

Prices as Signals: Joint Design of Information and Pricing

Zhicheng Du¹, Yanwei Sun², Zihe Wang³, and Jiahua Wu⁴

¹Renmin University of China, duzhicheng@ruc.edu.cn

²Imperial Business School, yanwei@imperial.ac.uk

³Renmin University of China, wang.zihe@ruc.edu.cn

⁴Imperial College London, j.wu@imperial.ac.uk

Abstract

A large body of literature studies the fundamental trade-offs of product information provision. However, much of this work treats prices as non-informative: pricing affects demand, but is not modeled as a signal about quality. To fill this gap, we study the joint design of information and pricing in markets for vertically differentiated products whose quality is privately observed by sellers. Each seller commits ex ante to a joint information–pricing mechanism that maps quality into a distribution over signal–price pairs, allowing the posted price itself to also convey information. We begin with a *monopoly setting*. We show that information design is payoff-irrelevant for the seller: *any* information structure attains the optimal revenue when paired with a price proportional to the induced posterior-mean quality. The monopoly analysis allows the buyer to be privately informed about her willingness to pay. We then turn to *oligopoly competition* in which sellers simultaneously choose such joint information–pricing mechanisms. We establish that the family of transparent contingent-pricing mechanisms is best-response-closed: against any profile of rival mechanisms, each seller has a best response that fully reveals quality and posts a deterministic quality-dependent price. In the resulting symmetric equilibrium, the pricing rule, which may not be monotone, coincides with the bid-shading function in a private-values first-price auction. When the quality distribution is log-concave, this pricing rule is strictly increasing, and the price alone reveals quality: the market needs no disclosure beyond the price tag. This equilibrium outcome is efficient: the buyer purchases the highest-quality product almost surely, and social welfare is first-best. Compared with fixed-price benchmarks with restricted disclosure, pricing flexibility triggers *over-competition*: each seller adopts transparent contingent pricing willingly, yet in thick markets every seller earns less than under a fixed-price regime, while buyers gain and social welfare rises. Finally, we discuss extensions with ex ante asymmetric sellers and with privately informed buyers.

1 Introduction

Posted prices do more than transfer surplus between buyers and sellers. When product quality is uncertain, a price can also supplement what a buyer learns from descriptions, inspections, reviews,

and other disclosures. Firms may be unable to commit literally to an immutable price scheme, but their pricing policies often remain stable over economically meaningful horizons, all the more so as pricing is delegated to algorithms. Buyers can therefore learn how a seller maps product characteristics into prices by combining prior purchase experience with public information and the seller’s reputation. Once that mapping becomes predictable, the price of a newly listed product provides a basis for inferring quality that has not been disclosed directly. The economically relevant commitment is thus not that a seller can never revise his prices, but that persistent pricing practices make current prices interpretable in light of past behavior.

Prices, however, are only one of the channels through which buyers learn. Consumers also encounter carefully curated disclosures: specifications, images, feature comparisons, warranties, and return policies, spread across retailer websites, search engines, social media, and online marketplaces. The rise of AI-powered shopping assistants further amplifies both the volume and the salience of such seller-provided information. Emerging “shopping agents” embedded in chat interfaces can generate product recommendations and present structured summaries of attributes and reviews, effectively intermediating how consumers compare and evaluate options.

The practical importance of product information provision has motivated a large literature investigating its fundamental trade-offs. In monopoly settings, additional disclosure can sharpen screening and price discrimination, but may also deter consumers who learn an unfavorable match (Bergemann et al., 2015). In oligopoly settings, competition reshapes these incentives: firms may disclose aggressively to steal demand, or strategically obfuscate to soften price competition (Hwang et al., 2026).

However, much of this literature treats pricing as *orthogonal* to information provision: prices affect payoffs and choice, but are not modeled as signals that are deliberately correlated with product quality and interpreted in light of a seller’s pricing practice. This abstraction excludes the inference described above. Once prices carry information, the seller must design disclosure and pricing jointly, because the buyer’s belief responds to both instruments at once.

To fill this gap, we study the joint design of pricing and information disclosure in markets for vertically differentiated products with uncertain quality. Each seller first commits to a *joint information–pricing mechanism*, a rule that maps quality into a joint distribution over signal–price pairs. Such joint commitment is increasingly literal in online marketplaces: automated pricing systems and pre-specified product-display pipelines jointly determine both the prices buyers face and the information they see, with little manual intervention at the level of an individual transaction. A priori, permitting arbitrary joint design of signals and prices creates a very large strategy space: sellers can garble quality in subtle ways, correlate price with disclosure, and randomize across different message formats. We show that in both monopoly and competition, this complexity largely *collapses*.

Monopoly. We begin with a monopoly setting in which a single seller faces a buyer whose willingness to pay may be privately known. The seller commits to an arbitrary joint information–pricing mechanism before observing his realized quality. Our main monopoly result shows that information design is payoff-irrelevant once prices can track beliefs: *any* information structure, from no disclosure to full disclosure, attains the optimal revenue when paired with a price proportional to the induced posterior-mean quality. Importantly, two endpoints of this continuum stand out as simple direct forms. With no disclosure, beliefs never move, and the rule collapses to a single posted price. With full disclosure, the posterior mean equals realized quality, and the rule becomes transparent contingent pricing: a deterministic price that scales with quality. This result is robust to allowing the buyer to have private information about willingness to pay.

Competition. We then turn to oligopoly competition with multiple sellers who simultaneously choose information–pricing mechanisms. The buyer observes the realized signal–price pair from each seller and chooses the option that delivers the highest nonnegative net value. In this environment, we find that, holding rivals fixed, a seller affects the buyer only through the net utility he can deliver (expected quality minus price). We show that, given *any* other sellers’ strategies, each seller has a best response within the family of *transparent contingent pricing* (TCP), under which the seller fully reveals his realized quality and posts a deterministic quality-dependent price. We then fully characterize the resulting TCP equilibrium strategy, which coincides with the bid-shading function from a private-values first-price auction. This auction connection emerges because sellers effectively “bid” by leaving surplus to the buyer in order to increase their chance of winning the sale. This equilibrium outcome is efficient: the buyer purchases the highest-quality product almost surely, the outside option is not chosen, and social welfare is first-best.

The equilibrium characterization yields additional implications. First, when the equilibrium pricing rule is strictly increasing, the price alone reveals quality: the buyer simply inverts the rule, and the price tag replaces disclosure altogether. Monotonicity, however, is not automatic: equilibrium prices need *not* be monotone in realized quality: higher-quality sellers may optimally charge lower prices in regions where incremental quality provides little competitive differentiation. We provide an exact condition on the quality distribution under which equilibrium prices are increasing in quality, and a convenient sufficient condition based on the log concavity of the prior distribution; outside this class, implementing the transparent outcome requires a disclosure channel alongside prices. Second, competition pushes prices down rapidly as the number of sellers grows. We characterize the rate at which equilibrium markups shrink in thick markets and show how the limiting surplus is divided: buyer surplus approaches the top of the quality support while per-seller revenues vanish much faster.

Decision-space comparisons. Because many platforms and regulators restrict either pricing flexibility or disclosure rules, we compare equilibrium outcomes across different seller decision

spaces. Specifically, we contrast our unrestricted joint design environment with two canonical fixed-price benchmarks that are widely used in models of price competition with quality heterogeneity. Under *no disclosure with a fixed price*, products are ex-ante homogeneous to the buyer and the interaction reduces to Bertrand competition (Tirole, 1988, Ch. 5), yielding a zero price and full market coverage but poor selection among heterogeneous qualities. Under *full disclosure with a fixed price*, the buyer is perfectly informed and would choose the best quality conditional on trade, but price rigidity generates an exclusion region: trade fails whenever the best available realized quality falls below the common posted price, a force closely related to full-information price competition with vertical differentiation (Perloff and Salop, 1985). We show that, in contrast, the unrestricted joint design environment achieves full market coverage and the first-best allocation: pricing flexibility eliminates exclusion, and disclosure (endogenously) supports efficient selection.

