Discussion Paper Series – CRC TR 224

Discussion Paper No. 700
Projects B03 and B04

Designing Vertical Differentiation with Information

Christoph Carnehl[1]
Anton Sobolev[2]
Konrad Stahl[3]
André Stenzel[4]

February 2026 (First Version : August 2025)

[1] Bocconi University, Department of Economics and IGIER. Email: christoph.carnehl@unibocconi.it
[2] University of Mannheim and MaCCI. Email: anton.sobolev@uni-mannheim.de
[3] University of Mannheim, CEPR, CESifo, MaCCI and ZEW. Email: konrad.stahl@uni-mannheim.de.
[4] Bank of Canada. Email: astenzel@bankofcanada.ca

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Collaborative Research Center Transregio 224 - www.crctr224.de
Rheinische Friedrich-Wilhelms-Universität Bonn - Universität Mannheim

Designing Vertical Differentiation with Information*

Christoph Carnehl[†] Anton Sobolev[‡] Konrad Stahl[§] André Stenzel[¶]

February 17, 2026

Abstract

We study information design in a vertically differentiated market. A third party publicly discloses information about the product qualities of two competing firms. More precise information improves consumer matching but increases perceived differentiation, enabling firms to raise prices. Disclosing the product ranking alone suffices to maximize industry profits in a fully covered market. Consumer surplus, however, is maximized by a rank-preserving policy that withholds any information that overturns the prior ranking, as gains from price competition outweigh losses from allocative inefficiency. The conflict between profit- and consumer-optimal policies persists in settings with endogenous participation and nonlinear or asymmetric costs.

JEL Classification Numbers: D43, D82, L13, L15.
Keywords: Information Design, Vertical Product Differentiation, Quality Rankings, Competition

*We are grateful for comments we have received from Gregorio Curello, Volker Nocke and Nicolas Schutz, and seminar participants at Universidad de los Andes (Chile) and the Chile Strategy Conference 2025. Anton Sobolev and Konrad Stahl gratefully acknowledge support by the Deutsche Forschungsgemeinschaft through CRC TR 224 (projects B03 and B04).

[†]Bocconi University, Department of Economics and IGIER, christoph.carnehl@unibocconi.it
[‡]University of Mannheim and MaCCI, anton.sobolev@uni-mannheim.de
[§]University of Mannheim, CEPR, CESifo, MaCCI and ZEW, konrad.stahl@uni-mannheim.de.
[¶]Universidad de los Andes, Chile, astenzel@uandes.cl

1 Introduction

In many markets, consumers would agree on the ranking of products if they had full information about their qualities. In practice, however, such information is rarely available, and consumers instead rely on third-party sources to form beliefs about the products they consider purchasing. Anticipating how information shapes consumer beliefs and demand, sellers incorporate it into their pricing decisions. In such an environment, it is not clear who benefits from public information provision.

We study markets in which the informing third party is a certifier, a platform on which the sellers are active, or a regulatory institution. They may provide information in various forms, such as direct product comparisons, search rankings, or product ratings, thus informing both sides of the market about the quality of the products offered.[1] Indeed, widely accessed third-party evaluations such as Consumer Reports in the United States or reports by the German Stiftung Warentest focus largely on the comparative evaluation of *relative* qualities, instead of a descriptive comparison of products along dimensions evaluated differently across consumers.

A trade-off arises because information simultaneously allows consumers to make better decisions and firms to condition their prices on the revealed information, potentially harming consumers. To analyze this trade-off, we build on a parsimonious model of vertical product differentiation à la Shaked and Sutton (1982), in which two firms offer products of different qualities to consumers who are heterogeneous in their taste for quality. Firms' qualities are ex ante unknown to both consumers and firms, possibly correlated and asymmetrically distributed, and a third party designs an experiment that publicly discloses information about the firms' qualities. The experiment design is fully flexible and can range from being completely uninformative to fully revealing product qualities. Our setting of mutual, rather than consumer-only, uncertainty is especially relevant for new, complex, or hard-to-test products, where the firms are uncertain about their relative qualities and thus, demand.[2]

In a baseline setting with full market coverage, we show that the relevant property of experiments is their informational content about the *ranking* of product qualities, rather than cardinal quality levels or differences. Firms always prefer experiments that are *rank-revealing*, that is, experiments that reveal the true quality ranking of their products. In contrast, consumer surplus is maximized by *rank-preserving* experiments, that is, experi-

[1]In most real-world cases, firms can easily obtain access to information even if it is primarily targeted at consumers.

[2]While firms are informed about their products' characteristics and intended positioning, they are often uncertain about how consumers will evaluate their product relative to competitors' products and whether it will ultimately be perceived as superior. This uncertainty may arise from imperfect knowledge of consumer preferences, heterogeneity in how products are used, and limited information about which product attributes will be most salient in consumers' purchase decisions. Likewise, consumers do not know which product is of higher quality, which naturally motivates a role for public information disclosure in shaping demand and competitive outcomes.

ments that do not overturn the prior ranking of product qualities. There is thus a clear tension between the impact of information on consumers and firms. The detrimental impact of information provision on consumer surplus raises doubts about the desirability of consumer-driven certification in vertically differentiated markets.

The mechanisms underlying our results are as follows. When signals are uninformative, consumers rank products based on their prior beliefs, and firms compete on the basis of perceived ex ante vertical differentiation. When signals are informative about the true product ranking, perceived differentiation increases. This, in turn, relaxes competition and allows firms to raise prices. At the same time, information improves product matching for heterogeneous consumers, which benefits them. While these effects suggest an ambiguous impact of information revelation to consumers, we show that the positive competitive effect always dominates.

Consumers therefore prefer less information about product qualities, even when it leads to a suboptimal allocation of consumers across products. We obtain this result even in environments where information would a priori seem most valuable, such as when product qualities are negatively correlated. By contrast, firms always prefer public provision of information about the product ranking. This is because increased differentiation relaxes competition and raises expected prices, including for the low-quality firm.

Since consumers and producers value information in opposing ways, the partial welfare outcome, comprising consumer and producer surplus, depends on their relative magnitudes as well as the welfare weights assigned to each. We show that maximizing unweighted welfare calls for disclosure of the products' quality ranking. Under welfare standards that place greater weight on consumer surplus, however, the absence of information provision about the ranking can be optimal.

We show that our main insights extend to the important case where consumers may decide against purchase. Under an extensive margin, the firms can no longer extract the rents from inelastically demanding consumers, but are faced with outside competition. To keep the analysis in this extension tractable and focused, we restrict attention to a state space with two asymmetric qualities and a uniformly distributed consumer preference parameter. In this setup, rank revelation—which coincides with full information revelation in the asymmetric two-state case—maximizes industry profits, while the consumer-optimal information structure maximizes information revelation conditional on preserving the prior ranking between the two products.[3] We also show that the tension between firms' desire for public information provision and its detrimental impact on consumers carries over to settings in which firms' production costs are non-linear or asymmetric. Most importantly, we show that the consumer-optimal experiment is typically rank-preserving in both cases.

Overall, our analysis highlights the adverse effects of public information provision on

[3] Relative to the baseline result, the latter modification is because information revelation subject to this condition keeps prices low while maximizing market participation, which increases consumer surplus.

consumers in markets with vertical differentiation. Partial information provision may therefore be desirable, even when they generate allocative inefficiencies.The contrast with settings featuring horizontal differentiation (see, e.g., Armstrong and Zhou, 2022; Bergemann et al., 2025) reinforces the need for consumer-oriented policies to be tailored to the specific market environment.

Below, we embed our argument into the extant literature. The model is specified in Section 2. In Section 3, we analyze the model. We discuss extensions, including the possibility for consumers not to purchase as well as alternative cost function specifications, in Section 4.

Related Literature We contribute to a recent and growing literature on information design in industrial organization broadly; see Bergemann and Morris (2019), Bergemann and Bonatti (2019), and Kamenica (2019) for surveys on information design in general. More specifically, we relate to the literature on the degree of information provision (see, for example, Hopenhayn and Saeedi, 2023), and to quality disclosure by a certifier as surveyed by Dranove and Jin (2010).

Most closely related to our work are Armstrong and Zhou (2022) and Roesler and Szentes (2017). The latter show that the consumer-optimal signal structure in a monopoly involves posterior distributions featuring partial learning. We find no learning beyond initial product rankings to be the consumer-optimal signal structure, as it stifles competition—an effect absent in their monopoly model.

As Armstrong and Zhou (2022), we focus on a competitive duopoly and, in line with earlier studies such as Anderson and Renault (2009), on a central provider of information rather than decentralized information provision, e.g., by firms. While the key trade-off—more information results in less competition and higher pricesbut allows consumers to choose the "right" alternative more frequently—is similar to Armstrong and Zhou (2022) and Biglaiser et al. (2025), we consider vertical instead of horizontal differentiation and allow for ex-ante asymmetric firms. In line with idiosyncratic consumer preferences, Armstrong and Zhou (2022) consider personalized provision of information, while we focus on public information provision with common product rankings across consumers. The latter appears more suitable when information is about common, objective product components.

These differences in the setup are economically relevant. We show that to maximize firm profits it is sufficient to reveal only product rankings. However, after any information provision, consumers remain heterogeneous, allowing both firms to retain positive market shares. Thus, in contrast to Armstrong and Zhou (2022), the firm-optimal policy does not induce allocative efficiency. Moreover, under the consumer-optimal policy no information about the quality ranking beyond the prior is disclosed. This also holds in the empirically relevant situation in which there is an extensive margin. These findings are related in spirit to those in Lewis and Sappington (1994), who consider a monopolistic supplier. It is also worth noting that if firms are perceived as ex-ante symmetric (as in Armstrong and Zhou, 2022), not revealing any information would be weakly optimal for consumer surplus.

From a welfare perspective, we show that welfare-optimal and seller-optimal rather than consumer-optimal information structures tend to be aligned in our setup. This is in contrast to Bergemann et al. (2025), who consider a market for horizontally differentiated products and the provision of information about the buyers' valuation.

Condorelli and Szentes (2025) similarly focus on the comparison of welfare- and consumer-optimal allocations in a model addressing the two-sided matching design problem of a platform that has superior information about its users but does not control their bargaining. They show that a platform can use the matching procedure to garble the information of sellers in a way that offsets their bargaining power at the cost of creating sorting inefficiencies compared to the first-best outcome. This is in line with the allocative distortion implied by the consumer-optimal information structure in our setting.

While their paper and ours address similar economic trade-offs in the design of information environments, they do so in different and thus complementary settings. We focus on third-party information provision and competition after information provision, whereas they study platform-side matching technologies that determine consumers' consideration sets in digital markets and their implications for post-matching outcomes. In particular, consumers in their model are presented with just one take-it-or-leave-it offer after the matching. Our results complement their analysis by showing that a trade-off between buyer surplus and allocative efficiency arises when there is the potential for information provision by third parties.

Hopenhayn and Saeedi (2023) primarily focus on the effect of information as mitigating adverse selection, by which low-quality firms crowd out high-quality firms' sales by pooling with them when consumers do not know the difference. Their attribution of the surplus generated by information to consumers vs. producers varies with concavity vs. convexity of the market supply function.[4] Adverse selection does not figure in our setup.

Finally, quality certification—an early example is Albano and Lizzeri (2001), a recent one Zapechelnyuk (2020)—may be focused primarily on the incentives to affect the suppliers' provision of higher quality. This is not the issue at stake here. We take quality as given and focus on the consequences of a third party informing both buyers and suppliers about quality differences.

[4] The finding that information may hurt consumers when the supply function is sufficiently convex is also present in Schlee (1996).

2 Model

There are two firms, indexed by $i \in \{A, B\}$, each producing a product of different quality, and competing in prices. The quality of firm i's product is $q_i \in \mathcal{Q} \equiv \{q^1, \ldots, q^n\}$, where $q^n > q^{n-1} > \ldots > q^1 \geq 0$. We assume that the production cost is independent of quality and normalized to zero. The state of the market, defined by the pair of product qualities, is denoted by $\omega \in \Omega \equiv \{(q_A, q_B) | q_i \in \mathcal{Q}\}$. Let $\lambda_{k\ell} \in [0, 1)$ be the common prior probability that the quality of firm A is q^k and the quality of firm B is q^ℓ, where $k, \ell = 1, \ldots, n$. We assume that there is initial uncertainty about the ranking of the firms' qualities, that is, that the prior satisfies $\mathbb{P}[q_A > q_B] \in (0, 1)$.

There is a unit mass of consumers, each demanding one unit of the good. Consumers differ in their willingness to pay for quality θ, which is distributed according to the CDF $F : [\underline{\theta}, \overline{\theta}]$, with $\underline{\theta} \geq 0$. We assume that it admits density $f(\theta) > 0$, which is twice differentiable and log-concave on $[\underline{\theta}, \overline{\theta}]$. To ensure an interior solution of the pricing subgame, we additionally assume that $\underline{\theta} < 1/f(\underline{\theta})$.

A consumer's valuation from purchasing a good of quality q at price p, given her willingness to pay for quality θ, is $\theta q - p$. For now, we assume that the outside option $\underline{v}$ of consumers is sufficiently low (e.g., $\underline{v} = -\infty$), ensuring that the market is fully covered. We discuss the implications of relaxing this assumption in Section 4.1.

We assume that there is an information designer who chooses an experiment that generates a public signal about the firms' qualities. We separately analyze two distinct objectives of the information designer: *(i)* the maximization of total industry profit, and *(ii)* the maximization of consumer surplus.[5] Given the finite state space, we can restrict attention to finite signal structures without loss of generality. Formally, an experiment $\mathcal{E}$ is represented by a finite set of signals Σ and a collection of conditional probabilities $\{s(\cdot|\omega)\}_{\omega \in \Omega}$ over Σ, where $s(\sigma|\omega)$ is the probability that signal $\sigma \in \Sigma$ realizes conditional on state $\omega \in \Omega$.

The timing of the game unfolds as follows. The information designer commits to an experiment. Then, the public signal is realized according to the experiment and the latent underlying state. In response to the signal, firms simultaneously set prices p_A and p_B. Consumers, observing both the signal and the prices, make their purchase decisions. We characterize the subgame perfect equilibria of the game.

3 Analysis

We begin by characterizing equilibria in the pricing subgame for a given experiment and a realized public signal about the firms' qualities. Then, using backward induction, we

[5]When the information designer's objective is a convex combination of *(i)* and *(ii)*, the results follow directly from the analysis of the boundary cases, as we show that consumer surplus can be represented as an affine function of industry profits.

characterize the necessary and sufficient conditions for an experiment to maximize total industry profits and consumer surplus, respectively.

