Theoretical Economics 19 (2024), 1701–1755 1555-7561/20241701

The persuasion duality

Piotr Dworczak
Department of Economics, Northwestern University

Anton Kolotilin
School of Economics, UNSW Business School

We present a unified duality approach to Bayesian persuasion. The optimal dual variable, interpreted as a price function on the state space, is shown to be a supergradient of the concave closure of the objective function at the prior belief. Strong duality holds when the objective function is Lipschitz continuous.

When the objective depends on the posterior belief through a set of moments, the price function induces prices for posterior moments that solve the corresponding dual problem. Thus, our general approach unifies known results for one-dimensional moment persuasion, while yielding new results for the multidimensional case. In particular, we provide a condition for the optimality of convex-partitional signals, derive structural properties of solutions, and characterize the optimal persuasion scheme when the state is two-dimensional and the objective is quadratic.

Keywords. Bayesian persuasion, information design, duality theory, price function, moment persuasion, convex partition.

JEL classification. D82, D83.

1. Introduction

Kamenica and Gentzkow (2011) show that the optimal signal in a Bayesian persuasion problem concavifies the objective function in the space of posterior beliefs over the state (see Bergemann and Morris (2019) and Kamenica (2019) for excellent overviews of the burgeoning literature on Bayesian persuasion). Although conceptually attractive, concavification is not always a tractable approach. Thus, several recent papers (see Kolotilin (2018), Dworczak and Martini (2019), Dizdar and Kováč (2020), Galperti, Levkun, and Perego (2024), and Kolotilin, Corrao, and Wolitzky (2024)) used duality theory to characterize the optimal signal.

Piotr Dworczak: piotr.dworczak@northwestern.edu
Anton Kolotilin: akolotilin@gmail.com
We thank Isa Chavez, Elliot Lipnowski, Benny Moldovanu, Jacopo Perego, Xiaoyun Qiu, David Rahman, Doron Ravid, Chris Shannon, Jakub Steiner, Marciano Siniscalchi, Bruno Strulovici, and Alexander Wolitzky for helpful comments and suggestions. Dworczak thanks the University of Oxford and the University of Zurich for hospitality and support during the academic year 2021–2022, when some of the work on this paper was undertaken; he also gratefully acknowledges the support of the Alfred P. Sloan Fellowship. Kolotilin gratefully acknowledges support from the Australian Research Council Discovery Early Career Research Award DE160100964.

In this paper, we present a unified duality approach to the Bayesian persuasion problem. Our approach builds on and extends the geometric duality of Gale (1967). The primal and the dual problems correspond to finding, respectively, the concave closure and the concave envelope of the objective function. We show that the optimal dual variable is a supergradient of the concave closure of the objective function at the prior belief (Section 3). Moreover, the dual variable can be represented as a price function on the state space. Because concave functions on finite-dimensional spaces have a supergradient at any interior point, strong duality always holds when the state space is finite. It may fail, however, when the state space is infinite; we prove that strong duality holds if the objective function is Lipschitz.

If the objective function depends only on a finite set of moments of the posterior distribution (the "moment persuasion" case analyzed in Section 4), prices for states induce prices for moments. The resulting price function is convex, lies above the graph of the objective function, and exhibits all other properties of the optimal dual variable known from the analysis of one-dimensional moment persuasion. Thus, our results generalize and unify the duality results established by Kolotilin (2018), Dworczak and Martini (2019), Dizdar and Kováč (2020), and Kolotilin, Corrao, and Wolitzky (2024) for the one-dimensional case. When the state space is multidimensional or the objective function depends on more than one moment, our generalized duality approach yields new results and insights. If the objective function is differentiable, the price function can be constructed explicitly as the upper envelope of hyperplanes that are tangent to the objective function at the conjectured support of moments. Using this construction, we derive a necessary and sufficient condition for the optimality of a convex-partitional signal (an extension of the one-dimensional notion of a monotone-partitional signal), and establish a multidimensional analog of the bipooling result due to Arieli, Babichenko, Smorodinsky, and Yamashita (2023) and Kleiner, Moldovanu, and Strack (2021).

We use these tools to characterize the optimal signal in the classical model of Rayo and Segal (2010) in which the state is two-dimensional and the objective function is a quadratic form (Section 5). We show that the "bait and switch" pooling strategy of Rayo and Segal results from a trade-off between the conflicting goals of disclosing as much information as possible about a sum of two variables, while disclosing as little information as possible about their difference. Under regularity conditions, duality permits us to represent the optimal signal as a convex partition of the two-dimensional state space into negative-sloped line segments. That is, the optimal signal discloses a weighted sum of the two dimensions, with a weight that may depend on the induced posterior moment. We further characterize cases in which the weight is constant, such as when the optimal signal is a sum of the two dimensions.

A contemporaneous paper (Malamud and Schrimpf (2022)) also made progress on analyzing multidimensional moment persuasion, relying on different tools.[1] While

[1]Using the theory of real analytic functions, Malamud and Schrimpf establish a remarkably powerful result that, under a regularity condition on the prior and the objective function, there exists an optimal deterministic signal. This result forms the foundation of their analysis. Relying on metric geometry and the theory of the Hausdorff dimension, they show that optimal signals correspond to low-dimensional manifolds.

some of our results in Section 4.4.1 are related to theirs, we believe the two approaches to be complementary: for example, Malamud and Schrimpf allow the state space to be noncompact, while we cover cases when optimal signals are nondeterministic. The precise relationship to this and other papers is discussed in more detail throughout the paper in the context of specific results.

We briefly note that—despite our focus on Bayesian persuasion as the leading application—the methods we develop can be applied in any problem in which a linear objective is maximized over distributions of posteriors subject to a Bayes-plausibility constraint. Such optimization programs arise in various models with multiple interacting Receivers and in the analysis of rational-inattention and information-acquisition problems. We further discuss alternative applications and directions for future research in Section 6.

2. Model

Let (Ω, ρ) be a compact metric space, where ρ is a metric on Ω. We will also refer to Ω as a measurable space, in which case the σ-algebra should be understood as the Borel σ-algebra induced by the metric ρ. The set of Lipschitz functions on Ω, denoted by $\mathrm{Lip}(\Omega)$, is the set of functions $p : \Omega \to \mathbb{R}$ such that

$$\|p\|_L := \sup\left\{\frac{\left|p(\omega) - p(\omega')\right|}{\rho(\omega, \omega')} : \omega, \omega' \in \Omega, \omega \neq \omega'\right\} < \infty.$$

A function $p \in \mathrm{Lip}(\Omega)$ is L-Lipschitz if $\|p\|_L \leq L$. Let $\mathrm{Lip}_1(\Omega)$ denote the set of 1-Lipschitz functions on Ω.

Let $M(\Omega)$ be the set of finite signed Borel measures on Ω, and $\Delta(\Omega)$ be the subset of probability measures. On the linear space $M(\Omega)$, we define the Kantorovich–Rubinstein norm: for each $\mu \in M(\Omega)$,

$$\|\mu\|_{\mathrm{KR}} := \left|\mu(\Omega)\right| + \sup\left\{\int_\Omega p(\omega)\, \mathrm{d}\mu(\omega) : p \in \mathrm{Lip}_1(\Omega),\, p(\omega_0) = 0\right\},$$

where ω_0 is an arbitrary fixed element of Ω. Since (Ω, ρ) is a compact metric space, Theorem 6.9 and Remark 6.19 in Villani (2009) yield that $\|\cdot\|_{\mathrm{KR}}$ metrizes the weak* topology on $\Delta(\Omega)$ and that $(\Delta(\Omega), \|\cdot\|_{\mathrm{KR}})$ is a compact metric space. Let $\Delta(\Delta(\Omega))$ be the set of Borel probability measures on $\Delta(\Omega)$, endowed with the Kantorovich–Rubinstein distance. Then $\Delta(\Delta(\Omega))$ is also a compact metric space.

We now formally define the persuasion problem, as in Kamenica and Gentzkow (2011). The state space is Ω, and there is a prior belief $\mu_0 \in \Delta(\Omega)$. An objective function $V : \Delta(\Omega) \to \mathbb{R}$ is bounded and upper semicontinuous. We will be imposing increasingly stronger assumptions on V to derive increasingly stronger results throughout the paper.

The persuasion problem is to find a *distribution of posterior beliefs* $\tau \in \Delta(\Delta(\Omega))$ to[2]

$$\begin{aligned} &\text{maximize} \int_{\Delta(\Omega)} V(\mu)\, \mathrm{d}\tau(\mu) \\ &\text{subject to} \int_{\Delta(\Omega)} \mu\, \mathrm{d}\tau(\mu) = \mu_0. \end{aligned} \tag{P}$$

We will denote by $\mathcal{T}(\mu_0)$ the set of feasible distributions of posteriors, that is,

$$\mathcal{T}(\mu_0) = \left\{\tau \in \Delta\big(\Delta(\Omega)\big) : \int_{\Delta(\Omega)} \mu\, \mathrm{d}\tau(\mu) = \mu_0\right\}.$$

We define the *concave closure* of V at μ_0 to be the value of the persuasion problem:

$$\widehat{V}(\mu_0) := \sup_{\tau \in \mathcal{T}(\mu_0)} \int_{\Delta(\Omega)} V(\mu)\, \mathrm{d}\tau(\mu).$$

That is, $\widehat{V}(\mu_0)$ is the supremum of $z \in \mathbb{R}$ over all (z, μ_0) that can be expressed as a convex combination of $(V(\mu), \mu)$ with $\mu \in \Delta(\Omega)$.[3]

The dual problem is to find a *price function* $p \in \mathrm{Lip}(\Omega)$ to

$$\begin{aligned} &\text{minimize} \int_{\Omega} p(\omega)\, \mathrm{d}\mu_0(\omega) \\ &\text{subject to } V(\mu) \leq \int_{\Omega} p(\omega)\, \mathrm{d}\mu(\omega) \quad \text{for all } \mu \in \Delta(\Omega). \end{aligned} \tag{D}$$

We will denote by $\mathcal{P}(V)$ the set of feasible price functions, that is,[4]

$$\mathcal{P}(V) = \left\{p \in \mathrm{Lip}(\Omega) : V(\mu) \leq \int_{\Omega} p(\omega)\, \mathrm{d}\mu(\omega) \text{ for all } \mu \in \Delta(\Omega)\right\}.$$

We define the *concave envelope* of V at μ_0 to be the value of the dual problem:

$$\overline{V}(\mu_0) := \inf_{p \in \mathcal{P}(V)} \int_{\Omega} p(\omega)\, \mathrm{d}\mu_0(\omega).$$

By Definition 7.4 in Aliprantis and Border (2006), the concave envelope of V at μ_0 is the infimum of values taken at μ_0 by all continuous affine functions on $M(\Omega)$ that bound V

[2] Formally, $\int_{\Delta(\Omega)} \mu\, \mathrm{d}\tau(\mu) = \mu_0$ is understood as $\int_{\Delta(\Omega)} \mu(B)\, \mathrm{d}\tau(\mu) = \mu_0(B)$ for all measurable $B \subset \Omega$. The same comment applies whenever we integrate functions with values in the space of measures. An alternative approach is to use the Bochner integral instead of the familiar Lebesgue integral.

[3] Kamenica and Gentzkow (2011) define the concave closure of V as the smallest concave function that lies above V. Instead, we defined it as the value of the persuasion problem. In the general case of compact metric Ω, the equivalence of these definitions follows from Proposition 3 in the Online Appendix of Kamenica and Gentzkow (2011).

[4] In an earlier draft Dworczak and Kolotilin (2019), we considered a dual problem with a continuous p (but not necessarily Lipschitz). While that approach allowed for strong duality to hold under slightly more permissive assumptions, we could not find any economic applications exploiting that additional generality. The current formulation, inspired by a comment from Doron Ravid, leads to a more elegant exposition.

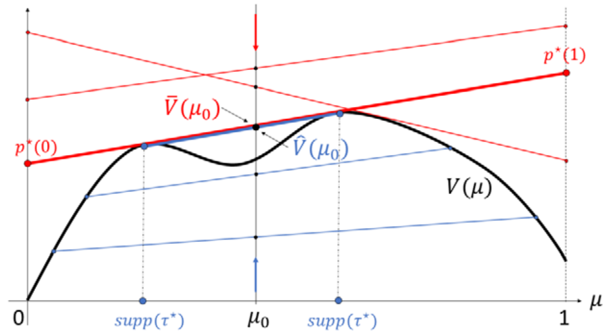

FIGURE 1. The concave closure and the concave envelope at μ_0 in the binary-state case. The concave closure is the inner construction of the convex hull of the graph of V: It involves maximizing the value at μ_0 over all convex combinations of points on the graph of V (exemplified by blue lines in the figure). The concave envelope is the outer construction of the convex hull of the graph of V: It involves minimizing the value at μ_0 over all affine functions lying above the graph of V (exemplified by the red lines in the figure).

from above on $\Delta(\Omega)$. Our definition is equivalent: By Theorem 0 in Hanin (1992),[5] the space dual to $(M(\Omega), \|\cdot\|_{\mathrm{KR}})$ is the space $\mathrm{Lip}(\Omega)$, modulo the constant functions. Hence, any continuous linear function on $(M(\Omega), \|\cdot\|_{\mathrm{KR}})$ can be represented as $\int_\Omega p(\omega)\,\mathrm{d}\mu(\omega)$ for some $p \in \mathrm{Lip}(\Omega)$.[6] The construction of the concave closure and the concave envelope are illustrated in Figure 1 (in the binary-state case).

We interpret the persuasion problem as a linear production problem of Gale (1960). The states are economic resources, and the probability measure μ_0 is a producer's endowment of resources. The set $\Delta(\Omega)$ is the set of linear production processes available to the producer. A process $\mu \in \Delta(\Omega)$ operated at unit level consumes the measure μ of resources and generates income $V(\mu)$. A production plan τ describes the level at which each process μ is operated. The primal problem is for the producer to find a production plan that exhausts the endowment μ_0 and maximizes the total income.

To interpret the dual problem, imagine that there is a wholesaler who wants to buy out the producer. The wholesaler sets a unit price $p(\omega)$ for each resource ω. The producer's (opportunity) cost of operating a process μ at unit level is thus $\int_\Omega p(\omega)\,\mathrm{d}\mu(\omega)$. A price function p is feasible for the wholesaler if the income generated by each process of the producer is not greater than the cost of operating the process, which makes the producer willing to sell all the resources. The dual problem is for the wholesaler to find feasible prices that minimize the total cost of buying up all the resources.[7]

[5] Hanin (1992) credits the result to Kantorovich and Rubinstein (1958). The version of the result that we use is formulated in Exercise 8.10.143 in Bogachev (2007); see also Theorem 7.3 in Edwards (2011).

[6] The distinction between affine and linear functions is immaterial here since a continuous affine function $\int_\Omega p(\omega)\,\mathrm{d}\mu(\omega) + c$ coincides with the continuous linear function $\int_\Omega (p(\omega) + c)\,\mathrm{d}\mu(\omega)$ on $\Delta(\Omega)$.

[7] In the persuasion context, similar interpretations of the dual variable as a price function appear in Dworczak and Martini (2019), Galperti, Levkun, and Perego (2024), and Kolotilin, Corrao, and Wolitzky (2024).

3. Duality

In this section, we establish weak and strong duality for the persuasion problem:

- **Weak duality** states that $\widehat{V}(\mu_0) \leq \overline{V}(\mu_0)$, that is, the concave closure is bounded above by the concave envelope.
- **No duality gap** requires the equality $\widehat{V}(\mu_0) = \overline{V}(\mu_0)$, that is, the concave closure and the concave envelope coincide.
- **Primal and dual attainment** additionally require existence of solutions to the primal and the dual problems, respectively. We use the term **strong duality** when both primal and dual attainment (and hence also no duality gap) hold.[8]

For the case of a binary state, these properties are illustrated in Figure 1. Weak duality follows from the fact that any red line (any affine function lying above the graph of V) achieves a higher value at μ_0 than any blue line (any convex combination of points on the graph of V). No duality gap states that the infimum over values that red lines can take at μ_0 is equal to the supremum over values that the blue lines can take at μ_0. Finally, strong duality requires that these extrema are attained (as depicted by the bold red and blue lines in the figure).

Weak duality serves as a verification tool. If we can find a feasible $\tau \in \mathcal{T}(\mu_0)$ and a feasible $p \in \mathcal{P}(V)$ such that $\int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\tau(\mu) = \int_{\Omega} p(\omega)\,\mathrm{d}\mu_0(\omega)$, then τ is optimal. Within our interpretation, weak duality states that the total income generated by the producer cannot exceed the total cost of the resources under feasible prices, which make the producer willing to sell the resources. Thus, if there exists a plan for the producer and feasible prices for the wholesaler that equalize the total income with the total cost, then this plan must be optimal for the producer, and the prices must be optimal for the wholesaler. However, weak duality does not guarantee that such solutions can be found.

No duality gap ensures that the bound imposed by weak duality is tight. Thus, a feasible $\tau \in \mathcal{T}(\mu_0)$ is optimal if and only if it achieves the value of the concave envelope $\overline{V}(\mu_0)$. The absence of a duality gap still does not guarantee that the optimality of τ can be verified by finding a feasible price function p.

Finally, primal and dual attainment ensure that the solutions to both the primal and the dual problems exist, and hence optimality of the primal solution can be demonstrated by exhibiting a dual solution. Within our interpretation, strong duality states that there exists a feasible plan for the producer and feasible prices for the wholesaler such that the cost of each operated process is equal to the income it generates. In the remainder of this section, we establish weak duality, no duality gap, primal attainment, and—under additional conditions—dual attainment.

Theorem 1 (Weak Duality). $\widehat{V}(\mu_0) \leq \overline{V}(\mu_0)$.

[8]The exact use of these terms varies across authors. For example, Villani (2009) uses the term strong duality to refer to primal attainment and no duality gap. Our convention is consistent with the economics literature where strong duality typically includes existence of solutions to the dual problem (see Daskalakis, Deckelbaum, and Tzamos (2017) and Kleiner and Manelli (2019) for recent examples).

Proof. The proof is relegated to Appendix A.1. □

As the (standard) proof reveals, weak duality does not even require the weak assumptions on V that we imposed (it is only needed that the primal and the dual problems are well-defined). Under our assumptions, weak duality is subsumed by the following stronger claim.

Theorem 2 (No Duality Gap and Primal Attainment). *There is no duality gap,*

$$\widehat{V}(\mu_0) = \overline{V}(\mu_0), \tag{O}$$

and the value of the concave closure $\widehat{V}(\mu_0)$ is attained by some feasible $\tau \in \mathcal{T}(\mu_0)$.

Proof. The proof is relegated to Appendix A.3. □

The primal problem (P) corresponds to maximizing an upper semicontinuous function V over the compact set of feasible distributions $\mathcal{T}(\mu_0)$, so existence of a solution follows from the Weierstrass theorem. No duality gap is a consequence of hyperplane separation. However, instead of explicitly relying on a version of the hyperplane separation theorem, we show that the second concave conjugate (double Legendre transform) of the concave closure equals the concave envelope. The Fenchel–Moreau theorem then establishes the absence of a duality gap (O). Theorem 2 thus implies that the concave closure and the concave envelope coincide, and hence we can use the two notions interchangeably.[9]

One consequence of duality in the persuasion setting is that we can provide a verification result for the persuasion problem and its dual. Within our interpretation, a feasible plan and supporting prices are optimal if and only if the cost of each operated process is equal to the income it generates.

Corollary 1 (Complementary Slackness). *Distribution $\tau \in \mathcal{T}(\mu_0)$ and price $p \in \mathcal{P}(V)$ are optimal solutions to (P) and (D), respectively, if and only if*

$$V(\mu) = \int_\Omega p(\omega)\, \mathrm{d}\mu(\omega), \quad \textit{for all } \mu \in \operatorname{supp}(\tau). \tag{C}$$

Proof. The proof is relegated to Appendix A.4. □

In applications, Corollary 1 can be used to infer properties of solutions to the persuasion problem. However, for this approach to be applicable, we must ensure that a solution to the dual problem exists. Our final goal is to establish conditions under which dual attainment holds. Contrary to previous results, additional regularity conditions on V are needed.

We say that $\widehat{V}$ is *superdifferentiable* at μ_0 if there exists a continuous linear function H on $M(\Omega)$ (called a *supporting hyperplane* of $\widehat{V}$ at μ_0) represented by $p \in \operatorname{Lip}(\Omega)$ (called

[9] When Ω is finite, this follows from Corollary 12.1.1 in Rockafellar (1970).

a *supergradient* of $\widehat{V}$ at μ_0) such that $\widehat{V}(\mu_0) = H(\mu_0)$ and $\widehat{V}(\mu) \le H(\mu) = \int_\Omega p(\omega)\,\mathrm{d}\mu(\omega)$ for all $\mu \in \Delta(\Omega)$. Note that the concave closure $\widehat{V}$ is a concave function. When Ω is finite, a concave function on $\Delta(\Omega)$ is also continuous on the interior of the domain, and hence it is superdifferentiable at all interior points (Theorems 7.12 and 7.24 in Aliprantis and Border (2006)). Interior points in case of finite Ω correspond to priors μ_0 that have full support on Ω. However, when Ω is infinite, the set of probability measures $\Delta(\Omega)$ has an empty (relative) interior—any $\mu_0 \in \Delta(\Omega)$ is a boundary point. As a result, the hyperplane separating $(\mu_0, \widehat{V}(\mu_0))$ from the graph of $\widehat{V}$ may be vertical, and hence the required linear function H may fail to exist.[10]

Following Gale (1967), we say that $\widehat{V}$ has *bounded steepness* at μ_0 if there exists a constant L such that

$$\frac{\widehat{V}(\mu) - \widehat{V}(\mu_0)}{\|\mu - \mu_0\|_{\mathrm{KR}}} \le L, \quad \text{for all } \mu \in \Delta(\Omega).$$

Intuitively, bounded steepness says that the marginal increase in the value of the persuasion problem is bounded above for a small perturbation of the prior.

Theorem 3 (Dual Attainment). *The following statements are equivalent:*

(i) The problem (D) has an optimal solution.

(ii) $\widehat{V}$ is superdifferentiable at μ_0.

(iii) $\widehat{V}$ has bounded steepness at μ_0.

Proof. The proof is relegated to Appendix A.5. □

Equivalence of conditions *(ii)* and *(iii)* is established by the duality theorem in Gale (1967), which we can apply because we represented the space of distributions as a normed space (by using the Kantorovich–Rubinstein norm).[11] Equivalence of conditions *(i)* and *(ii)* follows from the fact that continuous linear functions on $M(\Omega)$ can be identified with Lipschitz functions on Ω. Intuitively, superdifferentiability of $\widehat{V}$ at the prior means that we can find a supporting hyperplane at μ_0. Due to the representation theorem, a supporting hyperplane can be identified with a Lipschitz price function on the state space. By definition of a supporting hyperplane, this price function is feasible and touches the graph of $\widehat{V}$ at μ_0—it must therefore be optimal by weak duality (Theorem 1). This argument shows that the optimal price function is in fact a supergradient of the concave closure $\widehat{V}$ at the prior μ_0.

Geometrically, any price function p defines a hyperplane in $\Delta(\Omega) \times \mathbb{R}$ by specifying what values it takes on extreme points $(\delta_\omega, p(\omega))$ (as depicted by the red lines in Figure 1). The price function p is feasible for (D) if the hyperplane lies above V on $\Delta(\Omega)$.

[10] For an analogy, consider a concave and continuous function $f(x) = \sqrt{x}$ on $[0, 1]$. This function is not superdifferentiable at the boundary point $x = 0$ because the supporting hyperplane would have to be vertical.

[11] Holmes (1975) and Gretsky, Ostroy, and Zame (2002) extend Gale's theorem from normed spaces to locally convex spaces, which may be useful for future generalizations of our results.

The dual problem is to find a hyperplane that lies above V and whose value at the prior μ_0 is minimized. Thus, the optimal hyperplane supports $\widehat{V}$ at μ_0, and the optimal price $p^\star(\omega)$ of each state ω is the value of the supporting hyperplane at the Dirac probability measure δ_ω at ω.

While Theorem 3 provides a necessary and sufficient condition for dual attainment, the condition is stated in terms of a nonprimitive object, the concave closure of V. Next, we present a useful sufficient condition on the primitive objective function V.