The comparison also highlights that efficiency and revenue need not move together. Allowing quality-contingent pricing intensifies the “race to win” by enabling sellers to tailor how much surplus they leave to the buyer, which can compress equilibrium markups relative to fixed-price competition in thick markets. As a result, for sufficiently many sellers (under mild regularity conditions), per-seller revenue under full disclosure with a fixed price is higher than under the unrestricted joint design environment, even though the latter is always more efficient. *This reversal is over-competition: adopting transparent contingent pricing is individually optimal for every seller, yet once all sellers do so, each earns less than under price rigidity.* Buyer surplus shows the mirror pattern: in thin markets the zero-price regime can dominate, while in thick markets the value of information becomes first-order and the buyer prefers regimes that rank realized quality.

We also show that the shape of the quality distribution determines which competitive friction binds, competition against the outside option or competition against rivals, and can thereby reverse revenue and surplus rankings across decision spaces. To illustrate these forces concretely, we analyze a one-parameter power-law distribution and provide a uniform-duopoly example that also includes an intermediate regime in which sellers can design information but must post a single fixed price, as in Hwang et al. (2026).

Extensions. Finally, we extend the competitive analysis along two dimensions. First, we allow sellers to be ex-ante asymmetric. There also exists an equilibrium in which all sellers use transparent contingent pricing, even with heterogeneous quality distributions. Second, we allow the buyer to be privately informed about her valuation. In this case, transparent contingent pricing need not be an equilibrium. We show that whenever an equilibrium exists, one exists in *garbled contingent pricing* (GCP) mechanisms: each seller chooses a distribution over posterior means (a mean-preserving contraction of the prior, in the sense of Blackwell, 1951) and then posts a deterministic price as a function of the realized posterior mean. We also identify a tractable duopoly example with a uniform quality prior and a binary buyer type, in which full revelation together with a linear price

schedule remains an equilibrium.

Roadmap. Section 3 characterizes optimal joint information–pricing mechanisms in the monopoly environment. Section 4 studies oligopoly competition and characterizes the transparent contingent pricing equilibrium and its asymptotic behavior. Section 5 compares equilibrium outcomes across decision spaces that restrict pricing flexibility or disclosure. Section 6 extends the oligopoly analysis to asymmetric sellers and privately informed buyers, and Section 7 concludes. All proofs are provided in the appendix.

2 Related Literature

Information Design. We model strategic disclosure using the Bayesian persuasion framework (Kamenica and Gentzkow, 2011), which has inspired a large literature on information design (see Bergemann and Morris, 2019; Kamenica, 2019; Candogan, 2020 for surveys). Relative to the canonical formulation, we adopt a slightly expanded perspective in which a seller commits not only to what information to disclose, but to an *information–pricing* joint mechanism. Equivalently, one may treat the seller’s revealed “signal” as a pair consisting of a message and an associated price quote.

Our oligopoly analysis in Section 4 contributes to the literature on competitive information design, a branch of information design that studies strategic information provision by multiple senders competing to influence a receiver’s decision. This literature can be broadly divided into two strands: (i) each sender discloses information about *his own* state (Boleslavsky and Cotton, 2015; Jain and Whitmeyer, 2019; Au and Kawai, 2020; Du et al., 2024); and (ii) all senders disclose information about a *common* state (Gentzkow and Kamenica, 2016, 2017; Ravindran and Cui, 2020; Hossain et al., 2024). Our work belongs to the first strand. Each seller privately observes the realized quality of his product, and his disclosure policy depends only on this privately observed state. Our main departure from the existing literature is that each seller chooses a *joint information and pricing mechanism*, which maps realized product quality into a joint distribution over *signal–price* pairs.

Information Design in Monopoly In the canonical monopoly–pricing setting without information design, a buyer has a private valuation and the seller maximizes revenue by posting the *Myerson price* (Myerson, 1981; Maskin and Riley, 1984). Bergemann et al. (2015) extend this framework by introducing a third-party information designer who can achieve objectives beyond seller revenue maximization through market segmentation.

In practice, sellers may also be privately informed, in the sense that they privately know the qualities of their products. The buyer’s utility depends on both product quality and her private willingness to pay. Agrawal et al. (2023) consider such a setting and focus on a monopolist who lacks

ex-ante knowledge of the buyer’s type distribution and must learn it through repeated interactions. We also consider the monopoly setting. By contrast, we assume the distribution is known and focus on the seller’s most general decision problem: a mapping from quality to a joint distribution over signals and prices. We show that such a broad decision space collapses during optimization, and the seller’s optimal expected revenue can be achieved by an arbitrary information structure paired with a price proportional to the posterior mean.

[Ling et al. \(2024\)](#) also study the joint design of information and pricing in a dynamic monopoly environment, showing that when the seller can price dynamically, the optimal information policy is fully revealing, which is closely aligned with our findings in Section 3. Their setting differs from ours in that they consider a continuous-time monopoly with finite inventory and socially learning customers who update beliefs from an initial public signal and subsequent purchase histories.

Information Design in Oligopoly Our competition analysis follows the Bertrand framework, in which multiple sellers compete for the sale by simultaneously setting prices. When sellers offer identical products, price competition drives equilibrium prices down to marginal cost. [Hotelling \(1929\)](#) introduces product differentiation in a one-dimensional spatial framework, whereas [Perloff and Salop \(1985\)](#) studies it in a discrete-choice framework with idiosyncratic consumer tastes. Both show that product differentiation softens competition and enables sellers to sustain prices above marginal cost. Closer to our vertical-quality environment, [Spulber \(1995\)](#) studies homogeneous-good Bertrand competition in which each firm privately observes its own cost, so the lowest-price firm serves the entire market. Using techniques from low-bid auctions, Spulber characterizes the unique symmetric equilibrium: firms price above cost by a markup that shrinks as competition intensifies. In the special case of constant marginal cost and perfectly inelastic demand, his equilibrium is, in effect, the procurement counterpart of our TCP schedule, with a low cost playing the role of a high quality. The resemblance ends where our question begins, however: Spulber’s information environment is fixed—firms choose only prices, consumers make no inference, and there is nothing to disclose. Our sellers instead start from the far larger space of arbitrary joint signal–price mechanisms, and our central result is that they lose nothing by forgoing all of it for transparent contingent pricing. [Wolinsky \(1983\)](#) studies price competition in a market where each firm chooses product quality *endogenously* at different costs and then posts a *deterministic* price under imperfect consumer information. Consumers do not observe quality directly; instead, they infer it through the posted price, which serves as the firm’s signaling instrument (on top of any background, product-specific noisy information consumers may obtain). They show that a fulfilled-expectations separating equilibrium exists in which each quality level is associated with a unique posted price. In contrast, in our setting, quality is exogenously random, and each seller chooses a joint information–pricing mechanism (allowing for randomized disclosure and pricing). The canonical price-signaling literature develops the no-commitment side of this question. [Milgrom and Roberts \(1986\)](#) show

that a monopolist introducing a new experience good may signal quality through its introductory price together with dissipative advertising, with the direction of the price distortion determined by how production costs vary with quality. [Bagwell and Riordan \(1991\)](#) show that when higher quality carries a higher unit cost, the monopolist signals it with a supra-monopoly introductory price that declines as more consumers become informed. In both models the seller cannot commit to a pricing rule, so separation rests on costly distortions and on refinements that discipline off-path beliefs. Our sellers instead commit ex-ante to an observable mechanism, which turns prices from costly signals into exact, distortion-free carriers of information.

More recently, the literature has shifted toward studying how sellers endogenously determine product differentiation through strategic signaling. [Armstrong and Zhou \(2022\)](#) extend the Hotelling model by introducing a centralized information designer who solves for optimal information structures. In contrast, based on the Perloff-Salop model, [Hwang et al. \(2026\)](#) consider a game where each seller individually controls signaling about his own product quality. In their setting, sellers can only post a uniform price and cannot tailor it based on the realized state. They provide a structural characterization of the symmetric pure-price equilibrium, in which sellers generally refrain from full revelation and set a fixed price pinned down by a first-order condition. Closest in spirit is [Ivanov \(2013\)](#), who studies sellers that simultaneously choose how much information to reveal about their own product and a uniform price. There, the hidden variable is the buyer’s idiosyncratic match value (unknown even to the seller), disclosure is restricted to a one-parameter rotation-ordered family of signals, and the price is a non-contingent scalar that carries no information; full disclosure emerges as the unique symmetric equilibrium only when the number of sellers is sufficiently large, and it fails even in duopoly. We instead let each seller commit to an arbitrary joint distribution over signal–price pairs, so the posted price is itself an informative, quality-contingent signal, and transparent contingent pricing arises for every number of sellers.

