



Munich Personal RePEc Archive

Disclosure and Pricing of Attributes

Smolin, Alex

University of Bonn

October 2017

Online at <https://mpra.ub.uni-muenchen.de/84132/>
MPRA Paper No. 84132, posted 23 Jan 2018 15:26 UTC

Disclosure and Pricing of Attributes*

Alex Smolin

October 3, 2017

Abstract

A monopolist seller owns an object that has several attributes. A buyer is privately informed about his tastes and uncertain about the attributes. The seller can disclose attribute information to the buyer in a form of a statistical experiment. The seller offers a menu of call options varying in upfront payments, experiments, and strike prices. I study revenue-maximizing menus and show that optimal experiments belong to a simple class of linear disclosures. I fully characterize an optimal menu for a class of single-minded buyers. Surprisingly, the menu is nondiscriminatory and can be implemented by a single partial disclosure followed by a posted price.

Keywords: attributes, information design, mechanism design, private disclosure, call options, multidimensional screening, demand transformation

JEL Codes: D42, D82, D83

*Institute for Microeconomics, University of Bonn, alexey.v.smolin@gmail.com. I thank Daniel Krämer for motivating discussions and Dirk Bergemann for productive conversations. I am grateful to participants of research seminars at Toulouse and Mannheim, as well as at the SAET 2017 conference.

1 Introduction

In many important markets, sellers have considerable control over information available to their buyers. Real estate brokers can control the extent of property investigation and the documentation they supply; car retailers can limit the duration of test drives and amount of technical pre-sale support; professional networks can decide what parts of a job candidate's profile to reveal to recruiting companies. In all of these markets, the products (i.e., real property, car, meeting with a job candidate) are characterized by multiple attributes that appeal to different types of buyers. To maximize the revenue, the sellers need to understand what attribute information to provide, how to price their products, and whether and how to price the information provided.

In this paper, I develop a framework to study information disclosure and pricing of multi-attribute products. I consider a monopolist seller who has an indivisible object for sale to a single buyer. The object has several attributes, and the buyer is uncertain about their values. The seller does not know what attributes the buyer likes or how much. These preferences are the buyer's private information and constitute the buyer's type. The seller controls pricing and, importantly, can disclose attribute information to the buyer.

The seller offers a menu of call options that differ in their informativeness. Each menu item consists of a nonrefundable upfront payment, an attribute information, and a strike price for the object. I model attribute information as an arbitrary statistical experiment informative about attributes. Information control enables price discrimination. By varying the upfront payment, the experiment, and the strike price, the seller can screen buyer's types. This menu mechanisms provide a natural and simple framework for studying interaction between information disclosure and pricing. I study and characterize revenue-maximizing menus.

To see how the seller can benefit from attribute disclosure, consider the example with the object having two attributes. The attributes are distributed uniformly and independently, the first attribute over $[0, 1]$ and the second attribute over $[0, 2]$. The buyer can be one of two types, each equally likely. The first type values only the first attribute and the second type values only the second attribute, both with unit intensity. Under no disclosure, the types' ex-ante valuations are $1/2$ and 1 . The optimal no-disclosure price is $1/2$ resulting in the revenue $1/2$. The seller can increase the revenue by providing partial disclosure. Let the seller inform the buyer whether the first attribute is above $1/2$. The first type valuation splits into $1/4$ and $3/4$, equally likely; the second type valuation remains at 1 . The optimal partial-disclosure price is $3/4$ resulting in the revenue $9/16 > 1/2$. The seller is able to increase revenue by effectively targeting the first type with the disclosure. I show that, in

fact, this simple nondiscriminatory mechanism also dominates full disclosure and is optimal among all menu mechanisms in this example.

I study general optimal mechanisms in Section 3. I show that if there is only one attribute, then an optimal mechanism is no disclosure followed by a single posted price. The argument is analogous to that of [Esó and Szentes \(2007\)](#). Even if the seller could condition the price on the attribute, optimally she would always sell the object to the same types. No disclosure followed by the posted price that targets these types implements the optimal allocation and delivers the optimal revenue.

The analysis is richer in the presence of multiple attributes. I show that an optimal menu can involve discriminatory pricing and partial disclosure. The class of all possible disclosures is very large, covering all stochastic mappings from a continuum of attribute vectors. However, I show optimal experiments belong to a simple class of linear disclosures that inform the buyer whether a linear combination of attributes is above or below some threshold. This allows me to connect the seller's problem to a multiproduct bundling problem.

In Section 4, I study the leading case of a single-minded buyer, in which all types value only one attribute. This type class allows for both horizontal heterogeneity between attribute cohorts and vertical heterogeneity within attribute cohorts, yet is sufficiently tractable. I start with a simpler case of orthogonal types that features only one type per attribute and, hence, allows for only horizontal heterogeneity. I characterize an optimal menu and show it is nondiscriminatory and can be implemented by a free-of-charge disclosure followed by a posted price. I proceed by studying a general case that allows for both vertical and horizontal heterogeneity. This analysis features a continuum of types within each attribute cohort. Surprisingly, I show that for a broad class of type distributions, the optimal menu remains nondiscriminatory and again can be implemented by a single disclosure followed by a posted price.

I conclude with discussion of my findings in Section 5. First, I illustrate how optimal disclosure transforms the demand curve. I show that by providing partial attribute information the seller can target specific types and hence rotate the demand curve locally. It resonates with the analysis of global demand rotations of [Johnson and Myatt \(2006\)](#). Second, I show full disclosure is detrimental to the seller and the seller can benefit from conditioning the price on the information disclosed. These results contrast with those of [Esó and Szentes \(2007\)](#) and highlight the qualitative difference between our frameworks.

Related Literature The paper contributes to the literature on private disclosure and pricing. One strand of this literature focuses on nondiscriminatory mechanisms in which a single disclosure is followed by a posted price. [Lewis and Sappington \(1994\)](#) introduce

these mechanisms in a setting where a buyer has no prior information. They find optimal disclosure within a simple parameterized class and show that it is generally extreme, either full or no disclosure. [Bergemann and Pesendorfer \(2007\)](#) further observe that, if it is always efficient to trade, then no disclosure dominates any other possible disclosure simply because the seller can extract full expected surplus.¹ [Johnson and Myatt \(2006\)](#) extend the analysis to settings in which the buyer has prior information. They focus on disclosures that correspond to global rotations of a demand curve and show that, once again, extreme disclosures are optimal. My paper contributes to this literature by showing that if the product has several attributes, then partial attribute disclosure can dominate both full and no disclosure even if it is always efficient to trade (Section 5.1).

At the same time, when the buyer has private information, it is natural to study discriminatory mechanisms and how they can be used to screen the buyer’s types. In an influential paper, [Esó and Szentes \(2007\)](#) study the settings in which the seller provides what may be seen as “valuation-rank” disclosure. Information conveyed by a valuation-rank disclosure corresponds to statements such as, “Your valuation is in your x-th percentile,” uniform across all types. It is plausible in situations in which the types capture additive valuation components. In such settings, [Esó and Szentes \(2007\)](#) show that the seller optimally provides full disclosure and, moreover, does not benefit from conditioning the price on the information disclosed.² [Li and Shi \(2017\)](#) obtain different results in the framework of what can be seen as “valuation-level” disclosure. Information conveyed by a valuation-level disclosure corresponds to statements such as, “Your valuation is above x,” uniformly across all types. It is plausible in common value settings in which the types represent private information about the object. In such settings, [Li and Shi \(2017\)](#) show that the seller optimally provides only partial disclosure but are not able to identify an optimal mechanism.

My paper contributes to this literature by introducing attribute disclosure, which is natural in many settings yet not covered by existing models. I identify the class of optimal experiments and fully characterize optimal mechanisms when the buyer is single-minded. The results highlight that optimal disclosure can be partial and the seller can strictly benefit from conditioning the price on the information disclosed (Section 5.2).

Finally, the paper builds on several existing frameworks. Multiattribute utility function follows [Lancaster \(1966\)](#). The timing is analogous to sequential screening of [Courty and Li \(2000\)](#). Disclosure policy in a form of a statistical experiment is a feature of information design and Bayesian persuasion (e.g., [Rayo and Segal \(2010\)](#) and [Kamenica and Gentzkow](#)

¹See, however, [Anderson and Renault \(2006\)](#), who show that optimal disclosure is partial if the purchase is associated with search costs and the seller cannot commit to prices.

²[Esó and Szentes \(2017\)](#) generalize the latter finding to dynamic environments.

(2011)). The analysis relies on machinery of mechanism design (e.g., Myerson (1981, 1982)). Discriminatory menu of experiments appeared already in my previous work (Bergemann, Bonatti, and Smolin (2017a)). There, the seller offers only information products, not an object. The buyer’s value for information is determined by an outside decision problem and his action is not contractable.

