

Arieli, Itai; Babichenko, Yakov; Smorodinsky, Rann; Yamashita, Takuro

Article

Optimal persuasion via bi-pooling

Theoretical Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Arieli, Itai; Babichenko, Yakov; Smorodinsky, Rann; Yamashita, Takuro (2023) :
Optimal persuasion via bi-pooling, Theoretical Economics, ISSN 1555-7561, The Econometric
Society, New Haven, CT, Vol. 18, Iss. 1, pp. 15-36,
<https://doi.org/10.3982/TE4663>

This Version is available at:

<https://hdl.handle.net/10419/296403>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/4.0/>

Optimal persuasion via bi-pooling

ITAI ARIELI

Faculty of Industrial Engineering and Management, Technion–Israel Institute of Technology

YAKOV BABICHENKO

Faculty of Industrial Engineering and Management, Technion–Israel Institute of Technology

RANN SMORODINSKY

Faculty of Industrial Engineering and Management, Technion–Israel Institute of Technology

TAKURO YAMASHITA

Toulouse School of Economics, University of Toulouse Capitole

Mean-preserving contractions are critical for studying Bayesian models of information design. We introduce the class of *bi-pooling policies*, and the class of *bi-pooling distributions* as their induced distributions over posteriors. We show that every extreme point in the set of all mean-preserving contractions of any given prior over an interval takes the form of a bi-pooling distribution. By implication, every Bayesian persuasion problem with an interval state space admits an optimal bi-pooling distribution as a solution, and conversely, for every bi-pooling distribution, there is a Bayesian persuasion problem for which that distribution is the *unique* solution.

KEYWORDS. Bayesian persuasion, mean preserving contraction, bi-pooling, information disclosure, price function, signaling, extreme points.

JEL CLASSIFICATION. C72, D82, D83.

1. INTRODUCTION

The recent prolific literature on information design, known as Bayesian persuasion, studies optimal information disclosure policies when the informed player has commitment power.

Characterizing an optimal information disclosure policy when the state space is *large* turns out to be a hard problem. Several recent papers study different aspects of

Itai Arieli: iarieli@technion.ac.il

Yakov Babichenko: yakovbab@technion.ac.il

Rann Smorodinsky: rann@technion.ac.il

Takuro Yamashita: takuro.yamashita@tse-fr.eu

This research is supported by the European Research Council (starting grant #714693), ANR grant (ANR-17-EURE-0010, Investissements d’Avenir program), Israel Science Foundation grants #2018889 and #2030524, United States–Israel Binational Science Foundation grants #2016734 and #2026924, German–Israel Foundation grant I-1419-118.4/2017, Ministry of Science and Technology grant #19400214, Technion VPR grants, and the Bernard M. Gordon Center for Systems Engineering at the Technion. The authors thank Piotr Dworczak and the anonymous reviewers for valuable comments.

persuasion with large state spaces (e.g., Candogan (2019a,b), Deniz and Kováč (2019), Dworzak and Kolotilin (2019), Dworzak and Martini (2019), Gentzkow and Kamenica (2016), Kolotilin (2018), Kolotilin, Mylovanov, Zapechelnyuk, and Li (2017), Kolotilin and Wolitzky (2020), Yamashita (2018)).¹ The model studied and motivated in these papers, which is the one we study, considers an interval state space, and assumes that the receiver's action, and hence the sender's utility, is a function of the posterior mean over the state space. One application of this model is to study a sender who has partial information over an underlying binary state space. To see this, one should think of the unit interval as the set of all probabilities assigned to one of the two states.²

To gain insights into the structure of an optimal persuasion policy in this model, Dworzak and Martini (2019) formulate the dual problem associated with the sender's optimization problem. They interpret the dual problem as a "persuasion economy" and its solution, dubbed the *price function*, as the corresponding Walrasian equilibrium of the economy. Their main result is that there is no duality gap between the two problems.³ The important implication of this result is that the price function can be used to verify the optimality of a given arbitrary persuasion policy. In other words, given a candidate persuasion policy one can use the aforementioned price function to verify whether it is indeed optimal or not.

Our contribution In this paper, we introduce the class of *bi-pooling policies*, and the corresponding class of *bi-pooling distributions* over posteriors these policies induce. Our main technical contribution is to show that every extreme point in the set of all mean-preserving contractions of any given prior takes the form of a bi-pooling distribution. By implication, every Bayesian persuasion problem admits an optimal bi-pooling distribution as a solution, and conversely, for every bi-pooling distribution, there is a Bayesian persuasion problem for which that distribution is the *unique* solution.

A bi-pooling policy is quite simple. This policy partitions the state space (the unit interval) into two sets, one of which is a union of disjoint open intervals. Whenever the state is in one such interval, the sender reveals which interval it is in and provides an additional (possibly uninformative) binary signal (hence the name "bi-pooling"). In the complement of the union of intervals, the sender completely reveals the state; see Figure 1.

Based on our characterization result, we obtain some additional insights into the optimal signaling policies. For example, we show that there always exists a monotone solution for such linear (i.e., posterior mean-based) Bayesian persuasion problems. More precisely, a monotone signaling policy is such that for any two distinct states the conditional distribution over the posterior means of a higher state first-order stochastically

¹Large state spaces are commonly used also in information design applications; see, for example, Roesler and Szentes (2017), Harbaugh and Rasmusen (2018), Boleslavsky, Hwang, and Kim (2019).

²Under this interpretation, if the receiver knows that the sender's belief about the state is distributed according to a distribution P , then his subjective belief about the state will be $\mathbb{E}[P]$. Hence, the sender's utility is a function of the posterior mean, as required.

³In recent work, Dworzak and Kolotilin (2019) extend the duality approach beyond the single-dimension setting while Deniz and Kováč (2019) use the duality approach of Dworzak and Martini (2019) to provide a simpler proof of their results while allowing for weaker assumptions.




FIGURE 1. A bi-pooling policy. The policy bi-pools the interval $(0.1, 0.6)$ into the two posterior means $\{0.2, 0.4\}$ by sending an additional binary signal. The policy pools the interval $(0.8, 1)$ into the posterior mean 0.9 without sending an additional signal. In the complementary intervals $([0, 0.1]$ and $[0.6, 0.8])$, the policy is fully revealing.

dominates that of the lower state. As argued in Mensch (2019), there are various persuasion environments where monotone signals are compelling. This may be due to regulatory measures or moral hazard issues. For example, in the credit rating framework of Goldstein and Leitner (2018), banks whose assets are being rated may sue if they believe that they have been unfairly treated by the rating agency for giving them a lower credit rating than banks with riskier assets. More than that, nonmonotonic signals may introduce an incentive for banks to take unnecessary risks in order to get a higher rating. Similar considerations come up in other applications such as grading schemes for students.

In addition, we consider a variety of economic applications where our approach is useful in identifying the optimal signaling policies. We leverage the duality result of Dworzak and Martini (2019), which states that the optimal policy can be found by identifying its corresponding *price function* (under certain regularity conditions). Often, however, the set of potential price functions is too large to easily identify the right one. Our results about feasible bi-pooling policies are shown to be useful in imposing significant restrictions on the structure of feasible price functions, so that the search for the right price function (and hence for the optimal policy) can become much more straightforward. In particular, we revisit the framework of Kolotilin, Mylovanov, and Zapechelnyuk (2022), where the receiver has an outside option whose value is her private information. We provide an intuition to their characterization of the optimal policy for S-shaped utility functions. We extend their framework to multiple receivers, yielding an M-shaped utility function, and we characterize the corresponding optimal policy.

Related results Some variations of our main results have been developed independently and almost concurrently. Candogan (2019a) restricts attention to the case where the sender's utility is increasing and piecewise constant and proves that this admits an optimal bi-pooling policy. Candogan (2020) dispenses with the monotonicity assumption. Kleiner, Moldovanu, and Strack (2021) independently characterize the extreme points of the set of all mean-preserving contractions for a given distribution as the set of all mean-preserving bi-pooling distributions, and make the connection with the persuasion problem as well as other economic models. None of the papers mentioned above

show that all the extreme points are *exposed*, namely, that any bi-pooling policy serves as a unique optimizer for some persuasion problems.⁴

In a different model where the state space is small but the receiver has private information, Guo and Shmaya (2019) show the optimality of nested interval policies. Interestingly, such policies with just two nested intervals give rise to bi-pooling policies (see Lemma 4 below).

2. MODEL AND MAIN RESULTS

We consider a persuasion model where the state space is the interval $[0, 1]$ with a nonatomic common prior $F \in \Delta([0, 1])$ that has full support.⁵ The sender knows the realized state and the receiver is uninformed. Prior to the realization of the state, the sender commits to a *signaling policy* $\pi : [0, 1] \rightarrow \Delta(S)$, where S is an arbitrary measurable space. Once the state $\omega \in [0, 1]$ is realized, the sender sends a signal $s \in S$ according to the committed signaling policy $\pi(\omega)$. Observing s , the receiver forms a posterior belief about the state, and consequently, its posterior mean. Without loss of generality, we may assume that $S = [0, 1]$, and that the posterior mean of the state, given signal s , is s itself. Under this normalization, it is immediate that the distribution of the posterior mean s given the signal policy π , denoted by $F_\pi \in \Delta([0, 1])$, is a *mean-preserving contraction* of F (henceforth MPC).⁶ Let $\text{MPC}(F)$ represent the set of all mean-preserving contractions of F . It is well known that for any $G \in \text{MPC}(F)$, there exists a signaling policy π that *implements* G in the sense that $F_\pi = G$ (e.g., Gentzkow and Kamenica (2016), Kolotilin (2018)). In general, there may be multiple signaling policies that implement the same $G \in \text{MPC}(F)$.