[Li et al. \(2025\)](#) study a vertically differentiated duopoly in which a third-party designer commits to a public disclosure policy about qualities before firms price: revealing only the quality ranking softens price competition and maximizes industry profit, while full concealment is consumer-optimal. The contrast with our setting is twofold: their disclosure is centralized and public, and their firms hold no private information—the public signal is all anyone observes—whereas our sellers each privately observe quality and individually design joint information–pricing mechanisms. Our work extends this literature by endowing sellers with full flexibility in signaling and state-dependent pricing. Unlike [Hwang et al. \(2026\)](#), we show that when sellers have full flexibility, *full information* together with deterministic state-dependent pricing forms an equilibrium.

Joint Design of Information and Pricing. To the best of our knowledge, the papers closest to ours are [Lingenbrink and Iyer \(2019\)](#) and [Sun et al. \(2024\)](#), which likewise allow the most general form of *joint* information–pricing design: a mechanism mapping states to a joint distribution

over signal–price pairs. [Lingenbrink and Iyer \(2019\)](#) consider information design and pricing in an $M/M/1$ queueing setting, and show that the optimal revenue can be achieved by either information design with a single posted price or by transparent contingent pricing. Our monopoly analysis delivers a similar message in a stripped-down environment: both mechanisms attain the optimal revenue.

[Sun et al. \(2024\)](#) study a single sender facing multiple receivers, where the sender chooses a mechanism mapping the state into a joint distribution over signals and prices. In their setting, a coordination friction across receivers makes information design with a single price outperform transparent contingent pricing. Both papers, however, feature a single seller or sender: neither studies the joint design problem under competition. By contrast, in our oligopoly competition model, we show that transparent contingent pricing is a weakly dominant strategy in the full joint mechanism space. In this sense, transparent contingent pricing is “better” than information design with a single posted price in competitive markets.

3 Monopoly

We first study the monopoly benchmark and show that information design is payoff-irrelevant: any information structure attains the optimal revenue when paired with a price proportional to the posterior mean, and the price alone can then replace the signal (Theorem 1 and Corollary 1).

3.1 The Setup

There is one seller (he) who sells a product to a unit-demand buyer (she). The seller’s product has quality Q , drawn from prior distribution $F \in \Delta([0, 1])$.¹ The buyer has an outside option normalized to 0 and a private type θ , drawn from distribution $G \in \Delta([\underline{\theta}, \bar{\theta}])$ where $[\underline{\theta}, \bar{\theta}] \subseteq \mathbb{R}_+$.²

As for utilities, if the buyer purchases the product of realized quality q at price p , her payoff is $\theta q - p$ and the seller’s payoff is p ; otherwise, both receive 0. We assume both parties are risk neutral, and ties are broken in the seller’s favor: if the buyer is indifferent, she buys.

Seller’s strategy. The seller can use the price itself as a signal of product quality, thereby jointly designing information disclosure and pricing. Specifically, we take the signal space $\mathcal{S}$ to be standard Borel and the price space to be the compact interval $\mathcal{P} := [0, \bar{p}]$, where $\bar{p} > \bar{\theta}$. The seller commits ex-ante to a mechanism which maps realized qualities to a joint distribution over signal–price pairs. Denote the space of such joint mechanisms by

$$\mathcal{M} := \left\{ (\mathcal{S}, \kappa) : \mathcal{S} \text{ is standard Borel and } \kappa : [0, 1] \rightarrow \Delta(\mathcal{S} \times \mathcal{P}) \text{ is a stochastic kernel} \right\}. \quad (1)$$

¹Throughout the paper, $\Delta(\cdot)$ denotes the set of all probability distributions; unless specified otherwise, a distribution is represented by its cumulative distribution function (CDF).

²The type θ represents the buyer’s marginal rate of substitution between utility and quality ([Tirole, 1988](#)).

Given any realized quality q , the seller sends the buyer a signal-price pair (s, p) with probability $\kappa((s, p) \mid q)$. When there is no risk of confusion, we write κ to denote the entire mechanism $(\mathcal{S}, \kappa)$.

Buyer's action. Due to the buyer's quasi-linear utility, both the seller and the buyer care about the posterior belief only through its mean. Upon observing the pair (s, p) , the buyer forms the posterior mean $m := \mathbb{E}[Q \mid (s, p)]$ via Bayes' rule, and purchases the product, rather than choosing the outside option, if and only if $\theta \geq p/m$.³ Under the seller-preferred tie-breaking, the purchase probability is denoted by $\bar{G}(p/m) := \Pr_G\{\theta \geq p/m\}$.

Seller's problem. The seller's problem is to find the best joint mechanism $\kappa^* \in \mathcal{M}$ to maximize the revenue, i.e., κ^* is a solution to the following problem

$$\kappa^* \in \operatorname{argmax}_{\kappa \in \mathcal{M}} R(\kappa) := \mathbb{E}_{q \sim F, (s, p) \sim \kappa(\cdot \mid q)} \left[p \cdot \bar{G}\left(\frac{p}{m}\right) \right]. \quad (\mathcal{P}_{\text{mono}})$$

3.2 Prices Suffice as Signals under Optimal Mechanisms

We show that any feasible information structure can be paired with a pricing rule that depends only on the realized posterior mean to attain the optimal revenue.

Theorem 1 (Optimal mechanism under monopoly). *An optimal solution to Problem $\mathcal{P}_{\text{mono}}$ combines an arbitrary information structure with a pricing rule that depends only on the posterior mean: for every posterior mean m , an optimal pricing rule is*

$$p^*(m) = \theta^* m \quad \text{where} \quad \theta^* \in \operatorname{argmax}_{\theta \in [\underline{\theta}, \bar{\theta}]} \theta \bar{G}(\theta).$$

Theorem 1 shows that *the seller's choice of information structure does not affect the optimal revenue*, provided that the posted price adjusts proportionally to the realized posterior mean. The next corollary establishes a stronger implementation result: the auxiliary signal can be suppressed without changing the buyer's purchase decision or the seller's revenue.

Corollary 1 (Price sufficiency under monopoly). *Suppose $\theta^* > 0$, and fix any information structure paired with the pricing rule $p^*(m) = \theta^* m$ of Theorem 1. Suppressing the auxiliary signal and posting only the price leaves the buyer's purchase decision and the seller's revenue unchanged: because p^* is one-to-one in m , the posted price reveals the realized posterior mean, which is the only payoff-relevant statistic for the buyer.*

Price sufficiency is a statement about the auxiliary *signal*, not about quality: the price reveals exactly as much about quality as the chosen information structure discloses. Under full disclosure, the price is one-to-one in realized quality and reveals it; under no disclosure, the single posted price $\theta^* m_0$ reveals only the prior mean—consistent with the optimality of no disclosure plus a single

³When $m = 0$, we adopt the convention $p/m := +\infty$ for $p > 0$ and $0/0 := 0$: the buyer never pays a positive price for a product known to be worthless, and at a zero price the seller's revenue is zero either way, so the convention is immaterial for revenue.

price. What the corollary rules out is any role for a *separate* signal alongside the price.

Price sufficiency means that restricting attention to mechanisms in which the posted price is the only observable message entails no loss in optimal revenue. This is an optimal-implementation result: under the monopolist's optimal joint design of information and pricing, active disclosure becomes redundant, because the committed pricing rule allows the posted price to serve simultaneously as a transfer and as a sufficient signal of all payoff-relevant information.

