Consumer Information and the Limits to Competition

Mark Armstrong Jidong Zhou*

July 2021

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 induce pure-strategy pricing equilibria, we derive signal structures which are optimal for firms and those which are optimal for consumers. The firm-optimal policy amplifies underlying product differentiation, thereby relaxing competition, while ensuring consumers purchase their preferred product, thereby maximizing total welfare. The consumer-optimal policy dampens differentiation, which intensifies competition, but induces some consumers to buy their less-preferred product. Our analysis sheds light on the limits to competition when the information possessed by consumers can be designed flexibly.

Keywords: Information design, Bertrand competition, product differentiation, online platforms.
JEL classification: D43, D83, L13

1 Introduction

Firms in many markets supply differentiated products and compete in prices. Consumers might not always possess all information about product attributes before they

*Armstrong: Department of Economics, Oxford University (mark.armstrong@economics.ox.ac.uk); Zhou: School of Management, Yale University (jidong.zhou@yale.edu). We are grateful to the coeditor, four anonymous referees, Jason Abaluck, Heski Bar-Isaac, Dirk Bergemann, Glenn Ellison, Joe Farrell, Nima Haghpanah, Hanna Halaburda, Ju Hu, Bruno Jullien, Fei Li, Stephen Morris, Xiaosheng Mu, Barry Nalebuff, Andrew Rhodes, David Sappington, Armin Schmutzler, Xianwen Shi, Philipp Strack, Jean Tirole, John Vickers, Weijie Zhong and various seminar audiences for helpful comments. Armstrong thanks the European Research Council for financial support from Advanced Grant 833849.

make their purchase decision, and market performance depends on how much information consumers possess about products. If little such information is available, consumers cannot meaningfully compare products and will regard them as close substitutes. This induces firms to compete fiercely with low prices, but consumers often end up purchasing a mismatched product. When more product information is available, the quality of the consumer-product match improves, but firms enjoy greater market power and set higher prices. This broad trade-off between match quality and price is well known, but there has been little systematic investigation of the interaction between the amount of product information consumers have and the intensity of competition. For example, is more information for consumers always associated with increased market power of firms? Can biased information provision which favors one firm over others help relax price competition? At which point is the trade-off between match quality and market price optimally resolved for consumers? More generally, what kind of information provision is optimal for firm profit and what kind is optimal for consumer surplus? This paper aims to answer these questions.

These questions are also important when a third-party can design the information environment that consumers face. For instance, consumers often gather information and make their purchases via a platform (e.g., Amazon or Expedia) which hosts competing firms. The platform can control many aspects of the information environment, such as how much detail about product characteristics is revealed, whether to post customer reviews and/or its own reviews, whether to offer personalized recommendations, how flexibly consumers can filter and compare products, and so on. How should the platform design its information environment if it aims to maximize industry profit, or the profit of selected firms, or consumer surplus, or a weighted sum of profit and consumer surplus? Alternatively, when some products (e.g., insurance plans) are difficult for consumers to compare, regulators sometimes intervene by specifying how firms should disclose product information, aiming to enhance consumers' ability to select their preferred product. But is a more transparent information environment always better for consumers, given the potentially adverse effect on price competition when more detailed product information is revealed?[1]

Formally, as described in section 2, we study an *ex ante* symmetric duopoly market where two firms each costlessly supply a single variety of a product and compete in

[1] See, e.g., https://go.cms.gov/2U2lZJe for information transparency regulation in the US health insurance market. Sometimes the regulator may even require firms to offer "standardized" products to facilitate comparison. This involves a similar trade-off to information regulation: more standardized products are close substitutes and so facilitate competition on price, but prevent some consumers with particular preferences from obtaining a product tailored to those preferences. See, e.g., Ericson and Starc (2016) for a discussion and empirical analysis of this issue in the context of health insurance.

prices. Consumers are initially uncertain about their preferences for the varieties, but before purchase they receive a private signal of these preferences. Consumers then update their beliefs about their preferences and make their choice given the pair of prices offered by firms. The information environment or *signal structure*, which governs the mapping between consumers' true preferences and the signals they receive, is publicly known to both firms and consumers. We wish to understand how the signal structure affects competition and welfare. In particular, we explore the limits to competition in this market, and derive the signal structures that induce the highest profit for firms and those that generate the highest surplus for consumers.[2]

Solving this problem at the most general level is challenging. Consumer preferences are generally two-dimensional in our duopoly setup, and current understanding of information design in such cases is limited.[3] For this reason, we mostly focus on the case when the outside option for consumers is irrelevant, and so only the relative valuation between the two products—a scalar variable—matters for consumer decisions. Even with scalar heterogeneity, under some signal structures the only equilibria in the pricing game between firms involve mixed strategies, which can be hard to deal with. For this reason, we mostly focus on signal structures which induce a pure-strategy equilibrium in the pricing game.

Our analysis starts in section 3.1 where we study a stand-alone pricing problem which does not concern information design *per se*. We describe the set of relative valuation distributions, if any, which support a given pair of prices as equilibrium prices. It turns out that a price pair can be implemented if and only if the valuation distribution lies between two *bounds*, which are determined by each firm's no-deviation condition. Using these bounds we characterize the set of possible pure-strategy equilibrium prices when any valuation distribution is allowed.

In the information design problem, however, only some posterior distributions can be induced by signal structures. Using the bounds approach, we show in section 3.2 for both firms and consumers that any asymmetric signal structure which treats the two firms differently is dominated by a symmetric signal structure. Surprisingly, this

[2]Another interpretation of the consumer-optimal problem is that sometimes consumers can commit to how much product information to acquire, or use, before firms make their pricing decisions. For instance, a consumer could delegate her purchase decision to an agent who commits to focus more on price than other product characteristics in order to stimulate competition between suppliers.

[3]It is well known that the posterior (expected) consumer valuation distribution induced by any signal structure is a mean-preserving contraction of the underlying prior distribution. However, a multidimensional mean-preserving contraction currently has no simple characterization, in contrast to the scalar case where Rothschild and Stiglitz (1970) applies. See section 7.2 of Dworczak and Martini (2019) for discussion of this point.

is true even for an individual firm, i.e., the firm which is favored by an asymmetric signal will prefer some symmetric signal structure which treats it and its rival the same. Intuitively, when a firm is treated unfavorably with asymmetric signals, it has an incentive to set a low price. Our result shows that this force is sufficiently strong so that even the favored firm suffers from this fierce competition. This result suggests that firms have congruent interests when it comes to the design of consumer information, and that biased recommendations are not the best approach to improve profit even for an individual firm.

The bounds are used in sections 3.3 and 3.4 to characterize the (necessarily symmetric) firm-optimal and consumer-optimal information policies. The firm-optimal signal structure amplifies perceived product differentiation in order to relax competition, and does so by reducing the likelihood that consumers are near-indifferent between products after receiving their signal. Compared to full information disclosure, the firm-optimal information policy can significantly improve industry profit. This shows that providing consumers with more product information need not raise firms' market power. The firm-optimal policy typically enables consumers to buy their preferred product, in which case total welfare is also maximized. The consumer-optimal signal structure by contrast dampens perceived product differentiation in order to stimulate competition, and does so by increasing the number of consumers who are near-indifferent between products. In the consumer-optimal policy, a consumer with strong preferences can buy her preferred product for sure, but a consumer with weaker preferences receives less precise information and may end up with the inferior product. Thus, in contrast to the firm-optimal policy, the consumer-optimal policy does not maximize total welfare. The consumer-optimal policy typically induces rather low prices, suggesting that for information design with two firms the price competition effect tends to outweigh the match quality effect for consumers.

In section 3.5 we characterize all combinations of industry profit and consumer surplus which can be implemented using some choice of signal structure. Any profit between zero and the firm-optimal profit level is feasible. The efficient (welfare-maximizing) frontier is achievable if the profit is at least equal to that obtained under full disclosure; otherwise, a degree of mismatch is needed to induce fierce competition and efficiency is not feasible. This analysis can be used to solve the Ramsey problem, i.e., the information policy which maximizes a weighted sum of industry profit and consumer surplus.

In section 4 we discuss two extensions to this analysis. First, in section 4.1 we explore whether considering the wider class of signals which allow mixed-strategy pricing equilibria can improve outcomes. We show that allowing symmetric mixed-strategy equilibria could at best improve consumer surplus only slightly.

Second, in section 4.2 we extend the analysis to allow for more than two firms and for the outside option to be relevant for consumers. Optimal information design in this general framework is a hard problem, and instead we study the relative performance of three simple policies. We show that the "top product" policy, which informs consumers of their most preferred product but nothing else, can sometimes enable firms to achieve first-best profit (i.e., the profit equal to maximum total welfare). Moreover, even if this policy does not achieve first-best profit, it typically performs better than full disclosure of information for firms and bounds industry profit away from zero regardless of the number of firms. When there are *many* firms, the trade off for consumers between match quality and a low price vanishes. For instance, a policy of full disclosure can approximately achieve the first-best for consumers when there are many firms. A third policy, which we term the "top two" policy, informs consumers of their best *two* products (without ranking them). This policy induces firms to set price equal to marginal cost, and when there are many firms it causes negligible mismatch and so also induces approximately the first-best outcome for consumers. For any finite number of firms this "top two" policy performs better for consumers than both full disclosure and the "top product" policy.

Related literature. One strand of the relevant literature considers a monopolist's incentive to provide information to enable consumers to discover their valuation for its product. An early paper on this topic is Lewis and Sappington (1994), who study a monopoly market and show, within the class of "truth or noise" signal structures, how 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). Importantly, they allow for general signal structures as in the more recent Bayesian persuasion literature and show that firm-optimal information disclosure takes the coarse form whereby a consumer is informed merely whether her valuation lies above a threshold.

Roesler and Szentes (2017) study the signal structure which is best for consumers in a monopoly model. They show that the optimal signal structure can be found within the class of posterior distributions which induce unit-elastic demand functions. (These unit-elastic demand functions play a similar role to the posterior distribution bounds in our analysis.) They show that partial rather than complete learning is optimal for consumers, and that the optimal information structure induces *ex ante* efficient trade and maximizes total welfare.[4] In their setup, where trade is always efficient, the firm-optimal signal structure is simply to disclose no information at all, in which case the firm can extract 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. (Indeed in our duopoly model we 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. With competition, the consumer-optimal policy also exhibits some significant differences from the monopoly case in Roesler and Szentes. For example, it usually causes product mismatch so that the allocation is sometimes inefficient, the induced residual demand for each firm is unit-elastic for upward but not downward price deviations, and the consumer-optimal signal structure is not the least profitable policy for firms.

Our paper concerns information design in an oligopoly setting. Most of the previous research on this topic studies the "decentralized" disclosure policies of individual firms. For example, Ivanov (2013) studies competitive disclosure when each firm decides how much information about its own product to release and what price to charge. He focuses on information structures which rotate demand as in Johnson and Myatt (2006), and shows that full disclosure is the only symmetric equilibrium when the number of firms is large.[5] Hwang, Kim, and Boleslavsky (2020) show that the same result holds if general signal structures are allowed (and more generally they are able to show that increasing the number of firms induces each firm to reveal more information). Intuitively, with many firms, a consumer's valuation for the best rival product (if other firms fully disclose their information) is high. To compete for the consumer, a firm discloses all information as that is the policy which maximizes the posterior probability she has a high valuation.[6]

Instead of studying equilibrium disclosure by individual firms, though, we focus on a "centralized" design problem, such as when a platform mediates the information flow

[4] Condorelli and Szentes (2020) study the related problem of how to choose the demand curve to maximize consumer surplus, given the monopolist chooses its price optimally in response. (The consumer accurately observes her realized valuation in this model.) Choi, Kim, and Pease (2019) extend Roesler and Szentes (2017) to the set-up of Anderson and Renault (2006), and derive the consumer-optimal policy in the context of a search good.

[5] Bar-Isaac, Caruana, and Cunat (2012) study competitive product design within a sequential search market. They consider designs which rotate demand, and show that a reduction in the search cost induces more firms to choose niche product design (which can be interpreted as full information disclosure in the context of information design when consumers have a common prior).

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

from products to consumers. This enables us to discuss signals which reflect *relative* valuations across products, such as the "top product" signal. This more general signal structure introduces a number of additional features. In our framework, for instance, full information disclosure is not the firm-optimal policy even with many firms, and the "top product" policy sometimes yields first-best profit for firms (which can never be achieved with a decentralized design). In addition, when decentralized signals are not fully revealing, there are welfare losses since consumers sometimes choose a less preferred product, while with a centralized structure it is possible to have a coarse signal structure (e.g., the "top product" signal) which nevertheless maintains efficiency. Finally, Hwang *et al.* (2020) show that the equilibrium price (and hence profit) often falls relative to the full-information price when individual firms choose their disclosure policy non-cooperatively, while by construction the centralized firm-optimal policy must boost industry profit.

Other papers have also studied "centralized" aspects of the design of consumer information. Anderson and Renault (2009) study comparative advertising in a duopoly model where each firm unilaterally chooses between fully disclosing its own product information, fully disclosing information about both products, or disclosing nothing. Among other results, they make the point that disclosing more information improves match quality but also softens price competition. Jullien and Pavan (2019) make a similar point when they study how platforms' information management affects their competition in two-sided markets. Both papers consider specific classes of signal structures. When general signal structures are allowed, as we show in this paper, the monotonic relationship between the amount of information and the degree of price competition fails. Dogan and Hu (2021) study the consumer-optimal disclosure policy 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. The information policy is chosen by a third party, and the disclosure is only about each individual product (and so there is no disclosure about relative valuations across products). Because the reservation value in this search framework is static, their problem is similar to the monopoly problem studied by Roesler and Szentes (2017). Moscarini and Ottaviani (2001) study a duopoly model of price competition similar to ours, where the consumer receives a private signal of her relative valuation for the two products. A major difference, however, is that they assume the relative valuation is binary and the signal is binary and symmetric across states. In that simple setting, they can allow firms to be *ex ante* asymmetric, but the pricing equilibrium often involves complicated mixed strategies. Moscarini and Ottaviani (2001) also point out that both firm profit and consumer welfare can vary non-monotonically with the amount of information consumers have, but they do not explicitly study the optimal signal structure for firms or consumers.

More broadly, our paper belongs to the recent literature on Bayesian persuasion and information design. See Kamenica and Gentzkow (2011) for a pioneering paper in this literature, and Bergemann and Morris (2019) and Kamenica (2019) for recent surveys. Bergemann, Brooks, and Morris (2015) study third-degree price discrimination by a monopolist. In contrast to the consumer-side design problem in Roesler and Szentes (2017) and our model, their paper considers a firm-side design problem where signals are sent to the firm about consumer preferences, and consumers accurately know their valuation from the start. A given signal structure corresponds to a particular partition of consumers. If all ways to partition consumers are possible, the paper shows that any combination of profit (above the no-discrimination benchmark) and consumer surplus which sum to no more than maximum total welfare can be implemented. Elliot, Galeotti, and Koh (2020) extend Bergemann *et al.* to the case of competition with product differentiation, and derive conditions under which flexible market segmentation can earn firms their first-best profit. Li (2020) solves the consumer-optimal problem in the case of competition with product differentiation by applying techniques from Bergemann, Brooks, and Morris (2015). He shows that the optimal policy is for firms to see a public signal which induces a pure-strategy pricing equilibrium, and involves partitioning consumers into sub-markets within each of which one firm is "dominant" and other firms set price equal to marginal cost. Unlike our consumer-optimal policy, with this policy consumers always buy their preferred product.[7]

2 The model

A risk-neutral consumer wishes to buy a single unit of a differentiated product which is costlessly supplied by two risk-neutral firms, 1 and 2. The consumer's valuation for the unit from firm $i = 1, 2$ is denoted $v_i \geq 0$. The consumer is initially uncertain about her valuations and holds a prior belief about the joint distribution of (v_1, v_2).[8]

[7] Similar problems are also studied in the situation where firms supply a homogeneous product but consumers consider different sets of firms. See, for instance, Armstrong and Vickers (2019), Bergemann, Brooks, and Morris (2020), and Shi and Zhang (2020). In many circumstances, market segmentation also depends on how much preference information consumers are willing to share with firms. See, e.g., Ali, Lewis, and Vasserman (2020) and Ichihashi (2020) for relevant works which consider general information structures. More broadly, the firm-side design problem is also related to the literature on information design in auctions where how much information competing bidders receive on their valuations for the object affects their bidding strategy and welfare (see, e.g., Bergemann, Brooks, and Morris (2017), and Du (2018)).