The sender's indirect utility is denoted by $u : [0, 1] \rightarrow \mathbb{R}$, where $u(x)$ is the sender's expected utility in case the receiver's posterior mean is x . We assume that u is upper semi-continuous and refers to the pair (F, u) as a *persuasion problem*. The sender's problem takes the following simple form:

$$\max_{G \in \text{MPC}(F)} \mathbb{E}_{x \sim G}[u(x)]. \quad (1)$$

It is also well known that a distribution $G \in \Delta([0, 1])$ lies in $\text{MPC}(F)$ if and only if it satisfies:

- For all $x \in [0, 1]$: $\int_0^x G(x) dx \leq \int_0^x F(x) dx$, and
- $\int_0^1 G(x) dx = \int_0^1 F(x) dx$ (equivalently, $\mathbb{E}_{x \sim G}[x] = \mathbb{E}_{x \sim F}[x]$),

where we abuse notation and use $F(x)$, $G(x)$ to denote $F([0, x])$ and $G([0, x])$, the corresponding CDFs.

⁴Kleiner, Moldovanu, and Strack (2021) show that all the extreme points of the set for of mean-preserving spreads are exposed.

⁵Recall that the support of a distribution F is the smallest closed set that has probability one. The choice of the unit interval is inspired by our motivation of a partially informed sender. The results generalize to any compact interval.

⁶Equivalently, F is a *mean-preserving spread* of F_π .

Note that each of the constraints induces a convex and compact (in the weak* topology) subset of distributions. Therefore, $\text{MPC}(F)$ is the intersection of convex and compact subsets and, in turn, is itself convex and compact.

REMARK 1. A possible interpretation of our model is that the true economic state is binary (0 or 1) but the sender is only partially informed, and so his belief about the state follows $F \in \Delta([0, 1])$ from the ex ante viewpoint. In this case, the players naturally care (only) about the posterior mean.

For a distribution $H \in \Delta([0, 1])$ and a measurable set $C \subseteq [0, 1]$ we denote by $H|_C$ the distribution of $h \sim H$ conditional on the event that $h \in C$.

Within the class of mean-preserving contractions of the prior F , we single out the class of *bi-pooling distributions*. These are distributions that apply a contraction only to some collection of intervals. Each interval is contracted to at most two atoms.

DEFINITION 1. A distribution $G \in \text{MPC}(F)$ is called a *bi-pooling distribution (with respect to F)* if there exists a collection of pairwise disjoint open intervals $\{(\underline{y}_i, \bar{y}_i)\}_{i \in \mathcal{A}}$ such that⁷

- For every $i \in \mathcal{A}$, $G((\underline{y}_i, \bar{y}_i)) = F((\underline{y}_i, \bar{y}_i))$ and $|\text{supp}(G|_{(\underline{y}_i, \bar{y}_i)})| \leq 2$.
- $G|_{[0, 1] \setminus \bigcup_{i \in \mathcal{A}} (\underline{y}_i, \bar{y}_i)} = F|_{[0, 1] \setminus \bigcup_{i \in \mathcal{A}} (\underline{y}_i, \bar{y}_i)}$.

For $i \in \mathcal{A}$ such that $|\text{supp}(G|_{(\underline{y}_i, \bar{y}_i)})| = 2$ we call $(\underline{y}_i, \bar{y}_i)$ a *bi-pooling interval*. Whenever $|\text{supp}(G|_{(\underline{y}_i, \bar{y}_i)})| = 1$ we call $(\underline{y}_i, \bar{y}_i)$ a *pooling interval*. In the case where all intervals are pooling intervals, we say that G is a *pooling distribution (with respect to F)*.

To motivate our interest in bi-pooling distributions, consider the following example.

EXAMPLE 1. Consider the persuasion problem (F, u) , where $F = U[0, 1]$ is the uniform distribution over $[0, 1]$ and $u : [0, 1] \rightarrow \mathbb{R}$ is an arbitrary function satisfying $u(\frac{1}{3}) = u(\frac{2}{3}) = 0$ and $u(x) < 0$ for $x \notin \{\frac{1}{3}, \frac{2}{3}\}$. The only way for the sender to receive a utility of 0 is by inducing two posterior means $\frac{1}{3}$ and $\frac{2}{3}$.

This can be done using a binary signal space $S = \{s_1, s_2\}$ by sending the signal s_1 with probability one over the interval $(\frac{1}{12}, \frac{7}{12})$ and the signal s_2 with probability one over $[0, \frac{1}{12}] \cup [\frac{7}{12}, 1]$. The posterior mean that s_1 generates is $\frac{1}{3}$ and the posterior mean s_2 generates is $\frac{2}{3}$. In addition, this policy is a bi-pooling policy for the singleton collection of intervals $\{[0, 1]\}$. \diamond

One may well ask whether an alternative policy that pools each of two intervals into their respective means could implement the same distribution. It turns out that there is no $a \in (0, 1)$ such that $\frac{1}{3} = \mathbb{E}_F[x|x \in [0, a]]$ and $\frac{2}{3} = \mathbb{E}_F[x|x \in (a, 1)]$, where the expectation conditions must be satisfied because G is a mean-preserving contraction of F and the mass of the interval $[0, a]$ ($(a, 1]$) is mapped into $\frac{1}{3}$ ($\frac{2}{3}$).

⁷We identify open intervals with nonempty subsets of $[0, 1]$ of the form $(a, b) \cap [0, 1]$, where (a, b) is an interval on the real line. For ease of notation, we let $(\underline{y}_i, \bar{y}_i)$ denote these intervals.

The example demonstrates that the simple class of *pooling* distributions (i.e., those that use pooling intervals only) is insufficient for optimal persuasion. Is the set of bi-pooling distributions sufficiently rich to solve any persuasion problem? Our main result answers this question in the affirmative.

THEOREM 1. *Every persuasion problem (F, u) admits an optimal bi-pooling distribution.*

The intuition behind Theorem 1 is simple. Recall the problem formulation (1). Note that the set of $\text{MPC}(F)$ is a convex and compact (in the weak* topology) subset of $\Delta([0, 1])$ and the objective function (the sender's expected utility) is a linear upper semicontinuous functional over this set. Therefore, by Bauer's maximum principle (see, e.g., Theorem 7.69 in Aliprantis and Border (2006)) we are guaranteed that one of its maximizers of (1) is an extreme point of $\text{MPC}(F)$. The proof of Theorem 1 then follows directly from the following proposition.⁸

PROPOSITION 1. *The set of extreme points of $\text{MPC}(F)$ is precisely the set of bi-pooling distributions.*

The proof of Proposition 1, along with all the other proofs, is relegated to the Appendix. The idea behind the proof of the proposition is as follows. Let $G \in \text{MPC}(F)$ and let $[a, b] \subset [0, 1]$ be an interval such that $\int_0^a G(x) dx = \int_0^a F(x) dx$, $\int_0^b G(x) dx = \int_0^b F(x) dx$, and $\int_0^c G(x) dx < \int_0^c F(x) dx$ for all $c \in (a, b)$. Assume that the support of G over $[a, b]$ contains three or more points; then one can decompose G into two distributions G_1 and G_2 , both in $\text{MPC}(F)$. Therefore, any such interval induced from an extreme point of $\text{MPC}(F)$ must contain at most two points in its support. The conclusion that an extreme point must be bi-pooling follows from this observation combined with some straightforward topological arguments.

The aforementioned decomposition is done as follows. Let $x_1 < x_2 < x_3$ be three points in the support of G . To obtain G_1 , remove an infinitesimal probability mass assigned to x_2 and split it between x_1 and x_3 such that the expectation over the whole interval is maintained. To obtain G_2 , do the opposite transformation, namely, remove an infinitesimal probability mass that is assigned to x_1 and x_3 and add it to x_2 in such a way that the expectation of G is maintained. If the transformed masses are small enough, then both G_1 and G_2 are guaranteed to be in $\text{MPC}(F)$ and in addition $\alpha G_1 + (1 - \alpha)G_2 = G$ for some $\alpha \in (0, 1)$.

So far, we have argued that any persuasion problem in the domain we study admits an optimal bi-pooling distribution. Can we reduce the set of signal distributions further and maintain a similar result? The answer is negative, as we now proceed to show.

THEOREM 2. *For every bi-pooling distribution $G \in \text{MPC}(F)$ there exists a continuous utility function u for which G is the unique optimal solution of⁹ (1). That is, every extreme point of $\text{MPC}(F)$ is exposed.*

⁸For a parallel result, see Theorem 2 in Kleiner, Moldovanu, and Strack (2021).

⁹Recall that the uniqueness of an optimal distribution does not imply the uniqueness of an optimal signaling policy. This is made explicit in Example 3.




FIGURE 2. An illustration of the functions p (dotted) and u (solid).

For a given prior F and a corresponding bi-pooling distribution G , the proof of Theorem 2 instructs us how to construct a utility function u such that G is the unique optimum for the persuasion problem (F, u) . The proof uses the notion of price function due to Dworzak and Martini (2019), which we will discuss in Section 4. We demonstrate the idea behind the proof of Theorem 2 in the following example.

