Consumer Information and the Limits to Competition*

Mark Armstrong
Department of Economics
Oxford University

Jidong Zhou
School of Management
Yale University

October 2019

Abstract

This paper studies competition between firms when consumers observe a private signal of their preferences over products. Within the class of signal structures which allow pure-strategy Bertrand equilibria, we derive signal structures which are optimal for firms and those which are optimal for consumers. The firm-optimal signal structure amplifies the underlying product differentiation, thereby relaxing competition, while ensuring that consumers purchase their preferred product, thereby maximizing total welfare. The consumer-optimal structure dampens product differentiation, which intensifies competition, but induces some consumers with weak preferences over products to purchase their less-preferred product.

Keywords: Information design, Bertrand competition, product differentiation, online platforms.

1 Introduction

Information flows between firms and consumers affect firm competition and market performance. Information travels in both directions between the two sides of the market. Firms are able to obtain information about consumer preferences from data brokers, social media, past interaction with customers, and so on. This enables them (if permitted) to make personalized offers and price discriminate in a targeted manner. On the other side, consumers have several ways to gather information about the products

*We are grateful to Dirk Bergemann, Stephen Morris, Xiaosheng Mu, Barry Nalebuff, John Vickers, Weijie Zhong and various seminar audiences for helpful comments. Armstrong thanks the European Research Council for financial support from Advanced Grant 833849.

they might buy, and they often rely on online platforms such as search engines and product comparison websites to obtain information about their likely preferences over the various products. The precision of their information about products affects both the quality of the consumer-product match and the intensity of competition between firms. In this paper we study the second of these information flows.

In more detail, we study a symmetric duopoly market where two firms each costlessly supply a single variety of a product and compete in prices in one-shot Bertrand fashion. A consumer wishes to purchase a unit of at most one of these two product varieties. She is initially uncertain about her preferences for the varieties, but before purchase she receives a private signal of these preferences. For example, she might be informed which product she will prefer (but not the precise valuation of either product), so that the products are *ranked*. The consumer then updates her beliefs about her preferences and makes her choice given the pair of prices offered by firms. The signal structure, which governs the relationship between the consumer's true preferences and the signal she receives, is common knowledge. We wish to understand how the signal structure affects competition and welfare. In particular, we explore the limits to competition in this market: which signals induce the highest profit for firms and which generate the highest surplus for consumers?

One interpretation of the model is that consumers gather information and make purchases via a platform which provides them with product information. The platform can choose several aspects of its information disclosure to consumers, such as how detailed is the product information it displays, whether to post customer reviews or its own reviews, whether to offer personalized recommendations, and how flexibly consumers can filter and compare products. Some platforms choose to reveal little information about products, as when for instance they offer consumers a list of hotels of specified type for a given price, but the consumer only discovers which hotel she will be allocated once she has paid. Given this flexibility over the information released to consumers, we impose few restrictions on the signal structure. Whether the platform wishes to maximize firm profits, consumer surplus, or total welfare will depend in large part on which side(s) of the market it can levy fees and the form of those fees (which is something we do not consider in this paper).

Section 2 introduces the model, and in section 3 we show that the *rank* signal allows firms to obtain first-best profit whenever preferences are sufficiently dispersed *ex ante*. That is, the rank signal enables consumers to buy their preferred product, which maximizes total surplus, and with dispersed preferences it is an equilibrium for firms to charge consumers their reservation price for the preferred product. In such cases the rank signal is therefore the signal structure which maximizes industry profit. Except for the trivial case where products are perfect substitutes *ex ante*, though, it is not possible for consumers to obtain first-best surplus. Information which enables consumers to buy their preferred product also gives market power to firms and induces prices above marginal cost. Thus consumers face a trade-off between low prices and the ability to buy the better product.

Beyond situations where first-best profit is feasible, attempts to derive an optimal second-best signal structure encounter two problems. First, consumer preferences are generally two-dimensional, and current understanding of optimal information design in such cases is limited.[1] For that reason, from section 4 onwards we find ways to simplify the model so that relevant consumer heterogeneity is scalar. Second, even with scalar heterogeneity some posterior distributions are such that the only equilibria in the pricing game between firms involve mixed strategies, and often highly complex mixed strategies.[2] Conceivably, some distributions might be such that no equilibrium, even in mixed strategies, exists. For this reason, until section 6 we focus on signal structures which induce a pure strategy equilibrium in the pricing game.

In section 4 we simplify the model so that the outside option for consumers is not relevant for the analysis. (In essence this requires that the minimum utility from each product be sufficiently large.) This implies that only the *difference* in valuations for the two products—a scalar variable—matters for consumer decisions. Our approach is to find which signal structures (if any) can support given prices as equilibrium prices. It turns out that a price pair can be supported in equilibrium if the posterior preference distribution induced by the signal structure lies between two bounds. Posterior distributions which correspond to one or other of these bounds (or which comprise a

[1] It is well known that the posterior consumer valuation distribution induced by any information structure is a mean-preserving contraction of the underlying prior distribution. However, unlike the scalar case, a mean-preserving contraction has no simple characterization when consumer heterogeneity is multidimensional.

[2] For example, when the first-best profit is *not* feasible, the rank signal structure induces a binary posterior distribution for valuations with which there is no pure strategy Bertrand equilibrium.

combination of the two bounds) then play an important role in characterizing the optimal signal structures. In section 4.1 we restrict attention to symmetric signals which induce the two firms to offer the same price in equilibrium, which simplifies the analysis, and in section 4.2 we demonstrate that in regular cases neither firms nor consumers can do better if asymmetric signals and prices are implemented. (Surprisingly, asymmetric signals which favour one firm cannot improve that firm's profit relative to the optimal symmetric signal structure.)

The firm-optimal signal structure amplifies "perceived" product differentiation. If no information is disclosed the consumer regards products as perfect substitutes, leading perfect price competition and the lowest possible profit for firms. But full information disclosure is not optimal for firms, and a coarser signal structure typically does better. (The *rank* signal which sometimes yields first-best profit is an extreme instance of this.) The firm-optimal signal structure amplifies the extent of product differentiation by reducing the likelihood that consumers are near-indifferent between products. In regular cases, the firm-optimal signals allow consumers to buy their preferred product for sure, in which case total surplus is also maximized. The consumer-optimal signal structure needs to balance the trade-off between match quality and price. In the consumer-optimal signal structure, a consumer with strong preferences can buy her preferred product for sure, but a less choosy consumer will receive less precise information so that she regards the products as close substitutes and may end up with the inferior product. In other words, the consumer-optimal signal structure sacrifices the match quality of less choosy consumers to intensify price competition. In contrast to the firm-optimal signal structure, product mismatch means that the consumer-optimal signal structure does not maximize total welfare.

In section 4 the outside option was not relevant for consumers, and so the constraint on a firm raising its price was that the consumer would buy from its rival. By contrast, when the first-best profit was feasible in section 3 the constraint on raising price was that the consumer would exit the market. We bridge the gap between these extreme situations in section 5 using a framework in which both constraints play a role. (The framework is a variant of the Hotelling model, where valuations are negatively correlated.) Since the consumer-optimal policy induces low prices in the market, the presence of an outside option makes relatively little difference for the design of the consumer-optimal policy. However, the high prices typically seen with the firm-optimal policy are often constrained by the outside option, and the optimal policy then induces a posterior distribution such that no consumer regards the two products as close substitutes.

Section 6 discusses the extent to which the wider class of signals which allow the use of mixed-strategy pricing equilibria can improve outcomes. We show that allowing mixed-strategy equilibria can at best improve consumer surplus slightly. In particular, we construct an upper bound on consumer surplus and show that the upper bound is close to the maximum consumer surplus with the pure-strategy restriction.

Related literature. There is a large literature on information disclosure by firms themselves. If a firm has private information about its product quality, voluntary disclosure is often the outcome due to "unravelling" (see, e.g., Milgrom and Roberts (1986)). Another strand of the literature considers a firm's incentives to provide product information (which is not its private information) to enable consumers to discover their valuation for the product. An early paper on this topic is Lewis and Sappington (1994). They study a monopoly market and show that, within a particular class of signal structures, it is optimal for the firm either to disclose no information or all information. Johnson and Myatt (2006) derive a similar result for a more general class of information structures which induce rotations of the demand curve. Anderson and Renault (2006) argue that partial information disclosure before consumers search can be optimal for a monopolist if consumers need to pay a search cost to buy the product (in which case they learn their valuation automatically). In particular, they allow for a general signal structure as in the later Bayesian persuasion literature and show that firm-optimal information disclosure takes the coarse form whereby a consumer is informed whether or not her valuation lies above a threshold.

Roesler and Szentes (2017) study the signal structure which is optimal for consumers (rather than the firm) in a monopoly model. They show that partial rather than complete learning is optimal for consumers, that the optimal information structure induces efficient trade, and that the induced demand function takes the unit-elastic form whereby the firm is indifferent over a range of prices.[3] In their setup, where the

[3] There are a number of extensions of Roesler and Szentes (2017) within the monopoly framework, including Choi, Kim, and Pease (2019), Hinnosaar and Kawai (2018), and Terstiege and Wasser (2019). For instance, Choi, Kim, and Pease (2019) study the consumer-optimal information structure in the

product always has a positive surplus, it is clear that the firm-optimal signal structure is to disclose no information at all, so that the firm can obtain all surplus by charging a price equal to the expected valuation. With competition, however, this is no longer true since without any information consumers regard the firms' products as perfect substitutes and firms earn zero profit.[4] (Indeed in our duopoly model we will show that disclosing no information is nearly optimal for consumers, rather than firms.) Therefore, the firm-oriented problem is more interesting and challenging in our setting with competition than it is with monopoly. The consumer-optimal design with competition also exhibits some significant differences to the monopoly case. For example, it usually causes product mismatch so that the allocation is not efficient, the induced residual demand for each firm is unit-elastic for upward but not downward price deviations, and the consumer-optimal signal structure is no longer the worst signal structure for firms. Some of our results employ a similar technique used in Roesler and Szentes (2017), in that we can focus on a class of parameterized information structures. Otherwise, though, our analysis relies on the "bounds" approach we develop, which has the benefit of showing the connection between the firm-optimal and the consumer-optimal problems.

Previous research on product information disclosure and competition includes Ivanov (2013), who studies disclosure in a random-utility model where each firm individually decides how much information on its own product to release and what price to charge. He focuses on the same class of information structures as in Johnson and Myatt (2006), and shows that full disclosure is the only symmetric equilibrium when there is a large enough number of firms.[5] More recently, Boleslavsky, Hwang, and Kim (2019) show that the same result holds if general signal structures are allowed.[6] Intuitively, with a large number of firms, a consumer's valuation for the best rival product (if other firms fully disclose their information) is almost the highest possible valuation. To compete

setup of Anderson and Renault (2006) with search goods.

[4] Lewis and Sappington (1994) mention this point in their concluding discussion.

[5] Bar-Isaac, Caruana, and Cuñat (2012) study competitive product design which can be interpreted as competitive information disclosure. They consider a sequential search framework and assume information disclosure rotates the demand curve as in Johnson and Myatt (2006). They show that a reduction in the search cost induces more firms to choose niche product design/full information disclosure.

[6] Using techniques developed in Dworczak and Martini (2019), Boleslavsky, Hwang, and Kim (2019) are also able to characterize the disclosure equilibrium with any number of firms.

for the consumer, a firm provides full information in the hope that the consumer will discover a high value for its product as well.[7]

Instead of studying equilibrium disclosure by individual firms, we focus on a centralized design problem (e.g., by an information platform). This allows us to discuss signals which reflect *relative* valuations across products (e.g., which rank the products), while in this previous research each firm only controls its own product information.[8] In our framework full information disclosure is typically not the firm-optimal design even with many firms. Also in the context of centralized design, Dogan and Hu (2019) study the consumer-optimal problem in a sequential search framework with many firms. Consumers receive a signal of their valuation for a particular product only when they visit its seller. In particular, there is no information disclosure about the relative valuations across (unsampled) products. Because the reservation value in this search framework is static, this problem is related to the monopoly problem with a deterministic outside option as studied by Roesler and Szentes (2017). Another related paper is Moscarini and Ottaviani (2001). They study a duopoly model of price competition similar to ours, where a buyer receives a signal of her relative valuation for the two competing products. (They allow the two firms to be asymmetric.) The major difference, however, is that they assume both the relative valuation and the signal are binary variables. In such a setting the pricing equilibrium often involves quite complicated mixed strategies.

More broadly, our paper belongs to the recent literature on Bayesian persuasion and information design. See Kamenica and Gentzkow (2011) for a seminal paper in this literature, and Bergemann and Morris (2019) and Kamenica (2019) for recent surveys. Among its many applications to various topics, its method and insights have been used to revisit classic problems within Industrial Organization. For instance, Bergemann, Brooks, and Morris (2015) study third-degree price discrimination by a monopolist. (In contrast to our model, this paper considers signals sent to the firm about consumer preferences, and consumers accurately know their valuations from the start.) If all

[7] A similar result appears in other recent works which study competitive disclosure but without price competition, such as Au and Kawai (2018), and Board and Lu (2018).

[8] Anderson and Renault (2009) study comparative advertising in a duopoly model where each firm can choose to fully disclose its own product information, both products' information, or nothing. Among other results, they make the point that diclosing more information improves match quality but also softens price competition. (Jullien and Pavan (2019) make a similar point in a duopoly model of two-sided markets.)

ways to partition consumers are possible, the paper shows that any combination of profit (above the no-discrimination benchmark) and consumer surplus in the welfare triangle can be implemented by some signal structure. They study the set of demand functions which induce a given profit-maximizing price, which is a convex set, and the "extremal" members of this set are unit-elastic demand functions. As with our bounds, these extremal distributions play an important role in their analysis.[9]

2 The model

A risk-neutral consumer wishes to buy a single unit of a differentiated product costlessly supplied by two risk-neutral firms, 1 and 2. The consumer's valuation for the unit from firm $i = 1, 2$ is denoted $v_i \geq 0$, and her outside option is sure to have payoff zero. If p_i is firm i's price, the consumer wishes to buy from firm i if $v_i - p_i \geq v_j - p_j$ and $v_i - p_i \geq 0$, so that she prefers the net surplus from firm i to that available from firm j or from the outside option. (She wishes to consume the outside option if $v_1 < p_1$ and $v_2 < p_2$.)

Throughout the paper we assume that firms are symmetric *ex ante*, in the sense that the underlying distribution for consumer preferences $v = (v_1, v_2)$ is symmetric between v_1 and v_2. We assume that the support of v lies inside the square $[V, V + \Delta]^2$. Here, $V \geq 0$ represents the "basic utility" from any product, while $\Delta \geq 0$ captures the extra utility a consumer might obtain from the ideal product. (If Δ is small then the products are nearly perfect substitutes.) Let $\mu = \mathbb{E}[v_i]$ denote the expected valuation of either product, and in addition write

$$\mu_H = \mathbb{E}[\max\{v_1, v_2\}] \ ; \ \mu_L = \mathbb{E}[\min\{v_1, v_2\}] \ . \tag{1}$$

Thus μ_H denotes the expected valuation of the preferred product, while μ_L is the expected valuation of the less preferred product. (They are related by $\mu_L + \mu_H = 2\mu$.)

[9] Elliot and Galeotti (2019) consider a duopoly version of Bergemann *et al.* and assume each firm has some "captive" consumers who buy only from the firm they know. They show that if each firm has enough captive consumers, information design can earn firms the first-best profit. Ali, Lewis, and Vasserman (2019) depart from the information design approach in Bergemann *et al.* by considering a disclosure game with verifiable information where consumers choose how much information about their preferences they disclose to firms. They show that consumer control of their information tends to improve their welfare relative to full or no information disclosure.

Finally, write

$$\delta = \mu_H - \mu = \mathbb{E}[\max\{v_1 - v_2, 0\}] \tag{2}$$

for the incremental expected surplus from choosing the consumer's preferred product rather than a random product.