[8] Our model can also be interpreted as having a continuum of consumers who have preferences (v_1, v_2) independently drawn from the same prior distribution.

Throughout the paper we assume that firms are symmetric *ex ante*, in the sense that the prior distribution for (v_1, v_2) is symmetric between v_1 and v_2. We assume that the support of the prior distribution lies inside the square $[V, V+1]^2$. Here, $V \geq 0$ represents the "basic utility" from either product, while 1 is the normalized range of product valuations. Let $\mu \equiv \mathbb{E}[v_i]$ denote the expected valuation of either product (or a random product).

The consumer has an outside option which is sure to have payoff zero. Until section 4.2, however, we assume this outside option is irrelevant and the consumer always purchases one of the two products. (As shown in Lemma 0 in the appendix, this is the case if $V > 2$ and we restrict attention to signals that induce pure-strategy pricing equilibria.) In this case the consumer cares only about the difference in her valuations, $x \equiv v_1 - v_2$, and so consumer heterogeneity becomes one-dimensional. Let $F(x)$ be the prior distribution for x which is symmetric around 0 and has support in $[-1, 1]$.

Before purchase the consumer observes a private signal of x. The signal is generated according to a signal structure $\{\sigma(s|x), S\}$, where S is the signal space and $\sigma(s|x)$ specifies the distribution of signal s when the true preference parameter is x. We assume the signal structure is common knowledge to the consumer and to both firms, and determined before firms choose prices. This assumption is plausible when the signal structure is publicly chosen by some third-party such as a platform or a regulator as a relatively long-term decision. After observing signal s, the consumer updates her belief about her relative preference x to, say, $F_s(x)$. Given the consumer is risk neutral, only the expected x given s, $\mathbb{E}_{F_s}[x]$, matters for her choice. Let $G(x)$ be the distribution of this expected x, and which we refer to as the consumer's *posterior* distribution. It is jointly determined by the prior distribution for x and the signal structure. Even though some signal structures may be hard to implement in practical terms, this approach is a general way to model the information environment and enables us to explore the limits of competition when information can be flexibly designed.[9]

Firms know the consumer's prior and the signal structure, and hence know the posterior distribution $G(x)$, but they do not observe her private signal, and so they each choose a single price regardless of the signal received by the consumer. We assume that prices are always accurately observed by the consumer. If p_i is firm i's price,

[9] One interpretation of the signal structure in our setup is that an information designer (e.g., a platform) might know a consumer's preferences (e.g., due to its past interactions with the consumer) and can send the consumer a signal contingent on those preferences (e.g., personalized recommendations), but it commits to the signal structure before seeing the consumer's preferences. The more standard, also less informationally demanding, interpretation from the Bayesian persuasion literature is that the signal structure is an information experimentation device, and even consumers with the same preferences may learn different noisy signals from the same information environment.

then the consumer prefers to buy from firm 1 if her expected x after seeing a signal satisfies $x > p_1 - p_2$, prefers to buy from firm 2 if $x < p_1 - p_2$ (and is indifferent when $x = p_1 - p_2$). Firms set prices simultaneously to maximize their expected profit, and we use Bertrand-Nash equilibrium as the solution concept for the pricing game.

We aim to investigate how the signal structure affects the intensity of competition and the ability of the consumer to buy her preferred product. In particular, we search for those signal structures which maximize industry profit and those which maximize consumer surplus. As discussed in the introduction, for tractability we focus on signal structures which induce a *pure-strategy* pricing equilibrium. (We will discuss signal structures which might induce mixed-strategy equilibria in section 4.1.)

2.1 Examples of signals

It is useful initially to consider some simple signal structures. To illustrate these, assume here that the prior distribution is uniform, so $F(x) = \frac{1}{2}(1 + x)$ for $x \in [-1, 1]$.

Full information disclosure: Here the signal perfectly reveals the true preference x, e.g., where $s \equiv x$, and so the posterior and prior distributions for x coincide. The equilibrium is symmetric and each firm sets a price equal to 1, in which case industry profit is also 1 (given full consumer participation). The consumer always buys her preferred product, which maximizes total welfare, but the market price is relatively high. The consumer's expected surplus is $(\mu + \frac{1}{4}) - 1 = \mu - \frac{3}{4}$, where μ is the expected match utility of a random product, and $\frac{1}{4} = \mathbb{E}_F[\max\{x, 0\}]$ is the match efficiency improvement under full disclosure relative to random purchase.[10]

No information disclosure: In this case the signal is uninformative and the posterior G is degenerate at $x = 0$. In particular, the consumer views the two products as perfect substitutes. Both the equilibrium price and industry profit are 0, and consumer surplus is simply μ since the consumer buys a product randomly. This consumer surplus is higher than that under full information disclosure, and so the disutility she incurs from buying a random product is outweighed by the low price she pays. (As we will see in section 4.2, when there are more than two firms the "top-two" signal structure can achieve the same pricing outcome but without causing so much product mismatch.)

"Truth or rank" signal: Here the signal accurately informs the consumer of her x with probability $\theta < 1$ and otherwise only informs her whether $x > 0$ or $x < 0$, i.e.,

[10]More explicitly, total welfare is $\mathbb{E}[\max\{v_1, v_2\}] = \mathbb{E}[v_2 + \max\{v_1 - v_2, 0\}] = \mu + \mathbb{E}_F[\max\{x, 0\}]$, where the first two expectation operators $\mathbb{E}$ are with respect to the joint distribution of (v_1, v_2).

which product she prefers. (Suppose the rank signals are distinguishable from the truth signals.) Then the posterior distribution G has two mass points at $-\frac{1}{2}$ and $\frac{1}{2}$ (with mass $\frac{1-\theta}{2}$ at each) and is otherwise the same as the prior but with a reduced density $\frac{\theta}{2}$. In this case, if θ is sufficiently close to 1 (so the two mass points do not have too much weight), a symmetric equilibrium price exists and is equal to $\frac{1}{\theta} > 1$, and so firms earn more than with full disclosure.[11] (Consumer surplus, however, is lower than with full disclosure because in either case the consumer buys her preferred product.)

The first two of these examples illustrate the consumer's trade-off between match quality and the intensity of competition, and how the latter effect can dominate. The first and third examples suggest that one way to soften competition and enhance industry profit is to reduce the chance that the consumer is indifferent between the two products, and more information for consumers is not always associated with more market power of firms. The three examples together imply that full information disclosure is optimal neither for firms nor for the consumer. As we will see in section 4.2, these insights carry over to the case when there are more than two firms.

All these signal structures are symmetric in the sense that they induce a symmetric posterior G. More generally, as shown in this final example, a signal structure can induce an asymmetric posterior and make *ex ante* symmetric products appear asymmetric.

"Truth or biased recommendation" signal: Modify the previous example so that when the consumer is not provided with full information she learns whether $x > b$ or $x < b$ for some $b \in (0, 1)$. (An interpretation is that firm 1's product is recommended if $x > b$ and otherwise firm 2's product is recommended.) The posterior now has a mass point with mass $\frac{1}{2}(1-\theta)(b+1)$ at $\frac{1}{2}(b-1)$ and a mass point with mass $\frac{1}{2}(1-\theta)(1-b)$ at $\frac{1}{2}(b+1)$. One can check that given $b \in (0, 1)$, if θ is sufficiently close to 1 there is an asymmetric pure-strategy equilibrium with prices $p_1 = \frac{1}{\theta}(1-\Delta)$ and $p_2 = \frac{1}{\theta}(1+\Delta)$, where $\Delta \equiv \frac{b}{3}(1-\theta) > 0$. Compared to the symmetric case with $b = 0$, the favored firm 2 raises its price while firm 1 lowers its price. Firm 2's equilibrium market share rises to $\frac{1}{2}(1+\Delta)$ and firm 1's falls to $\frac{1}{2}(1-\Delta)$, and so the biased recommendation helps firm 2 but harms firm 1 relative to the previous symmetric example. Industry profit is $\frac{1}{\theta}(1+\Delta^2)$, which is even greater than in the previous example. (The consumer then must suffer from this biased recommendation as it harms match efficiency and total welfare.) This final example suggests that the use of asymmetric signals might be a

[11] This is the case if $\theta \geq 3 - \sqrt{5} \approx 0.76$. If θ is below the threshold, the incentive for each firm to undercut and steal the rival's consumers on the mass points becomes so strong that there is no pure-strategy equilibrium. When $\theta = 0$, we have the "top product" signal structure which will be discussed in section 4.2.

useful tool to improve industry profit. However, as we will see in section 3.2, this turns out not to be the case once we allow for more general signal structures.

2.2 Preliminaries

To study general signal structures we use the following well-known results to simplify the problem. 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.,

$$\int_{-1}^{x} G(\tilde{x})d\tilde{x} \leq \int_{-1}^{x} F(\tilde{x})d\tilde{x} \text{ for } x \in [-1,1], \text{ with equality at } x = 1\ . \tag{1}$$

Moreover, any G which is an MPC of F can be generated by means of some signal structure.[12] Therefore, instead of analyzing the signal structure directly, we can work with the posterior distribution G and look for a posterior which is optimal for firms or consumers subject only to the MPC constraint (1). Figure 1 depicts some examples of MPCs of the uniform prior (where the dashed line represents the uniform prior CDF and the bold curve shows the posterior CDF).

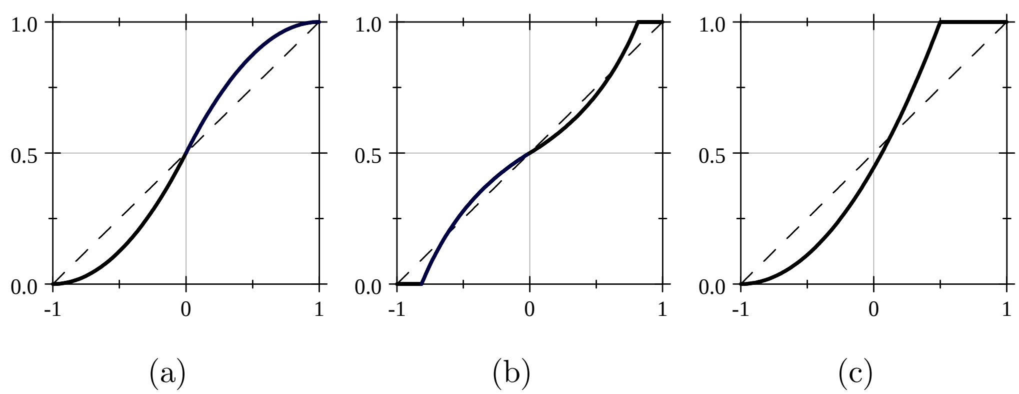

Figure 1: Some examples of MPCs

The following two observations will be useful: First, when G is symmetric as on Figures 1a and 1b, the mean-preserving requirement is satisfied automatically, and so the simpler condition

$$\int_{-1}^{x} G(\tilde{x})d\tilde{x} \leq \int_{-1}^{x} F(\tilde{x})d\tilde{x} \text{ for } x \in [-1,0] \tag{2}$$

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

is sufficient for G to be an MPC of F. Second, if a symmetric G crosses F at most once and from below in the negative range $x \in (-1, 0)$, as on Figures 1a and 1b, then

$$\int_{-1}^{0} G(\tilde{x}) d\tilde{x} \leq \int_{-1}^{0} F(\tilde{x}) d\tilde{x} \tag{3}$$

is sufficient for G to be an MPC of F.

We also need a measure of the efficiency of product choice corresponding to a (possibly asymmetric) posterior G. If the consumer always chooses her preferred product (i.e., chooses product 1 if the expected x is positive and product 2 otherwise, as would be the case if the two firms offered the same price), then as detailed in footnote 10 total welfare is

$$\mu + \mathbb{E}_G[\max\{x, 0\}] = \mu + \int_{-1}^{0} G(x) dx \ , \tag{4}$$

where the equality follows after integration by parts and uses the fact that G has a mean of zero (which implies $\int_{-1}^{1} G(x) dx = 1$). Since μ is the expected value of a random product, the integral term reflects the match efficiency improvement under posterior G relative to a policy of disclosing no information.

Define

$$\delta \equiv \mathbb{E}_F[\max\{x, 0\}] = \int_{-1}^{0} F(x) dx \ . \tag{5}$$

This parameter δ captures the underlying extent of product differentiation in the market. Since G is an MPC of F, condition (1) at $x = 0$ implies that (4) cannot exceed the maximum total surplus $\mu + \delta$. In other words, match efficiency cannot increase when the consumer observes a noisy signal of her preferences, as is intuitive. When the inequality (3) is strict—as in Figures 1a and 1c—mismatch occurs with the posterior G. There is no product mismatch when there is equality in (3). This is the case under full disclosure, and more generally when 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.

3 Optimal signal structures

We derive the firm and consumer optimal signal structures in two steps. First, we ignore the MPC constraint (1) and derive the constraints on G needed to implement a target pair of prices (p_1, p_2) in pure-strategy equilibrium. Second, we search for the prices, together with their supporting G, which are optimal for firms or the consumer, subject to the constraint that G is an MPC of F.

3.1 Bounds on posteriors

In this section we derive bounds on the distributions $G(x)$ which implement a given pair of prices (p_1, p_2) in pure-strategy equilibrium. The analysis here is distinct from the information design problem, and may have some independent interest for oligopoly theory.

Define

$$L_{p_1,p_2}(x) \equiv \begin{cases} 0 & \text{if } x \le \pi_1 - p_2 \\ 1 - \dfrac{\pi_1}{p_2 + x} & \text{otherwise} \end{cases} \tag{6}$$

and

$$U_{p_1,p_2}(x) \equiv \begin{cases} \dfrac{\pi_2}{p_1 - x} & \text{if } x < p_1 - \pi_2 \\ 1 & \text{otherwise ,} \end{cases} \tag{7}$$

where

$$\pi_i = \frac{p_i^2}{p_1 + p_2} \ . \tag{8}$$

(Here π_i in (8) is continuous at $p_1 = p_2 = 0$, and so we set $\pi_i = 0$ in this case.) The following lemma shows that these functions are the required bounds on G.

Lemma 1 *(i) A distribution G implements (p_1, p_2) as equilibrium prices if and only if*

$$L_{p_1,p_2}(x) \le G(x) \le U_{p_1,p_2}(x) \tag{9}$$

for all $x \in [-1, 1]$, where the two bounds are defined in (6) and (7) and π_i in (8) is firm i's equilibrium profit; furthermore $p_1 = p_2 > 0$ if and only if $G(0) = \frac{1}{2}$.
(ii) (p_1, p_2) can be equilibrium prices for some G if and only if $p_2 - 1 \le \pi_1$ and $p_1 - 1 \le \pi_2$ (which implies $|p_1 - p_2| \le 1$ and $p_i \le 2$).

The two bounds L_{p_1,p_2} and U_{p_1,p_2} arise from each firm's no-deviation requirement for equilibrium. For example, if firm 2 deviates to price $p_2' \neq p_2$, the consumer buys from firm 2 if $x \le p_1 - p_2'$. (If G has a mass point at $x = p_1 - p_2'$, the natural tie-breaking rule is that firm 2 serves all consumers at that mass point, since firm 2 can always achieve this outcome by offering a price slightly below p_2'.) Therefore, firm 2 has no incentive to deviate if and only if

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

for all p_2', where π_2 is firm 2's equilibrium profit and which is shown in the proof to take the form (8). By changing variables from p_2' to $x = p_1 - p_2'$, we can write this no-deviation requirement as $(p_1 - x)G(x) \le \pi_2$, or $G(x) \le U_{p_1,p_2}(x)$ for all $x \in [-1, 1]$. A similar argument implies that firm 1 has no incentive to deviate if and only if $G(x) \geq L_{p_1,p_2}(x)$. (See the proof of Lemma 1 for further details.) The way these two bounds are derived is similar to the monopoly analysis in Roesler and Szentes (2017): the upper bound corresponds to firm 2's unit-elastic demand curve while the lower bound corresponds to firm 1's. The two bounds are illustrated in Figure 2 below.