2 Model

A buyer decides whether to buy a single indivisible object from a seller. The object has a finite number J of characteristics or attributes. The attribute values constitute an *attribute vector* $x = (x_1, \dots, x_j, \dots, x_J) \in \mathcal{X} = \mathbb{R}_+^J$. The buyer’s preferences towards each attribute constitute the buyer’s *type* $\theta = (\theta_1, \dots, \theta_j, \dots, \theta_J) \in \Theta \subseteq \mathbb{R}_+^J$. The ex-post buyer’s valuation for the object comes from a multiattribute framework:

$$v(\theta, x) = \theta \cdot x = \sum_{j=1}^J \theta_j x_j. \quad (1)$$

The seller’s valuation for the object is fixed to be zero. Hence, the buyer always values the object more than the seller does, and it is always efficient to trade. Seller and buyer utilities are quasilinear in transfers.

Prior Information Attributes are continuously distributed over \mathcal{X} according to a cumulative distribution function G with full support. They can be correlated. The buyer and the seller are symmetrically informed about them. The type space Θ can be finite or infinite. The buyer’s type is his private information. From the seller’s perspective, the probability (probability density) of type θ is $f(\theta)$.

Information Disclosure The seller can disclose additional attribute information to the buyer. I model this information as a statistical experiment $E = (S, \pi)$ that consists of a signal set S and a likelihood function:

$$\pi : \mathcal{X} \rightarrow \Delta(S).$$

The experiment can be arbitrarily informative about the attributes. It can provide no information, or no disclosure, $\underline{E} \triangleq (\underline{S}, \underline{\pi})$, with \underline{S} being a singleton. It can fully reveal attributes, or provide full disclosure, $\overline{E} \triangleq (\overline{S}, \overline{\pi})$, with $\overline{S} = \mathcal{X}$ and $\overline{\pi}(x)$ placing probability 1 on $s = x$. Alternatively, it can provide partial information, for example, as illustrated in




Figure 1: Partially informative experiment with $J = 2$ and $S = (s_1, s_2, s_3, s_4)$. Colors indicate regions in which the corresponding signals are sent.

Figure 1. In this figure, after observing a signal s_1 the buyer knows that attributes belong to the red area but does not know to which part of it.

I highlight that attribute information affects valuation of different types differently according to the multiattribute valuation function (1). Consequently, attribute information cannot be represented by an experiment that informs the buyer directly about his valuation.

Selling Mechanisms The seller can price both the good and the information she provides. The seller designs a menu $\mathcal{M} = (\mathbf{E}, \mathbf{r}, \mathbf{t})$ to be offered to a buyer. It consists of a collection of experiments \mathbf{E} and tariff functions $\mathbf{r} : \mathbf{E} \rightarrow \mathbb{R}_+$ and $\mathbf{t} : \mathbf{E} \rightarrow \mathbb{R}_+$. The first tariff captures a price of information—an upfront payment paid irrespectively of a trade. The second tariff captures a price of the object—a strike price paid only if the trade occurs. Effectively, the menu is a collection of call options differing in monetary terms and information disclosure. The timing is as follows:

- the seller posts a menu \mathcal{M} ,
- the attribute vector x and the buyer's type θ are realized,
- the buyer chooses an experiment $E \in \mathbf{E}$ and pays the corresponding price $r(E)$,
- the buyer observes a signal s from the experiment E , and
- the buyer decides whether to buy the object at the strike price $t(E)$.

I highlight that the timing is analogous to that of [Li and Shi \(2017\)](#)'s discriminatory private disclosure. It assumes the seller commits to a menu before the realization of the attributes x and the type θ . The true attributes x and experiment realization s are not contractible.³ Sequential interactions between the players are excluded, so belief-elicitation schemes and scoring rules are not available.⁴ This setting provides a natural and simple framework for studying how information disclosure and pricing interact in designing screening mechanisms. My goal is to characterize a revenue-maximizing menu for the seller.

3 Disclosure and Pricing

3.1 Single Attribute

I begin with the basic case of a single attribute, $J = 1$, $x \in \mathbb{R}_+$. The buyer's type is one-dimensional, $\theta \in \Theta = \mathbb{R}_+$, and the buyer's valuation is

$$v(\theta, x) = \theta x.$$

This setting features vertical type heterogeneity. Informing the buyer about an attribute value x is equivalent to informing him about an attribute percentile $G(x)$. At the same time, the attribute percentile is equal to the valuation percentile from a perspective of any type. In other words, the valuation percentile for all types coincides, and informing about the attribute is equivalent to informing the types about this valuation percentile. It follows that a single-attribute disclosure can be approached as a valuation-rank disclosure of [Eső and Szentes \(2007\)](#), and I can apply their argument to solve the seller's problem.

The argument starts by considering a more beneficial setting for the seller in which she can condition payment on the experiment realization. In this case, the revelation principle implies I can focus on direct mechanisms in which the buyer reports his type θ and pays the upfront payment $r(\theta)$, the seller completely discloses the attribute x , and the trade happens with probability $q(x, \theta)$. The terms of trade determine the trade surplus

$$Y(\theta) \triangleq \int_{x \in \mathcal{X}} x q(x, \theta) dG(x)$$

that can be anywhere between 0 and $\mathbb{E}x$. I can then rewrite the seller's problem as maxi-

³For instance, the buyer cannot claim a refund ex post. See [Krähmer and Strausz \(2015\)](#), [Heumann \(2017\)](#) and [Bergemann, Castro, and Weintraub \(2017b\)](#) for recent studies on ex-post incentive constraints.

⁴[Krähmer \(2017\)](#) investigates usefulness of such schemes in screening problems with information design.

mizing the revenue directly over $r(\theta)$ and $Y(\theta)$ as

$$\max_{r(\theta), 0 \leq Y(\theta) \leq \mathbb{E}x} \int_{\theta \in \Theta} r(\theta) dF(\theta),$$

subject to incentive-compatibility constraints

$$\theta Y(\theta) - r(\theta) \geq \theta Y(\theta') - r(\theta'), \quad \forall \theta, \theta' \in \Theta,$$

and individual-rationality constraints

$$\theta Y(\theta) - r(\theta) \geq 0, \quad \forall \theta \in \Theta.$$

This problem is analogous to a canonical mechanism design problem of [Myerson \(1981\)](#). The optimal allocation $Y(\theta)$ is a step function, equal to 0 for $\theta < \theta^*$ and to $\mathbb{E}x$ for $\theta \geq \theta^*$. The corresponding optimal upfront payment $r(\theta)$ is equal to 0 for $\theta < \theta^*$ and equal to $r^* = \theta^* \mathbb{E}x$ for $\theta \geq \theta^*$.

The argument concludes by noting the optimal mechanism can be implemented by a posted price mechanism: providing no disclosure and charging a strike price r^* for the object. The posted price mechanism does not require the seller to observe the information disclosed and, hence, is optimal in the original problem as well.

Theorem 1. (Single Attribute) *If there is only one attribute, $J = 1$, then an optimal menu is a posted price mechanism with no disclosure, i.e., it contains a single item with zero upfront payment, $r = 0$, not informative experiment, $E = \underline{E}$, and a strike price t .*

There is a simple intuition behind the optimality of no disclosure if the seller can only use a nondiscriminatory mechanism consisting of a single disclosure followed by a posted price. Consider an arbitrary disclosure. Any signal realization s scales the demand with the proportionality coefficient equal to the expected attribute $\mathbb{E}[x | s]$. If the seller could observe this realization, she would charge a scaled price and received a scaled revenue. Importantly, the seller would sell the object to the same types irrespectively of the realization s . Hence, since any expectation under an information disclosure is a martingale, the seller would serve the same population share and charge, on average, the same price. The seller can do just as well using a posted price with no disclosure.

This argument does not allow the seller using discriminatory schemes with upfront payments. [Theorem 1](#) confirms that no disclosure is optimal even if the principal can do that. I highlight that this result requires no assumptions on type distribution F or attribute distribution G beyond the common knowledge of trade efficiency. Hence, very generally, if

the object has only one relevant attribute, then the seller has no incentives to provide any attribute information. However, as I show in the following sections, the analysis of multiple attributes is different.

