ..., \xi_m) \in \mathbb{R}^m$ denote the vector all possible sender payoffs indexed by type, and similarly $\boldsymbol{\rho} = (\rho_1, ..., \rho_m) \in \mathbb{R}^m$ denotes the vector of all possible receiver payoffs indexed by type. Also for convenience, we sometimes represent a state of nature θ as a $\{0,1\}$ matrix $M^\theta \in \mathbb{R}^{n\times m}$ with precisely a single 1 per row; in particular, for any $i \in [n]$, $j \in [m]$, $M^\theta_{ij} = 1$ if and only if action i has type j. We assume $\boldsymbol{\xi}$, $\boldsymbol{\rho}$, and $\boldsymbol{q}$ — the parameters describing the Bayesian persuasion problem in this setting — are given explicitly.

In this setting the number of states of nature is exponential in the number of actions, and therefore the natural representation of a signaling scheme involves an exponential number of variables. In fact, the natural linear program for the persuasion problem, described in Section 2.1, has an exponential number of both variables and constraints. Nevertheless, as mentioned in Section 2.1 we seek to efficiently implement an optimal (or near optimal) signaling scheme φ as an oracle which takes any state of nature θ as input and samples a signal $\sigma \sim \varphi(\theta)$.

Given a signaling scheme φ with signals $\Sigma = \{\sigma_1, \ldots, \sigma_n\}$, we let $\alpha_i = \sum_\theta \lambda(\theta)\varphi(\theta,\sigma_i)$ denote the probability of sending σ_i. Moreover, we define a matrix $M^{\sigma_i} \in \mathbb{R}^{n\times m}$ representing — absent a normalization factor — the marginal type distributions of the various actions conditioned on σ_i.

$$M^{\sigma_i} = \sum_\theta \lambda(\theta)\varphi(\theta,\sigma_i)M^\theta, \ \forall i \in [n]. \tag{3}$$

Note that each row of $\frac{1}{\alpha_i}M^{\sigma_i}$ is the posterior marginal distribution of the corresponding action conditioned on σ_i. Also note that each row of M^{σ_i} sums up to α_i since each row of M^θ sums up to 1. We call $\mathcal{M} = (M^{\sigma_1}, ..., M^{\sigma_n}) \in \mathbb{R}^{n\times m\times n}$ the *signature* of the signaling scheme φ. Notice that the sender's objective and the receiver's incentive compatibility constraints can both be expressed in term of the signature. In particular, using M_j to denote the j'th row of a matrix M, the incentive compatibility constraints can be stated as $\boldsymbol{\rho} \cdot M_i^{\sigma_i} \geq \boldsymbol{\rho} \cdot M_j^{\sigma_i}$ for all $i, j \in [n]$. By incentivizing the receiver to take action i upon receiving σ_i, the sender's expected utility would be $\sum_{i\in[n]} \boldsymbol{\xi} \cdot M_i^{\sigma_i}$.

We say $\mathcal{M} = (M^{\sigma_1}, ..., M^{\sigma_n}) \in \mathbb{R}^{n\times m\times n}$ is *realizable* if there exists a signaling scheme φ with $\mathcal{M}$ as its signature — i.e. satisfying Equation (3) for all $i \in [n]$. Clearly, not all $\mathcal{M} \in \mathbb{R}^{n\times m\times n}$ are realizable. In fact, the family of realizable signatures constitutes a polytope $\mathcal{P} \subseteq \mathbb{R}^{n\times m\times n}$, which has an extended formulation defined by the following system of exponentially many variables and linear constraints.

$$\begin{array}{ll} M^{\sigma_i} = \sum_\theta \lambda(\theta)\varphi(\theta,\sigma_i)M^\theta, & \text{for } i = 1, \ldots, n. \\ \sum_{i=1}^n \varphi(\theta,\sigma_i) = 1, & \text{for } \theta \in \Theta. \\ \varphi(\theta,\sigma_i) \geq 0, & \text{for } \theta \in \Theta, i \in [n]. \end{array} \tag{4}$$

As a result, the sender's optimization problem can be re-written as the following linear program:

$$\begin{array}{lll} \text{maximize} & \sum_{i=1}^n \boldsymbol{\xi} \cdot M_i^{\sigma_i} & \\ \text{subject to} & \boldsymbol{\rho} \cdot M_i^{\sigma_i} \geq \boldsymbol{\rho} \cdot M_j^{\sigma_i}, & \text{for } i, j \in [n]. \\ & \mathcal{M} = (M^{\sigma_1}, ..., M^{\sigma_n}) \in \mathcal{P} & \end{array} \tag{5}$$

3.1 Symmetry of the Optimal Signaling Scheme

We now show that there always exists a "symmetric" optimal signaling scheme if the action payoffs are i.i.d. Given a signature $\mathcal{M} = (M^{\sigma_1}, ..., M^{\sigma_n})$, it will sometimes be convenient to instead think of it as the set of pairs $\{(M^{\sigma_i}, \sigma_i)\}_{i\in[n]}$.

Definition 3.1. *(Symmetric Signaling Scheme) A signaling scheme φ with signature $\{(M^{\sigma_i}, \sigma_i)\}_{i\in[n]}$ is called* symmetric *if there exist $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^m$ such that $M_i^{\sigma_i} = \boldsymbol{x}$ for any $i \in [n]$ and $M_j^{\sigma_i} = \boldsymbol{y}$ for any $j \neq i$. The pair $(\boldsymbol{x}, \boldsymbol{y}) \in [0,1]^m \times [0,1]^m$ is called the* s-signature *of φ.*

In other words, a symmetric signaling scheme sends each signal with equal probability $||\boldsymbol{x}||_1$, and induces only two different posterior marginals for actions: $\frac{\boldsymbol{x}}{||\boldsymbol{x}||_1}$ for the recommended action, and $\frac{\boldsymbol{y}}{||\boldsymbol{y}||_1}$ for the others. Based on the above definition, we call $(\boldsymbol{x}, \boldsymbol{y})$ *realizable* if there exists a signaling scheme φ with $(\boldsymbol{x}, \boldsymbol{y})$ as its s-signature. The family of realizable s-signatures constitute a polytope $\mathcal{P}_s \subseteq [0,1]^m \times [0,1]^m$, with an extended formulation given below.

$$\begin{array}{ll} \boldsymbol{x} = \sum_\theta \lambda(\theta)\varphi(\theta,\sigma_i)M_i^\theta, & \text{for } i = 1,\ldots,n. \\ \boldsymbol{y} = \sum_\theta \lambda(\theta)\varphi(\theta,\sigma_i)M_j^\theta, & \text{for } j \neq i. \\ \sum_{i=1}^n \varphi(\theta,\sigma_i) = 1, & \text{for } \theta \in \Theta. \\ \varphi(\theta,\sigma_i) \geq 0, & \text{for } \theta \in \Theta, i \in [n]. \end{array} \tag{6}$$

We make two simple observations regarding realizable s-signatures. First, $||\boldsymbol{x}||_1 = ||\boldsymbol{y}||_1 = \frac{1}{n}$ for each $(\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{P}_s$, and this is because both $||\boldsymbol{x}||_1$ and $||\boldsymbol{y}||_1$ equal the probability of each of the n signals. Additionally, since the signature must be consistent with prior marginal distribution $\boldsymbol{q}$, we have

$$\boldsymbol{x} + (n-1)\boldsymbol{y} = \sum_{i=1}^n M_1^{\sigma_i} = \boldsymbol{q}. \tag{7}$$

We show that restricting to symmetric signaling schemes is without loss of generality.

Theorem 3.2. *When the action payoffs are i.i.d., there exists an optimal and incentive-compatible signaling scheme which is symmetric.*

To prove Theorem 3.2, we need two closure properties of optimal signaling schemes — with respect to permutations and convex combinations. We use π to denote a permutation of $[n]$, and let $\mathbb{S}_n$ denote the set of all such permutations. We define the permutation $\pi(\theta)$ of a state of nature $\theta \in [m]^n$ so that $(\pi(\theta))_j = \theta_{\pi(j)}$, and similarly the permutation of a signal σ_i so that $\pi(\sigma_i) = \sigma_{\pi(i)}$. Given a signature $\mathcal{M} = \{(M^{\sigma_i}, \sigma_i)\}_{i\in[n]}$, we define the permuted signature $\pi(\mathcal{M}) = \{(\pi M^{\sigma_i}, \pi(\sigma_i))\}_{i\in[n]}$, where πM denotes applying permutation π to the rows of a matrix M.

Lemma 3.3. *Assume the action payoffs are i.i.d., and let $\pi \in \mathbb{S}_n$ be an arbitrary permutation. If $\mathcal{M}$ is the signature of a signaling scheme φ, then $\pi(\mathcal{M})$ is the signature of the scheme φ_π defined by $\varphi_\pi(\theta) = \pi(\varphi(\pi^{-1}(\theta)))$. Moreover, if φ is incentive compatible and optimal, then so is φ_π.*

Proof. Let $\mathcal{M} = \{(M^\sigma, \sigma)\}_{\sigma\in\Sigma}$ be the signature of φ, as given in the statement of the lemma. We first show that $\pi(\mathcal{M}) = \{(\pi M^\sigma, \pi(\sigma))\}_{\sigma\in\Sigma}$ is realizable as the signature of the scheme φ_π. By definition, it suffices to show that $\sum_\theta \lambda(\theta)\varphi_\pi(\theta,\pi(\sigma))M^\theta = \pi M^\sigma$ for an arbitrary signal $\pi(\sigma)$.

$$\begin{aligned}
\sum_\theta \lambda(\theta)\varphi_\pi(\theta,\pi(\sigma))M^\theta &= \sum_\theta \lambda(\theta)\varphi(\pi^{-1}(\theta),\sigma)M^\theta && \text{(by definition of } \varphi_\pi) \\
&= \pi \sum_{\theta\in\Theta} \lambda(\theta)\varphi(\pi^{-1}(\theta),\sigma)(\pi^{-1}M^\theta) && \text{(by linearity of permutation)} \\
&= \pi \sum_{\theta\in\Theta} \lambda(\theta)\varphi(\pi^{-1}(\theta),\sigma)M^{\pi^{-1}(\theta)} \\
&= \pi \sum_{\theta\in\Theta} \lambda(\pi^{-1}(\theta))\varphi(\pi^{-1}(\theta),\sigma)M^{\pi^{-1}(\theta)} && \text{(Since } \lambda \text{ is i.i.d.)} \\
&= \pi \sum_{\theta'\in\Theta} \lambda(\theta')\varphi(\theta',\sigma)M^{\theta'} && \text{(by renaming } \pi^{-1}(\theta) \text{ to } \theta') \\
&= \pi M^\sigma && \text{(by definition of } M^\sigma)
\end{aligned}$$

Now, assuming φ is incentive compatible, we check that φ_π is incentive compatible by verifying the relevant inequality for its signature.

$$\boldsymbol{\rho} \cdot (\pi M^{\sigma_i})_{\pi(i)} - \boldsymbol{\rho} \cdot (\pi M^{\sigma_i})_{\pi(j)} = \boldsymbol{\rho} \cdot M_i^{\sigma_i} - \boldsymbol{\rho} \cdot M_j^{\sigma_i} \geq 0$$

Moreover, we show that the sender's utility is the same for φ and φ_π, completing the proof.

$$\boldsymbol{\xi} \cdot (\pi M^{\sigma_i})_{\pi(i)} = \boldsymbol{\xi} \cdot (M^{\sigma_i})_i$$

□

Lemma 3.4. *Let $t \in [0,1]$. If $\mathcal{A} = (A^{\sigma_1}, \ldots, A^{\sigma_n})$ is the signature of scheme φ_A, and $\mathcal{B} = (B^{\sigma_1}, \ldots, B^{\sigma_n})$ is the signature of a scheme φ_B, then their convex combination $\mathcal{C} = (C^{\sigma_1}, \ldots, C^{\sigma_n})$ with $C^{\sigma_i} = tA^{\sigma_i} + (1-t)B^{\sigma_i}$ is the signature of the scheme φ_C which, on input θ, outputs $\varphi_A(\theta)$ with probability t and $\varphi_B(\theta)$ with probability $1-t$. Moreover, if φ_A and φ_B are both optimal and incentive compatible, then so is φ_C.*

Proof. This follows almost immediately from the fact that (5) is a linear program, with a convex feasible set and a convex family of optimal solutions. We omit the straightforward details. □

Proof of Theorem 3.2

Proof. Given an optimal and incentive compatible signaling scheme φ with signature $\{(M^{\sigma_i}, \sigma_i)\}_{i\in[n]}$, we show the existence of a symmetric optimal and incentive-compatible scheme of the form in Definition 3.1. According to Lemma 3.3, for $\pi \in \mathbb{S}_n$ the signature $\{(\pi M^{\sigma_i}, \pi(\sigma_i))\}_{i\in[n]}$ — equivalently written as $\{(\pi M^{\sigma_{\pi^{-1}(i)}}, \sigma_i\}_{i\in[n]}$ — corresponds to the optimal incentive compatible scheme φ_π. Invoking Lemma 3.4, the signature

$$\{(A^{\sigma_i}, \sigma_i)\}_{i\in[n]} = \{(\frac{1}{n!} \sum_{\pi\in\mathbb{S}_n} \pi M^{\sigma_{\pi^{-1}(i)}}, \sigma_i)\}_{i\in[n]}$$

also corresponds to an optimal and incentive compatible scheme, namely the scheme which draws a permutation π uniformly at random, then signals according to φ_π.

