

MONOPOLISTIC SIGNAL PROVISION

LUIS RAYO

ABSTRACT. I consider the problem of a monopolist serving a population of consumers with hidden types that use the goods sold by the monopolist as a signaling device. The discrimination opportunities that arise include interfering with the transmission of information. In particular, by restricting the variety of signals, and forcing subsets of consumers to pool, the monopolist can extract additional information rents. The model delivers a theory of restricted variety in the supply of conspicuous goods.

1. INTRODUCTION

A classic problem in economics is that of a monopolist serving a population of consumers with unknown tastes –Mussa and Rosen [1978], Maskin and Riley [1984] (henceforth MRMR). These models are concerned with the standard case where each consumer’s valuation stems from the intrinsic properties of the goods she consumes, without reference to the consumption of others. In practice, while individuals care about intrinsic value, the consumption of many goods is also motivated by social interaction –Veblen [1899], Bernheim [1994], Pesendorfer [1995], Bagwell and Bernheim [1996], Becker and Murphy [2000]. Expensive goods, for example, can be a signal of wealth, and their consumption a means to increase social status. Presumably this is a central motivation behind the consumption of luxury goods ranging from cars to jewelry to first-class airplane tickets. In fact, it is rare to find luxury goods without intentional conspicuous features directly built into them.

This paper extends the MRMR problem to the case where the monopolist sells a menu of conspicuous goods that are used as a signal concerning the consumers’ hidden characteristics. I begin with a benchmark model where the monopolist offers a menu of goods that provide no intrinsic value. The good purchased by each consumer is observed by her peers, and is used to form a Bayesian posterior belief (a social perception) concerning the consumer’s type, e.g., her wealth. Consumers

Date: First Draft: 09/99. This Draft: 09/03.

The University of Chicago, GSB, lrayo@gsb.uchicago.edu. Acknowledgments to be added.

are assumed to care about this social perception.¹ The result is a signaling game where the monopolist has the power to select the set of available signals.

The main implication of the model is that the set of signals offered will not always be rich enough to allow for a full revelation of consumer types. Rather, the monopolist can potentially increase profits by restricting the variety of available goods in a way that forces subsets of consumers to pool. The thesis of the paper is that the source of variety restrictions is precisely a desire to interfere with the transmission of information. Whenever present, these variety restrictions will correspond to non-convexities in the menu of conspicuous goods. This contrasts with the menu of goods offered in MRMR, which, although includes downward distortions in quantity or quality, is always convex.

A simple example sketches the rationale behind these variety restrictions. Suppose a jeweler must select between two menus of diamond rings. The first menu consists of five rings with prices \$1, \$2, \$3, \$4, \$5, and the second menu simply excludes the intermediate \$3 option. In the absence of a signaling motive, this restriction will have only two effects: some of the wealthier consumers that would have purchased the \$3 ring will purchase the \$4 option instead (increasing profits), and some of the less wealthy will move down to the \$2 ring (reducing profits). From MRMR we know that, under single-crossing preferences, the combined effect over profits will be negative (or, more precisely, that there always exists a ring between \$2 and \$4 that would increase profits if offered).

However, when individuals care about the signaling consequences of their consumption, some additional effects will arise. Notice that once the \$3 ring is deleted, the average consumer of the \$4 ring will be less wealthy than before (due to the effect described above). Thus, the social image associated to this ring will deteriorate, and some of the previous wealthier \$4-ring consumers will prefer the \$5 ring instead (increasing profits). On the other hand, by the same token, the image associated to the \$2 ring will improve, motivating some previous \$1-ring consumers to buy the more expensive \$2 option (further increasing profits). These effects provide a new reason to limit variety.²

¹For example, a favorable social perception can lead to desirable outcomes such as access to a given social group, business partners, or even an attractive spouse.

²A more active way to increase profits is to increase the price of the rings that previously cost \$2 and \$5, which would still be purchased under the higher price due to the peer effects resulting from the deletion of the \$3 ring. The model will solve for the optimal strategy.

In section 5, I extend the model to allow for goods with intrinsic quality. The setup combines a version of MRMR (with unit demands and endogenous quality) and the benchmark model above. The novelty of the solution resides in the fact that pooling regions will give rise to discontinuities in the pricing schedule (i.e., variety gaps), as opposed to mere kinks (as in MRMR). Conversely, the presence of variety gaps will imply an underlying pooling of types, as opposed to discontinuities in the intrinsic quality of the goods. Section 6 presents some examples that are suggestive of variety restrictions in practice.

The study of discrimination behavior behind conspicuous goods is of interest because, in practice, these goods are commonly accompanied by some degree of market power. A luxury brand that is already well known, for example, will have an advantage as a signal provider vis a vis her little-known competitors. But acquiring this status will typically require large expenditures in advertisements and innovation, with a fixed-cost nature. Consumers also tend to coordinate around particular brands, favoring those that manage to become fashionable. The assumption of a single supplier is meant simply as a starting point towards understanding the discrimination opportunities that arise.

2. BENCHMARK MODEL

Consider a monopolist that sells signals to a continuum of consumers with unit mass. Consumers are characterized by their type θ , which is distributed over the interval $[\theta_L, \theta_H] \subset \mathbb{R}$ according to a smooth density $f > 0$, and distribution function F . These types are private information. A signal consists of an advertisement, observable by the general public (i.e., all consumers), that indicates how much a consumer paid the monopolist.

A consumer's payoff depends on the public's perception of her type, denoted by π , and on the amount she paid the monopolist, denoted by t . As formalized below, I assume that π simply corresponds to the average type across all consumers that, in equilibrium, purchase the same signal. Payoffs are given by

$$U(\pi, t; \theta) = \pi v(\theta) - t,$$

where θ stands for the true type, and $v > 0$ is smooth with $v' > 0$. Under these payoffs, the indifference curves will be single-crossing in (π, t) , allowing consumers to sort. For example, if θ represents wealth, then wealthier individuals are willing to pay more in order to be perceived as being rich (i.e., a higher level of π). On

the other hand, the marginal contribution of π to utility, $v(\theta)$, is constant for each type. Together with quasilinearity in $-t$, it implies that the indifference curves have a constant slope of $v(\theta)$. These assumptions are imposed for tractability.

The monopolist seeks to maximize total revenues

$$\int_{\theta_L}^{\theta_H} t(\theta) dF(\theta),$$

where $t(\theta)$ is the amount θ spends on her signal. The cost of providing a signal is assumed to be zero.

The model has three stages. In the first, the monopolist offers a menu $\mathcal{M} \subset \mathbb{R}$ of signals (i.e., a set of money amounts she is willing to advertise) that becomes common knowledge. Without loss of generality, $0 \in \mathcal{M}$. I also assume that \mathcal{M} is bounded from below (otherwise the monopolist would suffer unbounded losses).³ In the second stage, each consumer simultaneously purchases an item from this menu. Let $s(\theta)$ denote the signal selected by type θ . Finally, in the third stage, after observing stage-2 behavior, consumers form average perceptions π , about others' type θ , according to Bayes rule and prior knowledge of F . I assume this process results in an equilibrium defined as follows:

Definition 1. *Given \mathcal{M} , an equilibrium for the signaling game is a pair of functions $s : [\theta_L, \theta_H] \rightarrow \mathcal{M}$ and $\pi : \mathcal{M} \rightarrow [\theta_L, \theta_H]$, such that:*

- (i) $s(\theta) \in \arg \max_{x \in \mathcal{M}} \pi(x)v(\theta) - x$ for all θ , and
- (ii) $\pi(x) = E[\theta : s(\theta) = x]$ for all $x \in s([\theta_L, \theta_H])$.

³More generally, the monopolist could offer a menu of signals that include conspicuous features other than their price, e.g., different colors or styles (with no intrinsic value), while also keeping some prices secret. These schemes, however, cannot increase profits since, in equilibrium, consumers will only purchase those signals that, for any given price, lead to the highest level of π (implying that offering two different styles for the same price is redundant), and those signals that, for any given π , have the lowest price (implying that hiding prices is also redundant). In practice, of course, we do observe different colors and styles. These may be precisely a means to communicate price (as opposed to, say, a permanent price tag on a car), or may be intrinsically valuable, as modeled in section 5. On the other hand, although rarely observed, the monopolist could also offer lotteries of conspicuous goods. Since consumers are assumed to be risk neutral with respect to π , the only advantage of such lotteries would be the ability to partially pool subsets of types. However, precisely because of this linearity, partial pooling will always be (weakly) dominated by either full pooling or full separation, and therefore restricting to deterministic menus, which allow for both cases, is also without loss.

That is, given the perception function π , each consumer maximizes her payoff, and π is Bayes-consistent with consumer behavior. Notice that such an equilibrium corresponds to a Perfect Bayesian Equilibrium for an extended game in which consumers face a loss if their beliefs do not correspond to a Bayesian posterior.