3.2 The Buyer's Problem

I proceed with studying the general case of multiple attributes, $J \geq 2$. Consider the buyer's incentives when he chooses an experiment from a given menu. Let his type be θ . If he chooses an experiment $E = (S, \pi)$ then he pays the upfront price $r(E)$. Then, a signal s is realized according to the likelihood function $\pi(\theta)$. The realization results in the interim expectation

$$X(E, s) \triangleq \mathbb{E}[x \mid E, s].$$

Finally, the buyer decides whether to buy the object and does so optimally if and only if $\theta \cdot X - t(E)$ is greater than 0. Integrating over signal realizations, I can define the resulting *total trade probability* as

$$Q(E, \theta) \triangleq \mathbb{P}_E(\theta \cdot X - t(E) \geq 0).$$

The corresponding indirect utility of choosing an experiment E can be written as

$$U(E, \theta) = -r(E) + \mathbb{E}_E \max\{0, \theta \cdot X - t(E)\}, \quad (2)$$

Naturally, types seek information that most fits their interests. For example, if a type values only one attribute, then the type places no value on the experiments that provide no information about it. As a result, the types can disagree on the experiment ranking even if tariffs are the same. Faced with a menu, the types self-select into different items. This gives the seller opportunity to discriminate among them by carefully designing the menu.

3.3 The Seller's Problem

The seller's problem lies at the intersection of mechanism and information design because the seller can both control the information available to the buyer and charge monetary transfers. In principle, she can offer complex experiments in an attempt to better discriminate among types. However, in the next couple of sections, I show that an optimal class of experiments is simple and tractable.

I begin by appealing to the revelation principle and focusing on direct mechanisms $\mathcal{M} = (r(\theta), E(\theta), t(\theta))$, which ask the buyer his type and assign the experiment and the tariffs as functions of his report. The seller's revenue obtained from a particular type consists of

the upfront payment $r(\theta)$ and, if the buyer decides to purchase the object, trade revenue $t(\theta)$. The seller's problem can then be written as maximizing the total expected revenue

$$\max_{\{r(\theta), E(\theta), t(\theta)\}} \int_{\theta \in \Theta} (r(\theta) + Q(E(\theta), \theta) t(\theta)) dF(\theta) \quad (3)$$

subject to the incentive-compatibility constraints

$$U(E(\theta), \theta) \geq U(E(\theta), \theta'), \quad \forall \theta, \theta' \in \Theta, \quad (4)$$

and the individual-rationality constraints

$$U(E(\theta), \theta) \geq 0, \quad \forall \theta \in \Theta. \quad (5)$$

There are complications involved in the seller's problem. First, the buyer's type is a vector and multidimensional mechanism design problems are notoriously difficult. Second, the seller maximizes over all possible experiments, which are all possible stochastic mappings from the continuum of attribute vectors.

I begin approaching this problem by binding the size of the optimal signal sets. I follow the arguments of [Bergemann et al. \(2017a\)](#). For a given direct mechanism \mathcal{M} , I call an experiment $E(\theta)$ *responsive* if $S(\theta) = \{s^+, s^-\}$ and type θ , when choosing this experiment, purchases the good if and only if s^+ . Effectively, responsive experiments guide the buyer's action. I call the menu *responsive* if all of its experiments are responsive.

Proposition 1. (Responsive Menus) *The outcome of every menu can be replicated by a responsive menu.*

The proof is analogous to the argument of the revelation principle of [Myerson \(1982\)](#). Intuitively, an experiment should provide information minimal to guiding the decision of a truth-telling type. If the menu contains nonresponsive experiments, then the seller can replace them with responsive experiments that replicate the behavior of truth-telling types. After this modification, truth telling delivers the same payoff as before. Dishonesty, however, becomes weakly less appealing because the modified experiments are weakly less informative.

Focus on responsive menus allows characterizing every experiment $E(\theta)$ in the menu by a more tractable object. In particular, $E(\theta)$ can be characterized by its *trade function*,

$$q(\theta, x) \triangleq \Pr(s^+ \mid E(\theta), x).$$

The function defines a probability of trade given the object attributes. Effectively, it is a reduced form of the likelihood function $\pi(\theta)$. In the next section, I establish structural

properties of the optimal trade function and hence optimal experiments.

3.4 Linear Disclosure

Call a responsive experiment $E(\theta)$ a *linear disclosure* if its trade function is such that there exist $\alpha(\theta) \in \mathbb{R}^J$ and $b(\theta) \in \mathbb{R}$ such that

$$q(\theta, x) = \begin{cases} 1, & \text{if } \sum_{j=1}^J \alpha_j(\theta) x_j > b(\theta), \\ 0, & \text{if } \sum_{j=1}^J \alpha_j(\theta) x_j < b(\theta). \end{cases}$$

A linear disclosure informs the buyer whether a linear combination of attributes is above or below some threshold. Equivalently, its trade function is an indicator function of an attribute half-space. A linear disclosure is conditionally nonstochastic. In the case of a single attribute a linear disclosure corresponds to a binary monotone partition disclosure.

If all elements of the coefficient vector $\alpha(\theta)$ are positive, then a linear disclosure can be viewed as a reference disclosure that informs the buyer whether a hypothetical “reference” type $\theta' = \alpha(\theta)$ would like to buy the good at a price $b(\theta)$. If attributes are independently distributed, observing a “buy” recommendation uniformly increases attribute expectation, whereas observing a “no-buy” recommendation uniformly decreases it. In contrast, if the elements of a coefficient vector $\alpha(\theta)$ have different signs, then a linear disclosure can be viewed as a “comparative” disclosure between the attribute groups of different signs. If attributes are independently distributed, a “buy” recommendation increases the attribute expectations in one group and decreases them in the other group. In Figure 2 I illustrate these two types of linear disclosure in the case of two attributes.

Theorem 2. (Linear Disclosure) *Every experiment in an optimal responsive menu is a linear disclosure.*

I use a Lagrange approach to establish this result. By standard arguments, any solution to the seller’s problem is associated with a collection λ of Lagrange multipliers, one for each of the constraint (4) and (5). The optimal tariff and trade functions must maximize a Lagrange function, which is a linear combination of the objective function and the constraints, weighted by the corresponding multipliers. The Lagrange function can be written as:

$$L(r(\theta), q(\theta, x), t(\theta), \lambda) \sim \int_{\theta \in \Theta} \int_{x \in \mathcal{X}} (\Lambda(r, t, \lambda) \cdot x - B(r, t, \lambda)) q(\theta, x) dG(x) dF(\theta),$$

where Λ and B are functions of tariffs and Lagrange multipliers but not of the trade function. It follows that the optimal trade function maximizes the Lagrange function piecewise and




Figure 2: Linear disclosure $E(\theta)$ in the case of two attributes, $J = 2$. Colors indicate regions in which the corresponding recommendations are sent. Left: reference disclosure, $\alpha_1(\theta) > 0, \alpha_2(\theta) > 0, b(\theta) > 0$. Right: comparative disclosure, $\alpha_1(\theta) > 0, \alpha_2(\theta) < 0, b(\theta) > 0$.

thus corresponds to linear disclosure, precisely with coefficients Λ and B .⁵

Theorem 2 establishes that despite the complexity of the seller's problem, all optimal experiments belong to a very tractable class of linear disclosures. This simplifies the problem and allows relating it to the multiproduct monopolist problem. First, note that under the linear disclosure the total trade probability can be written as

$$Q(\theta) = \Pr(s^+ | E(\theta)) = \int_{x \in \mathcal{X}} q(\theta, x) dG(x).$$

Further, defining the *trade expectation* as

$$X(\theta) \triangleq \mathbb{E}[x | E(\theta), s^+] = \frac{\int_{x \in \mathcal{X}} xq(\theta, x) dG(x)}{\int_{x \in \mathcal{X}} q(\theta, x) dG(x)}.$$

Note that the trade probability and trade expectation pin down their counterparts: the no-trade probability $\Pr(s^- | E(\theta))$ is equal to $1 - Q(\theta)$ and the no-trade expectation $\mathbb{E}[x | E(\theta), s^-]$ is determined by the martingale property of expectations. It then follows that the pair $(Q(\theta), X(\theta))$ is a sufficient statistic for the seller problem's constraints. Moreover, the trade expectation pins down the corresponding linear disclosure.

⁵The Lagrange approach has been widely used in economic optimization problems. However, it is worth noting that in the information design literature [Matejka and McKay \(2014\)](#) similarly used it to identify an optimal class of experiments in a single-agent decision problem of information acquisition.

Proposition 2. (Maximal Probability) *For each $X \in \mathcal{X}$ there exists a unique linear disclosure E such that*⁶

$$\mathbb{E}[x \mid E, s^+] = X.$$

Moreover, the linear disclosure maximizes the probability of signal s^+ within the class of experiments such that $S = \{s^+, s^-\}$ and $\mathbb{E}[x \mid E, s^+] = X$.