EXAMPLE 2. Let $F = U([0, 1])$ be the uniform distribution over the unit interval. Let G be the distribution that is induced by the policy that bi-pools the interval $(0, 0.5)$ into the two points $\{0.2, 0.4\}$, fully reveals the state in the interval $[0.5, 0.7]$, and pools the interval $(0.7, 1)$ into the point 0.85. Note that this signaling policy induces a bi-pooling distribution in $\text{MPC}(F)$. To construct the utility function u , we begin with an arbitrary convex function $p : [0, 1] \rightarrow \mathbb{R}$ that is linear over $[0, 0.5]$ and $[0.7, 1]$, the bi-pooling and pooling intervals of G , and is strictly convex over $[0.5, 0.7]$. Let u be an arbitrary function satisfying $u(x) \leq p(x)$ with equality obtained only on the image in the support of G : $\{0.2, 0.4, 0.85\} \cup [0.5, 0.7]$ (see the illustration in Figure 2).¹⁰ \diamond

The arguments in the proof of Theorem 2 show that the unique optimizer of the persuasion problem (F, u) is indeed the bi-pooling distribution G .

We end this section by studying the support of bi-pooling distributions. Let G be such a distribution and let $\{(y_i, \bar{y}_i)\}_{i \in A}$ be the corresponding collection of pairwise disjoint open intervals. Any value in the complement of this collection is clearly in the support. In addition, whenever (y_i, \bar{y}_i) is a pooling interval the unique value in its support must be equal to $\mathbb{E}(F|_{(y_i, \bar{y}_i)})$. What can we say about the values in the support whenever $|\text{supp}(G|_{(y_i, \bar{y}_i)})| = 2$?

DEFINITION 2. We say that the support $\{z, \bar{z}\}$ is feasible for the interval (y, \bar{y}) if there exists a MPC of $F|_{(y, \bar{y})}$ whose support is $\{z, \bar{z}\}$.

¹⁰The convex function p equals the price function in Dworzak and Martini (2019) corresponding to the optimal signal distribution for the utility function u .

Note that the weights α and $1 - \alpha$ on the atoms at \underline{z} and \bar{z} are determined uniquely by the expectation requirement $\alpha\underline{z} + (1 - \alpha)\bar{z} = \mathbb{E}[F|_{(\underline{y}, \bar{y})}]$. The set of binary feasible supports is characterized in the following lemma.

LEMMA 1. *A support $\{\underline{z}, \bar{z}\}$ is feasible for the interval (\underline{y}, \bar{y}) if and only if $\underline{y} \leq \underline{z} \leq \mathbb{E}[F|_{(\underline{y}, \bar{y})}] \leq \bar{z} \leq \bar{y}$ and $\mathbb{E}[F|_{(C^{-1}(\underline{z}), \bar{y})}] \geq \bar{z}$, where $C(w) = \mathbb{E}[F|_{(\underline{y}, w)}]$ maps each value w to the conditional expectation of F given that the realized state lies in (\underline{y}, w) .*

Simply speaking, the condition requires us to consider the (unique) point $z = C^{-1}(\underline{z}) \in (\underline{y}, \bar{y})$ such that the conditional expectation over the interval (\underline{y}, z) is \underline{z} . For the support $\{\underline{z}, \bar{z}\}$ to be feasible, we must have that the conditional expectation of the remaining interval $[z, \bar{y})$ exceeds \bar{z} .

For the special case where F is a uniform distribution, this characterization takes the following very simple form.

COROLLARY 1. *If $F|_{(\underline{y}, \bar{y})}$ is a uniform distribution, then a support $\{\underline{z}, \bar{z}\}$ is feasible for the interval (\underline{y}, \bar{y}) if and only if $\underline{y} \leq \underline{z} \leq \frac{1}{2}(\underline{y} + \bar{y}) \leq \bar{z} \leq \bar{y}$ and $\bar{z} - \underline{z} \leq \frac{1}{2}(\bar{y} - \underline{y})$.*

3. IMPLEMENTING BI-POOLING DISTRIBUTIONS

In this section, we discuss possible implementations of bi-pooling distributions. Recall our terminology that π implements some $G \in \text{MPC}(F)$ if $F_\pi = G$. For $G \in \text{MPC}(F)$, let $\Pi(G)$ denote the family of signaling policies, π , for which $F_\pi = G$. It is well known that $\Pi(G) \neq \emptyset$ (e.g., [Gentzkow and Kamenica \(2016\)](#), [Kolotilin \(2018\)](#)). In fact, $\Pi(G)$ may not be a singleton, as demonstrated in the following example.

EXAMPLE 3. We denote by δ_c the Dirac measure on c . Consider the distribution $G = \frac{1}{2}\delta_{\frac{1}{3}} + \frac{1}{2}\delta_{\frac{2}{3}}$ that is an MPC of the uniform distribution $F = U([0, 1])$. As was suggested in [Example 1](#), one implementation of G can be obtained by sending the signal $\frac{1}{3}$ in the interval $(\frac{1}{12}, \frac{7}{12})$ and sending the signal $\frac{2}{3}$ in the remaining pair of intervals $[0, \frac{1}{12}] \cup [\frac{7}{12}, 1]$. An alternative implementation would be to send the signal $\frac{1}{3}$ in the two intervals $[0, \frac{5}{12}] \cup [\frac{11}{12}, 1]$ and to send the signal $\frac{2}{3}$ in the intervals $(\frac{5}{12}, \frac{11}{12})$. A third implementation would be to send the signal $\frac{1}{3}$ in the two intervals $(c, \frac{1}{2}] \cap [1 - c, 1]$ and the signal $\frac{2}{3}$ in the remaining pair of intervals $[0, c] \cup (\frac{1}{2}, 1 - c)$ for $c = (\sqrt{6} - \sqrt{5})/(\sqrt{24}) \approx 0.04$. \diamond

The class of policies that induce bi-pooling distributions is central to our work and is formally defined as follows.

DEFINITION 3. A signaling policy π is called a *bi-pooling policy* if there exists a collection of pairwise disjoint intervals $\{(\underline{y}_i, \bar{y}_i)\}_{i \in A}$ such that for every state $\omega \in (\underline{y}_i, \bar{y}_i)$ we have $\text{supp}(\pi(\omega)) \subset \{\underline{z}_i, \bar{z}_i\}$ for some pair of points $\underline{y}_i \leq \underline{z}_i \leq \bar{z}_i \leq \bar{y}_i$ (i.e., the policy bi-pools the interval $(\underline{y}_i, \bar{y}_i)$ into the pair of points $\{\underline{z}_i, \bar{z}_i\}$). For every $\omega \notin \bigcup_{i \in A} (\underline{y}_i, \bar{y}_i)$, the policy sends the signal $\pi(\omega) = \omega$ (i.e., it reveals the state). In the case where $\underline{z}_i = \bar{z}_i$ for all $i \in A$, we refer to π as a *pooling policy*.

It is quite straightforward to verify that the distribution over posterior means induced by a bi-pooling policy is a bi-pooling distribution and the one induced by a pooling policy is a pooling distribution.

Monotonic implementation As argued in the Introduction, a desired property of signaling policies is monotonicity. Simply speaking, monotonicity requires that higher states be mapped to higher posterior means. We now turn to show that bi-pooling distributions can be implemented by monotonic policies.

DEFINITION 4. A (possibly mixed) signaling policy, $\pi : [0, 1] \rightarrow \Delta([0, 1])$, is *monotonic* if $\pi(x)$ first-order stochastically dominates $\pi(y)$ for every $x \geq y$.

We revisit Example 3 and demonstrate a mixed and monotonic implementation. Consider a signaling policy that, over the interval $[0, \frac{2}{3}]$, sends a random signal whose realization is $\frac{1}{3}$ with probability $\frac{3}{4}$ and $\frac{2}{3}$ with probability $\frac{1}{4}$, while on the interval $(\frac{2}{3}, 1]$ it sends the signal $\frac{2}{3}$. It is easy to verify that this signaling policy implements the distribution G and, in addition, is monotonic. This signaling policy is a bi-pooling policy with a single bi-pooling interval, namely, the entire interval $[0, 1]$. Once again, the occurrence of a monotonic signaling policy in Example 3 is not a coincidence.

LEMMA 2. *Every bi-pooling distribution can be implemented by a monotonic bi-pooling signaling policy.*

The next proposition follows immediately from Lemma 2 and Theorem 1.

PROPOSITION 2. *Every persuasion problem admits an optimal (mixed) monotonic signaling policy.*

When we restrict attention to *pure* signaling policies, namely, policies of the form $\pi : [0, 1] \rightarrow [0, 1]$, then monotonicity is equivalent to pooling in the following sense.

LEMMA 3. *A persuasion problem (F, u) admits an optimal pure monotonic signaling policy if and only if it admits an optimal pooling policy.*

It turns out that not all persuasion problems admit a pure monotonic solution. In fact, Dworzak and Martini (2019) provide a sufficient condition on the utility function, u , for the existence of a pure monotonic solution. They show that this condition is also necessary if we require a monotonic solution *for all* priors. However, in a typical Bayesian persuasion problem a prior is exogenously given, whereas, in the present context, our analyses show that an optimal monotonic signaling policy always exists but not necessarily in the class of pure monotonic signaling policies.

Pure implementation Another natural way to implement a bi-pooling distribution is to consider a pure signaling policy that has the following *double-interval nested structure*: for each bi-pooling interval $(\underline{y}_i, \bar{y}_i)$, we can find a subinterval $(\underline{w}_i, \bar{w}_i) \subset (\underline{y}_i, \bar{y}_i)$ such that π is constant over the interval $(\underline{w}_i, \bar{w}_i)$ as well as over its complement $(\underline{y}_i, \bar{y}_i) \setminus (\underline{w}_i, \bar{w}_i)$.

