

# Selling Experiments\*

Dirk Bergemann<sup>†</sup>      Alessandro Bonatti<sup>‡</sup>      Alex Smolin<sup>§</sup>

July 23, 2015

## Abstract

A monopolist sells informative experiments to heterogeneous buyers who face a decision problem. Buyers differ in their prior information, and hence in their willingness to pay for additional signals. The monopolist can profitably offer a menu of experiments. Even under costless acquisition and degrading of information, the optimal menu is quite coarse. The seller offers at most two experiments, and we derive conditions under which flat pricing (one experiment) or discriminatory pricing (two experiments) is optimal.

**Keywords:** selling information, experiments, mechanism design, price discrimination, product differentiation.

**JEL Codes:** D42, D82, D83.

---

\*We thank Ben Brooks, Giacomo Calzolari, Gabriel Carroll, Gonzalo Cisternas, Jeff Ely, Emir Kamenica, Alessandro Lizzeri, Phil Reny, Mike Riordan, Maher Said, Juuso Toikka, Alex Wolitzky and seminar participants at Bocconi, Bologna, Chicago, Harvard, Mannheim, NYU, Toulouse, Vienna, and Yale for helpful comments.

<sup>†</sup>Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA, [dirk.bergemann@yale.edu](mailto:dirk.bergemann@yale.edu).

<sup>‡</sup>MIT Sloan School of Management, 100 Main Street, Cambridge MA 02142, USA [bonatti@mit.edu](mailto:bonatti@mit.edu).

<sup>§</sup>Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA, [alexey.smolin@yale.edu](mailto:alexey.smolin@yale.edu).

# 1 Introduction

We consider a monopolist who wishes to sell information about a payoff-relevant variable (the “state”) to a single buyer. The buyer faces a decision problem, and the seller has access to all the information that is relevant for solving it. In addition, the buyer is partially and privately informed about the relevant state variable. There are several possible ways to model demand for information. For example, the buyer’s private information may concern his beliefs over the state, the importance of the decision problem (the “stakes of the game”), or his preferences over certain actions. Within this context, we investigate the revenue-maximizing policy for the seller. How much information should the seller provide? And how should the seller price the access to the database?

We are initially motivated by the role of information in markets for online advertising. In that context, advertisers can tailor their spending to the characteristics of individual consumers. Large data holders compile databases with consumers’ browsing and purchasing history. Advertisers are therefore willing to pay in order to acquire information about each consumer’s profile. A contract between a buyer and a seller of data then specifies which consumer-specific attributes the seller shall release to the advertiser before any impressions are purchased. An alternate example is given by vendors of information about specific financial assets. In that case, the buyer could be an investment manager, who wishes to acquire a long or short position on a stock, based on its underlying fundamentals.

In all these examples, the monopolist sells information. In particular, the products being offered are experiments à la Blackwell, i.e. signals that reveal information about the buyer’s payoff-relevant state. As the buyer is partially informed, the value of any experiment depends on his type. The seller’s problem is then to screen buyers with heterogeneous private information by offering a menu of experiments. In other words, the seller’s problem reduces to the optimal *versioning* of information products.<sup>1</sup>

A long literature in economics and marketing has focused on the properties of information goods. This literature emphasizes how digitalized production and low marginal costs allow sellers to easily degrade (more generally, to customize) the attributes of such products (Shapiro and Varian, 1999). This argument applies even more forcefully to information products, i.e. experiments, and makes *versioning* an attractive price-discrimination technique (Sarvary, 2012). In this paper, we investigate the validity of these claims in a simple contracting environment.

---

<sup>1</sup>For example, consider the Consumer Sentiment Index released by Thomson-Reuters and the University of Michigan (see <http://www.sca.isr.umich.edu>), or PAWWS Financial Network’s portfolio accounting system (Shapiro and Varian, 1999), which were initially available in different versions, based on the timing of their release.

**Environment** The seller does not know the true state, but she can design any experiment ex-ante. The seller’s problem is therefore to design and to price an “information product line” to maximize expected profits. An information product line consists of a menu of experiments, and we characterize the optimal menu for the seller in this environment. To our knowledge, this is the first paper to analyze a seller’s problem of optimally “packaging” information in different versions.

The distinguishing feature of our approach to pricing information is that payments cannot be made contingent on the buyer’s actions or on the realized states. Consequently, the value of an information product for a given buyer can be computed independently of its price. This is in contrast with a contract specifying contingent payments for actions, where the marginal price influences the buyer’s behavior after observing a signal and hence his willingness to pay ex-ante. (Clearly, this also leaves open the question of how much more can be achieved in terms of profits in a richer contracting environment where actions and states are contractible.)

Finally, despite the buyer being potentially informed about his private beliefs, the analysis differs considerably from a belief-elicitation problem. Instead, we cast the problem into the canonical quality-pricing framework where the buyer’s demand for information is determined by his prior beliefs or tastes. However, important differences with the standard setting (Mussa and Rosen, 1978) will emerge due to the properties of information products.

**Results** Because information is only valuable if it affects the decision maker’s action, buyers with heterogeneous beliefs and tastes rank experiments differently. More precisely, all buyer types agree on the highest-value information structure (i.e., the perfectly informative experiment), but their ranking of *distorted* information structures differs substantially.

This peculiar property of information products induces a trade-off for the seller between the precision of an experiment (vertical quality) and its degree of targeting (horizontal positioning). At the same time, the asymmetry in buyers’ valuations allows the seller to extract more surplus. We formalize this intuition by characterizing the profit-maximizing nonlinear pricing scheme. The optimal menu for the seller does, in fact, exploit this asymmetry, and is consequently richer than in Mussa and Rosen (1978).

At the same time, bundling information is optimal quite generally. For distributions satisfying strong regularity conditions, the seller adopts flat pricing of the fully informative experiment. Thus, because of the linearity of expected utility in probabilities, we come to an intuition analogous to Riley and Zeckhauser (1983). However, when types correspond to interim beliefs (e.g., following observation of a private signal), it is quite natural to consider bimodal densities (e.g., as informative signals push mass to the tails). Therefore, the analysis of ironing is an economically relevant exercise, and not only a technical curiosity. When

ironing is required, we find that discriminatory pricing emerges naturally as part of the optimal menu.

Even in environments where virtual values are linear in the allocation, the seller can exploit differently informative signals. Thus, unlike in Myerson (1981) or Riley and Zeckhauser (1983), the seller can offer more than just the maximally informative experiment at a flat price. In particular, the optimal menu consists of (at most) two experiments: one is fully informative; and the other (if present) contains one signal that perfectly reveals one realization of the buyer’s underlying state. This property is best illustrated in a binary-type model, but holds more generally any time a seller has the ability to version its product along more than one dimension.

**Related Literature** This paper is tied to the literature on selling information. It differs substantially from classic papers on selling financial information (Admati and Pfleiderer, 1986, 1990), as well as from the more recent contributions of Esó and Szentes (2007b) and Hörner and Skrzypacz (2015). In our earlier work on markets for data and online advertising (Bergemann and Bonatti, 2015), we examined the problem of selling consumer-level information (e.g., as encoded in third-party cookies). Specifically, we considered the problem of a data provider that sells queries about individual consumers’ characteristics, and derived the optimal (linear) price for each query (“cookie”).