Proposition 2 is straightforward in the case of a single attribute. The multidimensional extension can be proven by a Lagrange method analogous to the argument of Theorem 2. It suggests that a linear disclosure can be viewed as a natural multidimensional extension of a binary monotone partition disclosure.

The result gives additional intuition behind the optimal menu. It states that for a given trade expectation, optimal experiments maximize trade probability and hence constrained efficient. I highlight the constrained efficiency could *a priori* be not optimal because the trade probability affects the no-trade expectations and the corresponding double-deviation constraints.

It follows that the trade expectation $X(\theta)$ can serve as a determinant of experiment $E(\theta)$. The total trade probability $Q(\theta)$ can be written as a function of the trade expectation as $Q(X(\theta)) = \Pr(s^+ \mid E(\theta))$. The so-defined function $Q(X)$ has the following properties.

Lemma 1. *The total trade probability $Q(X)$ has the following properties:*

1. $Q(X)$ is differentiable,
2. $Q(\mathbb{E}x) = 1$, $\lim_{X \rightarrow \partial\mathcal{X}} Q(X) = 0$, and
3. If attributes are independently distributed, then $\partial Q / \partial X_j \geq 0$ for $X_j \leq \mathbb{E}x_j$.

Naturally, inducing higher trade expectation comes at a cost of lower trade probability. The exact shape of $Q(X)$ depends on the attribute distribution G and is generally nonlinear.

Finally, the seller's problem can be written as maximization directly over the trade expectation $X(\theta)$ and tariffs $r(\theta), t(\theta)$ as maximizing the revenue

$$\max_{\{r(\theta), X(\theta), t(\theta)\}} \int_{\theta \in \Theta} (r(\theta) + Q(\theta) t(\theta)) dF(\theta) \quad (6)$$

⁶The uniqueness is defined up to a zero probability event.

subject to multiple incentive-compatibility constraints

$$\begin{aligned}
(\theta \cdot X(\theta) - t(\theta))Q(\theta) - r(\theta) &\geq (\theta \cdot X(\theta') - t(\theta'))Q(\theta') - r(\theta'), \\
(\theta \cdot X(\theta) - t(\theta))Q(\theta) - r(\theta) &\geq \theta \cdot (\mathbb{E}x - X(\theta')Q(\theta')) - (1 - Q(\theta'))t(\theta') - r(\theta'), \\
(\theta \cdot X(\theta) - t(\theta))Q(\theta) - r(\theta) &\geq \theta \cdot \mathbb{E}x - t(\theta') - r(\theta'), \\
(\theta \cdot X(\theta) - t(\theta))Q(\theta) - r(\theta) &\geq -r(\theta'),
\end{aligned}$$

the individual-rationality constraints

$$(\theta \cdot X(\theta) - t(\theta))Q(\theta) - r(\theta) \geq 0,$$

and the feasibility constraint

$$Q(\theta) = Q(X(\theta)), \quad \forall \theta, \theta' \in \Theta.$$

The multiple incentive-compatibility constraints ensure each type prefers truth telling to all double-deviating strategies: misreporting and following the recommendations of another experiment, “swapping” the buying decisions, always buying, or never buying, respectively. Deviations from θ to θ are included and ensure that the types are obedient on-path, after truth telling.

This analysis reduces the search for an optimal menu to a concrete multidimensional mechanism design problem. Even though the seller sells a single object, information disclosure allows him to control the buyer’s expectations $X(\theta)$ at the time of a purchase. However, unlike in the bundling problem, the seller cannot choose the expectations at will. They must be martingale so, for any increase in the expectations after recommendation to buy, the seller must compensate by decreasing the expectations after recommendation to not buy. Moreover, not any expectation martingale can be induced by a disclosure. The optimal trade probability is achieved by linear disclosure and captured by the trade probability function $Q(X(\theta))$.

Even the most basic multidimensional mechanism design problems are known to be prohibitively difficult.⁷ The current problem is further complicated by the presence of multiple double-deviation constraints and the non-linearity of the trade probability $Q(X(\theta))$. In the next section, I make progress by restricting my attention to a specific class of types. This class is sufficiently tractable to allow full characterization of an optimal menu, yet sufficiently rich to capture both vertical and horizontal type heterogeneity.

⁷Bergemann, Shen, Xu, and Yeh (2012) and Daskalakis, Deckelbaum, and Tzamos (2017) highlight the difficulties associated with the multiproduct monopolist problem.

4 Single-Minded Buyer

I call a type *single-minded* if it values only one attribute. For a generic single-minded type, the vector θ places a positive weight only on one dimension:

$$\theta = (0, \dots, 0, \theta_j, 0, \dots, 0).$$

Thus, single-minded types allow for a simpler notation. I can represent the types by J attribute cohorts Θ_j such that all types within the same cohort value the same attribute. I slightly abuse the notation and let the type subscript identify his attribute cohort and the type value identify the valuation intensity, so that $\Theta_j \subseteq \mathbb{R}_+$ and

$$v_j(\theta_j, x) = \theta_j x_j \quad \forall j, \theta_j \in \Theta_j.$$

I denote the frequency of a cohort Θ_j by $f(\Theta_j)$ and the type distribution within the cohort by $F_j(\theta_j)$.

A buyer is *single-minded* if all types $\theta \in \Theta$ are single-minded *and* attribute values are independently distributed, so that $x_j \sim G_j$ and $G(x) = \times_j G_j(x_j)$.⁸ The independence requirement is substantive. Starting with the general case, one can always redefine attributes as valuations of the corresponding types. In this formulation, each type naturally values only the attribute that arose from his original valuation. However, the so-defined attributes can be correlated with the correlation structure determined by the original attribute and type distributions.

If the buyer is single-minded, then the seller knows that the buyer values only one attribute but does not know which one or by how much. This type structure allows further narrowing of the class of optimal experiments. Because attributes are independently distributed, a type $\theta_j \in \Theta_j$ values only information about attribute j . Information about other attributes does not change his ex-ante valuation and has no value for him. This observation suggests an optimal way to screen single-minded types in a direct mechanism: if the buyer reports type $\theta_j \in \Theta_j$, then the seller should provide information only about attribute j . Providing any other information would make misreporting more appealing without adding value for truth telling. The following proposition confirms this intuition.

⁸The name is inspired by “single-minded” bidders studied in the literature on combinatoric auction design. A “single-minded” type values a specific attribute whereas a “single-minded” bidder values a specific bundle. See, for example [Lehmann, O’callaghan, and Shoham \(2002\)](#). I thank Laura Doval for drawing my attention to this helpful connection.

Proposition 3. (Discriminatory Disclosure) *If the buyer is single-minded, then in an optimal responsive menu, an experiment $E_j(\theta_j)$ is informative only about attribute j . That is, for all j , $\theta_j \in \Theta_j$, $q_j(\theta_j, x) = q_j(\theta_j, x')$ whenever $x_j = x'_j$.*

Proposition 3 establishes that every optimal experiment $E_j(\theta_j)$ provides information about a single dimension j . At the same time, by Theorem 2, $E_j(\theta_j)$ is a linear disclosure. However, any linear disclosure that is informative only about attribute j is effectively a binary monotone partition defined on this attribute.

Corollary 1. *If the buyer is single-minded, then in an optimal responsive menu, an experiment $E_j(\theta_j)$ is a binary monotone partition of attribute j .*

It follows that an optimal experiment $E_j(\theta_j)$ can be characterized by its threshold $b_j(\theta_j)$, so that it informs the buyer whether attribute j is above or below the threshold. By construction, the buyer should purchase the object in only one element of the partition. Incentive compatibility requires this element correspond to attributes above the threshold resulting in the trade function

$$q_j(\theta_j, x) = \begin{cases} 1, & \text{if } x_j > b_j(\theta_j), \\ 0, & \text{if } x_j < b_j(\theta_j). \end{cases}$$

As discussed in Section 3.4, it is more convenient to characterize the experiment by its trade expectation. Because an optimal experiment $E_j(\theta_j)$ is informative only about attribute j , trade expectations of all other attributes stay at their ex-ante values. Hence, I can characterize the experiment by the trade expectation of attribute j ,

$$X_j(\theta_j) = \mathbb{E}[x_j \mid x_j \geq b_j(\theta_j)].$$

Similarly, the total trade probability can be written as the function of that expectation $Q_j(X_j)$. The trade expectation X_j can take any value between $\mathbb{E}x_j$ and the upper bound of x_j , taken to be ∞ . As it increases from $\mathbb{E}x_j$ to ∞ , the corresponding disclosure threshold b_j increases from 0 to ∞ , and the total trade probability Q_j decreases from 1 to 0.