The proof reduces each realized signal–price pair (s, p) to the posterior mean m and the induced purchase cutoff $\tilde{\theta} = p/m$ for $m > 0$: the buyer purchases if and only if $\theta \geq \tilde{\theta}$, and the seller's conditional revenue is $m\tilde{\theta}\bar{G}(\tilde{\theta})$. This reduced form separates the posterior mean from the cutoff choice, so the seller optimally selects the cutoff $\theta^* \in \operatorname{argmax}_{\theta \in [\underline{\theta}, \bar{\theta}]} \theta\bar{G}(\theta)$ after every posterior realization and charges θ^*m . Further, by Bayes plausibility, every mechanism yields revenue at most $m_0\theta^*\bar{G}(\theta^*)$, and the proportional pricing rule attains this bound under any information structure.

Moreover, when $\theta^* > 0$, the rule $p^*(m) = \theta^*m$ is one-to-one in m . The buyer can therefore infer the posterior mean directly from the posted price. Because only the posterior mean enters her expected utility, a separate signal conveys no additional payoff-relevant information.

4 Competition

We now turn to oligopoly markets in which multiple sellers compete for a single buyer by simultaneously choosing joint mechanisms. First, we prove the existence of a symmetric equilibrium in which sellers adopt a *transparent contingent pricing* mechanism (Theorem 2). Unlike in monopoly, a price-only implementation requires the equilibrium pricing rule to be one-to-one, with strict log-concavity providing a sufficient condition (Proposition 1). Finally, we characterize the asymptotic equilibrium behavior as the number of competing sellers grows sufficiently large (Proposition 2).

4.1 The Setup

We consider $n \geq 2$ ex-ante homogeneous sellers, indexed by $i \in [n]$, who compete for a single unit-demand buyer. Each seller i offers one product with quality Q_i , which is privately observed by that seller. Qualities $Q_1, \dots, Q_n$ are i.i.d. from prior distribution $F \in \Delta([0, 1])$ with a strictly positive density f (i.e., ex-ante symmetric sellers). The buyer has an outside option normalized to 0. We assume that the buyer's type is degenerate, i.e., $\underline{\theta} = \bar{\theta}$, which we normalize to one without loss of generality; the buyer thus holds no private information. In Section 6, we relax the assumptions of ex-ante symmetry across sellers and a degenerate buyer type.

Seller's strategy. As in the monopoly setting, each seller *simultaneously* commits ex-ante to a mapping from his realized qualities to a joint distribution over signals and prices. Specifically, each seller chooses a strategy $\kappa_i \in \mathcal{M}$, where $\mathcal{M}$ is defined in (1); in particular, every posted price lies

in the compact set $\mathcal{P} = [0, \bar{p}]$, with $\bar{p} > \bar{\theta} = 1$ in the present section.

Buyer's action. Upon observing seller i 's realized signal-price pair (s_i, p_i) , the buyer forms a posterior belief about Q_i via Bayes' rule. Let $m_i := \mathbb{E}[Q_i \mid (s_i, p_i)]$ denote the corresponding posterior mean quality. The buyer's net utility from purchasing seller i 's product is $B_i := m_i - p_i$. The buyer chooses a seller from $\arg \max_{j \in [n]} B_j$ if $\max_{j \in [n]} B_j \geq 0$; otherwise she takes the outside option whose value is normalized to zero. Ties between the outside option and purchasing from one of the sellers are broken in favor of sellers. For ties among sellers, we adopt the convention that, when analyzing a fixed seller i , ties are resolved in favor of seller i ; the payoff expression below embeds this convention.⁴

Seller's payoff. For any seller i under a strategy profile (κ_i, κ_{-i}) , denote the highest net payoff among the other sellers by $B_{-i} := \max_{j \neq i} B_j$. Let

$$H(b) := \Pr\{B_{-i} \leq b\} = \Pr\{\max_{j \neq i} B_j \leq b\}$$

be the probability of beating all rivals (not including the outside option) when seller i offers the buyer a net utility of b . Since qualities are i.i.d. and each seller's mechanism depends only on his own quality, B_{-i} is independent of seller i 's signal-price pair (s_i, p_i) . Thus, seller i wins the game if and only if $B_i \geq \max\{0, B_{-i}\}$, so his ex-ante expected revenue is

$$R_i(\kappa_i, \kappa_{-i}) := \mathbb{E}[p_i \cdot \mathbf{1}\{B_i \geq \max\{0, B_{-i}\}\}] = \mathbb{E}[p_i \cdot \mathbf{1}\{B_i \geq 0\} \cdot H(B_i)].$$

Solution concept. We focus on *pure Nash equilibrium* (hereafter, equilibrium). A profile $\kappa^* = (\kappa_1^*, \dots, \kappa_n^*) \in \mathcal{M}^n$ forms an equilibrium if and only if, for every seller i and every alternative mechanism $\kappa_i \in \mathcal{M}$, strategy κ_i^* is a best response, i.e.,

$$R_i(\kappa_i^*, \kappa_{-i}^*) \geq R_i(\kappa_i, \kappa_{-i}^*).$$

That is, fixing the other sellers' committed mechanisms κ_{-i}^* , no seller can increase his expected revenue by unilaterally switching to a different joint information-pricing mechanism.

Timing. We summarize the timeline of the game as follows:

- (i) Each seller commits ex-ante to a joint information-pricing mechanism $\kappa_i \in \mathcal{M}$.
- (ii) Each seller privately observes his quality Q_i , drawn from the prior F . Each seller sends a signal-price pair (s_i, p_i) to the buyer according to the committed mechanism $\kappa_i(\cdot \mid Q_i)$.
- (iii) The buyer updates her posterior belief about each seller's product based on the observed

⁴The convention matters only off path, when rivals' mechanisms place an atom on the highest rival score: under a less favorable tie-breaking rule the seller's supremum need not be attained, and Theorem 2(i) should then be read in its ε -best-response form (see the footnote to the theorem). On path in the equilibrium of Theorem 2(ii), the induced score distribution is atomless, so seller-seller ties occur with probability zero and every tie-breaking rule yields the same outcome.

message (s_i, p_i) , through Bayes' rule.

- (iv) The buyer chooses the seller with the highest net utility or the outside option. Ties are broken according to the way mentioned above.

4.2 Equilibrium as Transparent Contingent Pricing

Before characterizing equilibrium in the enlarged mechanism space $\mathcal{M}$, we first establish that the added pricing flexibility is economically substantive. The following example shows that, in the uniform-duopoly environment of [Hwang et al. \(2026\)](#), the fixed-price equilibrium ceases to be stable once sellers can commit to quality-contingent pricing schemes.

Example 1 (Fixed-price equilibrium fails). *There are two sellers whose qualities are i.i.d. from the uniform distribution $U[0, 1]$. Sellers can design information but are restricted to fixed prices. [Hwang et al. \(2026\)](#) show that there exists a unique symmetric equilibrium, consisting of a partial-disclosure policy G^* and a fixed price $\bar{p} = \frac{3-\sqrt{5}}{2}$ (see their Example 1). Under this equilibrium, each seller's expected revenue is $5\sqrt{5} - 11$. However, this equilibrium does not survive in our enlarged space $\mathcal{M}$. Suppose seller 1 uses strategy $(G^*, \bar{p})$. Seller 2 deviates to the following pricing scheme:*

$$\tilde{p}(q) = \begin{cases} q, & \text{if } 0 \leq q < a, \\ \frac{q+a}{2}, & \text{if } a \leq q \leq 1, \end{cases} \quad \text{where } a = \frac{7-3\sqrt{5}}{2}.$$

Note that $\tilde{p}(\cdot)$ is strictly increasing on $[0, 1]$. Hence, after observing the posted price $\tilde{p}(q)$, the buyer can invert the pricing rule and recover the realized quality q . The posted price itself therefore serves as a fully revealing signal.

Seller 2's expected revenue from the deviation is $\frac{101-43\sqrt{5}}{24} > 5\sqrt{5} - 11$. □

Once sellers can commit to quality-contingent prices, the fixed-price equilibrium is no longer robust. Moreover, when the price scheme is strictly monotone, the posted price itself fully reveals the realized quality. Thus, prices serve not merely as transfers in this environment, but also as part of the information structure. At first glance, this richer decision space seems to introduce substantial strategic complexity. However, our main result shows that this complexity collapses in equilibrium. In the unrestricted mechanism space, there exists a symmetric equilibrium in which every seller adopts a *transparent contingent pricing* mechanism as defined below.