Observe that the i'th row of the matrix $\pi M^{\sigma_{\pi^{-1}(i)}}$ is the $\pi^{-1}(i)$'th row of the matrix $M^{\sigma_{\pi^{-1}(i)}}$. Expressing $A_i^{\sigma_i}$ as a sum over permutations $\pi \in \mathbb{S}_n$, and grouping the sum by $k = \pi^{-1}(i)$, we can write

$$\begin{aligned}
A_i^{\sigma_i} &= \frac{1}{n!} \sum_{\pi\in\mathbb{S}_n} [\pi M^{\sigma_{\pi^{-1}(i)}}]_i \\
&= \frac{1}{n!} \sum_{\pi\in\mathbb{S}_n} M_{\pi^{-1}(i)}^{\sigma_{\pi^{-1}(i)}} \\
&= \frac{1}{n!} \sum_{k=1}^{n} M_k^{\sigma_k} \cdot \left|\left\{\pi \in \mathbb{S}_n : \pi^{-1}(i) = k\right\}\right| \\
&= \frac{1}{n!} \sum_{k=1}^{n} M_k^{\sigma_k} \cdot (n-1)! \\
&= \frac{1}{n} \sum_{k=1}^{n} M_k^{\sigma_k},
\end{aligned}$$

which does not depend on i. Similarly, the j'th row of the matrix $\pi M^{\sigma_{\pi^{-1}(i)}}$ is the $\pi^{-1}(j)$'th row of the matrix $M^{\sigma_{\pi^{-1}(i)}}$. For $j \neq i$, expressing $A_j^{\sigma_i}$ as a sum over permutations $\pi \in \mathbb{S}_n$, and grouping the sum by $k = \pi^{-1}(i)$ and $l = \pi^{-1}(j)$, we can write

$$\begin{aligned} A_j^{\sigma_i} &= \frac{1}{n!} \sum_{\pi \in \mathbb{S}_n} [\pi M^{\sigma_{\pi^{-1}(i)}}]_j \\ &= \frac{1}{n!} \sum_{\pi \in \mathbb{S}_n} M_{\pi^{-1}(j)}^{\sigma_{\pi^{-1}(i)}} \\ &= \frac{1}{n!} \sum_{k \neq l} M_l^{\sigma_k} \cdot \left| \left\{ \pi \in \mathbb{S}_n : \pi^{-1}(i) = k, \pi^{-1}(j) = l \right\} \right| \\ &= \frac{1}{n!} \sum_{k \neq l} M_l^{\sigma_k} \cdot (n-2)! \\ &= \frac{1}{n(n-1)} \sum_{k \neq l} M_l^{\sigma_k}, \end{aligned}$$

which does not depend on i or j. Let

$$\begin{aligned} \boldsymbol{x} &= \frac{1}{n} \sum_{k=1}^{n} M_k^{\sigma_k}; \\ \boldsymbol{y} &= \frac{1}{n(n-1)} \sum_{k \neq l} M_l^{\sigma_k}. \end{aligned}$$

The signature $\{(A^{\sigma_i}, \sigma_i)\}_{i \in [n]}$ therefore describes an optimal, incentive compatible, and symmetric scheme with s-signature $(\boldsymbol{x}, \boldsymbol{y})$. ☐

3.2 A Simple $(1 - \frac{1}{e})$-Approximate Scheme

Theorem 3.2 shows that in order to compute an optimal signaling scheme, it suffices to optimize over symmetric signaling schemes. This permits re-writing our optimization problem (5) as follows, with variables $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^m$.

$$\begin{array}{ll} \text{maximize} & n\boldsymbol{\xi} \cdot \boldsymbol{x} \\ \text{subject to} & \boldsymbol{\rho} \cdot \boldsymbol{x} \geq \boldsymbol{\rho} \cdot \boldsymbol{y} \\ & \boldsymbol{x} + (n-1)\boldsymbol{y} = \boldsymbol{q} \\ & ||\boldsymbol{x}||_1 = \frac{1}{n} \\ & \boldsymbol{x}, \boldsymbol{y} \geq 0 \\ & (\boldsymbol{x}, \boldsymbol{y}) \in \mathcal{P}_s \end{array} \tag{8}$$

Recall that the constraints $||\boldsymbol{x}||_1 = \frac{1}{n}$ and $\boldsymbol{x} + (n-1)\boldsymbol{y} = \boldsymbol{q}$ are satisfied by symmetric schemes $\mathcal{P}_s$. Nevertheless, we list those constraints separately to aid in defining our relaxation of this optimization problem.

It is non-trivial to solve optimization problem (8) directly, since the last constraint involves an extended formulation with exponentially many variables and constraints, as described in Section 3.1. In this section we explicitly construct a signaling scheme by relaxing (8) by removing its last constraint, then constructing a scheme which approximately matches the optimal value of the relaxed optimization problem. The constructed signaling scheme has a particularly simple and instructive structure, and serves as a $(1 - \frac{1}{e})$ multiplicative approximation to the optimal signaling scheme assuming all the payoffs are nonnegative. Specifically, our scheme can either be interpreted as the combination of n independent signaling schemes, one per action, or as a direct incentive-compatible scheme as described in Section 2.1.

Theorem 3.5. *Algorithm 1 runs in poly(m, n) time, and serves as a $(1-\frac{1}{e})$-approximate signaling scheme for the Bayesian Persuasion problem with n i.i.d. actions, m types, and nonnegative payoffs.*

Proof. We start by relaxing Linear Program (8) by removing its last constraint. In particular, the following linear program has polynomially many variables and constraints, and therefore can be solved efficiently.

$$\begin{array}{ll} \text{maximize} & n\boldsymbol{\xi} \cdot \boldsymbol{x} \\ \text{subject to} & \boldsymbol{\rho} \cdot \boldsymbol{x} \geq \boldsymbol{\rho} \cdot \boldsymbol{y} \\ & \boldsymbol{x} + (n-1)\boldsymbol{y} = \boldsymbol{q} \\ & ||\boldsymbol{x}||_1 = \frac{1}{n} \\ & \boldsymbol{x}, \boldsymbol{y} \geq 0 \end{array} \tag{9}$$

Let $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ be the optimal solution to LP (9). We first consider each action separately. Fix an action i. Since $\boldsymbol{x}^* + (n-1)\boldsymbol{y}^* = \boldsymbol{q}$, we can construct two signals **HIGH** and **LOW** such that **HIGH** is sent with probability $||\boldsymbol{x}^*||_1 = \frac{1}{n}$ and induces the posterior $n\boldsymbol{x}^*$, while **LOW** is sent with probability $\frac{n-1}{n}$ and induces posterior $n\boldsymbol{y}^*$. This can be achieved as follows: if the action has type j, sample **HIGH** with probability x^*_j/q_j and **LOW** otherwise. Note that our **HIGH** / **LOW** terminology is justified by the incentive constraint $\boldsymbol{\rho} \cdot \boldsymbol{x}^* \geq \boldsymbol{\rho} \cdot \boldsymbol{y}^*$; i.e., the receiver has greater posterior expected utility from the action when the signal is **HIGH** than when it is **LOW**.

Zooming out, the signaling scheme we construct simply announces the state **HIGH** or **LOW** of each action independently. In particular, we construct an n-dimensional *binary* signal $\sigma \in \{\textbf{HIGH}, \textbf{LOW}\}^n$, where the i'th *component* of σ is the result of applying the above **HIGH** / **LOW** scheme to action i independently. Details are in Algorithm 1. Note that, as described, our signaling scheme involves 2^n signals, and is therefore not a direct scheme in the sense described in Section 2.1. However, as will become apparent shortly, an equivalent direct scheme would simply recommend an arbitrary action with state **HIGH** when at least one action is **HIGH**, and otherwise recommends an arbitrary action.

Algorithm 1 Independent Signaling Scheme

Input: sender payoff vector $\boldsymbol{\xi}$, receiver payoff vector $\boldsymbol{\rho}$, prior distribution $\boldsymbol{q}$

Input: an observed state of nature $\theta \in [m]^n$

Output: an n-dimensional binary signal $\sigma \in \{\textbf{HIGH}, \textbf{LOW}\}^n$

1: Compute the optimal solution $(\boldsymbol{x}^*, \boldsymbol{y}^*)$ of Linear Program (9).
2: For each action i independently, set component signal o_i to **HIGH** with probability $\frac{x^*_{\theta_i}}{q_{\theta_i}}$ and to **LOW** otherwise, where θ_i is the type of action i in the input state θ.
3: Return $\sigma = (o_1, ..., o_n)$.

To analyze the signaling scheme in Algorithm 1 let us think from the receiver's perspective. Upon receiving a binary signal $\sigma \in \{\textbf{HIGH}, \textbf{LOW}\}^n$, the receiver infers a posterior distribution $n\boldsymbol{x}^*$ for *every* action with component signal **HIGH**, and a posterior $n\boldsymbol{y}^*$ for every action with component signal **LOW**. This is simply a consequence of the independence of the actions and independence of our component signaling schemes.

When at least one component of our signal σ is **HIGH**, the receiver is indifferent between the corresponding **HIGH** actions; choosing such an action results in sender utility $n\boldsymbol{\xi} \cdot \boldsymbol{x}^* = Opt(LP\ (9))$. Since each action's component signal is **HIGH** with probability $1/n$, and since action payoffs and component signaling schemes are independent, the probability of at least one **HIGH** component signal is $1-(1-\frac{1}{n})^n \geq 1-\frac{1}{e}$. When all components of σ are **LOW**, the sender's utility is at least 0. Therefore, Algorithm 1 extracts at least a $(1-\frac{1}{e})$ fraction of $Opt(LP\ (9))$. This concludes our proof. $\square$

Remark 3.6. *Algorithm 1 signals independently for each action. This conveys an interesting conceptual message. That is, even though the optimal signaling scheme might induce posterior beliefs which correlate different actions, it is nevertheless true that signaling for each action independently yields an approximately optimal signaling scheme. As a consequence, collaborative persuasion by multiple parties (the senders), each of whom observes the payoff of one or more actions, is a task that can be parallelized, requiring no coordination when actions are identical and independent and only an approximate solution is sought. We leave open the question of whether this is possible when action payoffs are independently but not identically distributed.*

3.3 An Optimal Scheme

In this section, we exhibit a polynomial-time algorithm implementing an optimal signaling scheme when the action payoffs are i.i.d, with marginal distribution described explicitly. We make use of the ellipsoid method and a connection between symmetric signaling schemes and single-item auctions with i.i.d. bidders. In particular, we show that there is a one-to-one correspondence between symmetric signaling schemes and (a subset of) symmetric reduced forms of single-item auctions with i.i.d. bidders, defined as follows.

Definition 3.7 ([10])**.** *Consider a single-item auction setting with n i.i.d. bidders and m types for each bidder, where each bidder's type is distributed according to $\boldsymbol{q} \in \Delta_m$. An* allocation rule *is a randomized function A mapping a type profile $\theta \in [m]^n$ to a winner $A(\theta) \in [n] \cup \{*\}$, where $*$ denotes not allocating the item. We assume entries of θ are drawn i.i.d from $\boldsymbol{q}$, and say the allocation rule has* symmetric reduced form *$\boldsymbol{\tau} \in [0,1]^m$ if for each bidder $i \in [n]$ and type $j \in [m]$, τ_j is the conditional probability of i receiving the item given she has type j.*

When $\boldsymbol{q}$ is clear from context, we say $\boldsymbol{\tau}$ is *realizable* if there exists an allocation rule with $\boldsymbol{\tau}$ as its symmetric reduced form. We say an algorithm *implements* an allocation rule A if it takes as input θ, and samples $A(\theta)$. As usual, we use λ to denote the product distribution induced by $\boldsymbol{q}$ where convenient.

Theorem 3.8. *Consider the Bayesian Persuasion problem with n i.i.d. actions and m types, with parameters $\boldsymbol{q} \in \Delta_m$, $\boldsymbol{\xi} \in \mathbb{R}^m$, and $\boldsymbol{\rho} \in \mathbb{R}^m$ given explicitly. An optimal signaling scheme can be implemented in* $\text{poly}(m,n)$ *time.*

Theorem 3.8 is a consequence of the following set of lemmas. The proofs of these lemmas involve Border's characterization of the family of realizable reduced forms of a single item auction [10, 9], as well as its algorithmic properties (see e.g. [12, 2]).