A vast multiplicity of equilibria will typically arise (for example, there is always an equilibrium where all consumers purchase the same signal, and profits are non-positive). However, I impose the crucial assumption that the monopolist can coordinate consumers on one of her preferred equilibria.⁴

It will be convenient to reformulate the game as one of direct revelation, where consumers are invited to announce their types to the monopolist, who pre-commits to emit a signal $t(\cdot)$ (with slight abuse of notation) as a function of this announcement, and then perceptions are formed accordingly. The revelation principle then guarantees that the optimal voluntary (i.e., ‘0’ is always available) and truth-telling direct-revelation mechanism will indicate the highest possible payoff for the monopolist. Conversely, the outcome arising from this optimal mechanism will be implementable in the original formulation of the game by setting, for example, $\mathcal{M} = t([\theta_L, \theta_H]) \cup \{0\}$, and having consumers coordinate on the equilibrium actions $s = t$.

The monopolist’s problem therefore reduces to that of selecting the most profitable schedule t out of those that induce truth telling (constraint *(IC)* below) and participation in the mechanism (constraint *(P)*), while the social image follows from Bayesian inference (constraint *(BI)*):⁵

⁴The focus on pure strategy equilibria (where $s(\theta)$ is deterministic) is without loss. The reason is that at most a countable set of types can be induced to randomize in equilibrium (a consequence of single crossing). As a result, since types are atomless, these randomizations will have no impact over the monopolist’s payoff. Footnote 6 provides further detail.

⁵For the participation constraint, I assume without loss (given that the monopolist can select the equilibrium) that the image associated to non-participation, $\pi(0)$, is equal to θ_L .

In a paper posterior (although independent) to the present one, Damiano and Hao [2002] consider the problem of a monopolist controlling a matching market which is mathematically connected to the problem above, although such a market structure introduces matching feasibility constraints that generally differ from the Bayesian-inference constraints used herein.

$$\begin{aligned}
(1) \quad & \max_t \int_{\theta_L}^{\theta_H} t(\theta) dF(\theta) \\
(IC) \quad & s.t. \pi(t(\theta))v(\theta) - t(\theta) \geq \pi(t(\theta'))v(\theta) - t(\theta') \text{ for all } \theta, \theta', \\
(P) \quad & \pi(t(\theta))v(\theta) - t(\theta) \geq \theta_L v(\theta) \text{ for all } \theta, \\
(BI) \quad & \pi(x) = E[\theta' : t(\theta') = x] \text{ for all } x \in t([\theta_L, \theta_H]).
\end{aligned}$$

3. ANALYSIS

I begin by expressing problem (1) in a more tractable way. Lemma 1 is a straightforward consequence of single crossing. All omitted proofs are in the appendix.

Lemma 1. *In problem (1), all participation constraints are redundant, except that for $\theta = \theta_L$, which will hold with equality.*

An alternative to problem (1) is to view the monopolist as an active seller of the perceived type, instead of treating π simply as a consequence of t . In other words, we can reformulate her problem as one where she explicitly sells this social image with zero production costs, but is restricted to be consistent with Bayesian inference. Let $\varphi(\theta) := \pi(t(\theta))$ be the image sold to type θ in exchange for payment $t(\theta)$, which is a mapping from true types into perceived types. By selecting φ , the monopolist will decide how accurately types are revealed to society, while she may potentially limit accuracy in order to increase revenues. In direct revelation form, the monopolist will now invite consumers to announce their type in return for a social image $\varphi(\theta)$ and transfer $t(\theta)$. Using lemma 1, we can write the problem as

$$\begin{aligned}
(2) \quad & \max_{\varphi, t} \int_{\theta_L}^{\theta_H} t(\theta) dF(\theta) \\
(IC') \quad & s.t. \varphi(\theta)v(\theta) - t(\theta) \geq \varphi(\theta')v(\theta) - t(\theta') \text{ for all } \theta, \theta', \\
(P') \quad & \varphi(\theta_L)v(\theta_L) - t(\theta_L) = \theta_L v(\theta_L), \\
(BI') \quad & \varphi(\theta) = E[\theta' : t(\theta') = t(\theta)] \text{ for all } \theta.
\end{aligned}$$

This change in formulation will prove useful once two additional simplifications are made, allowing us to rewrite (2) only in terms of φ , while sending t to the background. First, constraint (BI') can be expressed in terms of φ alone as follows:

Lemma 2. *Suppose (IC') holds. Then, (BI') holds if and only if*

$$(TF) \quad \varphi(\theta) = E[\theta' : \varphi(\theta') = \varphi(\theta)] \text{ for all } \theta.$$

Condition (TF) simply requires that the social image sold to a given type must correspond to the average type of all consumers purchasing that same image in equilibrium. Whenever a schedule φ satisfies this condition, I will call it a “truthful filter.” Consider two examples: (i) $\varphi(\theta) = \theta$ for all θ (i.e., the identity function), and (ii) $\varphi(\theta) = E[\theta' : \theta' \leq \theta^*]$ for $\theta \leq \theta^*$ and $\varphi(\theta) = \theta$ for $\theta > \theta^*$. (i) represents the case where all types are fully revealed, and (ii) the case where types up to θ^* are pooled together, and full sorting occurs after that.

Second, by use of an envelope theorem (Milgrom and Segal [2002], see also Myerson [1981]), the problem can be further simplified by exchanging the (IC') and (P') constraints for an equation relating t and φ :

Lemma 3. *(Constraint Simplification.) In problem (2), (IC') and (P') are satisfied if and only if*

$$t(\theta) = \varphi(\theta)v(\theta) - \int_{\theta_L}^{\theta} \varphi(z)v'(z)dz - \theta_L v(\theta_L), \text{ and}$$

φ is non-decreasing.

By combining lemmas 2 and 3, the problem finally becomes one where the monopolist simply searches for the optimal non-decreasing truthful filter, while the transfer schedule t is sent to the background:

$$(3) \quad \max_{\varphi} \int_{\theta_L}^{\theta_H} \varphi(\theta)h(\theta)dF(\theta)$$

s.t. φ is non-decreasing, and

$$(TF) \quad \varphi(\theta) = E[\theta' : \varphi(\theta') = \varphi(\theta)] \text{ for all } \theta.$$

Where $h(\theta) := v(\theta) - v'(\theta)\frac{1-F(\theta)}{f(\theta)}$, and the objective is obtained by substituting for t from lemma 3, and changing the order of integration. A non-decreasing truthful filter will consist of a collection of flat and increasing intervals, where increasing intervals lie over the 45° line, and two separated flat intervals cannot have the same




FIGURE 1. A Non-Decreasing Truthful Filter

height. Figure 1 provides an example.⁶ A flat interval implies that these types are pooled together, while intervals over the 45° line imply separation.⁷

The shape of h will play a key role in the solution. In particular, an oscillating h will cause alternating pooling and separating regions. For expositional simplicity, I will restrict to the case where h oscillates at most finitely many times (i.e., h' changes sign finitely many times). When oscillations are infinite some technical nuisance is added but all the results extend in a straightforward way.

Due to the nature of the constraints, standard optimization techniques cannot be applied to this problem. Rather, we must consider whether each particular interval of types should be pooled together or not. Consider the decision to either fully pool or fully separate all types in a given interval (θ_1, θ_2) . The (virtual) revenue

⁶(Footnote 4 continued.) Under a mixed-strategy equilibrium, the allocation of social image $\pi(\theta)$ need not correspond to a truthful filter. However, since (IC) implies that only countably many types can be induced to randomize across signals, this social image must satisfy the (TF) constraint over a full-measure subset, and therefore cannot increase revenues beyond what pure-strategy equilibria can achieve.

⁷Lizzeri [1999] considers a monopolistic certification intermediary who can gather information from informed parties, and decides how to transmit it to an uninformed market. The emphasis is also on the manipulation of this information. But in Lizzeri's environment with a technology to acquire information (and without single-crossing preferences) the strategic opportunities for the monopolist are rather different, giving rise to an interesting result: the monopolist will report only whether or not the private type is above a given threshold.

associated to this interval is $\int_{\theta_1}^{\theta_2} \varphi(\theta)h(\theta)dF(\theta)$. Under full separation this revenue becomes $\int_{\theta_1}^{\theta_2} \theta h(\theta)dF(\theta)$, while pooling delivers $\int_{\theta_1}^{\theta_2} E[\theta \in (\theta_1, \theta_2)] h(\theta)dF(\theta)$. Notice that the difference is equal to $COV[\theta, h(\theta) : \theta \in (\theta_1, \theta_2)] \int_{\theta_1}^{\theta_2} dF(\theta)$, where the *COV* term is the covariance between θ and $h(\theta)$ within the interval (θ_1, θ_2) . Therefore, it will be desirable to separate whenever this covariance is positive, and vice versa.

Whenever h is increasing over (θ_1, θ_2) , for example, the covariance becomes positive and it will be optimal to separate. The (mathematical) reason for this is that pooling implies that all types at the low end of the interval receive a higher image relative to full separation, while the opposite occurs for the types at the high end of the interval, but these gains and losses are weighted in the monopolist's objective by h , which is larger for higher types. The opposite is true when h is decreasing. As a result, a necessary and sufficient condition for full separation of all types $[\theta_L, \theta_H]$ to be optimal is that h is always non-decreasing:

Proposition 1. *The identity truthful filter, $\varphi(\theta) = \theta$ for all θ , is optimal if and only if h is everywhere non-decreasing.*

Proof. For sufficiency (\Leftarrow), whenever h is everywhere non-decreasing the above covariance argument implies that any interval that is pooled together could instead be separated while weakly increasing the objective.