THEOREM 4 (Lipschitz Preservation). *Let V be Lipschitz on $\Delta(\Omega)$. Then $\widehat{V}$ is also Lipschitz on $\Delta(\Omega)$. Consequently, $\widehat{V}$ has bounded steepness at each $\mu_0 \in \Delta(\Omega)$.*

PROOF. The proof is relegated to Appendix A.2. □

COROLLARY 2 (Strong Duality). *When V is Lipschitz on $\Delta(\Omega)$, strong duality holds for the persuasion problem (P).*

While the statement of Theorem 4 may seem intuitive, its proof is quite involved in the general (infinite-dimensional) case.[12] Informally, we show that given two priors, μ_0 and η_0, and an optimal distribution $\tau \in \mathcal{T}(\mu_0)$, we can find a perturbation $\eta(\mu)$ of each posterior belief $\mu \in \operatorname{supp}(\tau)$ such that the perturbed posteriors $\eta(\mu)$ average out to η_0 under the distribution τ. Moreover, the average distance between the posteriors μ and their perturbations $\eta(\mu)$ is equal to the distance between μ_0 and η_0. This implies that the value of the persuasion problem under the prior μ_0 cannot exceed the value of the persuasion problem under the prior η_0 by more than $L\|\mu_0 - \eta_0\|_{\mathrm{KR}}$ when V is L-Lipschitz. Reversing the roles of μ_0 and η_0 leads to the desired conclusion.

To the best of our knowledge, Theorems 3 and 4 provide the first general dual attainment result for Bayesian persuasion.[13] Theorem 3 is mathematically more general than the existing strong duality results in the sense that it applies on a larger domain of problems; in fact, bounded steepness of the concave closure is shown to be necessary and sufficient for dual attainment so it must imply all existing sufficient conditions. However, verifying bounded steepness of the concave closure may be difficult in applications. Our Theorem 4 identifies Lipschitz continuity of V as a simple *sufficient* condition for strong duality; while this condition is stronger than the most permissive sufficient condition identified for one-dimensional moment persuasion (Dizdar and Kováč (2020)), it has the advantage of being fully universal—it applies to *any* persuasion problem.

[12]Theorem 4 extends Lemma 1 and Corollary 2 in Guo and Shmaya (2021) from the case of finite Ω to the general case. Theorem 1.17(f) in Laraki (2004) establishes a version of Theorem 4 for the total variation norm on $\Delta(\Omega)$, which is not suitable for our analysis because there is no tractable characterization of the space that is dual to $\Delta(\Omega)$ under the total variation norm.

[13]At the same level of generality, Section 8 of Dworczak and Martini (2019) establishes weak duality by defining a price function on the space of beliefs $\Delta(\Omega)$ and requiring it to be "outer-convex" (a relaxation of convexity). Theorems 3 and 4 demonstrate that such a price function exists when V is Lipschitz, and that the price function can in fact be taken to be *linear* on $\Delta(\Omega)$.

We conclude the section with an illustration of duality by studying conditions for optimality of two extreme information structures: full disclosure (distribution $\tau_F \in \mathcal{T}(\mu_0)$ uniquely characterized by attaching probability one to the set of Dirac probability measures on Ω) and no disclosure (distribution $\tau_N \in \mathcal{T}(\mu_0)$ that attaches probability one to the prior μ_0). We argue that strong duality makes the well-known sufficient conditions necessary.

Suppose that μ_0 has full support on Ω and let V be Lipschitz on $\Delta(\Omega)$ so that, by Theorems 3 and 4, dual attainment holds. Then full disclosure τ_F is optimal if and only if V lies below a linear function that passes through each extreme point $(\delta_\omega, V(\delta_\omega))$:

$$V(\mu) \leq \int_\Omega V(\delta_\omega)\, \mathrm{d}\mu(\omega) \quad \text{for all } \mu \in \Delta(\Omega). \tag{F}$$

No disclosure τ_N is optimal if and only if

$$V \text{ is superdifferentiable at } \mu_0. \tag{N}$$

Theorem 3 implies that the dual problem (D) has an optimal solution. Thus, by Corollary 1, a feasible distribution $\tau \in \mathcal{T}(\mu_0)$ is optimal if and only if the optimal price function $p \in \mathcal{P}(V)$ satisfies (C). The support of τ_F is the set of all Dirac probability measures δ_ω on Ω, so (C) simplifies to $p(\omega) = V(\delta_\omega)$ for all $\omega \in \Omega$. Thus, τ_F is optimal if and only if $V(\delta_\omega)$, treated as a function of ω, belongs to $\mathcal{P}(V)$—this simplifies to (F). Similarly, the condition for optimality of τ_N follows from the observation that feasibility of p along with (C) is equivalent to p being the supergradient of V at the prior, yielding (N).

Because sufficiency follows from weak duality, conditions (F) and (N) are sufficient even without the assumptions on V and μ_0. In Appendix B.1 of Dworczak and Kolotilin (2024), we show that these intuitive conditions are no longer necessary when dual attainment fails.

4. Moment persuasion

In this section, we apply the general duality approach developed in Section 3 to a persuasion problem in which the objective function depends on the posterior belief through a finite set of moments—what we refer to as "moment persuasion." This case arises naturally in persuasion problems in which the Sender's preferences only depend on the Receiver's action, and the Receiver's optimal action depends only on aggregate statistics such as the (potentially multivariate) mean, variance, or skewness of the posterior belief.[14] Multidimensionality allows for applications with multiple Receivers (under public communication), potentially caring about different moments of the public belief. For another example, suppose that there are $N+1$ primitive states of the world but a Sender only observes a partially revealing signal about the primitive state. The Sender sends a signal informative about her own posterior belief to a Receiver. As long as the Receiver maximizes expectation of a utility function that depends on the primitive state—by the law of iterated expectations—her payoff will only depend on the expectation of

[14] Even with a one-dimensional state, this nests the settings of Zhang and Zhou (2016) and Nikandrova and Pancs (2017), as well as a separable special case of Kolotilin, Corrao, and Wolitzky (2024).

the Sender's belief, which can be represented as an element of an N-dimensional simplex.[15] Finally, moment persuasion captures information acquisition problems for certain well-behaved utility functions of the agent acquiring information (e.g., representing mean-variance preferences).

Weak duality for (multidimensional) moment persuasion can be established directly and is often sufficient to solve instances of persuasion problems. However, our approach has two distinct advantages. First, by deriving duality for moment persuasion from the general case, we unify existing approaches (differing in the representation of the constraints in the moment persuasion problem), demonstrate how the dual variables in these alternative approaches relate to one another (Theorem 5), and extend them to the multidimensional case. More substantially—due to strong duality—we are able to derive general predictions about the *structure* of solutions (Theorems 6, 7, 8, as well as Propositions 1 and 2 in the application in Section 5). In particular, strong duality implies that the complementary slackness conditions (C) must always hold; even if the optimal p is unknown, these conditions impose restrictions on the optimal persuasion scheme.

4.1 *Formulation*

We assume that, given some underlying state space $\tilde{\Omega}$ and prior $\tilde{\mu}_0$,

$$\tilde{V}(\tilde{\mu}) = v\left(\int_{\tilde{\Omega}} m(\tilde{\omega})\,\mathrm{d}\tilde{\mu}(\tilde{\omega})\right), \quad \text{for all } \tilde{\mu} \in \Delta(\tilde{\Omega}),$$

for some measurable $m : \tilde{\Omega} \to \mathbb{R}^N$ and some real-valued function v. It will be convenient to redefine the state space as $\Omega = m(\text{supp}(\tilde{\mu}_0))$ with the prior μ_0 given by $\mu_0(B) = \tilde{\mu}_0(m^{-1}(B))$ for any measurable $B \subset \Omega$, so that

$$V(\mu) = v\left(\int_{\Omega} \omega\,\mathrm{d}\mu(\omega)\right), \quad \text{for all } \mu \in \Delta(\Omega).$$

We then define the space of "moments" X as the convex hull of Ω.[16] We assume that X is a compact convex set with nonempty interior[17] and that $v : X \to \mathbb{R}$ is Lipschitz with constant L.

The next lemma ensures that we can rely on dual attainment from Theorems 3 and 4.

Lemma 1. *If v is Lipschitz, then V is also Lipschitz.*

Proof. The proof is relegated to Appendix A.6. □

In moment persuasion, a distribution τ of posterior beliefs $\mu \in \Delta(\Omega)$ influences the objective only through the induced distribution of moments. By Strassen's theorem (e.g.,

[15] Arieli et al. (2023) offer an analogous interpretation of the one-dimensional moment persuasion problem.

[16] By redefining the state space, we have converted a general case of moment persuasion to a problem in which the objective function only depends on a multidimensional vector of posterior means.

[17] This is without loss of generality: As a convex set in $\mathbb{R}^N$, X has a nonempty relative interior, so we can always embed X in a (possibly lower-dimensional) Euclidean space such that X has a nonempty interior.

Theorem 7.A.1 in Shaked and Shanthikumar (1994)), a distribution $\pi_X \in \Delta(X)$ of moments is feasible (i.e., induced by some Bayes-plausible distribution of posterior beliefs) if and only if μ_0 is a mean-preserving spread of π_X. However, anticipating our results and following Kolotilin (2018), we will formulate the moment persuasion problem as optimization over joint distributions of moments and states. Formally, we call a distribution $\pi \in \Delta(X \times \Omega)$ *feasible*, denoted $\pi \in \Pi(\mu_0)$, if

$$\int_{X\times B} \mathrm{d}\pi(x, \omega) = \int_B \mathrm{d}\mu_0(\omega), \quad \text{for all measurable } B \subset \Omega,$$

$$\int_{B\times\Omega} (x-\omega)\, \mathrm{d}\pi(x, \omega) = 0, \quad \text{for all measurable } B \subset X.$$

The first equation is the Bayes-plausibility constraint, which says that the marginal distribution of states induced by π is μ_0. The second equation is the martingale constraint, which says that the conditional expectation $\mathbb{E}_\pi[\omega|x]$ induced by π is x.

We let π_X denote the marginal distribution of moments induced by π. The primal problem (P) simplifies to finding a joint distribution $\pi \in \Delta(X \times \Omega)$ to

$$\begin{aligned} &\text{maximize} \int_X v(x)\, \mathrm{d}\pi_X(x) \\ &\text{subject to } \pi \in \Pi(\mu_0). \end{aligned} \tag{$\mathrm{P_M}$}$$

When discussing intuitions, we will sometimes refer to π informally as a "signal."

4.2 *Prices for moments*

Our first major result of this section derives the implications of the general duality from Section 3 for the special case of moment persuasion.

Theorem 5. *Suppose that v is L-Lipschitz and fix an optimal solution $p : \Omega \to \mathbb{R}$ to the dual problem (D). There exists an extension $\bar{p} : X \to \mathbb{R}$ of p to X (i.e., p and $\bar{p}$ coincide on Ω) such that, for any optimal solution $\pi \in \Pi(\mu_0)$ to (P_M):*

(i) $\bar{p}$ is convex, L-Lipschitz, satisfies $\bar{p} \geq v$, and

$$\int_X v(x)\, \mathrm{d}\pi_X(x) = \int_\Omega \bar{p}(\omega)\, \mathrm{d}\mu_0(\omega);$$

(ii) there exists a measurable function $q : X \to \mathbb{R}^N$ such that $\|q(x)\| \leq L$ for all $x \in X$,

$$\bar{p}(y) = \sup_{x\in X} \big\{ v(x) + q(x) \cdot (y - x) \big\}, \quad \textit{for all } y \in X,$$

$$\bar{p}(\omega) = v(x) + q(x) \cdot (\omega - x), \quad \textit{for } \pi\textit{-almost all } (x, \omega).$$

Conversely, if there exists a feasible $\pi \in \Pi(\mu_0)$ and a price function $\bar{p} : X \to \mathbb{R}$ satisfying any one of conditions (i) or (ii), then π is optimal for (P_M). (The last claim is true under a weaker assumption that v is measurable and bounded.)

Theorem 5 provides sufficient and necessary conditions for optimality of a candidate solution $\pi \in \Pi(\mu_0)$. The main insight is that "prices for states" can be extended to "prices for moments." Additionally, condition *(i)* shows that optimal prices must be convex in moment persuasion. To see that intuitively, note that in our interpretation of the dual problem (D) from Section 2, a measure $\mu \in \Delta(\Omega)$ of resources and one unit of resource $x = \mathbb{E}_\mu[\omega]$ are now equivalent for the producer. If prices failed to be convex, the producer could sell at effectively higher prices by engaging in such "mean-preserving" transformations of the resources. Thus, the wholesaler offers convex prices to begin with.

Theorem 5 recovers (under a stronger assumption) the duality results for one-dimensional moment persuasion from Kolotilin (2018), Dworczak and Martini (2019), and Dizdar and Kováč (2020), and establishes strong duality for multidimensional moment persuasion. By providing the two conditions *(i)* and *(ii)* that are jointly necessary but individually sufficient, the theorem unifies two alternative approaches to moment persuasion. The price function from condition *(i)* is a direct analog of prices for moments in Dworczak and Martini (2019) who derive them as a multiplier on the mean-preserving spread constraint (represented in its integral form for the one-dimensional case). The price function from condition *(ii)*, along with the function q, are analogs of the dual variables from Kolotilin (2018) and Kolotilin, Corrao, and Wolitzky (2024) who derive them as multipliers on the two constraints defining the set $\Pi(\mu_0)$ of joint distributions of moments and states. In particular, q is the multiplier on the martingale constraint. Thus, the two existing duality formulations for moment persuasion are a consequence of two alternative representations of feasible distributions for the primal problem.[18] Theorem 5 shows that both formulations are a special case of our general duality, and that both can be extended to the multidimensional case.

Next, we give an overview of the proof of Theorem 5. Because we have guaranteed dual attainment (by the assumption that v is Lipschitz), there exists a solution p to the dual problem (D), and there is no duality gap: Equality (O) simplifies to

$$\int_X v(x)\,\mathrm{d}\pi_X(x) = \int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega),$$

for any π optimal for ($\mathrm{P_M}$). We can extend p (prices for states) from Ω to X (prices for moments) using the so-called "convex-roof" construction (Bucicovschi and Lebl (2013)):

$$\check{p}(x) := \inf\left\{\int_\Omega p(\omega)\,\mathrm{d}\mu(\omega) : \mu \in \Delta(\Omega), \int_\Omega \omega\,\mathrm{d}\mu(\omega) = x\right\}, \quad \text{for all } x \in X. \tag{R}$$

It is easy to show that $\check{p}$ is convex, $\check{p} \geq v$, and hence $\check{p}$ satisfies the constraint in (D). Moreover, by definition, $\check{p}$ is pointwise smaller than p on Ω. If we could show that $\check{p}$ is Lipschitz, then $\check{p}$ restricted to Ω would be a solution to the dual (D), and condition *(i)* in Theorem 5 would hold.

[18] In Appendix B.2 of Dworczak and Kolotilin (2024), we formally introduce the problem dual to ($\mathrm{P_M}$), show that the price function $\bar{p}$ from Theorem 5 is indeed a solution to that problem, and formalize the connection to previous duality formulations in Appendix B.3.

However, $\check{p}$ does not even have to be continuous when N—the dimension of the space of moments—is three or higher (even though p is Lipschitz).[19] There are moment-persuasion problems in which $p \in \text{Lip}(\Omega)$ solves (D) but its convex roof is discontinuous. Furthermore, for non-Lipschitz v, one can construct examples in which there does not exist *any* convex continuous extension of optimal prices for states to prices for moments. These cases help explain why our assumptions on the objective v are stronger than those imposed by Dworczak and Martini (2019) and Dizdar and Kováč (2020) in the one-dimensional case. In fact, the additional complications are a direct consequence of a multidimensional space of moments: It can be shown that $\check{p}$ is Lipschitz when Ω contains the boundary of X—a condition that holds trivially in the one-dimensional case.[20]

To circumvent these difficulties, we prove a lemma showing that the graph of $\check{p}$ can be separated by a hyperplane (with a properly bounded gradient, as captured by the function $q(x)$ from condition *(ii)*) from any point $(x, v(x))$ on the graph of the objective function v. We can then define a new price function $\bar{p}: X \to \mathbb{R}$ that is the supremum of all such hyperplanes. The resulting price function is a convex and Lipschitz extension of p that is "sandwiched" between $\check{p}$ and v. It follows that $\bar{p}$ solves (D) (viewed as a function on Ω) and that condition *(i)* of Theorem 5 holds. Additionally, using the function $q(x)$, we can show that the complementary-slackness condition (C) takes a particularly simple form described in condition *(ii)* of Theorem 5.

In the remainder of this section, we leverage Theorem 5 to derive structural properties of solutions to (P_M). Even though Theorem 5 guarantees existence of prices for moments, it does not provide a direct way to construct them. We show next that when v is continuously differentiable, we can take $q(x)$ from condition *(ii)* of Theorem 5 to be equal to the gradient of v at x on the support of any optimal π_X.

4.3 *Constructing solutions in the differentiable case*

To derive tighter implications of duality for the properties of optimal solutions, we further strengthen our assumptions on the objective function. We assume that v is continuously differentiable on X, and thus has a continuous gradient ∇v on X.[21] We will show that, in this case, solving the problem (P_M) can be reduced to finding the support of the optimal distribution of moments.

For any closed set $S \subset X$ (candidate support of the optimal distribution of moments), we define the function p_S on Ω by

$$p_S(\omega) := \max_{x \in S} \{ v(x) + \nabla v(x) \cdot (\omega - x) \}, \quad \text{for all } \omega \in \Omega. \tag{S}$$

[19]A careful reader might notice that this implies that some assumption of Berge's maximum theorem must be violated. Indeed, it turns out that the feasibility correspondence $\Phi(x) = \{\mu \in \Delta(\Omega) : \int_\Omega \omega \, d\mu(\omega) = x\}$ is not necessarily lower hemicontinuous in $\mathbb{R}^N$ for $N > 2$. However, because Φ is an upper hemicontinuous correspondence, $\check{p}$ is lower semicontinuous, by Lemma 17.30 in Aliprantis and Border (2006).

[20]Formal arguments supporting the claims made in this paragraph can be found in Appendix B.4 of Dworczak and Kolotilin (2024).

[21]We say that v is differentiable at $x \in X$ if there exists a gradient $\nabla v(x) \in \mathbb{R}^N$ such that $f(x+h) = f(x) + \nabla v(x) \cdot h + o(\|h\|)$ for all $h \in \mathbb{R}^N$ such that $x + h \in X$, in which case $\nabla v(x)$ is unique.

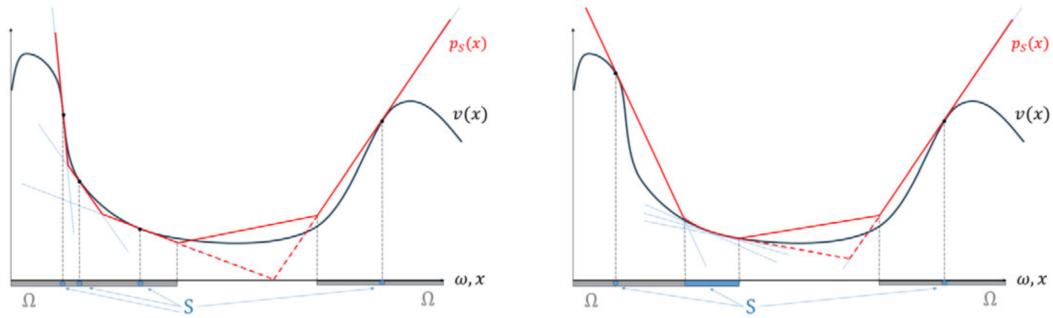

FIGURE 2. The construction of the function $p_S(x)$ in the one-dimensional case. The gray area in the x-axis represents the nonconvex domain Ω. The left panel depicts the price function induced by a suboptimal, discrete set S (indicated in blue)—the induced price function fails the first condition in (M). The right panel depicts the price function induced by a set S that satisfies condition (M) (for some π). The dashed red line depicts the extension of the function $p_S(\omega)$ from Ω to X obtained by applying formula (S) outside of Ω, while the red solid line is obtained by applying the convex-roof construction.

In case Ω is not convex, we extend p_S from Ω to X using the convex-roof construction:[22]

$$p_S(x) := \inf\left\{ \int_\Omega p_S(\omega)\, \mathrm{d}\mu(\omega) : \mu \in \Delta(\Omega), \int_\Omega \omega\, \mathrm{d}\mu(\omega) = x \right\}, \quad \text{for all } x \in X \setminus \Omega.$$

Figure 2 illustrates these definitions: The function $p_S(\omega)$ is found as the maximum over hyperplanes tangent to the graph of the function v at points in the set S. The convex-roof construction extends p_S from Ω to X by minimizing the value achieved at any $x \in X \setminus \Omega$ of any convex combination of points belonging to the graph of p_S on Ω.

Finally, for any feasible $\pi \in \Pi(\mu_0)$, consider the condition:

$$\begin{aligned} & p_S(x) \geq v(x), \quad \text{for all } x \in X, \\ & p_S(\omega) = v(x) + \nabla v(x) \cdot (\omega - x), \quad \text{for all } (x, \omega) \in \operatorname{supp}(\pi). \end{aligned} \tag{M}$$

The following theorem connects condition (M) to optimality of π.

THEOREM 6. *Suppose that v is continuously differentiable. A joint distribution $\pi \in \Pi(\mu_0)$ is an optimal solution to (P_M) if and only if condition (M) holds with $S = \operatorname{supp}(\pi_X)$.*

PROOF. The proof is relegated to Appendix A.8. □

Theorem 6 gives rise to a "guess and verify" procedure that can be used to identify optimal solutions to (P_M). The "guess" involves conjecturing the optimal support S of

[22] Note that because p_S is convex on Ω by definition, it does not matter whether we use the convex roof for $x \in X \setminus \Omega$ or for all $x \in X$. The reader might be surprised that we rely on the convex roof construction after arguing that it sometimes fails to properly extend prices for states to the prices for moments. And indeed, the price function p_S we construct does not necessarily satisfy all of the conditions of Theorem 5. Nevertheless, it turns out that p_S satisfies the conditions that are relevant for deriving properties of optimal solutions to (P_M), which is our ultimate goal.

moments. Fixing S, prices p_S can be computed mechanically, and then condition (M) becomes necessary and sufficient for optimality of π with support S.

In general, different solutions to (P_M) may have different supports S of posterior moments. However, duality implies that one can define a maximal set $S^\star$ of posterior moments that can be induced by an optimal signal. In other words, any optimal signal must induce posterior moments that belong to $S^\star$. Moreover, this set $S^\star$ can be easily found as long as we have one solution to (P_M)—we formalize this in the following remark.

Remark 1. Suppose that $\pi^\star \in \Pi(\mu_0)$ is optimal for (P_M), and let

$$S^\star = \left\{x \in X : p_{\text{supp}(\pi^\star_X)}(x) = v(x)\right\}.$$

Then $\pi \in \Pi(\mu_0)$ is optimal for (P_M) if and only if $\text{supp}(\pi_X) \subset S^\star$ and condition (M) holds with $S = S^\star$.[23]

Proof. The proof is relegated to Appendix A.8. □

Remark 1 is useful when proving uniqueness and characterizing properties of an optimal solution. We turn to these issues next.

4.4 *Structure of solutions*

In this subsection, we focus on deriving the implications of Theorem 6 for the structure of optimal solutions to (P_M). We provide a condition under which there exists a unique optimal solution π to (P_M) that partitions the state space into convex sets, and pools the states in each element of the partition. This is a natural extension of the idea of monotone-partitional solutions from one-dimensional moment persuasion to the multidimensional case. We also generalize a result proven by Arieli et al. (2023) and Kleiner, Moldovanu, and Strack (2021): In the one-dimensional case, there exists an optimal signal $\pi \in \Pi(\mu_0)$ with a bipooling structure. We derive a multidimensional analog of this property.

To simplify exposition and obtain tighter results, we assume that Ω is a convex set (so that $\Omega = X$). In Appendix A.9, we extend the analysis to the general case.

4.4.1 *Optimality of convex-partitional signals*

We first address the problem of when it is without loss of optimality to restrict attention to convex-partitional signals. Formally, we say that $\pi \in \Pi(\mu_0)$ is *convex-partitional* if there is a measurable function $\chi : \Omega \to X$ such that, for all measurable sets $A \subset X$ and $B \subset \Omega$,

$$\pi(A, B) = \int_B \mathbf{1}\left\{\chi(\omega) \in A\right\} \mathrm{d}\mu_0(\omega),$$

and the set $\chi^{-1}(x)$ is convex for all x. Intuitively, χ represents a distribution that pools all states in $\chi^{-1}(x)$ into the moment x.