Our approach also differs substantially from models of disclosure. In these models, the seller of a good discloses horizontal match-value information, in addition to setting a price. Several papers, among which Ottaviani and Prat (2001), Johnson and Myatt (2006), Esó and Szentes (2007a), Bergemann and Pesendorfer (2007), and Li and Shi (2013), have analyzed the problem from an ex-ante perspective. In these papers, the seller commits (simultaneously or sequentially) to a disclosure rule and to a pricing policy. More recent papers, among which Balestrieri and Izmalkov (2014), Celik (2014), Koessler and Skreta (2014), and Mylovanov and Tröger (2014) take an informed-principal perspective. Lizzeri (1999) considers vertical information acquisition and disclosure by a monopoly intermediary, and Abraham, Athey, Babaioff, and Grubb (2014) study vertical information disclosure in an auction setting.

Finally, commitment to a disclosure policy is also present in the literature on Bayesian persuasion, e.g. Rayo and Segal (2010) and Kamenica and Gentzkow (2011). However, these papers differ from our mainly because of (i) lack of transfers, and (ii) the principal derives utility directly from the agent’s action.

## 2 Model

We consider a model with a single agent (a buyer of information) facing a decision problem. We maintain throughout the paper the assumption that the buyer must choose between two actions.

$$a \in A = \{a_L, a_H\}.$$

In this section, we assume the relevant state for the buyer’s problem is also binary,

$$\omega \in \Omega = \{\omega_L, \omega_H\}.$$

The buyer’s objective is to match the state. In our applications, an advertiser wishes to purchase impressions only to consumers with a high match value; an investor wants to take a short or long position depending on the underlying asset’s value; and a manager wants to adopt the right business strategy.

We will consider for now a fully symmetric environment, and let the buyer’s ex-post utility  $u(a, \omega)$  from taking action  $a$  in state  $\omega$  be given by

$u(a, \omega)$	$a = a_L$	$a = a_H$
$\omega = \omega_L$	1	0
$\omega = \omega_H$	0	1

Note that with only two actions, it is without loss to assume that the state  $\omega$  equals the buyer’s payoff from taking the “high action”  $a_H$ , net of the payoff from choosing  $a_L$ .

The buyer’s private type is his interim belief  $\theta \in [0, 1]$ . This belief can be thought of as obtained from the private observation of an informative signal about the state. Therefore, the type space is a subset of the unit interval. We denote the buyer’s interim belief by

$$\theta = \Pr(\omega = \omega_H).$$

Thus, while we consider binary states and actions, we allow for a continuum of types for the consumer.<sup>2</sup>

The distribution of interim beliefs  $F(\theta)$  is common knowledge to the buyer and the seller, who share a common prior belief over the state

$$\mu = \mathbb{E}_F[\theta].$$

---

<sup>2</sup>In order to interpret the model as a continuum of buyers, we shall assume that states  $\omega$  are identically and independently distributed across buyers, and that buyers’ private signals are conditionally independent.

An experiment (i.e., an information structure)  $I \in \mathcal{I}$  consists of a set of signals and a likelihood function mapping states into signals.

$$I = \{S, \pi\} \quad \pi : \Omega \rightarrow \Delta S.$$

Signals are conditionally independent from the buyer's private information.

A strategy for the seller consists of a menu of experiments and associated tariff  $\mathcal{M} = \{\mathcal{I}, t\}$ , with

$$\mathcal{I} = \{I\} \quad t : \mathcal{I} \rightarrow \mathbb{R}^+.$$

With the conditional independence assumption, we are adopting the interpretation of a buyer querying a database, or request a diagnostic service. In particular, the buyer and the seller draw their information from independent sources. Under this interpretation, the seller does not know the realized state  $\omega$  at the time of contracting. The seller can, however, augment the buyer's original information with arbitrarily precise signals.

For instance, with the online advertising application in mind, the buyer is privately informed about the average returns to advertising. The seller can, however, improve the precision of his estimate consumer by consumer. The two parties therefore agree to a contract by which the seller discloses specific attributes of individual consumers upon the buyer's request. Thus, even if the seller is already endowed with a complete database, she does not know the realized state of the actual buyer at the time of contracting.

Finally, note that we distinguish the cost of acquisition of information (i.e. building the database) from duplication and distribution of the information, which we assume costless. The analysis could easily extended to a first stage where the seller invests in the maximal precision of the experiments, and to fixed or linear costs of information distribution.

To conclude the description of the model, we summarize the timing of the game: (i) the buyer observes an initial signal, and forms his interim belief  $\theta$ ; (ii) the seller offers a menu of experiments  $\mathcal{M}$ ; (iii) the buyer chooses an experiment  $I$ , and pays the corresponding price  $t$ ; (iv) the buyer observes a signal  $s$  from the experiment  $I$  (given the true state  $\omega$ ); and finally (v) the buyer chooses an action  $a$ .

## 3 The Seller's Problem

### 3.1 Buyer's Utility

We begin by defining the demand for information of each buyer type. Let  $u(\theta)$  denote the buyer's payoff under partial information

$$u(\theta) \triangleq \max_{a \in A} \mathbb{E}_\theta [u(a, \omega)].$$

The interim value of experiment  $I$  for type  $\theta$  is then equal to the net value of augmented information,

$$V(I, \theta) \triangleq \mathbb{E}_{I, \theta} [\max_{a \in A} \mathbb{E}_{s, \theta} [u(a, \omega)]] - u(\theta).$$

We now characterize the menu of experiments that maximizes the seller's profits. Because the Revelation Principle applies to this setting, we may state the seller's problem as designing a direct mechanism

$$\mathcal{M} = \{I(\theta), t(\theta)\}.$$

that assigns an experiment to each type of the buyer. The seller's problem consists of maximizing the expected transfers subject to incentive compatibility and individual rationality:

$$\begin{aligned} & \max_{\{I(\theta), t(\theta)\}} \int t(\theta) dF(\theta), \\ \text{s.t. } & V(I(\theta), \theta) - t(\theta) \geq V(I(\theta'), \theta) - t(\theta') \quad \forall \theta, \theta', \\ & V(I(\theta), \theta) - t(\theta) \geq 0 \quad \forall \theta. \end{aligned}$$

The seller's problem can be immediately simplified by taking advantage of the binary-action framework. In particular, we can reduce the set of optimal experiments to a very tractable class.

**Lemma 1** (Binary Signals). *Every experiment in an optimal menu consists of two signals only.*

The intuition for this result is straightforward: suppose the seller were to offer experiments with more than two signals; she could then combine all signals in experiment  $I(\theta)$  that lead to the same choice of action for type  $\theta$ ; clearly, the value of this experiment  $V(I(\theta), \theta)$  stays constant for type  $\theta$  (who does not modify his behavior); in addition, because the original signal is strictly less informative than the new one,  $V(I(\theta), \theta')$  decreases (weakly) for all  $\theta' \neq \theta$ .