3.1 Price Setting Stage

Consider the subgame in which the information designer has chosen an experiment $\mathcal{E} = \{s(\sigma|\omega)\}_{(\sigma,\omega)}$ and a public signal $\sigma \in \Sigma$ has realized. Define $\Delta q_{i,j}(\sigma)$ as the expected quality difference between firm i and firm j conditional on signal σ, i.e.,

$$\Delta q_{i,j}(\sigma) \equiv \mathbb{E}[q_i - q_j|\sigma].$$

If $\Delta q_{i,j}(\sigma) = 0$, then firms compete à la Bertrand and charge zero prices, resulting in zero profits. Otherwise, suppose, without loss of generality, that firm i has a strictly higher expected quality conditional on signal σ than firm j, i.e., $\Delta q_{i,j}(\sigma) > 0$. Let there exist a consumer of type $\hat{\theta} \in [\underline{\theta}, \overline{\theta}]$ who is indifferent between purchasing from firm i and firm j when the firms set prices p_i and p_j, respectively. Then, the indifference condition for this consumer is given by

$$\hat{\theta}\mathbb{E}[q_i|\sigma] - p_i = \hat{\theta}\mathbb{E}[q_j|\sigma] - p_j, \tag{1}$$

which implies that

$$\hat{\theta} = \frac{p_i - p_j}{\Delta q_{i,j}(\sigma)}. \tag{2}$$

Consumers with $\theta \geq \hat{\theta}$ buy from firm i, and the remaining consumers with $\theta < \hat{\theta}$ buy from firm j. The resulting firms' profit functions are:

$$\pi_i(p_i, p_j|\sigma) = \left(1 - F\left(\frac{p_i - p_j}{\Delta q_{i,j}(\sigma)}\right)\right) p_i, \tag{3}$$

$$\pi_j(p_j, p_i|\sigma) = F\left(\frac{p_i - p_j}{\Delta q_{i,j}(\sigma)}\right) p_j. \tag{4}$$

If the indifference condition (1) is solved for some $\theta > \overline{\theta}$, then firm i makes zero profits as all profits accrue to firm j. If instead it is solved for some $\theta < \underline{\theta}$, then all consumers buy from firm i, and firm j earns zero profits.

Since the density function f is log-concave, it follows that the firms' demand functions $F(\cdot)$ and $1 - F(\cdot)$ are log-concave (by Theorems 1 and 2 in Bagnoli and Bergstrom, 2005). Consequently, the profit functions given in equations (3) and (4) are also log-concave, and therefore quasi-concave. Thus, a solution to the system of first-order conditions together with equation (2) determines an equilibrium of the pricing subgame. The following lemma establishes the existence of a unique equilibrium of the pricing subgame for any given experiment $\mathcal{E}$ and realized signal σ. It also characterizes the corresponding equilibrium prices and the equilibrium location of the indifferent consumer.

Lemma 1. *Consider an experiment $\mathcal{E} = \{s(\sigma|\omega)\}_{(\sigma,\omega)}$ and a realized public signal $\sigma \in \Sigma$. If $\Delta q_{i,j}(\sigma) \geq 0$, then in the unique equilibrium of the pricing subgame, firms set prices*

$$p_i^*(\sigma) = \frac{1 - F(\hat{\theta})}{f(\hat{\theta})} \Delta q_{i,j}(\sigma) \quad \textit{and} \quad p_j^*(\sigma) = \frac{F(\hat{\theta})}{f(\hat{\theta})} \Delta q_{i,j}(\sigma),$$

where $\hat{\theta} \in (\underline{\theta}, \overline{\theta})$ uniquely solves

$$\hat{\theta} = \frac{1 - 2F(\hat{\theta})}{f(\hat{\theta})}.$$

Consumers with $\theta \in [\hat{\theta}, \overline{\theta}]$ buy from firm i, and the remaining consumers buy from firm j. Furthermore, $\hat{\theta}$ is independent of the signal σ.

The proof of Lemma 1 can be found in the Appendix. To see why the type of the indifferent consumer, $\hat{\theta}$, is invariant to the expected quality difference, define prices normalized by the expected quality gap, $p_i' = p_i/\Delta q_{i,j}(\sigma)$ and $p_j' = p_j/\Delta q_{i,j}(\sigma)$, as the new strategic variables of the firms. Then, the profit functions of firms i and j can be rewritten as $\Delta q_{i,j}(\sigma)(1 - F(p_i' - p_j'))p_i'$ and $\Delta q_{i,j}(\sigma)F(p_i' - p_j')p_j'$, respectively. Since $\Delta q_{i,j}(\sigma)$ enters the profit functions only as a multiplicative constant, the equilibrium normalized prices p_i' and p_j' do not depend on the expected quality gap. Therefore, $\hat{\theta} = p_i' - p_j'$ is independent of $\Delta q_{i,j}(\sigma)$ in equilibrium, and consequently also independent of the signal σ.

Given that the equilibrium normalized prices do not depend on the expected quality difference, the equilibrium prices, $p_i^*(\sigma)$ and $p_j^*(\sigma)$, increase linearly in $\Delta q_{i,j}(\sigma)$. The same applies to the equilibrium profits, which are given by:

$$\pi_i^*(\sigma) \equiv \frac{(1 - F(\hat{\theta}))^2}{f(\hat{\theta})} \Delta q_{i,j}(\sigma) \quad \text{and} \quad \pi_j^*(\sigma) \equiv \frac{F^2(\hat{\theta})}{f(\hat{\theta})} \Delta q_{i,j}(\sigma).$$

Intuitively, firms benefit from the reduced competitive pressure associated with increased quality differentiation. Indeed, the demand of the firm with higher expected quality becomes less elastic.[6] Furthermore, as price competition under vertical differentiation features strategic complementarity, the price increase of firm i triggers a strategic response of firm j to raise its price as well. Thus, even the firm with lower expected quality benefits from an increasing expected quality difference, $\Delta q_{i,j}(\sigma)$.

The resulting total industry profit is given by

$$\Pi(\sigma) \equiv \pi_A^*(\sigma) + \pi_B^*(\sigma) = \Psi |\Delta q_{A,B}(\sigma)|, \tag{5}$$

[6]The absolute value of the price elasticity of demand of firm i is $|\varepsilon_i| = \frac{f\left(\frac{p_i - p_j}{\Delta q_{i,j}(\sigma)}\right)}{1 - F\left(\frac{p_i - p_j}{\Delta q_{i,j}(\sigma)}\right)} \frac{p_i}{\Delta q_{i,j}(\sigma)}$. Both components are positive for the firm perceived to be of higher quality. As the hazard rate function, $f/(1-F)$, increases in its argument (and thus decreases in $\Delta q_{i,j}(\sigma)$), while the factor $p_i/\Delta q_{i,j}(\sigma)$ decreases as well, the absolute value of the elasticity decreases in $\Delta q_{i,j}(\sigma)$.

where

$$\Psi \equiv \frac{(1-F(\hat{\theta}))^2 + F^2(\hat{\theta})}{f(\hat{\theta})}. \tag{6}$$

Consumer surplus for a given signal σ can be written as follows:

$$CS(\sigma) \equiv \overline{T}\mathbb{E}[q_i|\sigma] + \underline{T}\mathbb{E}[q_j|\sigma] - \Pi(\sigma), \tag{7}$$

where

$$\overline{T} \equiv \int_{\hat{\theta}}^{\overline{\theta}} \theta dF(\theta) \quad \text{and} \quad \underline{T} \equiv \int_{\underline{\theta}}^{\hat{\theta}} \theta dF(\theta).$$

Let us define the ex-ante total industry profit and the ex-ante consumer surplus for a given experiment $\mathcal{E}$ as

$$\Pi^{\mathcal{E}} \equiv \sum_{\omega\in\Omega}\sum_{\sigma\in\Sigma} \lambda_\omega s(\sigma|\omega)\Pi(\sigma) \quad \text{and} \quad CS^{\mathcal{E}} \equiv \sum_{\omega\in\Omega}\sum_{\sigma\in\Sigma} \lambda_\omega s(\sigma|\omega) CS(\sigma).$$

In the following section, we provide separate characterizations of the experiments that maximize ex-ante total industry profit, $\Pi^{\mathcal{E}}$, and ex-ante consumer surplus, $CS^{\mathcal{E}}$.

3.2 Optimal Information Structures

To characterize which information structures are most beneficial for firms and consumers, respectively, we introduce two classes of experiments, namely *rank-revealing* and *rank-preserving* ones.

A *rank-revealing* experiment does not pool states in which firm qualities differ in their order; that is, it does not pool states in which $q_i > q_j$ with states in which $q_i < q_j$. Let $\mathbb{S}$ be the set of all rank-revealing experiments. Formally, we define rank-revealing experiments as follows.

Definition 1 (Rank-revealing signals and experiments)**.** *A signal σ generated by an experiment $\mathcal{E} = \{s(\cdot|\omega)\}_{\omega\in\Omega}$ is rank-revealing if and only if for every $k > \ell$, we have that $s\left(\sigma|\left(q^k, q^\ell\right)\right) > 0$ implies that $s\left(\sigma|\left(q^h, q^m\right)\right) = 0$ for every $h < m$. An experiment $\mathcal{E}$ is rank-revealing, that is, $\mathcal{E} \in \mathbb{S}$, if and only if all signals σ sent with positive probability are rank-revealing.*

Note that the fully informative experiment belongs to $\mathbb{S}$, implying that $\mathbb{S}$ is non-empty. Next, we introduce experiments that are *rank-preserving.* They never overturn the ranking of firms' prior expected qualities. Put differently, if firm i is ex ante perceived to be of strictly higher expected quality than firm j, there cannot be a signal realization $\sigma \in \Sigma$ such that $\mathbb{E}[q_i|\sigma] < \mathbb{E}[q_j|\sigma]$. Let $\mathbb{H}$ be the set of all rank-preserving experiments. Formally, we define rank-preserving experiments as follows.

Definition 2 (Rank-preserving experiments)**.** *An experiment* $\mathcal{E} = \{s(\cdot|\omega)\}_{\omega\in\Omega}$ *is rank-preserving, that is,* $\mathcal{E} \in \mathbb{H}$*, if and only if* $\mathbb{E}[q_i] > \mathbb{E}[q_j]$ *implies that* $\mathbb{E}[q_i|\sigma] \geq \mathbb{E}[q_j|\sigma]$ *for every signal* $\sigma \in \Sigma$*.*

Note that the uninformative experiment belongs to $\mathbb{H}$ by construction, implying that $\mathbb{H}$ is non-empty as well.

Industry profit-maximizing information design. For any rank-revealing experiment, we have by construction that, for all σ,

$$|\Delta q_{A,B}(\sigma)| = |\mathbb{E}[q_A - q_B|\sigma]| = \mathbb{E}[|q_A - q_B|\,|\sigma]. \tag{8}$$

This implies that the expected industry profit under a rank-revealing experiment $\mathcal{E} \in \mathbb{S}$ is given by:

$$\Pi^{\mathcal{E}} = \sum_{\omega\in\Omega}\lambda_\omega \sum_{\sigma\in\Sigma} s(\sigma|\omega)\Psi|\Delta q_{A,B}(\sigma)| \tag{9}$$

$$= \sum_{\omega\in\Omega}\lambda_\omega \sum_{\sigma\in\Sigma} s(\sigma|\omega)\Psi\mathbb{E}\left[|q_A - q_B|\,|\sigma\right] \tag{10}$$

$$= \Psi\mathbb{E}\left|q_A - q_B\right|, \tag{11}$$

where the last step follows from the law of iterated expectations.

We will next establish that any non-rank-revealing experiment attains lower industry profits. Towards this, consider any experiment $\mathcal{E}' \notin \mathbb{S}$, implying that there exists a signal σ' that does not fully reveal firms' quality ranking.[7] Applying Jensen's inequality, we obtain

$$|\Delta q_{A,B}(\sigma')| = |\mathbb{E}[q_A - q_B|\sigma']| < \mathbb{E}[|q_A - q_B|\,|\sigma'], \tag{12}$$

where the strict inequality follows from the fact that there are at least two nonzero terms with opposing signs in the summation over states. Note that (12) holds for any signal σ' that is non-revealing, while any revealing signal σ satisfies (8). Given the existence of at least one non-revealing signal realization, we obtain that

$$\Pi^{\mathcal{E}'} = \Psi\sum_{\omega\in\Omega}\lambda_\omega \sum_{\sigma'\in\Sigma'} s(\sigma|\omega')|\Delta q_{A,B}(\sigma')| \tag{13}$$

$$< \Psi\sum_{\omega\in\Omega}\lambda_\omega \sum_{\sigma'\in\Sigma'} s(\sigma|\omega')\mathbb{E}\left[|q_A - q_B|\,|\sigma'\right] \tag{14}$$

$$= \Psi\mathbb{E}\left|q_A - q_B\right| = \Pi^{\mathcal{E}}, \tag{15}$$

[7] Formally, this implies the existence of k', ℓ', h', m', such that $s'\left(\sigma'|\left(q^{k'}, q^{\ell'}\right)\right) > 0$ and $s'\left(\sigma'|\left(q^{h'}, q^{m'}\right)\right) > 0$ hold simultaneously with $q^{k'} > q^{\ell'}$ and $q^{h'} < q^{m'}$.

which holds for any $\mathcal{E} \in \mathbb{S}$. Given the non-emptiness of $\mathbb{S}$, we can conclude the following result about industry-profit maximizing information structures.

Proposition 1. *An experiment $\mathcal{E}$ maximizes expected industry profits if and only if it is rank-revealing, i.e., $\mathcal{E} \in \mathbb{S}$.*

Consumer surplus-maximizing information design. Consider a signal σ and suppose that $\Delta q_{i,j}(\sigma) \geq 0$. Then, consumer surplus given in equation (7) can be rewritten as follows:

$$CS(\sigma) = \frac{\overline{T} + \underline{T}}{2} \left(\mathbb{E}[q_A|\sigma] + \mathbb{E}[q_B|\sigma]\right) + \frac{\overline{T} - \underline{T}}{2} \left|\mathbb{E}[q_A|\sigma] - \mathbb{E}[q_B|\sigma]\right| - \Psi|\Delta q_{A,B}(\sigma)| \quad (16)$$

$$= \frac{\overline{T} + \underline{T}}{2}\mathbb{E}[q_A + q_B|\sigma] + \frac{\overline{T} - \underline{T}}{2}|\Delta q_{A,B}(\sigma)| - \Psi|\Delta q_{A,B}(\sigma)| \quad (17)$$

This expression allows us to intuitively disentangle the components affecting consumer surplus into allocation and expenditure effects. The first term represents the gross consumption surplus if consumers were randomly allocated across the two products in a uniform manner. The second term represents the change in gross consumption surplus—relative to the uniform random allocation—due to an allocation based on the difference in expected qualities: consumers with a high willingness to pay for quality ($\theta \geq \hat{\theta}$) gain from more frequently purchasing the high-quality product, while consumers with a low willingness to pay for quality ($\theta < \hat{\theta}$) lose, as they purchase the high-quality good less frequently. Finally, the third term captures the total expenditures of consumers.