[23] It is easy to see that $p_S \geq v$ in this case, so only the second condition in (M) is relevant.

Theorem 7. *Suppose that v is continuously differentiable and that μ_0 has a density on Ω with respect to the Lebesgue measure.*[24] *Suppose there do not exist distinct $x, y \in X$ with*

$$\nabla v(x) = \nabla v(y),$$
$$v(x) - \nabla v(x) \cdot x = v(y) - \nabla v(y) \cdot y,$$
$$\lambda v(x) + (1-\lambda) v(y) \geq v\big(\lambda x + (1-\lambda) y\big), \quad \textit{for all } \lambda \in [0, 1].$$

Then there is a unique optimal solution to (P_M), and that solution is convex-partitional.

Proof. The proof is relegated to Appendix A.10. □

Theorem 7 gives an easy-to-verify condition on the objective function v under which the optimal distribution is unique and convex-partitional. The condition can be seen as an extension of the affine-closure property from Dworczak and Martini (2019) that guarantees optimality of a monotone partition in the one-dimensional case.

In Appendix B.5 of Dworczak and Kolotilin (2024), we state a version of Theorem 7 that imposes a slightly weaker sufficient condition, which turns out to be necessary; if that weaker condition fails, then for at least some priors there exist optimal signals that are not convex-partitional. To the best of our knowledge, these results provide the most permissive conditions guaranteeing a convex-partitional signal for multidimensional moment persuasion. Prior to the current version of this paper, Malamud and Schrimpf (2022) obtained a stronger sufficient condition (requiring that $\nabla v(x) \neq \nabla v(y)$ for $x \neq y$).

In the remainder of this subsection, we give an overview of the proof of Theorem 7. The first part of the proof investigates the structure of optimal solutions, and does not rely on any of the assumptions of Theorem 7. Thus, our goal in the overview is to present these additional results; they will be useful for subsequent analysis. The second part of the proof gives an explicit construction of the elements of the optimal convex partition from Theorem 7.

We begin by introducing some additional notation. Fix an optimal solution $\pi^\star \in \Pi(\mu_0)$ to (P_M), and define the set $S^\star$ as in Remark 1:

$$S^\star := \big\{x \in X : p_{\mathrm{supp}(\pi^\star_X)}(x) = v(x)\big\}.$$

Recall that we can interpret $S^\star$ as the maximal set of posterior moments that can be induced by an optimal solution. To simplify notation, let $p^\star(x) := p_{S^\star}(x)$, for all $x \in X$. Next, we define the set Γ that encodes the second property in condition (M):

$$\Gamma := \big\{(x, \omega) \in S^\star \times \Omega : p^\star(\omega) = v(x) + \nabla v(x) \cdot (\omega - x)\big\}.$$

The set Γ is called the contact set in the linear programming literature. In light of Theorem 6 and Remark 1, a feasible $\pi \in \Pi(\mu_0)$ is optimal if and only if $\mathrm{supp}(\pi) \subset \Gamma$. Finally,

[24] The assumption that μ_0 is a continuous distribution allows us to circumvent the thorny issue of how to define a convex partition when there are atoms in the distribution of states—in this case, some of the atoms may need to be split among multiple elements of the partition.

we define the x-section of Γ,

$$\Gamma_x := \{\omega \in \Omega : (x, \omega) \in \Gamma\}.$$

Intuitively, the set Γ_x contains states that can appear together with x in the support of an optimal solution—states in Γ_x (and only these states) can be pooled into the moment x. Geometrically, Γ_x is the projection of the face of the epigraph of $p^\star$ exposed by the direction $(-1, \nabla v(x))$ on the state space, $\Gamma_x = \arg\max_{\omega\in\Omega}\{\nabla v(x) \cdot \omega - p^\star(\omega)\}$. A more intuitive statement of this property is that states can be pooled (in an arbitrary way as long as the induced posterior moments belong to $S^\star$) within regions where the price function is affine; at the same time, the optimal solution cannot pool together states that do not belong to a region on which $p^\star$ is affine. We can thus think of Γ_x as the "pooling region" of moment x.

The sets Γ_x can intersect in general. If $\omega \in \Gamma_x \cap \Gamma_y$, then ω could appear in the support of π both conditional on x and conditional on y—this is possible when the signal is random conditional on ω. However, an important consequence of the above geometric characterization is that each Γ_x is convex, and that $\mathrm{relint}(\Gamma_x) \cap \mathrm{relint}(\Gamma_y) \neq \emptyset$ implies $\Gamma_x = \Gamma_y$, where $\mathrm{relint}(\cdot)$ stands for the relative interior of a set. Thus, the set Γ generates a partition of Ω consisting of relatively open convex components $\{\mathrm{relint}(\Gamma_x)\}_{x\in S^\star}$ and the set of points on the boundaries of these components: $X \setminus \bigcup_{x\in S^\star} \mathrm{relint}(\Gamma_x)$. If $x \neq y$ implies that $\Gamma_x \neq \Gamma_y$, then π has a very simple structure: For any $x \in S^\star$, states in $\mathrm{relint}(\Gamma_x)$ are pooled together into the posterior mean x.

This is where the conditions of Theorem 7 come in. When the conditions on v hold, it is indeed true that $x \neq y$ implies that $\Gamma_x \neq \Gamma_y$. When μ_0 has a continuous distribution, we can ignore the measure-zero set of states on the boundaries of the convex elements of the partition. Thus, a convex-partitional signal is optimal. Moreover, the optimal $\chi : X \to X$ is uniquely determined, for μ_0-almost all $\omega \in \Omega$, by

$$\chi(\omega) = \{x \in S^\star : \omega \in \Gamma_x\} = \{x \in S^\star : \nabla p^\star(\omega) = \nabla v(x)\}.$$

We illustrate this discussion with an application in the next section.

4.4.2 *Beyond convex-partitional signals* In this subsection, we turn attention to the structure of solutions when the conditions of Theorem 7 fail. In the one-dimensional case, the bipooling result of Arieli et al. (2023) and Kleiner, Moldovanu, and Strack (2021) shows that even if no optimal signal is monotone-partitional, there still exist optimal signals with a relatively simple structure. Namely, the state space is partitioned into intervals, and conditional on any interval, an additional binary signal may be sent. We will derive a multidimensional version of this result. Our generalization is a direct consequence of duality, while Arieli et al. (2023) and Kleiner, Moldovanu, and Strack (2021) rely on an extreme-point characterization of optimal signals.

For a set $A \subset X$, let $\mathrm{cl}(A)$ denote the closure of A, and $\mathrm{ext}(A)$ denote the set of extreme points of the closed convex hull of A. Fixing a solution π to ($\mathrm{P_M}$), let

$$S_x := \mathrm{cl}(\mathrm{supp}(\pi_X) \cap \mathrm{relint}(\Gamma_x)),$$

for any $x \in \mathrm{supp}(\pi_X)$. Recall that Γ_x is the set of states that can be pooled into the posterior moment x by an optimal signal. Thus, conditional on x being the realized posterior moment under some optimal signal π, the set S_x contains all posterior moments in the support of π_X that could be generated by an optimal signal. For example, if the conditions of Theorem 7 hold, then the (unique) optimal signal π satisfies $S_x = \{x\}$ for almost all $x \in \mathrm{supp}(\pi_X)$. This means that any state in the support of the optimal signal conditional on x must be pooled into x; thus, the optimal signal is deterministic (and convex-partitional since each Γ_x is convex). The bipooling result of Arieli et al. (2023) and Kleiner, Moldovanu, and Strack (2021) in the one-dimensional case can be reformulated as stating that there exists an optimal solution such that S_x has at most two elements. That is, for any realized posterior moment x, there exists at most one other posterior moment $y \in \mathrm{supp}(\pi_X)$ such that $\Gamma_x = \Gamma_y$. In this case, states in the interval Γ_x can be pooled into either x or y, and we have $S_x = S_y = \{x, y\}$. The following result extends that conclusion to the multidimensional case.

Theorem 8. *Suppose that v is continuously differentiable and that μ_0 has a density on Ω with respect to the Lebesgue measure. There exists an optimal solution $\pi \in \Pi(\mu_0)$ to (P_M) such that $S_x = \mathrm{ext}(S_x)$ for π_X-almost all x.*

Proof. The proof is relegated to Appendix A.10. □

The conclusion $S_x = \mathrm{ext}(S_x)$ means that no posterior mean in S_x can be expressed as a convex combination of other posterior means in S_x. This generalizes the bipooling result of Arieli et al. (2023) and Kleiner, Moldovanu, and Strack (2021) because in the one-dimensional case, for any set $S \subset \mathbb{R}$, $|\mathrm{ext}(S)| \leq 2$. In higher dimensions, Theorem 8 guarantees that we can divide the state space into convex "pooling regions" (up to a measure-zero set) and find an optimal signal that only pools states inside pooling regions; moreover, the posterior moments induced from a given pooling region form a set that only consists of extreme points (of its own convex hull).

The proof of Theorem 8 relies on the fact that $\mathrm{supp}(\pi) \subset \Gamma$ is both necessary and sufficient for the optimality of $\pi \in \Pi(\mu_0)$. As shown in Section 4.4.1, Γ defines (up to a measure zero set) a convex partition of the state space, with a representative element Γ_x, which could in general coincide with Γ_y for $y \neq x$. That is, optimality of a signal requires that states in Γ_x are mapped only into posterior moments y for which $\Gamma_y = \Gamma_x$. We can modify the solution on Γ_x and it will remain optimal as long as we preserve the above property. Formally—to deal with the fact that sets Γ_x may have measure zero—we introduce an auxiliary optimization problem in which we minimize the average norm of induced posterior moments subject to maintaining the condition $\mathrm{supp}(\pi) \subset \Gamma$. The auxiliary problem then picks an optimal solution in which $S_x = \mathrm{ext}(S_x)$ must be satisfied, as otherwise the value of the auxiliary problem could be lowered by shifting probability mass toward some posterior mean $y \in S_x$ that can be expressed as a convex combination of other posterior means in S_x.

For one-dimensional problems, the geometric property $S_x = \mathrm{ext}(S_x)$ implies the cardinality restriction $|S_x| \leq 2$. This is no longer the case when the dimension N of the state space is two or more. In fact, one can construct an example in which S_x is infinite for any choice of optimal π.[25] The example implies that our result is tight if one works with the partition of the state space defined by the price function through the contact set Γ, as is implicitly assumed in our definition of S_x. However, that partition may sometimes be unnecessarily coarse; intuitively, the price function may be affine over a region that could be further subdivided into smaller "pooling regions" (sets of states that are only pooled with one another but not with states from other pooling regions). Obłój and Siorpaes (2017) and De March and Touzi (2019) show how to define the finest partition into pooling regions relying directly on the distribution of posterior moments.[26] If one defines an analog of S_x for the finest partition (by replacing Γ_x in the definition of S_x by the element of the finest partition containing x), then it might be possible to tighten the conclusion of Theorem 8, perhaps by showing that there are at most $N+1$ posterior means induced from every pooling region (as is loosely suggested by Carathéodory's theorem). Since duality does not seem immediately useful in pursuing this direction, we leave it for future research.

5. Application: Quadratic objective

In this section, we show how our duality approach developed in the preceding section can be used to solve a class of persuasion problems in which μ_0 has a density on Ω that is a compact convex set in $\mathbb{R}^2$ with nonempty interior (so that $\Omega = X$), the objective function depends on a pair of moments $x = (x_1, x_2)$, and $v(x)$ is a quadratic form: $v(x) = x\Lambda x^T$.

Variants of this model received considerable attention in the literature. The case $v(x) = x_1 x_2$ is equivalent to the model of Rayo and Segal (2010), who analyzed it under the assumption that Ω is a finite set. Nikandrova and Pancs (2017) studied this problem under the assumption that Ω is a strictly convex curve. These two papers mostly focus on deriving necessary conditions for optimality.[27] Tamura (2018) considers the case where v is a general quadratic form in $\mathbb{R}^N$ but imposes strong symmetry assumptions on the prior distribution. Kramkov and Xu (2022) consider a problem (inspired by the insider trading model of Rochet and Vila (1994)) that turns out to be mathematically equivalent to a generalized version of our problem where the assumption $\Omega = X$ is not imposed—their analysis is limited in its economic predictions since their methods are designed to handle even fairly pathological distributions of the state. Our marginal contribution is to provide a tighter characterization of optimal solutions for the well-behaved case when Ω is a compact convex set (i.e., when $\Omega = X$). Relative to Rayo and

[25] We provide one such example in Appendix B.6 of Dworczak and Kolotilin (2024).

[26] In the one-dimensional case, their construction can be understood through the integral characterization of mean-preserving spreads: An element of a partition (in this case, an interval) is pinned down by two consecutive points at which the integral constraint binds. In the multidimensional case, the construction is significantly more complicated since there exists no convenient representation of mean-preserving spreads.

[27] Rayo (2013) and Onuchic and Ray (2023) restrict attention to monotone partitional signals in the setting of Nikandrova and Pancs (2017).

Segal (2010) and Nikandrova and Pancs (2017), we show that a set of necessary conditions taken from these two papers become jointly sufficient for optimality in our case. Prior to the current version of this paper, Malamud and Schrimpf (2022) provided an alternative (less explicit) characterization of solutions under weaker assumptions.

We first argue that the case of a general quadratic form can easily be reduced to the special case $v(x) = x_1 x_2$. Indeed, for any quadratic form, there exists a basis such that the quadratic form is diagonal: $v(x) = \lambda_1 x_1^2 + \lambda_2 x_2^2$. If $\lambda_1, \lambda_2 \geq 0$ (resp., $\lambda_1, \lambda_2 \leq 0$), then full disclosure (resp., no disclosure) is optimal. If λ_1 and λ_2 have opposite signs, then there exists yet another basis such that $v(x) = x_1 x_2$, which we assume henceforth.

It is known from Rayo and Segal (2010) that the posterior means induced by an optimal signal must belong to a monotone set. Using duality, we can establish a stronger claim. Formally, we will say that a set $S \subset X$ is

- *monotone* if $(x_1 - y_1)(x_2 - y_2) \geq 0$, for all $x, y \in S$;
- *maximal monotone* in X if it is monotone, and for each $y \in X \setminus S$, there exists $x \in S$ such that $(x_1 - y_1)(x_2 - y_2) < 0$.
- *almost-maximal monotone* in X if it is monotone, compact, and, for each $y \in X \setminus S$, there exists $x \in S$ such that $(x_1 - y_1)(x_2 - y_2) \leq 0$.

Intuitively, a monotone set S in $\mathbb{R}^2$ has the property that if $x \in S$, then S cannot intersect the interiors of either the upper-left or the lower-right quadrants centered at x. A monotone set is maximal in X if it is not a proper subset of any monotone set in X. A maximal monotone set must be compact (when X is compact, as assumed). An almost-maximal monotone set S is a compact subset of a maximal monotone set S' such that $S' \setminus S$ is a collection of open line segments that are either horizontal or vertical.

Proposition 1. *If $\pi^\star \in \Pi(\mu_0)$ is optimal, then the support of moments* $\text{supp}(\pi_X^\star)$ *induced by $\pi^\star$ is an almost-maximal monotone set in X.*

Proof. Suppose that $\pi^\star \in \Pi(\mu_0)$ is optimal. To simplify notation, let $p^\star := p_{\text{supp}(\pi_X^\star)}$ as defined by (S). By Remark 1, $p^\star \geq v$ and $\text{supp}(\pi_X^\star) \subset S^\star$, where $S^\star = \{x \in X : p^\star(x) = v(x)\}$; moreover, $p^\star = p_{S^\star}$, and hence, since $\Omega = X$ and $v(x) = x_1 x_2$, we have, for all $x \in X$,

$$p^\star(x) = \max_{y \in S^\star} \{x_1 y_2 + x_2 y_1 - y_1 y_2\}.$$

We claim that the set $S^\star$ is monotone: Otherwise, we would have $x, y \in S^\star$ such that $(x_1 - y_1)(x_2 - y_2) < 0$, but then

$$p^\star(x) \geq x_1 x_2 - (x_1 - y_1)(x_2 - y_2) > x_1 x_2 = v(x),$$

contradicting that $x \in S^\star$. Next, we claim that the set $S^\star$ is maximal monotone in X. Otherwise, there would exist $x \in X \setminus S^\star$ such that $(x_1 - y_1)(x_2 - y_2) \geq 0$ for all $y \in S^\star$, and thus

$$p^\star(x) = \max_{y \in S^\star} \big\{ x_1 x_2 - (x_1 - y_1)(x_1 - y_2) \big\} \leq x_1 x_2 = v(x).$$

But then, since $p^\star \geq v$, we would have that $p^\star(x) = v(x)$, contradicting that $x \notin S^\star$.

Since $\text{supp}(\pi_X^\star) \subset S^\star$, and we have shown that $S^\star$ is a monotone set, $\text{supp}(\pi_X^\star)$ is also a monotone set. Finally, we claim that $\text{supp}(\pi_X^\star)$ is almost-maximal monotone in X. Otherwise, there would exist $x \in X$ such that $(y_1 - x_1)(y_2 - x_2) > 0$ for all $y \in \text{supp}(\pi_X^\star)$, which implies that (since $\text{supp}(\pi_X^\star)$ is compact)

$$p^\star(x) = \max_{y \in \text{supp}(\pi_X^\star)} \left\{x_1 x_2 - (x_1 - y_1)(x_1 - y_2)\right\} < x_1 x_2 = v(x),$$

contradicting that $p^\star \geq v$. □

In light of Remark 1, the proof of Proposition 1 implies that the optimal price function can always be derived from some candidate support S of the distribution of moments that is a maximal monotone set. A natural class of maximal monotone sets in X are graphs of continuous increasing functions. The main result of this section describes necessary and sufficient conditions for the optimality of a solution $\pi^\star \in \Pi(\mu_0)$ with $\text{supp}(\pi_X^\star)$ equal to the graph $\text{Gr}(f)$ of a given well-behaved function f. By Theorem 7, the unique solution is convex-partitional; the optimal partition divides Ω into negatively-sloped line segments; a line segment that induces the posterior mean $(t, f(t))$ has slope $-f'(t)$, as illustrated in Figure 3. These observations are formalized in the following proposition.

Proposition 2. *Let* $f : [\underline{x}_1, \overline{x}_1] \to \mathbb{R}$ *be a twice continuously differentiable function, with* $f'(t) > 0$ *for all* $t \in [\underline{x}_1, \overline{x}_1]$*, such that the graph* $\text{Gr}(f)$ *is a maximal monotone subset of* X*. An optimal* $\pi^\star \in \Pi(\mu_0)$ *induces a support of moments* $\text{supp}(\pi_X^\star)$ *equal to* $\text{Gr}(f)$ *if and only if* Ω *can be partitioned, up to a measure zero set,*[28] *into a collection of disjoint open line segments* $\{I_t\}_{t \in [\underline{x}_1, \overline{x}_1]}$ *such that:*

(i) $\mathbb{E}[\omega | \omega \in I_t] = (t, f(t))$*, for almost all* $t \in [\underline{x}_1, \overline{x}_1]$*;*[29]

(ii) $I_t = \text{relint}(\{\omega \in \Omega : t \in \arg\max_{s \in [\underline{x}_1, \overline{x}_1]}\{\omega_1 f(s) + \omega_2 s - s f(s)\}\})$*, for all* $t \in [\underline{x}_1, \overline{x}_1]$.

Whenever the above conditions hold, the optimal signal is convex-partitional and pools the states within each I_t*; moreover,* $I_t \subseteq \{\omega \in \Omega : \omega_2 = f(t) - f'(t)(\omega_1 - t)\}$*, for all* $t \in [\underline{x}_1, \overline{x}_1]$.

Proof. We will prove that existence of the required partition of Ω is sufficient for optimality of the corresponding $\pi^\star$. We relegate the more technical proof of necessity to Appendix A.11.

[28] That is, $\Omega \setminus \{\bigcup_{t \in [\underline{x}_1, \overline{x}_1]} I_t\}$ has zero (Lebesgue) measure.

[29] Since I_t has zero measure under the prior, $\mathbb{E}[\omega | \omega \in I_t]$ is formally defined almost everywhere via the conditional expectation of ω conditional on a σ-algebra generated by $\{I_t\}_{t \in [\underline{x}_1, \overline{x}_1]}$. We provide an explicit formula for the conditional expectation in Appendix B.7 of Dworczak and Kolotilin (2024).

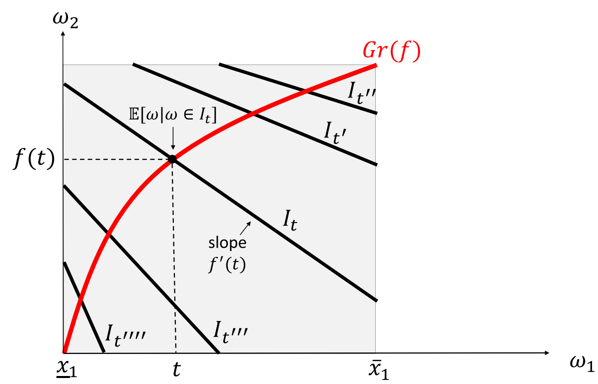

FIGURE 3. Illustration of Proposition 2: The optimal signal pools all states in each of the negatively sloped intervals I_t, and the resulting posterior means belong to $\mathrm{Gr}(f)$.

Suppose that there exists a collection of line segments $\{I_t\}_{t\in[\underline{x}_1,\overline{x}_1]}$ satisfying conditions *(i)* and *(ii)*. We can define $\pi^\star \in \Pi(\mu_0)$ as the convex-partitional signal that pools states in each I_t (it is irrelevant how the signal is defined for $\omega \in \Omega$ not belonging to any I_t). By condition *(i)*, the induced posterior-mean curve $\mathrm{supp}(\pi^\star_X)$ is equal to $\mathrm{Gr}(f)$. Following Section 4.3, define the price function

$$p_{\mathrm{Gr}(f)}(x) = \max_{y\in\mathrm{Gr}(f)} \big\{v(y) + \nabla v(y)\cdot(x-y)\big\} = \max_{t\in[\underline{x}_1,\overline{x}_1]} \big\{x_1 f(t) + x_2 t - t f(t)\big\}.$$

We will verify that condition (M) holds; optimality of $\pi^\star$ will then follow from Theorem 6. First, we argue that $p_{\mathrm{Gr}(f)}(x) \geq v(x)$, for all $x \in X$. It suffices to show that there exists a $t \in [\underline{x}_1, \overline{x}_1]$ such that $x_1 f(t) + x_2 t - t f(t) \geq x_1 x_2$, or equivalently, $(t - x_1)(f(t) - x_2) \leq 0$. The claim is obvious when $x \in \mathrm{Gr}(f)$, and follows from the fact that $\mathrm{Gr}(f)$ is maximal monotone in X when $x \in X \setminus \mathrm{Gr}(f)$. To complete the proof that (M) holds, note that, by condition *(ii)*, for almost all $\omega \in I_t$,

$$p_{\mathrm{Gr}(f)}(\omega) = v\big(x(t)\big) + \nabla v\big(x(t)\big)\cdot\big(\omega - x(t)\big) = \omega_1 f(t) + \omega_2 t - t f(t).$$

This shows that the equality in (M) holds for all $(x, \omega) \in \bigcup_{t\in[\underline{x}_1,\overline{x}_1]}((t, f(t)) \times I_t)$; by continuity, the equality extends to the closure of this set, which is $\mathrm{supp}(\pi^\star)$.

Finally, the inclusion $I_t \subseteq \{\omega \in \Omega : \omega_2 = f(t) - f'(t)(\omega_1 - t)\}$, for $t \in (\underline{x}_1, \overline{x}_1)$, follows from the observation that, by condition *(ii)*, the first-order condition $(\omega_1 - t)f'(t) + \omega_2 - f(t) = 0$ must hold for all $\omega \in I_t$.[30] For $t \in \{\underline{x}_1, \overline{x}_1\}$, the proof of the inclusion is more complicated, and thus relegated to Appendix A.11. □

Proposition 2 provides a tight characterization of optimal signals under the additional regularity requirement that the induced posterior mean curve is sufficiently regular (a graph of a twice differentiable function). If an optimal signal $\pi^\star$ induces

[30] This observation shows that it would suffice to require $\mathbb{E}[\omega_1|\omega \in I_t] = t$ in condition *(i)* in Proposition 1. Indeed, $(\omega_1 - t)f'(t) + \omega_2 - f(t) = 0$ for all $t \in (\underline{x}_1, \overline{x}_1)$ and $\omega \in I_t$ implies that $(\mathbb{E}[\omega_1|\omega \in I_t] - t)f'(t) + \mathbb{E}[\omega_2|\omega \in I_t] - f(t) = 0$, from which the second required equality $\mathbb{E}[\omega_2|\omega \in I_t] = f(t)$ follows.