Lemma 1 allows us to focus on experiments with two signals only. However, we are still

dealing with a rich type space. This means the seller designs a menu of experiments that differ in the informativeness about each state. We may represent each experiment as follows:

$$I(\theta) = \begin{array}{c|cc} & s_L & s_H \\ \hline \omega_L & \beta(\theta) & 1 - \beta(\theta) \\ \omega_H & 1 - \alpha(\theta) & \alpha(\theta) \end{array} .$$

Throughout the paper, we adopt the convention that  $\alpha(\theta) + \beta(\theta) \geq 1$  (else we should relabel the signals  $s_L$  and  $s_H$ ). We shall also refer to the difference in the conditional probabilities  $\alpha$  and  $\beta$  as the relative informativeness of an experiment.

We now derive the value of an arbitrary experiment. In particular, the value of experiment  $(\alpha, \beta)$  for type  $\theta$  is given by

$$V(\alpha, \beta, \theta) = [(\alpha - \beta)\theta + \beta - (1 - \theta) - \max\{0, 2\theta - 1\}]^+ . \quad (1)$$

While we are considering a natural type space, corresponding to the buyer's interim beliefs, the notions of a "high" and "low" type differ from the standard screening setting. In particular, due to the nature of the buyer's Bayesian problem, the most valuable type for the seller is the middle type  $\theta = 1/2$ . Conversely, the two extreme types  $\theta \in \{0, 1\}$  have no value of information.<sup>3</sup>

Figure 1 shows the value of information for two particular experiments.




Figure 1: Value of Information Structures  $(\alpha, \beta)$

The value of information (1) includes both level effects (terms depend on the allocation or type only) and interaction effects. In particular, the allocation and the buyer's type

---

<sup>3</sup>Buyer types with degenerate beliefs do not expect any contradictory signals to occur, and hence they are not willing to pay for such experiments. More generally, because contracting takes place ex-ante, the seller's profit does not depend on whether the buyer holds correct beliefs.



interact only through the difference in the experiment's relative informativeness  $\alpha - \beta$ . This is clear from Figure 1. A more optimistic type has a relatively higher value for experiments with a high  $\alpha$  because such experiments contain a signal that perfectly reveals the low state. Because this induces types  $\theta > 1/2$  to switch their action (compared to the status quo), these types have a positive value of information for any experiment with  $\alpha = 1$ .

Perhaps more importantly, the specific interaction of type and allocation in the buyer's utility means that the seller can increase the value of an experiment at the same rate for all types. In particular, increasing  $\alpha$  and  $\beta$  holding  $\alpha - \beta$  constant, and increasing the price at the same rate, the seller does not alter the attractiveness of the experiment for any buyer who is considering choosing it.

The next result allows us to further simplify the seller's optimal strategy.

**Lemma 2** (Partially Revealing Signals). *Every experiment in an optimal menu has  $\alpha = 1$  or  $\beta = 1$ .*

In other words, at least one signal perfectly reveals one state in any experiment part of an optimal menu. This result implies the allocation rule is one-dimensional, with

$$q(\theta) \triangleq \alpha(\theta) - \beta(\theta) \in [-1, 1]$$

measuring the relative informativeness of the experiment.

With this notation, two distinct information structures  $q \in \{-1, 1\}$  correspond to releasing no information to the buyer. (These are the two experiments that show the same signal with probability one.) We should also point out that (because of Lemma 2), a negative value of  $q$  implies  $\beta = 1$  and a positive  $q$  implies  $\alpha = 1$ . The fully informative experiment is given by  $q = 0$ . We summarize all optimal experiments in the tables below.

$$\begin{array}{c}
 I = \begin{array}{c|cc} & s_L & s_H \\ \hline \omega_L & 1 & 0 \\ \omega_H & -q & 1+q \end{array} & 
 I = \begin{array}{c|cc} & s_L & s_H \\ \hline \omega_L & 1-q & q \\ \omega_H & 0 & 1 \end{array} \\
 -1 \leq q \leq 0 & 0 \leq q \leq 1
 \end{array}$$

We may then rewrite the value of experiment  $q \in [-1, 1]$  for type  $\theta \in [0, 1]$  as follows:

$$V(q, \theta) = [\theta q - \max\{q, 0\} + \min\{\theta, 1 - \theta\}]^+. \quad (2)$$

It can be useful, at this point, to pause and discuss similarities between our demand function and those obtained in a traditional screening model (i.e., when the seller offers a physical

good).

All buyer types value the vertical “quality” of information structures, as measured by their participation constraint (which is reflected in the min term). Note, however, that the utility function  $V(q, \theta)$  has the single-crossing property in  $(\theta, q)$ . This indicates that buyers who are relatively more optimistic about the *high* state  $\omega_H$  assign a relatively higher value to information structures with a high  $q$ . In particular, very optimistic types have a positive willingness to pay for experiments with  $\alpha = 1$  because such experiments contain signals that perfectly reveal the *low* state  $\omega_L$ .

Figure 2 shows the value assigned to different experiments  $q$  by two types that are symmetric about  $1/2$ . The solid line represents  $V(q, 1/3)$ , while the dashed line represents  $V(q, 2/3)$ . Consistent with the earlier discussion, a buyer who assigns probability  $1/3$  to state  $\omega = 1$  is not willing to pay a positive price for experiments with a high  $q$ . Those experiments fully reveal state  $\omega = 0$ , in which case the buyer takes his default action  $a = 0$ , and are not sufficiently informative about state  $\omega = 1$  to induce the buyer to switch his action.




Figure 2: Value of Information Structures

Though the seller’s problem is reminiscent of classic nonlinear pricing, we uncover a novel aspect of horizontal differentiation. This feature is linked to the relative informativeness of an experiment. Furthermore, the quality and “positioning” of an information product cannot be chosen separately by the firm. The information nature of the good induces a technological constraint (which given by the formula for  $q$ ) that limits the asymmetric informativeness of an experiment, holding constant its quality level. In other words, it is difficult to imagine a non-information analog for our demand function.

To summarize, we present a canonical model for selling information that nonetheless differs from existing screening models along several dimensions: (a) buyers have type-dependent participation constraints; (b) experiment  $q = 0$  is the most valuable for all types; (c) a specific buyer type ( $\theta = 1/2$ ) always has the highest pay for any information structure; (d) buyers are horizontally differentiated with respect to the relative informativeness of experiments.

### 3.2 Incentive Compatibility

We now use the structure of the problem in order to derive a characterization of implementable allocations  $q(\theta)$ . In particular, the buyer's utility function in (2) has a downward kink in  $\theta$ . As discussed earlier, this follows from having an interior type assign the highest value to any allocation (and from the linearity of the buyer's problem).

Therefore, we compute the buyer's rents  $U(\theta)$  on  $[0, 1/2]$  and  $[1/2, 1]$  separately. We first recognize that the buyer's rents will be non-decreasing on the first subinterval and non-increasing on the second. Thus, the participation constraint will bind at  $\theta = 0$  and  $\theta = 1$ , if anywhere. Furthermore, types  $\theta = 0$  and  $\theta = 1$  have no value for any experiment, and must therefore receive the same utility.