Lemma 3.9. *Let $(\boldsymbol{x}, \boldsymbol{y}) \in [0,1]^m \times [0,1]^m$, and define $\boldsymbol{\tau} = (\frac{x_1}{q_1}, ..., \frac{x_m}{q_m})$. The pair $(\boldsymbol{x}, \boldsymbol{y})$ is a realizable s-signature if and only if (a) $||\boldsymbol{x}||_1 = \frac{1}{n}$, (b) $\boldsymbol{x} + (n-1)\boldsymbol{y} = \boldsymbol{q}$, and (c) $\boldsymbol{\tau}$ is a realizable symmetric reduced form of an allocation rule with n i.i.d. bidders, m types, and type distribution $\boldsymbol{q}$. Moreover, assuming $\boldsymbol{x}$ and $\boldsymbol{y}$ satisfy (a), (b) and (c), and given black-box access to an allocation rule*

A with symmetric reduced form $\boldsymbol{\tau}$, a signaling scheme with s-signature $(\boldsymbol{x}, \boldsymbol{y})$ can be implemented in $\text{poly}(n, m)$ *time.*

Proof. For the "only if" direction, $||\boldsymbol{x}||_1 = \frac{1}{n}$ and $\boldsymbol{x} + (n-1)\boldsymbol{y} = \boldsymbol{q}$ were established in Section 3.1. To show that $\boldsymbol{\tau}$ is a realizable symmetric reduced form for an allocation rule, let φ be a signaling scheme with s-signature $(\boldsymbol{x}, \boldsymbol{y})$. Recall from the definition of an s-signature that, for each $i \in [n]$, signal σ_i has probability $1/n$, and $n\boldsymbol{x}$ is the posterior distribution of action i's type conditioned on signal σ_i. Now consider the following allocation rule: Given a type profile $\theta \in [m]^n$ of the n bidders, allocate the item to bidder i with probability $\varphi(\theta, \sigma_i)$ for any $i \in [n]$. By Bayes rule,

$$\begin{aligned}\mathbf{Pr}[i \text{ gets item} | i \text{ has type } j] &= \mathbf{Pr}[i \text{ has type } j | i \text{ gets item}] \cdot \frac{\mathbf{Pr}[i \text{ gets item}]}{\mathbf{Pr}[i \text{ has type } j]} \\ &= nx_j \cdot \frac{1/n}{q_j} = \frac{x_j}{q_j}\end{aligned}$$

Therefore $\boldsymbol{\tau}$ is indeed the reduced form of the described allocation rule.

For the "if" direction, let $\boldsymbol{\tau}$, $\boldsymbol{x}$, and $\boldsymbol{y}$ be as in the statement of the lemma, and consider an allocation rule A with symmetric reduced form $\boldsymbol{\tau}$. Observe that A always allocates the item, since for each player $i \in [n]$ we have $\mathbf{Pr}[i \text{ gets the item}] = \sum_{j=1}^m q_j \tau_j = \sum_{j=1}^m x_j = \frac{1}{n}$. We define the direct signaling scheme φ_A by $\varphi_A(\theta) = \sigma_{A(\theta)}$. Let $\mathcal{M} = (M^{\sigma_1}, \ldots, M^{\sigma_n})$ be the signature of φ_A. Recall that, for $\theta \sim \lambda$ and arbitrary $i \in [n]$ and $j \in [m]$, $M^{\sigma_i}_{ij}$ is the probability that $\varphi_A(\theta) = \sigma_i$ and $\theta_i = j$; by definition, this equals the probability that A allocates the item to player i and her type is j, which is $\tau_j q_j = x_j$. As a result, the signature $\mathcal{M}$ of φ_A satisfies $M^{\sigma_i}_i = \boldsymbol{x}$ for every action i. If φ_A were symmetric, we would conclude that its s-signature is $(\boldsymbol{x}, \boldsymbol{y})$ since every s-signature $(\boldsymbol{x}, \boldsymbol{y}')$ must satisfy $\boldsymbol{x} + (n-1)\boldsymbol{y}' = \boldsymbol{q}$ (see Section 3.1). However, this is not guaranteed when the allocation rule A exhibits some asymmetry. Nevertheless, φ_A can be "symmetrized" into a signaling scheme φ'_A which first draws a random permutation $\pi \in \mathbb{S}_n$, and signals $\pi(\varphi_A(\pi^{-1}(\theta)))$. That φ'_A has s-signature $(\boldsymbol{x}, \boldsymbol{y})$ follows a similar argument to that used in the proof of Theorem 3.2, and we therefore omit the details here.

Finally, observe that the description of φ'_A above is constructive assuming black-box access to A, with runtime overhead that is polynomial in n and m. ◻

Lemma 3.10. *An optimal realizable s-signature, as described by linear program* (8)*, can be computed in* $\text{poly}(n, m)$ *time.*

Proof. By Lemma 3.9, we can re-write LP (8) as follows:

$$\begin{array}{ll} \text{maximize} & n\boldsymbol{\xi} \cdot \boldsymbol{x} \\ \text{subject to} & \boldsymbol{\rho} \cdot \boldsymbol{x} \geq \boldsymbol{\rho} \cdot \boldsymbol{y} \\ & \boldsymbol{x} + (n-1)\boldsymbol{y} = \boldsymbol{q} \\ & ||\boldsymbol{x}||_1 = \frac{1}{n} \\ & (\frac{x_1}{q_1}, \ldots., \frac{x_m}{q_m}) \text{ is a realizable symmetric reduced form} \end{array} \tag{10}$$

From [10, 9, 12, 2], we know that the family of all the realizable symmetric reduced forms constitutes a polytope, and moreover that this polytope admits an efficient separation oracle. The runtime of this oracle is polynomial in m and n, and as a result the above linear program can be solved in $poly(n, m)$ time using the Ellipsoid method. ◻

Lemma 3.11. *(See [12, 2]) Consider a single-item auction setting with n i.i.d. bidders and m types for each bidder, where each bidder's type is distributed according to $\boldsymbol{q} \in \Delta_m$. Given a realizable* *symmetric reduced form* $\boldsymbol{\tau} \in [0,1]^m$*, an allocation rule with reduced form* $\boldsymbol{\tau}$ *can be implemented in* $\mathrm{poly}(n, m)$ *time.*

We leave open whether the optimal signaling scheme can be computed efficiently when actions are independent but not identically distributed. The results of this section exploited crucially the symmetry of the optimal scheme, and extending these results to non-symmetric actions appears to be nontrivial.

4 Dispersion Testing and Inversion

Before tackling the general Bayesian Persuasion problem we define two abstract problems on vector spaces which we term *dispersion testing* and *dispersion inversion*, and connect them to signaling. For two distributions $\mathcal{A}$ and $\mathcal{B}$ supported on the same vector space $\mathcal{V}$, we say $\mathcal{A}$ is a *dispersion* of $\mathcal{B}$ (a.k.a. $\mathcal{A}$ is a *mean-preserving spread* of $\mathcal{B}$) if, for $\boldsymbol{b} \sim \mathcal{B}$, there is a random vector $\boldsymbol{z} \in \mathcal{V}$ potentially correlated with $\boldsymbol{b}$ such that $\boldsymbol{b} + \boldsymbol{z} \sim \mathcal{A}$ and $\mathbf{E}[\boldsymbol{z}|\boldsymbol{b}] = \mathbf{0}$. Equivalently, $\mathcal{A}$ is a dispersion of $\mathcal{B}$ if and only if the random variables $\boldsymbol{a} \sim \mathcal{A}$ and $\boldsymbol{b} \sim \mathcal{B}$ can be coupled[3] so that $\mathbf{E}[\boldsymbol{a}|\boldsymbol{b}] = \boldsymbol{b}$. We note that dispersion defines a partial order on distributions. When the vector space $\mathcal{V}$ is equipped with a norm $||.||$, we say $\mathcal{A}$ is an ϵ*-dispersion* of $\mathcal{B}$ if there is a distribution $\mathcal{C}$ such that $\mathcal{A}$ is a dispersion of $\mathcal{C}$, and moreover $\mathcal{C}$ is within earthmover distance ϵ of $\mathcal{B}$ with respect to the norm $||.||$.

The dispersion relation is related in a straightforward way to signaling. We think of the state of nature θ as a vector in $\mathcal{V}$ with distribution $\mathcal{A}$, and $\mathcal{B}$ as the desired distribution of posterior expectations. A signaling scheme $\varphi : \mathcal{V} \to \Sigma$ inducing the desired distribution of posterior expectations exists if and only if $\mathcal{A}$ is a dispersion of $\mathcal{B}$.

Lemma 4.1. *Let $\mathcal{A}$ and $\mathcal{B}$ be two distributions on a vector space $\mathcal{V}$. The following two statements are equivalent*

a. $\mathcal{A}$ is a dispersion of $\mathcal{B}$.

b. There is a set Σ and a randomized function $\varphi : \mathcal{V} \to \Sigma$ such that, when $\boldsymbol{a} \sim \mathcal{A}$, it holds that $\mathbf{E}[\boldsymbol{a}|\varphi(\boldsymbol{a})] \sim \mathcal{B}$.

Proof. For the first direction assume $\mathcal{A}$ is a dispersion of $\mathcal{B}$, and let $a \sim \mathcal{A}$ and $b \sim \mathcal{B}$ be coupled so that $\mathbf{E}[\boldsymbol{a}|\boldsymbol{b}] = \boldsymbol{b}$. Let $\Sigma = \mathcal{V}$ and $\varphi(\boldsymbol{a}) = \boldsymbol{b}$. We have $\mathbf{E}[\boldsymbol{a}|\varphi(\boldsymbol{a})] = \mathbf{E}[\boldsymbol{a}|\boldsymbol{b}] = \boldsymbol{b} \sim \mathcal{B}$.

For the second direction, take φ and Σ as given by (b). Let $\boldsymbol{a} \sim \mathcal{A}$, and let $\boldsymbol{b} = \mathbf{E}[\boldsymbol{a}|\varphi(\boldsymbol{a})]$. Note that $\boldsymbol{b} \sim \mathcal{B}$, as given by (b). Now, it remains to show that $\boldsymbol{a}$ and $\boldsymbol{b}$ are coupled to prove dispersion: $\mathbf{E}[\boldsymbol{a}|\boldsymbol{b}] = \mathbf{E}[\boldsymbol{a}|\,\mathbf{E}[\boldsymbol{a}|\varphi(\boldsymbol{a})]] = \mathbf{E}[\boldsymbol{a}|\varphi(\boldsymbol{a})] = \boldsymbol{b}$, as needed.[4] □

When a signaling scheme $\varphi : \mathcal{V} \to \Sigma$ satisfies (b) from Lemma 4.1, we say that φ *forges* $\mathcal{B}$ from $\mathcal{A}$. We can think of φ as *inverting* the mean-preserving spread applied to $\mathcal{B}$ to form $\mathcal{A}$. More generally, when φ forges from $\mathcal{A}$ a distribution $\mathcal{C}$ within earthmover distance ϵ from $\mathcal{B}$, we say φ ϵ*-forges* $\mathcal{B}$ from $\mathcal{A}$. From Lemma 4.1, $\mathcal{B}$ can be ϵ-forged from $\mathcal{A}$ if and only if $\mathcal{A}$ is an ϵ-dispersion of $\mathcal{B}$.

An instance of either the testing or inversion problems is given via two distributions $\mathcal{A}$ and $\mathcal{B}$ supported on a normed vector space, and two parameters $\epsilon_1 \geq 0$ and $\epsilon_2 > 0$. For testing, our goal is to distinguish the case in which $\mathcal{A}$ is an ϵ_1-dispersion of $\mathcal{B}$ from that in which $\mathcal{A}$ is not an

[3]Recall that a coupling of two random variables $\boldsymbol{a}$ and $\boldsymbol{b}$ is a distribution over pairs $(\boldsymbol{a}', \boldsymbol{b}')$ with the property that $\boldsymbol{a}' \stackrel{d}{=} \boldsymbol{a}$ and $b' \stackrel{d}{=} \boldsymbol{b}$, where $\stackrel{d}{=}$ denotes equality in distribution.