$\text{supp}(\pi_X^\star) = \text{Gr}(f)$, then it must have a simple convex-partitional structure in which only states belonging to negatively-sloped line segments I_t are pooled together. Moreover, the slopes of these line segments are uniquely pinned down by f. The full proof in Appendix A.11 additionally reveals that the closures of these line segments can only intersect at the endpoints; the endpoints can be found by solving the optimization problem in condition *(ii)* in Proposition 2.

As an illustration, we provide conditions under which it is optimal to reveal only some linear combination of ω_1 and ω_2. A simple implication of this characterization is that it is optimal to reveal $\omega_1 + \omega_2$ if the prior is symmetric around the line $\omega_2 = \omega_1$.

Proposition 3. *The joint distribution* $\pi \in \Pi(\mu_0)$ *induced by the disclosure of the realization of* $a\omega_1 + \omega_2$*, with* $a > 0$*, is optimal if and only if* $\text{supp}(\pi_X) \subset \{(t, at+b) : t \in \mathbb{R}\}$*, with* $b \in \mathbb{R}$*.*

Proof. *If.* Let $\pi \in \Pi(\mu_0)$ be induced by disclosure of the realization of $a\omega_1 + \omega_2$, and suppose that $\text{supp}(\pi_X) \subset \{(t, at+b) : t \in \mathbb{R}\}$. Note that π partitions Ω into parallel open line segments $\{I_t\}_{t \in [\underline{x}_1, \overline{x}_1]}$, where $I_t = \text{relint}(\{\omega \in \Omega : a\omega_1 + \omega_2 = 2at + b\})$, and the range $[\underline{x}_1, \overline{x}_1]$ is defined by the property that $(t, at+b) \in \Omega$. Since $\text{supp}(\pi_X) \subset \{(t, at+b) : t \in \mathbb{R}\}$, the induced posterior mean curve is a line segment with slope a that is a monotone maximal set in Ω. Finally, condition *(ii)* in Proposition 2 holds since

$$\Big\{\omega \in \Omega : t \in \operatorname*{arg\,max}_{s \in [\underline{x}_1, \overline{x}_1]} \big\{\omega_1(as+b) + \omega_2 s - s(as+b)\big\}\Big\} = \{\omega \in \Omega : a\omega_1 + \omega_2 = 2at + b\},$$

which is precisely our definition of I_t. Thus, Proposition 2 shows that π is optimal.

Only if. Here, we prove the necessity part under a regularity condition that the support of π_X corresponding to disclosure of the realization of $a\omega_1 + \omega_2$ is a twice continuously differentiable function f with $f' > 0$; we relegate the complete proof (without any regularity condition) to Appendix A.12. If disclosing $a\omega_1 + \omega_2$ is optimal, then the open line segments I_t partitioning Ω (whose existence is guaranteed by Proposition 2 under the regularity condition) must be parallel and have slope $-a$. But then we must have that $\omega_2 = f(t) - f'(t)(\omega_1 - t)$ if and only if $\omega_2 = 2at + b - a\omega_1$, for some b, which is only possible when $f(t) = at + b$. □

Proposition 3 showcases two ways in which Proposition 2 can be used. First, it can be applied to verify the optimality of a conjectured posterior mean curve. Once a posterior mean curve is fixed, Proposition 2 allows us to construct the unique candidate solution, and then check whether it is indeed optimal. Second, Proposition 2 provides a way to construct the optimal signal. Suppose that we partition Ω (up to a measure-zero set) into negatively-sloped open line segments in such a way that pooling the states within these line segments induces a posterior mean curve that is a graph of some continuous function f. Then this signal is optimal as long as condition *(ii)* holds. Moreover, if f is differentiable and the closures of these line segments are disjoint, then it suffices to verify that the slope of the line segment inducing posterior mean $(t, f(t))$ is $-f'(t)$.

Finally, we offer an intuition for our results. We can rewrite the objective function as

$$v(\omega) = \omega_1\omega_2 = \frac{1}{a}\left[\left(\frac{a\omega_1 + \omega_2}{2}\right)^2 - \left(\frac{a\omega_1 - \omega_2}{2}\right)^2\right].$$

Thus, intuitively, the objective is to disclose as much information as possible about $a\omega_1 + \omega_2$ while disclosing as little as possible about $a\omega_1 - \omega_2$. Typically, $a\omega_1 + \omega_2$ and $a\omega_1 - \omega_2$ will be correlated, leading to a trade-off. However, when $\mathbb{E}[a\omega_1 - \omega_2 | a\omega_1 + \omega_2] = \mathbb{E}[a\omega_1 - \omega_2]$, (so that disclosing $a\omega_1 + \omega_2$ does not change the expectation of $a\omega_1 - \omega_2$), it becomes optimal to disclose $a\omega_1 + \omega_2$. The condition $\text{supp}(\pi_X) \subset \{(t, at + b) : t \in \mathbb{R}\}$ states precisely that $\mathbb{E}[\omega_2 | a\omega_1 + \omega_2] = a\mathbb{E}[\omega_1 | a\omega_1 + \omega_2] + b$. Proposition 3 shows that this intuitive condition is not only sufficient but also necessary for the optimality of disclosing $a\omega_1 + \omega_2$. Note that no correlation between $a\omega_1 - \omega_2$ and $a\omega_1 + \omega_2$ requires that $a = \text{sd}(\omega_2)/\text{sd}(\omega_1)$ (where sd stands for standard deviation) implying that the optimal weight equalizes the contribution of the two states to the variability of the signal. The general case, covered by Proposition 2, can be understood as setting the weight a locally, as captured by the condition that the slope of I_t must be equal to $-f'(t)$.

6. Concluding remarks

We conclude with a few remarks on extensions and connections to other problems.

Potential applications Several other potential applications of persuasion duality are worth mentioning. Galperti, Levkun, and Perego (2024) show that duality can be used to quantify the value of "data records"; our results could thus be helpful in calculating that value. Benoît and Dubra (2011), Yang and Zentefis (2024), and Kolotilin and Wolitzky (2024) characterize the set of feasible distributions of posterior *quantiles*; it might be interesting to study the consequences of general duality for the special case of "quantile persuasion"—paralleling the developments for moment persuasion. Finally, a large literature on rational inattention and costly-information acquisition studies optimization problems in which a linear objective is maximized over distributions of posterior beliefs subject to Bayes' plausibility. Our analysis applies under the assumption that the cost of information satisfies posterior-separability (see, among many others, Caplin and Dean (2013, 2015), Gentzkow and Kamenica (2014), and Denti (2022)).

Additional constraints in the persuasion problem Doval and Skreta (2024), inspired by an earlier contribution of Le Treust and Tomala (2019), observe that many persuasion problems feature additional linear constraints (such as moral-hazard, inventive-compatibility, or capacity constraints) that modify the structure of optimal persuasion schemes. Our general duality approach easily accommodates a finite number M of additional linear constraints: In this case, there are M new prices that enter the objective function in the dual problem (D) (see an earlier version of the paper Dworczak and Kolotilin (2019), for details).

Such an extension could be useful in analyzing problems with a privately informed Receiver (see, among others, Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017) and Guo and Shmaya (2019)). Candogan and Strack (2023) point out that the one-dimensional moment persuasion problem with a privately informed Receiver reduces to the standard one-dimensional moment persuasion problem with additional linear constraints. It would be interesting to see if duality could be fruitfully applied to such a representation of the informed-Receiver problem.

Belief-based versus recommendation-based approach We have formulated the persuasion problem in terms of distributions of posterior beliefs. An alternative approach is to explicitly introduce a Sender and a Receiver, and maximize the Sender's utility from the realized state and the Receiver's action over joint distributions of states and recommendations, subject to Bayes' plausibility and an obedience constraints for the Receiver.

We first note that none of these two approaches is more general—it is in fact possible to reformulate the belief-based problem using the recommendation-based approach, and vice versa. To illustrate this point, suppose that Ω is a finite set. Consider a problem in which the Sender's and Receiver's utility functions are $w(a, \omega)$ and $u(a, \omega)$, respectively, where a is the action of the Receiver. Kamenica and Gentzkow (2011) show that this problem can be analyzed through the belief-based approach by defining $V(\mu) = \mathbb{E}_\mu[w(a^\star(\mu), \omega)]$, where $a^\star(\mu) \in \arg\max_{a\in A} \mathbb{E}_\mu[u(a, \omega)]$. Conversely, the problem we introduced in Section 2 is equivalent to a problem in which the action space is $A = \Delta(\Omega)$, the Sender's utility is given by $w(a, \omega) = V(a)$, and the Receiver's utility is $u(a, \omega) = 2a(\omega) - \sum_{\omega'\in\Omega} a^2(\omega')$. Indeed, given a posterior μ, the Receiver takes an action $a^\star(\mu) = \mu$, which maximizes his expected utility $\sum_{\omega\in\Omega}(2a(\omega)\mu(\omega) - a^2(\omega))$, and thus the objective function is $V(\mu)$.

In the context of moment persuasion, the two approaches are unified by Theorem 5 through the lens of duality—this is because the martingale constraint in the definition of the feasible set $\Pi(\mu_0)$ can be regarded as an obedience constraint for a Receiver with quadratic preferences who matches the action to the state (see Kolotilin (2018)). It is interesting to ask whether duality could similarly cast light on the relationship between the two approaches in more general contexts, such as a multidimensional version of the nonlinear persuasion problem considered by Kolotilin, Corrao, and Wolitzky (2024).

Multiple receivers Perhaps the biggest limitation of our setting is that it does not cover the case in which a Sender wishes to communicate privately with multiple interacting Receivers. Of course, our results do apply when the Sender is restricted to public signals, as in Inostroza and Pavan (2023). Moreover, our duality approach could be useful in analyzing private persuasion problems in conjunction with existing results. Mathevet, Perego, and Taneva (2020) show how to adapt the belief-based approach to persuasion in games, by decomposing a general signal into its public and (purely) private part. Our results apply to the optimal design of the public part of the signal. Additionally, in a recent contribution, Arieli, Babichenko, and Sandomirskiy (2023) apply transportation duality to cast light on the optimal design of the purely private signal—it is natural to ask whether our approach and theirs could be unified. Duality may also be useful within the recommendation-based approach to information design in games (introduced by Bergemann and Morris (2016), and Taneva (2019)). Galperti and Perego (2018) obtain strong duality under finite action and state spaces, while Smolin and Yamashita (2024) rely on weak duality in their analysis of "concave games." Obtaining conditions for strong duality in a general environment remains an open problem.

Appendix: Proofs

We will prove the results in Section 3 in a different order than they appear in Section 3. We first deal with weak duality and primal attainment, as their proofs are standard. We then prove Theorem 4. Finally, relying on Theorem 4, we prove Theorem 2 and Theorem 3.

A.1 *Proof of Theorem 1 and primal attainment*

We first prove Theorem 1. By the definition of the Lebesgue integral, τ belongs to $\mathcal{T}(\mu_0)$ if and only if for any measurable and bounded $p : \Omega \to \mathbb{R}$,

$$\int_{\Delta(\Omega)} \int_{\Omega} p(\omega)\, \mathrm{d}\mu(\omega)\, \mathrm{d}\tau(\mu) = \int_{\Omega} p(\omega)\, \mathrm{d}\mu_0(\omega).$$

Thus, for any $\tau \in \mathcal{T}(\mu_0)$ and any such p that additionally satisfies $V(\mu) \leq \int_{\Omega} p(\omega)\, \mathrm{d}\mu(\omega)$ for all $\mu \in \Delta(\Omega)$, we have

$$\int_{\Delta(\Omega)} V(\mu)\, \mathrm{d}\tau(\mu) \leq \int_{\Delta(\Omega)} \int_{\Omega} p(\omega)\, \mathrm{d}\mu(\omega)\, \mathrm{d}\tau(\mu) = \int_{\Omega} p(\omega)\, \mathrm{d}\mu_0(\omega).$$

Taking the supremum over $\mathcal{T}(\mu_0)$ on the left-hand side and the infimum over $\mathcal{P}(V)$ on the right-hand side (any $p \in \mathcal{P}(V)$ is measurable and bounded) yields the desired result.

Next, we prove primal attainment under the weaker assumption that V is bounded only from above, because this stronger version will be used in the proof of Theorem 8.

Lemma 2. *Let* $V : \Delta(\Omega) \to \mathbb{R} \cup \{-\infty\}$ *be bounded from above and upper semicontinuous. Then (P) has an optimal solution.*

Proof. Because the function $\tau \to \int_{\Delta(\Omega)} \mu\, \mathrm{d}\tau(\mu)$ is continuous, the feasible set $\mathcal{T}(\mu_0)$ is compact, being a closed subset of the compact set $\Delta(\Delta(\Omega))$. Moreover, $\mathcal{T}(\mu_0)$ is nonempty, as it contains the Dirac probability measure δ_{μ_0} at μ_0, which corresponds to no disclosure. Since V is bounded from above and upper semicontinuous, the function $\tau \to \int V(\mu)\, \mathrm{d}\tau(\mu)$ is also upper semicontinuous, and thus attains its maximum on the compact set $\mathcal{T}(\mu_0)$, by the Weierstrass theorem. Thus, an optimal solution $\tau^\star$ to the problem (P) exists. □

A.2 *Proof of Theorem 4*

We start with a key lemma.

Lemma 3. *Let* $\mu_0, \eta_0 \in \Delta(\Omega)$ *and* $\tau \in \mathcal{T}(\mu_0)$. *There exists a measurable function* $\eta : \Delta(\Omega) \to \Delta(\Omega)$ *such that*

$$\int_{\Delta(\Omega)} \eta(\mu)\, \mathrm{d}\tau(\mu) = \eta_0 \quad \textit{and} \quad \int_{\Delta(\Omega)} \left\| \mu - \eta(\mu) \right\|_{\mathrm{KR}}\, \mathrm{d}\tau(\mu) = \|\mu_0 - \eta_0\|_{\mathrm{KR}}.$$

Before proving Lemma 3, we show that it implies Theorem 4. First, since V is Lipschitz, it is upper semicontinuous, and hence, by Lemma 2, for any $\mu_0 \in \Delta(\Omega)$, there exists $\tau \in \mathcal{T}(\mu_0)$ that attains the concave closure of V at μ_0, so that $\widehat{V}(\mu_0) = \int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\tau(\mu)$. Next, since V is Lipschitz, there exists $L \in \mathbb{R}$, such that, for all $\mu_0, \eta_0 \in \Delta(\Omega)$, we have $|V(\mu_0) - V(\eta_0)| \leq L\|\mu_0 - \eta_0\|_{\mathrm{KR}}$. Then, using Lemma 3 to define the function η, we obtain

$$\begin{aligned}\widehat{V}(\mu_0) - \widehat{V}(\eta_0) &\leq \int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\tau(\mu) - \int_{\Delta(\Omega)} V\big(\eta(\mu)\big)\,\mathrm{d}\tau(\mu) \\ &\leq \int_{\Delta(\Omega)} L\big\|\mu - \eta(\mu)\big\|_{\mathrm{KR}}\,\mathrm{d}\tau(\mu) = L\|\mu_0 - \eta_0\|_{\mathrm{KR}}.\end{aligned}$$

By reversing the roles of μ_0 and η_0, we conclude that $\widehat{V}$ is Lipschitz (with constant L). Thus, it remains to prove Lemma 3.

Proof of Lemma 3. The idea behind the proof is to "perturb" each posterior belief $\mu \in \Delta(\Omega)$ (with η describing the perturbation function) in such a way that perturbed posteriors $\eta(\mu)$ average out to η_0, and the average distance between each posterior μ and its perturbation $\eta(\mu)$ is the same as the distance between the "priors" μ_0 and η_0. A naive approach would be to perturb each posterior μ by the same magnitude and in the same direction $\eta_0 - \mu_0$. However, this could easily take the perturbed beliefs outside the set $\Delta(\Omega)$. Thus, our construction is more complicated: We rely on a property of the Kantorovich–Rubinstein norm to find the transportation plan $\lambda \in \Delta(\Omega \times \Omega)$ that defines the distance between μ_0 and η_0; we then define the perturbation by conditioning on each realized posterior belief $\mu \in \Delta(\Omega)$, and using a properly constructed conditional transportation plan.

Since every norm is convex, we have $\|\mu_0 - \eta_0\|_{\mathrm{KR}} \leq \int_{\Delta(\Omega)} \|\mu - \eta(\mu)\|_{\mathrm{KR}}\,\mathrm{d}\tau(\mu)$ for any measurable function $\eta : \Delta(\Omega) \to \Delta(\Omega)$ such that $\int_{\Delta(\Omega)} \eta(\mu)\,\mathrm{d}\tau(\mu) = \eta_0$. Thus, it suffices to show the existence of a measurable function $\eta : \Delta(\Omega) \to \Delta(\Omega)$ satisfying the reverse inequality.

By the Kantorovich–Rubinstein theorem (Theorem 8.10.45 in Bogachev (2007)),

$$\|\mu_0 - \eta_0\|_{\mathrm{KR}} = \min_{\lambda \in \Lambda(\mu_0, \eta_0)} \int_{\Omega \times \Omega} \rho\big(\omega, \omega'\big)\,\mathrm{d}\lambda\big(\omega, \omega'\big),$$

where $\Lambda(\mu_0, \eta_0)$ is the set of probability measures $\lambda \in \Delta(\Omega \times \Omega)$ such that $\lambda(A \times \Omega) = \mu_0(A)$ and $\lambda(\Omega \times B) = \eta_0(B)$ for all measurable sets $A, B \subset \Omega$. In particular, the minimum is attained at some $\lambda \in \Lambda(\mu_0, \eta_0)$, which we fix for the remainder of the proof.

Define a probability measure $\sigma \in \Delta(\Delta(\Omega) \times \Omega)$ by $\sigma(M, A) = \int_M \mu(A)\,\mathrm{d}\tau(\mu)$ for all measurable $M \subset \Delta(\Omega)$ and $A \subset \Omega$. For all measurable $A \subset \Omega$, we have $\sigma(\Delta(\Omega), A) = \mu_0(A)$ because $\int_{\Delta(\Omega)} \mu\,\mathrm{d}\tau(\mu) = \mu_0$. For any probability measure on a product of two compact metric spaces, we can define its conditional measures (Theorem 10.4.5 in Bogachev (2007)). Since Ω and $\Delta(\Omega)$ are compact, there exists a measurable function $\omega \to \sigma(\cdot|\omega)$, from Ω into $\Delta(\Delta(\Omega))$, such that $\sigma(M, A) = \int_A \sigma(M|\omega)\,\mathrm{d}\mu_0(\omega)$ for all measurable $A \subset \Omega$ and $M \subset \Delta(\Omega)$. Similarly, there exists a measurable function $\omega \to \lambda(\cdot|\omega)$, from Ω into $\Delta(\Omega)$, such that $\lambda(A, B) = \int_A \lambda(B|\omega)\,\mathrm{d}\mu_0(\omega)$ for all measurable $A, B \subset \Omega$.

Define a probability measure $\zeta \in \Delta(\Omega \times \Omega \times \Delta(\Omega))$ by $\zeta(A, B, M) = \int_A \lambda(B|\omega)\sigma(M|\omega)\, d\mu_0(\omega)$ for all measurable $A, B \subset \Omega$, and $M \subset \Delta(\Omega)$. For all measurable $A, B \subset \Omega$, and $M \subset \Delta(\Omega)$, we have $\zeta(A, B, \Delta(\Omega)) = \lambda(A, B)$ and $\zeta(A, \Omega, M) = \sigma(M, A)$, by construction. Since $\Omega \times \Omega$ and $\Delta(\Omega)$ are compact, there exists a measurable function $\mu \to \zeta(\cdot|\mu)$, from $\Delta(\Omega)$ into $\Delta(\Omega \times \Omega)$, such that $\zeta(A, B, M) = \int_M \zeta(A, B|\mu)\, d\tau(\mu)$ for all measurable $A, B \subset \Omega$, and $M \subset \Delta(\Omega)$.

Finally, define a measurable function $\mu \to \eta(\mu)$, from $\Delta(\Omega)$ into $\Delta(\Omega)$, by $\eta(\mu)(B) = \zeta(\Omega, B|\mu)$ for all $\mu \in \Delta(\Omega)$ and all measurable $B \subset \Omega$. Notice that the conditional measure $\zeta(\cdot, \cdot|\mu)$ on $\Omega \times \Omega$ is a feasible transportation plan between μ and $\eta(\mu)$, for τ-almost all $\mu \in \Delta(\Omega)$. Indeed, $\zeta(\Omega, \cdot|\mu) = \eta(\mu)(\cdot)$ by construction, and for any measurable $A \subset \Omega$ and $M \subset \Delta(\Omega)$,

$$\int_M \zeta(A, \Omega|\mu)\, d\tau(\mu) = \zeta(A, \Omega, M) = \sigma(M, A) = \int_M \mu(A)\, d\tau(\mu),$$

establishing that $\zeta(\cdot, \Omega|\mu) = \mu(\cdot)$ for τ-almost all $\mu \in \Delta(\Omega)$.

To show that the constructed function η satisfies the required properties, note first that, for any measurable $B \subset \Omega$,

$$\int_{\Delta(\Omega)} \eta(\mu)(B)\, d\tau(\mu) = \int_{\Delta(\Omega)} \zeta(\Omega, B|\mu)\, d\tau(\mu) = \zeta\big(\Omega, B, \Delta(\Omega)\big) = \lambda(\Omega, B) = \eta_0(B),$$

and hence $\int_{\Delta(\Omega)} \eta(\mu)\, d\tau(\mu) = \eta_0$. Moreover,

$$\begin{aligned}\int_{\Delta(\Omega)} \|\mu - \eta(\mu)\|_{\mathrm{KR}}\, d\tau(\mu) &\le \int_{\Delta(\Omega)} \left[\int_{\Omega\times\Omega} \rho(\omega, \omega')\, d\zeta(\omega, \omega'|\mu)\right] d\tau(\mu) \\ &= \int_{\Omega\times\Omega} \rho(\omega, \omega')\, d\lambda(\omega, \omega') = \|\mu_0 - \eta_0\|_{\mathrm{KR}},\end{aligned}$$

where the inequality follows from the Kantorovich–Rubinstein theorem and the fact that the conditional measure $\zeta(\cdot, \cdot|\mu)$ on $\Omega \times \Omega$ is a feasible transportation plan between μ and $\eta(\mu)$, for τ-almost all $\mu \in \Delta(\Omega)$, as shown above. □

A.3 *Proof of Theorem 2*

Existence of an optimal solution to the primal problem follows from Lemma 2.

To prove the rest of the theorem, we introduce some basic tools from convex analysis used in the proof of the next lemma.[31] Let E be a normed vector space and $E^\star$ its topological dual space, that is, the space of all continuous linear functions on E. Let $\varphi : E \to \mathbb{R} \cup \{+\infty\}$ be an extended-valued function that is not identically $\{+\infty\}$. The Legendre transform of φ is the function $\varphi^\star : E^\star \to \mathbb{R} \cup \{+\infty\}$ given by

$$\varphi^\star(z^\star) = \sup_{z\in E}\big\{\langle z^\star, z\rangle - \varphi(z)\big\} \quad \text{for all } z^\star \in E^\star,$$

[31] See Chapter 1.4 in Brezis (2011) for further details.

where $\langle \cdot, \cdot \rangle$ is the duality product between E and $E^\star$. It is easy to verify that $\varphi^\star$ is convex, lower semicontinuous, and not identically $\{+\infty\}$. Next, define the function $\varphi^{\star\star} : E \to \mathbb{R} \cup \{+\infty\}$ as the Legendre transform of $\varphi^\star$, restricted from $E^{\star\star}$ to E,

$$\varphi^{\star\star}(z) = \sup_{z^\star \in E^\star} \{ \langle z^\star, z \rangle - \varphi^\star(z^\star) \} \quad \text{for all } z \in E.$$

Clearly, $\varphi^{\star\star}$ is a convex and lower semicontinuous function satisfying $\varphi^{\star\star}(z) \leq \varphi(z)$ for all $z \in E$. The Fenchel–Moreau theorem states that if $\varphi : E \to \mathbb{R} \cup \{+\infty\}$ is convex and lower semicontinuous, and not identically $\{+\infty\}$, then $\varphi^{\star\star} = \varphi$. We remark that the Fenchel–Moreau theorem is a consequence of an appropriate hyperplane separation theorem.[32]

We prove the theorem in two steps. First, we show the conclusion for Lipschitz objective functions. Here, we rely on the (already proven) Theorem 4. Second, we use an approximation argument to extend the conclusion to all bounded and upper semicontinuous objectives.