LEMMA 4. *Any bi-pooling distribution can be implemented by a bi-pooling signaling policy with a double-interval nested structure.*

As a corollary, we get the optimality of double-interval nested signaling policies.¹¹

COROLLARY 2. *Every persuasion problem (F, u) has an optimal bi-pooling policy that has a double-interval nested structure.*

Corollary 2 extends a key technical result from Candogan (2020) that proves the same result whenever u is piecewise constant.

4. SOLVING PERSUASION PROBLEMS

Dworczak and Martini Dworzak and Martini (2019) propose an elegant approach to dealing with persuasion problems in the same continuous-space setting as ours. A basic building block in their approach is the notion of *price function*.¹²

DEFINITION 5. A function $p : [0, 1] \rightarrow \mathbb{R}$ is called a *price function* of G in the problem (F, u) if the following conditions are satisfied:

- $p(x) \geq u(x)$ for every $x \in [0, 1]$.
- p is convex.
- $\text{supp}(G) \subset \{x : p(x) = u(x)\}$.
- $\mathbb{E}_{x \sim G}[p(x)] = \mathbb{E}_{x \sim F}[p(x)]$.

The two main theorems of Dworzak and Martini (2019) make a strong connection between price functions and optimal persuasion policies. For completeness, we now restate these results.

THEOREM 3 (Dworczak and Martini (2019)). *Let (F, u) be a persuasion problem and let $G \in \text{MPC}(F)$. If there exists a price function, p , of G then G is optimal. Furthermore, if u is regular, then there exists an optimal distribution G for which such a price function exists.*¹³

¹¹A recent paper by Candogan and Strack (2021) generalizes this result as well as a result of Guo and Shmaya (2019). They show that each state lies in an interval in which at most $n + 2$ messages are used, where n is the number of types of the receiver.

¹²This function corresponds to prices in an economy as defined in Dworzak and Martini (2019), and hence the term “price function.”

¹³A utility function u on the unit interval is *regular* if it is bounded, upper semicontinuous, and for some $\epsilon > 0$ it is Lipschitz continuous on $[0, \epsilon]$ and $[1 - \epsilon, 1]$. In their original proof, Dworzak and Martini (2019) make use of slightly stronger conditions. More recently, Deniz and Kováč (2019) provide a proof that makes use of this weaker notion of regularity.

Theorem 3 suggests that understanding the structure of price functions is instrumental for the construction of optimal persuasion policies in concrete settings. Proposition 2 of Dworzak and Martini (2019) furthermore argues that if a price function p is not linear then it must coincide with the utility function. Knowing that an optimal signaling policy can take the form of a bi-pooling distribution (Theorem 1) informs us about the structure of the corresponding price function, as the following corollary demonstrates.

COROLLARY 3. *Let (F, u) be a persuasion problem. Let G be an optimal bi-pooling distribution and let p be the corresponding price function of G (assuming it exists). Then p is linear over the pooling and bi-pooling intervals of G .*

Proposition 2 of Dworzak and Martini (2019) argues that if a price function p is linear over an interval $[a, b]$ and this is the largest interval (with respect to inclusion) in which p is linear, then the optimal distribution G maps the interval $[a, b]$ into itself.¹⁴ Furthermore, the price function is then tangential to the utility function in at least one point. In this respect, the following proposition extends Proposition 2 of Dworzak and Martini (2019) and argues that in such intervals the tangent points are either a single point at the expectation or there are two tangent points that form a *feasible pair*; see Definition 2 and Lemma 1.¹⁵

PROPOSITION 3. *Let (F, u) be a persuasion problem such that u is upper semicontinuous and there exists a sequence of points $0 = z_0 < z_1 < \dots < z_k = 1$ such that u is either concave or convex in the interval (z_{i-1}, z_i) for every $i \in [k]$. The price function $p : [0, 1] \rightarrow \mathbb{R}$ of an optimal bi-pooling distribution G is a concatenation of functions $p(x) = p_i(x)$ for $x \in [y_{i-1}, y_i]$, where $0 = y_0 < y_1 < \dots < y_m = 1$, $m \leq 2k$, and the functions $p_i : [y_{i-1}, y_i] \rightarrow \mathbb{R}$ satisfy*

- either $p_i(x) = u(x)$ (full revelation¹⁶)
- or $p_i(x) \geq u(x)$ is linear and tangential to u at the point $\mathbb{E}[F|_{(y_{i-1}, y_i)}]$ (pooling interval¹⁷)
- or $p_i(x) \geq u(x)$ is linear and tangential to u at two points $\underline{z}_i, \bar{z}_i$, where $\{\underline{z}_i, \bar{z}_i\}$ is feasible for the interval $[y_{i-1}, y_i]$ (bi-pooling interval).

Inspired by the third part of Proposition 3, we define a line p as a *bi-tangent* of a function u on the interval $[a, b]$ if it is tangential to u at two points. In the following subsections, we leverage our characterization of optimal persuasion policies as bi-pooling distributions for various concrete applications and use bi-tangents as a way to construct the price functions and the optimal policies.

¹⁴That is, the conditional expectation of G over $[a, b]$ equals the conditional expectation of F over $[a, b]$.

¹⁵For the uniform prior case $F = U([0, 1])$, the tangent points $\underline{z}_i, \bar{z}_i$ should satisfy $y_{i-1} \leq \underline{z} \leq \frac{1}{2}(y_{i-1} + y_i) \leq \bar{z} \leq y_i$ and $\bar{z}_i - \underline{z}_i \leq \frac{1}{2}(y_i - y_{i-1})$; see Corollary 1.

¹⁶In this case, u must be convex on $[y_{i-1}, y_i]$.

¹⁷For the uniform prior case $F = U([0, 1])$, the tangent point is $\frac{1}{2}(y_{i-1} + y_i)$.

4.1 *S-shaped utility*

Consider the setting of Kolotilin, Mylovanov, and Zapechelnyuk (2022). A receiver has to decide whether or not to accept a project whose quality is known to the sender. Upon acceptance, both the receiver and the sender enjoy a utility that is linear with respect to the quality. Upon rejection, the receiver's utility, denoted by V , is her private information while the sender's utility is zero. Kolotilin, Mylovanov, and Zapechelnyuk (2022) focus on the case where the distribution over the value of V has a unimodal density function. The unimodal density function induces an S-shaped indirect utility function (a convex interval followed by a concave interval).

One of the main results in Kolotilin, Mylovanov, and Zapechelnyuk (2022) states that an optimal signaling scheme for an S-shaped indirect utility function pools all high-quality signals, those above some threshold, while fully revealing low-quality signals, those below the threshold. We now demonstrate how our results can be used to provide the intuition behind this result.

Proposition 3 informs us that the corresponding price function is composed of three types of intervals: bi-pooling, pooling, and fully revealing. A bi-pooling interval corresponds to a bi-tangent interval of the price function with two *internal* tangential points (Lemma 1). Note that such a bi-tangent lies above the utility and hence touches the concave part of the utility at exactly one point. This implies that the other tangential point lies in the convex part of the utility. This point must be 0 (see Figure 3 (a)) for if it were not 0, then the bi-tangent would lie below the utility. Therefore, no bi-pooling interval exists in the optimal policy.

Moreover, since the tangent in a pooling interval of the price function lies above the utility, the tangent point must be in its concave part. However, since there is only one concave part in the utility function and since the price function itself is convex, no two such intervals can coexist. Hence, there is at most one pooling interval.

Finally, observe that full revelation cannot hold at any point of the concave regime since the price function is convex and it coincides with the utility function at full-revelation points. Therefore, the price function must take the form of a full-revelation




FIGURE 3. (a) The unique bi-tangent of an S-shape function. (b) The price function (i.e., the solution) of an S-shape function.

interval over the low values followed by a unique pooling interval over the high values; see Figure 3(b).

4.2 A technical example

Thus far, we have demonstrated how the machinery we develop allows us to solve the simple case of an S -shaped utility. We will next show further that our machinery is also helpful in solving more complicated problems where the optimal policy may entail some (possibly many) bi-pooling, pooling, and full-revelation regions.

The following example is provided solely as a means of demonstrating the contribution of Proposition 3 in solving for the optimal policy. The idea is to find a price function that corresponds to a bi-pooling distribution. Specifically, we consider all possible locations of bi-pooling intervals. For each such “suspicious” location, we attempt to construct a price function. A bi-pooling interval induces a bi-tangent interval of the price function. As the number of bi-tangents is typically small, this observation provides a practical way to start the search for the price function.¹⁸

EXAMPLE 4. Let $F = U([0, 1])$ and u be the utility function that is given in Figure 4. In this case, u has $k = 7$ interlaced concave and convex intervals and in the first interval u is convex.

We begin by drawing all the bi-tangents. In this case, we have three such lines; see Figure 5.

Note that the bi-tangent L_2 violates the feasibility of the tangent points $\{\underline{z}, \bar{z}\}$ (see the third item in Proposition 3), even in the case where the bi-pooling interval has the maximal possible length, that is, in the case of $[\underline{y}, \bar{y}] = [0, x]$, where x is the (nontangential) intersection of L_2 and u . Thus, a bi-pooling interval in L_2 is impossible. Similarly, a bi-pooling interval over the L_3 line is impossible.




FIGURE 4. The utility of the sender as a function of the expected value of ω for the receiver.

¹⁸By contrast, the number of tangents associated with pooling intervals is huge.




FIGURE 5. The supporting bi-tangents of u . A potentially feasible bi-tangent is drawn by a dotted line. Infeasible bi-tangents are drawn by solid lines.