It follows that to characterize an optimal menu, I need to find the optimal trade expectation functions $X_j(\theta_j)$ and tariff functions $r_j(\theta_j)$, $t_j(\theta_j)$. I start with a simpler case that does not incorporate the valuation heterogeneity within each attribute.

4.1 Orthogonal Types

I begin with the case in which each attribute cohort is a singleton, $\Theta_j = \{\theta_j\}$, so there is no vertical heterogeneity within each attribute and the number of types equals to the number of

attributes $|\Theta| = J$. In this case, without loss of generality, I can set all valuation intensities equal to 1,

$$\theta_j \equiv 1 \quad \forall j = 1, \dots, J.$$

Note that in this case, any two different types $\theta, \theta' \in \Theta$ are orthogonal to each other as vectors in \mathbb{R}^J so I refer to this case as the setting of *orthogonal types*. I will omit the dependence on the intensity within each attribute and differentiate types by subscripts.

The class of orthogonal types features particularly tractable incentive constraints. If type θ_j misreports, then he is offered an experiment tailored to another orthogonal type and hence not informative about attribute j . Hence, the type has no reason to act upon the experiment realization and the tightest incentive-compatibility constraint one in which he always buys. All others can be dropped. The seller's problem can be written as

$$\begin{aligned} & \max_{\{r_j, X_j, t_j\}_{j=1}^J} \sum_{j=1}^J f(\theta_j) (r_j + Q_j t_j) & (7) \\ \text{s.t.} \quad & (X_j - t_j) Q_j - r_j \geq \mathbb{E} x_j - t_k - r_k, \quad \forall j, k = 1, \dots, J, \\ & (X_j - t_j) Q_j - r_j \geq 0, \\ & X_j \geq \mathbb{E} x_j, \quad Q_j = Q_j(X_j). \end{aligned}$$

This problem resembles a one-dimensional mechanism design problem with the following important differences. First, each item in this problem features both horizontal and vertical components. The upfront payments r_j are purely vertical—all types value them the same. The experiments E_j and the associated trade expectations X_j are purely horizontal—they are valuable only to the type θ_j . The object strike prices t_j are mixed as they are paid only if the trade occurs, probability of which depends on a type. Second, the problem is non-linear. Not only are there products between the strike prices t_j and the trade probabilities Q_j , but also the trade probability functions $Q_j(X_j)$ are generically nonlinear as well.

Because of these differences, I cannot apply standard mechanism design techniques. Instead, I solve the problem in a sequence of simplifications. In the first step, I observe that using upfront payments r_j is detrimental to the seller. For any strictly positive r_j the seller can reduce the transfer and increase the strike price t_j to keep the total expected transfer $r_j + Q_j t_j$ the same. This change does not affect utilities of truth-telling types or the seller's revenue. However, it makes misreporting less appealing. Intuitively, by shifting the expected transfer towards the strike price the seller better discriminates against the types who always purchase the object.

In the second step, I use the special structure of incentive-compatibility constraints to show strike price discrimination is not profitable as well. Indeed, because all experiments off

the truth-telling path bring no value, an optimal deviation is to the types associated with the lowest price \underline{t} . If there is any strike price variation, there is a type θ_j with a price $t_j > \underline{t}$. This type's item is not attractive to any other type. Moreover, for the type to be willing to pay a higher price, the item must contain partial disclosure. This leads to a contradiction: the seller can simultaneously lower the price t_j and trade expectation X_j in such a way that the type's rents $(X_j - t_j) Q_j$ remain the same but the expected payment $Q_j t_j$ increases. Intuitively, the seller may not lose surplus on types irrelevant for incentives of the others.

Once I establish that optimal mechanism is nondiscriminatory, finding optimal experiments is straightforward. The seller should maximize trade probability by providing minimal information sufficient to convince the buyer to make a purchase, attribute by attribute. If the type θ_j is ex-ante sufficiently optimistic, $\mathbb{E} x_j \geq t$, then the seller should provide no attribute information, $E_j = \underline{E}$. Otherwise, the seller should increase the type's trade expectation up to the strike price, $X_j = t$. The following theorem summarizes the findings.

Theorem 3. (Optimal Menu, Orthogonal Types) *If the buyer is single-minded and the types are orthogonal, then an optimal responsive menu is characterized by the following properties:*

1. For any attribute j , $r_j = 0$,
2. For any attribute j , $t_j = t$, and
3. For any attribute j , E_j is a binary monotone partition of x_j such that $\mathbb{E}[x_j | E_j, s^+] = \max\{t, \mathbb{E} x_j\}$.

The optimal menu is illustrated in Figure 3. I highlight its notable features. First, the menu has the standard “no distortions at the top, no rents at the bottom” property. Namely, all types with the ex-ante valuation above the optimal price always buy the object, whereas all other types are indifferent to participating in the mechanism. In this way, “the top” and “the bottom” are not single types, but two type classes that partition the type space. Second, the optimal menu admits a nondiscriminatory indirect implementation. The seller can simply provide a single combined disclosure followed by the optimal posted price. Because each type values only one attribute, he will focus on the relevant attribute information. Third, for given attribute distributions $\{G_j\}$, the optimal mechanism can be found easily by a two-stage algorithm. In the first stage, for any fixed price t , the algorithm uses the third property of Theorem 3 to find optimal thresholds b_j and define the corresponding trade probability Q . The so-defined function $Q(t)$ is effectively a “modified” demand curve that accounts for optimal disclosure. In the second stage, the algorithm uses the demand curve to find an optimal price.




Figure 3: Optimal mechanism in the case of orthogonal types. Green color indicates attribute regions in which a purchase recommendation is sent, for types with partial disclosure. Blue color indicates the types' rent conditional on a trade, for types with no disclosure. Attributes are ordered by increasing ex-ante expectations.

4.2 General Case

I now extend the analysis to the general case of single-minded types and allow for the vertical heterogeneity within each attribute. In particular, I assume that for all j , the attribute cohorts admit an upper bound, $\Theta_j = [0, \bar{\theta}_j]$ and θ_j are continuously distributed over Θ_j according to a distribution function $F_j(\theta)$. For the main result of this section, I assume that each of the distributions $F_j(\theta)$ is log-concave.⁹

In this general case, for each attribute j , the seller needs to design the tariff and the trade expectation functions, $r_j(\theta_j)$, $t_j(\theta_j)$, $X_j(\theta_j)$. This requires incorporating more incentive constraints than in the case of orthogonal types, particularly, the constraints in which the type pretends to be another type within the same cohort and follows the experiment recommendation. Swapping the decision or never purchasing the object remains suboptimal, so those constraints may be omitted.

As before, I approach this problem in a sequence of simplifying steps. First, I invoke the same argument as in the case of orthogonal types to establish that upfront payments are not used in an optimal menu, $r_j(\theta_j) \equiv 0$. Indeed, if some $r_j(\theta_j) > 0$, then the seller can simultaneously decrease $r_j(\theta_j)$ and increase $t_j(\theta_j)$ to keep the expected payment $r_j(\theta_j) + Q_j(\theta_j)t_j(\theta_j)$ the same. This does not affect the revenue or incentive compatibility within each attribute cohort. However, it relaxes incentive constraints between different cohorts.

⁹The class of log-concave distributions includes normal, logistic, exponential, and uniform distributions.

Second, as in the case of a single attribute, I define the trade surplus as

$$Y_j(\theta_j) \triangleq X_j(\theta_j) Q_j(\theta_j) = \int_{b_j(\theta_j)}^{\infty} x_j dG_j(x_j).$$

By construction, $Y_j(\theta_j) \in [0, \mathbb{E}x_j]$. By standard arguments, incentive compatibility within the same cohort implies that $Y_j(\theta_j)$ is nondecreasing in θ_j . It in turn implies that $b_j(\theta_j)$ is nonincreasing; hence, $X_j(\theta_j)$ is nonincreasing and $Q_j(\theta_j)$ is nondecreasing in θ_j . That is, higher types must trade with higher probability but lower expectations.

The seller's problem can be written solely in term of the trade surplus.

Lemma 2. *The seller's problem can be rewritten as*

$$\max_{\{Y_j(\theta_j)\}_{j=1}^J} \sum_{j=1}^J f(\Theta_j) \int_0^{\bar{\theta}_j} \left(\theta_j - \frac{1 - F_j(\theta_j)}{f_j(\theta_j)} \right) Y_j(\theta_j) dF_j(\theta_j) \quad (8)$$

s.t. $Y_j(\theta_j)$ is non - decreasing, $Y_j(\theta_j) \in [0, \mathbb{E}x_j]$,

$$\int_0^{\bar{\theta}_j} Y_j(\theta_j) d\theta_j \geq \bar{\theta}_j \mathbb{E}x_j - \underline{t}(Y_1(\cdot), \dots, Y_J(\cdot)) \quad \forall j.$$

The objective function and the monotonicity constraints capture the incentive-compatibility constraints within each attribute cohort. They are derived by standard one-dimensional arguments. The integral constraints are novel and capture the incentive-compatibility constraints between different cohort. In particular, they require that the highest type within each cohort does not want to purchase the object at the minimal price.