For necessity (\Rightarrow), notice that it cannot be the case that an optimal φ lies over the 45° line over an interval where h is decreasing, because otherwise the objective could be increased by pooling this interval together. ■

Economic intuition is the following. Pooling an interval has three effects over revenues relative to separation. First, the consumers at the low end of the pool, which now receive a higher social image, can be charged a higher price. Second, the opposite occurs with the consumers at the high end of the pool. Third, all the higher-type consumers that manage to escape the pool can be charged a higher price in exchange for avoiding this low-image pool (i.e., their information rents are reduced). Due to single crossing, the combined impact of the first two effects is always negative. The third effect, however, will dominate as soon as h is decreasing over the relevant region. For this to occur we require that either v is sufficiently convex (v' is increasing) or that the inverse likelihood ratio $\frac{1-F}{f}$ is rapidly increasing (or a combination of the two). A convex v means that the consumers to the right of the pool have a high valuation for social image relative to the consumers inside the pool, implying that the third effect is large relative to the second. An increasing

inverse likelihood ratio, on the other hand, is associated to a sharply decreasing f and slowly increasing F . This will imply that the image associated to the pool will be low (as the mass is concentrated at the low end), allowing for a higher extraction of information rents from higher-end consumers, while at the same time there are plenty of these consumers to extract information rents from. For example, as discussed below, pooling will never be optimal at the high end of the type space since the third effect will disappear.


3.1. Pooling and Negative Covariance. Although a decreasing h will always lead to some type of pooling, the optimal pooling regions will not coincide with the regions over which h is decreasing. Consider, for example, the case depicted in figure 2 where h is decreasing over the interval (θ_A, θ_B) , and increasing elsewhere. Pooling the interval (θ_A, θ_B) alone is better than no pooling at all. However, it will be optimal to extend the pool in both directions. Suppose we extend the left end of the pool to $\theta_A - \epsilon$. Those new types in $(\theta_A - \epsilon, \theta_A]$ will now receive a higher image, while the opposite occurs to all types in (θ_A, θ_B) . But, extending the argument above, these gains and losses will be weighed by h , which is on average larger over $(\theta_A - \epsilon, \theta_A]$ than over (θ_A, θ_B) , implying that this change is beneficial. A symmetric argument applies over a neighborhood to the right of θ_B .

More generally, consider the net gains from pooling an arbitrary interval (θ_1, θ_2) relative to any type of separation within this interval (including full separation, or separation across some types and pooling across others). Pooling delivers $\int_{\theta_1}^{\theta_2} E[\theta \in (\theta_1, \theta_2)]h(\theta)dF(\theta)$, while separation under truthful filter φ delivers $\int_{\theta_1}^{\theta_2} \varphi(\theta)h(\theta)dF(\theta)$. The net gain from pooling is therefore $-COV[\varphi(\theta), h(\theta) : \theta \in (\theta_1, \theta_2)] \int_{\theta_1}^{\theta_2} dF(\theta)$. As a result, pooling will be optimal if and only if this covariance is non-positive for all possible non-decreasing truthful filters φ that lead to some type of separation. As shown below, a necessary and sufficient condition over h for this to be the case turns out to be the following:

Definition 2 (Condition *NC*). *h satisfies condition *NC* over the interval (θ_1, θ_2) if, for all $\theta \in (\theta_1, \theta_2)$, we have*

$$E[h(z) : z \in (\theta_1, \theta)] \geq E[h(z) : z \in (\theta, \theta_2)].$$

Figure 2 illustrates this condition. Notice that whenever h is decreasing over an interval, e.g., (θ_A, θ_B) , this condition will be readily satisfied. It will also hold under some types of oscillations. In the figure, for example, *NC* is satisfied over the

FIGURE 2. Condition NC

larger interval (θ_1, θ_2) . The reason is that the region where h is above its average lies to the left of the region where h falls below its average, suggesting already that pooling is desirable. More complicated oscillations are also allowed (as long as regions above average always precede regions below average). Condition NC –for (weakly) negative covariance– leads to the following characterization of covariance:

Lemma 4. h satisfies condition NC over (θ_1, θ_2) if and only if

$$COV [x(\theta), h(\theta) : \theta \in (\theta_1, \theta_2)] \leq 0$$

for every non-decreasing function $x : (\theta_1, \theta_2) \rightarrow \mathbb{R}$.

Due to this negative covariance, an optimal truthful filter will pool the largest, or “maximal,” intervals over which NC holds (such as interval (θ_1, θ_2) in figure 2).⁸ It will thus be useful to characterize the family of such intervals. Let the collection of open intervals \mathcal{P} , with typical element P , satisfy the following conditions:

- (1) For every interval (θ_1, θ_2) over which h is decreasing, there exists a $P \in \mathcal{P}$ such that $(\theta_1, \theta_2) \subset P$.
- (2) Over every $P \in \mathcal{P}$, h satisfies condition NC .
- (3) Every $P \in \mathcal{P}$ is “maximal,” i.e., there does not exist a larger interval $(\theta_1, \theta_2) \supset P$ over which h satisfies condition NC .

⁸Conversely, if an interval violates NC for some θ , splitting this interval into two pooling intervals will dominate pooling the full interval together.

Remark 1. *There exists a unique collection \mathcal{P} satisfying conditions (1)-(3). Moreover, the intervals in \mathcal{P} are disjoint and two intervals are never adjacent.*

The key behind the proof of remark 1 is that whenever two non-disjoint or adjacent intervals each satisfy NC , then the union of these intervals also satisfies NC , implying that the original intervals were not maximal to begin with. This same reasoning leads to uniqueness.

From condition NC we can infer the boundary conditions for every interval $P \in \mathcal{P}$. Letting $\inf P$ and $\sup P$ denote the extremes of P , these boundary conditions are given by

$$h(\inf P) \geq E[h(\theta) : \theta \in P] = h(\sup P),$$

with strict equality when $\inf P > \theta_L$.

Consider the first inequality. If $h(\inf P) < E[h(\theta) : \theta \in P]$, then NC would be violated for every θ close to $\inf P$. On the other hand, if $h(\inf P) > E[h(\theta) : \theta \in P]$, then P could be extended to the left while still satisfying NC (unless $\inf P = \theta_L$), a contradiction to the fact that P is maximal. A symmetric reasoning holds for the upper boundary, with the exception that it will never be the case that $\sup P = \theta_H$, and therefore we obtain an equality. The reason why $\sup P$ is always smaller than θ_H is that, over a neighborhood of θ_H , h is always increasing and larger than any value of h to the left of this neighborhood, which would violate NC .⁹ This will in turn imply that it is never optimal to pool the highest types.

The boundary conditions are illustrated in figure 2 for the maximal interval (θ_1, θ_2) . As shown in the next section, these conditions are closely related to the “ironing” conditions in the MRMR models.

Theorem 1 characterizes an optimal truthful filter, which consists of alternating pooling and separating intervals according to \mathcal{P} :

Theorem 1. *In the benchmark model, an optimal truthful filter is given by*

$$\varphi(\theta) = \begin{cases} E[\theta \in P] & \text{if } \theta \in P \text{ for some } P \in \mathcal{P}, \\ \theta & \text{otherwise.} \end{cases}$$

This need not be the uniquely optimal truthful filter. The reason is that, for some intervals in \mathcal{P} , the monopolist might be indifferent between either pooling the full intervals or separating some of their subsets from each other. For example, if h is flat

⁹This occurs because the second term in h will vanish as $\theta \rightarrow \theta_H$ and $F(\theta) \rightarrow 1$.

over a given interval, any policy for this interval will deliver the same payoff. (For regions outside \mathcal{P} , on the other hand, where h is increasing, full separation is always uniquely optimal.) In order to rule out this indifference within \mathcal{P} , we would need to rule out a “non-generic” type of h functions (such as functions containing flat regions) that allow for two consecutive intervals to satisfy condition NC , together with the boundary conditions above.

Definition 3 (Genericity). *The function h is generic if there do not exist three types $\theta_1 < \theta_2 < \theta_3$ such that:*

$$h(\theta_1) \geq h(\theta_2) = h(\theta_3) = E[h(\theta) : \theta \in (\theta_1, \theta_2)] = E[h(\theta) : \theta \in (\theta_2, \theta_3)],$$

with strict equality whenever $\theta_1 > \theta_L$.

A non-generic function would need to satisfy four equations with only three unknowns. When h is generic, it will have a negative covariance with any non-decreasing function (that is not constant) over any interval in \mathcal{P} .

Lemma 5. *Suppose h is generic. Then, for any $P \in \mathcal{P}$,*

$$COV [x(\theta), h(\theta) : \theta \in P] < 0$$

for any non-constant and non-decreasing function $x : P \rightarrow \mathbb{R}$.

Once this covariance is negative, it will be strictly optimal to fully pool every interval in \mathcal{P} , implying that the truthful filter described in theorem 1 will be uniquely optimal up to a measure-zero subset.¹⁰

Corollary 1. *Suppose h is generic. In the benchmark model, the uniquely optimal truthful filter (up to a measure-zero subset) is given by*

$$\varphi(\theta) = \begin{cases} E[\theta \in P] & \text{if } \theta \in P \text{ for some } P \in \mathcal{P}, \\ \theta & \text{otherwise.} \end{cases}$$

4. IMPLEMENTATION

Once we know the optimal allocation of social image φ , we can solve for the optimal pricing policy using the envelope theorem in lemma 3:

$$t(\theta) = \varphi(\theta)v(\theta) - \int_{\theta_L}^{\theta} \varphi(z)v'(z)dz - \theta_L v(\theta_L).$$

¹⁰This measure-zero subset corresponds to the extremes of the intervals in \mathcal{P} , which may be included in the pooling regions without affecting profits.