If we maintain the assumption that a bi-pooling interval indeed exists, then the corresponding tangent is L_1 and the corresponding tangent points lie in intervals 2 and 4. This leaves us with two possibilities. Either the bi-pooling interval is followed by a single pooling interval or it is followed by a pooling interval and later by a fully revealing interval. Searching for a price function that corresponds to the former case leads to a dead end and so we can conclude that it is the latter case that finally produces the price function as depicted in Figure 6.

As can be seen in the figure, the optimal policy in this example applies bi-pooling in the interval $(0, y)$ with the pair of points (z_1, z_2) in the support. It applies pooling in the interval (y, y') (in particular $z_3 = \frac{1}{2}(y' + y)$). Finally, it is fully revealing in the interval $[y', 1]$. \diamond




FIGURE 6. A price function of the utility.

APPENDIX A: PROOFS FOR SECTION 2

LEMMA 1. A support $\{\underline{z}, \bar{z}\}$ is feasible for the interval (y, \bar{y}) if and only if $\underline{y} \leq \underline{z} \leq \mathbb{E}[F|_{(y, \bar{y})}] \leq \bar{z} \leq \bar{y}$ and $\mathbb{E}[F|_{(C^{-1}(\underline{z}), \bar{y})}] \geq \bar{z}$, where $C(w) = \mathbb{E}[F|_{(y, w)}]$ maps each value w to the conditional expectation of F given that the realized state lies in (y, w) .

PROOF OF LEMMA 1. Let $(\underline{y}, \underline{z}, \bar{z}, \bar{y})$ be such that $\mathbb{E}_{x \sim F|_{(C^{-1}(\underline{z}), \bar{y})}}[x] \geq \bar{z}$. In the proofs of Lemma 2 and Lemma 4 below, we will show how to implement the corresponding binary-support distribution as a MPC of $F|_{(y, \bar{y})}$.

To prove the other direction, assume without loss of generality that $\underline{y} = 0$, $\bar{y} = 1$ and, therefore, $F|_{(y, \bar{y})} = F$. We ask first what is the largest probability that can be assigned to the atom \underline{z} by a binary support $G \in \text{MPC}(F)$. We claim this probability is at most $F((0, C^{-1}(\underline{z})))$. Let $G \in \text{MPC}(F)$ be the distribution that assigns the maximal weight and let $\pi : [0, 1] \rightarrow \Delta(\{\underline{s}, \bar{s}\})$ be a signaling policy that implements G such that $\mathbb{E}_{x \sim F}[x|\pi(x) = \underline{s}] = \underline{z}$. We contend that $\pi(x)(\underline{s}) = 1$ for F -almost every $x \in (0, C^{-1}(\underline{z}))$. To see this, note that one of the two possibilities must hold. Either $\int_{C^{-1}(\underline{z})}^1 \pi(x)(\underline{s}) dF(x) = 0$, or $\int_{C^{-1}(\underline{z})}^1 \pi(x)(\underline{s}) dF(x) > 0$.

In the former case, we have that the overall weight $\int_0^1 \pi(x)(\underline{s}) dF(x)$ that is assigned to \underline{s} is smaller than $F((0, C^{-1}(\underline{z})))$. Let $\pi' : [0, 1] \rightarrow \Delta(\{\underline{s}, \bar{s}\})$ be a signaling policy for which $\pi'(x)(\underline{s}) = 1$ for any $x \in (0, C^{-1}(\underline{z}))$ and $\pi'(x)(\bar{s}) = 1$ for any $x \in [C^{-1}(\underline{z}), 1)$. Note that $\mathbb{E}_{x \sim F, \pi'}[x|\pi(x) = \underline{s}] = \mathbb{E}_{x \sim F|_{(0, C^{-1}(\underline{z}))}}[x] = C(C^{-1}(\underline{z})) = \underline{z}$. Therefore, π' implements a binary distribution that assigns a probability of $F((0, C^{-1}(\underline{z})))$ to \underline{z} . This stands in contradiction to the fact that G assigns a maximal weight to \underline{z} .

Consider the latter case where $\int_{C^{-1}(\underline{z})}^1 \pi(x)(\underline{s}) dF(x) > 0$ and $\pi(x)(\underline{s}) < 1$ for an F positive measure of $x \in (0, C^{-1}(\underline{z}))$. We generate a new signaling policy π' from the signaling policy π as follows. We increase the weight $\pi(x)(\underline{s})$ to 1 for some positive measure of $x \in (0, C^{-1}(\underline{z}))$ for which $\pi(x)(\underline{s}) < 1$ and reduce the weight $\pi(x)(\underline{s})$ to zero for some positive measure of $x \in (C^{-1}(\underline{z}), 1)$ for which $\pi(x)(\underline{s}) > 0$. This can be done in such a way that $\int_0^1 \pi(x)(\underline{s}) dF(x) = \int_0^1 \pi'(x)(\underline{s}) dF(x)$. Since $\mathbb{E}_{X \sim F, \pi}[X|\pi(X) = \underline{s}] = \underline{z}$, it follows that $\mathbb{E}_{X \sim F, \pi'}[X|\underline{s}] < \underline{z}$. We can now generate a new signaling policy π'' from π' by increasing the weight $\pi'(x)(\underline{s})$ for some positive measure of $x \in (0, 1)$ for which $\pi(x)(\underline{s}) < 1$. This can be done in such a way that $\mathbb{E}_{x \sim F, \pi''}[x|\underline{s}] = \underline{z}$. Thus, overall, we obtain a signaling policy π'' that generates a distribution that assigns a higher weight to \underline{z} relative to G and, therefore, again reach a contradiction.

Thus, we must have that $\pi(x)(\underline{s}) = 1$ for F -almost every $x \in (0, C^{-1}(\underline{z}))$. We note that π must also satisfy $\pi(x)(\underline{s}) = 0$ for F -almost every $x \in [C^{-1}(\underline{z}), 1)$ as otherwise by the definition of $C^{-1}(\underline{z})$ we would have $\mathbb{E}_{X \sim F, \pi}[X|\underline{s}] > \underline{z}$. Therefore, overall we have shown that the distribution G assigns a weight of $F((0, C^{-1}(\underline{z})))$ to \underline{z} . We note that G is also the persuasion policy in $\text{MPC}(F)$ with the binary support $\{\underline{z}, \bar{z}\}$ that maximizes \bar{z} . Thus, by the above, $\bar{z} = \mathbb{E}_{x \sim F}[x|x \in (C^{-1}(\underline{z}), 1)]$. \square

We next turn to the proof of Proposition 1. Henceforth, for any distribution $G \in \Delta([0, 1])$ we abuse notation and use $G(x) = G([0, x])$ as the *cumulative distribution function* (CDF) of G . Let $r_G(x) = \int_0^x G(y) dy$. Recall that $G \in \text{MPC}(F)$ if and only if

$r_G(x) \leq r_F(x)$ for every $x \in [0, 1]$ and $r_G(1) = r_F(1)$. We start with the following auxiliary lemma.

LEMMA 5. *Let G be an extreme point of $\text{MPC}(F)$ and let $(\underline{z}, \bar{z}) \subseteq (0, 1)$ be a nonempty open interval. If G satisfies $r_G(\underline{z}) = r_F(\underline{z})$, $r_G(\bar{z}) = r_F(\bar{z})$, and $r_G(x) < r_F(x)$ for every $x \in (\underline{z}, \bar{z})$, then G has a support of at most 2 points over (\underline{z}, \bar{z}) .*

PROOF OF LEMMA 5. For simplicity, we assume that without loss of generality $(\underline{z}, \bar{z}) = (0, 1)$. Suppose by way of contradiction that $G \in \text{MPC}(F)$ contains at least three points in its support. We note that since F is a nonatomic distribution, by assumption, we must have that neither 0 nor 1 is an atom of G . Therefore, there exists a closed interval $[\underline{y}, \bar{y}] \subset (0, 1)$ such that the support of G contains at least three points in $[\underline{y}, \bar{y}]$. Hence, we can find three disjoint intervals $(\underline{y}_i, \bar{y}_i) \subset (0, 1)$ for $i = 1, 2, 3$ such that the following four properties hold: $\bar{y}_i < \underline{y}_{i+1}$ for $i = 1, 2$, $0 < \underline{y}_1$, $\bar{y}_3 < 1$, and $G((\underline{y}_i, \bar{y}_i)) = \alpha_i > 0$ for $i = 1, 2, 3$. Let $G_i = G|_{(\underline{y}_i, \bar{y}_i)}$ and let m_i be the expectation of G_i . Let $\beta \in (0, 1)$ be such that $\beta m_1 + (1 - \beta)m_3 = m_2$. For $\epsilon \in \mathbb{R}$, let H_ϵ be defined as follows:

$$H_\epsilon = G + \epsilon(\beta G_1 + (1 - \beta)G_3 - G_2).$$

The condition $\epsilon \leq \alpha_2$ guarantees that $H_\epsilon((\underline{y}_2, \bar{y}_2)) \geq 0$. Similarly, $\epsilon \geq -\frac{\alpha_1}{\beta}$ and $\epsilon \geq -\frac{\alpha_3}{1-\beta}$ guarantee that $H_\epsilon((\underline{y}_1, \bar{y}_1)) \geq 0$ and $H_\epsilon((\underline{y}_3, \bar{y}_3)) \geq 0$, respectively.

Therefore, H_ϵ is a well-defined probability measure for $\epsilon \leq \alpha_2$ and $\epsilon \geq \max\{-\frac{\alpha_1}{\beta}, -\frac{\alpha_3}{1-\beta}\}$.