[4]We use the fact that for two random variables $\boldsymbol{x}$ and $\boldsymbol{y}$ on a vector space, it is the case that $\mathbf{E}[\boldsymbol{x}|\,\mathbf{E}[\boldsymbol{x}|\boldsymbol{y}]] = \mathbf{E}[\boldsymbol{x}|\boldsymbol{y}]$. This is not too hard to show using the tower property of conditional expectations, though we omit the details.

$$
\begin{array}{lll}
\text{find} & x & \\
\text{subject to} & \sum_{\boldsymbol{b} \in \mathbf{supp}(\mathcal{B})} x(\boldsymbol{a}, \boldsymbol{b}) = \mathbf{Pr}_{\mathcal{A}}[\boldsymbol{a}], & \text{for } \boldsymbol{a} \in \mathbf{supp}(\mathcal{A}). \\
& \sum_{\boldsymbol{a} \in \mathbf{supp}(\mathcal{A})} x(\boldsymbol{a}, \boldsymbol{b}) = \mathbf{Pr}_{\mathcal{B}}[\boldsymbol{b}], & \text{for } \boldsymbol{b} \in \mathbf{supp}(\mathcal{B}). \\
& \sum_{\boldsymbol{a} \in \mathbf{supp}(\mathcal{A})} x(\boldsymbol{a}, \boldsymbol{b}) \cdot \boldsymbol{a} = \mathbf{Pr}_{\mathcal{B}}[\boldsymbol{b}] \cdot \boldsymbol{b}, & \text{for } \boldsymbol{b} \in \mathbf{supp}(\mathcal{B}). \\
& x(\boldsymbol{a}, \boldsymbol{b}) \geq 0, & \text{for } \boldsymbol{a} \in \mathbf{supp}(\mathcal{A}), \boldsymbol{b} \in \mathbf{supp}(\mathcal{B}).
\end{array}
$$

Figure 1: Dispersion Testing Linear Program when $\epsilon_1 = 0$

$(\epsilon_1 + \epsilon_2)$-dispersion of $\mathcal{B}$, and in the latter case output a dual certificate — whose form we will describe shortly — to that effect. For inversion, if our input is such that $\mathcal{A}$ is an ϵ_1-dispersion of $\mathcal{B}$, our goal is efficient implementation of signaling scheme φ which $(\epsilon_1 + \epsilon_2)$-forges $\mathcal{B}$ from $\mathcal{A}$. Specifically, we implement an oracle which takes as input a query $\boldsymbol{a} \in \mathcal{V}$ and outputs a sample from $\varphi(\boldsymbol{a})$, for some signaling scheme φ which $(\epsilon_1 + \epsilon_2)$-forges $\mathcal{B}$ from $\mathcal{A}$.

We restrict attention to the vector space $\mathbb{R}^d$ equipped with an arbitrary norm, which we denote by $||.||$. In other words, we work in the Banach Space $(\mathbb{R}^d, ||.||)$. We also use $||.||_*$ to denote the dual of our chosen norm.[5] Additionally, we assume that both $\mathcal{A}$ and $\mathcal{B}$ are supported on the unit ball with respect to norm $||.||$ (e.g. $[-1,1]^d$ in the case of the L^∞ norm). When both $\mathcal{A}$ and $\mathcal{B}$ have finite support and are represented explicitly, dispersion testing with $\epsilon = 0$ reduces to a simple linear feasibility problem, illustrated in Figure 1. Notably, a solution $\boldsymbol{x}$ to the feasibility problem describes a signaling scheme forging $\mathcal{B}$ from $\mathcal{A}$: on input $\boldsymbol{a}$, the scheme outputs $\boldsymbol{b}$ with probability $\frac{x(\boldsymbol{a},\boldsymbol{b})}{\Pr_{\mathcal{A}}[\boldsymbol{a}]}$.

Our case of interest is when $\mathcal{A}$ is implicitly described via a sampling oracle, while $\mathcal{B}$ is represented explicitly as a pair $(B, \boldsymbol{q})$ with $B \in \mathbb{R}^{d \times k}$ and $\boldsymbol{q} \in \Delta_k$, where the columns of the matrix B correspond to the support of $\mathcal{B}$, and $\boldsymbol{q}$ is a vector in the probability simplex assigning probabilities to the columns of B. In this case, we design a BPP algorithm for our testing problem which takes in parameters $\epsilon_1 \geq 0$ and $\epsilon_2, \delta > 0$, runs in time $\text{poly}(d, k, \frac{1}{\epsilon_2}, \log(\frac{1}{\delta}))$, and with probability $1 - \delta$ distinguishes between inputs in which $\mathcal{A}$ is an ϵ_1-dispersion of $\mathcal{B}$ and those in which $\mathcal{A}$ is not an $(\epsilon_1 + \epsilon_2)$-dispersion of $\mathcal{B}$; in the latter case, our algorithm outputs a dual certifier to the fact that $\mathcal{A}$ is not an ϵ_1-dispersion of $\mathcal{B}$, of the form $(W, \boldsymbol{z}) \in \mathbb{R}^{d \times k} \times \mathbb{R}^k$. For the inversion problem, our algorithm takes in the parameter $\epsilon_2 > 0$ and, assuming $\mathcal{A}$ is an ϵ_1-dispersion of $\mathcal{B}$ for some $\epsilon_1 \geq 0$, implements a signaling scheme which $(\epsilon_1 + \epsilon_2)$-forges $\mathcal{B}$ from $\mathcal{A}$ in time $\text{poly}(d, k, \frac{1}{\epsilon_2})$.

Our dispersion-testing algorithm is fairly simple, and proceeds as follows. The algorithm takes polynomially many samples $\widetilde{\mathcal{A}}$ from $\mathcal{A}$, then uses a convex program to test whether the empirical distribution $\widetilde{\mathcal{A}}$ is an $(\epsilon_1 + \frac{1}{2}\epsilon_2)$-dispersion of $\mathcal{B}$, using a dual of the convex program to certify negative answers. The inversion algorithm builds on the same convex program: given a query $\boldsymbol{a} \in \mathcal{V}$, it takes polynomially-many samples $\widetilde{\mathcal{A}}$ from $\mathcal{A}$ and uses the convex program to compute a description of signaling scheme φ approximately forging $\mathcal{B}$ from $\widetilde{\mathcal{A}} \cup \{\boldsymbol{a}\}$ — the empirical distribution augmented with the query. The inversion algorithm then signals $\varphi(\boldsymbol{a})$. We describe the details in Algorithms 2 and 3.

The correctness of our algorithms hinges on a primal/dual characterization of optimal dispersion, which we state first in the following two Lemmas.

Lemma 4.2. *The optimal value of convex program* (11) *is the minimum ϵ such that $\mathcal{A} = (A, \boldsymbol{p})$ is an ϵ-dispersion of $\mathcal{B} = (B, \boldsymbol{q})$. Moreover, the optimal solution $\widetilde{B}$ is such that $\mathcal{A}$ is a dispersion of $\widetilde{\mathcal{B}} = (\widetilde{B}, \boldsymbol{q})$, and $\text{EMD}(\widetilde{\mathcal{B}}, \mathcal{B}) = \epsilon$ via the transport map $\boldsymbol{b}_j \to \widetilde{\boldsymbol{b}}_j$.*

[5] Recall that the dual of a norm $||.||$ in $\mathbb{R}^d$ is defined as follows: $||\boldsymbol{x}||_* = \max\{\boldsymbol{x} \cdot \boldsymbol{y} : ||\boldsymbol{y}|| = 1\}$, where $\boldsymbol{x} \cdot \boldsymbol{y}$ denotes the standard inner product of $\boldsymbol{x}$ and $\boldsymbol{y}$.

Algorithm 2 Dispersion Testing Procedure

Parameter: $\epsilon_1 \geq 0$, $\epsilon_2, \delta > 0$

Input: Sample oracle for $\mathcal{A}$

Input: Explicit Distribution $\mathcal{B}$ with support $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_k$ and probabilities $q_1, \ldots, q_k$

Output: **ACCEPT** or $(\textbf{REJECT}, (W, \boldsymbol{z}))$, where $W \in \mathbb{R}^{d \times k}$ and $\boldsymbol{z} \in \mathbb{R}^k$

Sample $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_\ell \sim \mathcal{A}$, where $\ell = 288(\frac{1}{\epsilon_2})^2(\log \frac{2}{\delta} + k(d+1) \log \frac{33k}{\epsilon_2})$.

Solve optimization problem (11) and its dual (12) on $\widetilde{\mathcal{A}} := Uniform(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_\ell)$ and $\mathcal{B}$.

if The optimal value of (11) is greater than $\epsilon_1 + \epsilon_2/2$ **then**

  return **REJECT** and the dual solution $(W, \boldsymbol{z})$;

else

  return **ACCEPT**.

end if

Algorithm 3 Dispersion Inversion Procedure

Parameter: Parameter $\epsilon_2 > 0$

Parameter: Sample oracle for $\mathcal{A}$

Parameter: Explicit Distribution $\mathcal{B}$ with support $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_k$ and probabilities $q_1, \ldots, q_k$

Input: Vector $\boldsymbol{a} \in \mathbb{R}^d$

Output: Vector $\boldsymbol{b} \in \mathbb{R}^d$

Sample $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_\ell \sim \mathcal{A}$, where $\ell = 288(\frac{1}{\epsilon_2})^2 k(d+2) \log \frac{33k}{\epsilon_2} = \mathrm{poly}(d, k, \frac{1}{\epsilon_2})$.

Solve (11) on $uniform(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_\ell, \boldsymbol{a})$ and $\mathcal{B}$ to get solution $\boldsymbol{x}$.

return $\boldsymbol{b}_j \in supp(\mathcal{B})$ with probability proportional to $x(\ell+1, j)$.

$$
\begin{array}{lll}
\text{minimize} & \sum_{j=1}^{k} q_j ||\boldsymbol{b}_j - \widetilde{\boldsymbol{b}}_j|| & \\
\text{subject to} & \sum_{j=1}^{k} x(i,j) = p_i, & \text{for } i = 1, \ldots, \ell. \\
 & \sum_{i=1}^{\ell} x(i,j) = q_j, & \text{for } j = 1, \ldots, k. \\
 & \sum_{i=1}^{\ell} x(i,j) \cdot \boldsymbol{a}_i = q_j \cdot \widetilde{\boldsymbol{b}}_j, & \text{for } j = 1, \ldots, k. \\
 & x(i,j) \geq 0, & \text{for } i = 1, \ldots, \ell, j = 1, \ldots, k.
\end{array}
\tag{11}
$$

Figure 2: Approximate Dispersion Testing. Inputs are $\mathcal{A} = (A, \boldsymbol{p})$ and $\mathcal{B} = (B, \boldsymbol{q})$, with support sizes ℓ and k respectively. Variables are $x \in \mathbb{R}^{\ell \times k}$ and $\widetilde{\boldsymbol{b}}_1, \ldots, \widetilde{\boldsymbol{b}}_k \in \mathbb{R}^d$

$$
\begin{array}{lll}
\text{maximize} & \sum_{j=1}^{k} q_j(\boldsymbol{w}_j \cdot \boldsymbol{b}_j - 2z_j) - \mathbf{E}_{\boldsymbol{a} \sim \mathcal{A}} \left[\max_{j=1}^{k} (\boldsymbol{w}_j \cdot \boldsymbol{a} - 2z_j) \right] & \\
\text{subject to} & ||\boldsymbol{w}_j||_* \leq 1, & \text{for } j = 1, \ldots, k. \\
 & |z_j| \leq 1, & \text{for } j = 1, \ldots, k.
\end{array}
\tag{12}
$$

Figure 3: Dual formulation of Approximate Dispersion Testing. Inputs are distributions $\mathcal{A}$ and $\mathcal{B}$, where $\mathcal{B} = (B, \boldsymbol{q})$ has support size k. Variables are $W = [\boldsymbol{w}_1, \ldots, \boldsymbol{w}_k] \in \mathbb{R}^{d \times k}$ and $z_1, \ldots, z_k \in \mathbb{R}$.

Lemma 4.3. *Convex programs* (11) *and* (12) *are duals, and strong duality holds.*

The proofs of these Lemmas build on a number of structural properties of optimal dispersion which we develop in Section 4.1, and are hence deferred to Section 4.2. We note that, unlike the standard Lagrange dual, our dual convex program is distinguished by the fact that dual variables are constrained in norm, and hence Lemma 4.3 requires a specialized proof which constructs a 2-player game and invokes Sion's minimax theorem.

Using the above Lemmas, correctness and runtime for the two algorithms is summarized by the following theorems.

Theorem 4.4. *Fix the parameters $\epsilon_1, \epsilon_2, \delta$ of Algorithm 2. Let $\mathcal{A}$ and $\mathcal{B} = (B, \boldsymbol{q})$ denote arbitrary input distributions to Algorithm 2, and let k denote the support size of $\mathcal{B}$. The following hold for the output of Algorithm 2 with probability at least $1 - \delta$.*

1. *If $\mathcal{A}$ is an ϵ_1-dispersion of $\mathcal{B}$, then the algorithm accepts.*

2. *if $\mathcal{A}$ is not an $(\epsilon_1 + \epsilon_2)$-dispersion of $\mathcal{B}$, then the algorithm rejects.*

3. *If the algorithm rejects and outputs a certificate (W, z), then*

$$\sum_{j=1}^{k} q_j(\boldsymbol{w}_j \cdot \boldsymbol{b}_j - 2z_j) > \sum_{j=1}^{k} r_j(\boldsymbol{w}_j \cdot \boldsymbol{c}_j - 2z_j) \tag{13}$$

for every support-k distribution $\mathcal{C} = (C, \boldsymbol{r})$ such that $\mathcal{A}$ is an ϵ_1-dispersion of $\mathcal{C}$.

Moreover, Algorithm 2 can be implemented to run in time $\mathrm{poly}(d, k, \frac{1}{\epsilon_2}, \log(\frac{1}{\delta}))$ *when* $||.|| = ||.||_p$ *for some* $p \geq 1$.