In order to implement the optimal allocation in the original game, the monopolist can simply offer the menu of signals $\mathcal{M} = t([\theta_L, \theta_H]) \cup \{0\}$, and have consumers coordinate on the desired equilibrium. (Offering additional signals would be of no use and, as illustrated below, could even interfere with this equilibrium.) Recall that, whenever pooling takes place, φ will become discontinuous. From the envelope theorem, t will inherit any such discontinuities and, as a result, the image $t([\theta_L, \theta_H])$ will become non-convex. The resulting gaps imply a restricted variety of signals, forcing some consumers to pool.

This restricted variety would not arise under perfect competition. In a signaling context, perfect competition requires two things. First, any signal demanded should be available. This can be modeled, for example, by assuming that consumers have a technology to burn money in public, and social inferences are formed based on the amount of money that each consumer burns.¹¹ Second, and more elusively, equilibrium beliefs should actually reward expensive signals.¹² For example, there is always an equilibrium where no money is burnt and the belief associated to any positive level of burning is equal to θ_L . (More generally, any money-burning schedule t that satisfies the above envelope theorem for some truthful filter φ can be sustained as an equilibrium.) This aspect of perfect competition can be modeled by assuming that beliefs must satisfy an equilibrium refinement requiring that off-the-path beliefs are “reasonable,” in the sense that the beliefs associated to a given deviation are related to the types most likely to perform this deviation in the first place. An example of such a refinement is the *D1* criterion (Banks and Sobel [1987]), under which the unique signaling equilibrium leads to full separation (i.e., the identity truthful filter), and therefore implies that a convex set of signals is consumed in equilibrium.¹³

¹¹It is crucial that this money burning can be verified by the relevant audience. In practice this may be hard to achieve. A more effective signal may be, for example, a donation that is published in an alumni magazine, or a luxury car that can be easily displayed. In other words, frictions in the transmission of information may confer market power to well-known brands.

¹²In practice, fashionable brands may have an advantage over merely expensive brands, again leading to market power.

¹³In this context, under criterion *D1*, a signal t occurring off the path of play cannot lead to an inference equal to θ if there is a type $\theta' \neq \theta$ for which the set of inferences that would induce such a deviation to t is strictly larger than the set of inferences that would induce type θ to deviate. To see that *D1* rules out pooling, suppose there is an equilibrium in which types (θ_1, θ_2) are pooled while sending signal t . Under *D1*, a deviation to $t + \epsilon$ for any small $\epsilon > 0$ (which must occur off the path of play) must lead to an inference of at least θ_2 (since the set of inferences that would

This competitive allocation with full separation will maximize total surplus. The reason is that, under single crossing, pooling is always inefficient since it imposes a cost over the high types within the pool, relative to full separation, which is larger than the benefit received by the lower types in this pool. Of course, for this competitive surplus to be materialized we would require that, rather than literally burning resources, consumers transfer them to a third party.¹⁴

5. INTRINSIC QUALITY

In practice, conspicuous goods typically provide some intrinsic value in addition to their desirability as social signals. This section extends the model by introducing an intrinsic quality component $q \in \mathbb{R}_+$ in the consumers' preferences (e.g., a preference for a more accurate watch), while allowing the monopolist to produce goods in a range of intrinsic qualities with a corresponding production cost, as in MRMR. The problem becomes one of two-dimensional screening, where the monopolist sorts consumers by allocating both image φ and quality q .

A good is now a pair of intrinsic quality and price $\langle q, t \rangle$. For tractability, I assume that the payoff for consumer θ has a simple separable form:

$$U(\pi, q, t; \theta) = (\pi + q)v(\theta) - t,$$

which combines the assumptions that the marginal utilities of both social image and intrinsic quality are constant, and proportional to each other across types. After a normalization, this reduces to assuming that both these marginal utilities are equal to $v(\theta)$.

Associated to the level of q consumed by each individual, there is a production cost $c(q) \in \mathbb{R}_+$, where c is smooth, increasing, and strictly convex, with $c'(\infty) > v(\theta_H)$. Letting $\langle q, t \rangle(\theta)$ denote the good consumed by type θ , the payoff for the monopolist

induce θ_2 to deviate to $t + \epsilon$ is strictly larger than the corresponding set for any $\theta < \theta_2$). But this would induce some types in the pool to deviate to $t + \epsilon$ whenever ϵ is sufficiently small.

¹⁴The competitive process will nevertheless lead to a destruction of surplus whenever suppliers, in an effort to attract clients, over-invest in features that are intrinsically valuable to consumers (such as a disproportionate engine in a sports car). See Bagwell and Bernheim [1996], and Becker, Murphy and Glaeser [2000] for a formal argument. As a result, total surplus may happen to be larger under a monopolistic supply.

is now

$$\int_{\theta_L}^{\theta_H} \{t(\theta) - c(q(\theta))\} dF(\theta).$$

We must now specify the way inferences π are formed, which in turn will depend on the observability of q . For example, the number of jewels inside the machinery of a watch (which presumably contribute to its precision) are only observable if they are advertised on the exterior of the watch, but the size of a car is a conspicuous feature that will automatically affect the calculation of π . For concreteness, I will consider the case where q is always observable, so that π corresponds to the Bayesian posterior of type given q and t . The results will also hold for the alternative case.¹⁵

Analogous to the benchmark model, we can write the monopolist's problem as one of direct (truthful) revelation where she offers a three-dimensional schedule $\langle \varphi, q, t \rangle (\cdot)$ as a function of the announced type, subject to the corresponding participation and incentive compatibility constraints. The analysis in the benchmark model leading to the simplified problem (3), i.e., lemmas 1-3, can be generalized to this case in the following way (with the purpose of sending t to the background):

Define $x := \varphi + q$, and notice that this sum will now take the place previously occupied by φ in all the incentive and participation constraints. Lemmas 1 and 3 (which simplified these constraints) will therefore hold for $\varphi = x$. On the other hand, lemma 2 (which stated that the Bayesian inference constraint was equivalent to the condition that φ was a truthful filter) will now be affected by the fact that inferences are also influenced by the presence of q . The Bayesian inference constraint now becomes

$$(BI) \quad \varphi(\theta) = E[\theta' : \langle q, t \rangle(\theta') = \langle q, t \rangle(\theta)] \text{ for all } \theta.$$

This constraint will imply that whenever two different types receive the same social image, they must also receive the same intrinsic quality, which follows because a difference in quality would reveal their difference in type. Let (*Pool*) denote this constraint:

$$(Pool) \quad \varphi(\theta) = \varphi(\theta') \Rightarrow q(\theta) = q(\theta'), \text{ for all } \theta, \theta'.$$

Lemma 6 replaces lemma 2 while incorporating this new constraint.

¹⁵As in the benchmark model, adding additional features with no intrinsic value, such as color or style, will be of no use.

Lemma 6. *Under any incentive compatible schedule, (BI) is satisfied if and only if both (TF) and (Pool) hold.*

Using these results, the monopolist's problem can be expressed as

$$(4) \quad \max_{\varphi, q} \int_{\theta_L}^{\theta_H} \{ \langle \varphi + q \rangle (\theta) h(\theta) - c(q(\theta)) \} dF(\theta)$$

s.t. $\langle \varphi + q \rangle$ is non-decreasing,
(TF), and (Pool).

Problem (4), in sum, generalizes the benchmark problem (3) by adding the cost of intrinsic quality c , adding the (Pool) constraint, and replacing φ with $x = \varphi + q$. For concreteness I will assume that h is generic, so that any two solutions will coincide over a full measure subset. I focus on only one such solution, and for simplicity I refer to it as unique.

The complication in problem (4) arises from the fact that the joint monotonicity constraint will allow the social-image and intrinsic-quality schedules to interact in non-trivial ways. An increasing φ will allow for a decreasing q , and vice versa, so that the analysis above, or that in MRMR, which restrict φ and q to each be monotonic, cannot be directly applied.

Despite this fact, the solution will turn out to have a simple form, with schedules that are each non-decreasing. This can be readily shown for φ . In particular, whenever the schedule $\langle \varphi + q \rangle$ is increasing, the inference process implies that φ must also be increasing, and within regions where $\langle \varphi + q \rangle$ is constant, due to the convexity of c , it will be optimal for both φ and q to be constant:

Lemma 7. *In problem (4), the optimal truthful filter φ is non-decreasing.*

Moreover, the optimal menu will induce the same pooling regions as the optimal menu in the benchmark model, allowing for a generalization of theorem 1. It will be useful to define a partition of the type space into a collection of intervals \mathcal{I} , based on \mathcal{P} , as follows. Let \mathcal{S} denote the collection of maximal intervals complementary to \mathcal{P} , i.e., the largest intervals in $[\theta_L, \theta_H] \setminus \mathcal{P}$. Recall that, over these intervals, h will be increasing. Define $\mathcal{I} := \mathcal{S} \cup \mathcal{P}$, which partitions the type space into alternating intervals over which either h is increasing or condition *NC* holds.