We then apply the envelope theorem to each subinterval separately, and obtain two different expressions for the rent of type  $\theta = 1/2$ . Continuity of the rent function then implies

$$U(1/2) = U(0) + \int_0^{1/2} V_\theta(q, \theta) d\theta = U(1) - \int_{1/2}^1 V_\theta(q, \theta) d\theta.$$

While any type's utility can always be written in this form, the novel element of our model is that no further endogenous variables appear. For instance, in Mussa and Rosen (1978), the rent of the highest type  $U(1)$  depends on the allocation itself. This is not so in our context as a consequence of having two extreme types with zero value of information,  $U(0) = U(1)$ . Differentiating (2) and simplifying, we can express the above equation as

$$U(1/2) = \int_0^{1/2} (q(\theta) + 1) d\theta = - \int_{1/2}^1 (q(\theta) - 1) d\theta.$$

This is a key equation for the paper, which sets it apart from most other screening problems. As a consequence of the nature of the buyer's private information (i.e. his beliefs), incentive compatibility imposes an aggregate (integral) constraint on the allocation. We formalize this in the following result.

**Lemma 3** (Implementable Allocations). *The allocation  $q(\theta)$  is implementable if and only if*

$$q(\theta) \in [-1, 1] \text{ is non-decreasing,}$$

$$\text{and } \int_0^1 q(\theta) d\theta = 0.$$

The integral constraint is a requirement for implementability. As such it is not particularly meaningful to analyze the relaxed problem. This is in contrast with other instances of screening under integral constraints (e.g., constraints on transfers due to budget or enforceability, or capacity constraints). Finally, the resemblance to a persuasion budget constraint is purely cosmetic.

We can now state the seller's problem, and give its solution in the next section. (In the Appendix we characterize the transfers associated with allocation rule  $q(\theta)$  in the usual way.)

$$\max_{q(\theta)} \int_0^1 \left[ \left( \theta + \frac{F(\theta)}{f(\theta)} \right) q(\theta) - \max\{q(\theta), 0\} \right] f(\theta) d\theta, \quad (3)$$

$$\text{s.t. } q(\theta) \in [-1, 1] \text{ non-decreasing,}$$

$$\int_0^1 q(\theta) d\theta = 0.$$

## 4 Optimal Menu

We now fully solve the seller's problem (3) for the binary-state case. It can be useful to rewrite the objective with the density  $f(\theta)$  explicitly in each term:

$$\int_0^1 [(\theta f(\theta) + F(\theta)) q(\theta) - \max\{q(\theta), 0\} f(\theta)] d\theta.$$

This minor modification highlights two important features of our problem: (i) the constraint and the objective have generically different weights,  $d\theta$  and  $dF(\theta)$ ; and hence (ii) the problem is non separable in the type and the allocation, which interact in two different terms.

We therefore must consider the “virtual values” for each allocation  $q$  separately,

$$\phi(\theta, q) := \begin{cases} \theta f(\theta) + F(\theta) & \text{for } q < 0, \\ (\theta - 1)f(\theta) + F(\theta) & \text{for } q > 0. \end{cases}$$

The function  $\phi(\theta, q)$  takes on two values only due to the piecewise-linear objective function. The two values represent the marginal benefit to the seller (ignoring the constraint) of

increasing each type's allocation from  $-1$  to  $0$ , and from  $0$  to  $1$ , respectively.

We now let  $\lambda$  denote the multiplier on the integral constraint, and define the *ironed* virtual value for experiment  $q$  as  $\bar{\phi}(\theta, q)$ . Following Toikka (2011), we say that the allocation satisfies the *pooling property* if it is constant on any interval where the ironed virtual value is constant. We can then reduce the seller's problem to the following maximization program.

**Proposition 1** (Optimal Allocation Rule).

*The allocation  $q^*(\theta)$  is optimal if and only if there exists  $\lambda^* > 0$  s.t.  $q^*(\theta)$  solves*

$$\max_{q \in [-1, 1]} \left[ \int_{-1}^q (\bar{\phi}(\theta, x) - \lambda^*) dx \right] \text{ for all } \theta,$$

*has the pooling property, and satisfies the integral constraint in Lemma 3.*

The solution to the seller's problem is then obtained by combining standard Lagrange methods with the ironing procedure developed by Toikka (2011) that extends the approach of Myerson (1981). In particular, Proposition 1 provides a characterization of the general solution, and suggests an algorithm to compute it.

To gain some intuition for the shape of the solution, observe that the problem is piecewise-linear (but concave) in the allocation. Thus, absent the integral constraint, the seller would choose an allocation that takes values at the kinks, i.e.  $q^*(\theta) \in \{-1, 0, 1\}$  for all  $\theta$ . In other words, the seller would offer a one-experiment menu consisting of a flat price for the complete-information structure. It will indeed be optimal for the seller to adopt flat pricing in a number of circumstances. The main novel result of this section is that the seller can (sometimes) do better by offering one additional experiment.

**Proposition 2** (Optimal Menu).

*An optimal menu consists of at most two experiments.*

1. *The first experiment is fully informative.*
2. *The second experiment (contains a signal that) perfectly reveals one state.*

We now separately examine the solution under flat and discriminatory pricing.

## 4.1 Flat Pricing

We illustrate the procedure in an example where ironing is, in fact, not required. Let  $F(\theta) = \sqrt{\theta}$ , and consider the virtual values  $\phi(\theta, q)$  for  $q < 0$  and  $q \geq 0$  separately. The allocation that maximizes the expected virtual surplus in Proposition 1 assigns  $q^*(\theta) = -1$

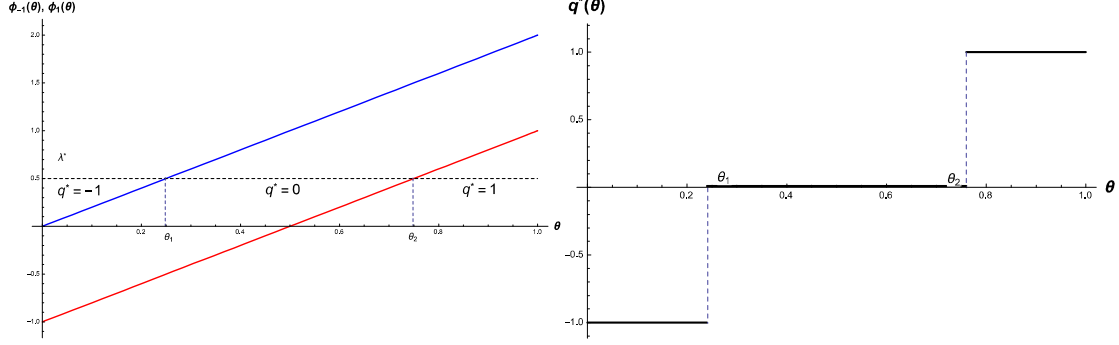


Figure 3: Uniform Distribution: Virtual Values and Optimal Allocation

to all types  $\theta$  for which  $\phi(\theta, -1)$  falls short of the multiplier  $\lambda$ ; it assigns  $q^*(\theta) = 0$  to all types  $\theta$  for which  $\phi(\theta, -1) > \lambda > \phi(\theta, 1)$ ; and  $q^*(\theta) = 1$  for all types  $\theta$  for which  $\phi(\theta, 1) > \lambda$ .