I now argue that the lowest price \underline{t} is offered to highest types $\bar{\theta}_j$. The argument and the result are analogous to that in the case of orthogonal types. Assume in an optimal mechanism some neighborhood of $\bar{\theta}_j$ is not offered the minimal price. Then, these types are not imposing externalities on other cohorts through the integral constraint. Moreover, to not go for the lowest price, these types should be offered some disclosure so $Y_j(\theta_j) < \mathbb{E}x_j$. It leads to a contradiction. The seller could marginally increase $Y_j(\theta_j)$ for these types, improving the revenue but not affecting any other constraint.

This observation allows to state a relaxed problem in which the monotonicity and integral constrains dropped but all high types are required to be offered the same minimal price. If type distributions are log-concave, then the solution to the relaxed problem is a single step function. It satisfies the original constraints and hence solves the original problem. The solution corresponds to only one item per attribute cohort served at the same price. The following theorem summarizes the findings.

Theorem 4. (Optimal Menu, Single-Minded Buyer) *If the buyer is single-minded and type distributions are log-concave, then an optimal menu is characterized by the following properties:*

1. For all j , $\theta_j \in \Theta_j$, $r(\theta_j) = 0$,
2. For all j , $\theta_j \in \Theta_j$, $t(\theta_j) = t$, and
3. For all j , $\theta_j \in \Theta_j$, $E_j(\theta_j) = E_j$ where E_j is a binary monotone partition of x_j .

It is worth pointing out that my analysis provides a partial characterization in the case of general distributions $F_j(\theta_j)$ as well. The first statement remains the same, upfront payments are not used with a single-minded buyer. However, the second and the third statements need to be modified to allow for limited price discrimination. In particular, the arguments of Samuelson (1984) and Bergemann et al. (2017a) can be applied to limit the number of optimal items to two per cohort. That is, the highest types are still offered the lowest price $\forall j, k$, $t(\bar{\theta}_j) = t(\bar{\theta}_k) = \underline{t}$ but per each attribute cohort there could be one more item that targets lower types.

5 Discussion

5.1 Demand Transformation

My analysis highlights that attribute disclosure can be profitably used to modify a demand curve. Consider the example in which there are two attributes $J = 2$, $x_1 \sim U[0, 1]$, $x_2 \sim U[0, 2]$, independently distributed and a continuum of single-minded types. Types $\theta_1 \in \Theta_1$ value only the first attribute and types $\theta_2 \in \Theta_2$ value only the second attribute. Each cohort is equally likely and within each cohort the types are uniformly distributed over $[0, 2]$ so the average type is equal to 1.

If the seller provides no disclosure, then she faces a piecewise-linear demand curve. The optimal price is $2/3$ with the corresponding revenue $1/3$. If the seller provides full disclosure, then the type valuations spread out. The demand decreases for lower prices and increases for higher prices. The overall effect is negative. The full-disclosure optimal price is around 0.82 with the corresponding revenue $0.28 < 1/3$. The left side of Figure 4 illustrates this case.

However, the seller can increase revenue by providing partial disclosure. The type distributions are regular so, by Theorem 4, the optimal mechanism can be implemented by combined partition disclosures followed by a single price. The optimal disclosure thresholds




Figure 4: Attribute disclosure and demand transformation. Left: demand curves under no disclosure and full disclosure. Right: demand curves under no disclosure and optimal disclosure. Vertical lines indicate revenue-maximizing prices.

can be calculated numerically to be $b_1 \simeq 0.27$, $b_2 = 0$. The seller optimally reveals whether the first attribute is above 0.27, yet provides no information about the second attribute. This disclosure targets the types that ex-ante value the object less. As a result, the demand decreases for low prices, increases for medium prices, and remains the same for high prices. The overall effect is positive. The disclosure increases demand even at the no-disclosure optimal price. The optimal-disclosure optimal price is around 0.80 with the corresponding revenue $0.35 > 1/3$. The right side of Figure 4 illustrates.

These findings resonate with the analysis of [Johnson and Myatt \(2006\)](#), who also study the impact of information disclosure on a demand curve. They restrict attention to disclosures that spread type valuations uniformly. Such disclosures translate into *global* rotations of the demand curve. [Johnson and Myatt \(2006\)](#) show that in many settings, the optimal global rotations are extreme and correspond to either no disclosure or full disclosure. No disclosure is associated with a mass market characterized by low price and high demand. Full disclosure is associated with a niche market characterized by high price and low demand.

In contrast, I show that attribute disclosure can rotate the demand curve *locally*. The local rotations correspond to partial disclosures that target specific types. These disclosures can outperform full and no disclosure in both mass and niche markets. I highlight that multiple attributes are required for this result. As shown in Section 3.1, no disclosure remains optimal in a one-dimensional framework.

5.2 Full and Public Disclosure

Esó and Szentes (2007) study discriminatory mechanisms in a “valuation-rank” framework and obtain two main qualitative results. First, they show that full information disclosure is generally optimal. Second, they show that the seller cannot benefit from conditioning the price on the disclosure realization.

In this section, I illustrate that both of these results do not hold in the attribute setting. Consider the example from the introduction in which there are two attributes $J = 2$, $x_1 \sim U[0, 1]$, $x_2 \sim U[0, 2]$, independently distributed. Types are orthogonal; θ_1 values only the first attribute and θ_2 values only the second attribute. The types are equally likely.

An optimal menu can be calculated by Theorem 3. The menu can be implemented by a single disclosure that reveals whether the first attribute is above or below $1/2$ followed by a posted price of $3/4$. This optimal mechanism obtains revenue $9/16$.

The revenue is strictly higher compared to that under full disclosure. In that case, if the seller does not use upfront payments, then the optimal strike price can be calculated to be $2/3$, translating into revenue of $1/3$. The seller can improve it by using upfront payments. The optimal mechanism can be easily calculated to be $r_1 = r_2 = 1/2$, $t_1 = t_2 = 0$. That is, anticipating full disclosure, the seller prefers to not discriminate between the types at all and sell the object in advance through an upfront fee. The corresponding maximal revenue under full disclosure is $1/2$ falling short off optimal $9/16$.

At the same time, the seller can strictly improve the revenue if she could condition the price on the disclosure realization. In fact, she can improve the revenue even under full disclosure. Consider the following mechanism. The seller provides full disclosure, observes the attributes, and chooses the price optimal for the realized valuation distribution. Under this scheme, the seller obtains the revenue,

$$\Pi = \int_0^1 \int_0^2 \frac{1}{2} \max \left\{ \min \{x_1, x_2\}, \frac{\max \{x_1, x_2\}}{2} \right\} dx_1 dx_2 = \frac{29}{48} > \frac{9}{16}.$$

Naturally, observing the disclosure realization reduces the buyer’s informational advantage and helps the seller to better screen the types.

These observations highlight the distinction between the attribute setting and the “valuation-rank” setting of Esó and Szentes (2007). In their framework, the seller informs the buyer about an “orthogonal shock” $\xi(\theta)$, defined as the type’s valuation percentile. By construction, these percentiles are uniformly distributed,

$$\xi(\theta) \sim U[0, 1] \quad \forall \theta \in \Theta.$$

The implicit assumption of the valuation-rank framework is that these shocks are in fact equal each other, $\xi(\theta) \equiv \xi$. This assumption is suitable for environments in which, conditional on the object's state, the buyer's type does not affect his valuation percentile. As [Esó and Szentes \(2007\)](#) discuss, these environments include the cases of additive valuations and Gaussian learning.

However, despite having the same distribution, the shocks $\xi(\theta)$ are generally different random variables. This distinction is particularly prominent in the case of a single-minded buyer with orthogonal types. Different shocks correspond to different attributes and thus are independently distributed. In the example above, these shocks are $\xi(\theta_1) = x_1$ and $\xi(\theta_2) = x_2/2$.