We study situations where before purchase the consumer observes a *private signal* of her preferences, v, rather than the preferences themselves. The signal is generated according to a signal structure $\{\sigma(s|v), S\}$, where S is a (sufficiently rich) signal space and $\sigma(s|v)$ specifies the distribution of signal s when the true preference parameter is v. We assume the signal structure is common knowledge to the consumer and to both firms, and determined before firms choose prices. After observing a signal s, the consumer updates her beliefs about her preferences v. Risk neutrality implies that only the expected v given s matters for the consumer's choice. The prior distribution for v and the signal structure jointly determine a new posterior distribution for (expected) v for the consumer. Since firms do not observe the consumer's private signal, they each choose a single price regardless of the signal received, and only the posterior distribution for v matters for their pricing decisions. Firms set prices simultaneously, and we use Bertrand-Nash equilibrium as the solution concept of the pricing game. Note that prices are accurately observed by the consumer in all cases, so that uncertainty concerns only the consumer's preferences.

To illustrate, consider these simple signal structures:

- Full information disclosure: here the signal perfectly reveals the true preferences, e.g., where $s \equiv v$, and so the posterior and prior distributions for v coincide.

- No information disclosure: here the signal is completely uninformative (i.e., the distribution of s does not depend on v) and the posterior distribution is a single point, $v = (\mu, \mu)$. In particular, the consumer views the two products as perfect substitutes and will choose to buy from the firm with the lower price (if the price is no higher than μ).

- Rank signal structure: here the signal informs the consumer which product she prefers but nothing else, so that $s \in \{s_1, s_2\}$ and she observes $s = s_1$ if $v_1 > v_2$ and $s = s_2$ if $v_2 > v_1$. (She sees either signal with equal probability in the knife-edge case $v_1 = v_2$.) In this case, the posterior distribution divides consumers into two groups: all consumers who see s_1 have expected valuation μ_H for product 1 and μ_L for product 2, while consumers who see s_2 have the reverse valuations.

We aim to investigate how the signal structure affects competition and the ability of consumers to buy their preferred product. In particular, we search for those signal structures which maximize industry profit and those which maximize consumer surplus. At a general level this appears to be an intractable problem, and in the following analysis we study various special cases of this framework. In the next section we discuss the easiest case to analyze, which is when the first-best profit can be achieved.

3 First-best outcomes

Suppose we find a signal structure which (i) maximizes total surplus (profit plus consumer surplus) and (ii) allocates all of that surplus to the firms in equilibrium. Then clearly no other signal structure can do better for firms (or do worse for consumers). If such a signal structure exists, its form is straightforward to derive. Since total surplus is maximized the consumer must always buy her preferred product, and since her surplus is fully extracted, she must only learn her expected valuation of the preferred product and pay a price equal to that valuation, i.e., $p_1 = p_2 = \mu_H$. For this to constitute an equilibrium, however, a firm cannot obtain more profit by deviating to a low enough price to attract those consumers who prefer the rival product, which requires deviating to price $p = \mu_L$, and thereby serving all consumers. Thus, if $\frac{1}{2}\mu_H \geq \mu_L$ the rank signal structure has equilibrium prices $p_1 = p_2 = \mu_H$ and fully extracts maximum surplus for firms. This discussion is formally stated in the following result.

Proposition 1 *If*

$$\mu_H \geq 2\mu_L \tag{3}$$

then the rank signal structure leads to an equilibrium which fully extracts maximum surplus for firms, and is therefore the signal structure which maximizes industry profit.

For instance, if v is uniformly distributed on the square $[0, \Delta]^2$ then $\mu_H = \frac{2}{3}\Delta$ and $\mu_L = \frac{1}{3}\Delta$ and so (3) is (just) satisfied. Since $\mu_L \geq V$ and $\mu_H \leq V + \Delta$, condition (3) requires $\Delta \geq V$, so that the range of valuations is large relative to the basic utility. Using (2), condition (3) can be written equivalently as

$$3\delta \geq \mu \ , \tag{4}$$

so that for a given mean μ first-best profit is more likely to be feasible when the expected absolute difference in valuations $|v_1 - v_2|$ is larger, i.e., when the difference $v_1 - v_2$ is more dispersed.

There are at least two ways in which this difference might be more dispersed. First, for a given marginal distribution for v_i if the joint distribution becomes less positively correlated then δ will rise. More precisely, let $F(v_1, v_2)$ and $\hat{F}(v_1, v_2)$ denote two joint CDFs for valuations with the same marginal distribution for v_i (and hence with the same mean μ), such that F is more correlated than $\hat{F}$ in the sense of Epstein and Tanny (1980), i.e., that $F \geq \hat{F}$. Since the function $\max\{v_1 - v_2, 0\}$ is "correlation averse" as defined in Epstein and Tanny, its expectation δ is higher with $\hat{F}$ than with F, and so if first-best profit is feasible with F it is also feasible with the less correlated distribution $\hat{F}$.

Second, if the two valuations are independently distributed, then a mean-preserving spread (MPS) in this distribution implies that δ rises.[10] To see this, note that if F is the CDF for each valuation v_i then an explicit formula for $\delta = \mu_H - \mu$ is

$$\delta = \int_0^\infty F(v)[1 - F(v)]dv \ .$$

Therefore, if $\hat{F}$ is another CDF which is a MPS of F, with corresponding $\hat{\delta}$, then we have

$$\begin{aligned} \hat{\delta} - \delta &= \int_0^\infty \left\{ \hat{F}(v)[1 - \hat{F}(v)] - F(v)[1 - F(v)] \right\} dv \\ &= \int_0^\infty (1 - F(v) - \hat{F}(v))(\hat{F}(v) - F(v))dv \\ &= \int_0^\infty (F'(v) + \hat{F}'(v)) \left(\int_0^v (\hat{F}(\tilde{v}) - F(\tilde{v}))d\tilde{v} \right) dv \geq 0 \ . \end{aligned}$$

Here, the final equality follows after integration by parts and uses the fact that F and $\hat{F}$ have the same mean, while the inequality follows from F and $\hat{F}$ being weakly increasing and the integral in $(\cdot)$ being non-negative due to the MPS assumption.

[10] Recall that $\hat{F}$ is an MPS of F if they have the same means and $\int_{-\infty}^{v} (\hat{F}(\tilde{v}) - F(\tilde{v}))d\tilde{v} \geq 0$ for all v.

We can use this second observation to obtain the following result:

Corollary 1 *Suppose that v_1 and v_2 are independently and identically distributed with density which weakly decreases over its support $[0, \Delta]$. Then condition (3) is satisfied and the rank signal structure generates first-best profits for the firms.*

Proof. Since the density is weakly decreasing, the associated CDF F is concave on $[0, \Delta]$ and let $\mu \leq \frac{1}{2}\Delta$ be its mean. Then the uniform distribution with support $[0, 2\mu]$ has the same mean and (just) satisfies condition (3). However, F is a MPS of this uniform distribution since it crosses the CDF for the uniform distribution once and from above (as it is concave and the uniform CDF is linear). Therefore (3) holds for CDF F. ■

If instead of duopoly the two products were jointly supplied by a multiproduct monopolist, the rank signal structure with associated prices $p_1 = p_2 = \mu_H$ always allows the firm to fully extract surplus, and so is the most profitable signal structure for the firm. Thus when condition (3) holds, the rank policy enables competing firms to "collude" and to achieve the monopoly profit.

One can also consider whether there exists a signal structure which maximizes total surplus and allocates it all to consumers in equilibrium. This would require that $p_1 = p_2 = 0$ are equilibrium prices, which in turn requires that consumers regard the two products as perfect substitutes. Except in the trivial case where the underlying products *are* perfect substitutes, i.e., when the prior distribution has $v_1 \equiv v_2$, if a signal structure induces consumers to view the products as identical, they are unable to choose the preferred product more than half the time. Therefore, there is a trade off for consumers between paying low prices and being able to buy their preferred product, and there is no signal structure which implements the first-best outcome for consumers.[11]

In this section we have derived the firm-optimal signal structure—which is the rank signal structure—when the valuation difference is dispersed relative to the mean valuation, in the sense that (3) holds. In this case, firms fully extract consumer surplus and the constraint that consumers not wish to consume their outside option always binds.

[11] This trade-off typically vanishes in the limit when there are *many* symmetric firms. With full information disclosure, consumers in a such a market can choose their preferred product *and* pay a price close to marginal cost, and so this policy achieves the first best for consumers.

In the remainder of the paper we discuss optimal signal structures when the first-best is not feasible, which is a considerably harder problem. In the next section we examine next simplest case, which is when the valuation difference is *concentrated* relative to the mean, in which case we can ignore consumer participation constraints altogether.

4 The market without an outside option

As discussed in the introduction, when first-best profit is not achievable attempts to derive an optimal second-best signal structure face two problems. First, preferences $v = (v_1, v_2)$ are generally two-dimensional, and current understanding of optimal information design in such cases is limited. For that reason, in the remainder of the paper we simplify the model so that relevant consumer information is only scalar. Second, even with scalar heterogeneity some posterior distributions are such that the only equilibria in the pricing game between firms involve mixed strategies, or conceivably no equilibrium (even in mixed strategies) exists at all.[12] For this reason, until section 6 we focus on signal structures which induce a pure strategy equilibrium in the pricing game.

In this section we suppose that consumer preferences are such that the outside option is never relevant for consumers. The advantage of this assumption is that only the difference in valuations, $x \equiv v_1 - v_2$, matters for consumer decisions and welfare, and so the relevant consumer heterogeneity is captured by this scalar variable. The following result provides a sufficient condition for the outside option to be irrelevant. (Note that we allow for signal structures which make the posterior market asymmetric.)

Lemma 1 *If $V > 3\Delta$, then under any signal structure which induces a pure strategy equilibrium, equilibrium prices are below V and all consumers obtain positive surplus from both firms.*

Proof. With any signal structure the maximum posterior valuation for a product does not exceed $V + \Delta$ and the minimum posterior valuation is at least V. Suppose a signal structure induces firms $i = 1, 2$ to offer respective prices p_1 and p_2 and to obtain

[12] By contrast, in the (single-product) monopoly problem with unit demand it is well understood that the firm optimally posts a single price.

profits π_1 and π_2. It is clear that neither p_1 nor p_2 can exceed $V + \Delta$.[13] Firm j will serve all consumers if it deviates to a low price p such that $V - p \geq V + \Delta - p_i$, i.e., if $p \leq p_i - \Delta$. (Since $p_i \leq V + \Delta$ the inequality $V - p \geq V + \Delta - p_i$ ensures that $V - p \geq 0$ so that all consumers prefer to buy from firm j than to buy nothing.) Therefore, we must have

$$p_i - \Delta \leq \pi_j \ . \tag{5}$$

If firms are labelled so $p_1 \geq p_2$, then the above inequality implies $p_1 - \Delta \leq \pi_2 \leq p_2$ so that the price difference $p_1 - p_2$ cannot exceed Δ. Adding the pair of inequalities (5) implies that

$$p_1 + p_2 - 2\Delta \leq \pi_1 + \pi_2 \leq p_1 \ ,$$

where the second inequality follows since industry profit cannot exceed the maximum price p_1, and so $p_2 \leq 2\Delta$. Since $p_1 \leq p_2 + \Delta$ it follows that $p_1 \leq 3\Delta$, and so when $3\Delta < V$ we must have $\max\{p_1, p_2\} < V$. ∎

Note that in the case of a symmetric equilibrium, the proof shows that $p_1 = p_2 \leq 2\Delta$, so that feasible symmetric prices lie in the interval $[0, 2\Delta]$. Note also that in the consumer-optimal policy we derive below the induced prices are low, which means that the above condition $V > 3\Delta$ can be considerably weakened.

For the remainder of section 4 suppose that both prices in any pure strategy equilibrium are less than V, so that the outside option is irrelevant in the sense that if at most one firm deviates from equilibrium all consumers continue to participate. The consumer prefers to buy from firm 1 if $x \equiv v_1 - v_2 > p_1 - p_2$, prefers to buy from firm 2 if $x < p_1 - p_2$ (and is indifferent when $x = p_1 - p_2$). Since in the underlying market v_1 and v_2 are symmetrically distributed, the scalar variable x is symmetrically distributed within the line segment $[-\Delta, \Delta]$, with CDF $F(x)$ say. Then δ in (2) takes the form

$$\delta = \mathbb{E}[\max\{v_1 - v_2, 0\}] = \int_0^{\Delta} x dF(x) = \int_{-\Delta}^{0} F(x) dx \tag{6}$$

[13]If firm j did choose price $p_j > V + \Delta$ in equilibrium, then no consumer will buy from it, and firm i acts as a monopolist and its optimal price must be $p_i \geq V > 0$. Then firm j can earn a strictly positive profit by deviating to a price slightly below p_i, as under any signal structure there must be a positive measure of consumers who weakly prefer product j over product i given the two products are symmetric *ex ante*.

where the final expression follows after integration by parts (which remains valid even if F is not continuous) and uses the fact x has a zero mean (which implies $\int_{-\Delta}^{\Delta} F(x)dx = \Delta$). Clearly $\delta \leq \frac{1}{2}\Delta$ for any symmetric F, while $\delta \leq \frac{1}{4}\Delta$ if F is convex in the range $[-\Delta, 0]$ and has no mass point at $x = -\Delta$.

Clearly, any signal of the two-dimensional preference parameter v in section 2 induces a signal s of the scalar preference parameter x (while any additional information in the signal plays no role for consumers or welfare). After observing signal s, the consumer updates her belief about her expected x. The prior distribution F and the signal structure jointly determine a signal distribution for the consumer, which further determines a posterior distribution for (expected) x which has CDF $G(x)$, say. For a given prior F, the only restriction on the posterior G imposed by Bayesian consistency is that it is a mean-preserving contraction (MPC) of F, i.e.,

$$\int_{-\Delta}^{u} G(x)dx \leq \int_{-\Delta}^{u} F(x)dx \text{ for } u \in [-\Delta, \Delta], \text{ with equality at } u = \Delta \ . \tag{7}$$

Moreover, any G which is an MPC of F can be generated by some signal structure (which is based on the scalar preference parameter x).[14] Therefore, instead of analyzing the signal structure directly (as we did with the first-best analysis), we work with the posterior distribution G subject only to the MPC constraint.

To fix ideas, Figure 1 depict various kinds of posterior distributions for x which are an MPC of the prior distribution (marked as dashed lines), here taken to be a uniform distribution on $[-1, 1]$. When G crosses F once and from below on $[-\Delta, \Delta]$, as on Figures 1a and 1c, then the necessary condition that G has the same mean as F, i.e., that there is equality at $u = \Delta$ in (7), is also sufficient for G to be an MPC of F. If the signal is symmetric and G crosses F at most once and from below in the negative range $x \in (-\Delta, 0)$, as on Figures 1a and 1b, then the necessary condition

$$\int_{-\Delta}^{0} G(x)dx \leq \int_{-\Delta}^{0} F(x)dx \equiv \delta \tag{8}$$

is sufficient for G to be an MPC of F.

[14] See, for example, Blackwell (1953), Rothschild and Stiglitz (1970), Gentzkow and Kamenica (2016), Roesler and Szentes (2017), and Kolotilin (2018).