Figure 3 considers the virtual values and multiplier  $\lambda^*$ . Figure 3 illustrates the resulting allocation rule. In order to satisfy the constraint, optimal value of the multiplier  $\lambda^*$  must identify two symmetric threshold types  $(\theta_1, \theta_2)$  that separate types receiving the efficient allocation  $q = 0$  from those receiving no information at all,  $q = -1$  or  $q = 1$ . It is then clear that, if virtual values are strictly increasing, the optimal menu is given by charging the monopoly price for the fully informative experiment.

The one-experiment result holds under weaker conditions than increasing virtual values. We now derive sufficient conditions under which the solution  $q^*$  takes values in  $\{-1, 0, 1\}$  only, i.e., conditions for the optimality of flat pricing.

**Proposition 3** (Flat Pricing).

*Suppose any of the following conditions hold:*

1.  $F(\theta) + \theta f(\theta)$  and  $F(\theta) + (\theta - 1)f(\theta)$  are strictly increasing;
2. the density  $f(\theta) = 0$  for all  $\theta > 1/2$  or  $\theta < 1/2$ ;
3. the density  $f(\theta)$  is symmetric around  $\theta = 1/2$ .

*The optimal menu contains only the fully informative experiment ( $q^* \equiv 0$ ).*

An implication of Proposition 3 is that the seller offers a second experiment *only if* ironing is required. At the same time, there exist examples with non-monotone virtual values and one-item menus. Symmetric distributions are one such instance: for any distribution function  $F(\theta)$ , e.g. hazard rate, the solution to the restricted problem on  $[0, 1/2]$  or  $[1/2, 1]$  is a cutoff policy. Because the cutoffs under a symmetric distributions are symmetric about  $1/2$ , it follows that the solutions to the two subproblems satisfy the integral constraint, and hence provide a tight upper bound to the seller's profits.

## 4.2 Discriminatory Pricing

The monotonicity conditions of Proposition 3 that guarantee increasing virtual values are not entirely appealing in our context. For example, starting from the common prior, if buyers observe binary signals, a bimodal distribution of beliefs would result with types holding beliefs above and below the mean  $\mu$ . In general, non monotone densities and distributions violating the standard monotonicity conditions are a quite natural benchmark. Therefore, ironing is not a technical curiosity in our case, but rather a technique that becomes unavoidable because of the features of the information environment.

We now illustrate the ironing procedure when virtual values are not monotone, and how it leads to a richer (two-item) optimal menu. Consider a bimodal distribution of types, which is given in this case by a linear combination of two Beta distributions. The probability density function and associated virtual values are given in Figure 4.




Figure 4: Probability Density Function and Virtual Values

Applying the procedure derived in Proposition 1, we consider the ironed versions of each virtual value, and we identify the equilibrium value of the multiplier  $\lambda^*$ . Notice that in this case the multiplier must be at the flat level of one of the virtual values: suppose not, apply the procedure from the regular case, and verify that it is impossible to satisfy the integral constraint.

Figure 5 illustrates the optimal two-item menu. Note that for types in the “pooling” region (approximately  $\theta \in [0.17, 0.55]$ ), the level of the allocation ( $q^* \approx -0.21$ ) is uniquely pinned down by the pooling property and by the integral constraint.

In both examples, extreme types with low value of information are excluded from purchase of informative signals. In the latter example, the monopolist is offering a second information structure that is tailored towards relatively lower types. This structure (with  $q < 0$ ) contains one signal that perfectly reveals the high state. This experiment is relatively unattractive for higher types, and it allows the monopolist to increase the price for the large mass of types

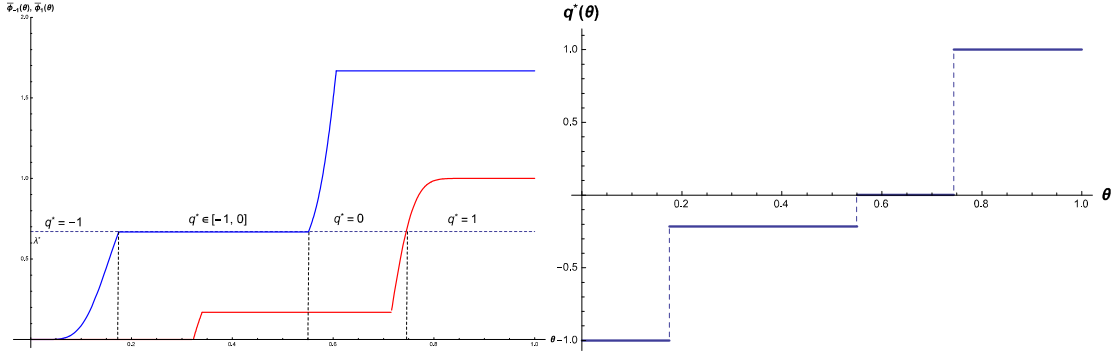


Figure 5: Ironed Virtual Values and Optimal Allocation

located around  $\theta \approx 0.7$ .

The properties of the optimal discriminatory pricing scheme reflect the fact that the type structure is quite different from the standard screening environment. Because of the Bayesian nature of the problem, the type space is not immediately ordered: there is a least informed type,  $\theta = 1/2$ , but not a most informed type. For instance, the distance  $|\theta - 1/2|$  is not sufficient to characterize the value of information of a given type  $\theta$ .

Consequently, type  $\theta = 1/2$  need not receive the most efficient information structure despite having the highest value of information. In the example above, inducing the types around  $\theta = 1/2$  to purchase the fully informative experiment would require imposing further distortions and charging a lower price for the second experiment. This would lead to a loss of revenue on the high-density types around  $\theta \approx 0.2$  that more than offsets the gain in revenue on the few types around  $1/2$ . More generally, whenever discriminatory pricing is optimal, the menu depends on the distribution of types in a rich way.

In the next subsection, we offer a precise characterization of the optimality of one- vs. two-item menus in a two-type environment.

### 4.3 Two Types

Consider the following two-type example. Let  $\theta \in \{0.2, 0.7\}$  with equal probability. The optimal menu is then given by  $q^*(\theta) \in \{-1/5, 0\}$ , with prices  $t^*(\theta) \in \{8/25, 3/5\}$ . In this example, the seller can offer the fully informative experiment  $q = 0$  to the type with the highest valuation (i.e.,  $\theta = 0.7$ ) and extract the buyer's entire surplus. In a canonical screening model, the seller would now have to exclude the lower type  $\theta = 0.2$ . However, when selling information, the monopolist can design another experiment with undesirable properties for the high type. In particular, the seller offers an experiment which is relatively more informative about the high state, and sets the price so to extract the low type's surplus.

The optimal menu is then characterized by the *most informative* such experiment the



seller can offer while extracting the entire surplus and without violating the high type’s incentive-compatibility constraint. Figure 6 illustrates the value of the two experiments offered by the monopolist as a function of the buyer’s type  $\theta \in [0, 1]$ .