We further note that by the definition of β ,

$$\begin{aligned} \int_0^1 (\beta G_1 + (1 - \beta)G_3 - G_2)(x) dx &= \int_{\underline{y}}^{\bar{y}} (\beta G_1 + (1 - \beta)G_3 - G_2)(x) dx \\ &= \beta(1 - m_1) + (1 - \beta)(1 - m_3) - (1 - m_2) = 0. \end{aligned}$$

Therefore, for every $x \notin (\underline{y}, \bar{y})$

$$\int_0^x H_\epsilon(x) dx = r_{H_\epsilon}(x) = r_G(x). \quad (2)$$

By assumption, $r_G(x) < r_F(x)$ for every $x \in (0, 1)$. Therefore, since $r_F(x) - r_G(x)$ is a continuous function, and since $[\underline{y}, \bar{y}] \subseteq (0, 1)$, we have that $r_F(x) - r_G(x) \geq \delta$ for every $x \in [\underline{y}, \bar{y}]$ for some $\delta > 0$. In addition, we have by the definition of H_ϵ that $|H_\epsilon(x) - G(x)| \leq \epsilon$ for every $x \in [0, 1]$. Therefore, $|r_{H_\epsilon}(x) - r_G(x)| < \epsilon$ for every $x \in [0, 1]$. Hence, there exists a small enough $\theta > 0$ such that for every $\epsilon \in [-\theta, \theta]$ the measure H_ϵ is well-defined and for every $x \in [\underline{y}, \bar{y}]$ it holds that $r_F(x) - r_{H_\epsilon}(x) \geq \frac{\delta}{2}$. Equation (2) implies that $r_F(x) > r_{H_\epsilon}(x)$ for every $x \in (0, 1)$.

Therefore, $H_\epsilon \in \text{MPC}(F)$ for every $\epsilon \in [-\theta, \theta]$. Let $H_1 = H_{-\theta}$ and $H_2 = H_\theta$. Note that $H_1 \neq H_2$ and that $G = \frac{H_1 + H_2}{2}$. This implies that G is not an extreme point of $\text{MPC}(F)$. \square

PROPOSITION 1. *The set of extreme points of $\text{MPC}(F)$ is precisely the set of bi-pooling distributions.*

PROOF OF PROPOSITION 1. We begin by showing that any bi-pooling distribution is an extreme point of the set $\text{MPC}(F)$. Let G be such a bi-pooling distribution with the corresponding countable collection of pooling and bi-pooling intervals $\{(y_i, \bar{y}_i)\}_{i \in A}$. Recall that $G(x) = F(x)$ for every $x \in K := [0, 1] \setminus (\bigcup_{i \in A} (y_i, \bar{y}_i))$.

By definition, the probabilities assigned to each interval of the form (y_i, \bar{y}_i) by F and by G are equal and so $r_G(\bar{y}_i) - r_G(y_i) = r_F(\bar{y}_i) - r_G(y_i)$ for any $i \in A$. In addition, $r_G(x) = r_F(x)$ for every $x \in K$. Now suppose by way of contradiction that G is not an extreme point of $\text{MPC}(F)$ and so there exist $G_1, G_2 \in \text{MPC}(F)$ such that $\frac{G_1 + G_2}{2} = G$. Since $G_j \in \text{MPC}(F)$, we must have that $r_{G_j}(x) \leq r_F(x)$ for $j = 1, 2$ and every $x \in [0, 1]$. Since $\frac{r_{G_1} + r_{G_2}}{2} = r_G$, it must hold that $r_{G_j}(x) = r_F(x)$ for $j = 1, 2$ and every $x \in K$. Since $G_j \neq G$, there must be some $i \in A$ and $x \in (y_i, \bar{y}_i)$ such that $G_j(x) \neq G(x)$ for $j = 1, 2$.

For $j = 1, 2$ denote $\tilde{G}_j = G_j|_{(y_i, \bar{y}_i)}$, and similarly $\tilde{G} = G|_{(y_i, \bar{y}_i)}$. Since $r_{G_j}(\bar{y}_i) = r_G(\bar{y}_i)$ and $r_{G_j}(y_i) = r_G(y_i)$, we must have that for $j = 1, 2$ the distributions \tilde{G}_j and \tilde{G} both have the same expectation and that $\frac{\tilde{G}_1 + \tilde{G}_2}{2} = \tilde{G}$. The support of \tilde{G} comprises at most two points, and both \tilde{G}_1, \tilde{G}_2 must have the same (two-point) support. However, then we must have $\tilde{G}_1 = \tilde{G}_2 = \tilde{G}$ as otherwise they could not have the same expected value, a contradiction.

Conversely, let G be an extreme point of $\text{MPC}(F)$. We let K be the set of points x such that $r_G(x) = r_F(x)$. K is a closed set that contains 0 and 1. The complement $K^c \cap [0, 1]$ is an open set and, therefore, can be written as a countable union of pairwise disjoint open intervals $\{(y_i, \bar{y}_i)\}_{i \in A}$. Hence, we must have that $y_i, \bar{y}_i \in K$ for every $i \in A$. Therefore, for every $i \in A$, we have $r_G(\bar{y}_i) = r_F(\bar{y}_i)$, $r_G(y_i) = r_F(y_i)$, and $r_G(x) < r_F(x)$ for every $x \in (y_i, \bar{y}_i)$.

In order to show that G is a bi-pooling distribution, it is sufficient to show that for any $i \in A$ there are at most 2 points in the support of $G_i = G|_{(y_i, \bar{y}_i)}$. Let $F_i = F|_{(y_i, \bar{y}_i)}$ and note that $G_i \in \text{MPC}(F_i)$. As G is an extreme point in $\text{MPC}(F)$, the distribution G_i must be an extreme point in $\text{MPC}(F_i)$ and so by Lemma 5 must have at most two points in its support. This completes the proof of the proposition. \square

THEOREM 2. *For every bi-pooling distribution $G \in \text{MPC}(F)$, there exists a continuous utility function u for which G is the unique optimal solution.*

PROOF OF THEOREM 2. Consider a bi-pooling distribution G , as characterized in Theorem 1, with $\{(y_i, \bar{y}_i)\}_{i \in A}$ as the sequence of intervals. Let $K = [0, 1] \setminus (\bigcup_{i \in A} (y_i, \bar{y}_i))$. Let $p : [0, 1] \rightarrow \mathbb{R}$ be a convex function such that p is strictly convex on K and it is linear on (y_i, \bar{y}_i) for every i . This function clearly exists and can be constructed by taking a strictly convex function $q : [0, 1] \rightarrow \mathbb{R}$ and changing it in (y_i, \bar{y}_i) by taking the linear interpolation between $q(y_i)$ and $q(\bar{y}_i)$. By construction, $F(y_i, \bar{y}_i) = G(y_i, \bar{y}_i)$, $\mathbb{E}_{X \sim F|_{(y_i, \bar{y}_i)}}[p(X)] = \mathbb{E}_{X \sim G|_{(y_i, \bar{y}_i)}}[p(X)]$ for every $i \in A$. In addition, $F(K) = G(K)$ and $\mathbb{E}_{X \sim F|_K}[p(X)] = \mathbb{E}_{X \sim G|_K}[p(X)]$. Therefore, $\mathbb{E}_{X \sim F}[p(X)] = \mathbb{E}_{X \sim G}[p(X)]$.

For every $i \in A$, let $\{w_i, \bar{w}_i\}$ be the support of G in (y_i, \bar{y}_i) .

Let u be any continuous function satisfying (1) $u(x) = p(x)$ for every $x \in K$ or $x \in \bigcup_{i=1}^n \{w_i, \bar{w}_i\}$; and (2) $u(x) < p(x)$ for every other x . Note that, by construction,

$\text{supp}(G) = \{x : u(x) = p(x)\}$. Theorem 1 in Dworczak and Martini (2019) implies that p is a price function for G with respect to the persuasion problem (F, u) and, therefore, G is an optimal distribution for (F, u) . We now claim that G is the unique optimal policy.

Assume by way of contradiction that H is another MPC of F such that $\mathbb{E}_{X \sim H}[u(X)] = \mathbb{E}_{X \sim G}[u(X)]$. By Bauer's maximum principle, we may assume, without loss of generality, that H is an extreme point of $\text{MPC}(F)$ and so is a bi-pooling distribution. Let $\{(\underline{z}_i, \bar{z}_i)\}_{i \in M}$ be the corresponding sequence of pooling and bi-pooling intervals. Since $H \in \text{MPC}(F)$ and p is a convex function, we have $\mathbb{E}_{X \sim H}[p(X)] \leq \mathbb{E}_{X \sim F}[p(X)] = \mathbb{E}_{X \sim G}[p(X)]$. Conversely, since $u \leq p$ we have

$$\mathbb{E}_{X \sim G}[p(X)] = \mathbb{E}_{X \sim G}[u(X)] = \mathbb{E}_{X \sim H}[u(X)] \leq \mathbb{E}_{X \sim H}[p(X)],$$

where the first equality follows since p is a price function for G . Therefore, $\mathbb{E}_{X \sim H}[p(X)] = \mathbb{E}_{X \sim G}[p(X)] = \mathbb{E}_{X \sim F}[p(X)] = \mathbb{E}_{X \sim H}[u(X)]$. So, p is also a price function for H .