Proof. First, we show that the optimal value of (11) converges rapidly in the number of samples — specifically, if we sample $\ell = 288(\frac{1}{\epsilon_2})^2(\log \frac{2}{\delta} + k(d+1)\log \frac{33k}{\epsilon_2}) = \mathrm{poly}(d, k, \frac{1}{\epsilon_2}, \log(\frac{1}{\delta}))$ times from $\mathcal{A}$ and solve (11) using the empirical distribution $\widetilde{\mathcal{A}}$ in lieu of $\mathcal{A}$, the empirical optimal value we get is within $\epsilon_2/2$ of the ground truth with probability at least $1 - \delta$. Invoking Lemma 4.3, we show this by analyzing the convergence of the dual (12). Note that the objective function of (12) has range $[-3, 3]$, and moreover it is Lipschitz continuous in each of its variables — specifically, it is 2-Lipschitz continuous in each of $z_1, \ldots, z_k \in \mathbb{R}$ with respect to the usual metric on $\mathbb{R}$, and it is 1-Lipschitz continuous in each of $\boldsymbol{w}_1, \ldots, \boldsymbol{w}_k \in \mathbb{R}^d$ with respect to the metric induced by the norm $||.||_*$.

By standard Hoeffding bounds and boundedness of the range of objective values, the empirical objective value for a fixed dual solution $(\widehat{W}, \widehat{\boldsymbol{z}})$ is accurate to within $\epsilon_2/4$ but for a probability of $2\exp(-\frac{1}{288}(\epsilon_2)^2\ \ell)$. The remainder of the proof follows from Lipschitz continuity of the objective function and boundedness of the feasible set, using a standard meshing argument and the union bound. We omit the (fairly standard) details, though mention that our meshing argument uses the fact that for every d-dimensional Banach space — in particular the space $(\mathbb{R}^d, ||.||_*)$ in which our variables $\boldsymbol{w}_1, \ldots, \boldsymbol{w}_k$ lie — the unit ball can be covered by at most $(1 + 2/\epsilon)^d$ ϵ-balls (see e.g. [14]).

Now, observe that (1) and (2) both follow from Lemma 4.2 and the rapid convergence to the optimal value proved above. We now prove (3). Assume that $\mathcal{B}$ is rejected by the algorithm; rapid convergence implies that, with probability $1 - \delta$, we have

$$\sum_{j=1}^{k} q_j(\boldsymbol{w}_j \cdot \boldsymbol{b}_j - 2z_j) - \mathop{\mathbf{E}}_{\boldsymbol{a} \sim \mathcal{A}}[\max_{j=1}^{k} \boldsymbol{w}_j \cdot \boldsymbol{a} - 2z_j] > \epsilon_1. \tag{14}$$

Now assume that $\mathcal{A}$ is an ϵ_1-dispersion of $\mathcal{C} = (C, \boldsymbol{r})$. Lemmas 4.2 and 4.3, combined with the observation that $(W, \boldsymbol{z})$ is also feasible for (12) when instantiated for $\mathcal{A}$ and $\mathcal{C}$, imply that

$$\sum_{j=1}^{k} r_j(\boldsymbol{w}_j \cdot \boldsymbol{c}_j - 2z_j) - \mathop{\mathbf{E}}_{\boldsymbol{a}\sim\mathcal{A}}[\max_{j=1}^{k} \boldsymbol{w}_j \cdot \boldsymbol{a} - 2z_j] \leq \epsilon_1 \tag{15}$$

Combining inequalities (14) and (15) yields (13).

For runtime, it is easy to verify that convex programs (11) and (12) each admit an efficient (exact) separation oracle for the feasible set, permit evaluation of the objective and its gradient (or a subgradient), and admit a suitable bounding ellipsoid, when $||.||$ is an L^p norm. Standard application of the ellipsoid method then permits solving both programs in time polynomial in d, k, ℓ, and the desired precision. This suffices for implementing Algorithm 2 efficiently. We omit the straightforward details. $\square$

Theorem 4.5. *Fix $\epsilon_1 \geq 0$ and $\epsilon_2 > 0$. Assume $\mathcal{A}$ is an ϵ_1-dispersion of $\mathcal{B}$. Algorithm 3 instantiated with $(\mathcal{A}, \mathcal{B})$ implements a signaling scheme which $(\epsilon_1+\epsilon_2)$-forges $\mathcal{B}$ from $\mathcal{A}$. Moreover, the algorithm can be implemented in time* $\mathrm{poly}(d, k, \frac{1}{\epsilon_2})$) *when $||.|| = ||.||_p$ for some $p \geq 1$.*

Proof. Similar to the proof of Theorem 4.4, we know that with an $\ell = 288(\frac{1}{\epsilon_2})^2 k(d+2)\log\frac{33k}{\epsilon_2} = \mathrm{poly}(d, k, \frac{1}{\epsilon_2})$, the value of (11) computed in Step 2 of Algorithm 3, is within $\epsilon_2/2$ of its limit — i.e. no more than $\epsilon_1 + \epsilon_2/2$ — with probability at least $1 - \epsilon_2/4$. Now consider the effect of the signaling scheme φ_x described by x — the solution computed in Step 2 — when applied to a vector $\boldsymbol{a}'$ uniformly chosen from $\widetilde{\mathcal{A}} \cup \{\boldsymbol{a}\}$. With probability at least $1 - \epsilon_2/4$ over choice of $\widetilde{\mathcal{A}}$ and $\boldsymbol{a}$, the distribution of the random variable $\boldsymbol{b}' = \mathbf{E}[\boldsymbol{a}'|\varphi_x(\boldsymbol{a}')]$, after conditioning on $\widetilde{\mathcal{A}}$ and $\boldsymbol{a}$, is $(\epsilon_1+\epsilon_2/2)$-close (in earthmover distance) to $\mathcal{B}$. Since $\boldsymbol{b}'$ is never more than a distance of 2 from any point in the support of $\mathcal{B}$ (recall we restricted ourselves to the unit ball), after removing the conditioning we are guaranteed that $\boldsymbol{b}'$ is within earthmover distance at most $\epsilon_1 + \epsilon_2/2 + 2\cdot\epsilon_2/4 = \epsilon_1 + \epsilon_2$ from distribution $\mathcal{B}$.

To complete the proof, we observe that the joint distribution of the input $\boldsymbol{a}$ and the signaling scheme φ_x is identical to the joint distribution of $\boldsymbol{a}'$ and φ_x.[6] By the principle of deferred decisions, applying the (random) signaling scheme φ_x to the input $\boldsymbol{a}$ yields an output which is $(\epsilon_1 + \epsilon_2)$-close to the distribution $\mathcal{B}$ in earthmover distance. This concludes the proof.

The runtime bound follows as in Theorem 4.4. $\square$

4.1 Structural Properties of Dispersion

In this subsection, we establish two structural properties of dispersion which lead to a proof of Lemmas 4.2 and 4.3, and also come in handy in Section 5. Our first structural lemma characterizes optimal approximate solutions to the dispersion problem. Specifically, when finding the closest (in terms of earth-mover distance) distribution to $\mathcal{B}$ that is dispersed by $\mathcal{A}$, we can restrict attention to distributions $\widetilde{\mathcal{B}}$ whose support is in a one-to-one correspondence with the support of $\mathcal{B}$, and for which the optimal transport map from $\widetilde{\mathcal{B}}$ to $\mathcal{B}$ simply sends all mass on $\widetilde{\boldsymbol{b}}_j$ to $\boldsymbol{b}_j$.

Lemma 4.6. *Let $\mathcal{A}$ and $\mathcal{B}$ be two distributions on $\mathbb{R}^d$, where $\mathcal{B} = (B, \boldsymbol{q})$ has support size k. If $\mathcal{A}$ is an ϵ-dispersion of $\mathcal{B}$, then $\mathcal{A}$ is a dispersion of some distribution $\widetilde{\mathcal{B}} = (\widetilde{B}, \boldsymbol{q})$ with support size k and* $\mathrm{EMD}(\widetilde{\mathcal{B}}, \mathcal{B}) = \sum_{j=1}^{k} q_j||\widetilde{\boldsymbol{b}}_j - \boldsymbol{b}_j|| \leq \epsilon$.

[6] Technically, this requires that the computation in step (2) does not depend on the order of the input $\widetilde{\mathcal{A}} \cup \{\boldsymbol{a}\}$. This can easily be guaranteed by applying a random permutation to the empirical distribution prior to solving the convex program.

Proof. Let $\mathcal{C}$ be a distribution such that $\mathcal{A}$ is a dispersion of $\mathcal{C}$, and $\mathrm{EMD}(\mathcal{C}, \mathcal{B}) \leq \epsilon$. Let $\boldsymbol{c} \sim \mathcal{C}$, and $\boldsymbol{z}_1$ be such that $\boldsymbol{c}+\boldsymbol{z}_1 \sim \mathcal{A}$ and $\mathbf{E}[\boldsymbol{z}_1|\boldsymbol{c}] = \mathbf{0}$ — such a random variable $\boldsymbol{z}_1$ exists by our assumption that $\mathcal{A}$ is a dispersion of $\mathcal{C}$. Denote $\boldsymbol{a} = \boldsymbol{c}+\boldsymbol{z}_1$. Let $\boldsymbol{z}_2$ be such that $\boldsymbol{c}+\boldsymbol{z}_2 \sim \mathcal{B}$, and $\mathbf{E}[||\boldsymbol{z}_2||] \leq \epsilon$ — such a random variable $\boldsymbol{z}_2$ exists by the fact that earthmover distance between $\mathcal{C}$ and $\mathcal{B}$ is at most ϵ. Denote $\boldsymbol{b} = \boldsymbol{c} + \boldsymbol{z}_2$. We may assume that $\boldsymbol{z}_1$ and $\boldsymbol{z}_2$ are conditionally independent given $\boldsymbol{c}$. Let $\widetilde{\boldsymbol{b}} = \mathbf{E}[\boldsymbol{c}|\boldsymbol{b}]$ — note that $\widetilde{\boldsymbol{b}}$ is a random variable support size k, taking on some value $\widetilde{\boldsymbol{b}}_j$ whenever $\boldsymbol{b} = \boldsymbol{b}_j$ (with probability q_j).

First, we show that $\boldsymbol{a}$ is a dispersion of $\widetilde{\boldsymbol{b}}$.

$$\begin{aligned}
\mathbf{E}[\boldsymbol{a}|\widetilde{\boldsymbol{b}} = \widetilde{\boldsymbol{b}}_j] &= \mathbf{E}[\boldsymbol{a}|\boldsymbol{b} = \boldsymbol{b}_j] \\
&= \mathbf{E}[\boldsymbol{c}|\boldsymbol{b} = \boldsymbol{b}_j] + \mathbf{E}[\boldsymbol{z}_1|\boldsymbol{b} = \boldsymbol{b}_j] \\
&= \widetilde{\boldsymbol{b}}_j + \mathbf{E}[\boldsymbol{z}_1|\boldsymbol{c} + \boldsymbol{z}_2 = \boldsymbol{b}_j] \\
&= \widetilde{\boldsymbol{b}}_j + \sum_{\boldsymbol{c}' \in \mathbf{supp}(\mathcal{C})} \mathbf{Pr}[\boldsymbol{c} = \boldsymbol{c}'|\boldsymbol{c} + \boldsymbol{z}_2 = \boldsymbol{b}_j]\, \mathbf{E}[\boldsymbol{z}_1|\boldsymbol{c} = \boldsymbol{c}', \boldsymbol{c} + \boldsymbol{z}_2 = \boldsymbol{b}_j] \\
&= \widetilde{\boldsymbol{b}}_j + \sum_{\boldsymbol{c}' \in \mathbf{supp}(\mathcal{C})} \mathbf{Pr}[\boldsymbol{c} = \boldsymbol{c}'|\boldsymbol{c} + \boldsymbol{z}_2 = \boldsymbol{b}_j]\, \mathbf{E}[\boldsymbol{z}_1|\boldsymbol{c} = \boldsymbol{c}', \boldsymbol{z}_2 = \boldsymbol{b}_j - \boldsymbol{c}'] \\
&= \widetilde{\boldsymbol{b}}_j + \sum_{\boldsymbol{c}' \in \mathbf{supp}(\mathcal{C})} \mathbf{Pr}[\boldsymbol{c} = \boldsymbol{c}'|\boldsymbol{c} + \boldsymbol{z}_2 = \boldsymbol{b}_j]\, \mathbf{E}[\boldsymbol{z}_1|\boldsymbol{c} = \boldsymbol{c}'] \\
&= \widetilde{\boldsymbol{b}}_j
\end{aligned}$$

Second, we show that the earthmover distance between $\widetilde{\boldsymbol{b}}$ and $\boldsymbol{b}$ is at most ϵ, by way of the trivial transport map which simply takes $\widetilde{\boldsymbol{b}}_j$ to $\boldsymbol{b}_j$.

$$\begin{aligned}
\epsilon &\geq \mathbf{E}[||\boldsymbol{z}_2||] \\
&= \mathbf{E}[||\boldsymbol{c} - \boldsymbol{b}||] \\
&= \sum_{j=1}^{k} q_j\, \mathbf{E}[||\boldsymbol{c} - \boldsymbol{b}_j|| \mid \boldsymbol{b} = \boldsymbol{b}_j] \\
&\geq \sum_{j=1}^{k} q_j ||\, \mathbf{E}[\boldsymbol{c}|\boldsymbol{b} = \boldsymbol{b}_j] - \boldsymbol{b}_j|| \\
&= \sum_{j=1}^{k} q_j ||\widetilde{\boldsymbol{b}}_j - \boldsymbol{b}_j||
\end{aligned}$$

□

Our second structural lemma shows that the family of distributions approximately dispersed by a distribution $\mathcal{A}$ form a convex set, when parametrized a certain way. Specifically, for a distribution $\mathcal{B} = (B, \boldsymbol{q})$ on $\mathbb{R}^d$ with $B = [\boldsymbol{b}_1, \ldots, \boldsymbol{b}_k] \in \mathbb{R}^{d \times k}$ and $\boldsymbol{q} \in \Delta_k$, we define an alternate representation $\mathrm{sig}(\mathcal{B})$, which we term its *signature*, as the pair $(\widehat{B}, \boldsymbol{q})$ where $\widehat{B} = [q_1\boldsymbol{b}_1, \ldots, q_k\boldsymbol{b}_k]$. We show that the family of all signatures of distributions ϵ-dispersed by $\mathcal{A}$ form a convex set.