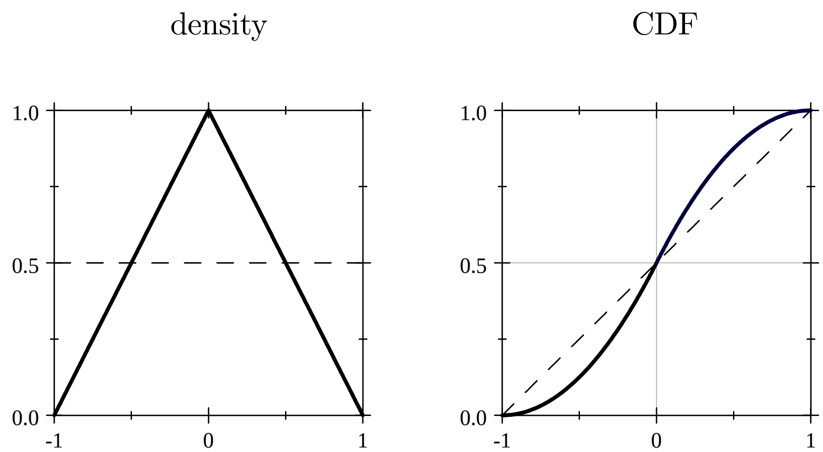

Figure 1a: MPC which increases density of consumers at $x = 0$

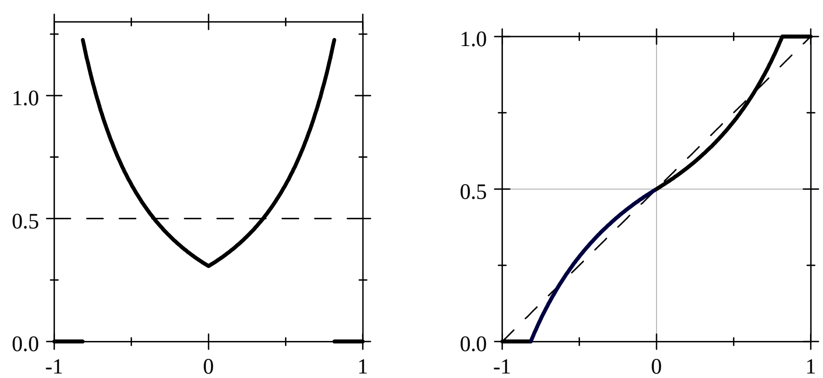

Figure 1b: MPC which reduces density of consumers at $x = 0$

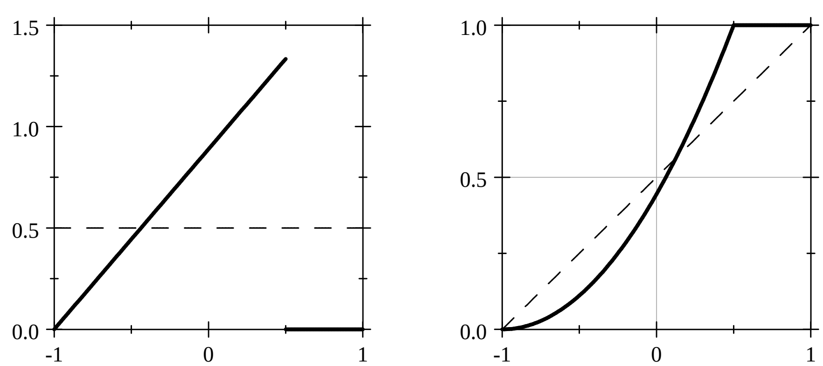

Figure 1c: Asymmetric MPC which shifts demand to one firm

It is useful to have a measure of the efficiency of product choice with a given signal structure corresponding to posterior G. With symmetric prices and full consumer participation, total surplus is the expected value of $\max\{v_1, v_2\}$, which can be written as

$$W_G = \mathbb{E}_G[\max\{v_1, v_2\}] = \mu + \mathbb{E}_G[\max\{v_1 - v_2, 0\}] = \mu + \int_{-\Delta}^{0} G(x)dx \ , \qquad (9)$$

where the final equality follows with similar logic to (6). Since G is a MPC of F, the necessary condition (8) shows that match efficiency cannot increase when the consumer observes a noisy signal of her preferences rather than her actual preferences, as is intuitive. When the equality (8) is strict—as in Figures 1a and 1c but not Figure 1b—then there is mismatch with the posterior G, due to the consumer sometimes buying the wrong product. In the extreme case where the signal is completely uninformative, we have $G = 0$ for $x < 0$ and $W_G = \mu$, the expected surplus from consuming a random product. There is no product mismatch when there is equality in (8). In terms of the signal structure this is the case when the range of signals seen when $x > 0$ does not overlap with the range of signals seen when $x < 0$, so that the consumer is fully informed about whether $x > 0$ or $x < 0$, even though she may not be fully informed about the magnitude of x.

Some of the most frequently-used signal structures induce consumers to concentrate *ex post* around $x = 0$, similarly to Figure 1a. This is so with a "truth-or-noise" structure, whereby the signal s is equal to the true x with some probability and otherwise the signal is a random realization of x, or when the distribution for x is "rotated" about $x = 0$ as studied by Johnson and Myatt (2006). Such signals induce a degree of mismatch. By contrast, a signal which accurately reveals to a consumer which product she prefers (so that (8) holds with equality) will necessary induce weakly *fewer* consumers around $x = 0$ *ex post*, as on Figure 1b.[15]

A leading case, which simplifies the following analysis and ensures all our major results, is when the prior distribution for x has a (symmetric) density which is *log-concave* on $[-\Delta, \Delta]$. In this case, as shown by Bagnoli and Bergstrom (2005), both $F(\cdot)$ and $1 - F(\cdot)$ are log-concave on $[-\Delta, \Delta]$, and in addition the underlying market (where consumers are fully informed about valuations) has a symmetric pure strategy

[15]This is because when (8) binds, the MPC constraint (7) requires that G lie weakly below F for x just below zero, in which case G is weakly steeper than F at $x = 0$.

equilibrium where the equilibrium price is $p_F = 1/(2f(0))$. A useful observation for later is the following:

Lemma 2 *If the prior density $f(x)$ is log-concave on $[-\Delta, \Delta]$, then the full-information price $p_F = 1/(2f(0))$ satisfies*

$$2\delta \leq p_F \leq 4\delta \ . \tag{10}$$

Proof. To show the upper bound, note f being symmetric and log-concave on $[-\Delta, \Delta]$ implies the density is single peaked so that F is convex for $x \leq 0$. In this case

$$\delta = \int_{-\Delta}^{0} F(x)dx \geq \int_{\frac{-1}{2f(0)}}^{0} (\tfrac{1}{2} + xf(0))dx = \frac{1}{8f(0)} = \tfrac{1}{4}p_F \ ,$$

where the inequality follows since F lies above its tangent at $x = 0$. To show the lower bound, note that

$$\tfrac{1}{2} = \int_{-\Delta}^{0} f(x)dx = \int_{-\Delta}^{0} \frac{f(x)}{F(x)}F(x)dx \geq \frac{f(0)}{F(0)} \int_{-\Delta}^{0} F(x)dx = 2f(0)\delta = \frac{\delta}{p_F} \ ,$$

where the inequality follows from F being log-concave (which in turn follows from f being log-concave). ■

In the next section we derive optimal policies using symmetric signal structures. When the underlying market has a log-concave density, we go on to show in section 4.2 that no individual firm nor consumers in aggregate can do better using asymmetric signals and prices.

4.1 Optimal symmetric signal structures

In this section, we focus on the relatively simple case of symmetric signal structures, where the posterior distribution G is symmetric, and study which symmetric prices can be implemented and which signal structures are best for firms and for consumers.

Having discussed the constraints on G imposed by Bayesian consistency, we turn next to the constraints on G needed to achieve a target symmetric price in pure strategy equilibrium. For $p = 0$ to be an equilibrium price, the consumer must regard the two varieties as identical, and this can happen only if G is degenerate at $x = 0$. In the following, we focus on positive prices. First note that to have a positive symmetric equilibrium price the distribution G cannot have an atom at $x = 0$, i.e., we must have $G(0) = \frac{1}{2}$, for otherwise a firm obtains a discrete jump in demand if it slightly undercuts its rival. Recall also that any symmetric equilibrium price satisfies $p \leq 2\Delta$.

Consider a candidate symmetric equilibrium price $p > 0$. If firm 2 deviates to price $p' \neq p$ the consumer buys from firm 2 if $x \leq p - p'$. (Thus we suppose that if G has a mass point at $x = p - p'$, firm 2 serves all consumers at that mass point. This is the natural tie-breaking rule given that the firm can achieve this outcome by charging a price slightly below p'.) Therefore, firm 2 has no incentive to deviate if and only if

$$p'G(p - p') \leq \tfrac{1}{2}p$$

holds for all p'. (The inequality holds with equality at $p' = p$ since $G(0) = \frac{1}{2}$.) By changing variables from p' to $x = p - p'$, we can write this requirement as

$$G(x) \leq U_p(x) \equiv \min\left\{1, \frac{p}{2\max\{0, p - x\}}\right\} . \tag{11}$$

(It is unprofitable for firm 2 to set a negative price p', and so there are restrictions on G only in the range where $p' = p - x > 0$, which is why there is $\max\{0, \cdot\}$ in the denominator. In addition, a CDF cannot exceed 1 which is why there is $\min\{1, \cdot\}$ in (11).) Likewise, for firm 1 to have no incentive to deviate we require $p'(1 - G(p' - p)) \leq \frac{1}{2}p$ for all (positive) p'.[16] Following the parallel argument to that for firm 2, this constraint can be written

$$G(x) \geq L_p(x) \equiv \max\left\{0, 1 - \frac{p}{2\max\{0, p + x\}}\right\} . \tag{12}$$

Notice that the two bounds are mirror images of each other, in the sense that $L_p(x) \equiv 1 - U_p(-x)$. Therefore, if a symmetric G lies between the bounds in the negative range $x \in [-\Delta, 0]$ it will lie between the bounds over the whole range $[-\Delta, \Delta]$. In summary, to induce the symmetric price $0 < p \leq 2\Delta$ in equilibrium it is necessary and sufficient that the posterior distribution G lie between the two bounds

$$L_p(x) \leq G(x) \leq U_p(x) \tag{13}$$

[16] As with firm 2, if x has an atom at $x = p' - p$ the natural tie-breaking assumption is that firm 1 serves all customers at x. Strictly speaking, as a CDF is right-continuous, here we need to interpret $1 - G(p' - p)$ as including any atom at $p' - p$. Since the bounds we derive are continuous functions, this point makes no difference to the analysis

for all $x \in [-\Delta, 0]$.

The lower bound L_p is increasing in x and begins to be positive at $x = -\frac{1}{2}p$ which exceeds $-\Delta$ given $p \leq 2\Delta$. Moreover, L_p is concave whenever L_p is positive. The upper bound U_p is increasing in x and reaches 1 at $x = \frac{1}{2}p \leq \Delta$, and it is convex when below 1. These two bounds coincide and equal $\frac{1}{2}$, and have the same slope $1/(2p)$, when $x = 0$. It follows that the lower bound is always below the upper bound, and so the set of G's satisfying (13) is not empty for $0 < p \leq 2\Delta$. Note also that both L_p and U_p rotate clockwise about the point $(0, \frac{1}{2})$ as p increases. Figure 2 illustrates the bounds (depicted as bold curves) for the three target prices, $p = 0.5$, $p = 1$ and $p = 2$ in the range of $[-\Delta, 0]$ when $\Delta = 1$. Another observation which will be often used in the subsequent analysis is that U_p is log-convex in x when it is below 1, and $1 - L_p$ is log-convex when L_p is positive.

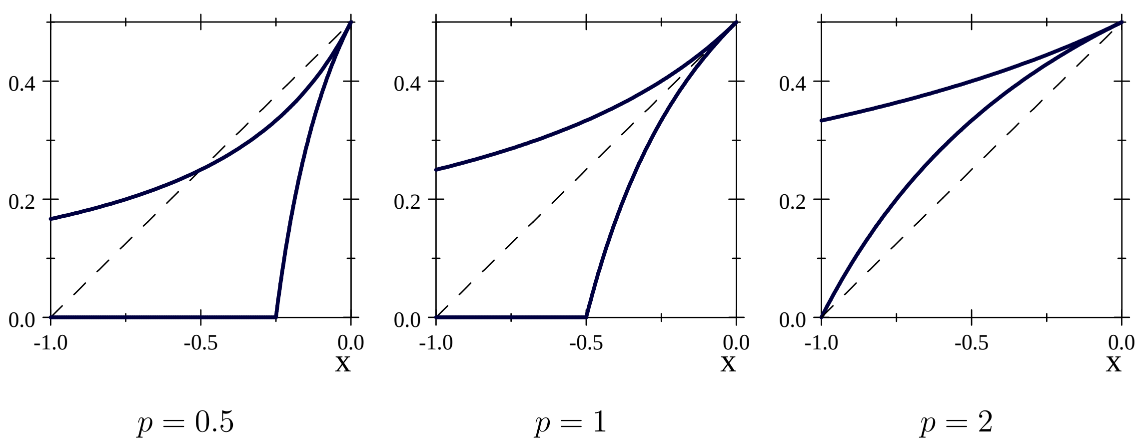

Figure 2: Bounds on G to implement prices $p = 0.5$, $p = 1$ and $p = 2$

Price p can be implemented with some signal structure provided a G can be found within the corresponding bounds and it is also an MPC of the prior. Consider a uniform prior distribution on $[-1, 1]$ as depicted by the dashed line in Figure 2. When $p = 0.5$, it can be implemented for instance by a symmetric G which equals the lower bound in the range of $[-1, 0]$. Such a G is clearly a MPC of the prior. When $p = 1$, it can be implemented for instance by a G which equals the prior (i.e. when full information is disclosed). However, it is impossible to implement price $p = 2$, as even the lower bound for $p = 2$ is above the dashed line and so cannot be an MPC of the prior.

In the following, we will largely focus on the case when the prior distribution has a density $f(x)$ which is log-concave on $[-\Delta, \Delta]$, although as we will discuss this analysis can be generalized to more general prior distributions.

Firm-optimal signal structure: Using the posterior bounds it is straightforward to derive the symmetric signal structure which maximizes profit in this market. Since there is full consumer participation, maximizing profit corresponds to finding the maximum price which can be implemented with an MPC of the prior. Notice that if a symmetric G which implements price p is a MPC of the prior, so is the symmetric posterior given by the associated lower bound L_p in the range of negative x. Therefore, to find the firm-optimal price we can restrict attention to the parameterized family of posteriors taking the form L_p for $x \in [-\Delta, 0]$.

Continue first the example with a uniform prior, as on Figure 2. Given that L_p is concave whenever it is positive, it crosses the (linear) prior CDF at most once and from below in the range of $[-1, 0]$. Therefore, from (8) a symmetric G which is equal to L_p for $x \leq 0$ is an MPC of the prior if and only if L_p has integral on $[-1, 0]$ no greater than δ, where $\delta = 1/4$ in this uniform example. Since L_p increases with p for $x \leq 0$, it is optimal to make this integral constraint bind, so that the optimal price p^* solves

$$\tfrac{1}{4} = \int_{-1}^{0} L_p(x)dx = \int_{-\frac{1}{2}p}^{0} \left(1 - \frac{p}{2(p+x)}\right) dx = \tfrac{1}{2}(1 - \log 2)p$$

and so

$$p^* = \frac{1}{2(1 - \log 2)} \approx 1.63 \; .$$

This optimal price is about 63 percent higher than the full-information outcome (where the equilibrium price was $p_F = 1$). The posterior distribution which implements this optimal price is what we depicted on Figure 1b above where the density is U-shaped. Intuitively, to soften price competition we reduce the number of marginal consumers around $x = 0$, as these consumers are the most price sensitive, and push consumers towards the two extremes insofar as this is feasible given the pure-strategy and the MPC constraints.

The same argument applies more generally whenever the lower bound L_p for each price crosses the prior CDF F at most once and from below in the range of $[-\Delta, 0]$. This is true, for example, if F is convex in the range of $[-\Delta, 0]$. More generally, since

$1 - L_p$ is log-convex when L_p is positive, this is the case if $1 - F$ is log-concave in the range of $[-\Delta, 0]$ (which includes the case where F is convex). A sufficient condition for $1 - F$ to be log-concave on $[-\Delta, 0]$ is that the density f be log-concave on $[-\Delta, \Delta]$. Then the lower bound is an MPC of the prior if and only if its integral over $[-\Delta, 0]$ is no greater than that of the prior (which is δ). Since the former integral increases in p, at the optimum there is equality in the two integrals (unless p reaches 2Δ first), so that the profit maximizing price p^* is

$$p^* = \frac{2\delta}{1 - \log 2} \ . \tag{14}$$

(This does not exceed 2Δ when the prior density is log-concave in which case $\delta \leq \frac{1}{4}\Delta$.) Clearly, the firm-optimal symmetric G, equal to the lower bound L_{p^*} for $x < 0$, is unique in this case. Also, given such a posterior distribution, when a firm unilaterally decreases its price to steal business from its rival, its demand is unit-elastic.