Let $B = [0, 1] \setminus (\bigcup_{i \in M} (\underline{z}_i, \bar{z}_i))$. Since $H|_{(\underline{z}_i, \bar{z}_i)}$ is an MPC of $F|_{(\underline{z}_i, \bar{z}_i)}$ for every $i \in M$ and since p is a convex function, we conclude that $\mathbb{E}_{X \sim H|_{(\underline{z}_i, \bar{z}_i)}}[p(X)] = \mathbb{E}_{X \sim F|_{(\underline{z}_i, \bar{z}_i)}}[p(X)]$ for every $i \in M$. This is possible only if p is linear on $(\underline{z}_i, \bar{z}_i)$ for every $i \in M$. Recall that by construction p is linear only on the intervals $(\underline{y}_i, \bar{y}_i)$ and so for every $i \in M$ we must have that $(\underline{z}_i, \bar{z}_i) \subseteq (\underline{y}_j, \bar{y}_j)$ for some $j \in A$. As p is also a price function for H we can use the above argument, by changing the roles of G and H , to conclude that any interval $(\underline{y}_i, \bar{y}_i)$ is contained in some interval $(\underline{z}_i, \bar{z}_i)$. Therefore, G and H are two bi-pooling policies with the same set of pooling and bi-pooling intervals. To show they are equal, it is enough to show that they agree on each such interval. Since $\mathbb{E}_{X \sim H}[u(X)] = \mathbb{E}_{X \sim G}[p(X)]$, we must also have that $\mathbb{E}_{X \sim H|_{(\underline{z}_i, \bar{z}_i)}}[u(X)] = \mathbb{E}_{X \sim G|_{(\underline{z}_i, \bar{z}_i)}}[p(X)]$ for every $i \in M$. This is possible only if the support of H in $(\underline{z}_i, \bar{z}_i)$ is contained in $\{\underline{w}_j, \bar{w}_j\}$. Thus, we must have that $(\underline{z}_i, \bar{z}_i) = (\underline{y}_j, \bar{y}_j)$ and $H = F$ on $(\underline{z}_i, \bar{z}_i)$. Therefore, $H = G$. \square

APPENDIX B: PROOFS FOR SECTION 3 AND SECTION 4

LEMMA 2. *Every bi-pooling distribution can be implemented by a monotonic bi-pooling signaling policy.*

PROOF OF LEMMA 2. Let G be a bi-pooling distribution. Any pooling interval unilaterally defines the signaling policy and so does the fully revealing part of G . Therefore, it is sufficient to show that we can implement $G|_{(\underline{y}_i, \bar{y}_i)}$ monotonically on any bi-pooling interval $(\underline{y}_i, \bar{y}_i)$.

Assume without loss of generality that $(\underline{y}, \bar{y}) = (0, 1)$, and let \underline{z} and \bar{z} be the atoms of G in $(0, 1)$.

We now show that there exists $\pi : (0, 1) \rightarrow \Delta(\{\underline{s}, \bar{s}\})$ such that (i) $\mathbb{E}_{\pi, X \sim F}[X|\underline{s}] = \underline{z}$, $\mathbb{E}_{\pi, X \sim F}[X|\bar{s}] = \bar{z}$ and (ii) for every $x, y \in (0, 1)$ such that $x \geq y$, it holds that $\pi(x)(\bar{s}) \geq \pi(y)(\bar{s})$.

To see this, for every $\beta \in [0, 1]$ define $\pi_\beta : (0, 1) \rightarrow \Delta(\{\underline{s}, \bar{s}\})$ as follows:

$$\pi_\beta(x) = \begin{cases} \beta \delta_{\underline{s}} + (1 - \beta) \delta_{\bar{s}} & \text{if } x \in [\underline{y}, C^{-1}(\underline{z})] \\ \delta_{\bar{s}} & \text{if } x \in (C^{-1}(\underline{z}), \bar{y}), \end{cases}$$

where we recall that $C(w) = \mathbb{E}_{x \sim F|_{(y,w)}}[x]$. We claim that π_β implements the distribution $\alpha\delta_{\underline{z}} + (1 - \alpha)\delta_{\bar{z}}$ for some $\beta \in (0, 1]$. To see this note that for $\beta > 0$, the signal \underline{s} is chosen with constant probability on the interval $(0, C^{-1}(\underline{z})]$ and with zero probability on $(C^{-1}(\underline{z}), 1)$. Therefore, the conditional expectation of the realized state given \underline{s} is

$$\mathbb{E}_{X \sim F}[X|X \in (0, C^{-1}(\underline{z}))] = C(C^{-1}(\underline{z})) = \underline{z}.$$

In addition, by construction, the conditional expectation of the realized state X given \bar{s} as a function of β is

$$\begin{aligned} \mathbb{E}_\pi[X|\bar{s}] &= \frac{F(C^{-1}(\underline{z}))(1 - \beta)}{F(C^{-1}(\underline{z}))(1 - \beta) + (1 - F(C^{-1}(\underline{z})))} \underline{z} \\ &+ \frac{(1 - F(C^{-1}(\underline{z})))}{F(C^{-1}(\underline{z}))(1 - \beta) + (1 - F(C^{-1}(\underline{z})))} \mathbb{E}_{X \sim F}[X|X \in (C^{-1}(\underline{z}), 1)]. \end{aligned} \quad (3)$$

That is, with conditional probability $\frac{F(C^{-1}(\underline{z}))(1 - \beta)}{F(C^{-1}(\underline{z}))(1 - \beta) + (1 - F(C^{-1}(\underline{z})))}$ the realization of \bar{s} is chosen from the interval $(0, C^{-1}(\underline{z})]$ and, therefore, has a conditional expectation \underline{z} , and with the complement probability the realization of \bar{s} is chosen from $(C^{-1}(\underline{z}), 1)$ and has a conditional expectation $\mathbb{E}_{X \sim F}[X|X \in (C^{-1}(\underline{z}), 1)]$.

This shows in particular that π_β is monotonic for every β . Since $\{\underline{z}, \bar{z}\}$ is feasible, the latter conditional expectation is at least \bar{z} . Note that $\mathbb{E}_\pi[X|\bar{s}]$ changes continuously in β . In addition, $\mathbb{E}_{\pi_0, X \sim F}[X|\bar{s}] = \mathbb{E}_{X \sim F}[X]$ for $\beta = 0$ and $\mathbb{E}_{\pi_1, X \sim F}[X|\bar{s}] = \mathbb{E}_{X \sim F}[X|X \in (C^{-1}(\underline{z}), 1)] \geq \underline{z}$. Therefore, by continuity $\mathbb{E}_{\pi_\beta, X \sim F}[X|\bar{s}] = \bar{z}$ for some value $\beta > 0$. This value β defines a lottery with expectation $m = \mathbb{E}_{X \sim F}[X]$. Hence, the lottery is $\alpha\delta_{\underline{z}} + (1 - \alpha)\delta_{\bar{z}}$, as desired. \square

LEMMA 3. *A persuasion problem (F, u) admits an optimal pure monotonic signaling policy if and only if it admits an optimal pooling policy.*

PROOF OF LEMMA 3. Let π be a pure monotonic signaling policy and let $A = \{y | \pi^{-1}(y) \text{ contains an open set}\}$. If $A = \emptyset$, then since π is increasing we must have that π is one to one, and hence fully revealing. This in particular means that it is a pooling policy. Assume $A \neq \emptyset$. In this case, the monotonicity of π guarantees that the collection of intervals $\{\pi^{-1}(y)\}_{y \in A}$ is pairwise disjoint. These intervals form pooling intervals for the corresponding policy. For $y \notin A$, we must have that $\pi^{-1}(y) = y$ is a singleton. Therefore, π fully reveals states outside $\bigcup_{y \in A} \pi^{-1}(y)$. The converse follows immediately from the definition of a pooling policy. \square

LEMMA 4. *Any bi-pooling distribution can be implemented by a bi-pooling signaling policy with a double-interval nested structure.*

PROOF OF LEMMA 4. It is sufficient to show that if $\{\underline{z}, \bar{z}\}$ is feasible for $\underline{z} \neq \bar{z}$, then we can find an interval (\underline{y}, \bar{y}) such that $\mathbb{E}_{X \sim F|_{(\underline{y}, \bar{y})}}[X] = \underline{z}$ and $\mathbb{E}_{X \sim F|_{[1, \underline{y}] \cup [\bar{y}, 1]}}[X] = \bar{z}$. Note that

since $\{\underline{z}, \bar{z}\}$ is feasible Lemma 1 implies that

$$\mathbb{E}_{X \sim F|_{[C^{-1}(\underline{z}), 1]}} \geq \bar{z}.$$

If the above inequality holds with equality, we can let $(y, \bar{y}) = (0, C^{-1}(\underline{z}))$.