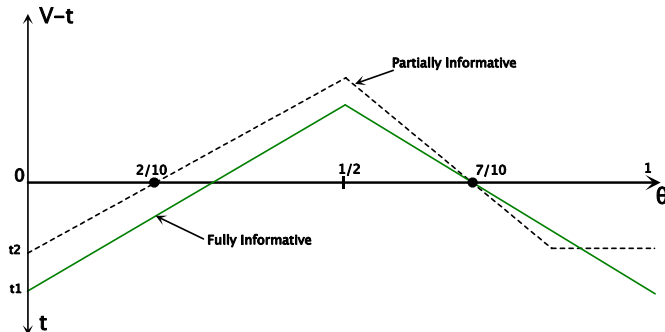


Figure 6: Net Value of Information Structures

More generally, with two types, we know the optimal menu contains either one or two experiments: if one only experiment is offered, one or both types may purchase. Offering two experiments is optimal only if the two types are *asymmetrically* located on opposite sides of  $1/2$ . (If they were at the same distance, the the seller would obtain the first-best profits.) Moreover, the allocation is characterized by “no distortion at the top,” and by full rent extraction whenever two experiments are offered. Finally, the type  $\theta$  closer to  $1/2$  buys the perfectly informative experiment. As we saw earlier, this is not true with more than two types.

Which distribution of types would the seller like to face? Notice that the horizontal differentiation aspect introduces a trade-off in the seller’s preferences between value of information and ability to screen different types (i.e., value creation vs. appropriation). In particular, screening becomes easier when types are located farther apart. In addition, if the common prior  $\mu$  is far from  $1/2$ , then farther-apart types are not necessarily better informed.

Consequently, the seller may benefit from a mean-preserving spread of  $F(\theta)$ . This seemingly counterintuitive result can occur when the seller serves the less informed type only, and the *ex ante* Blackwell more informative structure makes one buyer type less informed at the *interim* stage. While a spread can translate into higher profits for the seller, this does not imply she would like to give out free information (e.g., resulting in a distribution with four types).

We summarize our results with two types. Let  $\theta \in \{\theta_1, \theta_2\}$  with the corresponding frequency  $\gamma \triangleq \Pr(\theta = \theta_1)$ . We assume without loss that  $\theta_1 \leq 1/2$  and that the first type is

less informed, i.e.,  $|\theta_1 - 1/2| \leq |\theta_2 - 1/2|$ . Finally, we define the following threshold:

$$\bar{\gamma} \triangleq \frac{1 - \theta_2}{1 - \theta_1},$$

and we obtain the following result.

**Proposition 4** (Two Types). *The optimal menu with two types is the following:*

1. if  $|\theta_2 - 1/2| = |\theta_1 - 1/2|$ , then

$$q^*(\theta) \equiv 0;$$

2. if  $|\theta_2 - 1/2| > |\theta_1 - 1/2|$ , then

$$q^*(\theta_1) \neq q^*(\theta_2) \iff \gamma > \bar{\gamma};$$

3. if  $|\theta_2 - 1/2| > |\theta_1 - 1/2|$  and  $\theta_2 > 1/2 > \theta_1$ , then

$$0 = q^*(\theta_1) < q^*(\theta_2) < 1.$$

To conclude, we remark that the solution with two types can always be reconciled with the general case, and found using the integral constraint. In particular, because we can assume that the fully uninformative information structure is always present in the mechanism at zero price, we can construct the optimal allocation rule  $q^*(\theta)$  defined on the entire unit interval in order to satisfy the integral constraint. Not surprisingly then, the allocation rule resembles that of Figure 5, though the discreteness of this examples introduces an additional discontinuity.

## 5 Conclusions

We have examined the problem of a monopolist selling incremental information to privately informed buyers. The optimal mechanism involves at most two experiments, and we obtain sufficient conditions for one-item menus to be optimal. From the point of view of selling information, even under costless acquisition and free degrading, the optimal menu is quite coarse: this suggests a limited use of versioning, and the profitability of “minimal” distortions, in the absence of further, observable, heterogeneity among buyers or cost-efficiency reasons to provide impartially informative signals.

The comparative statics of the seller’s profits with respect to the distribution of types underscore a trade-off between value of information (to the buyer) and ability to screen

(for the seller). For instance, the ex-ante least informed types are not necessarily the most valuable, nor do they purchase the most informative signals in equilibrium. For the binary model, we have shown the equivalence between an environment with heterogeneous tastes for actions and one with heterogeneous beliefs. More work is required to clarify the role of orthogonal vs. correlated information that underscores the difference between preferences vs. beliefs heterogeneity with a continuum of states.

Further interesting extensions include studying the following: the optimal menu when the buyer is informed about an ex-ante type (e.g., about his private information structure, before observing any signals); the role of information-acquisition costs for the seller (which do not play a significant role if fixed or linear in precision, but may induce further cost-based screening if convex in the quality of the information released to the buyer); and the effect of competition among sellers of information (i.e. formalizing the intuition that each seller will be able to extract the surplus related to the innovation element of his database).

# Appendix

**Proof of Lemma 1.** Consider the following procedure. Fix any type  $\theta$  and experiment  $I$ . Let  $S_I^a$  denote the sets of the signals in experiment  $I$  that induce type  $\theta$  to choose action  $a$ . Thus,  $\cup_{a \in A} S_I^a = S_I$ . Construct the experiment  $I'$  as a recommendation for type  $\theta$  based on the experiment  $I$ ,  $S_{I'} = \{s_a\}_{a \in A}$  and

$$\pi_{I'}(s_a | \omega) = \int_{S_I^a} \pi_I(s | \omega) ds \quad \omega \in \Omega, a \in A.$$

By construction,  $I'$  induces the same outcome distribution for type  $\theta$  as  $I$  so  $V(I', \theta) = V(I, \theta)$ . At the same time,  $I'$  is a garbling of  $I$  so by Blackwell's theorem  $V(I', \theta') \leq V(I, \theta') \quad \forall \theta'$ .

We can use this procedure to construct for any feasible direct mechanism  $\{I(\theta), t(\theta)\}$  another feasible direct mechanism  $\{I'(\theta), t(\theta)\}$  with its experiments consisting of no more signals than the cardinality of action space  $A$ . Because we consider a binary setting, every experiment in an optimal menu consists of two signals only. ■

**Proof of Lemma 2.** Consider any individually rational and incentive compatible direct mechanism  $\mathcal{M} = \{\alpha(\theta), \beta(\theta), t(\theta)\}$ . For each  $\theta$  define  $\varepsilon(\theta) := 1 - \max\{\alpha(\theta), \beta(\theta)\}$ ,  $\alpha'(\theta) := \alpha(\theta) + \varepsilon(\theta)$ , and  $\beta'(\theta) := \beta(\theta) + \varepsilon(\theta)$ . It follows from the information value formula that

$$[V(\alpha'(\theta), \beta'(\theta), \theta) - \varepsilon(\theta)]^+ = V(\alpha(\theta), \beta(\theta), \theta).$$

Consequently, a direct mechanism  $\mathcal{M}' = \{\alpha'(\theta), \beta'(\theta), t(\theta) + \varepsilon(\theta)\}$  is also individually rational and incentive compatible. Moreover, for any type  $\theta$  either  $\alpha(\theta) = 1$  or  $\beta(\theta) = 1$  and transfers are weakly greater than in  $\mathcal{M}$ . ■

**Proof of Lemma 3.** Since each type's outside option coincides with the value of choosing an uninformative experiment, we drop the positivity qualifier in the formula for value function and set  $q(0) = -1$  and  $q(1) = 1$ .