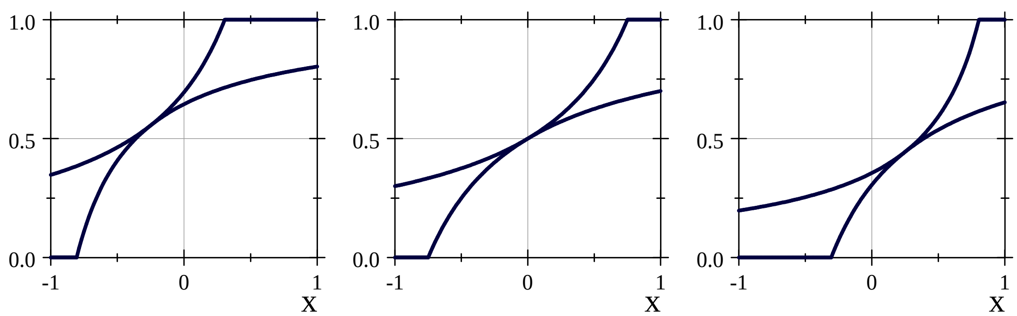

(a) $p_1 < p_2$ (b) $p_1 = p_2$ (c) $p_1 > p_2$

Figure 2: Bounds on G to implement pure-strategy equilibrium prices

The lower bound L_{p_1,p_2} increases with x and begins to be positive at $x = \pi_1 - p_2$ (which exceeds -1 given that Lemma 1 implies $p_2 - 1 \leq \pi_1$), and it is concave whenever it is positive. The upper bound U_{p_1,p_2} also increases with x and reaches 1 at $x = p_1 - \pi_2$ (which is below 1 given $p_1 - 1 \leq \pi_2$), and it is convex whenever it is less than 1. The two bounds coincide and have the same slope at $x = p_1 - p_2$.

Part (i) of Lemma 1 says that for a distribution G to induce a pair of equilibrium prices (p_1, p_2), it needs merely to lie between the two bounds, and in particular G can be irregular (e.g., the distribution can have atoms). But G has to be differentiable at $x = p_1 - p_2$ given the two bounds are tangent to each other at that point. In Figure 2a, we have $G(0) > \frac{1}{2}$ for any G above the lower bound, i.e., it is more likely that the consumer prefers firm 2's product, and that is why firm 2 sets a higher price; the opposite is true in Figure 2c.

From the equilibrium profit expression (8), firm i's equilibrium market share is $p_i/(p_1 + p_2)$. Hence, both equilibrium profits and market shares are determined entirely by equilibrium prices and do not depend separately on the shape of G, and the firm with the higher equilibrium price necessarily has higher market share and higher profit. Part (i) also shows that if the distribution is symmetric and induces a pure-strategy equilibrium, the equilibrium prices must be symmetric. (Notice, however, that an asymmetric distribution also induces symmetric prices if $G(0) = \frac{1}{2}$.)

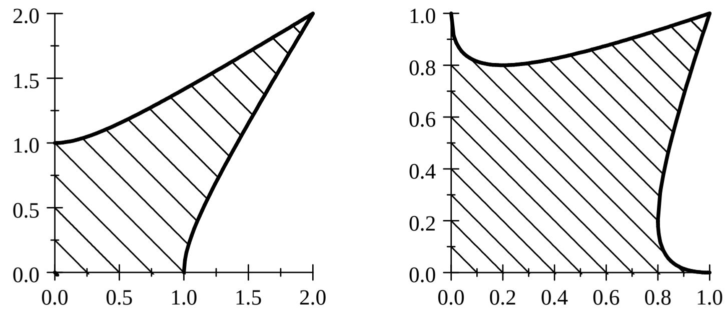

(a) prices (p_1, p_2) (b) profits (π_1, π_2)

Figure 3: Set of possible pure-strategy equilibrium outcomes

Part (ii) of Lemma 1 characterizes the set of possible pure-strategy equilibrium prices when G can be any distribution on $[-1, 1]$. This set is depicted as the shaded area in Figure 3a. In particular, the highest possible pure-strategy equilibrium price for either firm is 2 and this can be implemented by distributions lying between the bounds shown in Figure 4c below; the most asymmetric prices are $p_i = 1$ and $p_j = 0$ and these can be implemented by a distribution which is degenerate at $x = -1$ or degenerate at $x = 1$. Likewise, one can derive all possible pairs of profits (π_1, π_2) as described in Figure 3b.[13] It is clear that the best outcome for each firm is the symmetric outcome with $p_1 = p_2 = 2$ in which case each earns profit 1. This suggests that firms have congruent interests when it comes to jointly choosing the valuation distribution (e.g., via joint product design). No firm will be happy with an asymmetric distribution even if it is favored and earns more than its rival. As we will see in next section, a similar result holds when we consider the more constrained information design problem, when some (posterior) distributions are not feasible due to the MPC constraint.

The symmetric case $p_1 = p_2 = p$, shown on Figure 2b, is important in the following analysis. With symmetric prices the bounds simplify to

$$L_p(x) = \begin{cases} 0 & \text{if } x \leq -\frac{1}{2}p \\ 1 - \dfrac{p}{2(p+x)} & \text{otherwise} \end{cases} \tag{10}$$

[13]Figure 3a depicts the range of (p_1, p_2) satisfying $(p_j - 1)(p_i + p_j) \leq p_i^2$ for $i, j \in \{1, 2\}$ and $i \neq j$; Figure 3b depicts the range of (π_1, π_2) satisfying $\pi_j + \sqrt{\pi_1 \pi_2} \leq 1 + \pi_i$ for $i, j \in \{1, 2\}$ and $i \neq j$. (The latter inequality follows from $p_j - 1 \leq \pi_i$ and $p_j = \pi_j + \sqrt{\pi_1 \pi_2}$ as solved from (8).)

and

$$U_p(x) = \begin{cases} \dfrac{p}{2(p-x)} & \text{if } x < \frac{1}{2}p \\ \\ 1 & \text{otherwise .} \end{cases} \tag{11}$$

As indicated on Figure 2b above, these 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 $x \in [-1, 0]$, it will lie between the bounds over the whole range $[-1, 1]$. A property of the two bounds (10) and (11), illustrated in Figure 4, is that they rotate clockwise about the point $(0, \frac{1}{2})$ as p increases, and in particular both bounds increase with p for $x < 0$. Intuitively, to induce a higher price we need fewer price-sensitive consumers around $x = 0$, which requires the bounds to be flatter. (The dashed line in each plot indicates the CDF corresponding to the uniform prior and can be ignored for the current discussion.) Figure 4c also demonstrates why the highest possible price is 2: once the price exceeds 2, $U_p(1) < 1$ and likewise $L_p(-1) > 0$, and so these cannot bound a valid CDF.

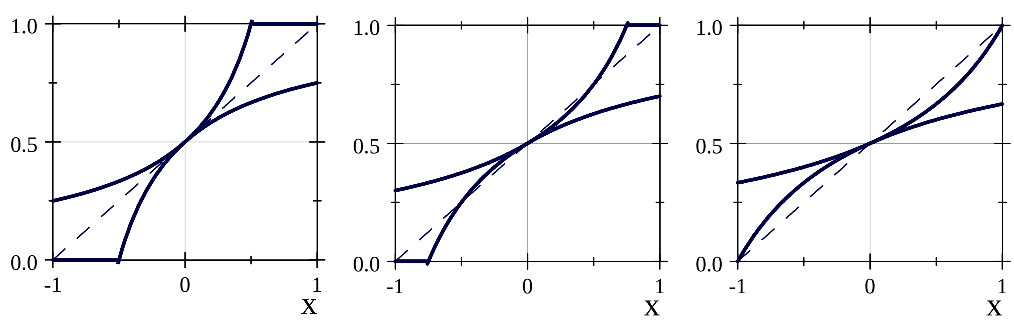

(a) $p = 1$ (b) $p = 1.5$ (c) $p = 2$

Figure 4: Bounds on G to implement symmetric equilibrium prices

In the following, we often consider the case when the market has a pure-strategy pricing equilibrium with a positive price in the benchmark with full disclosure. Lemma 1 implies that this is the case if and only if the (symmetric) prior distribution F lies between the bounds L_p and U_p for some $p > 0$. Since the two bounds are tangent at $x = 0$, F must be differentiable at that point and let $f(0)$ be the corresponding density. Then the symmetric equilibrium price under full information is $p_F = 1/(2f(0))$.

Condition 1 *F lies between the bounds L_{p_F} and U_{p_F} with $p_F = 1/(2f(0))$.*

Oligopoly models often invoke one of the following stronger conditions to ensure the existence of a pure-strategy equilibrium:

Condition 2 *(i)* $1 - F$ *is log-concave on* $[-1, 0]$ *and (ii)* F *is log-concave on* $[-1, 0]$.

Condition 3 F *has density* f *which is log-concave on* $[-1, 1]$.

It is known from Caplin and Nalebuff (1991) and Bagnoli and Bergstrom (2005) that Condition 3 implies both $1 - F$ and F are log-concave on $[-1, 1]$, and hence Condition 2.[14] Using the observation that U_p and $1 - L_p$ are log-convex whenever they are less than one, one can also see that Condition 2 implies Condition 1.[15] We will use some of these conditions in parts of the following analysis. (In particular, Condition 3 is sufficient for all of the main results in the remainder of section 3.)

3.2 The suboptimality of asymmetric posteriors

Ideally we would like to characterize the set of feasible price pairs which could be implemented via information design, starting from a given prior. That is, we would like to calculate how the set of prices in Figure 3a is modified once we impose the MPC constraint. This, however, is in general a hard problem. Fortunately, Lemma 2 below shows that when we derive optimal policies we can restrict attention to symmetric policies which treat each firm the same.

Lemma 2 *Relative to an asymmetric posterior which induces a pure-strategy equilibrium, there is (i) a symmetric posterior which induces a pure-strategy equilibrium with greater profit for each firm and (ii) a symmetric posterior which induces a pure-strategy equilibrium with greater consumer surplus. (The improvement in either case is strict if the asymmetric posterior induces asymmetric prices.)*

The result concerning profit states that even if a firm earns more than its rival under an asymmetric posterior, it will nevertheless prefer some symmetric posterior. Since a symmetric posterior induces a symmetric equilibrium, both firms then are better off and so industry profit improves. As with the case in the previous section where there was no information constraint, firms have congruent interests when it comes to the design of consumer information. Intuitively, when a firm is treated unfavorably

[14] A primitive condition for the difference $x = v_1 - v_2$ to have a log-concave density is that (v_1, v_2) has a log-concave joint density (which includes the case where v_1 and v_2 are i.i.d. with a log-concave density).

[15] Notice that at $x = 0$, both U_{p_F} and F equal $\frac{1}{2}$ and have the same slope $f(0)$. Then, if F is log-concave on $[-1, 0]$, we have $F \leq U_{p_F}$ over that range given the upper bound is log-convex. Similarly, when $1 - F$ is log-concave on $[-1, 0]$, using the fact that $1 - L_p$ is log-convex whenever L_p is positive, we can also see that $F \geq L_{p_F}$ over that range. Therefore, Condition 2 implies that F lies between L_{p_F} and U_{p_F}.

under an asymmetric signal structure, it has an incentive to set a low price. Lemma 2 shows that this force is sufficiently strong so that even the favored firm suffers from the resulting fierce competition. This suggests regardless of whether a designer (e.g., a platform) wants to maximize industry profit or one firm's profit (e.g., the profit of the platform's own product), "biased recommendations" achieved through asymmetric signals may not be the best approach. In the example with biased recommendations in section 2.1, we saw that relative to *certain* symmetric policies, asymmetric signals which favor one firm over the other can improve industry profit and the favored firm's profit. Lemma 2 demonstrates that this possibility does not arise when general information policies are feasible.

The proof of the result for profit is relatively straightforward. Suppose an asymmetric posterior distribution G induces equilibrium prices $p_1 < p_2$, say, in which case Lemma 1 shows that respective profits are $\pi_1 < \pi_2$ and G lies above the lower bound L_{p_1,p_2}. Consider a symmetric posterior which equals $L_{2\pi_2}$ for $x < 0$. This posterior induces a symmetric price $2\pi_2$ and earns each firm a profit π_2, and so it benefits firm 1 while leaving firm 2 no worse off. Routine algebra (see the appendix) shows that $L_{2\pi_2}(x) \le L_{p_1,p_2}(x)$, and so

$$\int_{-1}^{x} L_{2\pi_2}(\tilde{x})d\tilde{x} \le \int_{-1}^{x} L_{p_1,p_2}(\tilde{x})d\tilde{x} \le \int_{-1}^{x} G(\tilde{x})d\tilde{x} \le \int_{-1}^{x} F(\tilde{x})d\tilde{x} \ , \tag{12}$$

where the final inequality follows from G being an MPC of F. Therefore, from (2) this new symmetric posterior is also an MPC of F, and so is a feasible policy which induces weakly higher profit for each firm.

For the consumer, biased information provision from asymmetric signals can lead to lower match efficiency, and in addition there is inefficiency caused by price dispersion with asymmetric signals, where the consumer sometimes buys her less preferred product because it is cheaper. These negative match effects outweigh the potential benefit from lower prices in an asymmetric market. More specifically, if an asymmetric posterior yields relatively high profit which exceeds its match quality improvement relative to random match, it is dominated for the consumer by the (symmetric) policy of disclosing no information which induces random matching and zero profit. If an asymmetric posterior yields relatively low profit, we can construct a new symmetric posterior which yields the same profit but higher total welfare and so improves consumer surplus. Specifically, suppose the asymmetric posterior induces equilibrium prices $p_1 \ge p_2$ and industry profit π. Then we show in the proof that $U_{p_1,p_2} \le U_\pi$, i.e., the upper bound associated with asymmetric prices lies below that associated with those symmetric prices which yield the same industry profit. This is the crucial step in constructing a symmetric posterior which performs better for the consumer.

3.3 Firm-optimal policy

Using Lemma 2, we can restrict attention to symmetric posteriors and symmetric prices. Since there is full consumer participation, maximizing industry profit corresponds to maximizing the symmetric equilibrium price p. We therefore wish to find the highest p such that there exists a G lying between L_p in (10) and U_p in (11) which is an MPC of the prior F. Notice that, whenever a symmetric G is a solution to this problem which induces price p, a symmetric posterior which takes the form L_p for negative x (and so takes the form U_p for positive x) must be a solution as well.[16] Therefore, we can restrict attention to symmetric posteriors which take the form L_p for $x < 0$.

We first illustrate our approach by considering the uniform-prior example with $F(x) = \frac{1}{2}(1+x)$. First of all, it is clear from Figure 4c that it is impossible to implement $p = 2$ once the MPC constraint is introduced. (The lower bound L_2 lies above F when $x < 0$ and so cannot be an MPC of the prior.) Given that L_p is concave whenever it is positive, it crosses the linear prior at most once and from below in the range of negative x. Then according to (3) a symmetric G which is equal to L_p for negative x is an MPC of the prior if and only if $\int_{-1}^{0} L_p(x)dx \le \int_{-1}^{0} F(x)dx = \frac{1}{4}$. Since L_p increases with p for negative x, the optimal price p^* must make this inequality bind, so that

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

and so

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

This optimal price is about 63% higher than the full-information price $p_F = 1$ in this example. The optimal symmetric posterior equals L_{p^*} for negative x, and the way we derive the optimal price also implies that this is the unique optimal solution.[17] Notice also that (3) holds with equality when $G = L_{p^*}$, and so the firm-optimal posterior leads to perfect matching and so maximizes total welfare. (The consumer does not possess full information, but she has enough information to buy her preferred product.)