Lemma 4. *Let $V \in \mathrm{Lip}(\Delta(\Omega))$. Then (O) holds.*

Proof. Let $E = (M(\Omega), \|\cdot\|_{KR})$; then, as argued in the main text, $E^\star = \mathrm{Lip}(\Omega)$. Define the function φ on $M(\Omega)$ as

$$\varphi(\eta) = \begin{cases} -\sup\limits_{\tau \in \mathcal{T}(\eta)} \displaystyle\int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\tau(\mu), & \eta \in \Delta(\Omega), \\ +\infty, & \eta \notin \Delta(\Omega). \end{cases}$$

First, we note that φ is convex. Indeed, let $\eta_1, \eta_2 \in M(\Omega)$ and $\lambda \in (0,1)$. If $\eta_1, \eta_2 \in \Delta(\Omega)$, then, by Lemma 2, there exist $\tau_1 \in \mathcal{T}(\eta_1)$ and $\tau_2 \in \mathcal{T}(\eta_2)$ such that

$$\varphi(\eta_1) = -\int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\tau_1(\mu) \in \mathbb{R} \quad \text{and} \quad \varphi(\eta_2) = -\int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\tau_2(\mu) \in \mathbb{R}.$$

By the definition of $\mathcal{T}$,

$$\lambda\tau_1 + (1-\lambda)\tau_2 \in \mathcal{T}(\lambda\eta_1 + (1-\lambda)\eta_2)$$

and hence, by the definition of φ,

$$\varphi(\lambda\eta_1 + (1-\lambda)\eta_2) \leq -\int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}(\lambda\tau_1 + (1-\lambda)\tau_2) = \lambda\varphi(\eta_1) + (1-\lambda)\varphi(\eta_2).$$

If $\eta_1 \notin \Delta(\Omega)$ or $\eta_2 \notin \Delta(\Omega)$, then, trivially,

$$\varphi(\lambda\eta_1 + (1-\lambda)\eta_2) \leq \lambda\varphi(\eta_1) + (1-\lambda)\varphi(\eta_2) = +\infty.$$

Second, we note that $\varphi : M(\Omega) \to \mathbb{R} \cup \{+\infty\}$ is lower semicontinuous, because φ is Lipschitz on the compact set $\Delta(\Omega)$, by Theorem 4.

[32] Indeed, an earlier version of this paper Dworczak and Kolotilin (2019) contained a proof of strong duality that directly relied on a hyperplane separation theorem.

Let us compute the Legendre transform of φ. For each $g \in \mathrm{Lip}(\Omega)$,

$$\begin{aligned}
\varphi^\star(g) &= \sup_{\eta \in M(\Omega)} \left\{ \int_\Omega g(\omega)\,\mathrm{d}\eta(\omega) - \varphi(\eta) \right\} \\
&= \sup_{\eta \in \Delta(\Omega), \tau \in \mathcal{T}(\eta)} \left\{ \int_\Omega g(\omega)\,\mathrm{d}\eta(\omega) + \int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\tau(\mu) \right\} \\
&= \sup_{\eta \in \Delta(\Omega), \tau \in \mathcal{T}(\eta)} \left\{ \int_{\Delta(\Omega)} \left(\int_\Omega g(\omega)\,\mathrm{d}\mu(\omega) + V(\mu) \right) \mathrm{d}\tau(\mu) \right\} \\
&= \sup_{\eta \in \Delta(\Omega)} \left\{ \int_\Omega g(\omega)\,\mathrm{d}\eta(\omega) + V(\eta) \right\},
\end{aligned}$$

where the last equality follows from the fact that by treating $\widetilde{V}(\mu) := \int_\Omega g(\omega)\,\mathrm{d}\mu(\omega) + V(\mu)$ as an objective function, we obtain a persuasion problem in which we choose both a prior η and a distribution τ of posteriors, which averages out to the prior, so it is optimal to choose a prior $\eta \in \arg\max_{\mu \in \Delta(\Omega)} \widetilde{V}(\mu)$ and a degenerate distribution $\tau = \delta_\eta$.[33]

Let us finally compute $\varphi^{\star\star}(\mu_0)$,

$$\begin{aligned}
\varphi^{\star\star}(\mu_0) &= \sup_{p \in \mathrm{Lip}(\Omega)} \left\{ \int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega) - \varphi^\star(p) \right\} \\
&= \sup_{p \in \mathrm{Lip}(\Omega)} \left\{ \int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega) - \sup_{\eta \in \Delta(\Omega)} \left\{ \int_\Omega p(\omega)\,\mathrm{d}\eta(\omega) + V(\eta) \right\} \right\} \\
&= - \inf_{p \in \mathrm{Lip}(\Omega)} \left\{ \int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega) + \sup_{\eta \in \Delta(\Omega)} \left\{ V(\eta) - \int_\Omega p(\omega)\,\mathrm{d}\eta(\omega) \right\} \right\} \\
&= - \inf_{p \in \mathrm{Lip}(\Omega)} \left\{ \int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega) : \sup_{\eta \in \Delta(\Omega)} \left\{ V(\eta) - \int_\Omega p(\omega)\,\mathrm{d}\eta(\omega) \right\} = 0 \right\} \\
&= - \inf_{p \in \mathcal{P}(V)} \left\{ \int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega) \right\},
\end{aligned}$$

where the third equality follows from substituting p for $-p$ as the optimization variable, and the fourth follows, because for any fixed η, adding a constant to p does not change the value of the outer infimum—it is thus without loss of generality to normalize p by insisting that the inner supremum is equal to 0 (note that the inner supremum is attained and finite at each $p \in \mathrm{Lip}(\Omega)$). The Fenchel–Moreau theorem implies that $\varphi = \varphi^{\star\star}$, so (O) follows from $\varphi(\mu_0) = \varphi^{\star\star}(\mu_0)$. □

Lemma 5. *Let V be bounded and upper semicontinuous. Then (O) holds.*

Proof. This follows from a standard approximation argument as, for example, in the proof of Theorem 1.3 in Villani (2003). By Baire's theorem (see, e.g., Box 1.5 in Santambrogio (2015)), there exists a nonincreasing sequence of Lipschitz functions $V_k \in$

[33] This observation is also made in the proof of Theorem 2 in Dworczak (2020).

$\mathrm{Lip}(\Delta(\Omega))$ converging pointwise to V. That is, $V_k(\mu) \geq V_{k+1}(\mu)$ for all $\mu \in \Delta(\Omega)$ and $k \in \mathbb{N}$, and $\lim_{k\to\infty} V_k(\mu) = V(\mu)$ for all $\mu \in \Delta(\Omega)$. Let $\tau_k^\star$ denote an optimal solution to (P) with the objective function V_k. For each $k \in \mathbb{N}$, we have

$$\int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\tau^\star(\mu) \leq \inf_{p\in\mathcal{P}(V)} \int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega) \leq \inf_{p\in\mathcal{P}(V_k)} \int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega)$$
$$= \int_{\Delta(\Omega)} V_k(\mu)\,\mathrm{d}\tau_k^\star(\mu),$$

where the first inequality holds by Theorem 1, the second inequality holds by $\mathcal{P}(V_k) \subset \mathcal{P}(V)$ for $V_k \geq V$, and the equality holds by Lemma 4 for Lipschitz V_k. To establish (O) for upper semicontinuous V, it is thus sufficient to show that

$$\lim_{k\to\infty} \int_{\Delta(\Omega)} V_k(\mu)\,\mathrm{d}\tau_k^\star(\mu) \leq \int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\tau^\star(\mu).$$

Thanks to compactness of $\mathcal{T}(\mu_0)$, up to extraction of a subsequence, we can suppose that $\tau_k^\star$ converges weakly to some $\overline{\tau} \in \mathcal{T}(\mu_0)$. Then, for each $j \in \mathbb{N}$, we have

$$\lim_{k\to\infty} \int_{\Delta(\Omega)} V_k(\mu)\,\mathrm{d}\tau_k^\star(\mu) \leq \lim_{k\to\infty} \int_{\Delta(\Omega)} V_j(\mu)\,\mathrm{d}\tau_k^\star(\mu) = \int_{\Delta(\Omega)} V_j(\mu)\,\mathrm{d}\overline{\tau}(\mu),$$

where the first inequality holds because $V_k \leq V_j$ for $k \geq j$, and the equality holds because V_j is (Lipschitz) continuous and $\tau_k^\star \to \overline{\tau}$. Then, letting j go to infinity and invoking the monotone convergence theorem,

$$\lim_{j\to\infty} \int_{\Delta(\Omega)} V_j(\mu)\,\mathrm{d}\overline{\tau}(\mu) = \int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\overline{\tau}(\mu),$$

we obtain

$$\lim_{k\to\infty} \int_{\Delta(\Omega)} V_k(\mu)\,\mathrm{d}\tau_k^\star(\mu) \leq \int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\overline{\tau}(\mu) \leq \int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\tau^\star(\mu),$$

where the last inequality holds because $\tau^\star$ is an optimal solution to (P). This establishes (O) for upper semicontinuous V. As a by-product, it also shows the optimality of $\overline{\tau}$.[34] □

A.4 *Proof of Corollary 1*

By Theorem 2, $\tau \in \mathcal{T}(\mu_0)$ and $p \in \mathcal{P}(V)$ are optimal solutions to (P) and (D) if and only if

$$\int_{\Delta(\Omega)} V(\mu)\,\mathrm{d}\tau(\mu) = \int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega) \quad \Longleftrightarrow \quad \int_{\Delta(\Omega)} \left(V(\mu) - \int_\Omega p(\omega)\,\mathrm{d}\mu(\omega) \right) \mathrm{d}\tau(\mu) = 0.$$

Since the term in parenthesis is nonpositive for $p \in \mathcal{P}(V)$, it follows that $\tau(\Lambda) = 1$ where

$$\Lambda = \left\{ \mu \in \Delta(\Omega) : V(\mu) = \int_\Omega p(\omega)\,\mathrm{d}\mu(\omega) \right\} = \left\{ \mu \in \Delta(\Omega) : V(\mu) \geq \int_\Omega p(\omega)\,\mathrm{d}\mu(\omega) \right\}.$$

[34] In the persuasion literature, a similar argument appears in the proof of Theorem 1 in Dizdar and Kováč (2020) for the special case of one-dimensional moment persuasion.

The set Λ is closed because $V(\mu)$ is upper semicontinuous in μ and $\int_\Omega p(\omega)\,\mathrm{d}\mu(\omega)$ is continuous in μ, given that each $p \in \mathcal{P}(V)$ is Lipschitz continuous. Thus, $\operatorname{supp}(\tau) \subset \Lambda$ and (C) follows, since $\operatorname{supp}(\tau)$ is defined as the smallest closed set on which τ is concentrated.

A.5 *Proof of Theorem 3*

The duality theorem in Gale (1967) shows that $\widehat{V}$ is superdifferentiable at μ_0 if and only if $\widehat{V}$ has bounded steepness at μ_0. Thus, Theorem 3 follows from the following lemma.

Lemma 6. *There exists an optimal solution $p \in \mathcal{P}(V)$ to (D) if and only if $\widehat{V}$ is superdifferentiable at μ_0.*

Proof. If $\widehat{V}$ is superdifferentiable at μ_0, then, by the fact that $(M(\Omega), \|\cdot\|_{\mathrm{KR}})^\star = \operatorname{Lip}(\Omega)$, there exists $p \in \operatorname{Lip}(\Omega)$ such that

$$\widehat{V}(\mu_0) = \int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega) \quad \text{and} \quad \widehat{V}(\mu) \le \int_\Omega p(\omega)\,\mathrm{d}\mu(\omega), \quad \text{for all } \mu \in \Delta(\Omega).$$

Thus,

$$V(\mu) \le \widehat{V}(\mu) \le \int_\Omega p(\omega)\,\mathrm{d}\mu(\omega), \quad \text{for all } \mu \in \Delta(\Omega),$$

so $p \in \mathcal{P}(V)$ is an optimal solution to (D), by Theorem 1.

Conversely, if $p \in \mathcal{P}(V)$ is optimal, then we have $p \in \operatorname{Lip}(\Omega)$,

$$\overline{V}(\mu_0) = \int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega), \quad \text{and} \quad V(\mu) \le \int_\Omega p(\omega)\,\mathrm{d}\mu(\omega), \quad \text{for all } \mu \in \Delta(\Omega).$$

By the definition of the concave envelope,

$$\overline{V}(\mu) \le \int_\Omega p(\omega)\,\mathrm{d}\mu(\omega), \quad \text{for all } \mu \in \Delta(\Omega).$$

Therefore, by Theorem 2,

$$\widehat{V}(\mu_0) = \int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega), \quad \text{and} \quad \widehat{V}(\mu) \le \int_\Omega p(\omega)\,\mathrm{d}\mu(\omega), \quad \text{for all } \mu \in \Delta(\Omega).$$

Thus, p is a supergradient of $\widehat{V}$ at μ_0, and thus $\widehat{V}$ is superdifferentiable at μ_0 (simply define $H(\mu) = \int_\Omega p(\omega)\,\mathrm{d}\mu(\omega)$, which is a continuous linear function on $M(\Omega)$ because $p \in \operatorname{Lip}(\Omega)$). □

A.6 *Proof of Lemma 1*

Suppose that v is L-Lipschitz on $X \subset \mathbb{R}^N$. Since all norms are equivalent in an N-dimensional Euclidean space, without loss of generality, we endow $\mathbb{R}^N$ with the Euclidean norm,

$$\|x\| = \sqrt{\sum_{i=1}^N x_i^2}, \quad \text{for all } x \in \mathbb{R}^N.$$

For any μ, $\eta \in \Delta(\Omega)$, with $\mu \neq \eta$,

$$\frac{|V(\mu) - V(\eta)|}{\|\mu - \eta\|_{\mathrm{KR}}} = \frac{|v(\mathbb{E}_\mu[\omega]) - v(\mathbb{E}_\eta[\omega])|}{\|\mathbb{E}_\mu[\omega] - \mathbb{E}_\eta[\omega]\|} \frac{\|\mathbb{E}_\mu[\omega] - \mathbb{E}_\eta[\omega]\|}{\|\mu - \eta\|_{\mathrm{KR}}} \leq L \frac{\|\mathbb{E}_\mu[\omega] - \mathbb{E}_\eta[\omega]\|}{\|\mu - \eta\|_{\mathrm{KR}}}.$$

Because the function $f(\omega) = \omega_i$ is 1-Lipschitz,

$$|\mathbb{E}_\mu[\omega_i] - \mathbb{E}_\eta[\omega_i]| = \left| \int_\Omega \omega_i \,\mathrm{d}(\mu - \eta)(\omega) \right| \leq \|\mu - \eta\|_{\mathrm{KR}},$$

and thus

$$\|\mathbb{E}_\mu[\omega] - \mathbb{E}_\eta[\omega]\| = \sqrt{\sum_{i=1}^{N} (\mathbb{E}_\mu[\omega_i] - \mathbb{E}_\eta[\omega_i])^2} \leq \sqrt{N} \|\mu - \eta\|_{\mathrm{KR}},$$

showing that V is $L\sqrt{N}$-Lipschitz.

A.7 *Proof of Theorem 5*

By Lemma 1, we know that $V : \Delta(\Omega) \to \mathbb{R}$ is Lipschitz, since v is Lipschitz. It follows from Theorems 2, 3, and 4 that there exists a solution $p \in \mathrm{Lip}(\Omega)$ to the dual problem (D); moreover, since (P_M) is a special case of the general problem (P), $\pi \in \Pi(\mu_0)$ is then optimal for (P_M) if and only if

$$\int_X v(x) \,\mathrm{d}\pi_X(x) = \int_\Omega p(\omega) \,\mathrm{d}\mu_0(\omega).$$

Let $\check{p}$ be the convex roof extension of p from Ω to X, defined in the main text. By construction, $\check{p} \leq p$ on Ω. Moreover, the infimum in the definition of $\check{p}$ is attained because p is (Lipschitz) continuous on Ω and the set of feasible distributions is compact. Hence, for any $x \in X$, we can write $\check{p}(x) = \int_\Omega p(\omega) \,\mathrm{d}\mu_x(\omega)$ for some $\mu_x \in \Delta(\Omega)$ with $\int_\Omega \omega \,\mathrm{d}\mu_x(\omega) = x$. By the definition of $\check{p}$, for any $x, y \in X$, $\lambda \in (0, 1)$, we have

$$\lambda \check{p}(x) + (1 - \lambda)\check{p}(y) = \int_\Omega p(\omega) \,\mathrm{d}(\lambda\mu_x + (1 - \lambda)\mu_y)(\omega) \geq \check{p}(\lambda x + (1 - \lambda)y),$$

showing that $\check{p}$ is convex. Moreover, by feasibility of p, for any $x \in X$,

$$\check{p}(x) = \int_\Omega p(\omega) \,\mathrm{d}\mu_x(\omega) \geq V(\mu_x) = v(x).$$

Next, we prove a key lemma.

Lemma 7. *Let v be L-Lipschitz and $\check{p} \geq v$. There exists a measurable function $q : X \to \mathbb{R}^N$ such that $\|q(x)\| \leq L$ for all $x \in X$, and*

$$\check{p}(y) \geq v(x) + q(x) \cdot (y - x), \quad \textit{for all } y, x \in X.$$

Proof. Define

$$F(x) := \{r \in \mathbb{R}^N : \check{p}(y) \geq v(x) + r \cdot (y - x), \text{ for all } y \in X\},$$

and let

$$q(x) := \arg\min_{r \in F(x)} \|r\|, \quad \text{for all } x \in X.$$

Note that $F(x)$ is closed-valued and convex-valued. Thus, if $F(x)$ is nonempty, then $q(x)$ exists and is unique because $q(x)$ is the projection of 0 onto the nonempty closed convex set $F(x)$. If we can additionally prove that $\|q(x)\| \leq L$ for all $x \in X$, then q will be measurable by the measurable maximum theorem (Theorem 18.19 in Aliprantis and Border (2006)). To see that, note that the definition of q will not change if we additionally require that $\|r\| \leq L$, so that the correspondence $x \rightrightarrows F(x) \cap \{r \in \mathbb{R}^N : \|r\| \leq L\}$ is compact-valued and upper hemicontinuous (given that $\check{p}$ is lower semicontinuous and v is continuous), and thus measurable, by Theorem 18.20 in Aliprantis and Border (2006).

We deal with some easy cases first. If $0 \in F(x)$, then $q(x) = 0$ and $0 = \|q(x)\| \leq L$. Next, if $0 \notin F(x)$ but $\check{p}(x) = v(x)$, then we have, for any $y \in X$,

$$\check{p}(y) - \check{p}(x) \geq v(y) - v(x) \geq -L\|y - x\|,$$

because $\check{p} \geq v$ and v is L-Lipschitz. By the duality theorem in Gale (1967), $q(x)$ is well-defined and

$$\|q(x)\| = -\inf_{y \in X} \frac{\check{p}(y) - \check{p}(x)}{\|y - x\|} \leq L.$$

Thus, for the rest of the proof, we fix an arbitrary $x \in X$ such that $0 \notin F(x)$ and $\check{p}(x) > v(x)$.

We first show that $F(x)$ is nonempty. Because $\check{p}(x) > v(x)$, the point $(x, v(x))$ does not belong to the epigraph of $\check{p}$, defined as $\text{epi}(\check{p}) := \{(y, t) \in X \times \mathbb{R} : t \geq \check{p}(y)\}$. Note that $\text{epi}(\check{p})$ is closed and convex, because $\check{p}$ is lower semicontinuous (see footnote 19) and convex. By the separation theorem (e.g., Corollary 11.4.1 in Rockafellar (1970)), there exists $(\alpha, \beta) \in \mathbb{R}^N \times \mathbb{R}$ such that, for all $y \in X$ and $t \geq \check{p}(y)$,

$$\alpha \cdot y + \beta t > \alpha \cdot x + \beta v(x).$$

Clearly, $\beta \geq 0$; otherwise, the inequality would be violated for sufficiently large t. Moreover, $\beta \neq 0$; otherwise, the inequality would be violated for $(y, t) = (x, \check{p}(x))$. Thus, evaluating the inequality for $t = \check{p}(y)$, for all $y \in X$, proves that $-\alpha/\beta$ belongs to $F(x)$. Thus, $F(x)$ is indeed nonempty (and hence $q(x)$ is well-defined).

We now show that $\|q(x)\| \leq L$. Define the set

$$Y := \{y \in X : \check{p}(y) = v(x) + q(x) \cdot (y - x)\}.$$

Note that Y is nonempty: If there is no $y \in X$ such that $\check{p}(y) = v(x) + q(x) \cdot (y - x)$, then the constraint in the definition of $F(x)$ is slack, so it is possible to reduce $\|r\|$, contradicting that $q(x)$ is a minimizer (this step uses the fact that $\check{p}$ is lower semicontinuous).

Since $\check{p}$ is convex, the set Y is convex. Since $\check{p}(x) > v(x)$, the set Y cannot contain x. Also, let

$$E := \{e \in \mathbb{R}^N : e \cdot q(x) < 0\}.$$

We will prove that there exists $y^\star \in Y$ such that $e \cdot (y^\star - x) \geq 0$ for all $e \in E$. Suppose that such $y^\star$ does not exist. Since any such $y^\star$ must satisfy $y^\star - x = -tq(x)$ for some $t \geq 0$, we conclude that the compact convex set $Y - x := \{y - x : y \in Y\}$ and the closed convex cone $\{-tq(x) : t \geq 0\}$ must be disjoint. By the separation theorem (e.g., Corollary 11.4.1 in Rockafellar (1970)), there exists $e \in \mathbb{R}^N$ such that

$$\max_{y \in Y} e \cdot (y - x) < \inf_{t \geq 0} e \cdot \big(-tq(x)\big).$$

Notice that we must have $e \cdot q(x) \leq 0$, as otherwise the right-hand side is $-\infty$ and the inequality cannot hold. In fact, there exists $e \in \mathbb{R}^N$ such that $e \cdot q(x) < 0$, because we can always replace e with $e - \varepsilon q(x)$ for a sufficiently small $\varepsilon > 0$ without violating the above inequality, given that Y is compact. Since there is $e \in E$ such that $e \cdot (y - x) < 0$ for all $y \in Y$, there is $\delta > 0$ such that for all z in the δ-neighborhood of Y, we have $e \cdot (z - x) < 0$, and thus for all $\varepsilon > 0$,

$$v(x) + \big(q(x) + \varepsilon e\big) \cdot (z - x) < v(x) + q(x) \cdot (z - x).$$

Since $\check{p}(z) > v(x) + q(x) \cdot (z - x)$ for $z \notin Y$, and $\check{p}$ is convex and lower semicontinuous, there exists $\gamma > 0$ such that for all $z \in X$ outside the δ-neighborhood of Y, we have

$$\check{p}(z) > v(x) + q(x) \cdot (z - x) + \gamma.$$

Consequently, there exists a sufficiently small $\varepsilon > 0$ such that, for all $z \in X$,

$$\check{p}(z) > v(x) + \big(q(x) + \varepsilon e\big) \cdot (z - x).$$

This is a contradiction with the definition of $q(x)$. Indeed, the above inequality shows that $q(x) + \varepsilon e \in F(x)$ and, by the fact that $e \in E$ and $q(x) \neq 0$, we have $\|q(x) + \varepsilon e\| < \|q(x)\|$ for sufficiently small $\varepsilon > 0$.