There are also two other useful observations. First, since by construction (8) holds with equality, the firm-optimal signal structure induces no mismatch, and total surplus is maximized as well as profit. Second, the firm-optimal price p^* is considerably higher than the price p_F under full information disclosure, so that full information disclosure is not optimal for firms. Indeed, with a log-concave density (10) and (14) together imply that $p^* \geq \frac{1}{2(1-\log 2)} p_F \approx 1.63 \times p_F$, and so relative to full disclosure profits rise by at least 63% using the optimal signal structure.

We summarize this discussion in the following result:

Proposition 2 *Suppose the outside option is not relevant and the prior distribution has a log-concave density. Then:*
(i) the firm-optimal symmetric price p^* *is (14), which is at least 63% higher than the full-information price* p_F*, and it is uniquely implemented by the posterior* $G = L_{p^*}$*;*
(ii) with the firm-optimal symmetric signal structure there is no mismatch and total surplus is also maximized.

Note that familiar signal structures such as "truth-or-noise" or rotations in the distribution around $x = 0$ induce consumers to be *more* concentrated around $x = 0$, and so cannot be used to enhance profit relative to the full-information policy. Therefore, the use of unrestricted signal structures, which allow consumers to buy their preferred product, enables firms to do at least 63% better than they could do with these more *ad hoc* signal structures.

Beyond the simple case with a log-concave density, the firm-optimal price which can be implemented with some signal structure is the highest $p \leq 2\Delta$ such that L_p in (12) satisfies

$$\int_{-\Delta}^{u} L_p(x)dx \leq \int_{-\Delta}^{u} F(x)dx \text{ for } u \in [-\Delta, 0] \ .$$

In general L_p and F can cross multiple times in the range $[-\Delta, 0]$, in which case solving the optimal price can be less straightforward. Moreover, (8) might hold strictly and so there could be welfare loss associated with the firm-optimal signal structure. Figure 3 illustrates both points, where the prior shown as the dashed curve is initially convex and then concave. The highest price such that L_p is an MPC of the prior is shown as the solid curve, where the integrals of the two curves up to the crossing point a are equal (so with any higher price the MPC constraint would be violated). Here, since L_{p^*} lies below the prior for x above a, (8) holds strictly and there is some mismatch at the optimum.

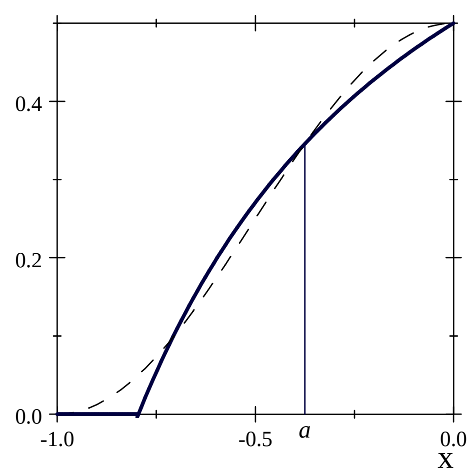

Figure 3: Firm-optimal posterior with a less regular prior

Consumer-optimal signal structure. We turn next to the problem of finding the best symmetric signal structure for consumers. Unlike firms, consumers do not care solely about the induced price but also about the reliability of the product match, and consumer surplus with posterior G and price p is $W_G - p$, where total surplus W_G is as given in (9).

Note first that δ is the incremental consumer benefit from buying the preferred product rather than a random product, and that in the information structure where consumers receive no product information they buy a random product at price zero. Therefore, for consumers to do better than the no-information policy, the price they pay cannot exceed δ. With a log-concave density, however, (10) shows that the full-information price satisfies $p_F \geq 2\delta$, and so the consumer-optimal price must be less than half of the full-information price.

To solve the consumer-optimal problem in more detail, we first find the highest possible G to maximize match efficiency for a given price p, subject to the bounds condition (13) and the MPC constraint, and we then identify the optimal price.

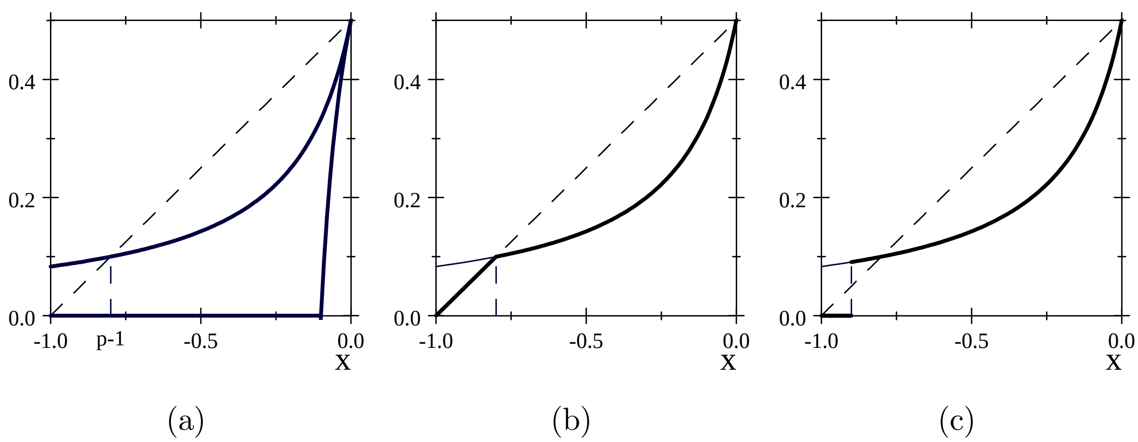

Figure 4: Consumer-optimal G for a given price p

Consider again the example where the prior distribution is uniform on $[-1, 1]$, where the equilibrium price with full information disclosure is $p_F = 1$. For consumers to do better than this policy, the induced price must be below 1 to counter-act any potential product mismatch, in which case the bounds look similar to the left-hand picture in Figure 2 and the upper bound is below the prior for x close to zero. (In fact, since this density is log-concave, we know the optimal price is below $\frac{1}{2}$.) Therefore, it is the *upper* bound (11) which will constrain G, rather than the lower bound which was relevant for the firm-optimal policy. Since the upper bound is convex, for any price $p < 1$ the upper bound cuts the prior CDF once and from above. Figure 4 illustrates how to maximize consumer surplus for a given price $p < 1$. Figure 4a shows the two bounds in (13) as bold curves, where the upper bound cuts the prior CDF at $x_p \equiv p - 1$ in this example. Given p, we wish to maximize the integral of G over $[-\Delta, 0]$, subject to lying between these bounds and the MPC constraint.

Two necessary conditions for G are that it satisfy the MPC constraint (7) at the intercept point x_p, i.e.,

$$\int_{-\Delta}^{x_p} G(x)dx \leq \int_{-\Delta}^{x_p} F(x)dx \ , \tag{15}$$

and that G lies below the upper bound for x above the intercept point. Clearly, the solution involves setting $G(x) = U_p(x)$ for x above the intercept point, since the MPC constraint (7) is surely satisfied for $x \geq x_p$ if (15) holds for any $G \leq U_p$. In addition, it is clear that (15) must bind. However, there are many ways to choose G such that this constraint binds, all of which yield the same consumer surplus. Figure 4b shows a convenient way to do this, which is to set G equal to the prior CDF, so that

$$G(x) = \min\{F(x), U_p(x)\} \ , \tag{16}$$

while Figure 4c depicts an alternative way to satisfy the constraint. Since there is a strict inequality in (8) for any optimal G, there is welfare loss at the consumer optimum and some consumers buy their less preferred product. However, those consumers with strong preferences, i.e., those with $x \leq x_p$, receive their preferred product for sure.[17]

Expression (16) implies that consumer surplus given $p \leq 1$ is

$$W_G - p = \mu + \int_{-\Delta}^{0} \min\{F(x), U_p(x)\}dx - p \ . \tag{17}$$

Its derivative with respect to p is

$$\int_{x_p}^{0} \frac{\partial U_p(x)}{\partial p} dx - 1 = \frac{1}{2}\left(\frac{p}{p - x_p} - \log \frac{p}{p - x_p} - 3\right) \ , \tag{18}$$

where x_p is the intercept point of $F(x)$ and $U_p(x)$. In the uniform example where $x_p = p - 1$, (18) becomes $\frac{1}{2}(p - \log p - 3)$ which decreases with p in the range $[0, 1]$. Then the optimal price is

$$p^* = \gamma \approx 0.05 \ ,$$

[17] Note that for price $p \leq 1$, the posterior (16) is above L_p. This is because at the full information price $p_F = 1$ we have $L_{p_F} \leq F$, and so the same is true for any lower price.

where γ is the root of $\gamma - \log \gamma = 3$ in the range $[0, 1]$.

Figure 5 depicts the consumer-optimal posterior distribution (when it takes the particular form in Figure 4b). Loosely speaking, the distribution for x is rotated around $x = 0$ so that number of price-sensitive consumers near $x = 0$ is amplified compared with prior distribution, and this forces firms to reduce their price in equilibrium.[18] Those consumers near $x = 0$ do not have strong preferences about which product they buy, and so there is only limit welfare loss due to product mismatch. Those consumers with very strong preferences, however, are sure to buy their preferred product and at a low price. Such a posterior distribution also implies that when a firm unilaterally increases its price its residual demand is unit-elastic. (This is opposite to the firm-optimal solution.)

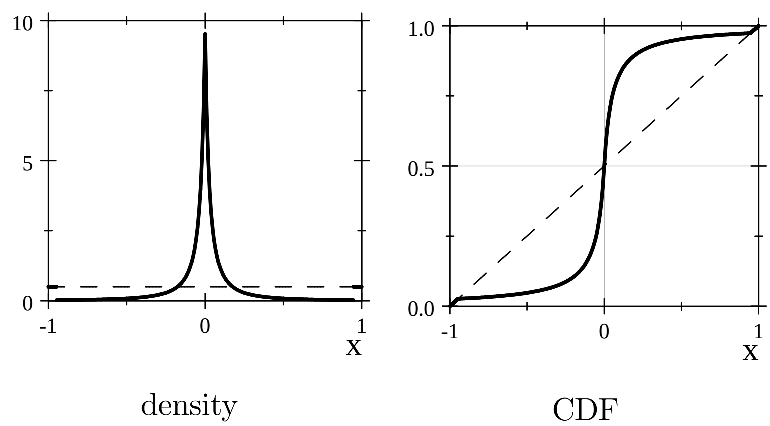

Figure 5: Consumer-optimal information structure

The same argument applies if the upper bound for each price below the full information price p_F crosses the prior CDF once and from above in the range $[-\Delta, 0]$. This is true, for example, when F is concave on $[-\Delta, 0]$. More generally, since U_p is log-convex, this is the case if F is log-concave on $[-\Delta, 0]$. A sufficient condition for F to be log-concave on $[-\Delta, 0]$ is that the density f be log-concave.[19]

[18] It is not a strict rotation in the sense of Johnson and Myatt (2006), since the posterior coincides with the prior at the extremes of the distribution.

[19] We also need to check that G in (16) is above the lower bound L_p for prices below p_F, which is ensured if F is above L_p or $1 - F$ is below $1 - L_p$. However, since $1 - L_p$ is log-convex, this is the case for all $p \leq p_F$ when $1 - F$ is log-concave.

With the optimal G in (16), the derivative of consumer surplus with respect to price is (18). Since $F(x_p) \equiv U_p(x_p)$ it follows that $\frac{p}{p-x_p} = 2F(x_p)$, and substituting this into (18) shows the derivative of consumer surplus with respect to price to be $\frac{1}{2}(2F(x_p) - \log(2F(x_p)) - 3)$. Note that the intercept point x_p increases with p given that the upper bound crosses F from above. The above derivative therefore decreases with p, and so the optimal intercept point x^* satisfies $2F(x^*) = \gamma \approx 0.05$, or $x^* = F^{-1}(\frac{1}{2}\gamma)$. The optimal price p^* then satisfies $\frac{p^*}{p^*-x^*} = 2F(x^*) = \gamma$, from which we obtain

$$p^* = \frac{-\gamma}{1-\gamma} x^* = \frac{-\gamma}{1-\gamma} F^{-1}(\tfrac{1}{2}\gamma) \ , \tag{19}$$

and that optimal consumer surplus is

$$\begin{aligned} \mu + \int_{-\Delta}^{x^*} F(x)dx + \int_{x^*}^{0} U_{p^*}(x)dx - p^* &= \mu + \int_{-\Delta}^{x^*} F(x)dx - p^* \left(1 + \tfrac{1}{2}\log \frac{p^*}{p^* - x^*}\right) \\ &= \mu + \int_{-\Delta}^{x^*} F(x)dx - x^* F(x^*) \\ &= \mu - \int_{-\Delta}^{x^*} x dF(x) \ . \end{aligned} \tag{20}$$

(Here, the second equality used $\frac{p^*}{p^*-x^*} = 2F(x^*) = \gamma$, the definition of γ, and (19), while the last equality follows by integration by parts.) For example, when the prior is uniform on $[-\Delta, \Delta]$, the optimal price in (19) is $p^* = \gamma\Delta$ and the optimal consumer surplus in (20) is $\mu + \frac{1}{4}(2\gamma - \gamma^2)\Delta$.

Two useful observations follow: First, the fraction of consumers sure to choose their preferred product, which is $2F(x^*)$, is equal to γ regardless of the prior (provided it has log-concave density). Second, with a log-concave density, F is convex in the range $[-\Delta, 0]$, and so $\frac{\gamma}{2} = F(x^*) \geq \frac{1}{2} + x^* f(0) = \frac{1}{2}(1 + \frac{x^*}{p_F})$, or $\frac{-x^*}{p_F} \geq 1 - \gamma$. Using (19), we see that $p^* > \gamma p_F$.

We summarise this discussion in the following result:

Proposition 3 *Suppose the outside option is not relevant and the prior distribution has a log-concave density. Then:*
(i) the consumer-optimal symmetric price p^* *is (19), which satisfies* $\gamma p_F < p^* < \frac{1}{2}p_F$*, and it is implemented by the posterior (16) and yields consumer surplus given in (20);*
(ii) with the consumer-optimal symmetric signal structure, only a fraction γ *of consumers are sure to buy their preferred product, so there is mismatch and total surplus is not maximized.*

Because the optimal price is low, the constraint that G should lie below the upper bound U_p is more important than the constraint that G is an MPC of F. (More precisely, in (16) $G = F$ for only about 5% of consumers, and otherwise it is equal to the upper bound U_p.) It follows that this consumer-optimal policy is closely approximated by the solution to a simpler optimization problem, which is to choose a symmetric distribution G in order to maximize equilibrium consumer surplus. In this alternative scenario, there is no "prior" and no MPC constraint, and we wish to choose the distribution for x, say within the support $[-\Delta, \Delta]$, which trades off the benefits of a low equilibrium price (which is implemented by a distribution concentrated around $x = 0$) and the benefits of being able to the choose the better of two products (which is greater when the distribution for x is more dispersed). The above discussion shows that the solution to this problem is to choose the price p to maximize

$$W_G - p = \mu + \int_{-\Delta}^{0} U_p(x)dx - p \tag{21}$$

instead of (17). Since the difference between (21) and (17) increases with p, the solution to this second problem involves a higher price than in Proposition 3. However, in most cases the difference is tiny.[20] For the same reason, the optimal policy depends only on the extreme tails of the underlying prior, as seen in expression (19), rather than its shape elsewhere.

Another implication of the low price in the consumer-optimal policy is that the conditions required for the validity of ignoring the consumer participation constraint are much less stringent than required for Lemma 1 (which was $V > 3\Delta$). More precisely, we claim that if $V > \delta$ then the consumer-optimal policy is as described in Proposition 3.[21] To see this, note that as we have pointed out before, for consumers to do better than the no-information policy, the industry profit π they generate cannot exceed δ. If some consumers do not participate under the optimal policy, the price p must exceed the minimum valuation V. If one firm deviates to a lower price V, it then sells to at least half the consumers and so it obtains deviation profit at least $\frac{1}{2}V$. Therefore,

[20] With support $[-1, 1]$, the price which maximizes (21) is $p \approx 0.055$, compared to $p^* = \gamma \approx 0.052$ for the example with a uniform prior.