6 Appendix

Proof of Proposition 1. Consider any menu $\mathcal{M} = (\mathbf{E}, \mathbf{r}, \mathbf{t})$. For any type θ , the menu induces the allocation distribution $\mu(\theta) : \mathcal{X} \rightarrow \Delta(A)$, $A = \{\text{"buy"}, \text{"not buy"}\}$, the expected upfront payment $\bar{r}(\theta)$ and the expected strike payment $\bar{t}(\theta)$, conditional on a trade. Consider a direct responsive menu $\mathcal{M}' = (E(\theta), r(\theta), t(\theta))$ with $r(\theta) = \bar{r}(\theta)$, $t(\theta) = \bar{t}(\theta)$, and $E(\theta) = \{A, \mu(\theta)\}$. If all types report truthfully and follow the recommendations, then the menu results in the same allocation distribution and the same expected payments as \mathcal{M} . Reporting truthfully and following the recommendations is incentive-compatible for the types because any deviation under \mathcal{M}' was available to them under \mathcal{M} .

Proof of Theorem 1. The argument is given in the text. The only difference from Myerson (1981)'s problem is that Y can take values in $[0, \mathbb{E}x]$, not in $[0, 1]$. However, it does not affect the bang-bang nature of the solution.

Proof of Theorem 2. The seller's problem can be written in terms of the trade function $q(\theta, x)$, as follows. The seller's objective is

$$\max_{(r(\theta), q(\theta, x), t(\theta))} \int_{\theta \in \Theta} \left(r(\theta) + t(\theta) \int_{x \in \mathcal{X}} q(\theta, x) dG \right) dF(\theta) \quad (9)$$

where the maximization goes over tariff functions and the trade function lies within its domain $q(\theta, x) \in [0, 1]$. The incentive-compatibility constraints can be rewritten in the following way: for all $\theta, \theta' \in \Theta$,

$$\int_{x \in \mathcal{X}} (\theta \cdot x - t(\theta)) q(\theta, x) dG(x) - r(\theta) \geq \int_{x \in \mathcal{X}} (\theta \cdot x - t(\theta')) \sigma_k(q(\theta', x)) dG(x) - r(\theta'), \quad (10)$$

where $k = 1, \dots, 4$ and $\sigma_k(q(\theta', x))$ is equal to $q(\theta', x)$, $1 - q(\theta', x)$, 1 , and 0 respectively. These constraints ensure each type prefers truth telling and following recommendations to any double-deviating strategy. Finally, the individual-rationality constraints can be rewritten as

$$\int_{x \in \mathcal{X}} (\theta \cdot x - t(\theta)) q(\theta, x) dG(x) - r(\theta) \geq 0, \quad \forall \theta \in \Theta. \quad (11)$$

Writing the seller's problem in the form of (9), (10), and (11) allows observing that both the objective function and constraints are linear in the trade probability function $q(\theta, x)$.

All conditions for the Lagrange approach are then satisfied (Luenberger (1969)), and I can apply the Lagrange approach as outlined in the main text. The optimality of linear

disclosure follows.

Proof of Proposition 2. The uniqueness is proven by contradiction. Assume there are two half-spaces A and B such that $\mathbb{E}[x | x \in A] = \mathbb{E}[x | x \in B] = X$ and $\Pr[x \in A \setminus B] > 0$. By the basic properties of conditional expectation, it follows that $\mathbb{E}[x | x \in A \cup B] = X$ and hence $\mathbb{E}[x | x \in A \setminus B] = X$. However, both $A \setminus B$ and B are convex, so $X \in A \setminus B$ yet $X \in B$. Contradiction.

The maximal probability claim is proven by the Lagrange method (Luenberger (1969)). Fix an arbitrary trade expectation $X \in \mathcal{X}$. The maximal trade function solves the following problem:

$$\begin{aligned} \max_{q: \mathcal{X} \rightarrow [0,1]} & \int q(x) dG(x) \\ \text{s.t.} & \int x_j q(x) dG(x) = X_j \int q(x) dG(x), \quad \forall j = 1, \dots, J. \end{aligned}$$

For any solution to this problem, there exist coefficients $\lambda \in \mathbb{R}^J$ such that the solution maximizes the Lagrangian

$$\mathcal{L} = \int q(x) dG(x) + \sum_{j=1}^J \lambda_j \left(\int x_j q(x) dG(x) - X_j \int q(x) dG(x) \right).$$

Rearranging the terms, the Lagrange maximization can be stated as:

$$\max_{\pi: \mathcal{X} \rightarrow [0,1]} \int \left(1 + \sum_{j=1}^J \lambda_j (x_j - X_j) \right) q(x) dG(x).$$

Maximizing pointwise, the unique solution to this problem (up to a zero probability event) is the linear disclosure

$$q(x) = \begin{cases} 1, & \text{if } \sum_{j=1}^J \lambda_j x_j > \sum_{j=1}^J \lambda_j X_j - 1, \\ 0, & \text{if } \sum_{j=1}^J \lambda_j x_j < \sum_{j=1}^J \lambda_j X_j - 1. \end{cases}$$

By the first part of the lemma, the disclosure is unique and translates into a maximal probability of trade $Q(X)$.

Proof of Lemma 1. (1) and (3): By Proposition 2, I can define $Q(X)$ as the maximal total trade probability given the trade expectation X . Moreover, conditions for the Envelope

Theorem hold, so

$$\frac{\partial Q}{\partial X_j} = \frac{\partial \mathcal{L}}{\partial X_j} = -\lambda_j Q(X),$$

where the Lagrange function is defined in the proof of Proposition 2. At the same time, it must be that

$$\mathbb{E} \left[x_j \mid \lambda_j x_j \geq \sum_{j' \neq j} \lambda_{j'} (X_{j'} - x_{j'}) + \lambda_j X_j \right] = X_j.$$

When $\{x_j\}_{j \in J}$ are independently distributed, the expectation is smaller than $\mathbb{E} x_j$ iff $\lambda_j < 0$ and greater than $\mathbb{E} x_j$ iff $\lambda_j > 0$. Hence, if $X_j < \mathbb{E} x_j$ then $\lambda_j < 0$ and $\partial Q / \partial X_j > 0$. If $X_j > \mathbb{E} x_j$ then $\lambda_j > 0$ and $\partial Q / \partial X_j < 0$.

(2): $Q(\mathbb{E} x) = 1$ achieved by no disclosure. As $X \rightarrow \partial \mathcal{X}$ it must be that the optimal cutoff is arbitrarily close to the boundary. The result follows from continuity of attribute distributions.

Proof of Proposition 3. Consider an arbitrary responsive experiment $E_j(\theta_j)$. For the contradiction, let $X_k(\theta_j) = \mathbb{E}[x_k \mid E, s^+] \neq \mathbb{E} x_k$ for some $k \neq j$. By Lemma 1 there exists a linear disclosure $E'_j(\theta_j)$ such that the relevant expectation remain the same $X'_j(\theta_j) = X_j(\theta_j)$, and all other trade expectations stay at their ex-ante values, $X'_k(\theta_j) = \mathbb{E} x_k$ for all $k \neq j$, and the trade probability is greater $Q(E'_j(\theta_j)) > Q(E_j(\theta_j))$. By uniformly reducing the likelihood of the signal s^+ one can obtain the experiment with the original trade probability $E''_j(\theta_j)$ so that $X''_j(\theta_j) = X_j(\theta_j)$, $Q(E''_j(\theta_j)) = Q(E_j(\theta_j))$, and $X'_k(\theta_j) = \mathbb{E} x_k$ for all $k \neq j$. Replacing $E_j(\theta_j)$ with $E''_j(\theta_j)$ does not change incentive compatibility within cohort Θ_j but relaxes the incentive-compatibility constraints of other cohorts.

Proof of Theorem 3. Consider the seller's problem (7) and an arbitrary solution. Define $\underline{t} = \min_{\theta} \{t(\theta)\}$. This is the price of the incentive-compatibility constraint. For the contradiction, assume $t(\theta) > \underline{t}$ for some type θ .

If $\mathbb{E} x_{j(\theta)} \geq \underline{t}$, then the incentive compatibility constraint is binding. Hence, $Q(\theta) t(\theta) = V(\theta) Q(\theta) - \mathbb{E} x_{j(\theta)} + \underline{t}$. For small $\varepsilon > 0$ consider a modified mechanism with $V'(\theta) = V(\theta) - \varepsilon$, $Q'(\theta) t'(\theta) = V'(\theta) Q'(\theta) - \mathbb{E} x_{j(\theta)} + \underline{t}$. Because Q_{θ} is continuous, the mechanism remains incentive compatible. Since $V(\theta) Q(\theta, V(\theta)) = \int_{\alpha_{\theta}(V)}^{\infty} x_{j(\theta)} dG$ is strictly decreasing in $V(\theta)$, the modified mechanism brings higher revenue. Contradiction.