Necessity. For any two types  $\theta_2 > \theta_1$  we have

$$\begin{aligned} V(q_1, \theta_1) - t_1 &\geq V(q_2, \theta_1) - t_2, \\ V(q_2, \theta_2) - t_2 &\geq V(q_1, \theta_2) - t_1, \\ V(q_2, \theta_2) - V(q_1, \theta_2) &\geq t_2 - t_1 \geq V(q_2, \theta_1) - V(q_1, \theta_1). \end{aligned}$$

It follows from the single-crossing property of  $V(q, \theta)$  that  $q_2 \geq q_1$  hence  $q(\theta)$  is increasing. Because the buyer's rent is non-decreasing (non-increasing) in  $\theta$  on  $[0, 1/2]$  and  $[1/2, 1]$

respectively, we can compute the function  $U(\theta)$  on  $[0, 1/2]$  and  $[1/2, 1]$  separately as

$$U(1/2) = U(0) + \int_0^{1/2} V_\theta(q, \theta) d\theta = U(1) - \int_{1/2}^1 V_\theta(q, \theta) d\theta.$$

By the envelope theorem  $V_\theta(q, \theta) = q + 1$  for  $\theta < 1/2$  and  $= q - 1$  for  $\theta > 1/2$ . Taking into account the boundary conditions  $U(0) = U(1) = 0$  we obtain

$$\int_0^1 q(\theta) d\theta = 0.$$

The corresponding transfers can be derived from the allocation rule as

$$\begin{aligned} U(\theta) &= V(q(\theta), \theta) - t(\theta) = 0 + \int_0^\theta q(\theta') d\theta' + \min\{\theta, 1 - \theta\} \\ t(\theta) &= q(\theta)\theta - \int_0^\theta q(\theta') d\theta' - \max\{q(\theta), 0\}. \end{aligned}$$

Sufficiency. Expected utility for a type  $\theta$  from reporting  $\theta'$  is

$$V(q(\theta'), \theta) - t(\theta') = (\theta - \theta')q(\theta') + \int_0^{\theta'} q(\theta) d\theta + \min\{\theta, 1 - \theta\}$$

which is maximized at  $\theta' = \theta$  by monotonicity of  $q(\cdot)$ ; incentive constraints are satisfied. At the same time  $U(\theta)$  is equal to zero for types 0 and 1 and is weakly positive for all others; participation constraints are satisfied. ■

**Proof of Proposition 1.** Consider the seller's problem (3). We first establish that the solution can be characterized through Lagrangean methods. For necessity, note that the objective is concave in the allocation rule; the set of non-decreasing functions is convex; and the integral constraint can be weakened to the real-valued inequality constraint

$$\int_0^1 q(\theta) d\theta \leq 0. \tag{4}$$

Necessity of the Lagrangean then follows from Theorem 8.3.1 in Luenberger (1969). Sufficiency follows from Theorem 8.4.1 in Luenberger (1969). In particular, any solution maximizer of the Lagrangean  $q(\theta)$  with

$$\int_0^1 q(\theta) d\theta = \bar{q}$$

maximizes the original objective subject to the inequality constraint

$$\int_0^1 q(\theta) d\theta \leq \bar{q}.$$

Thus, any solution to the Lagrangean that satisfies the constraint also solves the original problem.

Because the Lagrangean approach is valid, we can apply the results of Toikka (2011) to solve the seller's problem for a given value of the multiplier  $\lambda$  on the integral constraint. Write the Lagrangean as

$$\int_0^1 [(\theta f(\theta) + F(\theta)) q(\theta) - (\max\{q(\theta), 0\} + \lambda) f(\theta)] d\theta.$$

In order to maximize the Lagrangean subject to the monotonicity constraint, consider the *generalized virtual surplus*

$$\bar{J}(\theta, q) := \int_{-1}^q (\bar{\phi}(\theta, x) - \lambda^*) dx,$$

where  $\bar{\phi}(\theta, x)$  denotes the ironed virtual value for allocation  $x$ . Note that  $\bar{J}(\theta, q)$  is weakly concave in  $q$ . Because the multiplier  $\lambda$  shifts all virtual values by a constant, the result in Proposition 1 then follows from Theorem 4.4 in Toikka (2011). Finally, note that  $\bar{\phi}(\theta, q) \geq 0$  for all  $\theta$  implies the value  $\lambda^*$  is strictly positive (otherwise the solution  $q^*$  would have a strictly positive integral). Therefore, the integral constraint (4) binds. ■

**Proof of Proposition 2.** From the Lagrangean maximization, we have the following necessary conditions

$$q^*(\theta) = \begin{cases} -1 & \text{if } \bar{\phi}(\theta, -1) < \lambda^*, \\ \bar{q} \in [-1, 0] & \text{if } \bar{\phi}(\theta, -1) = \lambda^*, \\ 0 & \text{if } \bar{\phi}(\theta, -1) > \lambda^* > \bar{\phi}(\theta, 1), \\ \bar{q}' \in [0, 1] & \text{if } \bar{\phi}(\theta, 1) = \lambda^*, \\ 1 & \text{if } \bar{\phi}(\theta, -1) > \lambda^*, \end{cases}$$

and

$$\int_0^1 q^*(\theta) d\theta = 0.$$

If  $\lambda^*$  coincides with the flat portion of one virtual value, then by the pooling property of

Myerson (1981), the optimal allocation rule must be constant over that interval, and the level of the allocation is uniquely determined by the integral constraint. Finally, suppose  $\lambda^*$  equals the value of  $\bar{\phi}(\theta, q^*(\theta))$  over more than one flat portion of the virtual values  $\bar{\phi}(\theta, -1)$  and  $\bar{\phi}(\theta, 1)$ . Then we can focus without loss on the allocation  $q^*$  that assigns experiment  $q = 0$  or  $q \in \{-1, 1\}$  to all types in one of the two intervals. ■

**Proof of Proposition 3.** (1.) If  $F(\theta) + \theta f(\theta)$  and  $F(\theta) + (\theta - 1)f(\theta)$  are strictly increasing then the ironing is not required and it follows from the analysis in the text that the optimal solution has a single step at  $q = 0$ .

(2.) If all types are located at one side from  $1/2$  then the integral constraint has no bite since the allocation rule  $q(\theta)$  can always be adjusted on the other side to satisfy it. The unconstrained problem has a single step at  $q = 0$  that results in flat pricing.