We have thus proven that there exists $y^\star \in Y$ such that $e \cdot (y^\star - x) \geq 0$ for all $e \in E$. Since $e \cdot (y^\star - x) \geq 0$ for all $e \in E$ and Y does not contain x, it follows that there exists $t > 0$ such that $x - y^\star = tq$. Thus,

$$q(x) \cdot \big(x - y^\star\big) = \big\|q(x)\big\| \big\|x - y^\star\big\|.$$

And since $y^\star \in Y$, we have that

$$v(x) - \check{p}\big(y^\star\big) = q(x) \cdot \big(x - y^\star\big).$$

Putting these two equalities together, we conclude that

$$\big\|q(x)\big\| \big\|x - y^\star\big\| = v(x) - \check{p}\big(y^\star\big) \leq v(x) - v\big(y^\star\big) \leq L\big\|x - y^\star\big\|,$$

showing that $\|q(x)\| \leq L$. □

Fixing $q(x)$ from Lemma 7, we define

$$\bar{p}(y) := \sup_{x \in X} \{ v(x) + q(x) \cdot (y - x) \}, \quad \text{for all } y \in X.$$

Note that $\bar{p}$ is convex as a pointwise supremum of affine functions. It lies everywhere above v, by definition. Finally, we show that $\bar{p}$ is L-Lipschitz. Take any $y, z \in X$. Let x_n be a sequence of points in X that generate the supremum in the definition of $\bar{p}(y)$. Because X is compact and q is bounded, we can assume that x_n and $q(x_n)$ converge. Then we have that

$$\begin{aligned}
\bar{p}(y) - \bar{p}(z) &= \lim_{n \to \infty} \{ v(x_n) + q(x_n) \cdot (y - x_n) \} - \bar{p}(z) \\
&\leq \lim_{n \to \infty} \{ v(x_n) + q(x_n) \cdot (y - x_n) - v(x_n) - q(x_n) \cdot (z - x_n) \} \\
&= \lim_{n \to \infty} \{ q(x_n) \} \cdot (y - z) \leq L \| y - z \|.
\end{aligned}$$

Because y and z were arbitrary, this proves that $\bar{p}$ is L-Lipschitz.

Finally, notice that $\bar{p} \leq \check{p}$, by Lemma 7. Therefore, on Ω, we have that

$$\bar{p} \leq \check{p} \leq p.$$

Since $\bar{p}$ is Lipschitz, $\bar{p} \geq v$ and $\bar{p}$ is convex, it follows that $\bar{p}$ (restricted to Ω) is feasible for the dual (D); indeed, for any $\mu \in \Delta(\Omega)$,

$$\int_{\Omega} \bar{p}(\omega) \, \mathrm{d}\mu(\omega) \geq \bar{p}\left(\int_{\Omega} \omega \, \mathrm{d}\mu(\omega) \right) \geq v\left(\int_{\Omega} \omega \, \mathrm{d}\mu(\omega) \right) = V(\mu).$$

But since p solves the dual problem (D), we must have that $p = \bar{p}$ almost surely on Ω. Since both these function are (Lipschitz) continuous, we can conclude that p and $\bar{p}$ coincide on Ω. In particular, we have shown that $\bar{p}$ is convex and solves (D) when restricted to Ω.

Next, we prove that if $\pi \in \Pi(\mu_0)$ is optimal for ($\mathrm{P_M}$), then conditions *(i)* and *(ii)* hold. We have already shown that $\bar{p}$ is convex, Lipschitz, and satisfies $\bar{p} \geq v$. To complete the proof that condition *(i)* holds, note that

$$\int_X v(x) \, \mathrm{d}\pi_X(x) = \int_{\Omega} p(\omega) \, \mathrm{d}\mu_0(\omega) = \int_{\Omega} \bar{p}(\omega) \, \mathrm{d}\mu_0(\omega),$$

where the first equality is due to the absence of a duality gap (Theorem 2) and the second is by the fact that $p = \bar{p}$ on Ω. We can also prove that condition *(ii)* holds: $\bar{p}$ satisfies the required equality by definition when q is defined by Lemma 7; moreover,

$$\begin{aligned}
\int_{X \times \Omega} \big(v(x) + q(x) \cdot (\omega - x) \big) \, \mathrm{d}\pi(x, \omega) &= \int_X v(x) \, \mathrm{d}\pi_X(x) = \int_{\Omega} \bar{p}(\omega) \, \mathrm{d}\mu_0(\omega) \\
&= \int_{X \times \Omega} \bar{p}(\omega) \, \mathrm{d}\pi(x, \omega),
\end{aligned}$$

where the first and last equality follow from the feasibility of π, and the second equality was established above. Because, by definition, $\bar{p}(\omega) \geq v(x) + q(x) \cdot (\omega - x)$ for all (x, ω), we must have that for π-almost all (x, ω),

$$v(x) + q(x) \cdot (\omega - x) = \bar{p}(\omega).$$

It remains to show that any one of conditions *(i)* or *(ii)* imply optimality of $\pi \in \Pi(\mu_0)$. Note that we will not use the assumption that v is Lipschitz in that part of the proof.

Assume that condition *(i)* holds. Note that, under these assumptions, $\bar{p}$ is feasible for the dual (D) when viewed as a function on Ω (in particular, as shown previously, convexity and $\bar{p} \geq v$ imply that $\int_\Omega \bar{p}(\omega)\, \mathrm{d}\mu(\omega) \geq V(\mu)$, for all $\mu \in \Delta(\Omega)$). But then the fact that π_X achieves no duality gap means that π must be optimal.

Assume that condition *(ii)* holds. Note that under these assumptions, we have shown previously (using only the definition of $\bar{p}$ and the property that q is measurable with $\|q(x)\| \leq L$ for all $x \in X$) that $\bar{p}$ is feasible for the dual (D) on Ω. Moreover, by the last equation of condition *(ii)*,

$$\int_\Omega \bar{p}(\omega)\, \mathrm{d}\mu_0(\omega) = \int_{X \times \Omega} \big(v(x) + q(x) \cdot (\omega - x)\big)\, \mathrm{d}\pi(x, \omega) = \int_X v(x)\, \mathrm{d}\pi_X(x),$$

showing that $\bar{p}$ and π_X achieve no duality gap, and hence π is optimal.

A.8 *Proof of Theorem 6 and Remark 1*

Since v is continuously differentiable on the compact set X, it is L-Lipschitz on X where

$$L := \max_{x \in X} \|\nabla v(x)\| < \infty,$$

so all previous results apply. We now prove the two implications of the equivalence separately.

If. Fix $\pi \in \Pi(\mu_0)$, and let $S = \mathrm{supp}(\pi_X)$. The function p_S is convex (see footnote 22). Moreover, by condition (M), $p_S \geq v$. Thus, there exists a function q as in Lemma 7. Then, for any feasible $\tilde{\pi} \in \Pi(\mu_0)$, we have

$$\begin{aligned}
\int_{X \times \Omega} v(x)\, \mathrm{d}\tilde{\pi}(x, \omega) &= \int_{X \times \Omega} \big(v(x) + q(x) \cdot (\omega - x)\big)\, \mathrm{d}\tilde{\pi}(x, \omega) \\
&\leq \int_{X \times \Omega} p_S(\omega)\, \mathrm{d}\tilde{\pi}(x, \omega) = \int_\Omega p_S(\omega)\, \mathrm{d}\mu_0(\omega) = \int_{X \times \Omega} p_S(\omega)\, \mathrm{d}\pi(x, \omega) \\
&= \int_{X \times \Omega} \big(v(x) + \nabla v(x) \cdot (\omega - x)\big)\, \mathrm{d}\pi(x, \omega) = \int_{X \times \Omega} v(x)\, \mathrm{d}\pi(x, \omega),
\end{aligned}$$

showing that π is optimal. The inequality follows from Lemma 7. The second to last equality holds by condition (M). The remaining equalities follow from the feasibility of $\tilde{\pi}$ and π.

Only if. Fix an optimal distribution $\pi \in \Pi(\mu_0)$. By Theorem 5, there exists an optimal solution p to (D) and it is convex on Ω. Define the convex roof extension $\check{p}$ of p from Ω

to X, as in formula (R). For each $x \in X$, the infimum in the definition of $\check{p}(x)$ is attained at some $\mu_x \in \Delta(\Omega)$. By feasibility of p, for any $x \in X$,

$$\check{p}(x) = \int_\Omega p(\omega)\,\mathrm{d}\mu_x(\omega) \geq V(\mu_x) = v(x).$$

Consequently,

$$\int_X v(x)\,\mathrm{d}\pi_X(x) \leq \int_X \check{p}(x)\,\mathrm{d}\pi_X(x) \leq \int_\Omega \check{p}(\omega)\,\mathrm{d}\mu_0(\omega) = \int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega),$$

where the first inequality holds because $\check{p} \geq v$, the second inequality holds because $\check{p}$ is convex and μ_0 is a mean-preserving spread of π_X, and the equality holds because $\check{p}$ coincides with p on Ω, given that p is convex on Ω. Hence, condition *(i)* in Theorem 5 implies that all inequalities hold with equality,

$$\int_X v(x)\,\mathrm{d}\pi_X(x) = \int_X \check{p}(x)\,\mathrm{d}\pi_X(x) = \int_\Omega \check{p}(\omega)\,\mathrm{d}\mu_0(\omega).$$

Thus, $\pi_X(\check{S}) = 1$, where $\check{S} = \{x \in X : v(x) = \check{p}(x)\}$. Since X is closed, v is continuous, $\check{p}$ is lower semicontinuous (see footnote 19), and the set $\check{S}$ can be equivalently written as $\check{S} = \{x \in X : v(x) \geq \check{p}(x)\}$, it follows that the set $\check{S}$ is closed. Thus, $\mathrm{supp}(\pi_X) \subset \check{S}$.

Taking into account that v is continuously differentiable and $\check{p}$ is convex and satisfies $\check{p} \geq v$, we obtain that $\check{p}$ has a subgradient $\nabla v(x)$ at each $x \in \check{S}$, so, for all $y \in X$,

$$\check{p}(y) \geq \check{p}(x) + \nabla v(x) \cdot (y - x) = v(x) + \nabla v(x) \cdot (y - x).$$

Indeed, for $x \in \check{S}$, $y \in X$, and $\varepsilon > 0$, we have

$$\check{p}(y) - \check{p}(x) \geq \frac{1}{\varepsilon}\big(\check{p}\big(x + \varepsilon(y - x)\big) - \check{p}(x)\big) \geq \frac{1}{\varepsilon}\big(v\big(x + \varepsilon(y - x)\big) - v(x)\big),$$

where the first inequality is by convexity of $\check{p}$, and the second inequality is by $\check{p} \geq v$ and $\check{p}(x) = v(x)$. Taking $\varepsilon \downarrow 0$ yields that $\nabla v(x)$ is a subgradient of $\check{p}$ at $x \in \check{S}$.

Thus, since $\pi \in \Pi(\mu_0)$ and $p = \check{p}$ on Ω, we have

$$\int_\Omega p(\omega)\,\mathrm{d}\mu_0(\omega) \geq \int_{X\times\Omega} \big(v(x) + \nabla v(x) \cdot (\omega - x)\big)\,\mathrm{d}\pi(x, \omega) = \int_{X\times\Omega} v(x)\,\mathrm{d}\pi(x, \omega).$$

As shown above, the inequality holds with equality, so $\pi(\check{\Gamma}) = 1$, where

$$\check{\Gamma} = \big\{(x, \omega) \in \check{S} \times \Omega : \check{p}(\omega) = v(x) + \nabla v(x) \cdot (\omega - x)\big\}.$$

Note that the set $\check{\Gamma}$ is closed, given that $\check{S}$ and Ω are closed and ∇v and $\check{p}$ are continuous on X and Ω, respectively. Thus, $\mathrm{supp}(\pi) \subset \check{\Gamma}$. But then we have that, for all $\omega \in \Omega$,

$$p_{\mathrm{supp}(\pi_X)}(\omega) = \max_{x\in\mathrm{supp}(\pi_X)} \big\{v(x) + \nabla v(x) \cdot (\omega - x)\big\} = \check{p}(\omega),$$

where the first equality is by the definition of p_S, and the second equality is by $\mathrm{supp}(\pi) \subset \check{\Gamma}$. This shows that $p_{\mathrm{supp}(\pi_X)}(\omega) = \check{p}(\omega) = p(\omega)$ for $\omega \in \Omega$, and hence also that

$p_{\text{supp}(\pi_X)}(x) = \check{p}(x)$ for $x \in X$. Thus, we have shown that $p_{\text{supp}(\pi_X)}$ satisfies condition (M), which completes the proof of the theorem.

Finally, we explain why the above proof also implies Remark 1. First, note that in the "only if" part of the proof we established $p_{\text{supp}(\pi_X)} \equiv \check{p}$ for an arbitrary optimal π. It follows that $S^\star$, as defined in Remark 1, is equal to $\check{S}$ in the proof (note that $\check{S}$ does not depend on which optimal solution π we consider). Thus, we also have that $p_{S^\star} \equiv \check{p}$.

Fix a feasible $\pi \in \Pi(\mu_0)$. Suppose that π is optimal for (P_M). Then the "only if" part of the above proof shows that $\text{supp}(\pi_X) \subset \check{S}$ and $\text{supp}(\pi) \subset \check{\Gamma}$. As argued in the previous paragraph, we can replace $\check{S}$ with $S^\star$ and $\check{p}$ with $p_{S^\star}$, and hence condition (M) holds with $S = S^\star$. Conversely, if $\text{supp}(\pi_X) \subset S^\star$ and condition (M) holds with $S = S^\star$, then the "if" part of the proof shows that π is optimal for (P_M).

A.9 *Generalized analysis for Section 4.4*

In this Appendix, we set up generalized notation that agrees with the notation defined in Section 4.4 in the special case of convex Ω but may differ in the general case of non-convex Ω. In Appendix A.10, we use this generalized notation to prove Theorems 7 and 8 without assuming that Ω is convex.

It will be convenient to consider solutions $\pi \in \Pi(\mu_0)$ on the extended space $X \times X$ even though $\text{supp}(\pi) \subseteq X \times \Omega$. To make our notation more intuitive, we will use the symbols $x, y, z \in X$ to refer to moments, and $\omega \in X$ to refer to the "extended states."

For a closed set $S \subset X$, let $p_S : X \to \mathbb{R}$ be defined as in Section 4.3. Let $S^\star$ be defined as in Remark 1. Specifically, $S^\star$ is the closed subset of X such that

$$S^\star = \big\{x \in X : p_{S^\star}(x) = v(x)\big\},$$

and condition (M) holds with $S = S^\star$ (for any optimal solution π). Define the function $p^\star : X \to \mathbb{R}$,

$$p^\star(\omega) := \max_{x \in S^\star}\big\{v(x) + \nabla v(x) \cdot (\omega - x)\big\}, \quad \text{for all } \omega \in X.$$

Note that this definition agrees with the one introduced in Section 4.4 when $\Omega = X$ because $p^\star$ and $p_{S^\star}$ coincide on Ω; however, $p^\star$ and $p_{S^\star}$ may differ on $X \setminus \Omega$.

Define the contact set $\Gamma \subset X \times X$,

$$\Gamma := \big\{(x, \omega) \in S^\star \times X : p^\star(\omega) = v(x) + \nabla v(x) \cdot (\omega - x)\big\},$$

and its x-section,

$$\Gamma_x := \big\{\omega \in X : (x, \omega) \in \Gamma\big\}, \quad \text{for all } x \in S^\star.$$

To extend Theorem 7, we must first define convex-partitional signals for the case when Ω is not necessarily a convex set. To circumvent this difficulty, we define the partition on the convex hull of Ω (i.e., on X), and we require each element of the partition of X to be convex.[35] Formally, we say that $\pi \in \Pi(\mu_0)$ is *convex-partitional* if there is a measurable function $\chi : X \to X$ such that, for all measurable sets $A \subset X$ and $B \subset \Omega$,

$$\pi(A, B) = \int_B \mathbf{1}\{\chi(\omega) \in A\} \, \mathrm{d}\mu_0(\omega),$$

and, for all $x \in X$, the set $\chi^{-1}(x)$ is convex.

A.10 *Proof of Theorems 7 and 8*

In this Appendix, we rely on the general notation set up in Appendix A.9.

Before proceeding to the proofs of Theorems 7 and 8, we state and prove a key lemma. Define the correspondence $\mathcal{X} : X \rightrightarrows X$ by

$$\mathcal{X}(\omega) := \arg\max_{x \in S^\star} \{v(x) + \nabla v(x) \cdot (\omega - x)\}, \quad \text{for all } \omega \in X,$$

and fix any measurable selection $\chi : X \to X$ from $\mathcal{X}$, which exists by the measurable maximum theorem (Theorem 18.19 in Aliprantis and Border (2006)). We start with a key lemma that we will be using throughout.

Lemma 8.

(i) The function $p^\star$ is convex and Lipschitz on X. Moreover, $p^\star$ is differentiable at any $\omega \in \mathrm{int}(X)$ if and only if the set $\{\nabla v(x) : x \in \mathcal{X}(\omega)\}$ is a singleton, and in that case $\nabla p^\star(\omega) = \nabla v(x)$ for all $x \in \mathcal{X}(\omega)$.

(ii) The set $\Gamma \subseteq X \times X$ is closed. Its projection along the first coordinate is $S^\star$, and its projection along the second coordinate is X. For each $x \in S^\star$, Γ_x is a compact convex set such that $x \in \Gamma_x$ and

$$\Gamma_x = \arg\min_{\omega \in X} \{p^\star(\omega) - \nabla v(x) \cdot \omega\}.$$

Moreover, for any $x, y \in S^\star$, we have

(a) $\nabla v(x) = \nabla v(y) \implies \Gamma_x = \Gamma_y$;

(b) $\mathrm{relint}(\Gamma_x) \cap \mathrm{relint}(\Gamma_y) \neq \emptyset \implies \Gamma_x = \Gamma_y$;

(c) $\mathrm{relint}(\Gamma_x) \cap \Gamma_y \neq \emptyset \implies \Gamma_x \subset \Gamma_y$.

Proof. *(i)* Clearly, $p^\star$ is convex on X as a pointwise maximum of affine functions. Moreover, it is Lipschitz on X, because for any $\omega, \omega' \in X$,

$$\begin{aligned} p^\star(\omega) - p^\star(\omega') &\leq v(\chi(\omega)) + \nabla v(\chi(\omega)) \cdot (\omega - \chi(\omega)) - v(\chi(\omega)) - \nabla v(\chi(\omega)) \cdot (\omega' - \chi(\omega)) \\ &= \nabla v(\chi(\omega)) \cdot (\omega - \omega') \leq L\|\omega - \omega'\|, \end{aligned}$$

[35] To understand why we adopt this convention, consider the distribution π induced by no disclosure. Intuitively, pooling all states should correspond to a convex-partitional signal. However, the support of this distribution over states conditional on the induced moment is equal to Ω, and is hence not convex when Ω is not convex. We circumvent this by defining the partition on X; then the unique element of that partition corresponding to no disclosure is X itself, a convex set. And of course, this partition restricted to Ω still represents no disclosure.

with L defined (as in Appendix A.8) as the maximal value of the norm of the gradient of v on X.

The remainder of part *(i)* is a consequence of the envelope theorem. For $N = 1$, this follows immediately from Corollary 4 in Milgrom and Segal (2002). Below, we extend their analysis to the general case $N \geq 1$.

Suppose, by contradiction, that $p^\star$ is differentiable at $\omega \in \text{int}(X)$ but there exist $x, y \in \mathcal{X}(\omega)$ such that $\nabla v(x) \neq \nabla v(y)$. Denote $u := \nabla v(x) - \nabla v(y)$, so that $\nabla v(x) \cdot u > \nabla v(y) \cdot u$. Since $\omega \in \text{int}(X)$, we have $\omega \pm hu \in X$ for small enough $h > 0$. Moreover, by the definitions of $p^\star$ and $\mathcal{X}$,

$$\frac{p^\star(\omega + hu) - p^\star(\omega)}{h} \geq \nabla v(x) \cdot u \quad \text{and} \quad \frac{p^\star(\omega - hu) - p^\star(\omega)}{h} \geq -\nabla v(y) \cdot u,$$

and thus

$$-\lim_{h \downarrow 0} \frac{p^\star(\omega - hu) - p^\star(\omega)}{h} \leq \nabla v(y) \cdot u < \nabla v(x) \cdot u \leq \lim_{h \downarrow 0} \frac{p^\star(\omega + hu) - p^\star(\omega)}{h},$$

showing that $p^\star$ is not differentiable at ω.

Conversely, suppose that $\omega \in \text{int}(X)$ and $\{\nabla v(x) : x \in \mathcal{X}(y)\}$ is a singleton. Fix any $u \in \mathbb{R}^N$ and small enough $h'' > h' > 0$, so that $\omega + h'u$ and $\omega + h''u$ are both in X. By the definition of $p^\star$,

$$\nabla v\big(\chi(\omega + h'u)\big) \cdot u \leq \frac{p^\star\big(\omega + h''u\big) - p^\star\big(\omega + h'u\big)}{h'' - h'} \leq \nabla v\big(\chi(\omega + h''u)\big) \cdot u.$$

Taking the limit superior in this inequality as $h' \downarrow 0$ yields

$$\limsup_{h' \downarrow 0} \nabla v\big(\chi(\omega + h'u)\big) \cdot u \leq \frac{p^\star\big(\omega + h''u\big) - p^\star(y)}{h''} \leq \nabla v\big(\chi(\omega + h''u)\big) \cdot u.$$

Taking the limit inferior in the resulting inequality as $h'' \downarrow 0$ yields

$$\limsup_{h' \downarrow 0} \nabla v\big(\chi(\omega + h'u)\big) \cdot u \leq \lim_{h'' \downarrow 0} \frac{p^\star\big(\omega + h''u\big) - p^\star(\omega)}{h''} \leq \liminf_{h'' \downarrow 0} \nabla v\big(\chi(\omega + h''u)\big) \cdot u.$$

Since the limit superior is never smaller than the limit inferior, we conclude that the two limits coincide, and hence

$$\lim_{h \downarrow 0} \frac{p^\star(\omega + hu) - p^\star(\omega)}{h} = \lim_{h \downarrow 0} \nabla v\big(\chi(\omega + hu)\big) \cdot u.$$

Since the correspondence $\mathcal{X} : X \rightrightarrows X$ is upper hemicontinuous, a version of Berge's maximum theorem (see Lemma 17.30 in Aliprantis and Border (2006)) yields

$$\lim_{h \downarrow 0} \frac{p^\star(\omega + hu) - p^\star(\omega)}{h} = \lim_{h \downarrow 0} \nabla v\big(\chi(\omega + hu)\big) \cdot u \leq \max_{x \in \mathcal{X}(\omega)} \nabla v(x) \cdot u.$$

Since $\{\nabla v(x) : x \in \mathcal{X}(\omega)\}$ is a singleton, we have $\max_{x\in\mathcal{X}(\omega)} \nabla v(x)\cdot u = \nabla v(x)\cdot u$ for all $x \in \mathcal{X}(\omega)$. Finally, taking into account that, by the definition of $p^\star$, for any $x \in \mathcal{X}(\omega)$ and any small enough $h > 0$, we have

$$\nabla v(x)\cdot u \leq \frac{p^\star(\omega+hu) - p^\star(\omega)}{h},$$

it follows that

$$\lim_{h\downarrow 0} \frac{p^\star(\omega+hu) - p^\star(\omega)}{h} = \nabla v(x)\cdot u, \quad \text{for all } x \in \mathcal{X}(\omega),$$

showing that $p^\star$ is differentiable at y and $\nabla p^\star(\omega) = \nabla v(x)$ for all $x \in \mathcal{X}(\omega)$.

(ii) The set Γ is closed, because the function $p^\star(\omega) - v(x) - \nabla v(x)\cdot(\omega - x)$ is continuous in (x, ω) on $X \times X$. The projection of Γ along the second coordinate is X, because $(\chi(\omega), \omega) \in \Gamma$ for each $\omega \in X$. The projection of Γ along the first coordinate is $S^\star$ by the definition of $S^\star$ and the fact that Γ_x is nonempty, for any $x \in S^\star$, which is shown in the next paragraph.