This argument continues to hold provided that for each p, L_p crosses F at most once and from below in the range of negative x. This is true, for instance, when F is convex for negative x (i.e., the density f is weakly single-peaked at $x = 0$). More generally, we have the following result:

[16] Formally, we have $\int_{-1}^{x} L_p(\tilde{x})d\tilde{x} \le \int_{-1}^{x} G(\tilde{x})d\tilde{x} \le \int_{-1}^{x} F(\tilde{x})d\tilde{x}$ for $x \in [-1, 0]$, where the first inequality is because L_p is the lower bound of G and the second is because G is an MPC of F. Therefore, according to (2) a symmetric posterior which equals L_p for negative x must be an MPC of F.

[17] If there is another posterior which induces the same price p^* but differs from L_{p^*} on a positive measure of $x < 0$, then it must violate the MPC constraint.

Proposition 1 *Under Condition 2(i), the firm-optimal signal structure induces the symmetric equilibrium price*

$$p^* = \frac{2\delta}{1 - \log 2} \ , \tag{13}$$

where δ is defined in (5), and it is uniquely implemented by the symmetric posterior which is equal to L_{p^*} in the range of negative x, where L_p is defined in (10). With the firm-optimal signal structure there is no mismatch and total welfare is also maximized.

The firm-optimal posterior distribution depends on the prior only through the differentiation parameter δ. (A prior with a higher δ leads to a higher optimal price and so a flatter posterior.) Figure 5 illustrates the optimal posterior with a uniform prior (where the dashed lines are the prior density and CDF, respectively). In particular, the density is U-shaped to soften price competition, and consumers are pushed towards the two extremes insofar as this is feasible given the pure-strategy and MPC constraints.[18]

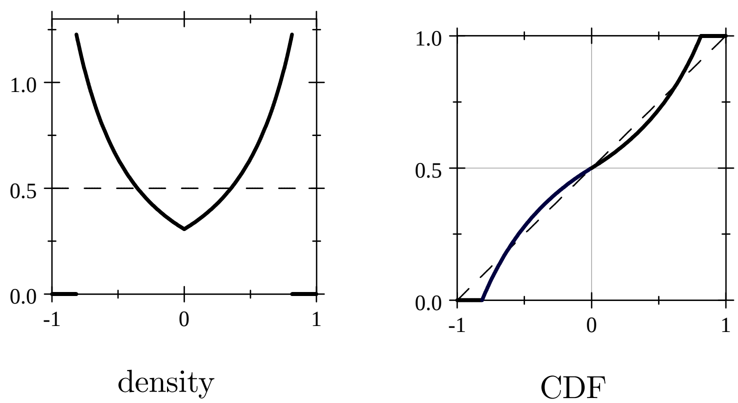

Figure 5: Firm-optimal posterior

The uniform example shows how the firm-optimal policy can significantly improve profit compared to the simple benchmark of full disclosure, and so giving more information to consumers does not always lead to greater market power for firms. More generally, with a log-concave prior density firms can achieve profit at least 63% higher than with full disclosure. The following result provides more details.

Corollary 1 *Under Condition 3, the firm-optimal price in (13) satisfies $\eta p_F \leq p^* \leq 2\eta p_F$ where $\eta = \frac{1}{2(1-\log 2)} \approx 1.63$.*

[18] With this posterior distribution, when a firm, say, firm 2 deviates to a price lower than p^*, its residual demand takes the form of the upper bound and so is unit-elastic (i.e., its profit is unchanged), since the upper bound was derived from firm 2's no deviation condition. While if it deviates to a price higher than p^*, its demand takes the form of the lower bound and so its profit strictly falls.

This result also shows that the advantage of considering general signal structures, relative to some frequently-used signal structures, can be significant. For instance, the often-used "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) induces a degree of mismatch so that consumers become more concentrated around $x = 0$, and so it cannot enhance profit relative to the full-information policy. The same is true more generally when the distribution for x is "rotated" about $x = 0$ as studied by Johnson and Myatt (2006). 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 with those more restricted signals.

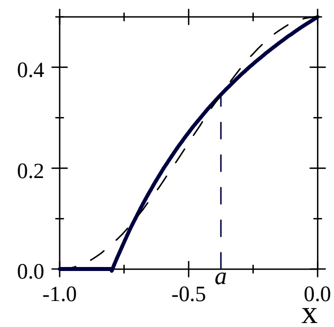

Figure 6: Firm-optimal posterior with a less regular prior

For priors outside the class covered by Proposition 1, the firm-optimal price is the highest p such that

$$\int_{-1}^{\tilde{x}} L_p(x)dx \le \int_{-1}^{\tilde{x}} F(x)dx \text{ for } \tilde{x} \in [-1, 0] \ .$$

In general, however, L_p and F can cross multiple times in the range of negative x, in which case it becomes somewhat harder to solve for the optimal price. Moreover, (3) might hold strictly and so there could be welfare loss associated with the firm-optimal signal structure. Figure 6 illustrates both points, where the prior shown as the dashed curve is initially convex and then concave in the range of negative x. 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, (3) holds strictly and there is some mismatch at the optimum.

Notice also that if the prior distribution is so dispersed that the MPC constraint never binds for any distribution which implements price $p \le 2$ (e.g., when the prior distribution is the binary distribution with support $\{-1, 1\}$), the bounds analysis in section 3.1 implies that the highest possible (pure-strategy) industry profit 2 is achievable.

3.4 Consumer-optimal policy

We turn next to the optimal information policy for the consumer. Unlike firms, the consumer does not care solely about the induced price but also about the reliability of the product match. When a posterior G induces a symmetric equilibrium with price p, the consumer always buys the product she prefers after seeing the signal, so from (4) her expected surplus is $\mu - p + \int_{-1}^{0} G(x)dx$. To maximize this, we first find the highest possible G to maximize match efficiency $\int_{-1}^{0} G(x)dx$ for a given price p, subject to the bounds condition $L_p \leq G \leq U_p$ and the MPC constraint, and then identify the optimal price.

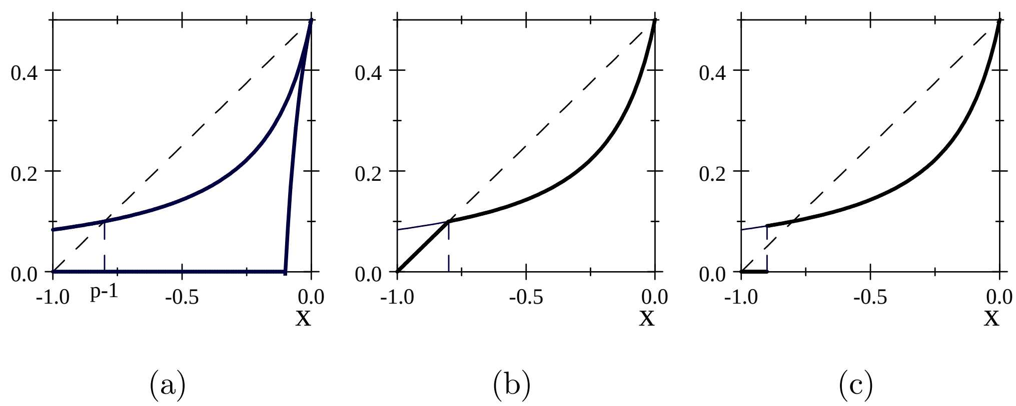

Figure 7: Consumer-optimal G for given price p

We illustrate our approach by considering the uniform example again. For the consumer to do better than with full information disclosure (when the equilibrium price is $p_F = 1$), the induced price must be below 1 to counteract any potential product mismatch. In this case the associated bounds are steep enough at $x = 0$ that they lie below the prior CDF for x close to zero. This is illustrated in Figure 7a, where the bold curves are the two bounds in the range of negative x. Since we want G to be as high as possible, it is now the upper bound U_p which will constrain G, rather than the lower bound which was relevant for the firm-optimal policy. Since the upper bound is convex and takes the form in (11), for any price $p < 1$ the upper bound cuts the prior CDF at the point $x_p \equiv p - 1 \in [-1, 0)$ and from above, again as illustrated in Figure 7a.

Therefore, for a given price $p < 1$, a consumer-optimal (symmetric) G is simply

$$G(x) = \min\{F(x), U_p(x)\} \tag{14}$$

for negative x, as shown as the bold curve in Figure 7b. First, this G is clearly between the two bounds associated with price p and is also an MPC of the prior.[19] Second, for $x \in [x_p, 0]$, this G already equals the upper bound; for $x < x_p$, this G equals F and so $\int_{-1}^{x_p} G(x)dx$ already reaches its maximum given the MPC constraint.

Notice, however, for $x < x_p$, the consumer-optimal G can take other forms provided that it is between the two bounds and satisfies $\int_{-1}^{x_p} G(x)dx = \int_{-1}^{x_p} F(x)dx$. Figure 7c depicts one alternative solution. It is also clear that for any consumer-optimal G, the strict inequality of (3) must hold because G is below F for $x \in [x_p, 0]$. Therefore, in contrast to the firm-optimal solution, there must be welfare losses at the consumer optimum and the consumer sometimes buys her less preferred product.

Expression (14) implies that the maximum consumer surplus for a given price $p < 1$ is

$$\mu + \int_{-1}^{0} \min\{F(x), U_p(x)\}dx - p \ . \tag{15}$$

One can check that the derivative of (15) with respect to p is $\frac{1}{2}(p - \log p - 3)$ when F is uniform, which decreases with p in the range $[0, 1]$. Therefore, the optimal price is $p^{**} \approx 0.05$ which is the root of $p - \log p = 3$ in the range $[0, 1]$, and a consumer-optimal symmetric posterior is given by expression (14) with $p = p^{**}$ for negative x.

The above argument based on a uniform prior continues to hold if for each price $p < p_F$, the upper bound U_p crosses F once and from above in the range of negative x. This is true, for example, when F is concave. More generally, we have the following result:

Proposition 2 *Under Conditions 1 and 2(ii), a consumer-optimal signal structure induces the symmetric equilibrium price*

$$p^{**} = \frac{-\gamma}{1-\gamma} F^{-1}(\tfrac{1}{2}\gamma) \ , \tag{16}$$

where $\gamma \approx 0.05$ *is the root of* $\gamma - \log \gamma = 3$ *in the range* $[0, 1]$*, and it can be implemented by a symmetric posterior which is equal to* $\min\{F(x), U_{p^{**}}(x)\}$ *in the range of negative* x*, where* U_p *is defined in (11). With the consumer-optimal signal structure there is mismatch and total welfare is not maximized.*

(Here Condition 1 is needed to ensure that the full-information benchmark has a pure-strategy equilibrium price p_F, so the argument that the consumer-optimal price must be below p_F continues to hold.)

[19] Note that at the full information price p_F we have $L_{p_F} \leq F$, and so $L_p < F$ for negative x when $p < p_F$ given the lower bound increases in price for negative x. Hence, for any price $p < p_F$, (14) must be above L_p.

Figure 8 depicts the consumer-optimal posterior distribution when the prior is uniform (and when the posterior takes the form in Figure 7b). The number of price-sensitive consumers near $x = 0$ is amplified compared with the prior distribution, and this forces firms to reduce their price in equilibrium. Those consumers near $x = 0$ do not have strong preferences about which product they buy, and so there is only limited 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.[20]

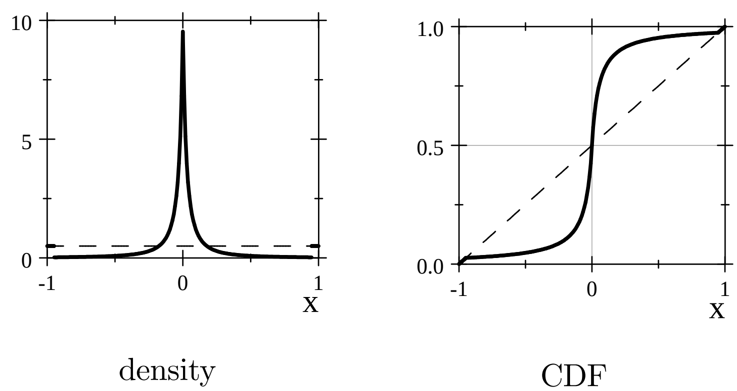

Figure 8: Consumer-optimal posterior

In this uniform example, the consumer-optimal policy induces a rather low price. This remains true more generally:

Corollary 2 *Under Condition 3, the consumer-optimal price in (16) satisfies* $\gamma p_F \leq p^{**} \leq \min\{\gamma, \frac{1}{2}p_F\}$ *where* $\gamma \approx 0.05$ *is defined in Proposition 2.*

Therefore, the insight from the no-disclosure example in section 2.1 is generally true that for the consumer the price effect appears more important than the match quality effect. When there are more firms, though, the match quality effect will become more important, and we will discuss this issue further in section 4.2.

For more general prior distributions, we can consider the family of posteriors illustrated by Figure 7c above. Given Lemma 2, it can be shown that the consumer-optimal policy can be implemented by a symmetric posterior which is equal to $G_p^m(x) = U_p(x)$ for $x \in [m, 0]$ and $= 0$ for $x < m$, where $m \in [-1, 0]$ is a constant. (This is in the spirit of Roesler and Szentes (2017).) The consumer problem is then to choose (p, m) to maximize $\mu + \int_{-1}^{0} G_p^m(x)dx - p$, subject to G_p^m being an MPC of F. This is a well-defined optimization problem, albeit hard to solve when the upper bound U_p crosses F

[20]In contrast to footnote 18, the consumer-optimal posterior distribution implies that when a firm unilaterally increases its price its residual demand is unit-elastic.

multiple times in the range of negative x. Nevertheless, this approach can be used to show that, unless the two products are perfect substitutes, disclosing no information cannot be optimal for the consumer. This implies that the consumer-optimal information policy is not the worst policy for firms, in contrast to the monopoly case in Roesler and Szentes (2017).

3.5 The limits to competition

Having discussed the signal structures which maximize profit and which maximize consumer surplus, we now describe the combinations of profit and consumer surplus which are feasible with some choice of signal structure. Let p^* be the firm-optimal price or industry profit. (This need not take the form (13) if the assumption in Proposition 1 does not hold.) We know that it is implemented by a symmetric posterior which equals L_{p^*} for negative x. Then any price or industry profit $p < p^*$ is feasible, since it can be implemented by a symmetric posterior which equals L_p for negative x. (This is because $L_p < L_{p^*}$ and so L_p is an MPC of the prior whenever L_{p^*} is.)

To derive all possible combinations of profit and consumer surplus, we first calculate the minimum possible consumer surplus for each industry profit $p \leq p^*$. As shown in the proof of Proposition 3 below, this minimum consumer surplus is achieved by a symmetric posterior. Then for a given $p \leq p^*$, the lowest value of consumer surplus is generated by the associated lower bound L_p. This is because consumer surplus is $\mu - p + \int_{-1}^{0} G(x)dx$ and so is lower with a smaller G, and the smallest possible G which induces symmetric price p is L_p (which is an MPC of F as explained above). The integral of $L_p(x)$ over $[-1, 0]$ is $\frac{1}{2}(1 - \log 2)p$ and so the minimum consumer surplus with profit $p \leq p^*$ is $\mu - \frac{1}{2}(1 + \log 2)p$.

We then calculate the maximum possible consumer surplus for each industry profit $p \leq p^*$. Lemma 2 shows that for the consumer any asymmetric posterior is dominated by some symmetric one (which may induce a different industry profit). What we need here is a stronger result that for the consumer any asymmetric posterior is dominated by a symmetric posterior which induces the *same* industry profit. This turns out to be true when the prior satisfies the assumptions used for Proposition 2. Once we rule out asymmetric posteriors, the analysis for the consumer-optimal problem immediately yields the maximum possible consumer surplus for a given price p.