Definition 1 (Transparent contingent pricing (TCP)). *A mechanism κ is called transparent contingent pricing if it (i) fully reveals the realized quality, and (ii) posts a deterministic quality-contingent price. Formally, $\mathcal{S} = [0, 1]$, and there exists a measurable function $p : [0, 1] \rightarrow \mathcal{P}$ such that*

$$\kappa((q, p(q)) \mid q) = 1, \quad \forall q \in [0, 1].$$

The seller separates the informational and pricing roles of the joint mechanism when adopting

the above kind of strategy. On the information side, he makes product quality fully transparent. On the pricing side, the seller extracts revenue through a deterministic price schedule $p(q)$ that adjusts to the realized quality.

This form of mechanism is common in markets where product attributes are explicitly displayed, and prices vary with those attributes. For example, online used-car platforms often truthfully disclose objectively verifiable attributes, such as mileage, model year, accident history, and inspection results, while allowing prices to adjust to the vehicle's underlying condition.

Theorem 2 (Equilibrium as transparent contingent pricing). *For the oligopoly market, we have*

- (i) *Given any other sellers' strategies (i.e., the competitive environment), each seller has a best response that falls within the family of transparent contingent pricing mechanisms.*⁵
- (ii) *There exists a symmetric equilibrium in which every seller fully reveals realized quality and adopts the deterministic quality-dependent pricing scheme:*

$$p^*(q) := \frac{\int_0^q F(t)^{n-1} dt}{F(q)^{n-1}}, \quad \forall q \in (0, 1]. \quad (2)$$

The net utility left for the buyer $b^(q) := q - p^*(q)$ is strictly increasing on $(0, 1]$, so the buyer almost surely purchases the highest-quality product and never selects the outside option.*

Hereafter, we refer to the equilibrium characterized above as the *TCP equilibrium*.

The theorem delivers two complementary conclusions. First, against any profile of rival mechanisms, a seller can attain a best response by fully revealing his realized quality and tailoring the price to both quality and the competitive environment. Second, these best responses are mutually consistent: there exists a symmetric equilibrium in which all sellers fully reveal quality and use the same closed-form pricing scheme. The key message is that *competition eliminates the strategic value of quality garbling*: regardless of rivals' strategies, a seller can disclose quality transparently and compete entirely through quality-contingent prices.

The force behind transparent contingent pricing is that, conditional on rivals' strategies, the buyer evaluates seller i only through the net utility she receives, $B_i = m_i - p_i$. Seller i is chosen if and only if $B_i \geq \max\{0, B_{-i}\}$, so any signal-price pair (s_i, p_i) matters only through the net utility it induces. Fix the rivals' mechanisms. Conditional on quality q_i , seller i faces a standard win-margin tradeoff: he chooses a buyer net utility $b \in [0, q_i]$ to maximize $(q_i - b)H(b)$, where $q_i - b$ is his margin conditional on being chosen and $H(b)$ is the probability of winning. Leaving more surplus to the buyer raises the probability of winning but reduces the price upon winning.

More generally, a joint information-pricing mechanism induces a distribution over buyer net utilities conditional on each q_i . The seller's expected revenue is linear in this distribution and can

⁵Under the seller-favorable tie convention of Section 4.1. Under an arbitrary tie-breaking rule, the statement holds in the following form: for every $\varepsilon > 0$, the seller has an ε -best response within the family, and an exact best response exists whenever the rivals' induced score distribution is atomless—as it is in the equilibrium of part (ii).

be represented as an average of $(q_i - b)H(b)$. Randomization or information garbling therefore offers no additional benefit once H is fixed: for each q_i , the seller can select a deterministic maximizer $b(q_i) \in [0, q_i]$. This choice is implemented directly through transparent contingent pricing by revealing q_i and charging $p_i(q_i) = q_i - b(q_i)$.

Moreover, Theorem 2 has several additional implications:

- (i) *A large mechanism space collapses.* Although each seller may commit to an arbitrary joint information–pricing mechanism, Theorem 2 shows that additional signaling or randomization provides no revenue advantage over transparent contingent pricing.
- (ii) *Sellers compete by leaving surplus.* Under transparent contingent pricing, sellers compete through the net utility left to the buyer. In symmetric equilibrium, $b^*(\cdot)$ is strictly increasing, so the highest-quality seller leaves the highest net utility and is chosen almost surely.
- (iii) *Connections to first-price auctions.* A seller with quality q_i chooses b by trading off the margin $q_i - b$ against the probability of winning, exactly as a bidder with value q_i chooses a bid in a first-price auction. Consequently, $b^*(q_i) = q_i - p^*(q_i)$ coincides with the symmetric equilibrium bidding function for n i.i.d. bidders with values distributed according to F , while $p^*(q_i)$ equals the corresponding amount of bid shading.

Remark 1 (Equilibrium multiplicity and selection). *Theorem 2 establishes existence, not uniqueness; other equilibria may arise. Nevertheless, transparent contingent pricing is a natural selection: against every rival profile, a seller attains his best-response payoff within this family (part (i)), so garbling cannot strictly improve his payoff on or off the equilibrium path. Within TCP, the score reduction links symmetric equilibria to those of a first-price auction. Characterizing multiplicity and its implications for buyer payoffs is a promising direction.*

Finally, it is worth noticing that the equilibrium pricing function $p^*(\cdot)$ in (2) need not be monotone in quality. Consequently, a seller with a higher realized quality may charge a lower price, and posted prices need not preserve the ranking of underlying qualities. The next result provides an exact characterization of price monotonicity.

4.3 Monotone Prices Suffice as Signals

When the equilibrium price schedule is one-to-one, the posted price itself reveals quality; otherwise, a separate quality signal remains necessary. This property need not hold in equilibrium in general, and we provide both an exact characterization and a convenient sufficient condition for it.

Proposition 1 (Monotone prices suffice as signals). *The equilibrium pricing rule p^* in (2) is strictly increasing everywhere if and only if $\int_0^q F(t)^{n-1} dt$ is strictly log-concave in q .⁶ In this case, the TCP equilibrium in Theorem 2 admits a price-only implementation: sellers need not send a*

⁶A positive function h is strictly log-concave on an interval if $\log h$ is strictly concave.

separate quality signal, because the buyer can recover q_i by inverting $p^*(q_i)$. A sufficient condition is that F is strictly log-concave on $(0, 1]$, which is implied by strict log-concavity of the density f .




Figure 1: Relationship between equilibrium pricing p^* and prior distribution F

Note. Consider the prior family $F(q) = q \cdot e^{\alpha(q^2-1)}$. For $\alpha = 0$, the log-concavity of $f(q)$ serves as a sufficient condition for the log-concavity of $F(q)$, ensuring a monotonic $p^*(q)$. For $\alpha = 0.5$, although the density $f(q)$ is no longer log-concave, the CDF $F(q)$ remains log-concave, which is sufficient to preserve the monotonicity of the equilibrium pricing rule. For $\alpha = 3.5$, both $f(q)$ and $F(q)$ violate log-concavity, and this loss of regularity leads to a non-monotone equilibrium pricing function.

Proposition 1 clarifies when equilibrium *prices* inherit the natural ordering of *quality*. In our equilibrium, a seller with realized quality q wins with probability $F(q)^{n-1}$ and earns expected revenue $\int_0^q F(t)^{n-1} dt$, so the posted price can be written as (2). A higher q increases both the numerator (the expected payoff from winning) and the denominator (the probability of winning). Nonmonotonicity arises when, in some region, the probability of winning increases *relatively faster* than the cumulative gains from winning, so the ratio (2) declines. In that case, a higher-quality seller optimally posts a lower price: to secure the sale, it is more effective to concede the surplus to the buyer for winning rather than relying on the incremental quality advantage. The log-concavity condition in Proposition 1 rules out exactly this reversal.