Consider now a relaxed approach to problem (4) where constraint (Pool) is ignored, and the schedule $\langle \varphi + q \rangle$ is only required to be non-decreasing within each

separate interval $I \in \mathcal{I}$:

$$(5) \quad \max_{\varphi, q} \int_{\theta_L}^{\theta_H} \varphi(\theta) h(\theta) dF(\theta) + \int_{\theta_L}^{\theta_H} \{q(\theta) h(\theta) - c(q(\theta))\} dF(\theta)$$

s.t. $\langle \varphi + q \rangle$ is non-decreasing over I for all $I \in \mathcal{I}$, and
 φ is non-decreasing and satisfies (TF) .

The solution to this relaxed problem (5) will turn out to satisfy all the constraints in problem (4).

First, we can readily obtain the values for $q(\theta)$ over all the intervals in \mathcal{S} . Over these intervals, the fully unconstrained optimal values for $q(\theta)$, namely $q(\theta) = \arg \max_{y \geq 0} \{yh(\theta) - c(y)\}$, are non-decreasing in θ (precisely because $h(\theta)$ is increasing in θ), and therefore satisfy all the constraints.

Lemma 8. *In the relaxed problem (5), for every interval $S \in \mathcal{S}$, and all $\theta \in S$, the optimal level of $q(\theta)$ solves $\max_{y \geq 0} \{yh(\theta) - c(y)\}$.*

Next, we can show that over each one of the remaining intervals $P \in \mathcal{P}$, the optimal q will be constant. Over these intervals, condition NC implies that a constant schedule $\langle \varphi + q \rangle$ will be preferred over an increasing one. Moreover, since c is convex, a constant $\langle \varphi + q \rangle$ can be most cheaply produced using a constant q .

Lemma 9. *In the relaxed problem (5), the optimal schedule q is constant over every interval $P \in \mathcal{P}$.*

Consider an interval $P \in \mathcal{P}$. Once we know that q must be constant over P , we can obtain the optimal level within this interval, $q(P)$, by solving: $\max_{q(P)} \int_P \{q(P)h(\theta) - c(q(P))\} dF(\theta) = \max_{y \geq 0} \{yE[h(\theta) : \theta \in P] - c(y)\}$.

Lemma 10. *In the relaxed problem (5), for every interval $P \in \mathcal{P}$, and all $\theta \in P$, the optimal level of $q(\theta)$ solves $\max_{y \geq 0} \{yE[h(\theta) : \theta \in P] - c(y)\}$.*

Consider the schedule q that results from combining lemmas 8 and 10. Due to the boundary conditions for the intervals $P \in \mathcal{P}$ (i.e., $h(\inf P) \geq E[h(\theta) : \theta \in P] = h(\sup P)$, with strict equality when $\inf P = \theta_L$), this schedule is continuous, and therefore non-decreasing over the entire type space. Moreover, given this schedule q , the optimal φ obtained in theorem 1 for the benchmark model (which pools an interval if and only if it belongs to \mathcal{P}) will be uniquely optimal in problem (5). The reason is that it uniquely maximizes the objective among all non-decreasing truthful

filters (Corollary 1, given genericity) while satisfying all the joint monotonicity constraints when combined with q .

We can now return to the original problem (4). Notice, from lemma 9, that the above schedules satisfy constraint (*Pool*), and since each schedule is non-decreasing over the entire type space, their sum is also non-decreasing. As a result, these schedules indeed solve problem (4).¹⁶

Theorem 2. *Suppose h is generic. In the model with intrinsic quality, the uniquely optimal schedule $\langle \varphi, q \rangle$ (up to a measure-zero subset) combines the optimal truthful filter φ from the benchmark model (theorem 1), and the intrinsic quality schedule q from lemmas 8 and 10.*

For this simple model, the optimal intrinsic quality schedule q corresponds to the optimal schedule from MRMR. Indeed, the *NC* conditions that characterize the pooling intervals become an equivalent way to characterize the ironing conditions in MRMR. While the ironing conditions focus on minimizing deviations from the schedule that maximizes virtual surplus, the *NC* conditions focus on covariance, an approach that also allows us to solve for the (discontinuous) allocation of social image.

The presence of a signaling motive, however, will lead to a pricing schedule that differs from that in MRMR. Once we know the optimal allocation of φ and q , the optimal pricing schedule follows from the envelope equation

$$t(\theta) = \langle \varphi + q \rangle (\theta)v(\theta) - \int_{\theta_L}^{\theta} \langle \varphi + q \rangle (z)v'(z)dz - \theta_L v(\theta_L).$$

This schedule will be discontinuous whenever the allocation $\langle \varphi + q \rangle$ is discontinuous. An example is given in figure 3. The left panel shows a pricing schedule, as a function of q , in the absence of a signaling motive. This continuous schedule corresponds to that in MRMR, where pooling simply leads to a kink (notice that the associated set of prices will be convex). The right panel shows the corresponding schedule when signaling is present, which is steeper than before, and jumps under pooling since this causes a discontinuity in φ . The result is a non-convex set of prices that limits the set of available signals, i.e., a restricted variety of goods.

¹⁶Notice that q may be flat (and equal to zero) at the low end of the type space even when this low end lies outside \mathcal{P} (e.g., when h is negative), implying that the converse of (*Pool*) need not hold.




FIGURE 3. Discontinuities \Rightarrow Menu Non-Convexities

As in the benchmark model, pooling will always be connected to discontinuous prices. But notice that this connection is no longer imposed by the structure of the model. In fact, pooling need not lead to a discontinuous allocation $\langle \varphi + q \rangle$ since q could always compensate for any discontinuity in φ . Conversely, any discontinuity in prices could be caused entirely by a discontinuity in q . However, the implication of theorem 2 is that price gaps cannot be explained by jumps in intrinsic quality, which is always continuous across types.¹⁷ Rather, these gaps follow from a desire to pool types together in the signaling game.

6. SOME EXAMPLES AND CONCLUSIONS

When a signaling motive is present, the value of a conspicuous good will depend both on its intrinsic properties and on the set of consumers purchasing this good in equilibrium. A discriminating monopolist will exploit this fact when designing and pricing her optimal menu of goods. The main implication of the model is that a restricted variety of goods that forces consumers to pool can lead to higher revenues, a phenomena that is absent when consumers care only about intrinsic characteristics. This result may help account for some peculiar pricing strategies

¹⁷The optimal q is continuous across types because, in terms of intrinsic quality, there is no reason why similar types should not receive a similar allocation. In contrast, discontinuities in social image become the only way to downward-distort its allocation in order to extract information rents.

that surround conspicuous goods in practice. I show some examples concerning luxury cars, watches, and charitable contributions.¹⁸

The 2003 BMW auto collection includes 24 models ranging from about \$28K to \$137K. But these models are unevenly distributed across the price range. In particular, the most expensive car is \$64.5K more expensive than the next model (which costs about half as much at \$72.5), while the average price gap among the remaining cars is only \$2K.¹⁹ Although it is unlikely that the supplier actually considers a detailed formal model when selecting this strategy, it suffices for the present argument that the supplier does take into account the fact that the value of each car will actually depend on the entire menu of cars being offered. For instance, adding an intermediate \$100K car may attract some lower-end consumers, but its very presence will also reduce the value of the \$72.5 car (in terms of social status), further reducing its sales in favor of cheaper models or even an outside option, a simple effect that favors a restricted variety.

Similar patterns also appear in luxury watches. Cartier's 2002 list price for its most expensive Jewelry watches (typically with exterior Diamonds) includes almost 140 models, but the watches are again unevenly distributed in the \$8K - \$330K price range. The price difference between the most and second most expensive is \$51K, and between the second and the third is \$67K, while the average gap among the others is merely \$1.5K. Introducing new models in this case seems almost trivial since a major component of the difference across models is simply the number of exterior diamonds. Price differences within particular models are even more pronounced. The Jewelry Pasha model (noticeable due to its circular case) comes in 18 models ranging from \$10K to \$221K, but the second most expensive is only \$89K.²⁰ As in the car example, a simple reason for not offering an intermediate

¹⁸The suppliers considered below presumably have some degree of market power arising either from a well-known and fashionable brand, or a market niche (in the case of charitable contributions). Suggestive of market power, profits on a luxury car may exceed \$20,000, while the profit per vehicle on mainstream cars is often as little as \$300 (The Economist, Jan. 4th-10th, 2003). BMW, for example, recently bought the rights to the well-known Rolls Royce name, requiring a tight bidding war with Volkswagen even though the existing Rolls Royce technology was seriously outdated.

¹⁹The 2000 Mercedes-Benz menu (with 20 models) displays a similar feature, with a \$43K price gap between the most and second most expensive cars, while the average gap among the remaining cars is \$3K (autos.yahoo.com, 09/2000, bmwusa.com, 01/2003).

²⁰The Rolex Cellini (as well as other Rolex models) follows an analogous pattern. Out of seven models, the cheapest six are no more than 25% away, while the most expensive is more than three

model is that its presence is likely to reduce the value of the \$89K watch (and lose customers beyond those that merely purchase the new option).