Based on this formulation of interim consumer surplus given σ, we can express ex-ante consumer surplus as

$$CS^{\mathcal{E}} = \frac{\overline{T} + \underline{T}}{2}\mathbb{E}[q_A + q_B] + \left(\frac{\overline{T} - \underline{T}}{2\Psi} - 1\right)\Pi^{\mathcal{E}}. \quad (18)$$

Thus, the effect of information on consumer surplus is directly related to how information affects total industry profits. Whether consumers' incentives regarding information are aligned with, or are opposed to firms' incentives depends on whether the sign of the factor multiplying profits in (18) is positive or negative.

As $\Psi > 0$, this sign is equivalent to that of $\frac{\overline{T} - \underline{T}}{2} - \Psi$, which captures the marginal effects of increased perceived quality differentiation in (17). Specifically, $\frac{\overline{T} - \underline{T}}{2}$ is the effect on consumers' gross surplus resulting from a marginal change in expected quality difference, while Ψ is the effect on industry profits—here equivalent to consumers' total expenditures—resulting from the same marginal change.

To see this, note that we can rewrite

$$\frac{\overline{T}-\underline{T}}{2}-\Psi=(1-F(\hat{\theta}))\left(\frac{\mathbb{E}[\theta|\theta\geq\hat{\theta}]}{2}-\frac{1-F(\hat{\theta})}{f(\hat{\theta})}\right)$$
$$-F(\hat{\theta})\left(\frac{\mathbb{E}[\theta|\theta\leq\hat{\theta}]}{2}+\frac{F(\hat{\theta})}{f(\hat{\theta})}\right). \tag{19}$$

In the first term of equation (19), $(1-F(\hat{\theta}))$ represents the mass of consumers purchasing the product perceived to be of higher quality. Because $\hat{\theta}$ remains constant following a marginal increase in perceived quality differentiation, the allocation of consumers across firms remains unaffected. Relative to the baseline surplus from the uniformly random allocation in the first component of equation (17), the marginal gain in gross consumer surplus is therefore equal to one half—the share of consumers re-allocated relative to the random allocation—times the expected willingness to pay for the marginal increase in quality of the consumers purchasing the product with higher perceived quality, $\mathbb{E}[\theta|\theta \geq \hat{\theta}]$. However, all consumers who purchase the high-quality product pay a higher price, which marginally increases by $(1-F(\hat{\theta}))/f(\hat{\theta})$.

Analogously, in the second part of equation (19), $F(\hat{\theta})$ is the mass of consumers purchasing the product perceived to be of lower quality, while $\mathbb{E}[\theta|\theta \leq \hat{\theta}]$ is the marginal loss in gross consumer surplus for those consumers reallocated relative to the random allocation, and $F(\hat{\theta})/f(\hat{\theta})$ is the marginal price increase which affects all consumers purchasing the product perceived to be of lower quality.

The following lemma shows that, following increased quality differentiation, the firms can always extract strictly more than the consumers can gain in terms of gross surplus. This implies that consumers' and firms' incentives regarding information provision are opposed.

Lemma 2. *It holds that* $\left(\overline{T}-\underline{T}\right)/2<\Psi$.

It directly follows from (18) and Lemma 2 that an experiment maximizes consumer surplus if and only if it minimizes industry profits. We next establish that any rank-preserving experiment minimizes industry profits. Towards this, consider a rank-preserving experiment $\mathcal{E}$. As the order of posteriors never changes, all $\Delta q_{A,B}(\sigma)$ have the same sign. It therefore holds for the expected industry profit that

$$\Pi^{\mathcal{E}}=\Psi\sum_{\omega\in\Omega}\lambda_{\omega}\sum_{\sigma\in\Sigma}s(\sigma|\omega)\left|\Delta q_{A,B}(\sigma)\right| \tag{20}$$

$$=\Psi\left|\sum_{\omega\in\Omega}\lambda_{\omega}\sum_{\sigma\in\Sigma}s(\sigma|\omega)\Delta q_{A,B}(\sigma)\right| \tag{21}$$

$$=\Psi\left|\mathbb{E}[q_A-q_B]\right|. \tag{22}$$

But for any experiment $\mathcal{E}' \notin \mathbb{H}$, there exists a signal σ' that changes the expected quality ranking. By Jensen's inequality, we therefore have that

$$\Pi^{\mathcal{E}'} = \Psi \sum_{\omega \in \Omega} \lambda_\omega \sum_{\sigma' \in \Sigma'} s(\sigma|\omega') \, |\Delta q_{A,B}(\sigma')| \tag{23}$$
$$= \Psi \sum_{\omega \in \Omega} \lambda_\omega \sum_{\sigma' \in \Sigma'} s(\sigma|\omega') \, |\mathbb{E}[q_A - q_B|\sigma']| \tag{24}$$
$$> \Psi \, |\mathbb{E}[q_A - q_B]| \tag{25}$$
$$= \Pi^{\mathcal{E}}, \tag{26}$$

where $\mathcal{E} \in \mathbb{H}$. As rank-preserving experiments minimize expected industry profits, they maximize expected consumer surplus. Given the non-emptiness of $\mathbb{H}$, we can conclude with the following proposition.

Proposition 2. *An experiment $\mathcal{E}$ maximizes expected consumer surplus if and only if it is rank-preserving, i.e., $\mathcal{E} \in \mathbb{H}$.*

4 Extensions

We discuss several extensions that relax assumptions made in the baseline model and showcase the robustness of the identified mechanisms and results. Specifically, we consider a setting in which the market is not fully covered, which introduces an extensive margin, and settings with alternative cost function specifications. We also discuss how changes in consumer preferences affect the results.

4.1 Allowing for an Extensive Margin

One might expect that our results are weakened if not nullified by allowing for an extensive margin: it limits the extractive power of the firms, as they compete against an elastic outside option rather than only against each other. In this section, we show that this expectation is incorrect. To keep the analysis tractable and focused, we restrict attention to perfectly negatively correlated qualities and uniformly distributed consumer types.

Let $\mathcal{Q} = \{\underline{q}, \bar{q}\}$ with $\underline{q} < \bar{q}$ and $\Omega = \{(\underline{q}, \bar{q}), (\bar{q}, \underline{q})\}$. Denote by λ the probability that $\omega = (\bar{q}, \underline{q})$, i.e., that firm A is of high quality. Further, we assume that $\theta \sim U : [0, 1]$. The value of the outside option is normalized to zero, i.e., $\underline{v} = 0$.

Within this framework, the information-design problem is characterized by choosing a distribution over a single posterior belief μ—the belief that firm one is the high-quality firm—subject to Bayes plausibility. Thus, given a posterior belief μ, consumers expect firm A (firm B) to be of higher quality if $\mu > (<)1/2$.

For any belief μ, an equilibrium in the pricing game exists and is unique (see Benassi et al., 2019, for a general treatment of the pricing game). The following lemma characterizes the unique equilibrium for a given belief μ (for the proof, see Appendix B).

Lemma 3. *For any belief $\mu \in [0, 1]$, the unique equilibrium features prices*

$$p_A^*(\mu) = \begin{cases} 2(\bar{q} - \underline{q})(2\mu - 1)\frac{\mu\bar{q}+(1-\mu)\underline{q}}{(5\mu-1)\bar{q}+(4-5\mu)\underline{q}}, & \text{if } \mu \geq 1/2 \\ (\bar{q} - \underline{q})(1 - 2\mu)\frac{(1-\mu)\bar{q}+\mu\underline{q}}{(4-5\mu)\bar{q}+(5\mu-1)\underline{q}}, & \text{if } \mu < 1/2 \end{cases} \tag{27}$$

$$p_B^*(\mu) = \begin{cases} (\bar{q} - \underline{q})(2\mu - 1)\frac{(1-\mu)\bar{q}+\mu\underline{q}}{(5\mu-1)\bar{q}+(4-5\mu)\underline{q}}, & \text{if } \mu \geq 1/2 \\ 2(\bar{q} - \underline{q})(1 - 2\mu)\frac{\mu\bar{q}+(1-\mu)\underline{q}}{(4-5\mu)\bar{q}+(5\mu-1)\underline{q}}, & \text{if } \mu < 1/2. \end{cases} \tag{28}$$

Industry Profits. Consider the total industry profits conditional on a posterior μ as $\Pi(\mu) = \pi_A^*(\mu) + \pi_B^*(\mu)$, where $\pi_i^*(\mu)$ are the profits induced by the equilibrium pricing strategies defined in Lemma 3. We derive the explicit expression for the corresponding industry profits in Appendix B. Here, we only observe that the industry profits are necessarily symmetric around $\mu = 1/2$ (as the model is symmetric in the firm identity). Moreover, industry profits are strictly increasing towards extreme beliefs—that is, strictly decreasing in μ for $\mu \in [0, 1/2)$ and strictly increasing in μ for $\mu \in (1/2, 1]$. Thus, the industry profit function has two (global) maxima at $\mu = 0$ and $\mu = 1$. It follows immediately that industry profits are maximized under full information revelation.

Proposition 3. *The perfectly informative experiment, which is the only rank-revealing experiment, uniquely maximizes industry profits when there is an extensive margin.*

As in the previous section, in which market demand was completely inelastic so that all consumers purchased, industry profits were maximized for rank-revealing information structures. In the specific case with perfectly negatively correlated binary qualities, the only rank-revealing information structure is full information revelation. The intuition follows again from the differentiation incentive of the firms. The more precisely informed the consumers are, the more differentiated they perceive the firms' products, and hence, the less competitive the market becomes.

Consumer Surplus. As before, we can express consumer surplus as being composed of three components: the gross surplus of consumers purchasing the perceived low-quality product, the gross surplus of consumers purchasing the perceived high-quality product, and the industry profits, which need to be subtracted from the first two components.

Building on the equilibrium prices in Lemma 3, we can obtain the consumer surplus function, which also has to be symmetric around $\mu = 1/2$. We provide the expression in Appendix B. We note here that consumer surplus strictly decreases as the belief approaches the extremes. Thus, it has a global maximum at $\mu = 1/2$ and two global minima

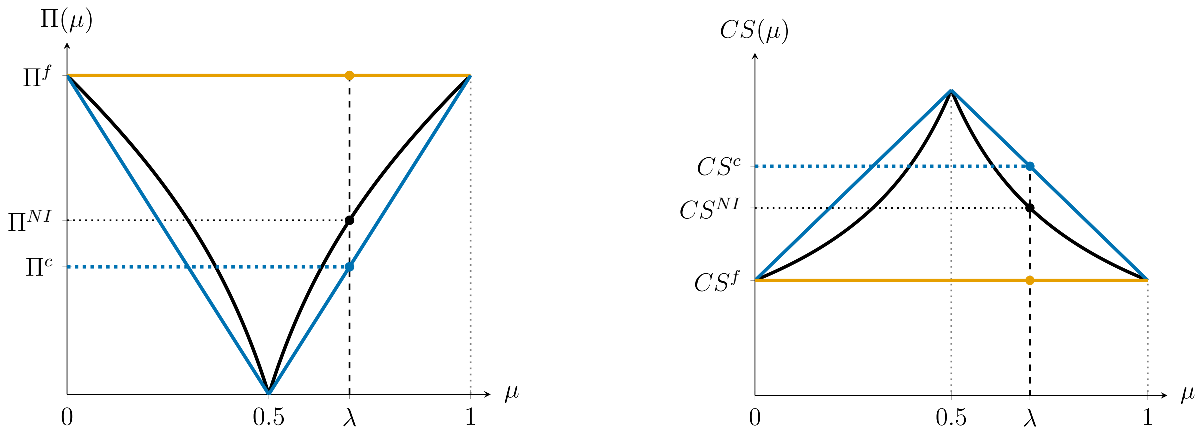

Figure 1: Industry profits and consumer surplus with an extensive margin.

The left panel depicts the industry profits $\Pi(\mu)$ in black. The orange line depicts the concavification of $\Pi(\mu)$. The blue line corresponds to the convexification of $\Pi(\mu)$. The right panel depicts the consumer surplus $CS(\mu)$. The blue line depicts the concavification $CS(\mu)$. The orange line depicts the convexification of $CS(\mu)$.
The graphs highlight three policies and their corresponding payoffs. (i) The orange points are associated with the industry-profit maximizing, fully informative policy (inducing Π^f, CS^f); (ii) the blue points are associated with the consumer-optimal maximally-informative, rank-preserving policy (inducing Π^c, CS^c); and (iii) the black points are associated with no information revelation (inducing Π^{NI}, CS^{NI}).

at $\mu = 0$ and $\mu = 1$. It follows that consumer surplus is minimized under the fully informative structure that maximizes industry profits. It turns out that the consumer surplus function is strictly convex in the belief within each subinterval $\mu \in [0, 1/2)$ and $\mu \in (1/2, 1]$. Therefore, the consumer-optimal information structure is such that it is (i) rank-preserving, and (ii) the most informative rank-preserving information structure.

If firm A is the a priori higher-quality firm (i.e., if $\lambda > 1/2$), this implies that consumers never believe that firm B is the higher-quality firm after any signal realization, and that the consumer-optimal experiment reveals firm A to be the high-quality firm sometimes and otherwise makes consumers indifferent between the two firms.

Proposition 4. *With an extensive margin, consumer surplus is uniquely maximized by the most informative rank-preserving experiment.*

There are two forces behind this result. First, as in the preceding analysis without the extensive margin, consumers benefit from not learning the ranking of the firms due to more intense price competition. Second, the consumers nevertheless benefit from some information revelation, as it affects participation in the market. Without an extensive margin, all consumers purchase under any rank-preserving information structure. With an intensive margin, however, the purchasing decision becomes relevant, and with that marginal information. Hence, among those information structures that lead to intense competition, more information benefits consumers by encouraging efficient selection of consumers into purchasing. Indeed, the existence of an outside option limits the rent-extraction ability of firms when their market power increases due to differentiation. The firms now have to balance inframarginal gains from higher prices against marginal losses due to consumers' non-participation. This new effect limits the losses in consumer surplus as firms become more differentiated, causing the consumer surplus function to become convex. This convexity, in turn, introduces the marginal benefit of information for consumers by limiting losses after a firm is perfectly revealed to be of high quality, but retaining the benefits of fierce competition after the uninformative signal is observed.

Note that due to the strict local concavity of industry profits in the rank-preserving regions, the consumer-optimal information structure minimizes industry profits. Similarly, that consumer surplus is decreasing toward extreme beliefs implies that the profit-maximizing, perfectly revealing information structure minimizes consumer surplus. Thus, we retain perfectly opposing incentives regarding information between firms and consumers in the model with an extensive margin.

4.2 Alternative cost function specifications

We next discuss the robustness of our findings to alternative cost function specifications. Specifically, we consider two deviations from our baseline model: non-linear (convex) cost functions and asymmetric cost functions. We analyze these specifications in detail in Appendices C.2 and C.3. As with an extensive margin, the marginal consumer's type is no longer necessarily independent of the revealed information. Instead, it becomes endogenous to the information structure, which substantially complicates the analysis. However, we show that, in both settings, the analyses boil down to one-dimensional information design problems, which depend only on the posterior mean of *quality differences*. This allows us to characterize profit- and consumer-surplus-maximizing information design by leveraging insights from the recent literature on duality results in information design (see, e.g., Dworczak and Martini, 2019; Dizdar and Kováč, 2020; Arieli et al., 2023).