Moreover, when p^* is strictly increasing, the posted price itself fully reveals quality: observing $p^*(q)$ allows the buyer to invert the pricing rule and recover the realized quality q . Separate disclosure is therefore redundant, and transparent contingent pricing reduces to pure price signaling. Unlike in classical price-signaling models, nothing here rests on off-path belief refinements or dissipative distortions: the pricing rule is committed and observable, so the buyer simply inverts it. Log-concavity of the prior is also a standard regularity condition in auction and mechanism-design models. Hence, in many commonly studied regular environments, **prices suffice as signals**: additional disclosure does not affect the induced allocation or equilibrium payoffs.

Figure 1 provides a concrete example in which monotonicity fails (e.g., the case $\alpha = 3.5$): when the quality distribution places substantial mass near the top, many sellers are clustered at high qualities, so incremental improvements in q provide limited differentiation. In that region, competition primarily operates through surplus concessions, which can translate into lower posted prices even as quality increases.

4.4 Asymptotic Equilibrium Behavior

A central question in competitive information design is how intensified competition reshapes the allocation of surplus and the role of information (see [Au and Kawai \(2020\)](#); [Hwang et al. \(2026\)](#)). This subsection studies the asymptotic behavior of the transparent contingent pricing equilibrium as the number of symmetric sellers grows.

We use the following notation. Let $Q_{(i;n)}$ denote the i^{th} -largest value among all realized qualities $Q_1, \dots, Q_n$ i.i.d. from prior F . Regarding the equilibrium outcome for sufficiently large n , we obtain the following results.

Proposition 2 (Asymptotic equilibrium behavior). *Suppose the prior density f is strictly positive, bounded, and continuous. For the equilibrium with n sellers obtained in Theorem 2, we have*

- (i) *The equilibrium pricing function under the equilibrium with n sellers, rewritten as $p_n^*(q)$, uniformly converges to the zero function at a rate of $\Theta(1/n)$; moreover, pointwise,⁷*

$$p_n^*(q) \sim \frac{1}{n} \cdot \frac{F(q)}{f(q)}, \quad \forall q \in (0, 1].$$

- (ii) *The buyer's ex-ante expected utility is denoted by $U_{\text{gen}}(n)$, which takes the form*

$$U_{\text{gen}}(n) = \mathbb{E}[Q_{(2;n)}], \quad \text{with} \quad 1 - U_{\text{gen}}(n) \sim \frac{1}{n} \cdot \frac{2}{f(1)}.$$

- (iii) *Each seller's ex-ante expected utility is denoted by $R_{\text{gen}}(n)$, which takes the form*

$$R_{\text{gen}}(n) = \frac{\mathbb{E}[Q_{(1;n)}] - \mathbb{E}[Q_{(2;n)}]}{n} \sim \frac{1}{n^2} \cdot \frac{1}{f(1)}.$$

- (iv) *The ex-ante expected social welfare, i.e., the sum of the buyer's and all sellers' utilities, is denoted by $SW_{\text{gen}}(n)$, which takes the form*

$$SW_{\text{gen}}(n) = \mathbb{E}[Q_{(1;n)}], \quad \text{with} \quad 1 - SW_{\text{gen}}(n) \sim \frac{1}{n} \cdot \frac{1}{f(1)}.$$

In sum, as the number of sellers grows, competition progressively erodes quality differentiation, and the market asymptotically approaches a Bertrand-like environment in which quality advantages generate vanishing markups.

First, part (i) shows that competition forces equilibrium prices toward zero as the number of sellers grows. The ratio $F(q)/f(q)$, which is the reciprocal of the reverse hazard rate, summarizes how much “nearby” competition a seller faces from slightly lower qualities: when $f(q)$ is large, many rivals are likely to have quality close to q from below, so the top seller cannot sustain much markup and must post a lower price to fend off nearby competitors. The key point is the $1/n$ factor: with more sellers, the sustainable markup shrinks proportionally to $1/n$, uniformly over $q \in [0, 1]$.

⁷Throughout this paper, the notation $x_n \sim y_n$ denotes asymptotic equivalence, defined as $\lim_{n \rightarrow \infty} \frac{x_n}{y_n} = 1$.

Second, part (ii) characterizes how the buyer benefits from this intensified competition. At equilibrium with n sellers, the buyer’s ex-ante expected utility equals the expected second-highest quality, which corroborates the revenue equivalence theorem (Myerson, 1981). To see this, note that the buyer’s utility in our setting corresponds to the auctioneer’s revenue in a first-price auction. By revenue equivalence, first-price and second-price auctions generate the same expected revenue for the auctioneer, and the second-price auction’s payment is the second-highest *value*—here, the second-highest quality. In other words, full pricing flexibility hands the buyer exactly the surplus she would obtain from running a second-price auction among the sellers. Intuitively, the buyer can credibly compare many close substitutes, so the winning seller must price aggressively, leaving the buyer with almost the entire surplus when n is large enough. The constant $f(1)$ matters because it captures how much probability mass lies near the top of the support: when $f(1)$ is larger, it is more likely that multiple sellers draw qualities very close to 1, which further increases competitive pressure and pushes the buyer’s payoff closer to 1.

Finally, parts (iii) and (iv) describe how the remaining surplus is split and how welfare converges. Because the best and second-best qualities become very close when n is large, total seller revenue is tiny and each individual seller’s expected utility is even smaller. Meanwhile, social welfare is simply the expected best available quality since the buyer always purchases and chooses the seller with the highest quality.

5 Comparison with Fixed-Price Benchmarks

We now compare our oligopoly setting to two canonical benchmarks that are widely used in price competition with quality differentiation: (i) *no disclosure with fixed prices*, and (ii) *full disclosure with fixed prices*. These benchmarks are also natural from a regulatory perspective: platforms and regulators routinely constrain sellers’ *pricing flexibility* (e.g., minimum-advertised-price policies, price-parity clauses, and marketplace rules mandating a uniform price across listings) and their *information power* (e.g., mandatory disclosure and standardized labeling rules for product quality).

In summary, our comparison clarifies how constraining sellers along these two dimensions changes market outcomes. Most notably, granting every seller full pricing flexibility can improve market efficiency (Proposition 3), yet, in sufficiently thick markets, it *lowers* each seller’s revenue relative to the fully constrained benchmark featuring full disclosure and fixed pricing (Proposition 4). Although such flexibility is attractive to an individual seller, its market-wide adoption intensifies competition and further compresses equilibrium revenues, leaving all sellers collectively worse off—an *over-competition* effect to which we return after Example 2.

Specifically, we compare equilibrium outcomes using four *ex-ante* metrics: (i) *Market coverage* MC: the probability that the buyer purchases from some seller (rather than the outside option); (ii) *Buyer surplus* U: the buyer’s expected equilibrium net utility; (iii) *Per-seller revenue* R: a

representative seller's expected equilibrium revenue; (iv) *Social welfare* $SW := U + n \cdot R$. For the unrestricted mechanism space $\mathcal{M}$ with n sellers, we denote the resulting equilibrium outcomes by $MC_{\text{gen}}(n)$, $SW_{\text{gen}}(n)$, $U_{\text{gen}}(n)$, and $R_{\text{gen}}(n)$, respectively.

Benchmark 1: No disclosure with fixed prices. Let $\mathcal{M}_{\text{nf}}$ be the decision space, where sellers are prohibited from any disclosure and can only post a constant price that is independent of quality. Since sellers cannot affect the buyer's beliefs, all sellers are ex-ante identical from the buyer's perspective. The induced pricing game therefore reduces to canonical *Bertrand competition* with homogeneous goods (e.g., (Tirole, 1988, Ch. 5)), which has a unique symmetric equilibrium price $p_{\text{nf}}^* = 0$. Hence, trade always occurs, but the outcome may not be efficient because the buyer does not have enough information to select the best product. As for equilibrium outcomes,

$$MC_{\text{nf}}(n) = 1, \quad U_{\text{nf}}(n) = \mathbb{E}[Q], \quad R_{\text{nf}}(n) = 0, \quad SW_{\text{nf}}(n) = \mathbb{E}[Q].$$

Benchmark 2: Full disclosure with fixed prices. Let $\mathcal{M}_{\text{ff}}$ be the decision space, where sellers are compelled to provide full disclosure and post a constant price. Given a common price p , the buyer purchases if and only if $Q_{(1;n)} \geq p$ and, conditional on purchasing, selects the highest-quality seller. This benchmark closely relates to the *Perloff–Salop model* (Perloff and Salop, 1985), augmented with an outside option. Our comparisons below rely on the following maintained assumption.