Finally, a well-known example of restricted variety occurs when universities report charitable contributions (Glazer and Konrad [1996], Harbaugh [1998]). These contributions are typically reported by categories (e.g., Mr. X contributed between \$10K and \$25K), as opposed to the exact amount (in practice, almost every donor contributes at the lowest end of these categories), leading to large variety gaps. For instance, the last four lower ends for the annual contributions to Stanford University are typically \$10K, \$25K, \$50K and \$100K (e.g., it is impossible to purchase a \$75K signal). A more extreme example comes from the University of Chicago's GSB campaign to fund a new \$125 million building. Donations are again reported in intervals, but in this case the low ends of the last four intervals are \$500K, \$1 million, \$5 million, and \$10 million.²¹ Consistent with the motive of extracting information rents, and thus full separation at the top, variety restrictions are not observed for the largest donors, who are typically given a category of their own, or even receive special recognition, say, by naming a building after them (and not after smaller donors as well).²²

As a starting point, the present model endowed the signal provider with full market power, which in a signaling environment is tantamount to the power of manipulating beliefs (subject only to Bayesian inference constraints). Although

times the second one. One can visually notice these gaps based on 'discontinuous' increases in the number and size of exterior diamonds, jumping in some cases from around 50 (e.g., placed on a thin band around the case) to several hundred covering the entire surface, bracelet included. List prices for Cartier, 11/2002, and Rolex, 03/1999, provided by Berger Jewelry.

²¹The Stanford Fund, 1998-99, gsbwww.uchicago.edu/campaign, 01/2003.

The model cannot account for the use of multiple consecutive categories. Rather, it would prescribe alternating regions with exact reporting. (Would this be tasteless?) However, the signaling argument can explain why categories can improve on exact reporting to begin with, and also why some of these categories are so large, e.g., why Stanford can potentially lose revenues by adding an intermediate \$75K option.

²²Harbaugh [1998] models the use of categories using an alternative approach, which does not concern the problem of signaling a hidden type. Instead, he assumes that donations map into a social "prestige" according to an exogenous function. In contrast to a signaling structure, the prestige associated with a given donation is independent of the particular set of consumers making that same donation in equilibrium (which is the focus of the present discussion). Moreover, his model deals with a finite set of types, so categories will no longer imply a bunching of types (in fact, all of the categories that are used to prove that exact reporting is suboptimal contain one type at most).

analytically convenient, this assumption abstracts from the issue of how market power arises in the first place, or how multiple signal suppliers would interact. The study of such matters is left for future work.

7. APPENDIX: OMITTED PROOFS

Proof of lemma 1. Redundancy of the participation constraints follows from combining (IC) with θ_L 's participation constraint. Equality for the lowest type's participation constraint follows from the fact that, if it had slack, the monopolist could raise the entire t schedule by some small $\epsilon > 0$, which would produce higher profits while satisfying all the constraints. ■

Proof of lemma 2. From (IC'), for any θ, θ' , $t(\theta') = t(\theta) \Leftrightarrow \varphi(\theta') = \varphi(\theta)$. Therefore, $E[\theta' : t(\theta') = t(\theta)] = E[\theta' : \varphi(\theta') = \varphi(\theta)]$ for all θ . ■

Proof of lemma 3. For necessity (\Rightarrow), the first equation follows from theorem 2 of Milgrom and Segal (2002), which implies that the equilibrium payoff for the consumers, $\varphi(\theta)v(\theta) - t(\theta)$, must be absolutely continuous in θ , a consequence of (IC'). That φ is non-decreasing follows directly from (IC').

For sufficiency (\Leftarrow), we have that (P') is trivially satisfied. For (IC'), let $\theta > \theta'$. From the first equation, $t(\theta) - t(\theta') = \varphi(\theta)v(\theta) - \varphi(\theta')v(\theta') - \int_{\theta'}^{\theta} \varphi(z)v'(z)dz \leq \varphi(\theta)v(\theta) - \varphi(\theta')v(\theta') - \int_{\theta'}^{\theta} \varphi(\theta')v'(z)dz = [\varphi(\theta) - \varphi(\theta')]v(\theta)$, where the inequality follows from the fact that φ is non-decreasing. The previous relations imply that type θ will not profit from announcing θ' instead of θ . The argument for $\theta < \theta'$ is symmetric. ■

Proof of lemma 4. I begin with necessity (\Rightarrow). Let x be an arbitrary non-decreasing function, and let $\bar{h} := E[h(\theta) : \theta \in (\theta_1, \theta_2)]$. Define recursively, for each $n = 0, 1, 2, \dots$, two functions, $h^n, x^n : (\theta_1, \theta_2) \rightarrow \mathbb{R}$, and three finite collections of intervals, $\{A_1^n, A_2^n, \dots\}$, $\{B_1^n, B_2^n, \dots\}$, and $\{C_1^n, C_2^n, \dots\}$, as follows:

- $h^0 \equiv h, x^0 \equiv x$.
- $\{A_1^n, A_2^n, \dots\}$ equals the collection of maximal intervals in (θ_1, θ_2) over which $h^n(\theta) \geq \bar{h}$, ordered so that $A_i^n < A_{i+1}^n$.
- $\{B_1^n, B_2^n, \dots\}$ equals the complementary set of maximal intervals in (θ_1, θ_2) , also ordered so that $B_i^n < B_{i+1}^n$.

(Notice that the intervals in these two sets must alternate and, since h satisfies NC, the left-most interval must be A_1^n , and the right-most interval will be an element of $\{A_1^n, A_2^n, \dots\}$ only when $h \equiv \bar{h}$ over such interval.)

- $C_i^n = A_i^n \cup B_i^n$, or $C_i^n = A_i^n$ if A_i^n is the right-most interval.
- And, for every $n \geq 1$, i , and $\theta \in C_i^{n-1}$,

$$h^n(\theta) = E [h^{n-1}(z) : z \in C_i^{n-1}], \quad x^n(\theta) = x^{n-1}(\max\{z \in A_i^{n-1}\}).$$

(Notice that every h^n also satisfies condition *NC*.)

Since every step merges at least two intervals, there exists an n such that $C_1^{n-1} = (\theta_1, \theta_2)$, and therefore $h^n(\cdot) \equiv \bar{h}$. Let N be the smallest integer such that this holds. We now have

$$\begin{aligned} \text{COV}[x(\theta), h(\theta) : \theta \in (\theta_1, \theta_2)] &= \int_{\theta_1}^{\theta_2} x^0(\theta)[h^0(\theta) - \bar{h}]dF(\theta) \\ &\leq \int_{\theta_1}^{\theta_2} x^1(\theta)[h^1(\theta) - \bar{h}]dF(\theta) = \int_{\theta_1}^{\theta_2} x^1(\theta)[h^1(\theta) - \bar{h}]dF(\theta) \leq \dots \\ &\leq \int_{\theta_1}^{\theta_2} x^N(\theta)[h^N(\theta) - \bar{h}]dF(\theta) = \int_{\theta_1}^{\theta_2} x^N(\theta)[h^N(\theta) - \bar{h}]dF(\theta) = 0. \end{aligned}$$

For sufficiency (\Leftarrow), suppose condition *NC* is violated for some θ . Then, any “one-step” function x that has a constant value over (θ_1, θ) , and a higher constant value over $[\theta, \theta_2)$, will have a positive covariance with h over (θ_1, θ_2) . ■

Proof of remark 1. I begin with existence. Notice that the collection of all open intervals over which h is decreasing satisfies conditions (1) and (2). Denote this collection by \mathcal{D} . Now, for each interval $D \in \mathcal{D}$, let $P(D)$ be the largest open interval containing D while satisfying condition *NC*. ($P(D)$ exists because the set of intervals satisfying these constraints is compact, i.e., it corresponds to a closed and bounded set in \mathbb{R}^2 .) But notice that the collection $\{P(D) : D \in \mathcal{D}\}$ satisfies conditions (1)-(3).²³

For uniqueness, suppose towards a contradiction that both \mathcal{P} and $\mathcal{P}' \neq \mathcal{P}$ satisfy (1)-(3). Since the complement of $\cup_{P \in \mathcal{P}} P$ consists of intervals over which h is increasing, it cannot be the case that some $P' \in \mathcal{P}'$ lies outside $\cup_{P \in \mathcal{P}} P$, or vice versa (otherwise condition (2) would not hold). Therefore, there must exist a pair of intervals $P \in \mathcal{P}$, and $P' \in \mathcal{P}'$, such that $P \cap P' \neq \emptyset$, and $P \neq P'$. Suppose without loss that $P \leq P'$. I now show that $P \cup P'$ satisfies condition *NC*, a contradiction to (3). Let $\theta \in P \cup P'$. We have two cases to consider, according to whether or not $\theta \in P \cap P'$. Suppose first that this is the case. From condition (2), applied to both

²³Notice also that we might have $P(D) = P(D')$ for two different intervals $D, D' \in \mathcal{D}$.

P and P' , we obtain the desired inequality:

$$\begin{aligned} \int_{\inf P}^{\theta} h(z)dF(z : z \in (\inf P, \theta)) &\geq \int_{\inf P}^{\sup P} h(z)dF(z : z \in (\inf P, \sup P)) \geq \\ \int_{\inf P'}^{\sup P'} h(z)dF(z : z \in (\inf P', \sup P')) &\geq \int_{\theta}^{\sup P'} h(z)dF(z : z \in (\theta, \sup P')). \end{aligned}$$

The case where $\theta \notin P \cap P'$ is similar, and left to the reader.