Non-linear production costs In Appendix C.2, we study a setting in which firms have symmetric convex production costs given by $C(x) = \frac{c}{2}x^2$, and in which consumer tastes are distributed according to the uniform distribution, $\theta \sim U[0,1]$. For any experiment, the equilibrium industry profits are strictly convex in the absolute value of the posterior mean of the quality difference, which we illustrate in Figure C.1. Thus, industry profits are maximized when the quality differential is fully revealed. We formalize this observation in Proposition C.1.

This result showcases that the mechanism identified in the baseline analysis—firms benefit from information revelation that increases the perceived differentiation—carries over. It is no longer sufficient, however, to reveal information about the ranking of firms' qualities alone. Instead, firms benefit from the additional revelation of the cardinal difference. This is because more precise information enables firms to more efficiently resolve the trade-off between allocative efficiency and increasing marginal production costs. We furthermore show in Appendix C.2 that this effect similarly dominates in terms of total welfare, which is maximized under full revelation of the quality differential (see Proposition C.2).

Regarding consumer surplus, we find that, as before, consumer surplus is single-peaked in the posterior quality gap, with the peak attained when firms are perceived as symmetric $\Delta q_{A,B} = 0$. Moreover, it is convex and increasing for $\Delta q_{A,B} < 0$ and convex and decreasing for $\Delta q_{A,B} > 0$. From this characterization, we can deduce the qualitative properties of the consumer-optimal information policy directly, which we formally establish in Proposition C.3. First, it is straightforward that consumer surplus is maximized under the non-informative experiment whenever the prior belief is that firms' expected qualities are identical, that is, whenever the firms are ex-ante symmetric. Otherwise, the optimal experiment is a censorship experiment in the sense of Kolotilin et al. (2022), in which the quality gap $\Delta q_{A,B}$ is fully revealed on one side of a cutoff quality gap, while realizations on the other side are pooled together into symmetric posterior quality beliefs.

In particular, this implies that the consumer-optimal experiment remains rank-preserving. Only states that rank firms according to the prior are fully revealed and the pooling signal is constructed in such a way that the expected quality gap is zero conditional on observing the pooling signal. Moreover, the tension between what is beneficial for firms and consumers continues to apply.

Asymmetric production costs Towards asymmetric production costs, we consider a setting in which firms have asymmetric but constant marginal costs c_A and c_B with $c_A > c_B \geq 0$. In Appendix C.3, we derive the equilibrium of the pricing subgame and characterize industry profits and consumer welfare as a function of the posterior expected quality difference. Restricting attention to uniformly distributed consumer tastes for quality, $\theta \sim U[0, 1]$, we show that industry profits are maximized by an experiment that pools interior—in terms of the underlying expected quality difference—states for which the cost difference dominates the quality difference, while states outside this interior interval are fully revealed (see Proposition C.1). This follows from a convex-concave-convex profit shape, where the concave shape for quality gaps around zero (that is, symmetric beliefs about the firms' qualities) results from monopolization due to more efficient production technologies.[8] We obtain the characterization of the industry-profit maximizing experiment using the techniques developed in Dworczak and Martini (2019) adapted to the finite state case (see Dizdar and Kováč, 2020).

These results are economically intuitive. Pooling states in which the cost asymmetry dominates avoids belief configurations in which firms compete heavily. This would occur when the cost and quality differences offset each other, that is, when the low-cost firm has

[8]It should be noted that analogous convex-concave-convex shapes and characterizations of optimal experiments can be obtained without monopolization in a model with asymmetric quadratic cost functions. Hence, the results do not rely on the monopolization region but on competition between firms with asymmetric cost structures when their product qualities are relatively symmetric.

a lower expected quality. In addition, the revelation of the full quality difference improves the allocative efficiency and thus provides an increased opportunity for rent extraction. Notably, the results reinforce the findings of our baseline setting: firms benefit when consumers perceive them as differentiated. Differentiation in turn can arise both from different expected qualities and from the underlying production cost differences when quality differences are comparatively small.

Consumers, in contrast, continue to benefit from experiments that limit information about firm qualities. While the analysis becomes substantially more complicated, the non-informative experiment continues to maximize consumer surplus if the prior is such that firms would be induced to compete fiercely. Crucially, the consumer-optimal experiment is rank-preserving for a wide range of prior expectations about the quality gap between firms, as we establish in Proposition C.2, and always features pooling of non-degenerate sets of quality differences. This once again showcases that consumers benefit from limits to the revealed information.

4.3 Non-multiplicative utility functions

Finally, we consider departures from the utility specification in which consumers' willingness to pay for quality and quality itself interact multiplicatively. While this specification is standard in the literature (dating back to Mussa and Rosen, 1978), it is a crucial ingredient to our analysis as it implies the Shaked and Sutton (1982)-type independence between the marginal consumer's type and firms' quality perceptions in the baseline setting. As such, it is a natural question to explore what happens under non-multiplicative utility specifications.

Towards this, we establish in Appendix D that the marginal consumer type $\hat{\theta}$ is independent of the information structure if and only if the utility function is linear in quality, up to a monotone transformation of quality. Under a non-multiplicative setting, the analysis therefore becomes substantially more complicated. In particular, it is not a priori clear whether the marginal type can even be represented as a function of the expected quality difference. If this were not the case, the problem would not be a one-dimensional information design problem as in the case of non-linear and asymmetric costs, but instead become an information design problem in which the full posterior distribution of qualities matters.[9]

Overall, our representation result in Proposition D.1 offers a sharp and informative characterization of why the widely used assumption of multiplicative consumer preferences is so powerful in the vertical differentiation literature. At the same time, we view it as

[9]It is straightforward to establish that this complication necessarily arises in settings with continuous qualities: there, the marginal type $\hat{\theta}$ can never be represented as a function of the expected quality difference, as it is either independent of the firms' qualities—as in our baseline model—or depends on the full posterior distribution of qualities.

an invitation for future research to develop new tools and insights for richer settings that move beyond this benchmark.

5 Concluding Remarks

We analyze a model in which duopolistic sellers offer vertically differentiated products to heterogeneous buyers. A third party provides public information about firms' product qualities. This information helps consumers make better decisions about which product to purchase, but also allows firms to adjust their prices to updated consumer beliefs. The latter is particularly important, as information may relax competition by increasing perceived product differentiation.

We show that, in a fully covered market, the primary criterion determining the welfare consequences of information structures is their content regarding product rankings. In this setting, it is therefore not necessary to provide informational content regarding specific quality levels or quality differences.

Our main finding is that buyer-optimal and seller-optimal information structures are opposite to each other. Sellers prefer the quality ranking of their products to be disclosed, while buyers prefer the prior ranking to be preserved even though it may be ex post incorrect. These results obtain both when the market is fully covered and when an extensive margin is considered. In the latter case, the consumer-optimal information structure is such that it is maximally informative subject to preserving the prior ranking. We also show that the tension between consumer-optimal and firm-optimal information structures extends to settings with non-linear or asymmetric production costs.

Our results have important implications for the design of third-party public information, such as search rankings and ratings on platform markets, or information disseminated by consumer organizations, certifiers, or regulators. More generally, they raise questions about consumer-induced certification of product qualities or rankings.

Appendix

A Proofs

Proof of Lemma 1. Assume for a contradiction that there is no indifferent consumer in equilibrium. Then, a firm serving the entire market can slightly increase its price and achieve greater profits, leading to a contradiction. Therefore, in any equilibrium, there exists an indifferent consumer.

Suppose that firms set prices p_i and p_j and there is an indifferent consumer $\hat{\theta} = (p_i - p_j)/\Delta q_{i,j}(\sigma) \in [\underline{\theta}, \overline{\theta}]$. In the main text we established that the profit functions given in equations (3) and (4) are quasi-concave and therefore it is sufficient to consider the system of the first-order conditions given by

$$\begin{aligned} 0 &= 1 - F(\hat{\theta}) - \frac{f(\hat{\theta})p_i}{\Delta q_{i,j}(\sigma)} \\ 0 &= F(\hat{\theta}) - \frac{f(\hat{\theta})p_j}{\Delta q_{i,j}(\sigma)}. \end{aligned}$$

Solving the system with respect to p_i and p_j, we obtain

$$p_i = \frac{1 - F(\hat{\theta})}{f(\hat{\theta})}\Delta q_{i,j}(\sigma) \quad \text{and} \quad p_j = \frac{F(\hat{\theta})}{f(\hat{\theta})}\Delta q_{i,j}(\sigma).$$

Plugging the expressions for p_i and p_j into $\hat{\theta} = (p_i - p_j)/\Delta q_{i,j}(\sigma)$, we find that the type of the indifferent consumer, $\hat{\theta}$, satisfies the following equation

$$\hat{\theta} - \frac{1 - 2F(\hat{\theta})}{f(\hat{\theta})} = 0. \tag{29}$$

As f is a log-concave function, $(1 - F)/f$ is a non-increasing function and F/f is a non-decreasing function. Thus, the expression in equation (29) is non-decreasing in $\hat{\theta}$. Since $\underline{\theta} < 1/f(\underline{\theta})$, the intermediate value theorem guarantees the existence of a unique $\hat{\theta} \in (\underline{\theta}, \overline{\theta})$ that solves equation (29). □

Proof of Lemma 2. Plugging in the expressions for $\overline{T}, \underline{T}, \Psi$, we have:

$$\overline{T} - \underline{T} - 2\Psi = \int_{\hat{\theta}}^{\overline{\theta}} \theta dF(\theta) - \int_{\underline{\theta}}^{\hat{\theta}} \theta dF(\theta) - 2\frac{(1 - F(\hat{\theta}))^2 + F^2(\hat{\theta})}{f(\hat{\theta})}. \tag{30}$$

Integrating by parts the first two integrals, we obtain:

$$\begin{aligned}\overline{T}-\underline{T}-2\Psi &= \hat{\theta}-\hat{\theta}F(\hat{\theta})+\int_{\hat{\theta}}^{\overline{\theta}}(1-F(\theta))d\theta \\ &\quad -\left(\hat{\theta}F(\hat{\theta})-\int_{\underline{\theta}}^{\hat{\theta}}F(\theta)d\theta\right)-2\frac{(1-F(\hat{\theta}))^2+F^2(\hat{\theta})}{f(\hat{\theta})} \qquad (31)\\ &= \int_{\underline{\theta}}^{\hat{\theta}}F(\theta)d\theta+\int_{\hat{\theta}}^{\overline{\theta}}(1-F(\theta))d\theta+\hat{\theta}(1-2F(\hat{\theta}))-2\frac{(1-F(\hat{\theta}))^2+F^2(\hat{\theta})}{f(\hat{\theta})}. \qquad (32)\end{aligned}$$

Plugging in $\hat{\theta}=(1-2F(\hat{\theta}))/f(\hat{\theta})$ into the third term, we obtain:

$$\begin{aligned}\overline{T}-\underline{T}-2\Psi &= \int_{\underline{\theta}}^{\hat{\theta}}F(\theta)d\theta+\int_{\hat{\theta}}^{\overline{\theta}}(1-F(\theta))d\theta+\frac{(1-2F(\hat{\theta}))^2}{f(\hat{\theta})}-2\frac{(1-F(\hat{\theta}))^2+F^2(\hat{\theta})}{f(\hat{\theta})} \qquad (33)\\ &= \int_{\underline{\theta}}^{\hat{\theta}}F(\theta)d\theta+\int_{\hat{\theta}}^{\overline{\theta}}(1-F(\theta))d\theta-\frac{1}{f(\hat{\theta})}. \qquad (34)\end{aligned}$$

By Theorem 1 and Theorem 3 of Bagnoli and Bergstrom (2005), log-concavity of the density function $f(\theta)$ implies that the functions $G(\theta)\equiv\int_{\underline{\theta}}^{\theta}F(t)dt$ and $H(\theta)\equiv\int_{\underline{\theta}}^{\theta}(1-F(t))dt$ are also log-concave on $[\underline{\theta},\overline{\theta}]$. By the definition of log-concavity, we have that $GG''-(G')^2\leq 0$ and $HH''-(H')^2\leq 0$, implying that

$$\begin{aligned}\int_{\underline{\theta}}^{\hat{\theta}}F(\theta)d\theta &\leq \frac{F^2(\hat{\theta})}{f(\hat{\theta})}, \qquad (35)\\ \int_{\hat{\theta}}^{\overline{\theta}}(1-F(\theta))d\theta &\leq \frac{(1-F(\hat{\theta}))^2}{f(\hat{\theta})}. \qquad (36)\end{aligned}$$

Applying these inequalities, we finally have that

$$\begin{aligned}\overline{T}-\underline{T}-2\Psi &\leq \frac{1}{f(\hat{\theta})}\left((1-F(\hat{\theta}))^2+F^2(\hat{\theta})-1\right) \qquad (37)\\ &= -\frac{2}{f(\hat{\theta})}(1-F(\hat{\theta}))F(\hat{\theta}) \qquad (38)\\ &< 0, \qquad (39)\end{aligned}$$

where we used the fact that $\theta\in(\underline{\theta},\overline{\theta})$ to obtain the last strict inequality. □

B Extensive Margin

Proof of Lemma 3. Suppose that $\mu > 1/2$. Then, firm A is the high-quality firm in expectation. With an extensive margin, there will be two marginal consumers, $\hat{\theta}_h$ and $\hat{\theta}_l$, both in $[0,1]$. Existence and uniqueness of equilibrium follow from Benassi et al. (2019). We conjecture that the equilibrium is such that $\hat{\theta}_h > \hat{\theta}_l \geq 0$, and verify it in the following. The marginal consumers define the firms' demand functions:

$$D_A(\hat{\theta}_h) = 1 - \hat{\theta}_h \tag{40}$$

$$D_B(\hat{\theta}_h, \hat{\theta}_l) = \hat{\theta}_h - \hat{\theta}_l. \tag{41}$$

Further, the marginal consumers are implicitly defined by the indifference conditions

$$\hat{\theta}_h(\mu\bar{q} + (1-\mu)\underline{q}) - p_A = \hat{\theta}_h((1-\mu)\bar{q} + \mu\underline{q}) - p_B \tag{42}$$

$$\hat{\theta}_l((1-\mu)\bar{q} + \mu\underline{q}) - p_B = 0. \tag{43}$$

From here, we obtain the explicit expressions for the marginal consumers

$$\hat{\theta}_h = \frac{p_A - p_B}{(2\mu - 1)(\bar{q} - \underline{q})} \tag{44}$$

$$\hat{\theta}_l = \frac{p_B}{(1-\mu)\bar{q} + \mu\underline{q}} \tag{45}$$

and the corresponding profit functions

$$\pi_A(p_A, p_B; \mu) = \left(1 - \frac{p_A - p_B}{(2\mu - 1)(\bar{q} - \underline{q})}\right) p_A \tag{46}$$

$$\pi_B(p_A, p_B; \mu) = \left(\frac{p_A - p_B}{(2\mu - 1)(\bar{q} - \underline{q})} - \frac{p_B}{(1-\mu)\bar{q} + \mu\underline{q}}\right) p_B. \tag{47}$$

The profit functions π_i are strictly concave in their own price p_i for any p_{-i}. Hence, first-order conditions are sufficient to find the unique best reply functions:

$$p_A^{BR}(p_B) = \frac{p_B + (2\mu - 1)(\bar{q} - \underline{q})}{2} \tag{48}$$

$$p_B^{BR}(p_A) = p_A \frac{(1-\mu)\bar{q} + \mu\underline{q}}{2(\mu\bar{q} + (1-\mu)\underline{q})}. \tag{49}$$

Solving this system of equations yields the prices stated in Lemma 3. By equilibrium uniqueness, the analogous argument for $\mu < 1/2$, and the standard Bertrand price competition logic for $\mu = 1/2$, the characterization follows.