Fix any $x \in S^\star$. We have

$$\begin{aligned}\Gamma_x &= \{\omega \in X : p^\star(\omega) = v(x) + \nabla v(x)\cdot(\omega - x)\} \\ &= \{\omega \in X : p^\star(\omega) \leq v(x) + \nabla v(x)\cdot(\omega - x)\},\end{aligned}$$

where the first equality is by the definition of Γ and Γ_x, and the second equality is by the definition of $p^\star$, which, in particular, implies that

$$p^\star(\omega) \geq v(x) + \nabla v(x)\cdot(\omega - x), \quad \text{for all } \omega \in X.$$

Thus, the set Γ_x is compact and convex, as it is a sublevel set of the (Lipschitz) continuous and convex function $p^\star(\omega) - v(x) - \nabla v(x)\cdot(\omega - x)$ (viewed as a function of ω). Moreover, we have $x \in \Gamma_x$, because

$$v(x) = p_{S^\star}(x) \geq p^\star(x) \geq v(x),$$

where the equality is by $x \in S^\star$, the first inequality is by the definition of $p_{S^\star}$, and the last inequality is by the definition of $p^\star$ and the fact that $x \in S^\star$. Since $p^\star(x) = v(x)$, we have

$$p^\star(\omega) \geq p^\star(x) + \nabla v(x)\cdot(\omega - x), \quad \text{for all } \omega \in X,$$

and thus

$$\Gamma_x = \arg\max_{\omega\in X}\{\nabla v(x)\cdot\omega - p^\star(\omega)\}.$$

We have thus shown that Γ_x is the projection along the first coordinate of the face of the epigraph of $p^\star$ exposed by the direction $(-1, \nabla v(x))$. Then implication (a) is immediate, whereas implications (b) and (c) follow from Corollary 18.1.2 and Theorem 18.1 in Rockafellar (1970). For completeness, we provide short self-contained proofs of (b) and (c). To show (c), let $\omega \in \operatorname{relint}(\Gamma_x) \cap \Gamma_y$. Since Γ_x is convex, for any $\omega' \in \Gamma_x$ with $\omega' \neq \omega$, there exists $\omega'' \in \Gamma_x$ and $\lambda \in (0, 1)$ such that $\omega = \lambda\omega' + (1-\lambda)\omega''$. Next, by the definition of $p^\star$, we have

$$p^\star(\omega') \geq v(y) + \nabla v(y) \cdot (\omega' - y) \quad \text{and} \quad p^\star(\omega'') \geq v(y) + \nabla v(y) \cdot (\omega'' - y).$$

Both inequalities must hold with equality, as otherwise we would have

$$p^\star(\omega) \geq \lambda p^\star(\omega') + (1-\lambda)p^\star(\omega'') > v(y) + \nabla v(y) \cdot (\omega - y),$$

contradicting that $\omega \in \Gamma_y$. Since ω' is arbitrary, we get $\Gamma_x \subset \Gamma_y$, proving (c). To prove (b), notice that if $\text{relint}(\Gamma_x) \cap \text{relint}(\Gamma_y) \neq \emptyset$, then $\text{relint}(\Gamma_x) \cap \Gamma_y \neq \emptyset$ and $\text{relint}(\Gamma_y) \cap \Gamma_x \neq \emptyset$, implying that $\Gamma_x \subset \Gamma_y$ and $\Gamma_y \subset \Gamma_x$, and thus $\Gamma_x = \Gamma_y$. □

In the remainder, we complete the proofs of Theorems 7 and 8. To deal with the general case in which $\Omega \neq X$, we follow Appendix A.9 and consider solutions defined on the larger space X rather than on Ω. All the notation used in the following proof completions is then defined as in Appendix A.9, and becomes consistent with the notation used in the main text under the assumption that Ω is convex (so that $\Omega = X$).

Completion of the proof of Theorem 7 Let $\tilde{X}$ be the set of interior points of X where $p^\star$ is differentiable. The set of boundary points of the convex set X is Lebesgue-negligible, by Theorem 1 in Lang (1986). The set of interior points of X where $p^\star$ is not differentiable is Lebesgue-negligible by Rademacher's theorem (Theorem 10.8 in Villani (2009)). Thus, taking into account that μ_0 has a density on X, the set $\tilde{X}$ has full measure under μ_0: $\mu_0(\tilde{X}) = 1$.

Fix $\omega \in \tilde{X}$. We claim that $|\mathcal{X}(\omega)| = 1$. Suppose, by contradiction, that there exist distinct $x, y \in \mathcal{X}(\omega)$. Since $\omega \in \text{int}(X)$ and $p^\star$ is differentiable at ω, part *(i)* of Lemma 8 yields

$$\nabla p^\star(\omega) = \nabla v(x) = \nabla v(y).$$

In turn, part *(ii)* of Lemma 8 yields $x \in \Gamma_x$, $y \in \Gamma_y$, and $\Gamma_x = \Gamma_y$, and thus, given that $p^\star$ is affine on Γ_x by the definition of Γ_x, we have $p^\star(y) = p^\star(x) + \nabla p^\star(\omega) \cdot (y - x)$ or, equivalently,

$$v(x) - \nabla v(x) \cdot x = v(y) - \nabla v(y) \cdot y.$$

Next, for all $\lambda \in [0, 1]$, we have $p_{S^\star}(\lambda x + (1-\lambda)y) = \lambda v(x) + (1-\lambda)v(y)$ as follows from

$$\begin{aligned}\lambda v(x) + (1-\lambda)v(y) &= \lambda p^\star(x) + (1-\lambda)p^\star(y) = p^\star(\lambda x + (1-\lambda)y) \\ &\leq p_{S^\star}(\lambda x + (1-\lambda)y) \leq \lambda p_{S^\star}(x) + (1-\lambda)p_{S^\star}(y) \\ &= \lambda v(x) + (1-\lambda)v(y),\end{aligned}$$

where the first equality is by $x \in \Gamma_x$ and $y \in \Gamma_y$, the second equality is by affinity of $p^\star$ on the convex set $\Gamma_x = \Gamma_y$, the first and second inequality follow from the definition of $p_{S^\star}$, and the last equality is by $p_{S^\star} = v$ on $S^\star$. Thus, since $p_{S^\star} \geq v$ on X, we get

$$\lambda v(x) + (1-\lambda)v(y) \geq v(\lambda x + (1-\lambda)y), \quad \text{for all } \lambda \in [0, 1].$$

This contradicts the conditions of the theorem. Thus, $\mathcal{X}(\omega)$ is a singleton $\{\chi(\omega)\}$ for each $\omega \in \tilde{X}$, where $\chi(\omega)$ is determined by

$$\{\chi(\omega)\} = \{x \in S^\star : \omega \in \Gamma_x\} = \{x \in S^\star : \nabla p^\star(\omega) = \nabla v(x)\}.$$

The first equality is by the definition of $\mathcal{X}$, and the second is by part *(i)* of Lemma 8.

Finally, for any optimal $\pi \in \Pi(\mu_0)$, we have

$$1 = \pi(\Gamma) = \pi\left(\bigcup_{\omega \in \tilde{X}} (\{\chi(\omega)\} \times \{\omega\})\right),$$

where the first equality is by Remark 1, and the second equality is by $\Gamma = \bigcup_{\omega \in X}(\mathcal{X}(\omega) \times \{\omega\})$, $\mathcal{X}(\omega) = \{\chi(\omega)\}$ for $\omega \in \tilde{X}$, and $\mu_0(\tilde{X}) = 1$. Since $\chi(\omega)$ is determined by $p^\star$ for μ_0-almost all $\omega \in X$, and $p^\star$ is independent of π, we conclude that π is uniquely determined by

$$\pi(A, B) = \int_B \mathbf{1}\{\chi(\omega) \in A\}\, \mathrm{d}\mu_0(\omega), \quad \text{for all measurable } A \subset X \text{ and } B \subset X.$$

Completion of the proof of Theorem 8 Fixing any solution to the primal problem (P_M) and the corresponding price function, define the set $S^\star$, the contact set Γ, and the sets Γ_x as in Appendix A.9. Recall that $S_x = \mathrm{cl}(\mathrm{supp}(\pi_X) \cap \mathrm{relint}(\Gamma_x))$. By Theorem 1 in Larman (1971), X can be partitioned (up to a measure zero set) into a collection of disjoint (relatively) open sets $\Xi = \{\mathrm{relint}(\Gamma_x)\}_{x \in S^\star}$ (where we ignore duplicates whenever $\Gamma_x = \Gamma_y$ for $x \neq y$).

Consider an auxiliary problem of finding a joint distribution $\pi \in \Pi(\mu_0)$ to maximize $\int_{X \times X} w(x, \omega)\, \mathrm{d}\pi(x, \omega)$, where

$$w(x, \omega) := \begin{cases} -\|x\|^2, & (x, \omega) \in \Gamma, \\ -\infty, & (x, \omega) \in (X \times X) \setminus \Gamma. \end{cases}$$

Note that $\int_{X \times X} w(x, \omega)\, \mathrm{d}\pi(x, \omega)$ is finite for $\pi \in \Pi(\mu_0)$ if and only if $\mathrm{supp}(\pi) \subset \Gamma$, which in turn is equivalent to optimality of $\pi \in \Pi(\mu_0)$ for the primary problem. Since w is upper semicontinuous and bounded from above, by Lemma 2, there exists an optimal solution $\pi \in \Pi(\mu_0)$ to the auxiliary problem, which is also optimal for the primal problem (P_M). We fix such $\pi \in \Pi(\mu_0)$.

Intuitively, the auxiliary problem selects a solution to the primal problem (P_M) that minimizes the average norm of the induced posterior means. The rest of the proof shows that if the set S_x induced by π differs from $\mathrm{ext}(S_x)$ on a positive measure set of $x \in \mathrm{supp}(\pi_X)$, we would obtain a contradiction with π solving the auxiliary problem. While this conclusion is intuitive, the details of the proof are complicated by the fact that the selection induced by the auxiliary problem may be "local" in the sense that it affects the structure of the solution on uncountably many measure-zero sets. Our strategy is to decompose the distribution π into conditional distributions conditional on each induced $\mathrm{relint}(\Gamma_x)$.

Note that we can treat the set Ξ as a measurable space, endowing it with the Borel σ-algebra generated by the Hausdorff metric. We can then define π_Ξ to be the probability distribution over Ξ induced by π: For any measurable subset $A \subset \Xi$,

$$\pi_\Xi(A) := \pi\big(\{(x, \omega) \in X \times X : \omega \in \text{relint}(\Gamma_x), \text{relint}(\Gamma_x) \in A\}\big).$$

By the disintegration theorem (e.g., Theorem 2.3 in Caravenna and Daneri (2010)), there exists a measurable function $\xi \mapsto \pi(\cdot|\xi)$ from Ξ to $\Delta(X \times X)$ such that for every "test function" $h \in C(X \times X)$, we have

$$\int_{X\times X} h(x, \omega)\, \mathrm{d}\pi(x, \omega) = \int_\Xi \int_{X\times X} h(x, \omega)\, \mathrm{d}\pi(x, \omega|\xi)\, \mathrm{d}\pi_\Xi(\xi).$$

Let $\pi_X(\cdot|\xi)$ and $\pi_\Omega(\cdot|\xi)$, for $\xi \in \Xi$, denote the marginal distributions of x and ω (i.e., the first and second coordinate, respectively) induced by $\pi(\cdot|\xi)$. Then, for π_X-almost all $x \in X$, we have

$$\text{supp}\big(\pi_\Omega\big(\cdot|\,\text{relint}(\Gamma_x)\big)\big) \subset \text{cl}\big(\text{relint}(\Gamma_x)\big), \tag{A.1}$$

$$\text{supp}\big(\pi_X\big(\cdot|\,\text{relint}(\Gamma_x)\big)\big) = S_x, \tag{A.2}$$

$$\int_{A\times X} (\omega - x)\, \mathrm{d}\pi\big(x, \omega|\,\text{relint}(\Gamma_x)\big) = 0, \quad \text{for all measurable } A \subset X, \tag{A.3}$$

$$\int w(x, \omega)\, \mathrm{d}\pi\big(x, \omega|\,\text{relint}(\Gamma_x)\big) \ge \int w(x, \omega)\, \mathrm{d}\tilde{\pi}(x, \omega), \quad \text{for all } \tilde{\pi} \in \Pi\big(\pi_\Omega\big(\cdot|\,\text{relint}(\Gamma_x)\big)\big), \tag{A.4}$$

where the first three properties follow from definitions, and the last inequality must be true because otherwise we would have a contradiction with the definition of π as the solution to the auxiliary problem.

Toward a contradiction, suppose that there exists a π_X-positive-measure set of points x such that $S_x \neq \text{ext}(S_x)$; that is, there exist distinct $x^0, x^1, \dots, x^n \in S_x$ such that $x^0 = \lambda^1 x^1 + \dots + \lambda^n x^n$, where $\lambda^1, \dots, \lambda^n > 0$ and $\lambda^1 + \dots + \lambda^n = 1$. (We suppress the dependence of these variables on x.) By condition (A.2), since $x^1, \dots, x^n \in S_x$, for all $i = 1, \dots, n$, and $\delta > 0$, we have $\pi_X(B_\delta(x^i)|\,\text{relint}(\Gamma_x)) > 0$, where $B_\delta(x^i)$ denotes an open ball with radius δ centered at x^i. To simplify notation, let $\pi^i_\delta(\cdot)$ denote the conditional probability measure on X induced from $\pi_X(\cdot|\,\text{relint}(\Gamma_x))$ by conditioning on the event $B_\delta(x^i)$. There exists a sufficiently small δ such that for some $\lambda^1_\delta, \dots, \lambda^n_\delta > 0$ with $\lambda^1_\delta + \dots + \lambda^n_\delta = 1$, we have $x^0 = \lambda^1_\delta x^1_\delta + \dots + \lambda^n_\delta x^n_\delta$ where $x^i_\delta = \int_X x\, \mathrm{d}\pi^i_\delta(x)$. Finally, by condition (A.3), for some sufficiently small $\epsilon > 0$, there exists $\tilde{\pi} \in \Pi(\pi_\Omega(\cdot|\,\text{relint}(\Gamma_x)))$ such that, for all measurable $A \subset X$,

$$\tilde{\pi}_X(A) = \pi_X\big(A|\,\text{relint}(\Gamma_x)\big) + \varepsilon\delta_{x^0} - \varepsilon \sum_i \lambda^i_\delta \pi^i_\delta(A),$$

where δ_{x^0} denotes the Dirac measure at x^0. Intuitively, $\tilde{\pi}_X$ modifies $\pi_X(\cdot|\,\text{relint}(\Gamma_x))$ by transferring some mass from the neighborhoods of points x^i into x^0. But then, by

Jensen's inequality, and relying on conditions (A.1) and (A.2) to ensure that $\text{supp}(\pi(\cdot|\text{relint}(\Gamma_x)) \subset \Gamma$ and $\text{supp}(\tilde{\pi}) \subset \Gamma$, we have

$$\int_{X\times X} w(x,\omega)\,\mathrm{d}\tilde{\pi}(x,\omega) - \int_{X\times X} w(x,\omega)\,\mathrm{d}\pi\big(x,\omega|\,\text{relint}(\Gamma_x)\big)$$
$$= \varepsilon\bigg(\sum_i \lambda_\delta^i \int_X x^2\,\mathrm{d}\pi_\delta^i(x) - \big(x^0\big)^2\bigg) \geq \varepsilon\bigg(\sum_i \lambda_\delta^i \big(x_\delta^i\big)^2 - \big(x^0\big)^2\bigg) > 0,$$

yielding a contradiction with (A.4).

A.11 *Proof of Proposition 2*

In this Appendix, we prove the necessity part of Proposition 2. Fix an optimal $\pi^\star \in \Pi(\mu_0)$. Since μ_0 has a density and $\nabla v(x) = (x_2, x_1) \neq (y_2, y_1) = \nabla v(y)$ for $x \neq y$, Theorem 7 implies that $\pi^\star$ is the unique optimal signal, and that it is convex-partitional. Suppose that $\text{supp}(\pi_X^\star)$ is the graph of the function f, as described in the proposition.

By the definition of Γ_x from Section 4.4, for each $t \in [\underline{x}_1, \overline{x}_1]$,

$$\Gamma_{(t,f(t))} = \{\omega \in \Omega : t \in \underset{s\in[\underline{x}_1,\overline{x}_1]}{\arg\max}\big\{\omega_1 f(s) + \omega_2 s - sf(s)\big\}.$$

First, consider $t \in (\underline{x}_1, \overline{x}_1)$. The necessary first-order condition yields $\omega_2 = f(t) - f'(t)(\omega_1 - t)$ for all $\omega \in \Gamma_{(t,f(t))}$. Define, for all $t \in [\underline{x}_1, \overline{x}_1]$,

$$\underline{l}_t := \min_{\omega\in X}\{\omega_1 - t\},$$
$$\text{subject to } \omega_2 = f(t) - f'(t)(\omega_1 - t),$$
$$\omega_2 + \frac{(t-\omega_1)\big(f(t)-\omega_2\big)}{s-\omega_1} \leq f(s), \quad \text{for all } s \in (\omega_1, \overline{x}_1],$$

and

$$\overline{l}_t := \max_{\omega\in X}\{\omega_1 - t\},$$
$$\text{subject to } \omega_2 = f(t) - f'(t)(\omega_1 - t),$$
$$\omega_2 + \frac{(t-\omega_1)\big(f(t)-\omega_2\big)}{s-\omega_1} \geq f(s), \quad \text{for all } s \in [\underline{x}_1, \omega_1).$$

Notice that $(t + \underline{l}_t, f(t) - f'(t)\underline{l}_t)$ and $(t + \overline{l}_t, f(t) - f'(t)\overline{l}_t)$ are the points in $\Gamma_{(t,f(t))}$ with the lowest and highest first coordinate. To see this, consider $\omega \in \Gamma_{(t,f(t))}$ with $t > \omega_1$ (and thus $f(t) - \omega_2 = -f'(t)(t-\omega_1) < 0$) and notice that, for $s \leq \omega_1$, we have $f(s) \leq f(t) < \omega_2$; thus,

$$(t-\omega_1)\big(f(t)-\omega_2\big) < 0 \leq (s-\omega_1)\big(f(s)-\omega_2\big).$$

Consequently, $\omega \in \Omega$ with $\omega_1 < t$ belongs to $\Gamma_{(t,f(t))}$ if and only if

$$\omega_2 = f(t) - f'(t)(\omega_1 - t),$$

$$\omega_2 + \frac{(t-\omega_1)\big(f(t)-\omega_2\big)}{s-\omega_1} \le f(s), \quad \text{for all } s \in (\omega_1, \overline{x}_1].$$

Since $(t, f(t)) \in \Gamma_{(t,f(t))}$, it follows that $(t+\underline{l}_t, f(t)-f'(t)\underline{l}_t)$ is indeed the point in $\Gamma_{(t,f(t))}$ with the lowest first coordinate. An analogous argument shows that $(t+\overline{l}_t, f(t)-f'(t)\overline{l}_t)$ is the point in $\Gamma_{(t,f(t))}$ with the highest first coordinate. Finally, since, by Lemma 8, $\Gamma_{(t,f(t))}$ is convex, it follows that

$$\Gamma_{(t,f(t))} = \text{cl}(I_t) := \big\{\omega \in \Omega : \omega_1 = x_1 + l, \omega_2 = f(x_1) - f'(x_1)l, l \in \big[\underline{l}(x_1), \overline{l}(x_1)\big]\big\}.$$

It turns out that the above condition also holds for $x \in \text{supp}(\pi^\star_X)$ with $x_1 \in \{\underline{x}_1, \overline{x}_1\}$. However, the proof of that fact is significantly more complicated.

Lemma 9. *$\Gamma_{(t,f(t))} = cl(I_t)$ for $t \in \{\underline{x}_1, \overline{x}_1\}$.*

Proof. See Appendix A.11.1. □

By Lemma 9, we can conclude that $\Gamma_{(t,f(t))} = \text{cl}(I_t)$ for each $t \in [\underline{x}_1, \overline{x}_1]$. Since the projection of the contact set Γ along the second coordinate is $X = \Omega$, it follows that $\Omega = \bigcup_{t\in[\underline{x}_1,\overline{x}_1]} \text{cl}(I_t)$. Define $I_t = \text{relint}(\text{cl}(I_t))$, for $t \in [\underline{x}_1, \overline{x}_1]$.[36] By part *(ii)(b)* in Lemma 8, for $t \neq s$, the open line segments I_t and I_s do not intersect. In fact, part *(ii)(c)* in Lemma 8 yields a stronger conclusion that, for $t \neq s$, the closed line segments $\text{cl}(I_t)$ and $\text{cl}(I_s)$ can intersect only at a common endpoint. Thus, as in the proof of Theorem 7, invoking Theorem 1 in Larman (1971), we conclude that $\Omega \setminus \{\bigcup_{t\in[\underline{x}_1,\overline{x}_1]} I_t\}$ has zero (Lebesgue) measure. In sum, we have established that there exists a collection $\{I_t\}_{t\in[\underline{x}_1,\overline{x}_1]}$ of open disjoint line segments that partition Ω, up to a measure-zero set.

Condition *(i)* then follows directly from the above characterization of the optimal signal $\pi^\star$ and the assumption that $\text{supp}(\pi^\star_X) = \text{Gr}(f)$. Condition *(ii)* follows from the definition of $\text{cl}(I_t)$. Moreover, the inclusion $I_t \subseteq \{\omega \in \Omega : \omega_2 = f(t) - f'(t)(\omega_1 - t)\}$, for $t \in [\underline{x}_1, \overline{x}_1]$, follows directly from the fact that $\Gamma_{(t,f(t))} = \text{cl}(I_t)$ for each $t \in [\underline{x}_1, \overline{x}_1]$. This finishes the proof of the proposition.

A.11.1 *Proof of Lemma 9* We start by proving yet another lemma.

Lemma 10. *There exists $\varepsilon > 0$ such that*

$$\omega_2 + \frac{(x_1-\omega_1)\big(f(x_1)-\omega_2\big)}{y_1-\omega_1} < f(y_1),$$

for all $x_1 \in [\underline{x}_1, \overline{x}_1]$, $\omega_1 \in [x_1-\varepsilon, x_1]$, $y_1 \in (\omega_1, x_1) \cup (x_1, \overline{x}_1]$, and $\omega_2 = f(x_1) - f'(x_1)(\omega_1 - x_1)$.

Proof. Since f' and f'' are continuous and $f' > 0$ on the compact set $[\underline{x}_1, \overline{x}_1]$, we have $\underline{f'} = \min_{\tilde{x}_1\in[\underline{x}_1,\overline{x}_1]} f'(\tilde{x}_1) > 0$ and $\underline{f''} = \min_{\tilde{x}_1\in[\underline{x}_1,\overline{x}_1]} f''(\tilde{x}_1) \in \mathbb{R}$. Thus, there exists $\varepsilon > 0$

[36] Note that I_t is a point when $\text{cl}(I_t)$ is degenerate, since a point is a relatively open set.

such that $2\underline{f}' + \varepsilon\underline{f}'' > 0$. Fix such ε. By direct calculation,

$$\omega_2 + \frac{(x_1 - \omega_1)\big(f(x_1) - \omega_2\big)}{y_1 - \omega_1} < f(x_1) + f'(x_1)\varepsilon - \frac{f'(x_1)\varepsilon^2}{y_1 - x_1 + \varepsilon},$$

for all $x_1 \in [\underline{x}_1, \overline{x}_1]$, $\omega_1 \in (x_1 - \varepsilon, x_1]$, $y_1 \in (\omega_1, x_1) \cup (x_1, \overline{x}_1]$, and $\omega_2 = f(x_1) - f'(x_1)(\omega_1 - x_1)$. Thus, it suffices to show that

$$f(x_1) + f'(x_1)\varepsilon - \frac{f'(x_1)\varepsilon^2}{y_1 - x_1 + \varepsilon} < f(y_1),$$

for all $x_1 \in [\underline{x}_1, \overline{x}_1]$ and $y_1 \in (x_1 - \varepsilon, x_1) \cup (x_1, \overline{x}_1]$.