Proposition 3 *Any profit between* 0 *and* p^* *is feasible. For a given profit* $p \in [0, p^*]$, *the lowest possible consumer surplus is* $\mu - \frac{1}{2}(1 + \log 2)p$*; if Conditions 1 and 2(ii) hold, the highest possible consumer surplus is* $\mu + \delta - p$ *if* $p \geq p_F$ *and is (15), which is concave in* p*, if* $p < p_F$*.*

For a given profit $p \leq p^*$, the minimum and the maximum possible consumer surplus are respectively generated by two symmetric posteriors which lie within the bounds L_p and U_p and which are an MPC of the prior. Therefore, any value of consumer surplus in between can be generated by a linear combination of these two symmetric posteriors. That is, Proposition 3 also determines the whole set of feasible profit and consumer surplus, and considering asymmetric signals does not change this set at all. The shaded area in Figure 9 illustrates the combinations of profit and consumer surplus in the uniform prior example.

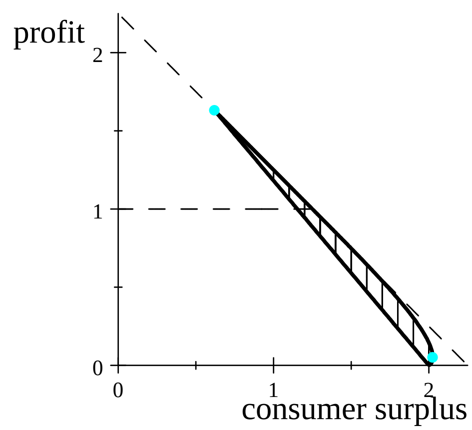

Figure 9: Feasible combinations of profit and consumer surplus

The two dots indicate the firm-optimal and the consumer-optimal outcome respectively. The higher dashed line is the efficient frontier with maximum total welfare $\mu + \delta$ (where $\delta = \frac{1}{4}$ and for convenience we have set $\mu = 2$). The higher bold curve shows the feasible outer frontier. Here, the full-information price is $p_F = 1$, and for prices above p_F the feasible frontier coincides with the efficient frontier. For lower prices, the feasible frontier lies strictly inside the efficient frontier and it is concave. Intuitively, to intensify price competition, information policy needs to induce mismatch which leads to more consumers who are near-indifferent between products. Figure 9 remains qualitatively the same for other prior distributions. In particular, due to competition it is impossible in our setup to have both low profit and low consumer surplus simultaneously, in contrast with the corresponding figure for monopoly in Roesler and Szentes (2017, Figure 1).

This discussion enables us to solve the Ramsey problem of maximizing a weighted sum of profit and consumer surplus. If the weight on profit is higher than that on consumer surplus, it is clear from Figure 9 that the solution simply coincides with the firm-optimal policy. While, as is more usual, if the weight on consumer surplus is higher the solution lies on the concave part of the outer frontier, where the optimal price is below p_F but above the consumer-optimal price p^{**} and is lower when the weight on consumer surplus is greater. This also implies that full information disclosure is optimal only in the knife-edge case when we wish to maximize (unweighted) total welfare.

4 Extensions and discussions

We have so far focused on the situation with two firms, where the outside option never binds for the consumer and where signal structures were required to induce a pure-strategy pricing equilibrium. In this section, we report the progress we can make when these restrictions are relaxed.

4.1 Allowing mixed pricing strategies

When signal structures induce mixed-strategy pricing equilibrium, the bounds approach in section 3 does not apply. We have not been able to find a way to solve the optimal information policy when mixed-strategy equilibria are allowed. Instead, 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, and show that this upper bound is close to the maximum consumer surplus available with pure strategies.

Consider the model introduced in section 2 with a zero outside option and (v_1, v_2) lying inside $[V, V+1]^2$. The following proposition reports the main result.

Proposition 4 *Suppose $V \geq 1$ and $v_1 - v_2$ has a log-concave density. Then the maximum consumer surplus available using only pure strategies attains at least 98.4% of the maximum consumer surplus available across all symmetric signal structures which induce a symmetric pure or mixed-strategy equilibrium.*

This result demonstrates that it is not possible that the consumer does significantly better if the class of (symmetric) signal structures is broadened to permit mixed pricing strategies in equilibrium. (Note, however, we have not found an example where the use of mixed strategies improves consumer surplus at all.) 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 the consumer.

Ideally one would also like to obtain a tight upper bound on the profit obtained using either pure or mixed strategies, and see how closely the optimal profit under the pure-strategy restriction approaches such an upper bound. This appears to be a harder problem, though, and we leave it for future work.

4.2 More than two firms

We now explore the general case with an arbitrary number of firms and an outside option (of zero) which might possibly bind in equilibrium. Consider $n \geq 2$ symmetric firms, each of which supplies one variety of a differentiated product. Let $\mathbf{v} = (v_1, \ldots, v_n) \in \mathbb{R}^+$ denote the consumer's valuations for the n varieties, and let $F(\mathbf{v})$ be the symmetric joint prior distribution. (We ignore the trivial case when the n varieties are identical.) Before her purchase, the consumer receives a private signal of $\mathbf{v}$ which is generated by a public signal structure $\{\sigma(\mathbf{s}|\mathbf{v}), \mathbf{S}\}$. Let $G(\mathbf{v})$ be the joint posterior distribution of the expected valuations after receiving the signal. As before, G is an MPC of F (in the sense that they have the same mean and $\int \phi dG \leq \int \phi dF$ for any convex function ϕ whenever the integrals exist), and any G which is an MPC of F can be generated by some signal structure. (See for instance Blackwell (1953) and Strassen (1965).) It would be interesting to know how the number of firms affects optimal information design, but such analysis would require consideration of multi-dimensional consumer heterogeneity even in situations where the outside option could be ignored. As discussed in the introduction, current understanding of information design in multi-dimensional environments is limited.

Here instead we compare the performance of three simple signal structures. First, there is the "full disclosure" policy where consumers are accurately informed of their preferences $\mathbf{v} = (v_1, \ldots, v_n)$. Second, the "top product" signal structure informs the consumer which product i is her best match but nothing else. Third, the "top two" signal structure informs the consumer which two products i and j are her best matches but nothing else. (With duopoly this corresponds to a completely uninformative signal.) We will argue that out of these three policies, firms tend to most prefer the "top product" policy, and under certain conditions this policy even achieves the first-best outcome for firms. The consumer, by contrast, tends to prefer the "top two" policy, and in the limit with many firms both this policy and full disclosure achieve approximately the first-best outcome for the consumer.

Recall that $\mu \equiv \mathbb{E}[v_i]$, and let $\mu_{n:n} \equiv \mathbb{E}[\max\{v_i\}]$ denote the expected valuation of the top product while $\mu_{<n:n}$ denotes the expected valuation of a product which is *not* the top product. These expectations are related as

$$\mu = \tfrac{1}{n}\mu_{n:n} + (1 - \tfrac{1}{n})\mu_{<n:n} \ . \tag{17}$$

Note that $\mu_{n:n}$ is the maximum total welfare in the market. We say that first-best industry profit is achieved if the maximum total welfare is achieved and is fully extracted by firms.[21]

[21] Notice that when first-best industry profit is achieved, the consumer has zero surplus in which

With a general joint valuation distribution $F(\mathbf{v})$, analysis of price competition under the policy of full disclosure can be complicated. The best understood scenario is the "i.i.d." case studied in Perloff and Salop (1985). They assume that valuations v_i are independently and identically distributed across the n products and have a log-concave density function, and that the outside option is irrelevant. In this i.i.d. case, the price competition has a unique symmetric pure-strategy pricing equilibrium. (The details of the equilibrium price and the condition for irrelevant outside option in this case are reported in the proof of Proposition 5 below.) In the following result, when we consider the full-disclosure policy we focus on this i.i.d. case.

Proposition 5 *(i) The "top product" signal structure achieves first-best industry profit* $\mu_{n:n}$ *if*

$$\mu_{n:n} \geq n\mu_{<n:n} \ . \tag{18}$$

More generally, this signal structure generates equilibrium industry profit at least

$$\mu_{n:n} - \mu_{<n:n} \ , \tag{19}$$

and performs better for firms than the "top two" signal structure, and also better than full disclosure in the i.i.d. case.
(ii) The "top two" signal structure performs better for the consumer than the "top product" signal structure, and also better than full disclosure in the i.i.d. case.

The reason that the "top product" signal achieves first-best profit when condition (18) holds is straightforward.[22] If its rivals choose the first-best price $p = \mu_{n:n}$, a firm obtains profit $\frac{1}{n}\mu_{n:n}$ when it chooses the same price, as it then sells to that portion of consumers who prefer its product to any other. All other consumers value its product at $\mu_{<n:n}$, and the firm could attract all such consumers if it reduced its price to $\mu_{<n:n}$ (or slightly below this). Condition (18) ensures that this is not a profitable deviation, in which case all firms choosing $p = \mu_{n:n}$ constitutes an equilibrium. For instance, when $n = 2$ and (v_1, v_2) is uniformly distributed on the square $[0, 1]^2$, then $\mu_{<n:n} = \frac{1}{3}$ and $\mu_{n:n} = \frac{2}{3}$ and (18) is (just) satisfied. However, condition (18) tends to be less likely to

case the outside option must bind. Therefore, the approach used in section 3, which is based on the assumption that the outside option is irrelevant, cannot be used to find instances where first-best profit is achieved.

[22] If instead of oligopoly the n products were jointly supplied by a multiproduct monopolist, the top-product signal with associated prices $p = \mu_{n:n}$ would allow the firm to fully extract surplus regardless of whether (18) holds. See Ichihashi (2020) for a related observation.

hold when there are more firms in the market. For example, when v_i has a finite upper bound, it must fail when n is sufficiently large.[23]

When (18) fails to hold, the price competition under the "top product" signal has no pure-strategy equilibrium. In any mixed-strategy equilibrium, industry profit is at least (19).[24] To see that, suppose a firm chooses price equal to (19), or slightly less. Then when the consumer is informed that this firm supplies her top product, she will buy from that firm, even if all other firms choose price zero. (This is because she values each rival product at $\mu_{<n:n}$, and so she prefers the top product at this positive price to any other product even when offered for free. In addition, since (19) is below $\mu_{n:n}$, the consumer prefers to buy this product than to buy nothing.) Therefore, by choosing this price the firm is ensured profit of at least $\frac{1}{n}(\mu_{n:n} - \mu_{<n:n})$, which is therefore a lower bound on each firm's profit and which proves the claim. This lower bound can be substantial even when there are many firms. For example, when each v_i is i.i.d. uniform on $[0, 1]$, the lower bound (19) is $\frac{n}{2(n+1)}$, which is half of the maximum welfare $\mu_{n:n}$ for any n. This contrasts with the literature discussed in the introduction in which firms can disclose information only about their own product, where with many firms firms disclose all information and industry profit converges to zero.

The "top two" signal induces perfect competition between firms, and each firm chooses price $p = 0$ in equilibrium and industry profit is zero. (If its rivals choose $p = 0$ then a firm has no demand if it chooses a positive price, since even when the consumer has the firm's product in her top two, there is another firm in her top two which she regards as a perfect substitute and which has a lower price. More generally, this Bertrand argument shows there can be no equilibrium in which a firm chooses a positive price with positive probability.) Hence, the "top two" policy must be worse for firms than the "top product" policy, as reported in part (i). When (18) does not hold, the profit comparison with the policy of full disclosure is more delicate. Intuitively, though, the "top product" policy eliminates the possibility that the consumer regards the best and the second best product as close substitutes, and so it is more able to relax price competition than full disclosure.

The consumer faces a trade off between price and match quality. For any finite n, it is impossible for the consumer to achieve the first-best outcome. This is simply because zero pricing requires the consumer regard all the products as identical but in that case perfect match cannot be guaranteed. The "top two" signal structure, as explained

[23] When v_i is i.i.d. and has a uniform distribution on $[0, 1]$, (18) fails with more than two firms. However, for a distribution with a heavy tail it can hold for a large n. For example, if v_i is i.i.d. from the Pareto distribution $F(x) = 1 - x^{-\alpha}$ with $\alpha > 1$, (18) holds for $n \leq 38$ if $\alpha = 1.2$.

[24] In our working paper, Armstrong and Zhou (2021), we constructed an equilibrium in mixed strategies, which is quite complex, when condition (18) fails to hold.

before, induces zero pricing, and meanwhile it enables the consumer to buy one of her top two products. So intuitively it should perform well for the consumer, especially when the number of firms is large. The result that the "top two" signal works better for the consumer than the "top product" signal is easy to see. Under the former signal, the consumer always participates as the price is zero, and so consumer surplus is

$$\tfrac{1}{2}(\mu_{n:n} + \mu_{n-1:n}) \ , \tag{20}$$

where $\mu_{n-1:n}$ is the expected valuation for the second best product; using the lower bound (19), it follows that consumer surplus with the latter signal is at most $\mu_{<n:n}$. Since both $\mu_{n:n}$ and $\mu_{n-1:n}$ are greater than $\mu_{<n:n}$ (with the former strictly greater), it follows from (20) that consumers are better off with the "top two" signal than with the "top product" signal, as reported in part (ii) of the result.

The comparison for the consumer between the "top two" signal and full disclosure in the i.i.d. case is less obvious. When there are many firms, under either information policy the trade off between price and match quality vanishes, and both of them achieve approximately the first-best outcome for the consumer (under a mild tail condition of the valuation distribution as specified in the proof). Under the "top two" signal, the price is always zero and the mismatch from buying one of the top two products vanishes as $n \to \infty$; under full disclosure, the consumer always buys the best matched product and the price goes to zero as $n \to \infty$. However, for any finite n, it turns out the price effect dominates the match quality effect, and the "top two" signal performs strictly better for the consumer as reported in part (ii). This simple policy may have additional advantages over full disclosure if the need to choose among many products involves high information processing costs for the consumer.

5 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. As with many other optimal policies in economics, the optimal information policies derived in this paper might be "too complicated" to implement in practice. Nevertheless, these optimal policies provide useful guidance for which kinds of informations help boost profits and which help 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 welfare too. In particular, the "top product" policy which only informs consumers of their best match can sometimes be optimal for firms. The consumer-optimal policy, 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 achieved by inducing a degree of product mismatch, however, and so the policy does not maximize total welfare. The "top two" policy which informs consumers of their best two products (but without ranking them) can be close to be optimal for consumers.

Besides the extensions discussed in section 4, there are other interesting extensions to explore. One would be to consider situations where firms were asymmetric *ex ante*, including the case of vertical differentiation where one firm was known to provide a higher match utility than its rival. One way to model this asymmetry is to consider an asymmetric prior distribution F, and then one could investigate whether the optimal information policy maintains, amplifies or reduces this prior asymmetry, and whether firms continue to have aligned interests over the design of consumer information. The bounds analysis, which is independent of the information problem, still applies in this more general problem, but the principal new challenge stems from dealing with the MPC constraint when the prior is asymmetric.

Another possible extension would be to allow consumers to be heterogeneous *ex ante*. For instance, a consumer's valuation v_i for product i might be decomposed as $v_i = a_i + b_i$, where consumers know the vector $(a_1, ...)$ from the start, from other information sources, and there is scope to manipulate information only about the vector $(b_1, ...)$. If there was enough heterogeneity in $(a_1, ...)$, then one might be able to rule out mixed pricing strategies in equilibrium, rather than assuming them away as we mostly did in this paper.

Finally, it would be valuable to embed this analysis within a framework in which the "information designer" is modelled explicitly as an economic agent. Platforms typically compete with each other to provide intermediation services. If a profit-maximizing platform chooses what product information to reveal to consumers, and also chooses its fees to each side of the market, then the relative competitive intensity among platforms on the two sides of the market and the platform's equilibrium fee structure will presumably affect whether its information policy is focussed more on delivering firm profit or consumer surplus.

Technical Appendix

A condition for the outside option to be irrelevant. Recall that in our model the consumer's valuations for the two varieties are (v_1, v_2) and they lie inside $[V, V+1]^2$. This remains the same for her expected valuations for the two varieties after she observes a signal of her true preferences. If we focus on pure-strategy pricing equilibria, the following lemma reports a condition for the zero outside option to be irrelevant.