For later reference, we obtain as equilibrium profits:

$$\pi_A^*(\mu) = \begin{cases} 4(2\mu-1)(\bar{q}-\underline{q})\left(\frac{\mu\bar{q}+(1-\mu)\underline{q}}{(5\mu-1)\bar{q}+(4-5\mu)\underline{q}}\right)^2, & \text{if } \mu \geq 1/2 \\ (1-2\mu)(\bar{q}-\underline{q})\frac{(\mu\bar{q}+(1-\mu)\underline{q})((1-\mu)\bar{q}+\mu\underline{q})}{\left((4-5\mu)\bar{q}+(5\mu-1)\underline{q}\right)^2}, & \text{if } \mu < 1/2 \end{cases}$$

$$\pi_B^*(\mu) = \begin{cases} (2\mu-1)(\bar{q}-\underline{q})\frac{(\mu\bar{q}+(1-\mu)\underline{q})((1-\mu)\bar{q}+\mu\underline{q})}{\left((5\mu-1)\bar{q}+(4-5\mu)\underline{q}\right)^2}, & \text{if } \mu \geq 1/2 \\ 4(1-2\mu)(\bar{q}-\underline{q})\left(\frac{(1-\mu)\bar{q}+\mu\underline{q}}{(4-5\mu)\bar{q}+(5\mu-1)\underline{q}}\right)^2, & \text{if } \mu < 1/2 \end{cases}$$

$$\Pi(\mu) = \begin{cases} (\bar{q}-\underline{q})(2\mu-1)(\mu\bar{q}+(1-\mu)\underline{q})\frac{\bar{q}+\underline{q}+3(\mu\bar{q}+(1-\mu)\underline{q}))}{\left((1-5\mu)\bar{q}+(5\mu-4)\underline{q}\right)^2}, & \text{if } \mu \geq 1/2 \\ (\bar{q}-\underline{q})(1-2\mu)((1-\mu)\bar{q}+\mu\underline{q})\frac{\bar{q}+\underline{q}+3((1-\mu)\bar{q}+\mu\underline{q}))}{\left((5\mu-4)\bar{q}+(1-5\mu)\underline{q}\right)^2}, & \text{if } \mu < 1/2 \end{cases} \tag{50}$$

For consumer surplus, we obtain in equilibrium

$$\begin{aligned} CS(\mu) &= \begin{cases} (1-\hat{\theta}_h)\frac{1+\hat{\theta}_h}{2}\left(\mu\bar{q}+(1-\mu)\underline{q}\right)+(\hat{\theta}_h-\hat{\theta}_l)\frac{\hat{\theta}_h+\hat{\theta}_l}{2}\left(\mu\underline{q}+(1-\mu)\bar{q}\right), & \text{if } \mu \geq 1/2 \\ (1-\hat{\theta}_h)\frac{1+\hat{\theta}_h}{2}\left((1-\mu)\bar{q}+\mu\underline{q}\right)+(\hat{\theta}_h-\hat{\theta}_l)\frac{\hat{\theta}_h+\hat{\theta}_l}{2}\left((1-\mu)\underline{q}+\mu\bar{q}\right), & \text{if } \mu < 1/2 \end{cases} \\ &\quad - \Pi(\mu) \\ &= \begin{cases} \frac{(\mu\bar{q}+(1-\mu)\underline{q})^2((5-\mu)\bar{q}+(4+\mu)\underline{q})}{2\left(\bar{q}(1-5\mu)+(5\mu-4)\underline{q}\right)^2}, & \text{if } \mu \geq 1/2 \\ \frac{((1-\mu)\bar{q}+\mu\underline{q})^2((4+\mu)\bar{q}+(5-\mu)\underline{q})}{2\left(\bar{q}(5\mu-4)+(1-5\mu)\underline{q}\right)^2}, & \text{if } \mu < 1/2. \end{cases} \end{aligned} \tag{51}$$

□

Proof of Proposition 3. We first show that industry profits are increasing and strictly concave in μ for $\mu \in (1/2, 1]$. Recall that the industry profits in this case are (see (50))

$$\Pi(\mu) = (\bar{q}-\underline{q})(2\mu-1)(\mu\bar{q}+(1-\mu)\underline{q})\frac{\bar{q}+\underline{q}+3(\mu\bar{q}+(1-\mu)\underline{q}))}{\left((1-5\mu)\bar{q}+(5\mu-4)\underline{q}\right)^2}. \tag{52}$$

For our purposes, we can ignore the factor $(\bar{q}-\underline{q})$. For convenience, denote $r \equiv \underline{q}/\bar{q} \in (0,1)$, $A(\mu) \equiv r+(1-r)\mu$ and $d(\mu) \equiv 1-4r-5(1-r)\mu$. Factoring out $(\bar{q})^2$ in the numerator and denominator, we obtain

$$\frac{\Pi(\mu)}{\bar{q}-\underline{q}} = \frac{(2\mu-1)(A(\mu))(1+r+3(A(\mu))))}{d(\mu)^2}. \tag{53}$$

Note that $d'(\mu) = -5(1-r) < 0$ as well as $d(\mu = 1/2) = -\frac{3}{2}(1+r) < 0$. Hence, the denominator is strictly positive and the function is smooth for all $\mu \in (1/2, 1]$.

Defining $B(\mu) \equiv 1+4r+3(1-r)\mu$ and $C(\mu) \equiv 1+7r+6(1-r)\mu$, straightforward computations deliver for the first and second derivatives

$$\left(\frac{\Pi(\mu)}{\bar{q}-\underline{q}}\right)' = \frac{2d(\mu)A(\mu)B(\mu) + (1-r)(2\mu-1)(d(\mu)C(\mu) + 10A(\mu)B(\mu))}{d(\mu)^3} \tag{54}$$

$$\left(\frac{\Pi(\mu)}{\bar{q}-\underline{q}}\right)'' = -\frac{2(1-r)(1+r)^2\,(17(1-r)\mu + 11 + 28r)}{d(\mu)^4} < 0, \tag{55}$$

where the sign of the second derivative follows directly by noting that all terms in the numerator are strictly positive and that the denominator has an even power. Thus, industry profits are strictly concave in the belief μ for all $\mu \in (1/2, 1]$.

Given the concavity of industry profits, we only need to show that the first derivative at $\mu = 1$ is positive to conclude that industry profits are strictly increasing in μ for $\mu \in (1/2, 1]$.

Note the following evaluations

$$A(\mu = 1) = 1, \quad B(\mu = 1) = 4 + r, \quad C(\mu = 1) = 7 + r, \quad d(\mu = 1) = -4 + r,$$

and conclude that for $\mu = 1$ all terms in both the numerator and the denominator are finite and non-zero. Evaluating the overall expression at $\mu = 1$ then yields

$$\left.\left(\frac{\Pi(\mu)}{\bar{q}-\underline{q}}\right)'\right|_{\mu=1} = \frac{20 - r(1 - r(10 + r))}{(4-r)^3} > 0. \tag{56}$$

Hence, we can conclude due to concavity that industry profits are increasing in μ for $\mu \in (1/2, 1]$.

The analogous reasoning by symmetry yields that the industry profits are strictly decreasing and concave in μ for all $\mu \in [0, 1/2)$. A simple computation shows that the left- and right-limit at $\mu = 1/2$ coincide. Thus, $\Pi(\mu)$ is continuous on its entire domain, and we obtain $\Pi(\mu = 1/2) = 0$.

These properties imply that industry profits have to global maxima at $\mu = 0$ and $\mu = 1$ and a global minimum at $\mu = 1/2$. Thus, the concavification of industry profits is the horizontal line segment connecting $(0, \Pi(0))$ and $(1, \Pi(1))$. It follows from the standard concavification argument, as in e.g., Kamenica and Gentzkow (2011); Aumann et al. (1995) that the optimal policy is fully revealing and induces the posterior beliefs $\mu = 0$ and $\mu = 1$ (with probability $(1-\lambda)$ and λ, respectively). □

Proof of Proposition 4. We again derive properties of the consumer surplus function $CS(\mu)$ from (51) for the case of $\mu \in (1/2, 1]$ and the analogous properties for the case $\mu \in [0, 1/2)$ follow by symmetry. Recall $A(\mu) = r + (1-r)\mu$ and $d(\mu) = 1 - 4r - 5(1-r)\mu$

and let $E(\mu) \equiv 5 + 4r - (1-r)\mu$. Then, we obtain[10]

$$CS(\mu) = \frac{A(\mu)^2 E(\mu)}{2d(\mu)^2}\bar{q} \tag{57}$$

$$CS'(\mu) = \frac{(1-r)A(\mu)\big[d(\mu)(2E(\mu) - A(\mu)) + 10A(\mu)E(\mu)\big]}{2d(\mu)^3}\bar{q} \tag{58}$$

$$= \frac{(1-r)(\mu + r - \mu r)P(\mu, r)}{2\,d(\mu)^3}\bar{q} < 0, \text{ where} \tag{59}$$

$$P(\mu, r) = 10(1 + r + \mu r(1-\mu)) + 5\mu^2(1 + r^2) + 7r(1 - \mu r) + 12r^2 - 3\mu > 0 \tag{60}$$

$$CS''(\mu) = \frac{(1-r)^2(1+r)^2\big[47(\mu(1-r) + r) + 5(1+r)\big]}{d(\mu)^4}\bar{q} > 0, \tag{61}$$

where the sign of the inequality in the third line follows from $d(\mu) < 0$ and the odd power in the denominator.

Consumer surplus is continuous at $\mu = 1/2$, as the left- and right-limits at $\mu = 1/2$ coincide. Hence, it follows that consumer surplus is strictly increasing and strictly convex in μ for all $\mu \in [0, 1/2)$ and strictly decreasing and strictly convex in μ for all $\mu \in (1/2, 1]$. Further, consumer surplus attains a global maximum at $\mu = 1/2$ and two global minima at $\mu = 0$ and $\mu = 1/2$. Hence, its concavification is given by the two line segments from $(0, CS(0))$ to $(1/2, CS(1/2))$ and $(1/2, CS(1/2))$ to $(1, CS(1))$. Inspecting the concavification at $\mu = \lambda$ for $\lambda > 1/2$ pins down the maximal consumer surplus, which is supported by an information structure that induces the posterior $\mu = 1/2$ and $\mu = 1$ subject to Bayes' plausibility.[11] Hence, the optimal experiment is the most informative rank-preserving information structure. In particular, if firm A is of high quality, it is revealed as such with probability $\frac{2\lambda-1}{\lambda}$ and pooled with firm B with probability $\frac{1-\lambda}{\lambda}$. firm B is never recommended as the higher-quality firm. □

C Alternative Cost Functions

In this Appendix, we study variations of our baseline model with respect to the production costs. We rely on technical results from the information design literature, which we summarize in Appendix C.1. Appendix C.2 analyzes the case of nonlinear symmetric production costs, while Appendix C.3 studies the setting with linear asymmetric production costs.

[10]Note that on their respective supports $\mu \in (1/2, 1]$, $d(\mu)$ is strictly negative, and that $A(\mu)$ and $E(\mu)$ are strictly positive, implying that consumer surplus is strictly positive and smooth on its support.

[11]The case $\lambda < 1/2$ is analogous.

C.1 Preliminaries: Price Functions and Optimal Experiments

In both alternative cost specifications studied in Appendices C.2 and C.3, we show that the objective function of the information designer, who maximizes either industry profits or consumer surplus, can be expressed as a function of the posterior mean of the quality difference. Therefore, to characterize optimal experiments, we can employ the duality approach developed by Dworczak and Martini (2019); Dizdar and Kováč (2020). In this Appendix, we outline this approach using the notation of our paper.

We denote realized quality differences by $\delta \equiv q_A - q_B$ with $\delta \in \mathcal{D} \equiv \{q^i - q^j : \ q^i, q^j \in \mathcal{Q}\}$, where we arrange the elements in increasing order and write $\mathcal{D} = \{\delta_1, \ldots, \delta_m\}$, the convex hull of the support by $\overline{\mathcal{D}} \equiv co(\mathcal{D}) = [\delta_1, \delta_m]$ and the probability mass function of δ at $\delta_0 \in \mathcal{D}$ by $\lambda^\delta(\delta_0) \equiv \sum_{(i,j):q^i - q^j = \delta_0} \lambda_{ij}$.

A sender chooses a public experiment generating a signal σ. We assume that the sender's payoff function depends solely on the posterior mean $\Delta q(\sigma) \equiv \mathbb{E}[\delta \mid \sigma]$, and we denote it by $v = v(\Delta q(\sigma)) : \overline{\mathcal{D}} \to \mathbb{R}$. The choice of an experiment is equivalent to choosing a distribution G of posterior means Δq supported on $\overline{\mathcal{D}}$ such that G is a mean-preserving contraction (MPC) of λ^δ. Thus, the optimal experiment solves

$$\sup_{G \in \mathrm{MPC}(\lambda^\delta)} \mathbb{E}_G[v(\Delta q)]. \tag{62}$$

We define the set of *price functions* on $\overline{\mathcal{D}}$ as

$$\mathcal{P}(v) \equiv \left\{\rho : \overline{\mathcal{D}} \to \mathbb{R} \text{ such that } \rho \text{ is convex and } \rho(\Delta q) \geq v(\Delta q) \ \forall \Delta q \in \overline{\mathcal{D}}\right\}. \tag{63}$$

The price function is derived from the dual of the optimal persuasion problem and can be used to characterize the optimal experiment, as shown in the following lemma, which restates Theorem 3 in Dworczak and Martini (2019).