Finally, that the intervals in \mathcal{P} are disjoint and never adjacent follows from the same reasoning as above, i.e., the union of two intervals in \mathcal{P} that intersect, or are adjacent, would also satisfy NC , a contradiction to (3). ■

Proof of theorem 1. Suppose \mathcal{P} is non-empty, otherwise the theorem would follow from proposition 1. Let \mathcal{S} denote the collection of maximal intervals complementary to \mathcal{P} , i.e., the largest intervals in $[\theta_L, \theta_H] \setminus \mathcal{P}$, and let $\mathcal{I} = \mathcal{S} \cup \mathcal{P}$ (which constitutes a partition of $[\theta_L, \theta_H]$). Also let Φ be the set of truthful filters φ satisfying the following condition:

$$\text{For every } P \in \mathcal{P}, \theta \in P, \text{ and } \theta' \notin P, \text{ we have } \varphi(\theta) \neq \varphi(\theta').$$

I proceed in two steps. I first show that every truthful filter φ' outside Φ is dominated by some truthful filter within this set, and then show that the truthful filter defined in the theorem, which belongs to Φ , weakly dominates every other element of Φ .

Let $\varphi' \notin \Phi$, and suppose it maximally pools the collection of intervals \mathcal{T} . Suppose we replace this filter with φ , which, instead, maximally pools the collection of smaller intervals $\mathcal{T} \cap \mathcal{I} := (T \cap I)_{T \in \mathcal{T}, I \in \mathcal{I}}$, and is otherwise equal to θ . Notice that φ belongs to Φ . Also, from the properties of \mathcal{P} , for every $T \in \mathcal{T}$, both the average $\bar{h}(T \cap I) := E[h(\theta) : \theta \in T \cap I]$ and the expectation $E[\theta \in T \cap I]$ must be increasing in I . The change in the objective is given by

$$\begin{aligned} \sum_{T \in \mathcal{T}} \int_T \{\varphi(\theta) - \varphi'(\theta)\} h(\theta) dF(\theta) &= \sum_{T \in \mathcal{T}} \sum_{I \in \mathcal{I}} \int_{T \cap I} \{\varphi(\theta) - \varphi'(\theta)\} h(\theta) dF(\theta) = \\ \sum_{T \in \mathcal{T}} \sum_{I \in \mathcal{I}} \{E[\theta \in T \cap I] - E[\theta \in T]\} \bar{h}(T \cap I) \int_{T \cap I} dF(\theta) &> 0, \end{aligned}$$

where the inequality follows because each term in the sum, which is proportional to the covariance between $E[\theta \in T \cap I]$ and $\bar{h}(T \cap I)$ across intervals I , is always non-negative, and any term in this sum becomes positive whenever T intersects

more than one interval I (which must occur for at least one interval $T \in \mathcal{T}$ since $\varphi' \notin \Phi$). This concludes the first step of the proof.

For the second step, let φ be as defined in the theorem, and let φ' be any other element in Φ . We can write

$$(*) \quad \int_{\theta_L}^{\theta_H} \{\varphi(\theta) - \varphi'(\theta)\} h(\theta) dF(\theta) = \\ \sum_{P \in \mathcal{P}} \int_P \{E[\theta \in P] - \varphi'(\theta)\} h(\theta) dF(\theta) + \sum_{S \in \mathcal{S}} \int_S \{\theta - \varphi'(\theta)\} h(\theta) dF(\theta).$$

Consider the first sum. Since $\varphi' \in \Phi$, for every $P \in \mathcal{P}$, we must have $\int_P \varphi'(\theta) dF(\theta) = E[\theta \in P] \int_P dF(\theta)$, and therefore each term in the sum is proportional to $-COV[\varphi'(\theta), h(\theta) : \theta \in P]$. But, from lemma 4, this quantity is non-negative.

Consider now the second sum. Notice that, for any $S \in \mathcal{S}$, the problem of maximizing the integral $\int_S \varphi'(\theta) h(\theta) dF(\theta)$, subject to $\varphi' \in \Phi$, is identical to the original problem (3) with $[\theta_L, \theta_H] = S$. From proposition 1, this problem is solved by setting $\varphi'(\theta) = \theta$ for all $\theta \in S$. As a result, each term in this second sum must also be non-negative. ■

Proof of lemma 5. Let $P = (\theta_1, \theta_2) \in \mathcal{P}$, and suppose x is non-constant and non-decreasing. Also let $\theta^* \in (\theta_1, \theta_2)$ be such that either $x(\theta) > (\leq) x(\theta^*)$ for all $\theta > (<) \theta^*$, or $x(\theta) \geq (<) x(\theta^*)$ for all $\theta > (<) \theta^*$. Suppose towards a contradiction that $COV[x(\theta), h(\theta) : \theta \in (\theta_1, \theta_2)] = 0$. From the algorithm in the proof of lemma 4, this requires that, for every $n \leq N - 2$, x^n is constant over C_i^n for all i . Therefore, $\theta^* = \min\{\theta \in C_i^{n-1}\}$ for some $i > 1$, and as a result $h(\theta^*) = \bar{h}$. This implies in turn that x^n is non-constant over (θ_1, θ_2) for all $n \leq N - 1$. So that, in particular, x^{N-1} is non-constant over $C_1^{N-1} = (\theta_1, \theta_2)$, and $\theta^* = \min\{\theta \in C_2^{N-2}\}$, implying that $C_1^{N-2} = (\theta_1, \theta^*)$ and $C_2^{N-2} = [\theta^*, \theta_2]$. From the latter equalities we must have: $h^{N-1}(\theta) = E[h(z) : z \in (\theta_1, \theta^*)]$ for all $\theta < \theta^*$, and $h^{N-1}(\theta) = E[h(z) : z \in [\theta^*, \theta_2]]$ for all $\theta \geq \theta^*$. But, since h is generic, and $h(\theta^*) = \bar{h}$, the boundary conditions for $(\theta_1, \theta_2) \in \mathcal{P}$ imply that the first expectation must be larger than the second, i.e., $h^{N-1}((\theta_1, \theta^*)) > h^{N-1}([\theta^*, \theta_2])$. By combining this with the fact that x^{N-1} is non-constant, we obtain a strict inequality in the final step of the algorithm, a contradiction:

$$\int_{\theta_1}^{\theta_2} x^{N-1}(\theta) [h^{N-1}(\theta) - \bar{h}] dF(\theta) < \int_{\theta_1}^{\theta_2} x^N(\theta) [h^{N-1}(\theta) - \bar{h}] dF(\theta) = 0.$$

■

Proof of Corollary 1. Consider the same proof of theorem 1, but invoking lemma 5 instead of lemma 4 after equation (*). Notice that, since h is generic, at least one of the summations in (*) must become positive whenever φ' differs from φ . ■

Proof of lemma 6. For both necessity and sufficiency it will suffice to show that, for any pair θ and θ' ,

$$(**) \quad \varphi(\theta) = \varphi(\theta') \Leftrightarrow \langle q, t \rangle(\theta) = \langle q, t \rangle(\theta').$$

For necessity (\Rightarrow), suppose towards a contradiction that (BI) holds while condition (**) fails in the “ \Rightarrow ” direction (the other direction cannot fail under (BI)), i.e., for some $\theta < \theta'$ we have that $\varphi(\theta) = \varphi(\theta')$ while $\langle q, t \rangle(\theta) \neq \langle q, t \rangle(\theta')$. From incentive compatibility (and single crossing), the sum $\langle \varphi + q \rangle$ must be non-decreasing, and $q(\theta)$ cannot be equal to $q(\theta')$ (otherwise the prices would also need to be equal). Therefore, $q(\theta) < q(\theta')$. But (from single crossing) the set of consumers purchasing $\langle \varphi, q \rangle(\theta)$ must lie below the set consuming $\langle \varphi, q \rangle(\theta')$, implying from (BI) that $\varphi(\theta) < \varphi(\theta')$, a contradiction.

For sufficiency (\Leftarrow), suppose (TF) and $(Pool)$ hold, while (**) does not. This failure must be in the “ \Leftarrow ” direction (since the opposite direction follows from $(Pool)$ and incentive compatibility). But, if for some $\theta \neq \theta'$ we have $\langle q, t \rangle(\theta) = \langle q, t \rangle(\theta')$ and $\varphi(\theta) \neq \varphi(\theta')$, then the schedule is not incentive compatible, a contradiction. ■

Proof of lemma 7. Since the schedule $\langle \varphi + q \rangle$ is non-decreasing, it will be composed of alternating intervals over which it is either increasing or constant. I proceed by showing that: (1) for all $\theta < \theta'$, $\langle \varphi + q \rangle(\theta) < \langle \varphi + q \rangle(\theta')$ implies $\varphi(\theta) < \varphi(\theta')$, and (2) whenever $\langle \varphi + q \rangle$ is constant over an interval, then φ will also be constant.

For the first claim, let $\theta < \theta'$ be such that $\langle \varphi + q \rangle(\theta) < \langle \varphi + q \rangle(\theta')$, and suppose towards a contradiction that $\varphi(\theta) \geq \varphi(\theta')$. The latter inequality, together with constraint (TF) , implies that either θ is pooled with some type larger than θ' , or θ' is pooled with some type smaller than θ , or both. But this contradicts the fact that $\langle \varphi + q \rangle$ is non-decreasing.