[21] Note that $\delta \leq \frac{1}{2}\Delta$, and so a sufficient condition to ignore the participation constraint is $V > \frac{1}{2}\Delta$. If the prior density for $x = v_1 - v_2$ is further log-concave, then F is convex on $[-\Delta, 0]$ and so $\delta \leq \frac{1}{4}\Delta$, in which a weaker sufficient condition is $V > \frac{1}{4}\Delta$.

equilibrium industry profit π exceeds V, which is not possible when $\pi \leq \delta$ and $\delta < V$. Therefore, there must be full participation at the consumer-optimal policy when $V > \delta$. It follows that the price itself must be below V, in which case the outside option does not then apply (even when one firm deviates to a higher price).

For more general prior distributions, the following method (in the spirit of Roesler and Szentes (2017)) can be used to simplify the original functional optimization problem into a parametrized optimization problem. Consider a symmetric posterior G described as the thin curve between the two bounds in Figure 6, which induces equilibrium price p and is an MPC of F. Then for each $m \in [-\Delta, 0]$ we can construct a new posterior

$$G_p^m(x) = \begin{cases} L_p(x) & \text{if } x < m \\ U_p(x) & \text{if } m \leq x \leq 0 \end{cases} \tag{22}$$

illustrated as the bold curves, which induces the same price p.

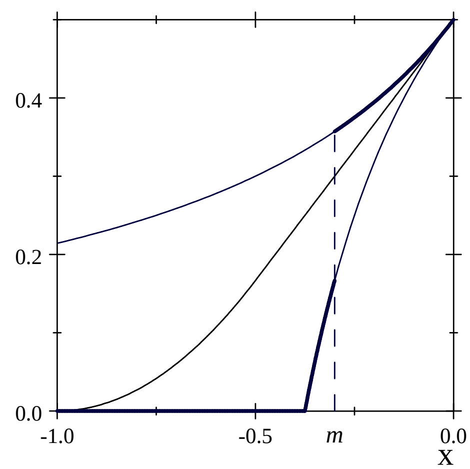

Figure 6: A class of candidate consumer-optimal G's

It is clear that there exists an m such that $\int_{-\Delta}^{0} G_p^m(x)dx = \int_{-\Delta}^{0} G(x)dx$, in which case G_p^m is an MPC of G (and hence an MPC of F) and yields the same consumer surplus as G. Therefore, we can look for the consumer-optimal posterior within the (two-dimensional) family of parameterized distributions G_p^m by solving the problem $\max_{p,m} \mu + \int_{-\Delta}^{0} G_p^m(x)dx - p$ subject to G_p^m being an MPC of F. For instance, using this method one can show that it is always optimal to disclose some information to consumers.

Corollary 2 *Except in the degenerate case where products are perfect substitutes (i.e., $x \equiv 0$), it is sub-optimal to disclose no information to consumers.*

Proof. It suffices to find a signal structure which is better for consumers than no information disclosure. Consider G_p^m defined in (22). Let $m = -\kappa p$, where $\kappa > e^2 - 1 \approx 6.4$ is a constant as p varies. Since F is not degenerate at $x = 0$, G_p^m is an MPC of F when $p > 0$ is sufficiently small. In addition, $L_p(x) = 0$ for $x \le m$. Therefore, consumer surplus with this policy is

$$\mu + \int_{-\Delta}^{0} G_p^m(x)dx - p = \mu + \int_{-\kappa p}^{0} U_p(x)dx - p = \mu + \tfrac{1}{2}p[\log(1+\kappa) - 2] \ .$$

The term $[\cdot]$ is positive by assumption, and so this policy is better for consumers than no information disclosure (since the latter corresponds to $p = 0$). ■

The welfare limits. Having discussed the signal structures which maximize profit and which maximize consumer surplus, we are in a position to describe the set of combinations of profit and consumer surplus which are feasible with some choice of symmetric signal structure. Suppose that the prior distribution has log-concave density. Any price (or profit) from zero to the firm-optimal price p^* in (14) can be acheived, and given such a price the range of consumer surplus which is possible is determined by the possible posteriors G which lie within the bounds (13) and which are an MPC of the prior. Since the set of such posterior distributions is convex, we merely need to determine the worst and the best consumer surplus for a given price, as any intermediate surplus can be acheived with a convex combination of the two extreme posteriors.

At the maximum price $p = p^*$ there is only one posterior which satisfies all constraints, which is the lower bound L_{p^*}. Therefore, there is a single feasible consumer surplus at this price, which is $\mu_H - p^*$. (Here, $\mu_H = \mathbb{E}_F[\max\{v_1, v_2\}] = \mu + \delta$.) Likewise, at the minimum price $p = 0$ the only posterior which satisfies all constraints corresponds to no information disclosure, and the only feasible consumer surplus is μ.

For a price $p \in (0, p^*)$ the lowest value of consumer surplus is generated by the lower bound L_p. (This is because consumer surplus is lower with a smaller G, and since L_{p^*} is a MPC of F, so is L_p for any lower price.) As above, the integral of $L_p(x)$ over $[-\Delta, 0]$ is $\frac{1}{2}(1 - \log 2)p$ and so the minimum consumer surplus with price p is the linear function $\mu - \frac{1}{2}(1 + \log 2)p$.

To derive the highest value of consumer surplus for a price $p \in (0, p^*)$, we need to deal with two cases. When $p \geq p_F$ (the full information price), it is possible to find a supporting posterior G between the bounds which is an MPC of F such that (8) holds with equality.[22] Since there is no mismatch with such a posterior, consumer surplus is $\mu+\delta-p$. For these prices, profit and maximum consumer surplus sum to maximum total surplus, $\mu + \delta$. When $p < p_F$, however, we have shown that the maximum consumer surplus corresponds to the posterior (16), and the resulting consumer surplus depends on the functional form of the prior F. For these prices, profit and maximum consumer surplus sum to less than maximum total surplus $\mu + \delta$ due to mismatch.

Figure 7 shows the feasible combinations of profit and consumer surplus lying between the two bold curves, in the case where the prior is uniform on $[-1, 1]$, when $\delta = \frac{1}{4}$, $p_F = 1$ and $p^* \approx 1.63$. (Note that we have normalized consumer surplus so that its minimum feasible value is zero.) The dashed line denotes the combinations of profit and consumer surplus which sum to maximum total surplus.

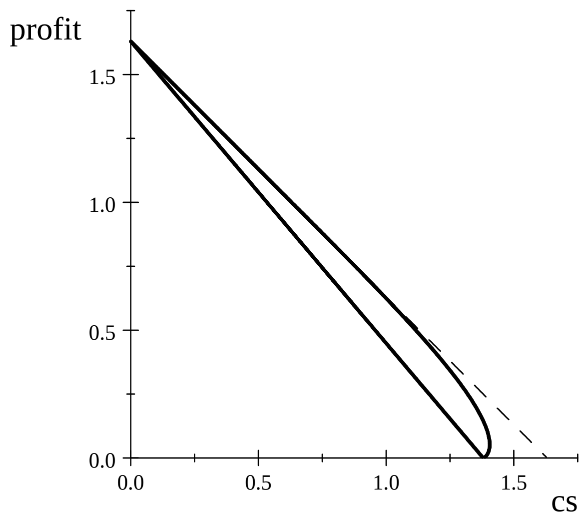

Figure 7: Feasible combinations of profit and consumer surplus

This contrasts with the corresponding figure for monopoly in Roesler and Szentes (2017, Figure 1) in a number of ways. First, minimum profit is zero here due to competition, whereas in the monopoly situation the smallest price achievable in equilibrium

[22]This is because given $p < p^*$ the lower bound L_p must have (8) hold with strict inequality, while given $p \geq p_F$ the log-convex upper bound U_p must be above the log-concave F in the range of $x < 0$. Then one way to construct such a G is similar to (22): $G(x) = L_p(x)$ for $x < m$ and $G(x) = U_p(x)$ for $x \geq m$, where m is chosen to make (8) bind.

is positive. Second, here the feasible combinations are close to being "zero-sum", in the sense that it is not feasible that consumer surplus is low while profit is also low, but it is not precisely zero-sum since consumer surplus is not maximized at a point where profit is minimized. Finally, the outer frontier of the feasible set is not always efficient, since to implement a price below the full-information price ($p_F = 1$ on the figure) we must introduce a degree of product mismatch.

4.2 Asymmetric signal structures

We now extend the analysis to allow for asymmetric signal structures, such as illustrated on Figure 1c above, which induce firms to choose distinct prices in equilibrium. This extension is important as, for instance, it could be possible to design consumer information in such a way that at least one firm can obtain higher profit than it did with the firm-optimal symmetric signals presented in Proposition 2. This higher profit might come either at the expense of its rival or as part of a reduction in competitive intensity due to asymmetric rivalry between firms. It might also be possible for consumers to prefer asymmetric prices: for a *fixed* distribution over x consumer surplus is convex and decreasing in the two firms' prices, so that consumers prefer distinct prices to a uniform price equal to the average of the two prices.

As with the symmetric analysis, our approach is to provide bounds on the posterior distributions $G(\cdot)$ which induce a given pair of positive prices (p_1, p_2) in equilibrium. As before, G can have no atom at $x = p_1 - p_2$, and firm 2's equilibrium market share is $G(p_1 - p_2)$. Since the posterior support of x must lie within $[-\Delta, \Delta]$, if π_i denotes firm i's equilibrium profit then (5) must hold. More generally, firm 2 can make no greater profit with another price p_2', so that

$$p_2' G(p_1 - p_2') \le \pi_2 \ .$$

As in section 4.1, changing variable to $x = p_1 - p_2'$ implies

$$G(x) \le \min\left\{1, \frac{\pi_2}{\max\{0, p_1 - x\}}\right\} \equiv U_{p_1,p_2}(x) \ . \tag{23}$$

(Here only the range $x \le p_1$ is relevant, for otherwise firm 2 offers a negative price.) The parallel argument for firm 1 yields

$$G(x) \ge \max\left\{0, 1 - \frac{\pi_1}{\max\{0, p_2 + x\}}\right\} \equiv L_{p_1,p_2}(x) \ . \tag{24}$$

Note that the lower bound L_{p_1,p_2} is increasing in x and begins to be positive at $x = \pi_1 - p_2$ which exceeds $-\Delta$ from (5). Moreover, similar as in the symmetric case, L_{p_1,p_2} is concave and $1 - L_{p_1,p_2}$ is log-convex in x whenever L_{p_1,p_2} is greater than zero. The upper bound U_{p_1,p_2} is increasing in x and reaches 1 at $x = p_1 - \pi_2$ which is below Δ from (5). Moreover, it is log-convex (and hence convex) when it is less than 1.

Since we must have $U_{p_1,p_2} \geq L_{p_1,p_2}$ to have a chance to implement these prices, and since the bounds coincide and equal firm 2's market share when $x = p_1 - p_2$, the two functions should have the same slope at $x = p_1 - p_2$, i.e., $\pi_2/p_2^2 = \pi_1/p_1^2$. If we write $s = 1 - G(p_1 - p_2)$ for firm 1's market share, so that $\pi_1/p_1 = s$ and $\pi_2/p_2 = 1 - s$, this then implies

$$s = \frac{p_1}{p_1 + p_2} \ , \tag{25}$$

and

$$\pi_i = \frac{p_i^2}{p_1 + p_2} \ . \tag{26}$$

In particular, equilibrium profits and market shares are determined entirely by equilibrium prices and do not depend separately on $G(x)$, and the firm with the higher equilibrium price necessarily has the higher market share (and hence the higher profit).

In sum, given a prior distribution F, a price pair (p_1, p_2) can be implemented with some signal structure if and only if a posterior G exists which is both (i) an MPC of F and (ii) lies between the bounds (23)–(24), where profits are (26).

To illustrate this discussion, consider the target prices are $p_1 = \frac{5}{4}$ and $p_2 = 1$, in which case profits in (26) are $\pi_1 = \frac{25}{36}$ and $\pi_2 = \frac{16}{36}$, and firm 1 has market share $\frac{5}{9}$. Then any posterior G which lies between the two bold curves on Figure 8 will implement this pair of prices. When the prior distribution is uniform on $[-1, 1]$, F is shown as the dashed line on Figure 8, and no MPC of F can lie within the bounds on the figure.[23] Therefore, with this prior the prices $p_1 = \frac{5}{4}$ and $p_2 = 1$ cannot be implemented in equilibrium with any signal structure.

[23] The easiest way to see this is to note that the area under the upper bound between -1 and 1 is lower than that of the prior CDF, which implies that the mean of a distribution with CDF below this upper bound is greater than zero, and so cannot be an MPC of F.

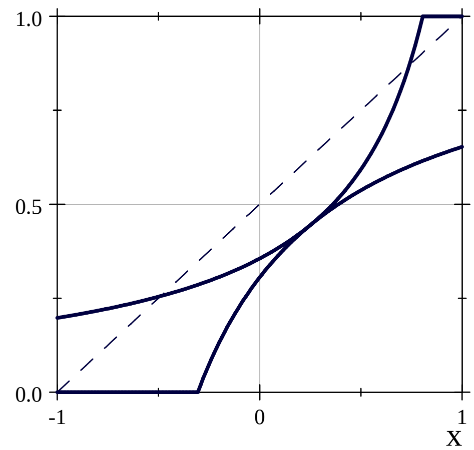

Figure 8: Bounds on G to implement prices $p_1 = \frac{5}{4}$ and $p_2 = 1$

Firm-optimal signal structure. A necessary condition for a potentially asymmetric posterior G to be an MPC of the prior F is (8), and in particular (8) needs to hold for the lower bound (24), so that when $\pi_1 \leq p_2$ we require

$$\delta \geq \int_{\pi_1 - p_2}^{0} \left(1 - \frac{\pi_1}{p_2 + x}\right) dx = p_2 - \pi_1 - \pi_1 \log \frac{p_2}{\pi_1} \ . \tag{27}$$

(If $\pi_1 \geq p_2$, so that the point where the lower bound reaches zero is positive, constraint (8) is sure to be satisfied.)

Consider the signal structure which maximizes firm 2's profit, π_2, say. Such a signal structure must induce profits such that $\pi_2 \geq \pi_1$ (and hence $p_2 \geq p_1$), for otherwise firm 2 would do better if the signal structure was reversed and it obtained firm 1's profit. Since we then have $\pi_1 \leq p_2$, the point where the lower bound reaches zero is negative and it is necessary that (27) be satisfied. If we write $r = \pi_2/\pi_1 \geq 1$ for the profit ratio, then (27) can be written as

$$\delta \geq \pi_2 \left(\frac{p_2}{\pi_2} - \frac{1}{r} - \frac{1}{r} \log \frac{p_2}{\pi_1}\right) = \pi_2 \left[1 + \frac{1}{\sqrt{r}} - \frac{1}{r} - \frac{1}{r} \log(r + \sqrt{r})\right]$$

where the equality follows after inverting the pair of equations (26) to obtain $p_2 = \pi_2 + \sqrt{\pi_1 \pi_2}$. The term $[\cdot]$ is equal to $1 - \log 2$ in the symmetric case $r = 1$ and is strictly greater than $1 - \log 2$ for $r > 1$.[24] It follows that any feasible profit π_2 cannot exceed

[24] This can be verified by using the concavity of $\log(\cdot)$ to show that $\log(r + \sqrt{r}) \leq \log 2 + \frac{1}{2}(r + \sqrt{r} - 2)$. The term $[\cdot]$ is initially increasing for $r \geq 1$ and then decreases asymptotically to 1.

$\delta/(1-\log 2)$. Since the two firms are symmetric *ex ante*, the same is true for firm 1's profit.

Proposition 2 showed that this upper bound on firm 2's profit could be achieved whenever the prior density was log-concave, and it was acheived simultaneously for both firms under a symmetric signal structure. Therefore, at least with a log-concave density, there is no way to design consumer information so that one firm (let alone both firms) is better off than with the optimal symmetric signal structure. Perhaps surprisingly, then, firms have perfectly congruent interests when it comes to the design of consumer information.