Lemma 4. *The value of the optimal experiment is equal to its dual value, that is,*[12]

$$\sup_{G \in MPC(\lambda^\delta)} \mathbb{E}_G[v(\Delta q)] = \inf_{\rho \in \mathcal{P}(v)} \mathbb{E}_{\lambda^\delta}[\rho(\Delta q)]. \tag{64}$$

Moreover, a candidate experiment $\mathcal{E}$ that induces a distribution G of posterior means is optimal if there exists a price function $\rho \in \mathcal{P}(v)$ such that:

$$supp(G) \subseteq \left\{\Delta q \in \overline{\mathcal{D}} : \ v(\Delta q) = \rho(\Delta q)\right\}, \tag{Contact}$$

$$\mathbb{E}[\rho(\Delta q) \mid \sigma] = \rho\big(\mathbb{E}[\Delta q \mid \sigma]\big) \quad a.s. \tag{Linearity}$$

Condition (Linearity) implies that pooling, i.e., generating a non-degenerate posterior belief, can occur only in regions where the price function is affine. Conversely, if ρ is strictly

[12] $MPC(\lambda^\delta)$ denotes the set of all mean-preserving contraction of the distribution λ^δ.

convex in a neighborhood of a posterior mean Δq, any signal inducing Δq must fully reveal the relevant states. Furthermore, if ρ is affine on an interval $[a, b]$ and $\rho(x) > v(x)$ in its interior, condition (Contact) implies that the posterior means must be supported on the endpoints $\{a, b\}$. In such intervals, the optimal experiment corresponds to *bi-pooling*, as described by Arieli et al. (2023).

C.2 Non-linear Production Costs

Towards analyzing a setting with non-linear cost functions, suppose that firms have symmetric and convex production costs given by $C(x) = \frac{c}{2}x^2$, where $c > 0$. We start by considering the problem of maximizing total industry surplus.

Consider an experiment $\mathcal{E}$ and a public signal realization $\sigma \in \Sigma$. Suppose first that $\Delta q_{i,j} \equiv \mathbb{E}[q_i - q_j|\sigma] > 0$. Firm i's profits are given by

$$\pi_i = p_i(1 - F(\hat{\theta})) - \frac{c}{2}(1 - F(\hat{\theta}))^2, \tag{65}$$

where $\hat{\theta} = (p_i - p_j)/\Delta q_{i,j}$ (for $\Delta q_{i,j} \neq 0$) remains unchanged relative to the baseline analysis. The derivative of firm i's profit function is given by

$$\pi_i' = (1 - F(\hat{\theta})) - p_i f(\hat{\theta})\frac{1}{\Delta q_{i,j}} + c(1 - F(\hat{\theta}))f(\hat{\theta})\frac{1}{\Delta q_{i,j}}. \tag{66}$$

Solving the FOC for the price, we obtain

$$p_i = \frac{(1 - F(\hat{\theta})) + c(1 - F(\hat{\theta}))f(\hat{\theta})\frac{1}{\Delta q_{i,j}}}{f(\hat{\theta})\frac{1}{\Delta q_{i,j}}} = c(1 - F(\hat{\theta})) + \frac{1 - F(\hat{\theta})}{f(\hat{\theta})}\Delta q_{i,j}. \tag{67}$$

Analogously, we obtain for firm j's profits

$$\pi_j = p_j F(\hat{\theta}) - \frac{c}{2}F^2(\hat{\theta}), \tag{68}$$

with derivative

$$\pi_j' = F(\hat{\theta}) - p_j f(\hat{\theta})\frac{1}{\Delta q_{i,j}} + cF(\hat{\theta})f(\hat{\theta})\frac{1}{\Delta q_{i,j}}. \tag{69}$$

Solving the FOC for the price, we obtain

$$p_j = cF(\hat{\theta}) + \frac{F(\hat{\theta})}{f(\hat{\theta})}\Delta q_{i,j}. \tag{70}$$

Using the characterization of prices from the FOCs yields

$$\hat{\theta} = \frac{c}{\Delta q_{i,j}}(1 - 2F(\hat{\theta})) + \frac{1 - 2F(\hat{\theta})}{f(\hat{\theta})} = \left(\frac{c}{\Delta q_{i,j}} + \frac{1}{f(\hat{\theta})}\right)(1 - 2F(\hat{\theta})). \tag{71}$$

With a slight abuse of notation, let $\Delta q \equiv \Delta q_{A,B} = \mathbb{E}[q_A - q_B|\sigma]$. It follows that total industry profits are given by

$$\begin{aligned}\Pi(\Delta q) &\equiv \frac{c}{2}(1-F(\hat{\theta}))^2 + \frac{c}{2}F^2(\hat{\theta}) + \frac{(1-F(\hat{\theta}))^2 + F^2(\hat{\theta})}{f(\hat{\theta})}|\Delta q| \\ &= \left((1-F(\hat{\theta}))^2 + F^2(\hat{\theta})\right)\left(\frac{c}{2} + \frac{|\Delta q|}{f(\theta)}\right). \qquad (72)\end{aligned}$$

To simplify the analysis, we restrict attention to uniformly distributed consumer tastes, $\theta \sim U[0,1]$. This implies that $\hat{\theta}$ is determined by the following equation:

$$\hat{\theta} = \left(1 + \frac{c}{|\Delta q|}\right)(1-2\hat{\theta}), \qquad (73)$$

or, equivalently,

$$\hat{\theta} = \frac{c + |\Delta q|}{2c + 3|\Delta q|} = \frac{1}{3}\left(1 + \frac{c}{2c+3|\Delta q|}\right). \qquad (74)$$

Note that for any $\Delta q \neq 0$, there is a unique solution $\hat{\theta} \in (0,1)$, which uniquely pins down the equilibrium prices p_i^* and p_j^*, as well as industry profits.

For $\Delta q = 0$, consider any price $p^* \in \left[\frac{c}{4}, \frac{3c}{4}\right]$. Clearly, if firms set p^*, then the resulting profit of each firm is weakly positive as $p^*\frac{1}{2} - \frac{c}{2}\left(\frac{1}{2}\right)^2 \geq \frac{c}{8} - \frac{c}{8} = 0$. Moreover, undercutting p^* is unprofitable as the profit from undercutting is bounded from above by

$$p^* - \frac{c}{2} \leq \frac{1}{2}\left(p^* - \frac{c}{4}\right) + \frac{1}{2}\left(p^* - \frac{3c}{4}\right) \leq \frac{1}{2}\left(p^* - \frac{c}{4}\right), \qquad (75)$$

where the last term is exactly the profit at price p^*. The second weak inequality follows from the assumption that $p^* \leq 3c/4$. Thus, any $p^* \in \left[\frac{c}{4}, \frac{3c}{4}\right]$ when $\Delta q = 0$ constitutes an equilibrium price charged by both firms. For comparison to the baseline model, we select the equilibrium in which both firms set a price $p^* = \frac{c}{2}$, which is the equilibrium that arises as the limit of the equilibrium in the vertically differentiated model when the quality differential $|\Delta q|$ tends to zero.

Combining the analysis for $\Delta q \neq 0$ and $\Delta q = 0$, we can express the equilibrium industry profit as

$$\begin{aligned}\Pi(\Delta q) &= \left((1-\hat{\theta})^2 + \hat{\theta}^2\right)\left(\frac{c}{2} + |\Delta q|\right) \qquad (76)\\ &= \frac{(c + 2|\Delta q|)(2c^2 + 6c|\Delta q| + 5|\Delta q|^2)}{2(2c+3|\Delta q|)^2}. \qquad (77)\end{aligned}$$

It is straightforward to show that the function Π is strictly convex in Δq.[13] By Jensen's

[13]Note that in the figures the functions appear piecewise linear. However, they are, in fact, strictly convex.

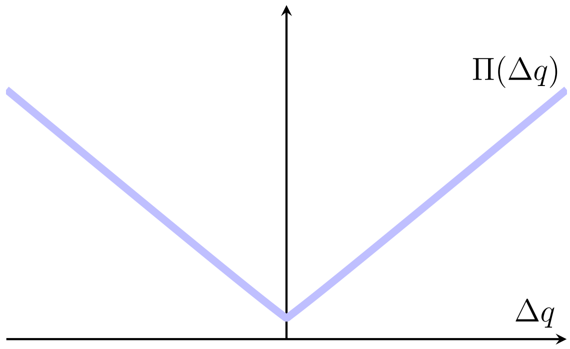

Figure C.1: Total Industry Profit $\Pi(\sigma)$, $C(x) = x^2/2$

inequality, the industry profits are maximized by the fully informative experiment.

Proposition C.1. *Suppose that $\theta \sim U[0,1]$ and that firms have symmetric and convex production costs given by $C(x) = \frac{c}{2}x^2$, with $c > 0$. Then, industry profit is maximized when the quality differential $q_A - q_B$ is fully revealed.*

Total Welfare and Consumer Surplus. Towards assessing consumer and total welfare, we continue to work with the uniform distribution on $[0,1]$. Suppose that $\Delta q = \mathbb{E}[q_A - q_B|\sigma] > 0$ for some signal $\sigma \in \Sigma$. Then, the total welfare for a given signal σ is given by

$$\begin{aligned} TW(\sigma) &\equiv \int_{\hat{\theta}}^{1} \theta \mathbb{E}[q_A|\sigma] d\theta + \int_{0}^{\hat{\theta}} \theta \mathbb{E}[q_B|\sigma] d\theta - \frac{c}{2}\left((1-\hat{\theta})^2 + \hat{\theta}^2\right) \\ &= \frac{1}{2}\mathbb{E}[q_B|\sigma] + \int_{\hat{\theta}}^{1} \theta \Delta q d\theta - \frac{c}{2}\left((1-\hat{\theta})^2 + \hat{\theta}^2\right) \end{aligned} \tag{78}$$

where $\hat{\theta}$ solves equation (73). If instead $\Delta q < 0$, then

$$\begin{aligned} TW(\sigma) &= \int_{\hat{\theta}}^{1} \theta \mathbb{E}[q_B|\sigma] d\theta + \int_{0}^{\hat{\theta}} \theta \mathbb{E}[q_A|\sigma] d\theta - \frac{c}{2}\left((1-\hat{\theta})^2 + \hat{\theta}^2\right) \\ &= \frac{1}{2}\mathbb{E}[q_B|\sigma] + \int_{0}^{\hat{\theta}} \theta \Delta q d\theta - \frac{c}{2}\left((1-\hat{\theta})^2 + \hat{\theta}^2\right). \end{aligned} \tag{79}$$

We define the function

$$H(\Delta q) \equiv \begin{cases} \frac{\Delta q}{2}(1-\hat{\theta}^2) - \frac{c}{2}\left((1-\hat{\theta})^2 + \hat{\theta}^2\right), & \text{if} \quad \Delta q \geq 0, \\ \frac{\Delta q}{2}\hat{\theta}^2 - \frac{c}{2}\left((1-\hat{\theta})^2 + \hat{\theta}^2\right), & \text{if} \quad \Delta q < 0. \end{cases} \tag{80}$$

It follows that an experiment maximizes $\mathbb{E}[TW(\sigma)]$ if and only if it maximizes $\mathbb{E}[H(\Delta q)]$ (subject to the same Bayes' plausibility constraints). As for consumer surplus, we have that an experiment maximizes $\mathbb{E}[TW(\sigma)]$ if and only if it maximizes $\mathbb{E}[H(\Delta q) - \Pi(\Delta q)]$.

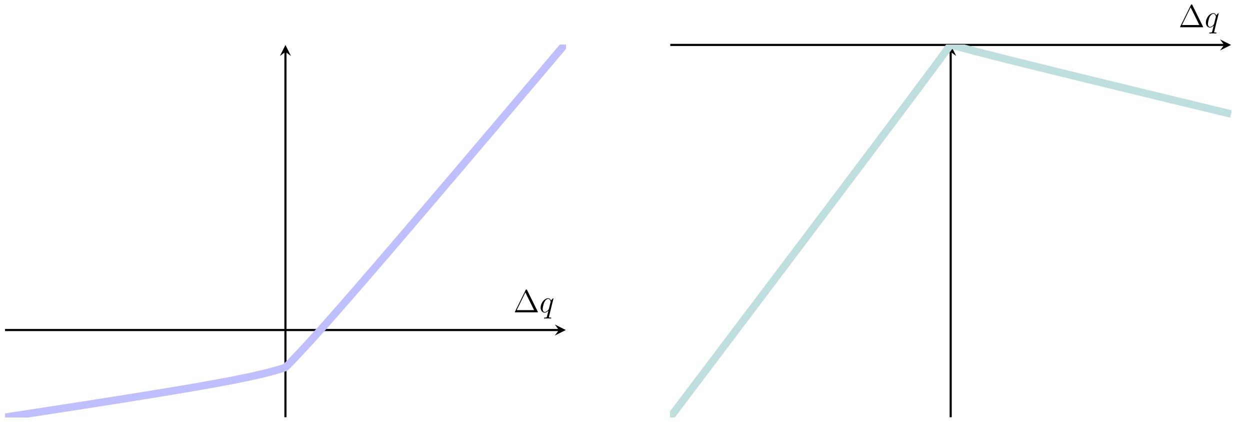

Figure C.2: Functions $H(\cdot)$ (left panel) and $H(\cdot) - \Pi(\cdot)$ (right panel) in Δq, $C(x) = x^2/2$

Proposition C.2. *Suppose that $\theta \sim U[0,1]$ and that firms have symmetric and convex production costs given by $C(x) = \frac{c}{2}x^2$, with $c > 0$. Then, total welfare is maximized when the quality differential $q_A - q_B$ is fully revealed.*

Proof. It is straightforward to show that the function H is strictly convex in Δq. By Jensen's inequality, total welfare is maximized by the fully informative experiment. □

Proposition C.3. *Suppose that $\theta \sim U[0,1]$ and that firms have symmetric and convex production costs given by $C(x) = \frac{c}{2}x^2$, with $c > 0$. Then, consumer surplus is always maximized by a rank-preserving experiment.*

Specifically, if firm i's prior expected quality is greater than that of firm j, then the optimal experiment is a censorship experiment: there exists a quality gap differential $\tilde{\delta} > 0$ such that all states with $q_i - q_j > \tilde{\delta}$ are fully revealed, and all other states are pooled into a single signal inducing symmetric quality beliefs. If instead the prior expected quality gap is zero (that is, the two firms are ex ante symmetric), then consumer surplus is maximized by the uninformative experiment.

Proof. It is straightforward to verify that the function $v(\cdot) \equiv H(\cdot) - \Pi(\cdot)$ is continuous and single-peaked in the posterior expected quality gap Δq. Moreover, $v(\cdot)$ is convex and increasing for $\Delta q < 0$, convex and decreasing for $\Delta q > 0$, and attains its maximum at $\Delta q = 0$ (see the right panel of Figure C.2).

Denote by δ^e the prior expected quality gap, $\delta^e = \mathbb{E}_{\lambda^\delta}[\delta]$, where $\delta = q_A - q_B$. If the prior mean of the quality gap is zero, $\delta^e = 0$, then $\mathbb{E}_G[v(\Delta q)]$ is maximized at the prior, $G = \lambda^\delta$, implying that the uninformative experiment maximizes consumer surplus.