If $\underline{f}'' \geq 0$, the inequality holds because the right-hand side $f(y_1)$ is convex in y_1 with derivative $f'(x_1)$ at x_1, while the left-hand side is strictly concave in y_1 with derivative $f'(x_1)$. So, assume that $\underline{f}'' < 0$ and denote

$$\hat{y}_1 = x_1 + \frac{f'(x_1) - \underline{f}'}{-\underline{f}''}.$$

Since $f''(y_1) \geq \underline{f}''$ and $f'(y_1) \geq \underline{f}'$, for all $y_1 \in [\underline{x}_1, \overline{x}_1]$, we have $f(y_1) \geq \underline{f}(y_1)$, where

$$\underline{f}(y_1) = \begin{cases} f(x_1) + f'(x_1)(y_1 - x_1) + \dfrac{\underline{f}''}{2}(y_1 - x_1)^2, & y_1 \leq \hat{y}_1, \\ f(x_1) + f'(x_1)(\hat{y}_1 - x_1) + \dfrac{\underline{f}''}{2}(\hat{y}_1 - x_1)^2 + \underline{f}'(y_1 - \hat{y}_1), & y_1 > \hat{y}_1. \end{cases}$$

Thus, it suffices to show that

$$f(x_1) + f'(x_1)\varepsilon - \frac{f'(x_1)\varepsilon^2}{y_1 - x_1 + \varepsilon} < \underline{f}(y_1). \tag{A.5}$$

By direct calculation, for $y_1 \in (x_1 - \varepsilon, x_1) \cup (x_1, \hat{y}_1]$, inequality (A.5) is equivalent to $2f'(x_1) + \underline{f}''(y_1 - x_1 + \varepsilon) > 0$, which holds if and only if it holds at $\hat{y}_1$. At $\hat{y}_1$, (A.5) simplifies to $f'(x_1) + \underline{f}' + \underline{f}''\varepsilon > 0$, which holds because $2\underline{f}' + \varepsilon\underline{f}'' > 0$. Again, by direct calculation, for $y_1 > \hat{y}_1$, inequality (A.5) is equivalent to

$$\frac{\big(f'(x_1) - \underline{f}'\big)^2}{2\big(-\underline{f}''\big)}(y_1 - x_1 + \varepsilon) + \underline{f}'(y_1 - x_1)(y_1 - x_1 + \varepsilon) - f'(x_1)(y_1 - x_1)\varepsilon > 0,$$

where the left-hand side is quadratic and convex in y_1. Moreover, the derivative at $y_1 = \hat{y}_1$ is positive because $3\underline{f}' + f'(x_1) + 2\underline{f}''\varepsilon > 0$, as follows from $2\underline{f}' + \varepsilon\underline{f}'' > 0$. Thus, the left-hand side is increasing in y_1 and inequality (A.5) holds for $y_1 > \hat{y}_1$, because it holds for $y_1 = \hat{y}_1$, as shown above. □

We are now ready to prove Lemma 9. We will focus on the case $t = \underline{x}_1$ since the other case is fully analogous. The necessary Kuhn–Tucker condition yields $\omega_2 \leq f(\underline{x}_1) - f'(\underline{x}_1)(\omega_1 - \underline{x}_1)$ for all $\omega \in \Gamma_{(\underline{x}_1, f(\underline{x}_1))}$. We claim that $\omega_2 \geq f(\underline{x}_1) - f'(\underline{x}_1)(\omega_1 - \underline{x}_1)$ for all

$\omega \in X$, and thus $\omega_2 = f(\underline{x}_1) - f'(\underline{x}_1)(\omega_1 - \underline{x}_1)$ for all $\omega \in \Gamma_{(\underline{x}_1, f(\underline{x}_1))}$, so $\Gamma_{(\underline{x}_1, f(\underline{x}_1))} = \mathrm{cl}(I_{\underline{x}_1})$, by the same argument as previously. Toward a contradiction, suppose that there exists $z \in X$ such that $z_2 < f(\underline{x}_1) - f'(\underline{x}_1)(z_1 - \underline{x}_1)$ and $z_1 < x_1$ (the case $z_1 > x_1$ is analogous and omitted). Since X is convex and the graph of f is a maximal monotone set in X, it follows that $z_2 > f(\underline{x}_1)$ and that there exists $\varepsilon > 0$ such that, for all $\omega_1 \in (\underline{x}_1 - \varepsilon, \underline{x}_1)$, points $(\omega_1, f(\underline{x}_1) - f'(\underline{x}_1)(\omega_1 - \underline{x}_1))$ and $(\omega_1, f(\underline{x}_1) - \iota(\omega_1 - \underline{x}_1))$ with $\iota = (z_2 - f(\underline{x}_1))/(x_1 - z_1) \in (0, f'(\underline{x}_1))$ belong to X. It is easy to see that, for all $\omega_1 < \underline{x}_1$ and $y_1 > x_1$, we have

$$\omega_2 - \iota(\omega_1 - x_1) - \iota \frac{(\underline{x}_1 - \omega_1)^2}{y_1 - \omega_1} < \omega_2 - f'(x_1)(\omega_1 - x_1) - f'(x_1)\frac{(\underline{x}_1 - \omega_1)^2}{y_1 - \omega_1}.$$

Thus, by Lemma 10, for sufficiently small $\varepsilon > 0$, points $(\omega_1, f(\underline{x}_1) - f'(\underline{x}_1)(\omega_1 - \underline{x}_1))$ and $(\omega_1, f(\underline{x}_1) - \iota(\omega_1 - \underline{x}_1))$ belong to Γ_x. But then Γ_x has a nonempty interior, and all points in the interior belong only to Γ_x, by Lemma 8. Consequently, since μ_0 has full support density on X,

$$\begin{aligned}&\int_{\Gamma_x} \big(\omega_2 - f(x_1) - f'(x_1)(\omega_1 - x_1)\big)\, \mathrm{d}\mu_0(\omega) \\ &\quad = \int_{\mathrm{int}(\Gamma_x)} \big(\omega_2 - f(x_1) - f'(x_1)(\omega_1 - x_1)\big)\, \mathrm{d}\mu_0(\omega) < 0,\end{aligned}$$

as the boundary of the convex set Γ_x has zero Lebesgue measure, by Theorem 1 in Lang (1986), and the integrand is strictly negative on the interior of Γ_x, as implied by the Kuhn–Tucker condition. This shows that any π supported on Γ cannot be in $\Pi(\mu_0)$, as it violates the second constraint in the definition of $\Pi(\mu_0)$—a contradiction.

A.12 *Proof of Proposition 3*

Suppose that $\pi \in \Pi(\mu_0)$, induced by the disclosure of the realization of $a\omega_1 + \omega_2$, is optimal. Define $\Theta = \{\theta = a\omega_1 + \omega_2 : \omega \in \Omega\}$. Since Ω is a compact convex set with a nonempty interior, we have $\Theta = [\underline{\theta}, \overline{\theta}]$ for some $\underline{\theta} < \overline{\theta}$. By Proposition 1, $\mathrm{supp}(\pi_X)$ is a monotone set. Thus, since μ_0 has full-support density on Ω, we have $\mathrm{supp}(\pi_X) = \{(x_1(\theta), x_2(\theta)) : \theta \in \Theta\}$ for some nondecreasing functions x_1 and x_2 satisfying $ax_1(\theta) + x_2(\theta) = \theta$ for all $\theta \in \Theta$ and $(x_1(\theta), x_2(\theta)) \in \mathrm{int}(\Omega)$ for almost all $\theta \in \Theta$. Note that x_1 is $1/a$-Lipschitz and x_2 is 1-Lipschitz, and thus $\tilde{\Theta} = \{\theta \in \Theta : (x_1(\theta), x_2(\theta)) \in \mathrm{int}(\Omega)\}$ is an open set of full measure.

Lemma 11. *For each $\theta \in \tilde{\Theta}$, there exists $\delta > 0$ such that, for all $\theta' \in (\theta - \delta, \theta + \delta)$,*

$$a\big(x_1(\theta') - x_1(\theta)\big) = x_2(\theta') - x_2(\theta) = \frac{1}{2}\big(\theta' - \theta\big).$$

Proof. Since $\theta \in \tilde{\Theta}$, there exists $\varepsilon > 0$ such that $\omega \in \mathrm{int}(\Omega)$ for all $\omega \in \mathbb{R}^2$ such that $\omega_1 \in (x_1(\theta) - \varepsilon, x_1(\theta) + \varepsilon)$ and $\omega_2 \in (x_2(\theta) - \varepsilon, x_2(\theta) + \varepsilon)$. Fix $\delta = \min\{\varepsilon/2, a\varepsilon/2\}$. We claim that for all $\theta' \in (\theta - \delta, \theta + \delta)$ and all $\omega' \in \mathbb{R}^2$ such that $\omega'_2 \in (x_2(\theta') - \delta, x_2(\theta') + \delta)$ and $a\omega'_1 + \omega'_2 = \theta'$, we have $x(\theta') \in \mathrm{int}(\Omega)$ and $\omega' \in \mathrm{int}(\Omega)$. Indeed, since x_1 and x_2 are nondecreasing and satisfy $ax_1(\theta') + x_2(\theta') = \theta'$, we have $x_1(\theta') \in (x_1(\theta) - \delta/a, x_1(\theta) + \delta/a)$

and $x_2(\theta') \in (x_2(\theta)-\delta, x_2(\theta)+\delta)$, so $x(\theta') \in \text{int}(\Omega)$. Next, since $a\omega_1' + \omega_2' = \theta' = ax_1(\theta') + x_2(\theta')$ and $\omega_2' \in (x_2(\theta') - \delta, x_2(\theta') + \delta)$, we have $\omega_2' \in (x_2(\theta') - \delta, x_2(\theta') + \delta) \subset (x_2(\theta) - 2\delta, x_2(\theta) + 2\delta)$ and $\omega_1' \in (x_1(\theta') - \delta/a, x_1(\theta') + \delta/a) \subset (x_1(\theta) - 2\delta/a, x_1(\theta) + 2\delta/a)$, so $\omega' \in \text{int}(\Omega)$.

Fix $\theta' \in (\theta-\delta, \theta+\delta)$ and an integer $n > 0$. For $i \in \{0, \dots, n\}$, define $\theta^i = \theta + (\theta' - \theta)i/n$, $\omega^{Li} = (x_1(\theta^i) - \delta/a, x_2(\theta^i) + \delta)$, and $\omega^{Ri} = (x_1(\theta^i) + \delta/a, x_2(\theta^i) - \delta)$. As shown in the previous paragraph, we have $x(\theta^i), \omega^{Li}, \omega^{Ri} \in \text{int}(\Omega)$ for all i. Next, by Theorem 6, we have, for all $i \in \{0, \dots, n\}$,

$$\begin{aligned}
&\big(x_1(\theta^i) - \omega_1^{Li}\big)\big(x_2(\theta^i) - \omega_2^{Li}\big) \le \big(x_1\big(\theta^{i+1}\big) - \omega_1^{Li}\big)\big(x_2\big(\theta^{i+1}\big) - \omega_2^{Li}\big) \\
&\iff \quad x_2\big(\theta^{i+1}\big) - x_2\big(\theta^i\big) \ge \frac{a\big(x_1\big(\theta^{i+1}\big) - x_1\big(\theta^i\big)\big)}{1 + \frac{a}{\delta}\big(x_1\big(\theta^{i+1}\big) - x_1\big(\theta^i\big)\big)},
\end{aligned}$$

and

$$\begin{aligned}
&\big(x_1(\theta^i) - \omega_1^{Ri}\big)\big(x_2(\theta^i) - \omega_2^{Ri}\big) \le \big(x_1\big(\theta^{i+1}\big) - \omega_1^{Ri}\big)\big(x_2\big(\theta^{i+1}\big) - \omega_2^{Ri}\big) \\
&\iff \quad x_2\big(\theta^{i+1}\big) - x_2\big(\theta^i\big) \le \frac{a\big(x_1\big(\theta^{i+1}\big) - x_1\big(\theta^i\big)\big)}{1 - \frac{a}{\delta}\big(x_1\big(\theta^{i+1}\big) - x_1\big(\theta^i\big)\big)}.
\end{aligned}$$

Since x_1 is $1/a$-Lipschitz, we have, for all $i \in \{0, \dots, n\}$,

$$\frac{a\big(x_1\big(\theta^{i+1}\big) - x_1\big(\theta^i\big)\big)}{1 + \frac{1}{n\delta}\big(\theta' - \theta\big)} \le x_2\big(\theta^{i+1}\big) - x_2\big(\theta^i\big) \le \frac{a\big(x_1\big(\theta^{i+1}\big) - x_1\big(\theta^i\big)\big)}{1 - \frac{1}{n\delta}\big(\theta' - \theta\big)}.$$

Summing over $i \in \{0, \dots, n-1\}$ gives

$$\frac{a\big(x_1\big(\theta'\big) - x_1(\theta)\big)}{1 + \frac{1}{n\delta}\big(\theta' - \theta\big)} \le x_2\big(\theta'\big) - x_2(\theta) \le \frac{a\big(x_1\big(\theta'\big) - x_1(\theta)\big)}{1 - \frac{1}{n\delta}\big(\theta' - \theta\big)}.$$

Since n is arbitrary, we have $x_2(\theta') - x_2(\theta) = a(x_1(\theta') - x_1(\theta))$. Taking into account that $ax_1(\theta) + x_2(\theta) = \theta$ and $ax_1(\theta') + x_2(\theta') = \theta'$ completes the proof of the lemma. □

Since $\tilde{\Theta}$ is an open set in $\mathbb{R}$, it is the union of at most countably many disjoint open intervals $(\underline{\theta}^i, \overline{\theta}^i)$. Lemma 11 implies that

$$a\big(x_1\big(\theta'\big) - x_1(\theta)\big) = x_2\big(\theta'\big) - x_2(\theta) = \frac{1}{2}\big(\theta' - \theta\big), \quad \text{for all } \theta', \theta \in (\underline{\theta}^i, \overline{\theta}^i).$$

Since $\tilde{\Theta}$ has full Lebesgue measure, it follows that $\text{cl}(\tilde{\Theta}) = \Theta$. Since x_1 and x_2 are (Lipschitz) continuous, we have

$$a\big(x_1\big(\theta'\big) - x_1(\theta)\big) = x_2\big(\theta'\big) - x_2(\theta) = \frac{1}{2}\big(\theta' - \theta\big), \quad \text{for all } \theta', \theta \in \Theta,$$

and thus $x_2(\theta) = ax_1(\theta) + b$ for all $\theta \in \Theta$ and some $b \in \mathbb{R}$.

References

Aliprantis, Charalambos D. and Kim Border (2006), *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer-Verlag, Berlin. [1704, 1708, 1714, 1735, 1741, 1742]

Arieli, Itai, Yakov Babichenko, and Fedor Sandomirskiy (2023). "Persuasion as transportation." Available at arXiv:2307.07672. [1726]

Arieli, Itai, Yakov Babichenko, Rann Smorodinsky, and Takuro Yamashita (2023), "Optimal persuasion via bi-pooling." *Theoretical Economics*, 18, 15–36. [1702, 1711, 1716, 1718, 1719]

Benoît, Jean-Pierre and Juan Dubra (2011), "Apparent overconfidence." *Econometrica*, 79, 1591–1625. [1725]

Bergemann, Dirk and Stephen Morris (2016), "Bayes correlated equilibrium and the comparison of information structures in games." *Theoretical Economics*, 11, 487–522. [1726]

Bergemann, Dirk and Stephen Morris (2019), "Information design: A unified perspective." *Journal of Economic Literature*, 57, 44–95. [1701]

Bogachev, Vladimir (2007), *Measure Theory*. Springer-Verlag, Berlin. [1705, 1728]

Brezis, Haim (2011), *Functional Analysis, Sobolev Spaces and Partial Differential Equations*. Springer, New York. [1729]

Bucicovschi, Orest and Jiri Lebl (2013), "On the continuity and regularity of convex extensions." *Journal of Convex Analysis*, 4, 1113–1126. [1713]

Candogan, Ozan and Philipp Strack (2023), "Optimal disclosure of information to privately informed agents." *Theoretical Economics*, 18, 1225–1269. [1725]

Caplin, Andrew and Mark Dean (2013), "Behavioral implications of rational inattention with Shannon entropy." NBER Working Paper No. 19318. [1725]

Caplin, Andrew and Mark Dean (2015), "Revealed preference, rational inattention, and costly information acquisition." *American Economic Review*, 105, 2183–2203. [1725]

Caravenna, Laura and Sara Daneri (2010), "The disintegration of the Lebesgue measure on the faces of a convex function." *Journal of Functional Analysis*, 258, 3604–3661. [1746]

Daskalakis, Constantinos, Alan Deckelbaum, and Christos Tzamos (2017), "Strong duality for a multiple-good monopolist." *Econometrica*, 85, 735–767. [1706]

De March, Hadrien and Nizar Touzi (2019), "Irreducible convex paving for decomposition of multidimensional martingale transport plans." *Annals of Probability*, 47, 1726–1774. [1720]

Denti, Tommaso (2022), "Posterior separable cost of information." *American Economic Review*, 112, 3215–3259. [1725]

Dizdar, Deniz and Eugen Kováč (2020), "A simple proof of strong duality in the linear persuasion problem." *Games and Economic Behavior*, 122, 407–412. [1701, 1702, 1709, 1713, 1714, 1732]

Doval, Laura and Vasiliki Skreta (2024), "Constrained information design." *Mathematics of Operations Research*, 49, 78–106. [1725]

Dworczak, Piotr (2020), "Mechanism design with aftermarkets: Cutoff mechanisms." *Econometrica*, 88, 2629–2661. [1731]

Dworczak, Piotr and Anton Kolotilin (2019), "The persuasion duality." Available at arXiv:1910.11392. [1704, 1725, 1730]

Dworczak, Piotr and Anton Kolotilin (2024), "The persuasion duality." Available at arXiv:1910.11392. [1710, 1713, 1714, 1717, 1720, 1722]

Dworczak, Piotr and Giorgio Martini (2019), "The simple economics of optimal persuasion." *Journal of Political Economy*, 127, 1993–2048. [1701, 1702, 1705, 1709, 1713, 1714, 1717]

Edwards, David A. (2011), "On the Kantorovich–Rubinstein theorem." *Expositiones Mathematicae*, 29, 387–398. [1705]

Gale, David (1960), *The Theory of Linear Economic Models*. McGraw-Hill, New York. [1705]

Gale, David (1967), "A geometric duality theorem with economic applications." *Review of Economic Studies*, 34, 19–24. [1702, 1708, 1733, 1735]

Galperti, Simone, Aleksandr Levkun, and Jacopo Perego (2024), "The value of data records." *Review of Economic Studies*, 91, 1007–1038. [1701, 1705, 1725]

Galperti, Simone and Jacopo Perego (2018), "A dual perspective on information design." Available at SSRN, https://ssrn.com/abstract=3297406. [1726]

Gentzkow, Matthew and Emir Kamenica (2014), "Costly persuasion." *American Economic Review*, 104, 457–462. [1725]

Gretsky, Neil E., Joseph M. Ostroy, and William R. Zame (2002), "Subdifferentiability and the duality gap." *Positivity*, 6, 261–274. [1708]

Guo, Yingni and Eran Shmaya (2019), "The interval structure of optimal disclosure." *Econometrica*, 87, 653–675. [1725]

Guo, Yingni and Eran Shmaya (2021), "Costly miscalibration." *Theoretical Economics*, 16, 477–506. [1709]

Hanin, Leonid G. (1992), "Kantorovich–Rubinstein norm and its application in the theory of Lipschitz spaces." *Proceedings of the American Mathematical Society*, 115, 345–352. [1705]

Holmes, Richard B. (1975), *Geometric Functional Analysis and Its Applications*, volume 24. Springer-Verlag, New York. [1708]

Inostroza, Nicolas and Alessandro Pavan (2023), "Adversarial coordination and public information design." Available at SSRN, https://ssrn.com/abstract=4531654. [1726]

Kamenica, Emir (2019), "Bayesian persuasion and information design." *Annual Review of Economics*, 11, 249–272. [1701]

Kamenica, Emir and Matthew Gentzkow (2011), "Bayesian persuasion." *American Economic Review*, 101, 2590–2615. [1701, 1703, 1704, 1726]

Kantorovich, Leonid V. and Gennadii Sh. Rubinstein (1958), "On a space of completely additive functions." *Vestnik Leningrad Univ. Math.*, 13, 52–59. (In Russian). [1705]

Kleiner, Andreas and Alejandro Manelli (2019), "Strong duality in monopoly pricing." *Econometrica*, 87, 1391–1396. [1706]

Kleiner, Andreas, Benny Moldovanu, and Philipp Strack (2021), "Extreme points and majorization." *Economic Applications. Econometrica*, 89, 1557–1593. [1702, 1716, 1718, 1719]

Kolotilin, Anton (2018), "Optimal information disclosure: A linear programming approach." *Theoretical Economics*, 13, 607–636. [1701, 1702, 1712, 1713, 1726]

Kolotilin, Anton, Roberto Corrao, and Alexander Wolitzky (2024), "Persuasion and matching: Optimal productive transport." *Journal of Political Economy* (forthcoming). [1701, 1702, 1705, 1710, 1713, 1726]

Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li (2017), "Persuasion of a privately informed receiver." *Econometrica*, 85, 1949–1964. [1725]

Kolotilin, Anton and Alexander Wolitzky (2024), "Distributions of posterior quantiles via matching." *Theoretical Economics*, 19, 1399–1413. [1725]

Kramkov, Dmitry and Yan Xu (2022), "An optimal transport problem with backward martingale constraints motivated by insider trading." *Annals of Applied Probability*, 32, 294–326. [1720]

Lang, Robert (1986), "A note on the measurability of convex sets." *Archiv der Mathematik*, 47, 90–92. [1744, 1750]

Laraki, Rida (2004), "On the regularity of the convexification operator on a compact set." *Journal of Convex Analysis*, 11, 209–234. [1709]

Larman, David G. (1971), "On a conjecture of Klee and Martin for convex bodies." *Proceedings of the London Mathematical Society*, 3, 668–682. [1745, 1748]

Le Treust, Maël and Tristan Tomala (2019), "Persuasion with limited communication capacity." *Journal of Economic Theory*, 184, 104940. [1725]

Malamud, Semyon and Andreas Schrimpf (2022), "Persuasion by dimension reduction." Available at arXiv:2110.08884. [1702, 1703, 1717, 1721]

Mathevet, Laurent, Jacopo Perego, and Ina Taneva (2020), "On information design in games." *Journal of Political Economy*, 128, 1370–1404. [1726]

Milgrom, Paul and Ilya Segal (2002), "Envelope theorems for arbitrary choice sets." *Econometrica*, 70, 583–601. [1742]

Nikandrova, Arina and Romans Pancs (2017), "Conjugate information disclosure in an auction with learning." *Journal of Economic Theory*, 171, 174–212. [1710, 1720, 1721]

Obłój, Jan and Pietro Siorpaes (2017), "Structure of martingale transports in finite dimensions." Available at arXiv:1702.08433. [1720]

Onuchic, Paula and Debraj Ray (2023), "Conveying value via categories." *Theoretical Economics*, 18, 1407–1439. [1720]

Rayo, Luis (2013), "Monopolistic signal provision." *B.E. Journal of Theoretical Economics*, 13, 27–58. [1720]

Rayo, Luis and Ilya Segal (2010), "Optimal information disclosure." *Journal of Political Economy*, 118, 949–987. [1702, 1720, 1721]

Rochet, Jean-Charles and Jean-Luc Vila (1994), "Insider trading without normality." *Review of Economic Studies*, 61, 131–152. [1720]

Rockafellar, Ralph Tyrell (1970), *Convex Analysis*. Princeton University Press, Princeton, New Jersey. [1707, 1735, 1736, 1743]

Santambrogio, Filippo (2015), *Optimal Transport for Applied Mathematicians*. Birkäuser, New York, NY. [1731]

Shaked, Moshe and J. George Shanthikumar (1994), *Stochastic Orders*. Springer, New York. [1712]

Smolin, Alex and Takuro Yamashita (2024), "Information design in smooth games." Available at arXiv:2202.10883. [1726]

Tamura, Wataru (2018), "Bayesian persuasion with quadratic preferences." Available at SSRN, https://ssrn.com/abstract=1987877. [1720]

Taneva, Ina (2019), "Information design." *American Economic Journal: Microeconomics*, 11, 151–185. [1726]

Villani, Cédric (2003), *Topics in Optimal Transportation*, volume 58 of Graduate Studies in Mathematics. American Mathematical Society, Providence, RI. [1731]

Villani, Cédric (2009), *Optimal Transport, Old and New*. Springer-Verlag, Berlin. [1703, 1706, 1744]

Yang, Kai Hao and Alex Zentefis (2024), "Monotone function intervals: Theory and applications." *American Economic Review* 114, 2239–2270. [1725]

Zhang, Jun and Junjie Zhou (2016) "Information disclosure in contests: A Bayesian persuasion approach." *Economic Journal*, 126, 2197–2217. [1710]

Co-editor Simon Board handled this manuscript.

Manuscript received 1 November, 2023; final version accepted 13 June, 2024; available online 25 June, 2024.