If $\mathbb{E} x_{j(\theta)} < \underline{t}$, then the individual-rationality constraint is binding. Hence, $t(\theta) = V(\theta)$. For small $\varepsilon > 0$, consider the modified mechanism with $t'(\theta) = t(\theta) - \varepsilon$, $V'(\theta) = V(\theta) - \varepsilon$. The mechanism remains incentive compatible. Again, because $V(\theta) Q(\theta, V(\theta))$ is strictly decreasing in $V(\theta)$, the modified mechanism brings higher revenue. Contradiction.

Now, consider optimal disclosure. By feasibility and individual rationality, $V(\theta) \geq \max\{t, \mathbb{E}x_{j(\theta)}\}$. If $V(\theta) > \max\{t, \mathbb{E}x_{j(\theta)}\}$, then for small $\varepsilon > 0$ the mechanism with $V'(\theta) = V(\theta) - \varepsilon$ is incentive compatible and increases probability of trade, $Q(\theta, V') > Q(\theta, V)$, and consequently, revenue. Contradiction.

Proof of Lemma 2. The seller's problem can be written as

$$\begin{aligned} & \max_{\{r_j(\theta_j), X_j(\theta_j), t_j(\theta_j)\}} \sum_{j=1}^J f(\Theta_j) \int_{\theta_j \in \Theta_j} (r_j(\theta_j) + Q_j(\theta_j) t_j(\theta_j)) dF_j(\theta_j) \\ \text{s.t. } & (\theta_j X_j(\theta_j) - t_j(\theta_j)) Q_j(\theta_j) - r_j(\theta_j) \geq (\theta_j X_j(\theta'_j) - t_j(\theta'_j)) Q_j(\theta'_j) - r_j(\theta'_j), \quad \forall j, \theta_j, \theta'_j \in \Theta_j, \\ & (\theta_j X_j(\theta_j) - t_j(\theta_j)) Q_j(\theta_j) - r_j(\theta_j) \geq \mathbb{E}x_j - t_k(\theta_k) - r_k(\theta_k), \quad \forall j, k, \theta_j \in \Theta_j, \theta_k \in \Theta_k, \\ & (\theta_j X_j(\theta_j) - t_j(\theta_j)) Q_j(\theta_j) - r_j(\theta_j) \geq 0, \\ & X_j(\theta_j) \geq \mathbb{E}x_j, \quad Q_j(\theta_j) = Q_j(X_j(\theta_j)), \quad \forall j, \theta_j \in \Theta_j. \end{aligned}$$

Define the expected transfer function

$$T_j(\theta_j) \triangleq Q_j(\theta_j) t_j(\theta_j).$$

I can then use standard one-dimension arguments within each cohort to establish the connection between the trade surplus and the expected transfer function. Incentive compatibility requires the slope of the indirect utility function $V'_j(\theta_j)$ be equal to $Y_j(\theta)$ almost everywhere. Hence, the indirect function is convex and, by the Envelope Theorem, the optimal transfers can be recovered to be

$$T_j(\theta_j) = \theta_j Y_j(\theta_j) - \int_0^{\bar{\theta}_j} Y_j(z) dz.$$

Individual rationality and incentive compatibility are satisfied by construction. However, incentive compatibility between different cohorts imposes one additional constraint,

$$U(E_j(\bar{\theta}_j), \bar{\theta}_j) = \int_0^{\bar{\theta}_j} Y_j(\theta_j) d\theta_j \geq \bar{\theta}_j \mathbb{E}x_j - \underline{t},$$

where \underline{t} is the minimal strike price in the menu determined by $\{Y_j\}_{j=1}^J$. The deviations from all other types $\theta_j \in \Theta$ follow because $V_j(\theta_j)$ is convex and grows slower than $\theta_j \mathbb{E}x_j$. Applying double integration to the objective function completes the derivation.

Proof of Theorem 4. The argument in the text establishes that all high types are offered the minimal price. The optimal mechanism should then solve the problem (8) with the

additional constraints that all high types are offered the same fixed price \underline{t}^* , and are served the fixed trade surplus $Y_j^*(\bar{\theta}_j)$. These constraints can be written as:

$$\int_0^{\bar{\theta}_j} Y_j(\theta_j) d\theta_j = Y_j(\bar{\theta}_j) - \underline{t}^* Q_j(Y_j(\bar{\theta}_j)),$$

$$Y_j(\bar{\theta}_j) = Y_j^*(\bar{\theta}_j).$$

As discussed in the text, I can consider a relaxed problem with the original integral constraints and the monotonicity constraints dropped. In this relaxed problem, by [Luenberger \(1969\)](#), there exist the Lagrange multipliers $\{\lambda_j\}$ such that optimal $Y_j(\theta_j)$ maximize the Lagrange function

$$\mathcal{L} \sim \sum_{j=1}^J f(\Theta_j) \int_0^{\bar{\theta}_j} \left(\theta_j - \frac{1 - F_j(\theta_j) - \lambda_j}{f_j(\theta_j)} \right) Y_j(\theta_j) dF_j(\theta_j)$$

over a domain $Y_j(\theta_j) \in [0, Y_j^*(\bar{\theta}_j)]$. If all type distributions are log-concave, then the integrand is increasing in θ_j . Hence, the optimal $Y_j(\theta_j)$ are bang-bang: $Y_j(\theta_j) = 0$ for $\theta_j < \theta_j^*$, $Y_j(\theta_j) = Y_j^*(\bar{\theta}_j)$ for $\theta_j > \theta_j^*$. This relaxed solution corresponds to a single item per each attribute cohort. Hence, the relaxed constraints are satisfied and the solution solves the original problem as well.

References

- ANDERSON, S. P. AND R. RENAULT (2006): “Advertising Content,” *American Economic Review*, 96, 93–113.
- BERGEMANN, D., A. BONATTI, AND A. SMOLIN (2017a): “The Design and Price of Information,” *American Economic Review*, forthcoming.
- BERGEMANN, D., F. CASTRO, AND G. Y. WEINTRAUB (2017b): “The Scope of Sequential Screening with Ex-Post Participation Constraints,” *Discussion paper*.
- BERGEMANN, D. AND M. PESENDORFER (2007): “Information Structures in Optimal Auctions,” *Journal of Economic Theory*, 137, 580–609.

- BERGEMANN, D., J. SHEN, Y. XU, AND E. YEH (2012): “Multi-Dimensional Mechanism Design with Limited Information,” in *Proceedings of the 13th ACM Conference on Electronic Commerce*, ACM, 162–178.
- COURTY, P. AND H. LI (2000): “Sequential Screening,” *Review of Economic Studies*, 67, 697–717.
- DASKALAKIS, C., A. DECKELBAUM, AND C. TZAMOS (2017): “Strong Duality for a Multiple-Good Monopolist,” *Econometrica*, 85, 735–767.
- ESÓ, P. AND B. SZENTES (2007): “Optimal Information Disclosure in Auctions and the Handicap Auction,” *Review of Economic Studies*, 74, 705–731.
- (2017): “Dynamic Contracting: An Irrelevance Theorem,” *Theoretical Economics*, 12, 109–139.
- HEUMANN, T. (2017): “Optimal Sequential Screening with Ex Post Participation Constraint,” *Discussion paper*.
- JOHNSON, J. P. AND D. P. MYATT (2006): “On the Simple Economics of Advertising, Marketing, and Product Design,” *American Economic Review*, 96, 756–784.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KRÄHMER, D. (2017): “Information Disclosure and Monopolistic Screening,” *Discussion paper*.
- KRÄHMER, D. AND R. STRAUZ (2015): “Optimal Sales Contracts with Withdrawal Rights,” *Review of Economic Studies*, 82, 762–790.
- LANCASTER, K. J. (1966): “A New Approach to Consumer Theory,” *Journal of Political Economy*, 74, 132–157.

- LEHMANN, D., L. I. O'CALLAGHAN, AND Y. SHOHAM (2002): "Truth Revelation in Approximately Efficient Combinatorial Auctions," *Journal of the ACM*, 49, 577–602.
- LEWIS, T. R. AND D. E. SAPPINGTON (1994): "Supplying Information to Facilitate Price Discrimination," *International Economic Review*, 309–327.
- LI, H. AND X. SHI (2017): "Discriminatory Information Disclosure," *American Economic Review*, forthcoming.
- LUENBERGER, D. G. (1969): *Optimization by Vector Space Methods*, John Wiley & Sons.
- MATEJKA, F. AND A. MCKAY (2014): "Rational Inattention to Discrete Choices: A New Foundation for The Multinomial Logit Model," *American Economic Review*, 105, 272–298.
- MYERSON, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.
- (1982): "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems," *Journal of Mathematical Economics*, 10, 67–81.
- RAYO, L. AND I. SEGAL (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118, 949–987.
- SAMUELSON, W. (1984): "Bargaining under Asymmetric Information," *Econometrica*, 995–1005.