Suppose that the prior mean is negative, $\delta^e < 0$. To construct the optimal experiment, we use the approach outlined in Section C.1. We first construct an experiment and then verify its optimality using the price-function approach. We proceed constructively and define an *upper censorship* experiment that reveals the most negative realizations of the quality gap and pools all quality gap realizations above a cutoff into one signal, inducing a posterior belief of zero for the quality gap. Then, we state the corresponding price function and verify its optimality using Lemma 4.

We denote $S_k \equiv \sum_{i=k}^{m} \lambda^\delta(\delta_i)\delta_i$ and $P_k \equiv \sum_{i=k}^{m} \lambda^\delta(\delta_i)$. Since $S_1 = \mathbb{E}_{\lambda^\delta}[\delta] = \delta^e < 0$ and $S_m > 0$, there exists a smallest index $k \in \{2, \ldots, m\}$ such that $S_k \geq 0$. By construction, $S_{k-1} < 0$, implying $\delta_{k-1} < 0$ and we can define

$$\alpha \equiv \frac{-S_k}{\lambda^\delta(\delta_{k-1})\delta_{k-1}} \in [0, 1] \tag{81}$$

derived from $\alpha\lambda^\delta(\delta_{k-1})\delta_{k-1} + S_k = 0$. Consider the following experiment:

- For each δ_i with $i \leq k-2$, fully reveal δ_i.
- For the boundary state δ_{k-1}, send the pooling signal σ^p with probability α and fully reveal δ_{k-1} with probability $1-\alpha$.
- For each δ_i with $i \geq k$, send the pooling signal σ^p.

Note first that the posterior mean after observing the pooling signal σ^p is by construction equal to zero. For any other signal, the state is fully revealed. Moreover, Bayes' plausibility is satisfied by construction.

Consider the slope

$$s \equiv \frac{v(0) - v(\delta_{k-1})}{0 - \delta_{k-1}} > 0 \tag{82}$$

and define the price function $\rho : \overline{\mathcal{D}} \to \mathbb{R}$ by

$$\rho(x) = \begin{cases} v(x), & \text{for } x < \delta_{k-1} \\ v(\delta_{k-1}) + s \cdot (x - \delta_{k-1}), & \text{for } x \geq \delta_{k-1}. \end{cases} \tag{83}$$

Note that the function ρ coincides with a strictly convex function up to δ_{k-1} and is linear for $x \geq \delta_{k-1}$. Since the linear part of ρ coincides with the secant line of the convex function v over the interval $[\delta_{k-1}, 0]$, it follows that ρ is convex. Furthermore, $\rho \geq v$ on the support. Hence, ρ is a feasible price function.

To verify the optimality of the constructed experiment, note that the support of posterior quality gaps under the proposed experiment is $\{\delta_1, \cdots, \delta_{k-1}\} \cup \{0\}$. States $\{\delta_1, \cdots, \delta_{k-1}\}$ are fully revealed and, by construction, ρ coincides with v on $[\delta_1, \delta_{k-1}]$. For the zero posterior gap, we also have that $\rho(0) = v(0)$, since ρ coincides with the secant line of v connecting δ_{k-1} and 0. Thus, condition (Contact) holds. To see that condition (Linearity) holds, note that for fully revealed signals, Jensen's inequality must bind as the support is a singleton. In the pooling region, $[\delta_{k-1}, \delta_m]$, the price function is linear, and hence, Jensen's inequality binds too. Thus, the price function corresponds to an optimal experiment, implying that the constructed upper censorship experiment maximizes $\mathbb{E}_G[H(\Delta q) - \Pi(\Delta q)]$ and therefore maximizes consumer surplus.

The analogous reasoning applies when the prior quality gap is strictly positive and delivers lower-censorship for the optimal experiment. □

C.3 Asymmetric Production Costs

Consider the case in which firms have asymmetric but constant marginal costs c_A and c_B. Without loss of generality, assume that $\Delta c \equiv c_A - c_B > 0$.

We begin by solving the pricing stage game. Consider any experiment $\mathcal{E}$ and a public signal realization $\sigma \in \Sigma$. Suppose first that $\Delta q_{i,j} \equiv \mathbb{E}[q_i - q_j|\sigma] > 0$ and that there exists an indifferent consumer $\hat{\theta} = (p_j - p_i)/\Delta q_{i,j}$, with $\hat{\theta} \in (0, 1)$. Then, the profit functions of firms i and j are respectively given by:

$$\pi_i(p_i, p_j|\sigma) = (p_i - c_i)(1 - F(\hat{\theta})), \tag{84}$$
$$\pi_j(p_j, p_i|\sigma) = (p_j - c_j)F(\hat{\theta}). \tag{85}$$

Solving the first-order conditions of firms i and j, we obtain

$$p_i = c_i + \frac{1 - F(\hat{\theta})}{f(\hat{\theta})}\Delta q_{i,j} \tag{86}$$

and

$$p_j = c_j + \frac{F(\hat{\theta})}{f(\hat{\theta})}\Delta q_{i,j}. \tag{87}$$

Plugging these two equations back into the expression for $\hat{\theta}$, we have:

$$\hat{\theta} - \frac{1 - 2F(\hat{\theta})}{f(\hat{\theta})} = \frac{c_i - c_j}{\Delta q_{i,j}}. \tag{88}$$

With a slight abuse of notation, define $\Delta q \equiv \Delta q_{A,B} = \mathbb{E}[q_A - q_B|\sigma]$. Since

$$\frac{c_A - c_B}{\Delta q_{A,B}} = \frac{c_B - c_A}{\Delta q_{B,A}}, \tag{89}$$

it follows that, for $\Delta q > 0$ and $\Delta q < 0$, the indifferent type $\hat{\theta}$ satisfies

$$\hat{\theta} - \frac{1 - 2F(\hat{\theta})}{f(\hat{\theta})} = \frac{\Delta c}{\Delta q}. \tag{90}$$

The corresponding industry profit is given by

$$\Pi(\Delta q) \equiv \frac{(1 - F(\hat{\theta}))^2 + F^2(\hat{\theta})}{f(\hat{\theta})}|\Delta q|. \tag{91}$$

Note that there exists an indifferent consumer whenever the quality difference is sufficiently large, i.e., whenever $\Delta q \leq -\Delta c$ or $\Delta q \geq \Delta c/2$. If instead $\Delta q \in (-\Delta c, \Delta c/2)$, then all consumers strictly prefer to buy from the more efficient firm (firm B). If $\Delta q \in (-\Delta c, 0]$, then firms set prices $p_A^* = p_B^* = c_A$ and all consumers buy from firm B. If instead $\Delta q \in (0, \Delta c/2)$, then firm A sets $p_A^* = c_A$ and firm B sets $p_B^* = c_A - \Delta q$ and serves all consumers.

Restricting attention to a uniform distribution of consumer tastes, $\theta \sim U[0,1]$, we obtain

$$\hat{\theta} = \frac{1}{3} + \frac{1}{3}\frac{\Delta c}{\Delta q}. \tag{92}$$

The industry profit in turn is given by

$$\Pi(\Delta q) = \begin{cases} \Delta c, & \text{if} \quad \Delta q \in (-\Delta c, 0], \\ \Delta c - \Delta q, & \text{if} \quad \Delta q \in (0, \Delta c/2], \\ ((1-\hat{\theta})^2 + \hat{\theta}^2)|\Delta q|, & \text{otherwise.} \end{cases} \tag{93}$$

Then, it is straightforward to show that the industry profit function Π is convex on $[\delta_1, 0]$ and on $[0, \delta_m]$. The non-convexity in industry profits arises due to the downward kink at $\Delta q = 0$.[14] Figure C.1 plots industry profits (solid line) as a function of the quality difference.

The industry-profit-maximizing information design problem takes the standard one-dimensional form:

$$\sup_{G \in MPC(\lambda^\delta)} \mathbb{E}_G[\Pi(\Delta q)], \tag{94}$$

where λ^δ denotes the prior distribution of the quality differences, and G is the induced distribution of posterior expectations of Δq (i.e., G is a mean preserving contraction of λ^δ).

Proposition C.1. *Suppose that $\theta \sim U[0,1]$ and firms have asymmetric but constant marginal costs c_A and c_B and $\Delta c = c_A - c_B > 0$. Then, industry profits are maximized by an experiment that pools quality gap realizations in some interval $[\delta', \delta'']$, where $\delta' \leq 0 \leq \delta''$, and fully reveals all realizations outside this interval.*

Proof. First, recall that industry profits are convex on $[\delta_1, 0]$ and on $[0, \delta_m]$. Second, recall that the non-convexity in industry profits derives from the downward kink at $\Delta q = 0$. It

[14] Note that the profit function is initially decreasing for $\Delta q > \Delta c/2$ and turns increasing eventually.

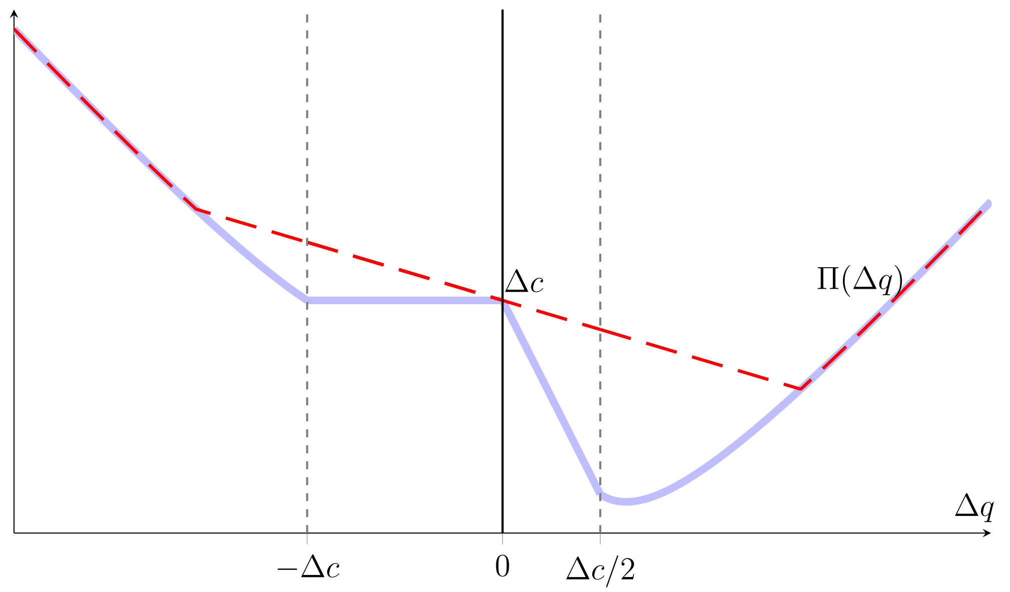

Figure C.1: Total Industry Profit $\Pi(\Delta q)$

follows from Dworczak and Martini (2019) that on any open interval on which an optimal price function ρ lies strictly above the objective function, that price function must be affine.

To show that a pooling region including $\Delta q = 0$ exists, we assume otherwise, that is, suppose that $\rho = \Pi$ on a two-sided neighborhood of zero. Then, ρ inherits the downward jump in slope at $\Delta q = 0$ from Π, contradicting convexity. Thus, we can conclude that there exists an open interval (δ', δ'') with $\rho > \Pi$ for all $\Delta q \in (\delta', \delta'')$ and $0 \in [\delta', \delta'']$.

Next, we show that there cannot be another pooling interval. Suppose otherwise, that is, that there is another interval $(\tilde{\delta}', \tilde{\delta}'') \in \overline{\mathcal{D}}$ with affine price function ρ and $\rho > \Pi$. As $0 \in [\delta', \delta'']$, it follows that $(\tilde{\delta}', \tilde{\delta}'') \subseteq (-\infty, 0)$ or $(\tilde{\delta}', \tilde{\delta}'') \subseteq (0, \infty)$. However, we know that Π is convex on each of these regions. Hence, we can construct (weak) improvements of the dual objective by choosing a modified feasible price function with $\hat{\rho}(\Delta q) = \Pi(q)$ on $[\tilde{\delta}', \tilde{\delta}'']$ and $\hat{\rho}(\Delta q) = \rho(\Delta q)$ otherwise.

Thus, the optimal price function is affine on a single interval $[\delta', \delta'']$. As any pooling signal must place support only on the affine region of the price function, it follows that any non-degenerate posterior must lie in $[\delta', \delta'']$. Condition (Linearity) forces in strictly convex regions with $\rho = \Pi$ fully revealing signals and on affine regions with $\rho = \Pi$ satisfy condition (Linearity) as well. Hence, states outside $[\delta', \delta'']$ can be fully revealed. All pooled states therefore lie in the interval $[\delta', \delta'']$, possibly with boundary randomization. $\square$

The proposition shows that industry profits are maximized by an experiment that reveals the quality difference for sufficiently large quality differentiation $|q_A - q_B|$ relative to the cost difference Δc, and pools intermediate states. Moreover, the pooling interval shrinks to a single point at zero as $\Delta c \to 0$.

The two differences from our baseline result are economically intuitive. First, pooling states in which the cost asymmetry dominates arise because they relax the fierce competition that obtains when quality differences are not large enough to overcome the efficiency advantage. Specifically, it pools states in which consumers believe that the low-cost firm has a small to moderate quality advantage together with those in which the low-cost firm provides a moderately lower-quality good. By doing so, the higher-cost firm will not compete as fiercely while still not serving any competitors, which it otherwise would in the region with $\Delta c/2 > q_A - q_B > 0$.

Second, industry profits are maximized not simply by revealing the rank order of firm qualities, but by fully revealing the intensity of the quality difference. This is due to the improvement of allocative efficiency and increased opportunity for rent extraction in the presence of asymmetric costs.

Consumer-optimal experiment. Let $TW(\sigma)$ denote the total welfare for a public signal realization $\sigma \in \Sigma$. Suppose first that $\Delta q \in (-\Delta c, \Delta c/2]$. Then, all consumers buy from firm B and the realized total welfare is $TW(\sigma) = \mathbb{E}[\theta q_B - c_B] = \frac{q_B}{2} - c_B$. If instead $\Delta q \notin (-\Delta c, \Delta c/2]$, then some consumers buy from firm A and some consumers buy from firm B. The resulting total welfare is given by

$$TW(\sigma) \equiv \begin{cases} \mathbb{E}[q_B|\sigma]/2 - c_B, & \text{if} \quad \Delta q \in (-\Delta c, \Delta c/2], \\ \mathbb{E}[q_B|\sigma]/2 - c_B + \mathbb{E}[\theta \Delta q - \Delta c | \theta \geq \hat{\theta}], & \text{if} \quad \Delta q > \Delta c/2, \\ \mathbb{E}[q_B|\sigma]/2 - c_B + \mathbb{E}[\theta \Delta q - \Delta c | \theta < \hat{\theta}], & \text{if} \quad \Delta q \leq -\Delta c. \end{cases} \tag{95}$$

We define

$$H(\Delta q) \equiv \begin{cases} 0, & \text{if} \quad \Delta q \in (-\Delta c, \Delta c/2], \\ \mathbb{E}[\theta \Delta q - \Delta c | \theta \geq \hat{\theta}], & \text{if} \quad \Delta q > \Delta c/2, \\ \mathbb{E}[\theta \Delta q - \Delta c | \theta < \hat{\theta}], & \text{if} \quad \Delta q \leq -\Delta c. \end{cases} \tag{96}$$

Since $TW(\sigma) = \mathbb{E}[q_B|\sigma]/2 - c_B + H(\Delta q)$, the law of iterated expectations implies that an experiment maximizes $\mathbb{E}[CS(\sigma)]$ if and only if it maximizes $\mathbb{E}[H(\Delta q) - \Pi(\Delta q)]$, where, in the latter expression, expectations are taken over Δq. Therefore, the consumer-surplus-maximizing information design problem takes the standard one-dimensional form.