(3.) If types are symmetrically distributed then the separately optimal menus for types  $\theta < 1/2$  and  $\theta > 1/2$  are the same. Since the profits in the jointly optimal menu cannot be higher than weighted sum of profits in the separate ones the result follows. ■

**Proof of Proposition 4.** Let  $\theta \in \{\theta_1, \theta_2\}$  with the corresponding frequency  $\gamma \triangleq \Pr(\theta = \theta_1)$ . We assume without loss that  $\theta_1 \leq 1/2$  and that the first type is less informed, i.e.  $|\theta_1 - 1/2| \leq |\theta_2 - 1/2|$ . Finally, define the following threshold:

$$\bar{\gamma} \triangleq \frac{1 - \theta_2}{1 - \theta_1}.$$

(1.) If  $\theta_2 \leq \theta_1 \leq 1/2$  then both types have the same default action and by Proposition 3 the optimal scheme is flat pricing. The value of the fully informative experiment to these types given by  $V(\theta, 0) = \min\{\theta, 1 - \theta\}$ . Therefore the seller compares the profits from non-exclusive vs. exclusive pricing, i.e.,

$$\theta_2 \text{ vs. } \gamma\theta_1.$$

Consequently, if  $\gamma > \theta_2/\theta_1$  only the less informed type is served at price  $t_1 = \theta_1$ . Otherwise, both types are served at price  $t_1 = t_2 = \theta_2$ .

(2.) If  $\theta_1 \leq 1/2 \leq \theta_2$  then types have different default actions. In an optimal menu  $q_1 = 0$  and  $q_2 \geq 0$ . Indeed, an optimal menu must contain the fully informative experiment,  $q = 0$ . Since  $q$  is monotone by Lemma 3 either  $q_1 < 0$  and  $q_2 = 0$  or  $q_1 = 0$  and  $q_2 \geq 0$ . However, if  $q_2 = 0$  then  $q_1 = 0$  since otherwise the seller would have to charge  $t_1 < t_2$  that would contradict optimality.

If  $\gamma < \bar{\gamma}$  then flat pricing is optimal, the seller serves both types with fully informative experiment,  $q_1 = q_2 = 0$  at price  $t_1 = t_2 = 1 - \theta_2$ . To the contradiction, assume that

$q_2 > 0$  so the type  $\theta_2$  is served with a skewed experiment. The seller can then improve the menu by increasing the quality of the skewed experiment and adjusting prices:  $q'_1 = q_1$ ,  $t'_1 = t_1 - \varepsilon(\theta_2 - \theta_1)$ ,  $q'_2 = q_2 - \varepsilon$ ,  $t'_2 = t_2 + \varepsilon(1 - \theta_2)$ ,  $\varepsilon > 0$ . It can be checked directly by (2) that the modified menu is individually rational and incentive compatible and improves profits by  $\varepsilon(1 - \theta_1)(\hat{\gamma} - \gamma)$  contradicting the optimality of  $q_2 > 0$ .

(3.) If  $\gamma > \hat{\gamma}$  then the discriminatory pricing is optimal with type  $\theta_1$  being served with fully informative experiment  $q_1 = 0$  at price  $t_1 = \theta_1$ . Full surplus is extracted from type  $\theta_1$  since otherwise the menu can be improved by the modification above with  $\varepsilon < 0$ . Note that the more informed type  $\theta_2$  would not buy  $q_1 = 0$  at price  $\theta_1$  anyways so he is served with a skewed experiment with optimal  $q_2, t_2$  solving the following problem:

$$\begin{aligned} & \max_{q_2 \in [0,1], t_2 \geq 0} t_2 \\ \text{s.t. } & \theta_1 q_2 - q_2 + \theta_1 - t_2 \leq 0, \quad (IC_{12}), \\ & \theta_2 q_2 - q_2 + 1 - \theta_2 - t_2 \leq 0, \quad (IR_2). \end{aligned}$$

Both constraints bind at the optimum and we obtain

$$\begin{aligned} q_2 &= \frac{\theta_1 + \theta_2 - 1}{\theta_2 - \theta_1}, \\ t_2 &= \frac{(1 - 2\theta_1)(1 - \theta_2)}{\theta_2 - \theta_1}. \end{aligned}$$

That ends the proof. ■



## References

- ABRAHAM, I., S. ATHEY, M. BABAIOFF, AND M. GRUBB (2014): “Peaches, Lemons, and Cookies: Designing Auction Markets with Dispersed Information,” Discussion paper, Microsoft Research, Stanford, and Boston College.
- ADMATI, A. R., AND P. PFLEIDERER (1986): “A monopolistic market for information,” *Journal of Economic Theory*, 39(2), 400–438.
- (1990): “Direct and indirect sale of information,” *Econometrica*, 58(4), 901–928.
- BALESTRIERI, F., AND S. IZMALKOV (2014): “Informed seller in a Hotelling market,” Discussion paper, HP Labs and New Economic School.
- BERGEMANN, D., AND A. BONATTI (2015): “Selling Cookies,” *American Economic Journal: Microeconomics*, forthcoming.
- BERGEMANN, D., AND M. PESENDORFER (2007): “Information Structures in Optimal Auctions,” *Journal of Economic Theory*, 137(1), 580–609.
- CELIK, L. (2014): “Information unraveling revisited: disclosure of horizontal attributes,” *Journal of Industrial Economics*, 62(1), 113–136.
- ESŐ, P., AND B. SZENTES (2007a): “Optimal information disclosure in auctions and the handicap auction,” *Review of Economic Studies*, 74(3), 705–731.
- (2007b): “The price of advice,” *Rand Journal of Economics*, 38(4), 863–880.
- HÖRNER, J., AND A. SKRZYPACZ (2015): “Selling Information,” *Journal of Political Economy*, forthcoming.
- JOHNSON, J. P., AND D. P. MYATT (2006): “On the Simple Economics of Advertising, Marketing, and Product Design,” *American Economic Review*, 96(3), 756–784.
- KAMENICA, E., AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101(6), 2590–2615.
- KOESSLER, F., AND V. SKRETA (2014): “Sales Talk,” Discussion paper, PSE and UCL.
- LI, H., AND X. SHI (2013): “Discriminatory Information Disclosure,” Discussion paper, University of British Columbia and University of Toronto.

- LIZZERI, A. (1999): “Information revelation and certification intermediaries,” *RAND Journal of Economics*, 30(2), 214–231.
- LUENBERGER, D. G. (1969): *Optimization by Vector Space Methods*. John Wiley & Sons.
- MUSSA, M., AND S. ROSEN (1978): “Monopoly and Product Quality,” *Journal of Economic Theory*, 18(2), 301–317.
- MYERSON, R. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- MYLOVANOV, T., AND T. TRÖGER (2014): “Mechanism Design by an Informed Principal: Private Values with Transferable Utility,” *Review of Economic Studies*, 81(4), 1668–1707.
- OTTAVIANI, M., AND A. PRAT (2001): “The value of public information in monopoly,” *Econometrica*, 69(6), 1673–1683.
- RAYO, L., AND I. SEGAL (2010): “Optimal Information Disclosure,” *Journal of Political Economy*, 118(5), 949–987.
- RILEY, J., AND R. ZECKHAUSER (1983): “Optimal Selling Strategies: When to Haggle, When to Hold Firm,” *Quarterly Journal of Economics*, 98, 267–290.
- SARVARY, M. (2012): *Gurus and Oracles: The Marketing of Information*. MIT Press.
- SHAPIRO, C., AND H. R. VARIAN (1999): *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business Press.
- TOIKKA, J. (2011): “Ironing without Control,” *Journal of Economic Theory*, 146(6), 2510–2526.