For the second claim, let (θ_1, θ_2) be a maximal interval over which the schedule $\langle \varphi + q \rangle$ is constant. Suppose towards a contradiction that φ is different across two (positive-measure) subsets of this interval, and therefore q is also different across these subsets (otherwise the sum $\langle \varphi + q \rangle$ would not be constant). Now let the monopolist pool all the types in (θ_1, θ_2) together by changing the schedules over this interval to $\varphi' = E[\varphi(z) : z \in (\theta_1, \theta_2)]$ and $q' = E[q(z) : z \in (\theta_1, \theta_2)]$. This change will satisfy all the constraints, and the change in the objective will be equal

to

$$\begin{aligned} & \int_{\theta_1}^{\theta_2} \{ \langle \varphi' + q' \rangle (\theta) h(\theta) - c(q') \} dF(\theta) - \int_{\theta_1}^{\theta_2} \{ \langle \varphi + q \rangle (\theta) h(\theta) - c(q(\theta)) \} dF(\theta) \\ &= \int_{\theta_1}^{\theta_2} c(q(\theta)) dF(\theta) - \int_{\theta_1}^{\theta_2} c(q') dF(\theta) > 0, \end{aligned}$$

where the inequality follows from the convexity of c , a contradiction. ■

Proof of lemma 9. Consider a schedule $\langle \varphi, q \rangle$ satisfying the constraints in problem (5), and such that q satisfies lemma 8 but is not constant over some $P \in \mathcal{P}$. I consider a specific change in this schedule, resulting in a new schedule $\langle \varphi', q' \rangle$, such that q' is constant over every interval in \mathcal{P} , and then show that this change will be profitable.

Let \mathcal{T} denote the collection of intervals T over which φ is constant. Suppose the new schedule φ' satisfies

$$\varphi'(\theta) = \begin{cases} E[\theta \in T \cap S] & \text{if } \theta \in T \cap S \text{ for some } T \in \mathcal{T} \text{ and } S \in \mathcal{S}, \\ E[\theta \in P] & \text{if } \theta \in P \text{ for some } P \in \mathcal{P}, \\ \varphi(\theta) = \theta & \text{otherwise,} \end{cases}$$

and q' satisfies

$$q'(\theta) = \begin{cases} q(\theta) & \text{if } \theta \in S \text{ for some } S \in \mathcal{S}, \\ E[q(\theta) : \theta \in P] & \text{if } \theta \in P \text{ for some } P \in \mathcal{P}. \end{cases}$$

Notice that these new schedules satisfy all the constraints in (5).

Consider now the change in the objective. In order to measure this change, it will be useful to decompose the shift from $\langle \varphi, q \rangle$ to $\langle \varphi', q' \rangle$ into a shift from $\langle \varphi, q \rangle$ to an intermediate schedule $\langle \varphi'', q' \rangle$, plus a shift from $\langle \varphi'', q' \rangle$ to $\langle \varphi', q' \rangle$. Define the intermediate schedule φ'' as follows:

$$\varphi''(\theta) = \begin{cases} \varphi(\theta) & \text{if } \theta \in S \text{ for some } S \in \mathcal{S}, \\ E[\varphi(\theta) : \theta \in P] & \text{if } \theta \in P \text{ for some } P \in \mathcal{P}. \end{cases}$$

In other words, the intermediate regime $\langle \varphi'', q' \rangle$ is equal to the original one, $\langle \varphi, q \rangle$, except for the fact that, within each interval in \mathcal{P} , both φ and q are replaced with their expected values. (φ'' will not satisfy the (TF) constraint, but it is only used as a device to calculate profits.)²⁴

²⁴Intuitively, the shift from $\langle \varphi, q \rangle$ to $\langle \varphi'', q' \rangle$ is profitable because it exploits condition NC within each $P \in \mathcal{P}$, as well as the fact that c is convex. While the shift from $\langle \varphi'', q' \rangle$ to $\langle \varphi', q' \rangle$

Consider first the change in the objective when shifting from $\langle \varphi, q \rangle$ to $\langle \varphi'', q' \rangle$. This change will only affect the intervals in \mathcal{P} . For each $P \in \mathcal{P}$, the change is equal to

$$\begin{aligned} & \int_P \{ \langle \varphi'' + q' \rangle (\theta) h(\theta) - c(q') \} dF(\theta) - \int_P \{ \langle \varphi + q \rangle (\theta) h(\theta) - c(q(\theta)) \} dF(\theta) \\ &= -COV[\langle \varphi + q \rangle (\theta), h(\theta) : \theta \in P] \int_P dF(\theta) + \int_P \{ c(q(\theta)) - c(q') \} dF(\theta) \geq 0. \end{aligned}$$

The inequality follows because both terms on the right are non-negative: The first term is non-negative because $\langle \varphi + q \rangle$ is non-decreasing and h satisfies condition *NC* over P (lemma 4). The second term is non-negative due to the convexity of c . Moreover, the inequality becomes strict for some $P \in \mathcal{P}$ whenever the claim in the lemma is violated over a positive-measure subset (due to the convexity of c).

Consider, finally, the change in the objective when shifting from $\langle \varphi'', q' \rangle$ to $\langle \varphi', q' \rangle$. After some algebra, this change is given by

$$\begin{aligned} \int_{\theta_L}^{\theta_H} [\varphi'(\theta) - \varphi''(\theta)] dF(\theta) &= \sum_{T \in \mathcal{T}} \left\{ \sum_{S \in \mathcal{S}} \{ E[\theta \in T \cap S] - E[\theta \in T] \} \int_{T \cap S} h(\theta) dF(\theta) \right. \\ &\quad \left. + \sum_{P \in \mathcal{P}} \{ E[\theta \in T \cap P] - E[\theta \in T] \} \frac{\mu(T \cap P)}{\mu(P)} \int_P h(\theta) dF(\theta) \right\}, \end{aligned}$$

where $\mu(A) := \int_A dF(\theta)$. It turns out the every term in this sum is non-negative. To see this, fix $T \in \mathcal{T}$ and define a new function $g : T \rightarrow \mathbb{R}$ such that

$$g(\theta) = \begin{cases} h(\theta) & \text{if } \theta \in T \cap S \text{ for some } S \in \mathcal{S}, \\ \int_P h(z) dF(z) / \mu(P) & \text{if } \theta \in T \cap P \text{ for some } P \in \mathcal{P}. \end{cases}$$

Notice from the boundary conditions for the intervals in \mathcal{P} that $g(\theta)$ is non-decreasing. Each term in brackets in the above sum over the intervals $T \in \mathcal{T}$ can now be written as

$$\sum_{I \in \mathcal{I}} \{ E[\theta \in T \cap I] - E[\theta \in T] \} \int_{T \cap I} g(\theta) dF(\theta),$$

which is proportional to the covariance between $E[\theta \in T \cap I]$ and $E[g(\theta) : \theta \in T \cap I]$ across intervals $I \in \mathcal{I}$, and is therefore non-negative. ■

is (weakly) profitable because it avoids any pooling across intervals in \mathcal{I} (over which the average level of h is increasing).

REFERENCES

- [1] Bagwell, L. S. and Bernheim, B. D. [1996], "Veblen Effects in a Theory of Conspicuous Consumption," *American Economic Review*, 86, 349-373.
- [2] Banks, J. S. and Sobel, J. [1987], "Equilibrium Selection in Signaling Games," *Econometrica*, 55, 647-661.
- [3] Becker, G. S. and Murphy, K. [2000], *Social Economics*, Belknap/Harvard.
- [4] Becker, G. S., Murphy, K. and Glaeser, E. [2000], "Social Markets and the Escalation of Quality: The World of Veblen Revisited," in Becker, G. S. and Murphy, K. [2000].
- [5] Bernheim, B. D. [1994], "A Theory of Conformity," *Journal of Political Economy*, 102, 841-877.
- [6] Damiano, E. and Hao, L. [2002], "Price Discrimination in Matching Markets," mimeo, University of Toronto.
- [7] Glazer, A. and Konrad, K. [1996], "A Signaling Explanation for Charity," *American Economic Review*, 86, 1019-1028.
- [8] Harbaugh, W. T. [1998], "What do Donations Buy? A Model of Philanthropy Based on Prestige and Warm Glow," *Journal of Public Economics*, 67, 269-284.
- [9] Lizzeri, A. [1999], "Information Revelation and Certification Intermediaries," *Rand Journal of Economics*, 30, 214-231.
- [10] Maskin, E. and Riley, J. [1984], "Monopoly with Incomplete Information," *Rand Journal of Economics*, 15, 171-196.
- [11] Milgrom, P. and Segal, I. [2002], "Envelope Theorems for Arbitrary Choice Sets," *Econometrica*, 70, 583-601.
- [12] Mussa, M. and Rosen, S. [1978], "Monopoly and Product Quality," *Journal of Economic Theory*, 18, 301-317.
- [13] Myerson, R. [1981], "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58-73.
- [14] Pesendorfer, W. [1995], "Design Innovation and Fashion Cycles," *American Economic Review*, 85, 771-792.
- [15] Veblen, T. [1899], *The Theory of the Leisure Class: an Economic Study of Institutions*, London: Unwin Books.