Lemma 0 *If $V > 2$, then for any distribution of valuations which induces a pure-strategy pricing equilibrium, equilibrium prices are no greater than* 2 *(so below V) and the consumer obtains positive surplus when she buys from either firm.*

Proof: Suppose a distribution of valuations induces a pure-strategy equilibrium where firm i, $i = 1, 2$, offers price p_i and sells with probability q_i. Note first that in equilibrium neither price can exceed the highest possible valuation $V + 1$.[25] Given firm i's price p_i, if firm $j \neq i$ deviates and sets a low price $p = p_i - 1$ (or slightly below it), it will serve the whole market as the maximum possible valuation difference between the two varieties is 1. (This deviation price is such that $V - p = V + 1 - p_i$. Given $p_i \leq V + 1$, this ensures that $V - p \geq 0$ so that all consumers prefer to buy from firm j than to buy nothing.) Therefore, in equilibrium we must have

$$p_i - 1 \leq p_j q_j \tag{21}$$

where $p_j q_j$ is firm j's equilibrium profit.

Using these two no-deviation conditions we show that in equilibrium it is impossible for any price to exceed 2. Suppose in contrast that $p_1 > 2$, say. Then since $p_1 - 1 \leq p_2 q_2$ and $q_2 \leq 1$ it follows that $p_2 > 1$. On the other hand, (21) implies that $p_1 + p_2 - 2 \leq p_1 q_1 + p_2 q_2$ and so

$$p_1(1 - q_1) + p_2(1 - q_2) \leq 2 \ .$$

If $p_2 \geq 2$, the left-hand side would be strictly greater than $2(2 - q_1 - q_2) \geq 2$ given $q_1 + q_2 \leq 1$, which would be a contradiction. Hence, if $p_1 > 2$ in equilibrium we have $1 < p_2 < 2$.

Since both prices are above 1, (21) implies $q_1, q_2 > 0$. Given $V > 2 > p_2$, the consumer always buys a product (i.e., $q_1 + q_2 = 1$) and obtains *strictly* positive surplus, i.e., the outside option never binds. This remains true if either firm unilaterally changes its price slightly. Let $H(x)$ be the CDF of $x \equiv v_1 - v_2$, and let $\pi_i = p_i q_i$ denote firm i's equilibrium profit. If firm 2 deviates to a price p_2' close to p_2, the consumer will buy its product if $v_2 - p_2' \geq v_1 - p_1$ or $x \leq p_1 - p_2'$. So a necessary condition for the equilibrium is

$$p_2' H(p_1 - p_2') \leq p_2 q_2 \tag{22}$$

[25] If $p_1 > V + 1$, say, then no consumer wishes to buy product 1 and firm 2 can operate as a monopolist, in which case it sets a price $p_2 \geq V$. But if firm 1 sets price $V - 1$ it serves the whole market. Since $V > 1$, this generates positive profit for firm 1, which is a contradiction.

for all p_2' close to p_2. (Here we suppose that if H has a mass point at $x = p_1 - p_2'$, 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_2'.) By changing variables from p_2' to $x = p_1 - p_2'$, (22) requires $H(x) \le \frac{p_2 q_2}{p_1 - x}$ for x sufficiently close to $p_1 - p_2$.

Similarly, a necessary condition for firm 1 to have no incentive to deviate is

$$p_1'[1 - H((p_1' - p_2)^-)] \le p_1 q_1$$

for all p_1' close to p_1, where $H(x^-) \equiv \lim_{z \nearrow x} H(z)$. (As with firm 2, if H has an atom at $x = p_1' - p_2$ the natural tie-breaking assumption is that firm 1 serves all consumers at x. The deviation demand is written as $1 - H((p_1' - p_2)^-)$ as a CDF is defined to be right-continuous.) Since $H(x) \ge H(x^-)$, this requires $H(x) \ge H(x^-) \ge 1 - \frac{p_1 q_1}{p_2 + x}$ for x sufficiently close to $p_1 - p_2$.

Notice that (a) the upper bound function $\frac{p_2 q_2}{p_1 - x}$ is differentiable and convex in x while the lower bound function $1 - \frac{p_1 q_1}{p_2 + x}$ is differentiable and concave in x, and (b) they coincide and are equal to q_2 at $x = p_1 - p_2$. Therefore, the above two necessary conditions for the equilibrium will be satisfied only if these two bounds are tangent to each other at $x = p_1 - p_2$. After differentiating the bounds, this implies $\frac{q_1}{p_1} = \frac{q_2}{p_2}$. Since $p_1 > 2 > p_2$, we must then have $q_1 > q_2$ and so $q_2 < \frac{1}{2}$. This, together with (21), implies that

$$p_1 - 1 \le p_2 q_2 < \tfrac{1}{2} p_2 < \tfrac{1}{2} p_1 \ ,$$

which implies $p_1 < 2$, which is a contradiction. Therefore, when $V > 2$, in any pure-strategy pricing equilibrium neither price is above 2. ■

The outside option is then irrelevant in the sense that if one firm deviates from equilibrium the consumer continues to buy one product and the market is fully covered. The condition of $V > 2$ essentially requires that valuations have a concentrated distribution in the sense that the range of valuations is small relative to the basic utility of the product. As we will show in Lemma 1 below, when the outside option is irrelevant, $p_1 = p_2 = 2$ can be equilibrium prices for some valuation distributions, and so the requirement $V > 2$ is tight.

Proof of Lemma 1. It is more convenient to prove this lemma by dealing with two cases separately:

Case 1: both prices are positive. (i) Let π_i, $i = 1, 2$, denote firm i's equilibrium profit when (p_1, p_2) are equilibrium prices. (We will show that π_i is given by (8) but for now these profits are undetermined.) As shown in the proof of Lemma 0, in any pure-strategy equilibrium we must have $p_i - 1 \leq \pi_j$ for $i \neq j$ to ensure that no firm has an incentive to deviate to a low price to capture all demand. We can therefore focus on prices and profits which satisfy these conditions.

We first show the "if and only if" condition on G for (p_1, p_2) to be the associated equilibrium prices. Prices (p_1, p_2) constitute an equilibrium with G if and only if neither firm has an incentive to deviate from their price. If firm 2 deviates to price $p_2' \neq p_2$, the consumer buys from firm 2 if $x \leq p_1 - p_2'$. Therefore, firm 2 has no incentive to deviate if and only if

$$p_2' G(p_1 - p_2') \leq \pi_2$$

holds for all p_2'. (The tie-breaking rule if G has a mass point at $x = p_1 - p_2'$ was explained in the proof of Lemma 0.) By changing variables from p_2' to $x = p_1 - p_2'$, we can write this no-deviation requirement as $(p_1 - x)G(x) \leq \pi_2$. This requires $G(x) \leq \frac{\pi_2}{p_1 - x}$ if $p_1 - x > \pi_2$ (or $x < p_1 - \pi_2$); otherwise any $G(x) \in [0, 1]$ works. Therefore, firm 2's no-deviation condition is equivalent to $G(x) \leq U_{p_1,p_2}(x)$ for all $x \in [-1, 1]$.

Likewise, firm 1 has no incentive to deviate if and only if

$$p_1'(1 - G((p_1' - p_2)^-)) \leq \pi_1$$

for all p'. (Again, the tie-breaking rule when G has a mass point at $x = p_1' - p_2$ and why the deviation demand is written as $1 - G((p_1' - p_2)^-)$ have been explained in the proof of Lemma 0.) By changing variables from p_1' to $x = p_1' - p_2$, this constraint can be written as $(p_2 + x)(1 - G(x^-)) \leq \pi_1$. This requires $G(x^-) \geq 1 - \frac{\pi_1}{p_2 + x}$ if $p_2 + x > \pi_1$ (or $x > \pi_1 - p_2$ which is no less than -1 given $p_2 - 1 \leq \pi_1$); otherwise any $G(x^-) \in [0, 1]$ works. Given the first part of the lower bound is continuous, we can replace $G(x^-)$ by $G(x)$, and so firm 1's no-deviation condition is actually equivalent to $G(x) \geq L_{p_1,p_2}(x)$ for any $x \in [-1, 1]$.

It remains to show that $U_{p_1,p_2} \geq L_{p_1,p_2}$ for any $x \in [-1, 1]$ (so that a G between the two bounds exists) if and only if π_i is defined as in (8). Let $q_i \equiv \pi_i/p_i$ denote firm i's market share, where $q_1 + q_2 = 1$ since there is full consumer participation. Notice that L_{p_1,p_2} is increasing and concave whenever it is positive, and U_{p_1,p_2} is increasing and convex whenever it is below 1. The two bounds coincide and equal firm 2's market share q_2 at $x = p_1 - p_2$. These properties of the two bounds imply that $U_{p_1,p_2} \geq L_{p_1,p_2}$ for any $x \in [-1, 1]$ if and only if the two bounds have the same slope at $x = p_1 - p_2$, which requires $\pi_2/p_2^2 = \pi_1/p_1^2$. Together with $\pi_i = p_i q_i$ and $q_1 + q_2 = 1$, this then implies $q_i = p_i/(p_1 + p_2)$ and $\pi_i = p_i^2/(p_1 + p_2)$ as in (8).

We now prove the symmetric-price result in (i). Notice that $U_{p_1,p_2}(0) = \frac{\pi_2}{p_1}$ and $L_{p_1,p_2}(0) = 1 - \frac{\pi_1}{p_2}$. The "only if" part is straightforward. If $p_1 = p_2$, then (8) implies that $U_{p_1,p_2}(0) = L_{p_1,p_2}(0) = \frac{1}{2}$, and so $G(0) = \frac{1}{2}$. To prove the "if" part, suppose

$G(0) = \frac{1}{2}$. If (p_1, p_2) are the equilibrium prices under G, the bounds condition (9) at $x = 0$ requires

$$1 - \frac{\pi_1}{p_2} \leq \frac{1}{2} \leq \frac{\pi_2}{p_1} ,$$

from which we have $\frac{p_2}{p_1} \leq 2q_1$ and $\frac{p_1}{p_2} \leq 2q_2$, where q_i is firm i's market share defined before. This, together with $q_1 + q_2 = 1$, implies $\frac{p_1}{p_2} + \frac{p_2}{p_1} \leq 2$. This can hold only if $p_1 = p_2$.

(ii) If (p_1, p_2) are the equilibrium prices under some distribution G, we must have $p_j - 1 \leq \pi_i$ as pointed out before, and part (i) implies that firm i's equilibrium profit π_i must take the form in (8). To prove the "if" part, notice that when $p_j - 1 \leq \pi_i$, we have $U_{p_1,p_2}(1) = 1$ and $L_{p_1,p_2}(-1) = 0$. Meanwhile, as shown in the proof of part (i) we have $U_{p_1,p_2} \geq L_{p_1,p_2}$ for any x when π_i takes the form in (8). Hence, there exists a distribution G which is between the two bounds. Then the desired result follows from part (i). Finally, from $p_i - 1 \leq \pi_j \leq p_j$, it is clear that the price difference cannot exceed 1. The result $p_i \leq 2$ has been proved in Lemma 0.

Case 2: at least one price is zero. (i) The argument in case 1 which derives the two bounds by using the two firms' no-deviation conditions does not rely on whether any price is zero. The only difference here is how to find the condition for the upper bound to be nowhere below the lower bound. Consider first the case of $p_1 > p_2 = 0$. (The case of $p_2 > p_1 = 0$ is similar.) In this case, firm 2 must make a zero profit (i.e., $\pi_2 = 0$). Then the lower bound is $1 - \frac{\pi_1}{x}$ if $x > \pi_1$ and 0 otherwise; the upper bound is 0 if $x < p_1$ and 1 otherwise. (The upper bound is now a step function. This is why the "tangency" argument in case 1 does not work here.) Given $\pi_1 \leq p_1$, one can see that the upper bound is nowhere below the lower bound if and only if $\pi_1 = p_1$. In the case of $p_1 = p_2 = 0$, both firms must earn a zero profit. Then both bounds from the no-deviation conditions become a step distribution function with a discontinuous point at $x = 0$. (The lower bound is left-continuous while the upper bound is right-continuous at $x = 0$.) Then the only CDF permitted by the two bounds is G which is degenerate at $x \equiv 0$. (Intuitively, $p_1 = p_2 = 0$ is an equilibrium if and only if consumers regard the two products as perfect substitutes. If there were any consumers with $x \neq 0$, a firm could profitably serve them even if its rival sets a zero price.)

(ii) If $p_1 = p_2 = 0$, they satisfy $p_j - 1 \leq \pi_i$ and is supported as equilibrium prices by a degenerate distribution at $x \equiv 0$. If $p_1 > p_2 = 0$, then $p_1 - 1 \leq \pi_2 = 0$ requires $p_1 \leq 1$. In this case, the two prices can be sustained in equilibrium by, for example, a degenerate distribution at $x \equiv p_1$. (The case of $p_2 > p_1 = 0$ is similar.) ■

Proof of Lemma 2. (i) First consider the firm problem. Suppose that an asymmetric posterior G induces an equilibrium with prices (p_1, p_2). Without loss of generality let

$p_1 \le p_2$, and so $\pi_1 \le \pi_2$ according to Lemma 1. (Figure 2a illustrates the two bounds on G in this case.) Now consider a symmetric CDF $\hat{G}(x)$ which equals $L_{2\pi_2}(x)$ for negative x, where $L_{2\pi_2}$ defined in (10) is the lower bound of the posteriors which induce a symmetric equilibrium with price $p = 2\pi_2$. We show below that $\hat{G}$ is an MPC of the prior, and so it is a legitimate posterior and induces a symmetric equilibrium in which each firm earns π_2. That is, $\hat{G}$ weakly outperforms G for each firm.

To prove $\hat{G}$ is an MPC of the prior it suffices to show $L_{2\pi_2} \le L_{p_1,p_2}$, in which case (12) in the main text applies. Expression $L_{2\pi_2}$ is positive and equal to $1 - \frac{\pi_2}{2\pi_2+x}$ when $x > -\pi_2$; $L_{p_1,p_2}(x)$ is positive and equal to $1 - \frac{\pi_1}{p_2+x}$ when $x > \pi_1 - p_2$. Given $\pi_1 + \pi_2 \le p_2$, the latter is positive over a wider range of x, and at the same time for $x > -\pi_2$ we have

$$\frac{\pi_1}{p_2 + x} \le \frac{\pi_1}{\pi_1 + \pi_2 + x} \le \frac{\pi_2}{2\pi_2 + x} \ ,$$

where the second inequality follows from $\pi_1 \le \pi_2$. Therefore, we have $L_{2\pi_2} \le L_{p_1,p_2}$. (With distinct prices $p_1 < p_2$, then $\pi_1 < \pi_2$ and $\pi_1 + \pi_2 < p_2$, and so $L_{2\pi_2} < L_{p_1,p_2}(x)$ whenever the latter is positive. The same argument then works for a symmetric price p which is slightly above $2\pi_2$. As a result, $\hat{G}$ strictly outperforms G for each firm.)

(ii) Next consider the consumer problem. Suppose that an asymmetric posterior G induces an equilibrium with prices (p_1, p_2) and industry profit π. It is now more convenient to label firms as $p_1 \ge p_2$, in which case we have $\pi_2 \le \frac{1}{2}\pi \le \pi_1$ and $\pi \le p_1$. (Figure 2c illustrates the two bounds on G in this case.) Total welfare in this equilibrium can be no greater than (4) since if the firms offer distinct prices there will be some mismatch, and consumers sometimes buy the wrong product.

If in (4) we have $\int_{-1}^{0} G(x)dx < \pi$ then consumer surplus under G is less than μ, and so G is dominated for consumers by the (symmetric) policy of disclosing no information and the claim holds. Assume instead that

$$\int_{-1}^{0} G(x)dx \ge \pi \ . \tag{23}$$

We show that there exists a feasible symmetric posterior which induces a symmetric equilibrium with the same industry profit π but weakly higher total welfare. (As with part (i), with distinct prices $p_1 > p_2$ the following argument implies that a strict improvement exists.)