Taking expectations in the expression for $H(\Delta q)$ with respect to the taste parameter θ, which is uniformly distributed on $U[0, 1]$, yields:

$$H(\Delta q) = \begin{cases} 0, & \text{if} \quad \Delta q \in (-\Delta c, \Delta c/2], \\ \Delta q \frac{1-\hat{\theta}^2}{2} - (1 - \hat{\theta})\Delta c, & \text{if} \quad \Delta q > \Delta c/2, \\ \Delta q \frac{\hat{\theta}^2}{2} - \hat{\theta} \Delta c, & \text{if} \quad \Delta q \leq -\Delta c. \end{cases} \tag{97}$$

Figure C.2 plots function $H(\cdot) - \Pi(\cdot)$ as a function of Δq.

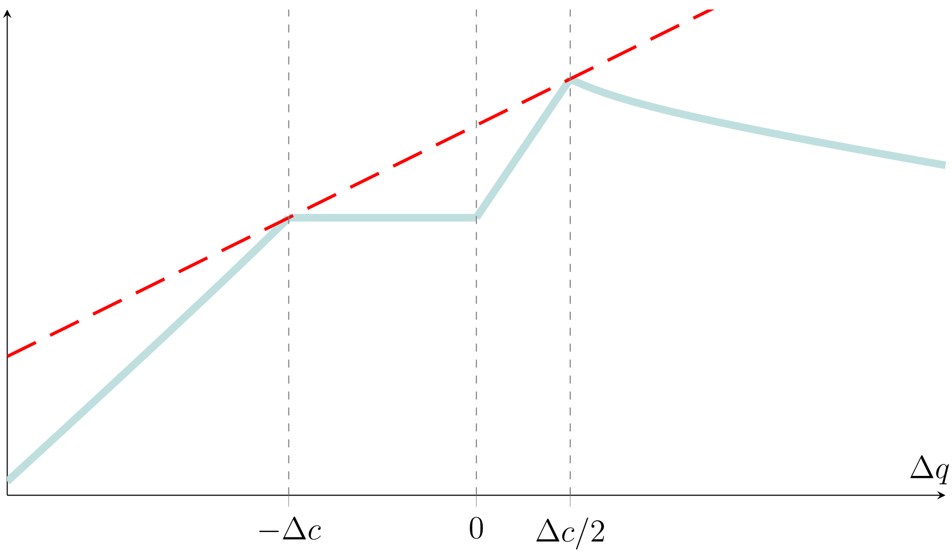

Figure C.2: Function $H(\Delta q) - \Pi(\Delta q)$

Denote by δ^e the prior expected quality gap, $\delta^e = \mathbb{E}_{\lambda^\delta}[q_A - q_B]$.

Proposition C.2. *Suppose that $\theta \sim U[0,1]$ and firms have asymmetric but constant marginal costs c_A and c_B and $\Delta c = c_A - c_B > 0$. The consumer-optimal experiment has the following features:*

(i) If the high-cost firm's expected quality advantage under the prior is equal to half the cost difference, that is, if $\delta^e = \Delta c/2$, consumer surplus is maximized by the uninformative experiment.

(ii) If the high-cost firm's expected quality advantage under the prior is greater than half the cost difference, that is, if $\delta^e > \Delta c/2$, consumer surplus is maximized by a rank-preserving lower censorship experiment.

(iii) If the high-cost firm's expected quality disadvantage under the prior is greater than the cost difference, that is, if $\delta^e < -\Delta c$, consumer surplus is maximized by a rank-preserving upper censorship experiment.

Proof. The objective function $H(\Delta q) - \Pi(\Delta q)$ is continuous and single-peaked in Δq: it is strictly convex and increasing on $[\delta_1, -\Delta c)$, constant on the interval $[-\Delta c, 0]$, linearly increasing on $[0, \Delta c/2]$, and convex and decreasing on the interval $[\Delta c/2, \delta_m]$. The objective attains its maximum at $\Delta q = \Delta c/2$ (see Figure C.2).

If the prior mean δ^e equals $\Delta c/2$, then $\mathbb{E}_G[H(\Delta q) - \Pi(\Delta q)]$ is maximized at the prior, i.e., $G = \lambda^\delta$. It follows that the uninformative experiment is optimal, establishing part (i).

Parts (ii) and (iii) follow from the same reasoning as the proof of Proposition C.3, since the prior mean δ^e lies in a strictly convex region of the objective function $H(\cdot) - \Pi(\cdot)$. □

If the prior mean is greater than $\Delta c/2$, then applying the duality theorem of Dworczak and Martini (2019), we have that $H(\Delta q) - \Pi(\Delta q)$ is maximized by a rank-preserving lower-censorship experiment, i.e., there exists some cutoff $\Delta q^* > 0$, such that the optimal experiment reveals states above Δq^* and pools the states below Δq^*, such that the expected posterior mean in the pooling interval equals the value of the function $\Delta c/2$. Similarly, if the prior mean is below $-\Delta c$, then a rank-preserving upper-censorship experiment maximizes consumer surplus. If the prior mean belongs to $[-\Delta c, \Delta c/2]$, then the optimal experiment is more involved and takes one of the following forms. First, it may bi-pool the entire state space of quality gaps, inducing a posterior–mean distribution supported on the two points $\{-\Delta c, \Delta c/2\}$ (see Arieli et al. (2023) for the details of optimal persuasion via bi-pooling). In Figure C.2, the dashed line represents the price function of an optimal bi-pooling distribution that is tangential to $H(\Delta q) - \Pi(\Delta q)$ at points $-\Delta c$ and $\Delta c/2$. Alternatively, the sender may choose an experiment that pools all sufficiently low states together, pools all sufficiently high states together, and fully reveals the intermediate range of states.

D Non-multiplicative Utility Functions

In this extension, we suppose that the type θ is distributed according to F with support $[0, 1]$. Suppose that the utility function of consumers is given by $u = u(\theta, q)$. We assume that $u \in \mathcal{C}^1$ in q and θ, $u'_1(\theta, q) > 0$ and $u'_2(\theta, q) > 0$. We also assume that the single-crossing condition holds, i.e., that $\frac{\partial^2 u}{\partial\theta\partial q} > 0$. This ensures existence, uniqueness, and monotonicity of the marginal consumer.

Suppose first that $q_A > q_B$ and define $\Delta p = p_A - p_B$. The marginal consumer type solves

$$u(\hat{\theta}, q_A) - u(\hat{\theta}, q_B) = \Delta p. \tag{98}$$

The first-order conditions of both firms are, respectively, given by

$$\frac{1 - F(\hat{\theta})}{f(\hat{\theta})} = p_A \frac{d\hat{\theta}}{dp_A} \tag{99}$$

$$\frac{F(\hat{\theta})}{f(\hat{\theta})} = -p_B \frac{d\hat{\theta}}{dp_B}. \tag{100}$$

Since $\frac{d\hat{\theta}}{dp_A} = \frac{d\hat{\theta}}{d(\Delta p)} = -\frac{d\hat{\theta}}{dp_B}$, we obtain that

$$\Delta p = \frac{1 - 2F(\hat{\theta})}{f(\hat{\theta})} \frac{d(\Delta p)}{d\hat{\theta}} = \frac{1 - 2F(\hat{\theta})}{f(\hat{\theta})} (u'_1(\hat{\theta}, q_A) - u'_1(\hat{\theta}, q_B)). \tag{101}$$

The equation determining the marginal consumer's type is

$$u(\hat{\theta}, q_A) - u(\hat{\theta}, q_B) = \frac{1 - 2F(\hat{\theta})}{f(\hat{\theta})}(u_1'(\hat{\theta}, q_A) - u_1'(\hat{\theta}, q_B)). \tag{102}$$

Rearranging, we obtain

$$\frac{f(\hat{\theta})}{1 - 2F(\hat{\theta})} = \frac{u_1'(\hat{\theta}, q_A) - u_1'(\hat{\theta}, q_B)}{u(\hat{\theta}, q_A) - u(\hat{\theta}, q_B)}. \tag{103}$$

Under uncertainty about (q_A, q_B), this equation instead takes the following form,

$$\frac{f(\hat{\theta})}{1 - 2F(\hat{\theta})} = \frac{\mathbb{E}[u_1'(\hat{\theta}, q_A) - u_1'(\hat{\theta}, q_B)|\sigma]}{\mathbb{E}[u(\hat{\theta}, q_A) - u(\hat{\theta}, q_B)|\sigma]}, \tag{104}$$

where σ is a realized public signal of some experiment $\mathcal{E}$. The following proposition characterizes the necessary and sufficient conditions under which $\hat{\theta}$ is independent of the information structure. Specifically, we show that this is the case if and only if the utility function is linear in quality, up to a monotone transformation of quality.

Proposition D.1. *The marginal consumer's type $\hat{\theta}$ is independent of the information structure (and a realized signal) if and only if there exist functions $a(\cdot) > 0$, $a' > 0$ and $b(\cdot)$ such that,*

$$u(\theta, q) = a(\theta)q + b(\theta).$$

Proof. If $\hat{\theta}$ is independent of the information structure, then it can be obtained by solving the subgame in which the experiment is fully revealing. The corresponding equation is:

$$\frac{f(\hat{\theta})}{1 - 2F(\hat{\theta})} = \frac{u_1'(\hat{\theta}, q_A) - u_1'(\hat{\theta}, q_B)}{u(\hat{\theta}, q_A) - u(\hat{\theta}, q_B)}.$$

This equation holds true for ay $q_A \neq q_B$. Note that the marginal consumer's type does not depend on q_A, q_B, $q_A \neq q_B$, if and only if there exist function $h(\theta)$ such that

$$u_1'(\theta, q_A) - h(\theta)u(\theta, q_A) = u_1'(\theta, q_B) - h(\theta)u(\theta, q_B).$$

Since this must be true for any q_A and q_B, we have that there exists a function $t(\theta)$ such that

$$u_1'(\theta, q) - h(\theta)u(\theta, q) = t(\theta).$$

In the following, we solve this ODE. Define $\mu(\theta) = \exp\left(-\int_\infty^\theta h(\tau)d\tau\right)$. Then,

$$\frac{d}{d\theta}(\mu(\theta)u(\theta, q)) = \mu(\theta)u_1'(\theta, q) - \mu(\theta)h(\theta)u(\theta, q) = \mu(\theta)t(\theta).$$

Integrating both sides from θ_0 to θ, we obtain:

$$u(\theta, q) = \left[\mu(\theta_0)u(\theta_0, q) + \int_{\theta_0}^{\theta} \mu(\tau)t(\tau)d\tau\right] \Big/ \mu(\theta).$$

We conclude that the marginal consumer's type does not depend on q_A, q_B if and only if $u(\theta, q) = a(\theta)v(q) + b(\theta)$. By relabeling of qualities $\tilde{q} = v(q)$, we obtain the result of the proposition.

□

References

Albano, G. L. and A. Lizzeri (2001): "Strategic certification and provision of quality," *International Economic Review*, 42, 267–283.

Anderson, S. P. and R. Renault (2009): "Comparative advertising: disclosing horizontal match information," *The RAND Journal of Economics*, 40, 558–581.

Arieli, I., Y. Babichenko, R. Smorodinsky, and T. Yamashita (2023): "Optimal persuasion via bi-pooling," *Theoretical Economics*, 18, 15–36.

Armstrong, M. and J. Zhou (2022): "Consumer information and the limits to competition," *American Economic Review*, 112, 534–577.

Aumann, R. J., M. Maschler, and R. E. Stearns (1995): *Repeated games with incomplete information*, MIT press.

Bagnoli, M. and T. Bergstrom (2005): "Log-concave probability and its applications," *Economic Theory*, 26, 445–469.

Benassi, C., A. Chirco, and C. Colombo (2019): "Vertical differentiation beyond the uniform distribution," *Journal of Economics*, 126, 221–248.

Bergemann, D. and A. Bonatti (2019): "Markets for information: An introduction," *Annual Review of Economics*, 11, 85–107.

Bergemann, D., B. Brooks, and S. Morris (2025): "On the alignment of consumer surplus and total surplus under competitive price discrimination," *American Economic Journal: Micro.*

Bergemann, D. and S. Morris (2019): "Information design: A unified perspective," *Journal of Economic Literature*, 57, 44–95.

Biglaiser, G., J. Gu, and F. Li (2025): "Information Acquisition and Product Differentiation Perception," *American Economic Journal: Microeconomics*, 17, 1–34.

Condorelli, D. and B. Szentes (2025): "Buyer-optimal platform design," *RAND Journal of Economics, forthcoming.*

Dizdar, D. and E. Kováč (2020): "A simple proof of strong duality in the linear persuasion problem," *Games and Economic Behavior*, 122, 407–412.

Dranove, D. and G. Z. Jin (2010): "Quality disclosure and certification: Theory and practice," *Journal of Economic Literature*, 48, 935–963.

Dworczak, P. and G. Martini (2019): "The simple economics of optimal persuasion," *Journal of Political Economy*, 127, 1993–2048.

Hopenhayn, H. and M. Saeedi (2023): "Optimal information disclosure and market outcomes," *Theoretical Economics*, 18, 1317–1344.

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, 2590–2615.

Kolotilin, A., T. Mylovanov, and A. Zapechelnyuk (2022): "Censorship as optimal persuasion," *Theoretical Economics*, 17, 561–585.

Lewis, T. R. and D. E. Sappington (1994): "Supplying information to facilitate price discrimination," *International Economic Review*, 309–327.

Mussa, M. and S. Rosen (1978): "Monopoly and product quality," *Journal of Economic Theory*, 18, 301–317.

Roesler, A.-K. and B. Szentes (2017): "Buyer-optimal learning and monopoly pricing," *American Economic Review*, 107, 2072–2080.

Schlee, E. E. (1996): "The value of information about product quality," *The RAND Journal of Economics*, 803–815.

Shaked, A. and J. Sutton (1982): "Relaxing price competition through product differentiation," *The Review of Economic Studies*, 49, 3–13.

Zapechelnyuk, A. (2020): "Optimal quality certification," *American Economic Review: Insights*, 2, 161–176.