Given the firm which charges a higher price has a larger market share, this profit result implies also that any price pair implemented with a feasible signal structure has both prices below the optimal symmetric price pair. Intuitively, the firm which is treated unfavorably under an asymmetric signal structure has an incentive to set a low price, and this force turns out to be sufficiently powerful so that the firm which is treated favorably will also reduce its price.

Consumer-optimal signal structure. Let p_1 and p_2 be the equilibrium prices induced with some signal structure. Similarly to (9), consumer surplus with these prices is

$$\mathbb{E}[\max\{v_1 - p_1, v_2 - p_2\}] = \mu - p_1 + \int_{-\Delta}^{p_1 - p_2} G(x)dx \ . \tag{28}$$

Suppose asymmetric prices are implemented, where without loss of generality we have $p_1 > p_2$. (Then firm 1 has a larger market share and so $\pi_2 < \frac{1}{2}p_2 < \frac{1}{2}p_1$.) As on Figure 8, the upper bound U_{p_1,p_2} satisfies $U_{p_1,p_2}(x) < \frac{1}{2}$ for $x \le p_1 - p_2$. Since the upper bound is log-convex and the prior F is log-concave given its density is log-concave, the upper bound must cross F once and from above in the range of $[-\Delta, p_1 - p_2]$. Let $\hat{x} < 0$ denote this intercept point. Then a similar argument to that which led to (17) shows that consumer surplus in (28) can be no greater than

$$\mu - p_1 + \int_{-\Delta}^{p_1 - p_2} \min\{F(x), U_{p_1,p_2}(x)\}dx \ . \tag{29}$$

Consider changing prices to a symmetric price pair $p_1 = p_2 = p$ such that the new upper bound crosses F at the *same* point $\hat{x}$. This implies that p satisfies

$$\frac{\pi_2}{p_1 - \hat{x}} = \frac{\frac{1}{2}p}{p - \hat{x}} \ ,$$

or

$$p = \frac{-\hat{x}\pi_2}{\frac{1}{2}(p_1 - \hat{x}) - \pi_2} < 2\pi_2 < p_2 \ .$$

Here, the first inequality follows since $\pi_2 < \frac{1}{2}p_1$ and the second inequality follows since $\pi_2 < \frac{1}{2}p_2$. Therefore, this uniform price is lower than p_2 (and hence lower than p_1).

The difference between expression (29) with the uniform price p and with original prices (p_1, p_2) is

$$p_1 - p + \int_{\hat{x}}^{0} \frac{\frac{1}{2}p}{p - x} dx - \int_{\hat{x}}^{p_1 - p_2} \frac{\pi_2}{p_1 - x} dx \ . \tag{30}$$

Note that the first integrand, $\frac{1}{2}p/(p-x)$, is greater than the second, $\pi_2/(p_1 - x)$, in the range $\hat{x} \leq x \leq 0$. (This is because there is equality by construction in the two terms when $x = \hat{x}$, and $(p - x)/(p_1 - x)$ decreases with x given that $p < p_1$.) Since we also have $p < p_2$, it follows that (30) is greater than

$$p_1 - p_2 - \int_{0}^{p_1 - p_2} \frac{\pi_2}{p_1 - x} dx > 0$$

where the inequality holds since the integrand (which is the upper bound) is less than 1.

We deduce that starting from any distinct prices (p_1, p_2), the upper bound on consumer surplus (29) increases if we instead implement this uniform price p. Since this upper bound is acheived with symmetric prices (given the log-concavity assumption), we see that consumer surplus cannot be increased by using asymmetric signals and prices. Intuitively, this is because an asymmetric signal structure causes mismatch, and it turns out that this negative effect outweighs the benefit from lower prices in an asymmetric market.

We summarise this section with the following formal result:

Proposition 4 *Suppose the outside option is not relevant and the prior distribution has a log-concave density. Then relative to the optimal symmetric signal structures in Propositions 2 and 3 the use of asymmetric signal structures cannot improve an individual firm's profit or aggregate consumer surplus.*

The analysis in this section is also useful for studying optimal policies when the underlying market is *asymmetric*. For example, the bounds (23)–(24) continue to apply, as does the expression for equilibrium profit in (26). Figure 8 illustrates the bounds, except that the prior no longer need pass through the point $(0, \frac{1}{2})$. However, the calculation of optimal signals becomes significantly more complicated, as the symmetric benchmark—which played an important role in both the firm and consumer analysis above—is no longer relevant.

5 A market with an outside option

In section 3 we showed how firms earn the first-best profit with the rank signal structure when consumer valuations were sufficiently dispersed, in which case the participation constraint for all consumers was binding. By contrast, section 4 studied the situation where valuations were sufficiently concentrated, in which case the participation constraint was irrelevant and second-best policies could be derived since only the (scalar) valuation difference $x = v_1 - v_2$ mattered. In this section we bridge the gap between these two situations by considering a case where consumer heterogeneity is actually one-dimensional.

There are a number of demand specifications where there is scalar heterogeneity. For instance, we could suppose that the valuation for one of the products is accurately known *ex ante*, while information about a second product (perhaps a new product) might be manipulated. This is a special case of the general setup when consumers are distributed on a vertical or horizontal segment in the valuation space. In this section, however, we maintain the assumption that products are symmetric and suppose that average valuations are the same for all consumers while there is uncertainty about relative preferences, as in a Hotelling-style market. This is a special case of the general setup when consumers are distributed on a diagonal segment of the form $v_1 + v_2 =$ *constant* in valuation space. More precisely, suppose a consumer values product 1 at $v_1 = 1 + \frac{1}{2}x$ and product 2 at $v_2 = 1 - \frac{1}{2}x$, where $x = v_1 - v_2 \in [-\Delta, \Delta]$ indicates her relative preference for product 1 and each consumer's average valuation $\frac{1}{2}(v_1 + v_2) \equiv 1$ is constant. Assume $\Delta \leq 2$ so that all consumers value both products. The prior distribution for $x \in [-\Delta, \Delta]$ is symmetric about zero and has CDF $F(x)$. For simplicity, in this section we focus on symmetric signal structures which induce a pure strategy pricing equilibrium.

We first show that even if the outside option binds, all consumers purchase in the firm-optimal or consumer-optimal solution. The argument for the consumer-optimal policy is simple. A consumer-optimal signal structure must be weakly better for consumers than no information disclosure, where consumers buy a random product at price zero and so consumer surplus is 1. Since the match efficiency improvement relative to random match is at most δ (which has the same definition as in (6)), firms cannot earn more than δ in the consumer-optimal solution, where $\delta \leq \frac{\Delta}{2} \leq 1$ given $\Delta \leq 2$. Suppose in contrast the market is not fully covered. Then the price must exceed 1 so that consumers around $x = 0$ do not buy. Then a feasible unilateral deviation is to charge at 1, in which case at least half of the consumers will buy from the deviating firm. Hence, each firm's equilibrium profit must be greater than $\frac{1}{2}$, and so industry profit exceeds 1 which is a contradiction.

The argument for the firm-optimal policy is more complicated, and we provide the details in the appendix.

Lemma 3 *A firm-optimal or consumer-optimal symmetric signal structure induces an equilibrium with full market coverage.*

We next extend the previous bounds analysis to the situation where the outside option may be relevant but the market is still fully covered in equilibrium. Consider a symmetric equilibrium price p, and suppose firm 2 deviates to p'. A type-x consumer will buy from firm 2 if and only if $1 - \frac{x}{2} - p' \geq \max\{0, 1 + \frac{x}{2} - p\}$, i.e., if $x \leq \min\{2(1 - p'), p - p'\}$. Hence, p is a full-coverage equilibrium price if and only if

$$p'G(\min\{2(1-p'), p - p'\}) \leq \tfrac{1}{2}p \tag{31}$$

holds for any p' and with equality at $p' = p$. To implement a price $p \leq 1 - \frac{\Delta}{2}$ (which is the lowest valuation for a product), the extensive margin $2(1 - p')$ does not matter and the bounds are (13) as before. To implement a higher price $p > 1 - \frac{\Delta}{2}$, we need to deal with the extensive margin explicitly. Note that if $p > 1$ then any consumers with posterior $x \approx 0$ will not participate. However, as with the rank signal, the signal could induce a gap in the posterior distribution around $x = 0$, in which case it is possible to have full coverage with a price $p > 1$.

For convenience, define

$$U_p^M(x) = \min\left\{1, \frac{p}{\max\{0, 2-x\}}\right\} \; ; \; U_p(x) = \min\left\{1, \frac{p}{2\max\{0, p-x\}}\right\} \; . \tag{32}$$

Here, $U_p(x)$ is the same upper bound as before, and $U_p^M(x)$ is the upper bound when the outside option binds. Notice that U_p^M and U_p intersect only once at $\tilde{x}_p \equiv 2(p-1)$ and $U_p^M > U_p$ if and only if $x < \tilde{x}_p$. (Note that $\tilde{x}_p \leq \Delta$ given p never exceeds $1 + \frac{\Delta}{2}$, the highest valuation for a product.) Using this notation, condition (31) can be written as

$$G(x) \leq \max\{U_p^M(x), U_p(x)\}$$

and

$$G(\min\{-\tilde{x}_p, 0\}) = \tfrac{1}{2} \ . \tag{33}$$

The qualitative form of the bounds depend on the size of p as shown in Figure 9 below. (Recall that when G is symmetric, the lower bound is the mirror image of the upper bound.) For price $1 - \frac{\Delta}{2} < p \leq 1$, we have $\tilde{x}_p \in (-\Delta, 0]$ so the upper bound takes the form of U_p^M for $x < \tilde{x}_p$ as illustrated in Figure 9a. The upper bound passes through the point $(0, \frac{1}{2})$, and the bounds conditions imply (33) which is now $G(0) = \frac{1}{2}$. In particular, the lower bound in the range $x \in [-\Delta, 0]$ is unchanged from (13). For a price $1 < p < \frac{4}{3}$, the bounds are shown in Figure 9b. We have $\tilde{x}_p \in (0, \frac{p}{2})$ where $\frac{p}{2}$ is the value of x where U_p reaches 1. The crucial difference is that now (33) implies $G(-\tilde{x}_p) = \frac{1}{2}$. This requires $G(x) = \frac{1}{2}$ for $x \in [-\tilde{x}_p, \tilde{x}_p]$, and so in this middle range there are no consumers and the upper bound and the lower bound coincide. Finally, for price $p \geq \frac{4}{3}$, we have $\tilde{x}_p > \frac{p}{2}$ and the middle range is so large that the bounds are as shown in Figure 9c. In particular, the lower bound for negative x is a step function with discontinuity at $-\tilde{x}_p$.

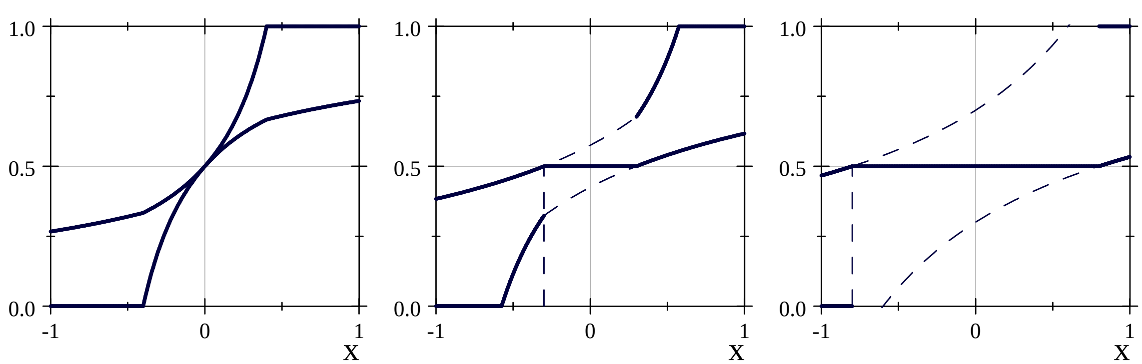

(a) $1 - \frac{\Delta}{2} < p \leq 1$ (b) $1 < p < \frac{4}{3}$ (c) $p \geq \frac{4}{3}$

Figure 9: Bounds on G to implement price $p > 1 - \frac{\Delta}{2}$

Using these bounds, the next result reports the firm-optimal solution when the prior has a log-concave density.[25]

Proposition 5 *When the prior distribution has a log-concave density, the firm-optimal solution involves no mismatch, and is as follows:*
(i) when $\delta \leq \frac{1}{2}(1 - \log 2) \approx 0.15$*, the firm-optimal price is* p^* *in (14), which satisfies* $p^* \leq 1$*, and is uniquely implemented by* L_{p^*}*;*
(ii) when $\frac{1}{2}(1 - \log 2) < \delta < \frac{1}{3}$*, the firm-optimal price* $p^* \in (1, \frac{4}{3})$ *solves*

$$\int_{-\Delta}^{0} \tilde{L}_p(x)dx = 1 - \tfrac{1}{2}p\left(1 + \log(\frac{4}{p} - 2)\right) = \delta \tag{34}$$

where

$$\tilde{L}_p(x) = \begin{cases} L_p(x) & \text{if } x < 2(1-p) \\ \frac{1}{2} & \text{if } 2(1-p) \leq x \leq 0 \end{cases}$$

is the lower bound shown in Figure 9b, and is uniquely implemented by $\tilde{L}_{p^*}$*;*
(iii) when $\delta \geq \frac{1}{3}$*, which corresponds to (4), the firm-optimal price is* $p^* = \mu_H = 1 + \delta$ *which earns firms the first-best profit and is implemented by the rank signal structure.*

Intuitively, if the prior distribution is sufficiently concentrated (in the sense that δ is small), the firm-optimal price must be low so that the outside option is irrelevant and the solution is the same as in Proposition 2. In contrast, if the prior distribution is sufficiently dispersed that (4) holds the first-best outcome is achievable. In between, the optimal solution is a mixture of these two cases. In all the three cases, there is no product mismatch and so total welfare is maximized as well.

To illustrate, consider the uniform example with support $[-\Delta, \Delta]$ and $\delta = \frac{\Delta}{4}$. When $\Delta \leq 2(1 - \log 2) \approx 0.61$, case (i) in Proposition 5 applies and the optimal G has a U-shaped density similar to Figure 1b before. When $2(1-\log 2) \leq \Delta \leq \frac{4}{3}$, case (ii) applies and the optimal G is as described in Figure 10 in the case $\Delta = 1$ (where $p^* \approx 1.235$ in (34)). The distribution has two symmetric mass points (represented as the dots on the density figure) and no consumers located between them. When Δ is larger, the optimal distribution has more weight on the two mass points, and as Δ approaches $\frac{4}{3}$ it converges to a binary distribution on $\{-\frac{\Delta}{2}, \frac{\Delta}{2}\}$ which is implemented by the rank structure and earns firms the first-best profit.

[25] As in section 4.1, the analysis can also be extended to the case with a more general prior.

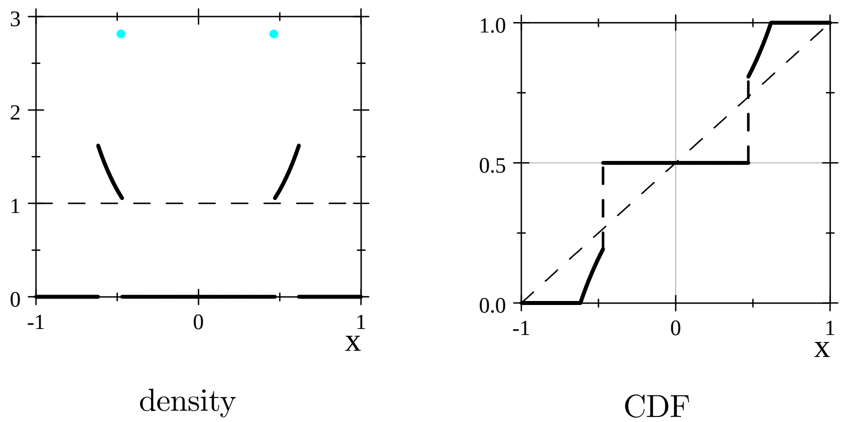

Figure 10: Firm-optimal G when $\Delta = 1$

The consumer-optimal policy is less affected by the presence of the outside option. As mentioned, the consumer-optimal price is no greater than δ. When the prior density is log-concave, we have $\delta \leq \frac{\Delta}{4}$, and so the consumer-optimal price is no greater than $1 - \frac{\Delta}{2}$ if $\Delta \leq \frac{4}{3}$, in which case Proposition 3 continues to apply. For larger Δ, we have $\delta \leq \frac{1}{2}$ given $\Delta \leq 2$, so we can focus on price $p \leq \frac{1}{2}$. Therefore, the posterior bounds take the form in Figure 9a, where the upper bound in the range $x \in [-\Delta, 0]$ is

$$\tilde{U}_p(x) = \max\{U_p^M(x), U_p(x)\} \ . \tag{35}$$

Notice that $\tilde{U}_p$ is log-convex and F is log-concave in the range of $x \leq 0$ given its density is log-concave, and since $p < \delta \leq \frac{1}{2}p_F$ the upper bound $\tilde{U}_p$ crosses F only once and from above. Therefore, the same analysis of the consumer problem in section 4.1 applies here, after replacing the upper bound U_p there by $\tilde{U}_p$.