Assumption 1 (Benchmark equilibrium in $\mathcal{M}_{\text{ff}}$). *For every $n \geq 2$, the game with decision space $\mathcal{M}_{\text{ff}}$ admits a symmetric pure-strategy equilibrium, and its price p_{ff}^* is the unique solution of the first-order condition (FOC_{ff}), displayed below.*

Under Assumption 1, the equilibrium price is characterized by

$$p_{\text{ff}}^* = \frac{(1 - F(p_{\text{ff}}^*)^n) / n}{F(p_{\text{ff}}^*)^{n-1} f(p_{\text{ff}}^*) + (n-1) \int_{p_{\text{ff}}^*}^1 F(x)^{n-2} f(x)^2 dx}. \quad (\text{FOC}_{\text{ff}})$$

Given such an equilibrium price p_{ff}^* , the four metrics are

$$\begin{aligned} MC_{\text{ff}}(n) &= \Pr\{Q_{(1;n)} \geq p_{\text{ff}}^*\} = 1 - F(p_{\text{ff}}^*)^n, \quad U_{\text{ff}}(n) = \mathbb{E}[(Q_{(1;n)} - p_{\text{ff}}^*)_+], \\ R_{\text{ff}}(n) &= \frac{p_{\text{ff}}^*}{n} \Pr\{Q_{(1;n)} \geq p_{\text{ff}}^*\} = \frac{p_{\text{ff}}^*}{n} (1 - F(p_{\text{ff}}^*)^n), \quad SW_{\text{ff}}(n) = \mathbb{E}[Q_{(1;n)} \cdot \mathbf{1}\{Q_{(1;n)} \geq p_{\text{ff}}^*\}]. \end{aligned}$$

Metric 1: Market coverage MC and social welfare SW. We first examine market coverage and social welfare. Together, these metrics summarize the efficiency of trade: market coverage captures the extensive margin of whether trade occurs, whereas social welfare additionally captures the allocative margin of whether the buyer purchases the highest-quality available product.

Proposition 3. *Suppose Assumption 1 holds. The following statements hold for market coverage and social welfare:*

- (i) $1 = \text{MC}_{\text{gen}}(n) = \text{MC}_{\text{nf}}(n) > \text{MC}_{\text{ff}}(n)$ for any n .
- (ii) $\text{SW}_{\text{gen}}(n) > \max\{\text{SW}_{\text{nf}}(n), \text{SW}_{\text{ff}}(n)\}$ for any n .
- (iii) There exists $N \in \mathbb{Z}^+$ such that $\text{SW}_{\text{ff}}(n) > \text{SW}_{\text{nf}}(n)$ for any $n > N$.

The comparison is driven by two frictions: *information* (whether the buyer can learn from sellers' signals and identify the best product) and *price rigidity* (whether sellers can tailor prices to their own realized qualities). In our unrestricted space $\mathcal{M}$, each realized quality q is mapped to a nonnegative buyer net utility $b^*(q) = q - p^*(q) \geq 0$, so the outside option is never chosen and $\text{MC}_{\text{gen}}(n) = 1$. Moreover, $b^*(\cdot)$ is strictly increasing, so the buyer selects the highest-quality seller almost surely. Because transfers cancel in welfare, this yields the first-best allocation and hence $\text{SW}_{\text{gen}}(n) = \mathbb{E}[Q_{(1;n)}]$.




Figure 2: Market coverage and social welfare across different settings (fix a and vary n)

Note. We adopt a prior $F(q) = q^a$ for all $q \in [0, 1]$ with a power parameter $a \in \{0.1, 1.0, 2.0\}$. It holds $\text{MC}_{\text{gen}} = \text{MC}_{\text{nf}} = 1$ and $\text{SW}_{\text{gen}} > \max\{\text{SW}_{\text{nf}}, \text{SW}_{\text{ff}}\}$ across all settings. For any fixed a , as n increases, MC_{ff} converges to 1 and SW_{ff} converges to SW_{gen} both with a rate depending on a .

By contrast, the two fixed-price benchmarks each remove one of these margins. Under $\mathcal{M}_{\text{nf}}$, market coverage is also one (Bertrand competition drives price to 0), but the buyer remains uninformed and therefore cannot systematically select the best realization; consequently, $\text{SW}_{\text{nf}}(n) = \mathbb{E}[Q] < \mathbb{E}[Q_{(1;n)}]$ for $n \geq 2$. Under $\mathcal{M}_{\text{ff}}$, the buyer is perfectly informed and would select the best quality *if trade occurs*, but the common posted price cannot adjust to realized quality. As a result, trade fails whenever $Q_{(1;n)} < p_{\text{ff}}^*$, which occurs with positive probability in any finite- n equilibrium. This is why $\text{MC}_{\text{ff}}(n) < 1$ and the corresponding welfare $\text{SW}_{\text{ff}}(n)$ falls short of the first best $\mathbb{E}[Q_{(1;n)}]$.

Finally, part (iii) formalizes that the welfare loss from price rigidity is primarily a *finite-competition* phenomenon. As n grows, the best available quality $Q_{(1;n)}$ concentrates near the top of the support, and competitive pressure in the fixed-price game makes market exclusion increasingly unlikely. We provide an illustration of the welfare and market-coverage comparison in Figure 2.

Metric 2: Per-seller revenue R. The seller revenue is shaped by how intensely sellers compete

for the sale. In our setting, a seller's pricing incentives are disciplined by two margins. At the *extensive margin*, the seller must beat the outside option to induce trade. At the *intensive margin*, the seller must beat rival offers to win the market. Both margins are governed by the same primitives: the number of sellers n and the distribution of qualities F .

Proposition 4. *Suppose Assumption 1 holds. For per-seller revenue, we have:*

- (i) $0 = R_{\text{nf}}(n) < \min\{R_{\text{ff}}(n), R_{\text{gen}}(n)\}$ for any n .
- (ii) Suppose the prior density f satisfies: (a) f is continuous and strictly positive on $[0, 1]$ (hence bounded and bounded away from zero); (b) f is continuously differentiable around $x = 1$; and (c) $f(1)^2 > f'(1)$, which holds whenever f is log-concave. There exists a threshold $N \in \mathbb{Z}^+$ such that $R_{\text{ff}}(n) \geq R_{\text{gen}}(n)$ for all $n \geq N$.

Part (i) is immediate. Under $\mathcal{M}_{\text{nf}}$, Bertrand competition forces $p_{\text{nf}}^* = 0$, so each seller's expected revenue is identically zero. By contrast, under $\mathcal{M}_{\text{ff}}$ and $\mathcal{M}$, trade in equilibrium (when it occurs) involves a strictly positive expected payment, so each seller earns strictly positive ex-ante revenue for any finite n .




Figure 3: Seller revenue comparison across different settings (fix a and vary n)

Note. We adopt a prior $F(q) = q^a$ for all $q \in [0, 1]$ with a power parameter $a \in \{0.1, 1.0, 2.0\}$. It holds $R_{\text{nf}} = 0$ across all settings. When a and n are both sufficiently small ($a = 0.1$, $n \leq 14$), $R_{\text{ff}} < R_{\text{gen}}$. For any fixed a , it emerges that $R_{\text{ff}} > R_{\text{gen}}$ as n becomes sufficiently large (larger than threshold N).