Consider a symmetric CDF

$$\hat{G}(x) = \begin{cases} 0 & \text{if } x < m \\ U_\pi(x) & \text{if } m \le x \le 0 \end{cases} \ ,$$

where $U_\pi(x)$ defined in (11) is the upper bound of posteriors which induce a symmetric equilibrium price $p = \pi$, and $m \in (-1, 0)$ is chosen so that

$$\int_{-1}^{0} G(x)dx = \int_{-1}^{0} \hat{G}(x)dx \ . \tag{24}$$

To see that such an m exists, note that given $\pi_2 \leq \frac{1}{2}\pi$ and $p_1 \geq \pi$ we have

$$G(x) \leq U_{p_1,p_2}(x) = \frac{\pi_2}{p_1 - x} \leq \frac{\pi/2}{\pi - x} = U_\pi(x) \tag{25}$$

for $x < 0$, since in this range these upper bounds are less than 1. Since $U_\pi(x) \leq \frac{1}{2}$ for $x < 0$ we have

$$\int_{-\pi/2}^{0} U_\pi(x)dx \leq \frac{\pi}{4} < \int_{-1}^{0} G(x)dx \ ,$$

where the second inequality is implied by (23). These two observations imply that (24) has a unique solution $m \in (-1, -\frac{\pi}{2})$, and $\hat{G}$ crosses G at most once and from below. The latter observation and (24) imply that

$$\int_{-1}^{\tilde{x}} \hat{G}(x)dx \leq \int_{-1}^{\tilde{x}} G(x)dx \leq \int_{-1}^{\tilde{x}} F(x)dx$$

for $\tilde{x} < 0$, where the second inequality follows because G is an MPC of F. Therefore, the symmetric $\hat{G}$ is an MPC of F and so is a feasible posterior.

Since $m < -\frac{\pi}{2}$ and $-\frac{\pi}{2}$ is the x at which the lower bound $L_\pi(x)$ becomes positive, $\hat{G}$ lies between L_π and U_π when $x < 0$. Hence, $\hat{G}$ induces a symmetric equilibrium with industry profit π. It also yields total welfare $\mu + \int_{-1}^{0} \hat{G}(x)dx$, which is equal to (4) given expression (24). Therefore, the symmetric posterior $\hat{G}$ induces at least as much consumer surplus as G. ■

Proof of Proposition 1. Notice that whenever L_p is positive, $1 - L_p = \frac{p}{2(p+x)}$ is log-convex. Therefore, when $1 - F$ is log-concave in the range of negative x, the single-crossing property as in the uniform example continues to hold. Then the same argument as in that example implies that the firm-optimal symmetric price must solve

$$\int_{-1}^{0} L_p(x)dx = \int_{-1}^{0} F(x)dx = \delta \ ,$$

from which we derive p^* in (13). This equality immediately implies that the symmetric posterior G which equals L_{p^*} for negative x is the unique optimal posterior, and it also leads to no mismatch and so maximizes total surplus. ■

Proof of Corollary 1. When f is log-concave, the full-information price is well defined and is equal to $p_F = 1/(2f(0))$. When f is log-concave, it must be weakly increasing and so F must be weakly convex in the range of negative x. Then we have

$$\delta = \int_{-1}^{0} F(x)dx \geq \int_{\frac{-1}{2f(0)}}^{0} (\tfrac{1}{2} + xf(0))dx = \frac{1}{8f(0)} = \tfrac{1}{4}p_F \ , \tag{26}$$

where the inequality follows since F lies above its tangent at $x = 0$. Using this result and (13), we have $p^* \geq \frac{1}{2(1-\log 2)} p_F$.

On the other hand, when f is log-concave, F is log-concave and so $\frac{F(x)}{f(x)}$ is increasing. Then we have

$$\delta = \int_{-1}^{0} F(x)dx = \int_{-1}^{0} \frac{F(x)}{f(x)} dF(x) \leq \frac{F(0)^2}{f(0)} = \tfrac{1}{2}p_F \ . \tag{27}$$

Using this result and (13), we have $p^* \leq \frac{1}{1-\log 2} p_F$. ■

Proof of Proposition 2. Following the same argument as in the main text, we know that the consumer-optimal symmetric price must be below the full-information price p_F. Notice that $U_p = \frac{p}{2(p-x)}$ is log-convex when it is less than one. Hence, when prior CDF F is log-concave in the range $[-1, 0]$, the upper bound U_p crosses F once and from above in that range for any $p < p_F$. Then the same argument used for the uniform example shows that an optimal G given $p < p_F$ is given by (14). This immediately implies that under the consumer-optimal solution there is mismatch and total welfare is not maximized.

With G in (14), consumer surplus is (15) and its derivative with respect to p is

$$\int_{x_p}^{0} \frac{\partial U_p(x)}{\partial p} dx - 1 = \frac{1}{2}\left(\frac{p}{p - x_p} - \log \frac{p}{p - x_p} - 3\right) \ , \tag{28}$$

where x_p is the intercept point of $F(x)$ and $U_p(x)$. Since $F(x_p) \equiv U_p(x_p)$, it follows that $\frac{p}{p-x_p} = 2F(x_p)$, and so (28) equals $\frac{1}{2}(2F(x_p) - \log(2F(x_p)) - 3)$. This decreases in p since the intercept point x_p increases with p (which is because the upper bound crosses F from above). Therefore, the optimal intercept point x^* satisfies $2F(x^*) = \gamma$, 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 $p^{**} = \frac{-\gamma}{1-\gamma} x^*$ and so (16). ■

Proof of Corollary 2. When f is log-concave, F is convex on $[-1, 0]$, and so it is below the linear line $\frac{1}{2}(1 + x)$. This implies $F^{-1}(\frac{1}{2}\gamma) \geq \gamma - 1$. Using this result and (16), we have $p^{**} \leq \gamma$. Also, p^{**} should never exceed δ. (This is because δ is the maximum match efficiency improvement relative to a random match. Given the consumer-optimal policy is better than no-information disclosure (in which case consumers buy a random product at a zero price), the consumer-optimal price should not exceed δ.) When f is log-concave, we have known from (27) that $\delta \leq \frac{1}{2}p_F$, and so $p^{**} \leq \frac{1}{2}p_F$.

On the other hand, when F is convex on $[-1,0]$, we have $\frac{\gamma}{2} = F(x^*) \geq \frac{1}{2} + f(0)x^* = \frac{1}{2}(1+\frac{x^*}{p_F})$, where the first equality is from the proof of Proposition 2 and the final equality is from $p_F = \frac{1}{2f(0)}$. This implies $\frac{-x^*}{p_F} \geq 1-\gamma$. Since $p^{**} = \frac{-\gamma}{1-\gamma}x^*$, we deduce $p^{**} \geq \gamma p_F$. ■

Proof of Proposition 3. We first derive the minimum possible consumer surplus for a given profit $p \leq p^*$. The following lemma is the key step:

Lemma 3 *If an asymmetric posterior induces a pure-strategy equilibrium, there exists a symmetric posterior which induces the same profit but weakly lower total welfare (and so weakly lower consumer surplus).*

Proof. Suppose that posterior G induces an asymmetric equilibrium with industry profit π. Suppose firms are labelled so that $p_1 \leq p_2$, so the bounds for G are as shown in Figure 2a. Let π_i be firm i's profit. Then we have $\pi_1 < \frac{1}{2}\pi < \pi_2$ and $p_1 < \pi < p_2$. Let $\Delta = p_1 - p_2 < 0$. Total welfare under G is

$$\mu + \int_{-1}^{0} G(x)dx + \int_{\Delta}^{0} xdG(x) = \mu + \int_{-1}^{\Delta} G(x)dx - \Delta G(\Delta) \ .$$

(For those consumers with expected $x \in (\Delta, 0)$, they buy the worse matched product 1 because of its lower price, which causes an efficiency loss $\int_{\Delta}^{0} xdG(x)$. The equality is from integration by parts.) Notice that $G(\Delta)$ is firm 2's equilibrium demand and according to Lemma 1 this should be greater than $\frac{1}{2}$.

Consider a symmetric posterior $\hat{G}(x) = L_\pi(x)$ for $x < 0$, where L_π is the lower bound for posteriors which induce a symmetric equilibrium with price π. Notice that $L_\pi(x) = 1 - \frac{\pi/2}{\pi + x} \leq L_{p_1,p_2}(x) = 1 - \frac{\pi_1}{p_2 + x}$, where L_{p_1,p_2} is the lower bound of G. Following a similar chain of inequalities as in (12), we can see that $\hat{G}(x)$ is an MPC of F. On the other hand, we have

$$\begin{aligned} \int_{-1}^{\Delta} G(x)dx - \Delta G(\Delta) &\geq \int_{-1}^{\Delta} L_\pi(x)dx - \Delta G(\Delta) \\ &\geq \int_{-1}^{\Delta} L_\pi(x)dx + \int_{\Delta}^{0} L_\pi(x)dx \\ &= \int_{-1}^{0} L_\pi(x)dx \ , \end{aligned}$$

where the first inequality used $G \geq L_{p_1,p_2} \geq L_\pi$ and the second used $G(\Delta) \geq \frac{1}{2} \geq L_\pi(x)$ for $x \leq 0$. Therefore, the symmetric posterior $\hat{G}(x)$ induces lower total welfare and so lower consumer surplus. ■

Therefore, we can focus on symmetric posteriors when we calculate the minimum possible consumer surplus for a given profit level, and the argument in the main text follows.

We then turn to the maximum possible consumer surplus for a given profit. Under Conditions 1 and 2(ii), the full-information price p_F is well defined and F is log-concave on $[-1, 0]$. When $p_F \leq p \leq p^*$, it is possible to find a supporting posterior G between the bounds L_p and U_p which is an MPC of F such that (3) holds with equality (i.e., total welfare is maximized). This is because given $p < p^*$ the lower bound L_p must have (3) hold with strict inequality, while given $p \geq p_F$ the upper bound U_p must be above F when $x < 0$ as $U_p \geq U_{p_F} \geq F$. One way to construct such a G is to choose $G(x) = L_p(x)$ for $x < m$ and $G(x) = U_p(x)$ for $m \leq x \leq 0$, where m is chosen to make (3) bind. (This G is an MPC of F since $L_p \leq L_{p^*}$ is an MPC of F.) Since there is no mismatch with such a posterior, this achieves the maximum consumer surplus $\mu + \delta - p$ for a given p.

To prove the result for industry profit $p < p_F$, we need the following result to rule out asymmetric posteriors:

Lemma 4 *When F is log-concave on $[-1, 0]$, if an asymmetric posterior induces a pure-strategy equilibrium with industry profit less than p_F, there exists a symmetric posterior which induces the same profit but weakly higher total welfare (and so weakly higher consumer surplus).*

Proof. Suppose that posterior G induces an asymmetric equilibrium with industry profit $\pi < p_F$. Suppose firms are labelled so that $p_1 \geq p_2$, so the bounds for G are as shown in Figure 2c. As in the proof of Lemma 2, $\Delta = p_1 - p_2 \geq 0$, $\pi_2 \leq \frac{1}{2}\pi \leq \pi_1$ and $\pi \leq p_1$, while total welfare is no greater than (4).

Consider the symmetric posterior $\hat{G}(x) = \min\{F(x), U_\pi(x)\}$ for $x < 0$, where U_π is the upper bound for posteriors which induce the symmetric equilibrium price $p = \pi$. Since $\pi < p_F$, the log-convex U_π must cross the log-concave F once and from above in the range $x < 0$. Since $\hat{G} \leq F$ it is an MPC of the prior and so is a legitimate posterior. It is also between the bounds U_π and L_π since $L_\pi < L_{p_F} \leq F$. Hence, $\hat{G}$ induces a symmetric equilibrium with profit π, the same as with G, and results in total welfare $\mu + \int_{-1}^{0} \hat{G}(x)dx$. If this total welfare is higher than (4), an upper bound for welfare with G, then $\hat{G}$ leads to higher consumer surplus.

Notice that U_{p_1,p_2}, the upper bound of G, is log-convex, and so it crosses F once at, say, $\hat{x} \in (-1, 0)$. Then

$$\begin{aligned}\int_{-1}^{0} G(x)dx &\leq \int_{-1}^{\hat{x}} F(x)dx + \int_{\hat{x}}^{0} U_{p_1,p_2}(x)dx \\ &= \int_{-1}^{0} \min\{F(x), U_{p_1,p_2}(x)\}dx \\ &\leq \int_{-1}^{0} \min\{F(x), U_{\pi}(x)\}dx \\ &= \int_{-1}^{0} \hat{G}(x)dx \ .\end{aligned}$$

Here the first inequality follows since G is an MPC of F and lies below U_{p_1,p_2}, and the second inequality follows from (25). This proves the result. ■

Therefore we can focus on symmetric posteriors when we calculate the maximum consumer surplus given industry profit $p < p_F$. Then the same analysis as in the consumer-optimal problem in section 3.4 applies, and the maximum consumer surplus, for a given $p < p_F$, equals (15). In this case total welfare is not maximized, and as shown in the proof of Proposition 2, the derivative of (15) with respect to p decreases in p. Hence, the welfare frontier is concave in p, and it reaches the optimum at the positive consumer-optimal price in (16). ■

Proof of Proposition 4. Let F denote the symmetric prior distribution of $x = v_1 - v_2$ and G denote a symmetric posterior distribution of (expected) x. Suppose there is a symmetric pure or mixed strategy pricing equilibrium under G. Consumer surplus under G is no greater than $\mu + \int_{-1}^{0} G(x)dx$ (which is total welfare when the consumer buys her preferred product) minus industry profit in that equilibrium. We first derive a lower bound on that industry profit:

Lemma 5 *Suppose $V \geq 1$ and let G be a symmetric distribution for $x = v_1 - v_2$. Then in any symmetric equilibrium (with pure or mixed strategies) industry profit is no lower than*

$$\max_{x\in[-1,1]} : \frac{-2xG(x)}{1-G(x)} \ . \tag{29}$$

Proof. Note that (29) is zero if and only if the distribution G is degenerate at $x = 0$, in which case equilibrium profit is also zero and the result holds. Suppose now that (29) is positive, and slightly abusing the notation denote its value by $p > 0$. Since x which solves (29) must be negative, we have $p \leq 2$. Suppose in contrast to the statement that there is an equilibrium where each firm obtains profit π^* strictly below

$p/2$. Firm 1, say, will never choose a price below π^* in this equilibrium (as then it obtains lower profit even if it serves all consumers). Since a firm's profit increases with its rival's price, firm 2's profit π^* is then at least equal to the maximum profit it can obtain if firm 1 chooses price π^*. Given $\pi^* < p/2 \le 1 \le V$, if firm 1 chooses price π^* the outside option is not relevant for consumers, regardless of the price chosen by firm 2. Hence firm 2's profit π^* satisfies

$$\pi^* \ge \max_{p'} : p' \times \Pr\{v_2 - p' \ge v_1 - \pi^*\} = \max_{p'} : p' G(\pi^* - p') = \max_{x \in [-1,1]} : (\pi^* - x) G(x) \ ,$$

where the final equality follows after changing to the variable $x = \pi^* - p'$. Thus for any $x \in [-1, 1]$ we have $(1 - G(x))\pi^* \ge -xG(x)$, in which case π^* is at least equal to $p/2$. As this contradicts our assumption, the result is proved. ■

Let p denote (29) for a given G. (The proof of Lemma 5 shows that $p \le 2$.) Then an upper bound on consumer surplus with posterior G is $\mu + \int_{-1}^{0} G(x)dx - p$. By construction, for any $x \in [-1, 0)$ we have $G(x)/(1 - G(x)) \le \frac{p}{-2x}$, or $G(x) \le \frac{p}{p-2x}$. (If the lower bound p in Lemma 5 is attained, then G should equal $\frac{p}{p-2x}$ for some $x < 0$.) Since G cannot exceed $\frac{1}{2}$ for $x \in [-1, 0]$, it follows that G in the negative range lies below the upper bound

$$G(x) \le \hat{U}_p(x) \equiv \min \left\{ \frac{1}{2}, \frac{p}{p - 2x} \right\} \ . \tag{30}$$

Here, the upper bound $\hat{U}_p$ increases with p and x, and reaches $\frac{1}{2}$ at $x = -\frac{1}{2}p \ge -1$.