Lemma 4.7. *Let $\mathcal{A}$ be a distribution on $\mathbb{R}^d$, and fix an integer k and a parameter $\epsilon \geq 0$. The set*

$$\left\{ \mathrm{sig}(B, \boldsymbol{q}) : \boldsymbol{q} \in \Delta_k, B \in \mathbb{R}^{d \times k}, \mathcal{A} \textit{ is a } \epsilon\textit{-dispersion of } (B, \boldsymbol{q}) \right\} \tag{16}$$

is a convex subset of $\mathbb{R}^{d \times k} \times \mathbb{R}^k$.

Proof. We use $\widehat{\boldsymbol{b}}_j$ to denote $q_j \boldsymbol{b}_j$. From Lemma 4.6, it follows that $\mathcal{A}$ is a ϵ-dispersion of $(B, \boldsymbol{q})$ if and only if the following system, with variables $\{x_j(\boldsymbol{a}) \in \mathbb{R} : j \in [k], \boldsymbol{a} \in \mathbf{supp}(\mathcal{A})\}$ and $\{\widehat{\boldsymbol{c}}_j \in \mathbb{R}^d : j \in [k]\}$, is satisfiable. Note that, in every feasible solution of this system, $(\widehat{\boldsymbol{C}}, \boldsymbol{q})$ is the signature of a distribution dispersed by $\mathcal{A}$.

$$
\begin{array}{ll}
\sum_j ||\widehat{\boldsymbol{c}}_j - \widehat{\boldsymbol{b}}_j|| \leq \epsilon & \\
\sum_{j=1}^k x_j(\boldsymbol{a}) = \mathbf{Pr}_{\mathcal{A}}[\boldsymbol{a}], & \text{for } \boldsymbol{a} \in \mathbf{supp}(\mathcal{A}). \\
\sum_{\boldsymbol{a} \in \mathbf{supp}(\mathcal{A})} x_j(\boldsymbol{a}) = q_j, & \text{for } j \in [k]. \\
\sum_{\boldsymbol{a} \in \mathbf{supp}(\mathcal{A})} x_j(\boldsymbol{a}) \cdot \boldsymbol{a} = \widehat{\boldsymbol{c}}_j, & \text{for } j \in [k]. \\
x_j(\boldsymbol{a}) \geq \mathbf{0}, & \text{for } \boldsymbol{a} \in \mathbf{supp}(\mathcal{A}), j \in [k].
\end{array}
\tag{17}
$$

Thinking of $\boldsymbol{q}$ and $\widehat{\boldsymbol{b}}_1, \ldots, \widehat{\boldsymbol{b}}_k$ as variables as well, the above is a system of linear equations and a convex inequality, which can be interpreted as an extended formulation for the set of signatures described in (16). In other words, the set (16) is the projection of polytope (17) onto a subset of its variables. The projection of a convex set onto a subspace is convex, completing our proof.

□

4.2 Proof of Lemmas 4.2 and 4.3

Lemma 4.2 follows immediately from Lemma 4.6. We devote the remainder of this section to proving Lemma 4.3.

In order to characterize the optimal value of (12), we define a two-player zero-sum game parametrized by distributions $\mathcal{A}$ and $\mathcal{B}$ over $\mathbb{R}^d$, where $\mathcal{B}$ has support size k. In this game, the maximizing player's strategy consists of the matrix $W = [\boldsymbol{w}_1, \ldots, \boldsymbol{w}_k]$ and the coefficients $z_1, \ldots, z_k$ (constrained as in (12)), while the minimizing player's strategy is a distribution $\widetilde{\mathcal{B}} = (\widetilde{B}, \widetilde{\boldsymbol{q}})$ with support size k such that $\mathcal{A}$ is a dispersion of $\widetilde{\mathcal{B}}$. The payoff of the maximizing player is given by the function f, defined below.

$$
f(W, \boldsymbol{z}, \widetilde{B}, \widetilde{\boldsymbol{q}}) := \sum_{j=1}^k q_j(\boldsymbol{w}_j \cdot \boldsymbol{b}_j - 2z_j) - \sum_{j=1}^k \widetilde{q}_j(\boldsymbol{w}_j \cdot \widetilde{\boldsymbol{b}}_j - 2z_j) \tag{18}
$$

To see the relationship with (12), we can think of $\widetilde{\mathcal{B}}$ in terms of the (possibly randomized) signaling scheme $\varphi : \mathbb{R}^d \to [k]$ which forges $\widetilde{\mathcal{B}}$ from $\mathcal{A}$ (whose existence is guaranteed by Lemma 4.1). Letting $\boldsymbol{a} \sim \mathcal{A}$, we have

$$
\begin{aligned}
f(W, \boldsymbol{z}, \widetilde{B}, \widetilde{\boldsymbol{q}}) &= \sum_{j=1}^k q_j(\boldsymbol{w}_j \cdot \boldsymbol{b}_j - 2z_j) - \sum_{j=1}^k \mathbf{Pr}[\varphi(\boldsymbol{a}) = j](\boldsymbol{w}_j \cdot \mathbf{E}[\boldsymbol{a}|\varphi(\boldsymbol{a}) = j] - 2z_j) && (19) \\
&= \sum_{j=1}^k q_j(\boldsymbol{w}_j \cdot \boldsymbol{b}_j - 2z_j) - \sum_{j=1}^k \mathbf{Pr}[\varphi(\boldsymbol{a}) = j]\, \mathbf{E}[\boldsymbol{w}_j \cdot \boldsymbol{a} - 2z_j|\varphi(\boldsymbol{a}) = j] && (20) \\
&\geq \sum_{j=1}^k q_j(\boldsymbol{w}_j \cdot \boldsymbol{b}_j - 2z_j) - \mathbf{E}[\max_{j=1}^k \boldsymbol{w}_j \cdot \boldsymbol{a} - 2z_j], && (21)
\end{aligned}
$$

where equality holds for the signaling scheme $\varphi(\boldsymbol{a}) = \operatorname{argmax}_{j=1}^k \boldsymbol{w}_j \cdot \boldsymbol{a}_j - 2z_j$. Therefore, the objective of (12) is simply the maximizing player's utility when he plays $(W, \boldsymbol{z})$ and the minimizing player best responds. Consequently, the optimal value of (12) is the maximin value of this game.

Note that f is linear in the maximizing player's strategy $(W, \boldsymbol{z})$, and the feasible such strategies form a compact, convex set. Moreover, f is linear in the minimizing player's strategy when said strategy is represented via the signature of $\widetilde{B}$, as defined in Section 4.1; the family of such signatures form a convex set by Lemma 4.7. Consequently, by Sion's minimax theorem, the value of the game is well-defined, and equals both its minimax and maximin values.

Let ν denote the value of the game (equal to the optimal value of the dual program (12)), and let ϵ^* denote the minimum ϵ such that $\mathcal{A}$ is an ϵ-dispersion of $\mathcal{B}$ (i.e. the optimal value of the primal program (11) by Lemma 4.2). We will complete the proof by showing that $\nu = \epsilon^*$. The first direction, namely $\nu \leq \epsilon^*$, is easy: let $\mathcal{B}^*$ be such that $\mathcal{A}$ is a dispersion of $\mathcal{B}^*$, and $\mathrm{EMD}(\mathcal{B}^*, \mathcal{B}) = \epsilon^*$. By Lemma 4.6, we assume $\mathcal{B}^*$ has support size k, is of the form $(B^*, \boldsymbol{q})$, and the optimal transportation from $\mathcal{B}^*$ to $\mathcal{B}$ is the trivial map $\boldsymbol{b}_j^* \to \boldsymbol{b}_j$. The following holds for every response $(W, \boldsymbol{z})$ of the maximizing player.

$$
\begin{aligned}
f(W, \boldsymbol{z}, B^*, \boldsymbol{q}) &= \sum_{j=1}^{k} q_j(\boldsymbol{w}_j \cdot \boldsymbol{b}_j - 2z_j) - \sum_{j=1}^{k} q_j(\boldsymbol{w}_j \cdot \boldsymbol{b}_j^* - 2z_j) \\
&= \sum_{j=1}^{k} q_j \boldsymbol{w}_j \cdot (\boldsymbol{b}_j - \boldsymbol{b}_j^*) \\
&\leq \sum_{j=1} q_j ||\boldsymbol{b}_j^* - \boldsymbol{b}_j|| && \text{(By Holder's inequality)} \\
&= \mathrm{EMD}(\mathcal{B}^*, \mathcal{B}) = \epsilon^*
\end{aligned}
$$

The second direction, $\nu \geq \epsilon^*$, is more involved. Let $(\widetilde{B}, \widetilde{\boldsymbol{q}})$ be an arbitrary strategy of the minimizing player. We will show that there is a response $(W^*, \boldsymbol{z}^*)$ such that $f(W^*, \boldsymbol{z}^*, \widetilde{B}, \widetilde{\boldsymbol{q}}) \geq \epsilon^*$. Specifically, let $\boldsymbol{w}_j^*$ be such that $\boldsymbol{w}_j^* \cdot (\boldsymbol{b}_j - \widetilde{\boldsymbol{b}}_j) = ||\boldsymbol{b}_j - \widetilde{\boldsymbol{b}}_j||$ and $||\boldsymbol{w}_j^*||_* = 1$ — the existence of such a vector follows from duality of the norms. We denote $J^+ = \{j \in [k] : \widetilde{q}_j > q_j\}$ and $J^- = \{j \in [k] : \widetilde{q}_j \leq q_j\}$. For $j \in J^+$, we let $z_j^* = \frac{1}{2}(||\widetilde{\boldsymbol{b}}_j|| + \boldsymbol{w}_j^* \cdot \widetilde{\boldsymbol{b}}_j)$. For $j \in J^-$, we let $z_j^* = \frac{1}{2}(\boldsymbol{w}_j^* \cdot \boldsymbol{b}_j - ||\boldsymbol{b}_j||)$. We now lowerbound the maximizing player's utility.

$$
\begin{aligned}
f(W^*, \boldsymbol{z}^*, \widetilde{B}, \widetilde{\boldsymbol{q}}) &= \sum_{j=1}^{k} q_j(\boldsymbol{w}_j^* \cdot \boldsymbol{b}_j - 2z_j^*) - \sum_{j=1}^{k} \widetilde{q}_j(\boldsymbol{w}_j^* \cdot \widetilde{\boldsymbol{b}}_j - 2z_j^*) \\
&= \sum_{j \in J^+} q_j \boldsymbol{w}_j^* \cdot (\boldsymbol{b}_j - \widetilde{\boldsymbol{b}}_j) + (\widetilde{q}_j - q_j)(2z_j^* - \boldsymbol{w}_j^* \cdot \widetilde{\boldsymbol{b}}_j) \\
&\quad + \sum_{j \in J^-} \widetilde{q}_j \boldsymbol{w}_j^* \cdot (\boldsymbol{b}_j - \widetilde{\boldsymbol{b}}_j) + (q_j - \widetilde{q}_j)(\boldsymbol{w}_j^* \cdot \boldsymbol{b}_j - 2z_j^*) \\
&= \sum_{j \in J^+} q_j ||\boldsymbol{b}_j - \widetilde{\boldsymbol{b}}_j|| + (\widetilde{q}_j - q_j)||\widetilde{\boldsymbol{b}}_j|| + \sum_{j \in J^-} \widetilde{q}_j ||\boldsymbol{b}_j - \widetilde{\boldsymbol{b}}_j|| + (q_j - \widetilde{q}_j)||\boldsymbol{b}_j||
\end{aligned}
$$