Assume otherwise that $\bar{z} < \mathbb{E}_{X \sim F|_{[C^{-1}(\underline{z}), 1]}}$. Define a function $h : [0, \underline{z}] \rightarrow [0, 1]$ such that $h(x)$ is the number $y > \underline{z}$ that satisfies

$$\mathbb{E}_{X \sim F|_{[x, y]}}[X] = \underline{z}.$$

In words, for every $x \in [0, \underline{z}]$ the value $h(x)$ is determined such that the conditional expectation of F given that the realized value lies in $[x, h(x)]$ is \underline{z} . Since F is non-atomic with a full support on $[0, 1]$, such a function h exists uniquely. Furthermore, h is clearly continuous and strictly increasing. Note that, by definition, $h(0) = C^{-1}(\underline{z})$. Since $\bar{z} < \mathbb{E}_{X \sim F|_{[C^{-1}(\underline{z}), 1]}$ and since

$$\begin{aligned} & F([0, C^{-1}(\underline{z})])\mathbb{E}_{X \sim F|_{[0, C^{-1}(\underline{z})]}[X] + F((C^{-1}(\underline{z}), 1])\mathbb{E}_{X \sim F|_{(C^{-1}(\underline{z}), 1]}[X] \\ &= F((0, C^{-1}(\underline{z})))\underline{z} + F([C^{-1}(\underline{z}), 1])\mathbb{E}_{X \sim F|_{[C^{-1}(\underline{z}), 1]} = \mathbb{E}_{X \sim F}[X] := m, \end{aligned}$$

we must have that $F([0, C^{-1}(\underline{z})]) > \frac{1-m}{\bar{z}-\underline{z}}$. In addition, it clearly holds by continuity considerations that as x approaches \underline{z} , $h(x)$ approaches \underline{z} , and $F((h(x), x))$ approaches zero. Therefore, by the mean value theorem, there must exist x' such that $F((x', h(x'))) = \frac{1-m}{\bar{z}-\underline{z}}$. We now let $(y, \bar{y}) = (x', h(x'))$.

Note that, by construction, $\mathbb{E}_{X \sim F|_{(y, \bar{y})}}[X] = \underline{z}$. Since $F((y, \bar{y})) = \frac{1-m}{\bar{z}-\underline{z}}$ and since the overall conditional expectation of the realized state is m , we must have that $\mathbb{E}_{X \sim F|_{[1, y] \cup [\bar{y}, 1]}}[X] = \bar{z}$. This completes the proof of the lemma. \square

PROPOSITION 3. *Let (F, u) be a persuasion problem such that u is upper semicontinuous and there exists a sequence of points $0 = z_0 < z_1 < \dots < z_k = 1$ such that u is either concave or convex in the interval (z_{i-1}, z_i) for every $i \in [k]$. The price function $p : [0, 1] \rightarrow \mathbb{R}$ of an optimal bi-pooling distribution G is a concatenation of functions $p(x) = p_i(x)$ for $x \in [y_{i-1}, y_i]$, where $0 = y_0 < y_1 < \dots < y_m = 1$, $m \leq 2k$, and the functions $p_i : [y_{i-1}, y_i] \rightarrow \mathbb{R}$ satisfy*

- either¹⁹ $p_i(x) = u(x)$,
- or $p_i(x) \geq u(x)$ is linear and tangential to u at the point $\mathbb{E}[F|_{(y_{i-1}, y_i)}]$,
- or $p_i(x) \geq u(x)$ is linear and tangential to u at two points $\underline{z}_i, \bar{z}_i$, where $\{\underline{z}_i, \bar{z}_i\}$ is feasible for the interval $[y_{i-1}, y_i]$; see Definition 2 and Lemma 1.

PROOF OF PROPOSITION 3. We first show that there exists an optimal bi-pooling policy that applies pooling or bi-pooling in at most $2k$ intervals. Let G be an optimal bi-pooling policy such that the set of intervals $\{(y_{\underline{i}}, \bar{y}_{\underline{i}})\}_{\underline{i} \in A}$ is minimal in the following sense. If

¹⁹In this case, u must be convex on $[y_{i-1}, y_i]$.

we take any interval $(\underline{y}_i, \bar{y}_i)$ and apply full revelation instead, then the resulting policy is no longer optimal. In addition, for two disjoint intervals $(\underline{y}_i, \bar{y}_i)$ and $(\underline{y}_{i'}, \bar{y}_{i'})$ such that $\bar{y}_i < \underline{y}_{i'}$ the optimal policy resulting from applying pooling in $(\underline{y}_i, \bar{y}_{i'})$ and keeping revelation outside $(\underline{y}_i, \bar{y}_{i'})$ as in G , leads to a suboptimal policy.

Let p be the price function of G . Let (z_j, z_{j+1}) be an interval in which u is convex. We now show that it does not contain an interval $(\underline{y}_i, \bar{y}_i) \subseteq (z_j, z_{j+1})$ for some $i \in A$. To see this, note that in this case we can remove interval i from A and apply full revelation in $(\underline{y}_i, \bar{y}_i)$. Since u is convex on $(\underline{y}_i, \bar{y}_i)$, this cannot decrease the utility of the sender. This stands in contradiction to the minimality assumption of G .

Next, consider an interval (z_j, z_{j+1}) in which u is concave. We contend that it contains at most one interval. Otherwise, we would have two disjoint intervals $(\underline{y}_i, \bar{y}_i)$ and $(\underline{y}_{i'}, \bar{y}_{i'})$ that are contained in (z_j, z_{j+1}) . Thus, since $p = u$ in the support of G , and since p is convex and u is concave on (z_j, z_{j+1}) , we must have that p and u are linear on the support of G over the intervals $(\underline{y}_i, \bar{y}_i)$ and $(\underline{y}_{i'}, \bar{y}_{i'})$. Thus, if $\bar{y}_i < \underline{y}_{i'}$, x is the smallest point in the support of G in $(\underline{y}_i, \bar{y}_i)$ and x' is the largest point in the support of G in $(\underline{y}_{i'}, \bar{y}_{i'})$, then u and p are linear on $[x, x']$. We could therefore apply pooling in $(\underline{y}_i, \bar{y}_{i'})$. By doing this, we take the probability mass on $[x, x']$ and shift it to a single point $y \in [x, x']$. Since u is linear on $[x, x']$, this would not decrease the utility of the sender. Again we arrived at a contradiction to the minimality of G .

Since there are k intervals (z_j, z_{j+1}) in which u is continuous and does not change convexity or concavity, we have k indices $i \in A$ for which $(\underline{y}_i, \bar{y}_i) \subseteq (z_j, z_{j+1})$ for some j . All other intervals $(\underline{y}_i, \bar{y}_i)$ intersect at least two intervals (z_j, z_{j+1}) . Therefore, we must have $|A| \leq 2k$, as desired.

We turn to the other statements in the proposition. The first case corresponds to a full information interval. The second case corresponds to a pooling interval. The third case corresponds to a bi-pooling interval (in the uniform case we can utilize Corollary 1 to deduce that $\bar{z}_i - \underline{z}_i \leq \frac{1}{2}(y_i - y_{i-1})$). \square

REFERENCES

- Aliprantis, Charalambos D. and Kim C. Border (2006), *Infinite Dimensional Analysis: A Hitchhikers Guide*. Springer, Berlin. [20]
- Boleslavsky, Raphael, Ilwoo Hwang, and Kyungmin Kim (2019), “Competitive advertising and pricing.” Emory University and University of Miami. [16]
- Candogan, Ozan (2019a), “Persuasion in networks: Public signals and k-cores.” Available at SSRN. [16, 17]
- Candogan, Ozan (2019b), “Optimality of double intervals in persuasion: A convex programming framework.” Available at SSRN. [16]
- Candogan, Ozan (2020), “Reduced form information design: Persuading a privately informed receiver.” Available at SSRN. [17, 24]
- Candogan, Ozan and Philipp Strack (2021), “Optimal disclosure of information to a privately informed receiver.” Available at SSRN 3773326. [24]

Deniz, Dizdar and Eugen Kováč (2019), “A simple proof of strong duality in the linear persuasion problem.” Available at SSRN 3426166. [16, 24]

Dworczak, Piotr and Anton Kolotilin (2019), “The persuasion duality.” Available at SSRN 3474376. [16]

Dworczak, Piotr and Giorgio Martini (2019), “The simple economics of optimal persuasion.” *Journal of Political Economy*, 127, 1993–2048. [16, 17, 21, 23, 24, 25, 32]

Gentzkow, Matthew and Emir Kamenica (2016), “A Rothschild–Stiglitz approach to Bayesian persuasion.” *American Economic Review*, 106, 597–601. [16, 18, 22]

Goldstein, Itai and Yaron Leitner (2018), “Stress tests and information disclosure.” *Journal of Economic Theory*, 177, 34–69. [17]

Guo, Yingni and Eran Shmaya (2019), “The interval structure of optimal disclosure.” *Econometrica*, 87, 653–675. [18, 24]

Harbaugh, Rick and Eric Rasmusen (2018), “Coarse grades: Informing the public by withholding information.” *American Economic Journal: Microeconomics*, 10, 210–235. [16]

Kleiner, Andreas, Benny Moldovanu, and Philipp Strack (2021), “Extreme points and majorization: Economic applications.” *Econometrica*, 89, 1557–1593. [17, 18, 20]

Kolotilin, Anton (2018), “Optimal information disclosure: A linear programming approach.” *Theoretical Economics*, 13, 607–635. [16, 18, 22]

Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li (2017), “Persuasion of a privately informed receiver.” *Econometrica*, 85, 1949–1964. [16]

Kolotilin, Anton, Tymofiy Mylovanov, and Andriy Zapechelnyuk (2022), “Censorship as optimal persuasion.” *Theoretical Economics*, 17, 561–585. [17, 26]

Kolotilin, Anton and Alexander Wolitzky (2020), “Assortative information disclosure.” [16]

Mensch, Jeffrey (2019), “Monotone persuasion.” Working Paper. [17]

Roesler, Anne-Katrin and Balázs Szentes (2017), “Buyer-optimal learning and monopoly pricing.” *American Economic Review*, 107, 2072–2080. [16]

Yamashita, Takuro (2018), “Optimal public information disclosure by mechanism designer.” Working Paper. [16]

Co-editor Marina Halac handled this manuscript.

Manuscript received 15 December, 2020; final version accepted 20 March, 2022; available online 1 April, 2022.