Given the prior distribution F, let τ_p denote the maximum match efficiency when the lower bound on industry profit is p, i.e., $\tau_p =: \max_G \int_{-1}^{0} G(x)dx$ subject to (i) G lying below the upper bound $\hat{U}_p(x)$ in (30) (and touching it at some $x < 0$) and (ii) G being a symmetric MPC of F. Then an upper bound on consumer surplus is $\mu + \max_p(\tau_p - p)$. Notice that for a given p, the optimization problem in defining τ_p is similar to the problem of finding the best symmetric G for consumers under the pure-strategy restriction in section 3.4, except that here is the relevant upper bound for G is $\hat{U}_p$ instead of U_p in (11) which for negative x equals $p/(2p - 2x)$ and lies below $\hat{U}_p$. However, for small p, which is usually the relevant case for the consumer-optimal solution, the two bounds are very close and for this reason the use of mixed strategies cannot significantly benefit the consumer.

The remaining task is to calculate τ_p, and this can be done in a manner similar to the way we found the consumer-optimal policy for a given price with pure strategies. If the prior has a log-concave density, then F is log-concave on $[-1, 0]$, while the upper bound $\hat{U}_p$ is log-convex in the range $[-1, -\frac{1}{2}p]$. For relatively small p, which will be the relevant case, the upper bound $\hat{U}_p$ therefore crosses the prior F twice.[26] See Figure 10 for an illustration when the prior is uniform. Let $\hat{x}_p$ denote the smaller of the two crossing points given p (i.e., the smaller solution to $\frac{p}{p-2x} = F(x)$). As in section 3.4, two necessary conditions for a feasible G are that it satisfies the MPC constraint at the intercept point $\hat{x}_p$, and that G lies below $\hat{U}_p$ for $x \in [\hat{x}_p, 0]$. The bold curve on Figure A1 is then a convenient candidate for the optimal G.

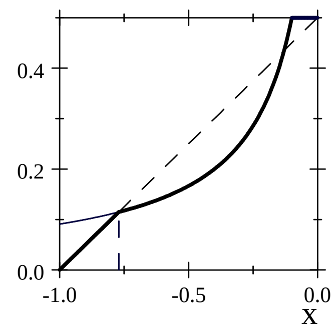

Figure 10: Consumer-optimal way to reach the profit lower bound p

We have not shown that the candidate G illustrated as the bold curve in Figure 10 is an MPC of F, as $\hat{U}_p$ is above F for x close to zero. Therefore, the resulting τ_p based on this G is an upper bound on the feasible match efficiency when the MPC constraint is fully considered.

As with expression (15), an upper bound on consumer surplus given p is therefore

$$\mu + \int_{-1}^{\hat{x}_p} F(x)dx + \int_{\hat{x}_p}^{0} \hat{U}_p(x)dx - p \ . \tag{31}$$

The derivative of this expression with respect to p is

$$\int_{\hat{x}_p}^{-\frac{1}{2}p} \frac{\partial \hat{U}_p(x)}{\partial p} dx - 1 = \frac{1}{2}\left(\frac{p}{p-2\hat{x}_p} - \log\frac{2p}{p-2\hat{x}_p} - \frac{5}{2}\right) \ .$$

This equals $\frac{1}{2}(F(\hat{x}_p) - \log(2F(\hat{x}_p)) - \frac{5}{2})$ by using $\frac{p}{p-2\hat{x}_p} = F(\hat{x}_p)$. Note that $\hat{x}_p$ increases with p given that $\hat{U}_p$ crosses F from above at the smaller of the two crossing points. This derivative therefore decreases with p, and so the point $\hat{x}^*$ which maximizes the upper bound (31) satisfies $F(\hat{x}^*) = \hat{\gamma}$, where $\hat{\gamma} \approx 0.043$ is the root of $\gamma - \log(2\gamma) = \frac{5}{2}$.

[26] To see that with the p which maximizes $\tau_p - p$ the upper bound $\hat{U}_p$ crosses the prior F, we argue as follows. For any p we must have $\tau_p \leq \delta$ since G is an MPC of F. Let $\tilde{p}$ denote the price such that $\hat{U}_{\tilde{p}}$ just touches F. Then setting $G \equiv F$ solves the stated problem for $\tau_{\tilde{p}}$, in which case $\tau_{\tilde{p}} = \delta$. For $p > \tilde{p}$, when the upper bound $\hat{U}_p$ lies strictly above F, we must have $\tau_p - p \leq \delta - p < \tau_{\tilde{p}} - \tilde{p}$. As claimed, then, the p which maximizes $\tau_p - p$ is no greater than $\tilde{p}$ and so the upper bound crosses F.

Evaluating the upper bound (31) at this crossing point $\hat{x}^*$ shows the maximum consumer surplus upper bound to be

$$\mu + \int_{-1}^{\hat{x}^*} F(x)dx - \hat{x}^* F(\hat{x}^*) = \mu - \int_{-1}^{\hat{x}^*} xdF(x) \ . \tag{32}$$

We next compare this with the optimal consumer surplus under the pure-strategy restriction as described in Proposition 2. Note that Corollary 2 implies that the consumer-optimal price p^{**} is (much) less than 1, and so the assumption $V \geq 1$ implies the outside option is not relevant in the consumer-optimal policy. Proposition 2 implies that optimal consumer surplus is

$$\begin{aligned} \mu + \int_{-1}^{x^*} F(x)dx + \int_{x^*}^{0} U_{p^{**}}(x)dx - p^{**} &= \mu + \int_{-1}^{x^*} F(x)dx - p^{**}\left(1 + \tfrac{1}{2}\log \frac{p^{**}}{p^{**} - x^*}\right) \\ &= \mu + \int_{-1}^{x^*} F(x)dx - x^* F(x^*) \\ &= \mu - \int_{-1}^{x^*} xdF(x) \ . \end{aligned} \tag{33}$$

(Here, the second equality used $\frac{p^{**}}{p^{**}-x^*} = 2F(x^*) = \gamma$, the definition of γ, and (16), while the last equality follows by integration by parts.) This is the same expression as (32) but using $x^* < \hat{x}^*$.

Finally, we show that consumer surplus with pure strategies in (33) comes close to reaching the upper bound in (32). The ratio of (33) to (32) is

$$\frac{\mu - \int_{-1}^{x^*} xdF(x)}{\mu - \int_{-1}^{\hat{x}^*} xdF(x)} = \frac{\mu - x^* F(x^*) + \int_{-1}^{x^*} F(x)dx}{\mu - \hat{x}^* F(\hat{x}^*) + \int_{-1}^{\hat{x}^*} F(x)dx} > \frac{\mu - x^* F(x^*) + \int_{-1}^{x^*} F(x)dx}{\mu - x^* F(\hat{x}^*) + \int_{-1}^{x^*} F(x)dx}$$

$$> \frac{\mu - x^* F(x^*)}{\mu - x^* F(\hat{x}^*)} > \frac{\mu + F(x^*)}{\mu + F(\hat{x}^*)} \geq \frac{1 + F(x^*)}{1 + F(\hat{x}^*)} = \frac{1 + \frac{1}{2}\gamma}{1 + \hat{\gamma}} \approx 0.984 \ .$$

Here, the first inequality uses $\int_{x^*}^{\hat{x}^*} F(x)dx < (\hat{x}^* - x^*)F(\hat{x}^*)$, the third inequality uses $-1 < x^* < 0$, and the final inequality uses the fact that $V \geq 1$ implies $\mu \geq 1$. Thus, when the prior has log-concave density the maximum consumer surplus attainable with pure strategies attains at least 98.4% of the consumer surplus which could be available when mixed pricing strategies were permitted. ■

Proof of Proposition 5: (i) All the results have been explained in the main text, except for the comparison between the "top product" signal structure and full disclosure in the "i.i.d." case. Let us first report the equilibrium price under full disclosure. Suppose (with some abuse of notation) that each v_i has CDF $F(v_i)$ and density $f(v_i)$ on the support support $[\underline{v}, \bar{v}]$, where $\bar{v} = \infty$ is allowed. When the outside option is irrelevant and the density $f(\cdot)$ is log-concave, the unique equilibrium price in the Perloff-Salop model is

$$p_F = \frac{1}{n \int_{\underline{v}}^{\overline{v}} f(v) dF(v)^{n-1}} = \frac{1}{n(n-1) \int_0^1 f(F^{-1}(t)) t^{n-2} dt} , \tag{34}$$

where the second equality is from changing the integral variable, and this price decreases with n and converges to 0 as $n \to \infty$ if

$$\lim_{v \to \overline{v}} \frac{1 - F(v)}{f(v)} = 0 . \tag{35}$$

(For this analysis, see for instance Lemma 1 in Zhou (2017).) Therefore, the outside option is indeed irrelevant for consumers if $\underline{v} > p_F$ when $n = 2$. Note that when the outside option is irrelevant, p_F is also the industry profit under full disclosure.

The profit lower bound (19) with the "top product" signal is

$$\mu_{n:n} - \mu_{<n:n} = \frac{n}{n-1}(\mu_{n:n} - \mu) = \frac{n}{n-1} \int_{\underline{v}}^{\overline{v}} [F(v) - F(v)^n] dv = \frac{n}{n-1} \int_0^1 \frac{t - t^n}{f(F^{-1}(t))} dt ,$$

where the first equality follows from the identity (17) and the final equality follows after changing variables.

We first show that $\mu_{n:n} - \mu_{<n:n} \geq p_F$ if $n \geq 3$. Note that $\mu_{n:n} - \mu_{<n:n} \geq p_F$ if

$$n^2 \left(\int_0^1 \frac{t - t^n}{f(F^{-1}(t))} dt \right) \left(\int_0^1 f(F^{-1}(t)) t^{n-2} dt \right) \geq 1 . \tag{36}$$

The Cauchy-Schwartz inequality implies that the product of the two integrals is at least

$$\left(\int_0^1 \sqrt{(t - t^n) t^{n-2}} dt \right)^2 = \left(\int_0^1 \sqrt{t^{n-1}(1 - t^{n-1})} dt \right)^2 .$$

Therefore, (36) holds if

$$n \int_0^1 \sqrt{t^{n-1}(1 - t^{n-1})} dt \geq 1 .$$

Rewrite the left-hand side as

$$n \int_0^1 \sqrt{t^{n-1}(1 - t^{n-1})} dt = \frac{2n}{n+1} \int_0^1 \sqrt{1 - t^{n-1}} dt^{(n+1)/2} = \frac{2n}{n+1} \int_0^1 \sqrt{1 - z^{2\frac{n-1}{n+1}}} dz ,$$

where the last step follows from changing variable, which is therefore increasing in n. With $n = 3$ it equals $\frac{3}{2} \int_0^1 \sqrt{1 - z} dz = 1$. Therefore, (36) holds for all $n \geq 3$.

The remaining case is $n = 2$. In this case the lower bound (19) is not a tight enough bound on profit, and we need to calculate equilibrium profit with the "top product" policy when $n = 2$. To do that, however, we need to rely on the mixed-strategy characterization presented in our working paper, Armstrong and Zhou (2021, pp. 53-54), and the analysis there can be used to show that industry profit with this policy is greater than $2(\mu_{n:n} - \mu_{<n:n})$ when $n = 2$. This in turn is greater than the profit under full disclosure if

$$8\left(\int_0^1 \frac{t-t^2}{f(F^{-1}(t))}dt\right)\left(\int_0^1 f(F^{-1}(t))dt\right) \geq 1 \ .$$

By Cauchy-Schwartz again, the left-hand side is greater than

$$8\left(\int_0^1 \sqrt{t(1-t)}dt\right)^2 = \frac{\pi^2}{8} > 1$$

as required, where π is the mathematical constant.

(ii) The first part of the result has been proved in the main text. We now compare consumer surplus between the "top two" signal and full disclosure in the i.i.d. case. Under full disclosure, the consumer always buys her preferred product, so her expected surplus is $\mu_{n:n} - p_F$. As explained in the main text, with the "top-two" policy consumer surplus is (20). Therefore, the consumer prefers this policy to full disclosure if $\frac{1}{2}(\mu_{n:n} + \mu_{n-1:n}) > \mu_{n:n} - p_F$, i.e., if $\mu_{n:n} - \mu_{n-1:n} < 2p_F$. (Both sides converge to zero as $n \to \infty$ under condition (35).) The CDF of the best match is $F(v)^n$ and the CDF of the second-best match is $F(v)^n + n(1 - F(v))F(v)^{n-1}$. Therefore,

$$\mu_{n:n} - \mu_{n-1:n} = n\int_{\underline{v}}^{\overline{v}} (1 - F(v))F(v)^{n-1}dv = n(\mu_{n:n} - \mu_{n-1:n-1}) \ .$$

That is, the expected gap between the best match and the second-best match equals n times the expected difference between the best match among n options and the best match among $n-1$ options. Therefore, it suffices to show $\mu_{n:n} - \mu_{n-1:n-1} < \frac{2}{n}p_F$. In fact, a stronger version of this inequality has been shown previously to be true when the density f is log-concave, namely

$$\mu_{n:n} - \mu_{n-1:n-1} \leq \frac{p_F}{n} \ . \tag{37}$$

(See Anderson, de Palma, and Nesterov (1995) and Tan and Zhou (2021) for this result.) This completes the proof. ■

References

Ali, N., G. Lewis, and S. Vasserman (2020): "Voluntary Disclosure and Personalized Pricing," mimeo.

Anderson, S., A. de Palma, and Y. Nesterov (1995): "Oligopolistic Competition and the Optimal Provision of Products," *Econometrica*, 63(6), 1281–1301.

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.

Armstrong, M., and J. Vickers (2019): "Discriminating against Captive Customers," *American Economic Review: Insights*, 1(3), 257–272.

Armstrong, M., and J. Zhou (2021): "Information Structure and the Limits to Competition," Cowles Foundation Discussion Paper no. 2269.

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. Cunat (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.

——— (2017): "First Price Auctions with General Information Structures: Implications for Bidding and Revenue," *Econometrica*, 85(1), 107–143.

——— (2020): "Search, Information and Prices," *Journal of Political Economy*, forthcoming.

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.

Caplin, A., and B. Nalebuff (1991): "Aggregation and Social Choice: A Mean Voter Theorem," *Econometrica*, 59(1), 1–23.

Choi, M., K. Kim, and M. Pease (2019): "Optimal Information Design for Search Goods," *American Economic Review P&P*, 109(5), 550–556.

Condorelli, D., and B. Szentes (2020): "Information Design in the Hold-up Problem," *Journal of Political Economy*, 128(2), 681–709.

Dogan, M., and J. Hu (2021): "Consumer Search and Optimal Information," mimeo.

Du, S. (2018): "Robust Mechanisms Under Common Valuation," *Econometrica*, 86(5), 1569–1588.

Dworczak, P., and G. Martini (2019): "The Simple Economics of Optimal Persuasion," *Journal of Political Economy*, 127(5), 1993–2048.

Elliot, M., A. Galeotti, and A. Koh (2020): "Market Segmentation Through Information," mimeo.

Ericson, K., and A. Starc (2016): "How Product Standardization Affects Choice: Evidence from the Massachusetts Health Insurance Exchange," *Journal of Health Economics*, 50, 71–85.

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

Hwang, I., K. Kim, and R. Boleslavsky (2020): "Competitive Advertising and Pricing," mimeo.

Ichihashi, S. (2020): "Online Privacy and Information Disclosure by Consumers," *American Economic Review*, 110(2), 569–595.

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.

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

Li, W. (2020): "Using Information to Amplify Competition," mimeo.

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

Perloff, J., and S. Salop (1985): "Equilibrium with Product Differentiation," *Review of Economic Studies*, 52(1), 107–120.

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.

Shi, X., and J. Zhang (2020): "Welfare of Price Discrimination and Market Segmentation in Duopoly," mimeo.

Strassen, V. (1965): "The Existence of Probability Measures with Given Marginals," *Annals of Mathematical Statistics*, 36(2), 423–439.

Tan, G., and J. Zhou (2021): "Consumer Heterogeneity and Inefficiency in Oligopoly Markets," mimeo.

Zhou, J. (2017): "Competitive Bundling," *Econometrica*, 85(1), 145–172.