We will now argue that the resulting expression is the cost of a certain transport map from $\widetilde{\mathcal{B}}$ to $\mathcal{B}$, from which it follows that $f(W^*, \boldsymbol{z}^*, \widetilde{B}, \widetilde{\boldsymbol{q}}) \geq \epsilon^*$. For $j \in J^+$, the term $q_j ||\boldsymbol{b}_j - \widetilde{\boldsymbol{b}}_j|| + (\widetilde{q}_j - q_j)||\widetilde{\boldsymbol{b}}_j||$ is the cost of sending q_j probability mass from $\widetilde{\boldsymbol{b}}_j$ to $\boldsymbol{b}_j$, plus the cost exporting the excess $(\widetilde{q}_j - q_j)$ mass from $\widetilde{\boldsymbol{b}}_j$ to $\mathbf{0}$. For $j \in J^-$, the term $\widetilde{q}_j ||\boldsymbol{b}_j - \widetilde{\boldsymbol{b}}_j|| + (q_j - \widetilde{q}_j)||\boldsymbol{b}_j||$ is the cost of sending $\widetilde{q}_j$ mass from $\widetilde{\boldsymbol{b}}_j$ to $\boldsymbol{b}_j$, plus the cost of importing the deficit $q_j - \widetilde{q}_j$ from $\mathbf{0}$ to $\boldsymbol{b}_j$. Since the total excess

$$
\begin{array}{lll}
\text{maximize} & \sum_{i=1}^{n} q_i s_i(i) & \\
\text{subject to} & r_i(i) \geq r_i(j), & \text{for } i,j \in [n] \text{ with } i \neq j. \\
& \lambda \text{ is an } \epsilon'\text{-dispersion of } ([(\boldsymbol{s}_1, \boldsymbol{r}_1), \ldots, (\boldsymbol{s}_n, \boldsymbol{r}_n)], q) & \\
& \boldsymbol{s}_i \in [-1,1]^n, & \text{for } 1 \leq i \leq n. \\
& \boldsymbol{r}_i \in [-1,1]^n, & \text{for } 1 \leq i \leq n. \\
& \boldsymbol{q} \in \Delta_n &
\end{array}
\tag{22}
$$

Optimal Signaling Optimization Problem

probability mass for J^+ equals the total deficit mass for J^-, the expression does evaluate the cost of a transport map from $\widetilde{\mathcal{B}}$ to $\mathcal{B}$, as needed.

5 The General Persuasion Problem

We turn our attention to the Bayesian Persuasion problem where the payoffs of different actions are arbitrarily correlated, and the joint distribution λ is presented as a black-box sampling oracle. We assume that payoffs are bounded, and our approximation guarantees will be additive relative to the range of payoffs. Without loss of generality, we normalize payoffs in order to simplify our exposition; In particular, we assume the black-box distribution λ is supported on $\Theta \subseteq [-1,1]^n \times [-1,1]^n$. In addition to permitting additive approximation in our objective function (i.e. the sender's expected payoff), we also permit approximate incentive compatibility, in the additive sense described in Section 2.1. We prove the following theorem.

Theorem 5.1. *Consider the Bayesian Persuasion problem in the black-box oracle model with n actions and payoffs in $[-1,1]$, and let $\epsilon > 0$ be a parameter. An ϵ-optimal and ϵ-incentive compatible direct signaling scheme can be implemented in $\mathrm{poly}(n, \frac{1}{\epsilon})$ time.*

In proving the above theorem, we employ dispersion testing and inversion as subroutines. Since states of nature lie in $[-1,1]^{2n}$, we instantiate the machinery of Section 4 with $\mathcal{V} = \mathbb{R}^{2n}$ and $||.|| = ||.||_\infty$. Our algorithm will in essence approximately optimize over realizable distributions of posterior expectations — i.e. distributions approximately dispersed by λ — using the dispersion test from Section 4 as an approximate separation oracle. After the optimization yields an approximately realizable distribution of posterior expectations, we use the dispersion inversion procedure to produce a signaling scheme.

The optimization step of our algorithm applies the ellipsoid method, as described in Section 2.3, to approximately solve the optimization problem (22) for a parameter $\epsilon' = 1/\mathrm{poly}(\frac{1}{\epsilon}, n)$ to be set later. Note that (22) slightly relaxes our optimization problem to optimizing over distributions ϵ'-dispersed by λ, and this is necessary for technical reasons related to establishing a lower-bound on the volume of the feasible set, as required by the ellipsoid method. Whereas (22) is not a convex optimization problem as written, it is equivalent to the "convexified" optimization problem (23), which works with the *signature* of the distribution of posterior expectations, as defined in Section 4.1. Lemma 4.7 implies that (23) is indeed a convex program. We use a one-sided-error separation oracle (Definition 2.1) for the feasible set which simply employs a dispersion test (Algorithm 2) with parameters $\epsilon_1 = \epsilon_2 = \epsilon'$, and $\delta = \frac{\epsilon'}{\mathrm{poly}(n, \frac{1}{\epsilon'})}$ sufficiently small so as to guarantee success for all calls to the dispersion test throughout the ellipsoid method's execution with probability $1-\epsilon'$. In our separation oracle, the other constraints are checked explicitly.

$$
\begin{array}{lll}
\text{maximize} & \sum_{i=1}^{n} \widehat{s}_i(i) & \\
\text{subject to} & \widehat{r}_i(i) \geq \widehat{r}_i(j), & \text{for } i,j \in [n] \text{ with } i \neq j. \\
 & \lambda \text{ is an } \epsilon'\text{-dispersion of } \mathrm{sig}^{-1}([(\widehat{\boldsymbol{s}}_1, \widehat{\boldsymbol{r}}_1), \ldots, (\widehat{\boldsymbol{s}}_n, \widehat{\boldsymbol{r}}_n)], \boldsymbol{q}) & \\
 & -q_i \leq \widehat{s}_i(j) \leq q_i, & \text{for } i,j \in [n]. \\
 & -q_i \leq \widehat{r}_i(j) \leq q_i, & \text{for } i,j \in [n]. \\
 & \boldsymbol{q} \in \Delta_n &
\end{array} \tag{23}
$$

Optimal Signaling Convex Program

Lemma 5.2. *There is an algorithm for approximately solving* (22)*, parametrized by* $\epsilon' > 0$*, with the following guarantees:*

- *The algorithm runs in time* $\mathrm{poly}(n, \frac{1}{\epsilon'})$
- *The algorithm outputs* $\mathcal{B} = ([(\boldsymbol{s}_1, \boldsymbol{r}_1), \ldots, (\boldsymbol{s}_n, \boldsymbol{r}_n)], \boldsymbol{q})$ *satisfying all non-dispersion constraints of* (22)
- *With probability* $1 - \epsilon'$*, both the following hold:*
  - λ *is a* $2\epsilon'$*-dispersion of* $\mathcal{B}$
  - $\mathcal{B}$ *is* ϵ'*-optimal; i.e.,* $\sum_i q_i s_i(i)$ *is within an additive* ϵ' *of the optimal value of* (22)

Proof. We apply the ellipsoid method, as described in Section 2.3, to (23). For the OSO, we use Algorithm 2 with parameters $\epsilon_1 = \epsilon_2 = \epsilon'$, and δ polynomially small in n and $\frac{1}{\epsilon'}$ in order to guarantee a probability of success of $1 - \epsilon'$. Note that the certificate output by the dispersion test, as described in Theorem 4.4, actually describes a separating hyperplane in signature space. Other constraints are checked explicitly. The lemma follows from Theorems 4.4 and 2.3, modulo bounds on the volume of the feasible set and a circumscribing ellipsoid.

Observe that the feasible set of (23) lies in the affine set $\mathbb{R}^{2n^2} \times \mathbf{aff}(\Delta_n) \subseteq \mathbb{R}^{2n^2+n}$. As our bounding ellipsoid $\mathcal{E}$, we take the projection of the radius $2n^2 + n$ ball $B(0, 2n^2 + n) \subseteq \mathbb{R}^{2n^2+n}$ onto the subspace $\mathbb{R}^{2n^2} \times \mathbf{aff}(\Delta_n)$. It is clear that the volume of $\mathcal{E}$, relative to the affine space in which it lies, is exponential in n.

To show that the volume of the feasible set of (23) is nontrivial (i.e. no worse than exponentially small in n and $\frac{1}{\epsilon'}$), we express the feasible set as the intersection of two sets $X, Y \subseteq \mathbb{R}^{2n^2+n}$, where X is the family of vectors satisfying the ϵ'-dispersion constraint in (23), and Y is the family of vectors satisfying all other constraints. We will show that Y has nontrivial volume and polynomially-upper-bounded diameter, and that X includes the intersection of Y and a Euclidean ball of nontrivial volume centered inside Y. To complete our proof we use the geometric fact, whose elementary proof we omit, that a convex set of nontrivial volume and polynomially-bounded diameter remains nontrivial if intersected with a ball of nontrivial volume centered within it.

Claim 5.3. *The set* Y *has volume at least* $\exp(-\mathrm{poly}(n))$ *relative to the affine space* $\mathbb{R}^{2n^2} \times \mathbf{aff}(\Delta_n)$*, and has Euclidean diameter at most* $5n^2$*.*

Proof. First we bound the volume. Let $a = ([(\widehat{\boldsymbol{s}}_1, \widehat{\boldsymbol{r}}_1), \ldots, (\widehat{\boldsymbol{s}}_n, \widehat{\boldsymbol{r}}_n)], \boldsymbol{q})$, where $\widehat{\boldsymbol{s}}_1 = \widehat{\boldsymbol{s}}_2 = \ldots = \widehat{\boldsymbol{s}}_n = \mathbf{0}_n$, $\widehat{\boldsymbol{r}}_i = \frac{1}{2n}\boldsymbol{e}_i$ for all i, and $\boldsymbol{q} = \frac{1}{n}\mathbf{1}_n$. It is easy to verify that $a \in Y$. Moreover, a moment's thought reveals that any $b \in \mathbb{R}^{2n^2} \times \mathbf{aff}(\Delta_n)$ with $||b - a||_2 \leq \frac{1}{2n}$ is also in Y — indeed, this follows almost immediately by observing that $||b - a||_\infty \leq \frac{1}{2n}$ and that all inequality constraints defining Y are sufficiently slack at a. Therefore, Y contains a Euclidean ball in $\mathbb{R}^{2n^2} \times \mathbf{aff}(\Delta_n)$ of radius at least $1/2n$. Such a ball has nontrivial volume, as needed.

For the diameter of Y, observe that $Y \subseteq [-1,1]^{2n^2} \times [0,1]^n$, and therefore has Euclidean diameter no more than $4n^2 + n < 5n^2$. $\square$

Claim 5.4. *There is a feasible solution a for (23) (i.e. in both X and Y) such that $b \in X$ for all $b \in \mathbb{R}^{2n^2} \times \Delta_n$ satisfying $||b - a||_2 \leq \frac{\epsilon'}{3n}$.*

Proof. Let $a = \mathrm{sig}([(\boldsymbol{s}_1, \boldsymbol{r}_1), \ldots, (\boldsymbol{s}_n, \boldsymbol{r}_n)], \boldsymbol{q})$ be an arbitrary solution to (23) with the property that λ is a dispersion of $([(\boldsymbol{s}_1, \boldsymbol{r}_1), \ldots, (\boldsymbol{s}_n, \boldsymbol{r}_n)], \boldsymbol{q})$; such a solution exists, induced by an arbitrary signaling scheme (say the trivial scheme which always outputs σ_{i^*} where i^* is the optimal uninformed action for the receiver). Let $b = \mathrm{sig}([(\boldsymbol{s}'_1, \boldsymbol{r}'_1), \ldots, (\boldsymbol{s}'_n, \boldsymbol{r}'_n)], \boldsymbol{q}')$ be as in the statement of the Lemma.

To show that $b \in X$ — i.e. λ is an ϵ'-dispersion of b — we will show that the earthmover distance between a and b with respect to the L^∞-norm is at most ϵ'. This is a consequence of the fact that the L^∞ earthmover distance between two distributions is at most a factor $3n$ greater than the Euclidean distance between their signatures, as we show below for a and b.

$$
\begin{aligned}
\underset{\infty}{\mathrm{EMD}}(a,b) &\leq \sum_{i=1}^{n} \min(q_i, q'_i) ||(\boldsymbol{s}_i, \boldsymbol{r}_i) - (\boldsymbol{s}'_i, \boldsymbol{r}'_i)||_\infty + 2 \sum_{i=1}^{n} |q_i - q'_i| \\
&\leq \sum_{i=1}^{n} ||(\boldsymbol{s}_i, \boldsymbol{r}_i) - (\boldsymbol{s}'_i, \boldsymbol{r}'_i)||_2 + 2\sqrt{n} ||\boldsymbol{q} - \boldsymbol{q}'||_2 \\
&\leq (2\sqrt{n} + n) ||b - a||_2 \\
&\leq 3n ||b - a||_2
\end{aligned}
$$

The first inequality follows from a particular transport map from the distribution described by a to the distribution described by b. Namely, $\min(q_i, q'_i)$ probability mass is moved directly from $(\boldsymbol{s}_i, \boldsymbol{r}_i)$ to $(\boldsymbol{s}'_i, \boldsymbol{r}'_i)$, and remaining probability mass of $\sum_i |q_i - q'_i|$ is routed as needed a distance of at most 2 (the L^∞ diameter of the space). The remaining inequalities follow from elementary manipulations. $\square$

Claim 5.4 implies that the feasible region of (23) includes the intersection of Y — a set with nontrivial volume and polynomially-bounded diameter according to Claim 5.3 — with a Euclidean ball of radius $\frac{\epsilon'}{O(n)}$. Invoking the geometric fact below, we conclude that the volume of the feasible region of (23) is no worse than $\exp(-\operatorname{poly}(n, \frac{1}{\epsilon'}))$, as needed to complete the proof of the Lemma.