To illustrate, consider again the uniform example with support $[-\Delta, \Delta]$. One can check that $F(\tilde{x}_p) \leq U_p(\tilde{x}_p)$ if $p \leq 2 - \Delta$, where recall $\tilde{x}_p = 2(p-1)$ is where U_p^M and U_p intersect. In this price range, U_p^M becomes irrelevant and the intercept point of F and $\tilde{U}_p$ in the range of negative x is the same as in section 4.1. Following the analysis there, the consumer-optimal price is $p^* = \gamma\Delta$, where recall $\gamma \approx 0.05$ is the solution to $\gamma - \log\gamma = 3$. This is indeed less than $2 - \Delta$ if $\Delta \leq \frac{2}{1+\gamma} \approx 1.9$. When $p > 2 - \Delta$, the intercept point of F and $\tilde{U}_p$ solves $F(x) = U_p^M(x)$, from which it follows

$x_p = \frac{1}{2}[2 - \Delta - \sqrt{(2-\Delta)^2 - 8\Delta(p-1)}]$. In this case, (18) becomes

$$\int_{x_p}^{0} \frac{\partial \tilde{U}_p(x)}{\partial p} dx - 1 = \log \frac{2 - x_p}{2 - x_p} + \frac{1}{2}\left(\frac{p}{p - x_p} - \log \frac{p}{p - x_p} - 3\right) .$$

This is positive at $p = 2 - \Delta$ if $\Delta > \frac{2}{1+\gamma}$ and must be negative at $p = 1$. For instance, when $\Delta = 2$, the consumer-optimal price is $p^* \approx 0.105$, but this is almost the same as $\gamma\Delta$. In other words, even in this case with large Δ the outside option has a negligible effect on the consumer-optimal policy.

6 Allowing mixed pricing strategies

It is hard to deal systematically with signal structures which induce mixed strategy pricing equilibrium, and the bounds approach we developed earlier in the paper does not applies with mixed strategy equilibria. However, in this section we derive an upper bound for consumer surplus across all symmetric signal structures which induce a symmetric pure or mixed strategy equilibrium. When the prior distribution is regular, we show this upper bound is close to the optimal consumer surplus under the pure strategy restriction. Intuitively, mixed strategy pricing usually does not intensify price competition and the resulting price dispersion further causes product mismatch, in which case it does not benefit consumers.

Consider the model introduced in section 2, where F denotes the prior distribution of $x = v_1 - v_2$ and G denotes its posterior distribution, where we assume $V \geq \Delta$. We first derive a lower bound for profit when the match efficiency of the signal structure is equal to $\tau \leq \delta$, so that the (symmetric) posterior G satisfies

$$\int_{-\Delta}^{0} G(x)dx = \tau \quad \text{and} \quad G \text{ is an MPC of } F \ . \tag{36}$$

When $\tau = 0$ we must have an equilibrium with zero profit, and in the following assume $\tau > 0$. We find a lower bound on a firm's profit given τ, denoted, $\underline{\pi}(\tau)$, by iterated deletion of dominated strategies.

Clearly, no firm will offer a negative price. The worst outcome for a firm (say, firm 2) is if its rival charges a zero price, so that firm 2's profit with G is at least $\max_p p \Pr[v_2 - p \geq v_1] = \max_p pG(-p) = \max_x (-x)G(x)$. (This is true regardless of whether the equilibrium involves pure or mixed strategies. The outside option plays no role here, since the rival offers positive surplus to consumers given its price is zero.) Define $\pi_0(\tau) = \inf_G \max_x (-x)G(x)$ over G which satisfy (36). Then $\pi_0(\tau) > 0$ is a lower bound on a firm's profit, and so no firm will charge a price below $\pi_0(\tau)$ in equilibrium. Note that (36) implies $\pi_0(\tau) \le \tau$, and since $\tau \le \frac{1}{2}\Delta$ it follows that $\pi_0(\tau) \le V$ given the assumption $V \ge \Delta$.

We can repeat the previous step assuming the rival firm charges price $p = \pi_0(\tau)$ (rather than $p = 0$). Since $\pi_0(\tau) \le V$, the rival offers positive surplus to consumers and we can ignore the outside option again. We obtain a new, and higher, lower bound $\pi_1(\tau) = \inf_G \max_x (\pi_0(\tau) - x)G(x)$ over G which satisfy (36). Suppose x and G solve this problem. If $x \ge 0$ then $\pi_1(\tau) \le \pi_0(\tau)G(x) \le \pi_0(\tau) \le \tau$. If $x < 0$ then $\pi_1(\tau) \le \tau + \frac{1}{2}\pi_0(\tau)$, where the inequality follows since $(-x)G(x) \le \tau$ and $\pi_0(\tau)G(x) \le \frac{1}{2}\pi_0(\tau)$. This implies $\pi_1(\tau) \le \frac{3}{2}\tau$. In either case we have $\pi_1(\tau) \le 2\tau \le \Delta \le V$.

Following this procedure recursively leads to an increasing sequence of lower bounds $\{\pi_k(\tau)\}_{k\ge 0}$. (Following the previous argument, if $\pi_{k-1}(\tau) \le 2\tau$ then also $\pi_k(\tau) \le 2\tau$, and so in each step the outside option plays no role.) Given profit is bounded this sequence must converge, and denote its limit by $\underline{\pi}(\tau)$ which solves the fixed-point equation

$$\underline{\pi}(\tau) = \inf_G \max_x (\underline{\pi}(\tau) - x)G(x) \tag{37}$$

subject to (36).

Finally, with a posterior G satisfying (36), total welfare is no greater than the full-coverage welfare in (9) which is $\mu + \tau$. Thefore, since any $\tau \in [0, \delta]$ can be chosen, consumer surplus can be no greater than $\mu + \max_{\tau\in[0,\delta]}[\tau - 2\underline{\pi}(\tau)]$. When the prior distribution of $v_1 - v_2$ is regular, it turns out that this upper bound has a simple expression.

Proposition 6 *Suppose $V \ge \Delta$ and $x = v_1 - v_2$ has log-concave density on $[-\Delta, \Delta]$. Then an upper bound for consumer surplus (across all symmetric signal structures which induce a symmetric pure or mixed strategy equilibrium) is given by expression (20) where $x^* < 0$ solves $F(x) - \log(2F(x)) = \frac{5}{2}$ so that $F(x^*) \approx 0.043$. The maximum consumer surplus which can be acheived using only pure strategy equilibria attains at least 98.4% of this upper bound.*

The consumer surplus upper bound has the same expression as the optimal consumer surplus under the pure strategy restriction derived in section 4.1, except that there x^* satisfies $F(x^*) \approx 0.026$ while here x^* satisfies $F(x^*) \approx 0.043$. Since both of these $F(x^*)$ are small, the difference between the two consumer surplus levels is also small. In other words, allowing mixed strategy pricing can increase consumer surplus only slightly. (However, we have not found an example where the use of mixed strategies improves consumer surplus at all.)

Ideally we would like to obtain a tight upper bound on profit as well, and see how close the optimal profit under the pure strategy restriction is relative to the upper bound. This appears to be a harder problem, and we leave it for future work.[26]

7 Conclusion

This paper has studied the limits to competition when product information possessed by consumers can be designed flexibly. Among signal structures which induce pure strategy pricing equilibrium, we derived the optimal policy for firms and for consumers. The firm-optimal signal structure amplifies perceived product differentiation by reducing the number of consumers who regard the products as close substitutes. The firm-optimal signal structure typically enables consumers to buy their preferred product, and so it maximizes total surplus as well. In particular, a rank information structure which only informs consumers of which product is a better match can sometimes be optimal for firms. The consumer-optimal signal structure, in contrast, dampens perceived product differentiation by increasing the number of marginal consumers and so implements a low price. This low price can only be acheived by inducing a degree of product mismatch, however, and so the policy does not maximizing total surplus.

One interesting extension to this analysis would be to consider the case where firms were asymmetric *ex ante*, and to investigate whether the optimal information policy amplifies or reduces this asymmetry. (The analysis in section 4.2 would be a useful ingredient for such an extension.)

Another, more ambitious, extension would be to investigate how the number of

[26] Under the rank signal structure, following Moscarini and Ottaviani (2001) we can characterize a symmetric mixed strategy pricing equilibrium when $\mu_H < 2\mu_L$. The profit there is less than the optimal profit derived in section 4.1 under the pure strategy restriction if the prior density is log-concave. However, the opposite can be true for an irregular prior when there is no pure strategy equilibrium with full information (e.g., when $v_1 - v_2$ is binary on $\{-\Delta, \Delta\}$).

rivals affects optimal information design. Full analysis of such an extension would likely require consideration of multi-dimensional signals, rather than the scalar analysis in this paper, even in situations where the outside option could be ignored. Nevertheless, preliminary observations about the n-firm case include: (i) the signal structure which informs consumers of their most preferred product but nothing else can sometimes enable firms to achieve the first-best outcome under a suitably modified version of (3); (ii) even if it does not achieve first-best profit, the same signal structure bounds industry profit away from zero regardless of the number of firms; (iii) a modified signal structure which informs consumers which are their best *two* products (without ranking that top two) will induce marginal cost pricing without sacrificing too much match quality, and can give nearly all the surplus to consumers when the number of firms is large.

References

Ali, N., G. Lewis, and S. Vasserman (2019): "Voluntary Disclosure and Personalized Pricing," mimeo.

Anderson, S., and R. Renault (2006): "Advertising Content," *American Economic Review*, 96(1), 93–113.

——— (2009): "Comparative Advertising: Disclosing Horizontal Match Information," *RAND Journal of Economics*, 40(3), 558–581.

Au, P., and K. Kawai (2018): "Competitive Information Disclosure by Multiple Senders," mimeo.

Bagnoli, M., and T. Bergstrom (2005): "Log-Concave Probability and its Applications," *Economic Theory*, 26(2), 445–469.

Bar-Isaac, H., G. Caruana, and V. Cuñat (2012): "Search, Design, and Market Structure," *American Economic Review*, 102(2), 1140–1160.

Bergemann, D., B. Brooks, and S. Morris (2015): "The Limits of Price Discrimination," *American Economic Review*, 105(3), 921–957.

Bergemann, D., and S. Morris (2019): "Information Design: A Unified Perspective," *Journal of Economic Literature*, 57(1), 44–95.

Blackwell, D. (1953): "Equivalent Comparisons of Experiments," *The Annals of Mathematical Statistics*, 24(2), 265–272.

Board, S., and J. Lu (2018): "Competitive Information Disclosure in Search Markets," *Journal of Political Economy*, 126(5), 1965–2010.

Boleslavsky, R., I. Hwang, and K. Kim (2019): "Competitive Advertising and Pricing," mimeo.

Choi, M., K. Kim, and M. Pease (2019): "Optimal Information Design for Search Goods," *American Economic Review P&P*, 109(5), 550–56.

Dogan, M., and J. Hu (2019): "Consumer Search and Optimal Information," mimeo.

Dworczak, P., and G. Martini (2019): "The Simple Economics of Optimal Persuasion," *Journal of Political Economy*, forthcoming.

Elliot, M., and A. Galeotti (2019): "Market Segmentation Through Information," mimeo.

Epstein, L., and S. Tanny (1980): "Increasing Generalized Correlation: A Definition and Some Economic Consequences," *Canadian Journal of Economics*, 13(1), 16–34.

Gentzkow, M., and E. Kamenica (2016): "A Rothschild-Stiglitz Approach to Bayesian Persuasion," *American Economic Review P&P*, 106(5), 597–601.

Hinnosaar, T., and K. Kawai (2018): "Robust Pricing with Refunds," mimeo.

Ivanov, M. (2013): "Information Revelation in Competitive Markets," *Economic Theory*, 52(1), 337–365.

Johnson, J., and D. Myatt (2006): "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review*, 96(3), 756–784.

Jullien, B., and A. Pavan (2019): "Information Management and Pricing in Platform Markets," *Review of Economic Studies*, 86(4), 1666–1703.

Kamenica, E. (2019): "Bayesian Persuasion and Information Design," *Annual Review of Economics*, 11, 249–272.

Kamenica, E., and M. Gentzkow (2011): "Bayesian Persuasion," *American Economic Review*, 101(6), 2590–2615.

Kolotilin, A. (2018): "Optimal Information Disclosure: A Linear Programming Approach," *Theoretical Economics*, 13(2), 607–635.

Lewis, T., and D. Sappington (1994): "Supplying Information to Facilitate Price Discrimination," *International Economic Review*, 35(2), 309–327.

Milgrom, P., and J. Roberts (1986): "Relying on the Information of Interested Parties," *RAND Journal of Economics*, 17(1), 18–32.

Moscarini, G., and M. Ottaviani (2001): "Price Competition for an Informed Buyer," *Journal of Economic Theory*, 101(2), 457–493.

Roesler, A., and B. Szentes (2017): "Buyer-Optimal Learning and Monopoly Pricing," *American Economic Review*, 107(7), 2072–2080.

Rothschild, M., and J. Stiglitz (1970): "Increasing Risk: I. A Definition," *Journal of Economic Theory*, 2(3), 225–243.

Terstiege, S., and C. Wasser (2019): "Buyer-Optimal Robust Information Structures," mimeo.

Technical Appendix

Proof of Lemma 3. Here we prove that the market is fully covered in the firm-optimal solution. It suffices to show that for any signal structure which induces a partial-coverage equilibrium, there exists another signal structure which induces a full-coverage equilibrium with a strictly higher industry profit.

Consider a symmetric posterior distribution G which is an MPC of F and induces an equilibrium where each firm charges $p > 1$ and only a fraction $\alpha < 1$ of consumers buy. (If $p \leq 1$ all consumers would buy in equilibrium.) Notice that $\tilde{x}_p \equiv 2(p-1) > 0$

solves $1+\frac{x}{2}=p$, so consumers with $x \geq \tilde{x}_p$ buy from firm 1 and those with $x \leq -\tilde{x}_p$ buy from firm 2. Other consumers in the range of $(-\tilde{x}_p, \tilde{x}_p)$ are excluded from the market. Industry profit in this equilibrium must be no less than one, i.e., $\alpha p \geq 1$, since each firm could attract half the consumers by charging price 1.

Suppose firm 1 charges the equilibrium price p but firm 2 deviates to p'. A consumer of type x will buy from firm 2 if and only if $1-\frac{x}{2}-p' \geq \max\{0, 1+\frac{x}{2}-p\}$. This requires $x \leq \min\{2(1-p'), p-p'\}$. The no-deviation condition for the partial-coverage equilibrium is then $p'G(\min\{2(1-p'), p-p'\}) \leq \frac{1}{2}\alpha p$, with equality at $p'=p$. Changing variables yields

$$G(x) \leq \alpha \max\{U_p^M(x), U_p(x)\} \;\; \text{and} \;\; G(-\tilde{x}_p) = \frac{\alpha}{2} \; ,$$

where U_p^M and U_p are given in (32). Here, $U_p^M > U_p$ if and only if $x < \tilde{x}_p$. The upper bound passes through the point $(-\tilde{x}_p, \frac{\alpha}{2})$. For our purpose, we only need the lower bound which is the mirror image of the upper bound:

$$L_{\alpha,p}(x) = \begin{cases} 1-\alpha U_p(-x) & \text{if } x < -\tilde{x}_p \\ \max\{\frac{\alpha}{2}, 1-\alpha U_p^M(-x)\} & \text{if } -\tilde{x}_p \leq x < 0 \end{cases} .$$

In the following, we will use

$$L^-_{\alpha,p}(x) = \begin{cases} 1-\alpha U_p(-x) & \text{if } x < -\tilde{x}_p \\ \frac{\alpha}{2} & \text{if } -\tilde{x}_p \leq x < 0 \end{cases} \tag{39}$$

which is weakly lower than $L_{\alpha,p}(x)$.