Fact 5.5. *Let $S \subseteq \mathbb{R}^N$ be a convex set with diameter at most D, and let $\boldsymbol{x} \in S$. For every $\mu > 0$,*

$$\mathbf{vol}(S \cap B(\boldsymbol{x}, \mu)) \geq \left(\frac{\mu}{D}\right)^N \mathbf{vol}(S)$$

where $B(\boldsymbol{x}, \mu)$ denotes a Euclidean ball of radius μ centered at $\boldsymbol{x}$, and $\mathbf{vol}(.)$ *denotes the standard Lebesgue measure on* $\mathbf{aff}(S)$*.*

$\square$

For the second step of the algorithm, we use the dispersion inversion procedure from Section 4 to turn the output of the optimization step into a signaling scheme.

Lemma 5.6. *Let $\mathcal{B}$ be as described in Lemma 5.2 for $\epsilon' = \frac{\epsilon^2}{64n}$, assuming success of the optimization step of the algorithm. Given $\mathcal{B}$, an $\frac{\epsilon}{2}$-optimal and $\frac{\epsilon}{2}$-incentive compatible direct signaling scheme can be implemented in time* $\text{poly}(n, \frac{1}{\epsilon})$.

Proof. First, we implement a signaling scheme φ which $3\epsilon'$-forges $\mathcal{B}$ from λ. This proceeds via the inversion procedure in Algorithm 3 with $\epsilon_1 = 2\epsilon'$ and $\epsilon_2 = \epsilon'$, whose properties are summarized by Theorem 4.5. Since the objective function of (22) is clearly 1-Lipschitz with respect to L^∞-earthmover distance, and $\mathcal{B}$ is ϵ'-optimal, the signaling scheme φ is $4\epsilon'$-optimal.

However, despite $\mathcal{B}$ describing an incentive compatible scheme and φ $3\epsilon'$-forging $\mathcal{B}$, the scheme φ may not be approximately incentive compatible. This is a consequence of low-probability signals in $\mathcal{B}$, which may be distorted significantly by φ without contributing much to the earthmover distance between the $\mathcal{B}$ and the posterior distribution $\mathcal{C}$ exactly forged by φ. To remedy this, we postprocess φ into a scheme φ', which "discards" such low-probability signals in $\mathcal{B}$ in a careful manner. Call a signal σ_i *small* if $q_i < \frac{\epsilon}{8n}$, otherwise call it *big*. Since $\text{EMD}(\mathcal{B}, \mathcal{C}) \leq 3\epsilon'$, for every big signal σ_i, the incentive compatibility constraint is approximately satisfied by $\mathcal{C}$ up to an additive error of $\frac{3\epsilon'}{q_i} < 32n\frac{\epsilon'}{\epsilon} = \frac{\epsilon}{2}$. Our scheme φ' is defined as follows: On input θ, invoke $\varphi(\theta)$ to get a signal σ; if σ is big output σ, otherwise output $\sigma_{\text{argmax}_i r_i(\theta)}$ — the signal corresponding to the receiver's optimal action in state of nature θ.

By recommending the receiver's optimal action in place of small signals, the scheme φ' preserves the incentive-constraints of big signals in φ, and removes small signals from its portfolio of signals. Moreover, the total probability mass of small signals is at most $\frac{\epsilon}{8n} \times n = \frac{\epsilon}{8}$. Consequently, the earthmover distance between the distribution of posteriors $\mathcal{C}'$ forged by φ' and $\mathcal{C}$ is at most $\frac{\epsilon}{8}$, and since φ is $4\epsilon'$-optimal we get that φ' is within $(4\epsilon' + \frac{\epsilon}{8}) < \frac{\epsilon}{2}$ of optimality. □

Lemmas 5.2 and 5.6 yield an algorithm which implements, with probability $1 - \epsilon' > 1 - \frac{\epsilon}{2}$, a direct signaling scheme which is $\frac{\epsilon}{2}$-optimal and $\frac{\epsilon}{2}$-incentive compatible in the black box model. Removing the conditioning on the success of the optimization step, and noting that both the incentive constraints and our objective are 1-Lipschitz in the posterior distribution with respect to the earthmover distance, overall we implement an ϵ-optimal and ϵ-incentive compatible scheme, as needed to complete the proof of Theorem 5.1.

References

[1] G. A. Akerlof. The Market for 'Lemons': Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, August 1970. URL http://ideas.repec.org/a/tpr/qjecon/v84y1970i3p488-500.html.

[2] S. Alaei, H. Fu, N. Haghpanah, J. D. Hartline, and A. Malekian. Bayesian optimal auctions via multi- to single-agent reduction. In B. Faltings, K. Leyton-Brown, and P. Ipeirotis, editors, *ACM Conference on Electronic Commerce*, page 17. ACM, 2012. ISBN 978-1-4503-1415-2. URL http://dblp.uni-trier.de/db/conf/sigecom/sigecom2012.html#AlaeiFHHM12.

[3] R. Alonso and O. Camara. Persuading voters. Working paper, 2014.

[4] S. P. Anderson and R. Renault. Advertising content. *American Economic Review*, 96(1):93–113, 2006. doi: 10.1257/000282806776157632. URL http://www.aeaweb.org/articles.php?doi=10.1257/000282806776157632.

[5] G. Antioch. Persuasion is now 30 per cent of us gdp. *Economic Roundup*, (1):1–10, 2013. URL http://EconPapers.repec.org/RePEc:tsy:journl:journl_tsy_er_2013_1_1.

[6] D. Bergemann and S. Morris. The comparison of information structures in games: Bayes correlated equilibrium and individual sufficiency. Technical Report 1909R, Cowles Foundation for Research in Economics, Yale University, 2014.

[7] D. Bergemann and M. Pesendorfer. Information structures in optimal auctions. *Journal of Economic Theory*, 137(1):580 – 609, 2007. ISSN 0022-0531. doi: http://dx.doi.org/10.1016/j.jet.2007.02.001. URL `http://www.sciencedirect.com/science/article/pii/S0022053107000324`.

[8] D. Bergemann, B. Brooks, and S. Morris. The limits of price discrimination. Technical Report 1896RR, Cowles Foundation for Research in Economics, Yale University, 2014.

[9] K. Border. Reduced Form Auctions Revisited. *Economic Theory*, 31(1):167–181, April 2007. URL `http://ideas.repec.org/a/spr/joecth/v31y2007i1p167-181.html`.

[10] K. C. Border. Implementation of Reduced Form Auctions: A Geometric Approach. *Econometrica*, 59(4), 1991. ISSN 00129682. doi: 10.2307/2938181.

[11] I. Brocas and J. D. Carrillo. Influence through ignorance. *The RAND Journal of Economics*, 38(4):931–947, 2007. ISSN 1756-2171. doi: 10.1111/j.0741-6261.2007.00119.x. URL `http://dx.doi.org/10.1111/j.0741-6261.2007.00119.x`.

[12] Y. Cai, C. Daskalakis, and S. M. Weinberg. An algorithmic characterization of multi-dimensional mechanisms. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing*, STOC '12, pages 459–478, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1245-5. doi: 10.1145/2213977.2214021. URL `http://doi.acm.org/10.1145/2213977.2214021`.

[13] A. Chakraborty and R. Harbaugh. Persuasive puffery. Technical Report 2012-05, Indiana University, Kelley School of Business, Department of Business Economics and Public Policy, 2012. URL `http://EconPapers.repec.org/RePEc:iuk:wpaper:2012-05`.

[14] F. Cucker and D. Zhou. *Learning Theory: An Approximation Theory Viewpoint*. Cambridge Monographs on Applied and Computational Mathematics. Cambridge University Press, 2007. ISBN 9781139462860. URL `http://books.google.com/books?id=d8wmcLiuDtgC`.

[15] C. Daskalakis and S. M. Weinberg. Symmetries and optimal multi-dimensional mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 370–387, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1415-2. doi: 10.1145/2229012.2229042. URL `http://doi.acm.org/10.1145/2229012.2229042`.

[16] S. DellaVigna and M. Gentzkow. Persuasion: Empirical Evidence. *Annual Review of Economics*, (0), 2010. ISSN 1941-1383.

[17] S. Dughmi. On the hardness of signaling. In *Proceedings of the 55th Symposium on Foundations of Computer Science*, FOCS '14. IEEE Computer Society, 2014.

[18] S. Dughmi, N. Immorlica, and A. Roth. Constrained signaling in auction design. In *Proceedings of the Twenty-five Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14. Society for Industrial and Applied Mathematics, 2014.

[19] Y. Emek, M. Feldman, I. Gamzu, R. Paes Leme, and M. Tennenholtz. Signaling schemes for revenue maximization. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 514–531, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1415-2. doi: 10.1145/2229012.2229051. URL `http://doi.acm.org/10.1145/2229012.2229051`.

[20] P. Eso and B. Szentes. Optimal information disclosure in auctions and the handicap auction. *The Review of Economic Studies*, 74(3):pp. 705–731, 2007.

[21] M. Gentzkow and E. Kamenica. Competition in persuasion. Working Paper 17436, National Bureau of Economic Research, September 2011.

[22] M. Gentzkow and E. Kamenica. Costly persuasion. *American Economic Review*, 104(5):457–62, 2014. doi: 10.1257/aer.104.5.457. URL `http://www.aeaweb.org/articles.php?doi=10.1257/aer.104.5.457`.

[23] W. Gick and T. Pausch. *Persuasion by stress testing: Optimal disclosure of supervisory information in the banking sector*. Number 32/2012. Discussion Paper, Deutsche Bundesbank, 2012.

[24] I. Goldstein and Y. Leitner. Stress tests and information disclosure. 2013.

[25] M. Guo and A. Deligkas. Revenue maximization via hiding item attributes. *CoRR*, abs/1302.5332, 2013. URL `http://dblp.uni-trier.de/db/journals/corr/corr1302.html#abs-1302-5332`.

[26] J. P. Johnson and D. P. Myatt. On the simple economics of advertising, marketing, and product design. *American Economic Review*, 96(3):756–784, 2006. doi: 10.1257/aer.96.3.756. URL `http://www.aeaweb.org/articles.php?doi=10.1257/aer.96.3.756`.

[27] E. Kamenica and M. Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011. doi: 10.1257/aer.101.6.2590. URL `http://www.aeaweb.org/articles.php?doi=10.1257/aer.101.6.2590`.

[28] A. Kolotilin. Experimental design to persuade. *UNSW Australian School of Business Research Paper*, (2013-17), 2013.

[29] I. Kremer, Y. Mansour, and M. Perry. Implementing the "wisdom of the crowd". *Journal of Political Economy*, 122(5):988–1012, 2014.

[30] Y. Mansour, A. Slivkins, and V. Syrgkanis. Bayesian incentive-compatible bandit exploration. *arXiv preprint arXiv:1502.04147*, 2015.

[31] D. McCloskey and A. Klamer. One quarter of gdp is persuasion. *The American Economic Review*, 85(2):pp. 191–195, 1995. ISSN 00028282. URL `http://www.jstor.org/stable/2117917`.

[32] P. B. Miltersen and O. Sheffet. Send mixed signals: earn more, work less. In B. Faltings, K. Leyton-Brown, and P. Ipeirotis, editors, *ACM Conference on Electronic Commerce*, pages 234–247. ACM, 2012. ISBN 978-1-4503-1415-2. URL `http://dblp.uni-trier.de/db/conf/sigecom/sigecom2012.html#MiltersenS12`.

[33] Z. Rabinovich, A. X. Jiang, M. Jain, and H. Xu. Information disclosure as a means to security. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS),*, 2015.

[34] L. Rayo and I. Segal. Optimal information disclosure. *Journal of Political Economy*, 118(5): pp. 949–987, 2010. ISSN 00223808. URL http://www.jstor.org/stable/10.1086/657922.

[35] M. Spence. Job market signaling. *The Quarterly Journal of Economics*, 87(3):355–374, 1973. doi: 10.2307/1882010.

[36] C. Villani. *Optimal transport : old and new*. Grundlehren der mathematischen Wissenschaften. Springer, Berlin, 2009. ISBN 978-3-540-71049-3. URL http://opac.inria.fr/record=b1129524.

[37] S. M. Weinberg. *Algorithms for Strategic Agents*. PhD thesis, Massachusetts Institute of Technology, 2014.

[38] H. Xu, Z. Rabinovich, S. Dughmi, and M. Tambe. Exploring information asymmetry in two-stage security games. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2015.