Let $\hat{p} = \alpha p \geq 1$ and construct a new symmetric posterior which is equal to

$$L_{1,\hat{p}}(x) = \begin{cases} 1-U_{\hat{p}}(-x) & \text{if } x < -\tilde{x}_{\hat{p}} \\ \frac{1}{2} & \text{if } -\tilde{x}_{\hat{p}} \leq x < 0 \end{cases} \tag{40}$$

in the range of negative x. Note that this is the lower bound of posteriors which support a full-coverage equilibrium with price $\hat{p} \geq 1$. In the following, we show that $L_{1,\hat{p}}$ is a 'strict' MPC of G in the sense of $\int_{-\Delta}^{u} L_{1,\hat{p}}(x)dx < \int_{-\Delta}^{u} G(x)dx$ for any $u \in (-\Delta, 0]$. (Then a similar posterior associated with a price slightly above $\hat{p}$ must be an MPC of G.) Since $L^-_{\alpha,p} \leq G$, it suffices to show $L_{1,\hat{p}}$ is a 'strict' MPC of $L^-_{\alpha,p}$. One can check that $L_{1,\hat{p}}$ crosses $L^-_{\alpha,p}$ only once and from below in the range of negative x. Therefore, it suffices to show

$$\int_{-\Delta}^{0} L_{1,\hat{p}}(x)dx < \int_{-\Delta}^{0} L^-_{\alpha,p}(x)dx \; .$$

Using (39) and (40), one can rewrite this condition as

$$1-\tfrac{1}{2}\alpha p\times\left(1+\log\left(\frac{4}{\alpha p}-2\right)\right)<(2-\alpha)(1-\tfrac{1}{2}p)-\tfrac{1}{2}\alpha p\times\log\left(\frac{4}{\alpha p}-\frac{2}{\alpha}\right)\ ,$$

which further simplifies to

$$\alpha p\times\log\frac{2-p}{2-\alpha p}<2(\alpha-1)(p-1)\ .$$

Given $\log x\leq x-1$, a sufficient condition for the above inequality is

$$\frac{\alpha p}{2-\alpha p}>\frac{2(p-1)}{p}\ .$$

Since $\alpha p\geq 1$, we have $\frac{\alpha p}{2-\alpha p}\geq 1$. Therefore, the above condition holds if $1>\frac{2(p-1)}{p}$ or $p<2$. This must be true given $p<1+\frac{\Delta}{2}$ and $\Delta\leq 2$. This completes the proof. ■

Proof of Proposition 5. The lower bound for $x\in[-\Delta,0]$ across the three cases depicted in Figure 9 can be succinctly defined as

$$\tilde{L}_p(x)=\begin{cases}L_p(x) & \text{if } x<\min\{0,-\tilde{x}_p\}\\ \frac{1}{2} & \text{if } \min\{0,-\tilde{x}_p\}\leq x<0\end{cases}\ ,$$

and it increases with p. When the prior density is log-concave the prior CDF F is convex with $F(-\Delta)=0$. Therefore, $\tilde{L}_p$ crosses F at most once and from below in the range of negative x. This implies that the optimal posterior must take the form of the lower bound, and the optimal price p^* solves

$$\int_{-\Delta}^{0}\tilde{L}_p(x)dx=\delta\ . \tag{41}$$

(This implies there is no mismatch with the firm-optimal signal structure.)

We then have: (i) if $\delta\leq\int_{-\Delta}^{0}\tilde{L}_1(x)dx=\frac{1}{2}(1-\log 2)$, (41) has a unique solution $p^*=\frac{2\delta}{1-\log 2}\leq 1$ and $\tilde{L}_{p^*}$ takes the form in Figure 9a; (ii) if $\frac{1}{2}(1-\log 2)<\delta<\int_{-\Delta}^{0}\tilde{L}_{4/3}(x)dx=\frac{1}{3}$, (41) has a unique solution $p^*\in(1,\frac{4}{3})$ as defined in (34) and $\tilde{L}_{p^*}$ takes the form in Figure 9b; (iii) if $\delta\geq\frac{1}{3}$, which implies (4), (41) has a unique solution $p^*=1+\delta$ and $\tilde{L}_{p^*}$ takes the form in Figure 9c. ■

Proof of Proposition 6. Here we explain how to solve the optimization problem in (37). We first show that we can restrict attention to a family of parameterized distributions, similar to the method in Roesler and Szentes (2017). Consider a G which satisfies (36) and $\underline{\pi} \leq 2\tau$, and let $\pi = \max_x(\underline{\pi} - x)G(x)$. Then we have $\frac{\underline{\pi}}{2} \leq \pi \leq \tau + \frac{\underline{\pi}}{2}$. (The first inequality is because $\pi \geq \underline{\pi}G(0) \geq \frac{\underline{\pi}}{2}$, and the second follows with similar logic to that used to prove $\underline{\pi} \leq 2\tau$ in the main text.) Construct a new symmetric posterior, with parameters π and $m \geq 2\pi$, defined in the negative range as

$$G_\pi^m(x) = \begin{cases} 0 & \text{if } x < \underline{\pi} - m \\ \frac{\pi}{\underline{\pi} - x} & \text{if } x \in [\underline{\pi} - m, \underline{\pi} - 2\pi] \\ 1/2 & \text{if } x \in (\underline{\pi} - 2\pi, 0] \end{cases} \tag{42}$$

which has a mass point at $x = \underline{\pi} - m$. (See Figure 11 for an illustration of G_π^m as the bold curve, with the original G as the dashed curve.) Given $\frac{\underline{\pi}}{2} \leq \pi \leq \tau + \frac{\underline{\pi}}{2}$ and $\tau \leq \delta \leq \frac{\Delta}{2}$, we have $-\Delta \leq \underline{\pi} - 2\pi \leq 0$. It can be shown that there exists $\hat{m} \in [2\pi, \underline{\pi} + \Delta]$ such that (i) $\max_x(\underline{\pi} - x)G_\pi^{\hat{m}}(x) = \pi$, (ii) $\int_{-\Delta}^u G_\pi^{\hat{m}}(x)dx \leq \int_{-\Delta}^u G(x)dx$ for all $u \leq 0$, and (iii) $\int_{-\Delta}^0 G_\pi^{\hat{m}}(x)dx = \tau$.[27] In other words, $G_\pi^{\hat{m}}$ also satisfies (36) and generates the same maximum value π. Therefore, to solve (37) we can restrict attention to the class of distributions (42).

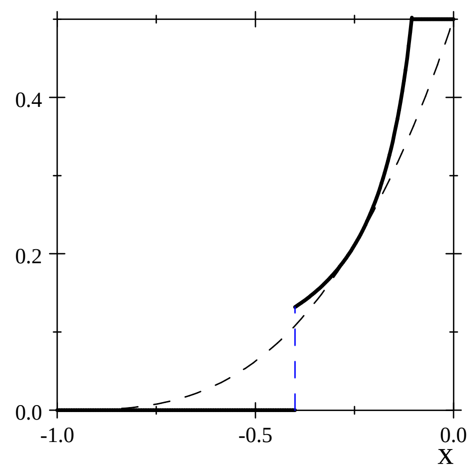

Figure 11: G and G_π^m

Condition (iii) implies that the required $\hat{m}$ satisfies

$$\int_{-\Delta}^0 G_\pi^{\hat{m}}(x)dx = \pi\left(1 + \log\frac{\hat{m}}{2\pi}\right) - \frac{\underline{\pi}}{2} = \tau\ ,$$

[27]Property (i) is easy to verify for any m from the construction of G_π^m. To see properties (ii) and (iii), notice that for $x < 0$, $(\underline{\pi} - x)G(x) \leq \pi$ implies $G(x) \leq \frac{\pi}{\underline{\pi} - x}$. Then G_π^m is above G in the range $[\underline{\pi} - m, 0]$ and below G in the range $[-\Delta, \underline{\pi} - m)$. Notice also that as m varies from 2π to $\underline{\pi} + \Delta$, $\int_{-\Delta}^0 G_\pi^m(x)dx$ increases from $\int_{-\Delta}^0 G_\pi^{2\pi}(x)dx = \pi - \frac{\underline{\pi}}{2} \leq \tau$ to $\int_{-\Delta}^0 G_\pi^{\underline{\pi}+\Delta}(x)dx \geq \int_{-\Delta}^0 G(x)dx = \tau$. Hence, there exists $\hat{m} \in [2\pi, \underline{\pi} + \Delta]$ such that $\int_{-\Delta}^0 G_\pi^{\hat{m}}(x)dx = \tau$, which is property (iii). At the same time, since G_π^m crosses G from below only once, property (iii) implies property (ii).

and so the member of the family (42) which satisfies condition (iii) has m given by

$$m(\pi) = 2\pi \exp\left(\frac{2\tau + \underline{\pi}}{2\pi} - 1\right) \ . \tag{43}$$

From now on, regard m as the function of π defined in (43). One can check that m decreases in π when $\pi \leq \tau + \frac{\underline{\pi}}{2}$. Therefore, as π decreases the range of x in which G_{π}^{m} is positive expands, but at the same time in this range G_{π}^{m} becomes lower. Given $\underline{\pi}$ and τ, we look for the smallest π such that G_{π}^{m} in (42) is an MPC of F.

Notice that G_{π}^{m} is increasing and log-convex in x in the range of $[\underline{\pi}-m, \underline{\pi}-2\pi]$, while the prior CDF $F(x)$ is log-concave given its density is log-concave. Therefore, G_{π}^{m} and F cross at most twice in the range of $[\underline{\pi}-m, \underline{\pi}-2\pi]$. At the smallest π, they must cross twice. (Otherwise, either G_{π}^{m} is everywhere above F in the range of $[\underline{\pi}-m, \underline{\pi}-2\pi]$, or it crosses F only once and from below. In either case we can reduce π further without violating the MPC constraint given $\int_{-\Delta}^{0} G_{\pi}^{m}(x)dx = \tau \leq \delta = \int_{-\Delta}^{0} F(x)dx$.) Let x_1 be the smaller of the two roots of $\frac{\pi}{\underline{\pi}-x} = F(x)$. Then G_{π}^{m} is an MPC of F if and only if $\int_{-\Delta}^{x_1} G_{\pi}^{m}(x)dx \leq \int_{-\Delta}^{x_1} F(x)dx$. The optimal π (which needs to be $\underline{\pi}$ according to the definition of $\underline{\pi}$ in (37)) should make this inequality bind. That is, $\underline{\pi}$ satisfies

$$\underline{\pi} \log \frac{m(\underline{\pi})}{\underline{\pi} - x_1} = \int_{-\Delta}^{x_1} F(x)dx \ .$$

Using (43) this equation can be written as

$$\tau = \underline{\pi}\left(\tfrac{1}{2} - \log \frac{2\underline{\pi}}{\underline{\pi} - x_1}\right) + \int_{-\Delta}^{x_1} F(x)dx \ ,$$

which gives $\underline{\pi}(\tau)$ in (37) implicitly as a function of τ.

Recall that our upper bound on consumer surplus is $\mu + \max_{\tau \in [0,\delta]}[\tau - 2\underline{\pi}(\tau)]$. Therefore, we wish to maximize

$$\mu + \tau - 2\underline{\pi} = \mu - \underline{\pi}\left(\tfrac{3}{2} + \log \frac{2\underline{\pi}}{\underline{\pi} - x_1}\right) + \int_{-\Delta}^{x_1} F(x)dx$$

by choosing $\underline{\pi}$. The definition of x_1 implies $\underline{\pi} = -\frac{x_1 F(x_1)}{1-F(x_1)}$. Substituting this into the above yields

$$\mu + \tau - 2\underline{\pi} = \mu + \frac{x_1 F(x_1)}{1 - F(x_1)}\left(\tfrac{3}{2} + \log(2F(x_1))\right) + \int_{-\Delta}^{x_1} F(x)dx \ , \tag{44}$$

and we wish to choose x_1 to maximize the right-hand side of (44). However, not all $x_1 < 0$ are feasible. Since x_1 is the smaller root of $\frac{\pi}{\underline{\pi}-x} = F(x)$, $\frac{\pi}{\underline{\pi}-x}$ is flatter than $F(x)$

at $x = x_1$, which requires $\frac{\underline{\pi}}{(\underline{\pi}-x_1)^2} < f(x_1)$, or equivalently $F(x_1) + \frac{x_1 f(x_1)}{1-F(x_1)} < 0$. The derivative of (44) with respect to x_1 is

$$\frac{1}{1-F(x_1)}\left(F(x_1)+\frac{x_1 f(x_1)}{1-F(x_1)}\right)\times\left(\tfrac{5}{2}-F(x_1)+\log(2F(x_1))\right) \ .$$

The first part is negative for any feasible x_1, and the second part increases with x_1, being negative at $x_1 = -\Delta$ and positive at $x_1 = 0$. Therefore, (44) is maximized at x^* which solves $F(x) - \log(2F(x)) = \frac{5}{2}$ or $F(x^*) \approx 0.043$.[28] Using this first-order condition, (44) simplifies to $\mu - x^*F(x^*) + \int_{-\Delta}^{x^*} F(x)dx = \mu - \int_{-\Delta}^{x^*} xdF(x)$, which is the same expression as in (20), but with a different x^*.

Finally, we show that consumer surplus with pure strategies comes close to reaching this upper bound. Let $x_H < 0$ be the optimal x for the upper bound, i.e., $F(x_H) \approx 0.043$, and let $x_L < x_H$ be the optimal x in the optimal consumer surplus with the pure-strategy restriction, i.e., $F(x_L) \approx 0.026$. The ratio of consumer surplus with pure strategies to this upper bound is

$$\frac{\mu-\int_{-\Delta}^{x_L} xdF(x)}{\mu-\int_{-\Delta}^{x_H} xdF(x)} = \frac{\mu - x_LF(x_L)+\int_{-\Delta}^{x_L}F(x)dx}{\mu - x_HF(x_H)+\int_{-\Delta}^{x_H}F(x)dx} > \frac{\mu - x_LF(x_L)+\int_{-\Delta}^{x_L}F(x)dx}{\mu - x_LF(x_H)+\int_{-\Delta}^{x_L}F(x)dx}$$

$$> \frac{\mu - x_LF(x_L)}{\mu - x_LF(x_H)} > \frac{\mu+\Delta F(x_L)}{\mu+\Delta F(x_H)} \geq \frac{\Delta+\Delta F(x_L)}{\Delta+\Delta F(x_H)} = \frac{1+F(x_L)}{1+F(x_H)} \approx 98.4\% \ .$$

Here, the first inequality uses $\int_{x_L}^{x_H} F(x)dx < (x_H - x_L)F(x_H)$, the third inequality uses $-\Delta < x_L < 0$, and the final inequality uses the fact that $V \geq \Delta$ implies $\mu \geq \Delta$. ■

[28] We need to verify that $F(x^*) + \frac{x^*f(x^*)}{1-F(x^*)} < 0$ so that x^* is feasible. The log-concavity of F implies $\log F(x) + (-x)(\log F(x))' \geq \log\frac{1}{2}$ for $x < 0$, from which we derive $xf(x) \leq F(x)\log(2F(x))$. Then $F(x) + \frac{xf(x)}{1-F(x)} \leq F(x)(1 + \frac{\log(2F(x))}{1-F(x)})$, and the right-hand side of this is negative at $x = x^*$.