Part (ii) highlights a seller-side cost of expanded pricing flexibility. When n is large, the outside option becomes increasingly irrelevant in $\mathcal{M}_{\text{ff}}$ because $Q_{(1;n)}$ concentrates near the top of the support, so the equilibrium is primarily pinned down by inter-seller competition. In this regime, allowing transparent contingent pricing in $\mathcal{M}$ intensifies the “race to win” by letting sellers tailor the buyer's net utility to realized quality (equivalently, sellers compete in the surplus they leave the buyer). As a result, equilibrium markups are compressed more strongly than under a fixed-price constraint, and per-seller revenue eventually becomes weakly lower under $\mathcal{M}$ than under $\mathcal{M}_{\text{ff}}$. We provide further numerical illustration in Figure 3.

In Figure 3, we compare revenues across different values of n under a power-law distribution. The figure suggests that when both a and n are small, the per-seller revenue under full disclosure

with a fixed price (the space $\mathcal{M}_{\text{ff}}$) is lower than the revenue in our setting with a general decision space. The intuition is that for small a and n , R_{ff} is depressed by price rigidity, which induces substantial market exclusion (the extensive margin dominates). As n grows and exclusion becomes negligible, the intensive margin dominates, and the revenue ranking in part (ii) emerges.

Metric 3: Buyer surplus U . The buyer surplus is shaped by two forces. First, it depends on the *quality* of the purchased product, which is higher when the buyer can identify and select the best realization among sellers. Second, it depends on the *equilibrium payment*, which is endogenously disciplined by competition.

Proposition 5. *Suppose Assumption 1 holds and the density f is continuous and strictly positive on $[0, 1]$. For buyer surplus, the following statements hold:*

(i) *There exists a threshold $N_1 \in \mathbb{Z}^+$ such that*

$$U_{\text{nf}}(n) \geq U_{\text{gen}}(n), \quad \forall n < N_1 \quad \text{and} \quad U_{\text{nf}}(n) \leq U_{\text{gen}}(n), \quad \forall n > N_1.$$

(ii) *There exists a threshold $N_2 \in \mathbb{Z}^+$ such that $U_{\text{nf}}(n) < U_{\text{ff}}(n)$ for any $n > N_2$.*

(iii) *Suppose f is continuously differentiable around 1 and $f(1)^2 > f'(1)$, which holds whenever f is log-concave. There exists a threshold $N_3 \in \mathbb{Z}^+$ such that $U_{\text{gen}}(n) > U_{\text{ff}}(n)$ for all $n > N_3$.*

The proposition highlights a “thin-market vs. thick-market” contrast. Under $\mathcal{M}_{\text{nf}}$, Bertrand competition drives price to zero, this zero-price advantage usually dominates when n is small: with few sellers, the expected gain from identifying the best realization is limited, so paying nothing usually outweighs the informational benefit.

As n increases, however, the value of information becomes first-order. In both $\mathcal{M}_{\text{ff}}$ and $\mathcal{M}$, the buyer can effectively rank sellers by realized qualities, so selection improves with market thickness. At the same time, competition among sellers suppresses equilibrium prices (and under $\mathcal{M}_{\text{ff}}$ also reduces the incidence of market exclusion). These forces eventually overturn the $\mathcal{M}_{\text{nf}}$ advantage, leading to parts (i) and (ii).

Finally, part (iii) compares the two *full-information* regimes. Even though both $\mathcal{M}_{\text{ff}}$ and $\mathcal{M}$ allow the buyer to identify the top quality, $\mathcal{M}$ gives sellers additional pricing flexibility. This flexibility intensifies inter-seller competition—sellers can tailor prices to realized quality and effectively compete in the net utility they leave the buyer—which compresses markups relative to a fixed-price constraint. Under mild regularity, this pro-competitive effect dominates for large n , so the buyer eventually prefers $\mathcal{M}$ over $\mathcal{M}_{\text{ff}}$. Figure 4 provides a numerical illustration for finite n under a power-law distribution.

Benchmark 3: Arbitrary disclosure with fixed prices. Beyond the two polar benchmarks $\mathcal{M}_{\text{nf}}$ and $\mathcal{M}_{\text{ff}}$, Hwang et al. (2026) studies the model of *arbitrary disclosure with fixed prices*, in which sellers can commit to an arbitrary information structure but are restricted to a single




Figure 4: Buyer surplus comparison across different settings (fix a and vary n)

Note. We adopt a prior $F(q) = q^a$ for all $q \in [0, 1]$ with a power parameter $a \in \{0.1, 1.0, 2.0\}$. It holds U_{nf} is constant across all settings. When n becomes sufficiently large, the inequality $U_{gen} > U_{ff} > U_{nf}$ holds, with the gap $|U_{gen} - U_{ff}|$ vanishing. When n is sufficiently small, U_{nf} becomes the largest.

constant price that is independent of quality. Denote this decision space by $\mathcal{M}_{if}$. The equilibrium in this space is intricate and does not admit a simple characterization, especially when buyers have outside options, as in our setting; see [Hwang et al. \(2026\)](#). Thus, we use an example to illustrate how allowing information design without pricing flexibility compares to the regimes studied above.

Example 2 (Uniform duopoly). *Consider a duopoly with two sellers whose qualities are independently and identically distributed as $Q_1, Q_2 \sim \mathcal{U}[0, 1]$. Table 1 summarizes equilibrium outcomes under the four decision spaces (our general space and three fixed-price benchmarks).*

Table 1: Comparison of equilibrium outcomes ($n = 2, F = \mathcal{U}[0, 1]$)

Space	p^*	MC	U	R	SW
$\mathcal{M}_{nf}$	0	1.000	0.500	0.000	0.500
$\mathcal{M}_{ff}$	0.414	0.828	0.276	0.172	0.619
$\mathcal{M}_{if}$	0.382	0.944	0.269	0.180	0.630
$\mathcal{M}$	$q/2$	1.000	0.333	0.167	0.667

In $\mathcal{M}_{nf}$, Bertrand competition drives the common posted price to 0, so trade occurs surely and the buyer obtains the mean quality $\mathbb{E}[Q] = 1/2$. In $\mathcal{M}_{ff}$, full transparency combined with a fixed posted price creates an exclusion region: whenever both realizations fall below the posted price, the buyer opts out, which lowers market coverage and welfare relative to regimes with flexible pricing.

The intermediate space $\mathcal{M}_{if}$ illustrates how information design can partially mitigate this exclusion effect even when prices are rigid. By pooling sufficiently low-quality realizations with better ones, sellers can raise the trade probability and achieve higher per-seller revenue than under $\mathcal{M}_{ff}$. At the same time, because the buyer's purchase decision need not be ex post incentive compatible, it generates lower surplus U than the full-disclosure benchmark.

Finally, in our general space $\mathcal{M}$, sellers can jointly adjust disclosure and prices to realized quality. As shown in Theorem 2, transparent contingent pricing delivers full market coverage and

the first-best allocation (the highest quality is purchased), yielding $\text{SW}_{\text{gen}} = \mathbb{E}[Q_{(1;2)}] = 2/3$. In this duopoly example, the added pricing flexibility also sharpens inter-seller competition and reduces per-seller revenue relative to $\mathcal{M}_{\text{ff}}$ and $\mathcal{M}_{\text{if}}$. For the buyer, $\mathcal{M}$ dominates the other informative regimes in Table 1, while $\mathcal{M}_{\text{nf}}$ yields the highest buyer surplus here because the zero-price effect is first-order in a thin market. $\square$

Table 1, together with Example 1, provides a clear illustration of the over-competition effect. Under the fixed-price regime $\mathcal{M}_{\text{if}}$, each seller benefits from unilaterally deviating to transparent contingent pricing, as his revenue increases from 0.180 to 0.202 (Example 1). However, when both sellers have this flexibility, each earns only 0.167, less than under either $\mathcal{M}_{\text{if}}$ (0.180) or mandated full disclosure (0.172). Thus, individually attractive flexibility becomes collectively harmful. Sellers face a prisoner’s dilemma in which price rigidity serves as a collective commitment device. Proposition 4(ii) shows that this effect is a general feature of sufficiently thick markets rather than an artifact of the example. A platform that mandates uniform pricing may therefore protect sellers from excessive competition, albeit at buyers’ expense.

6 Extensions to Oligopoly Competition

This section extends our oligopoly analysis along two dimensions. In Section 6.1, we show that transparent contingent pricing remains an equilibrium even when sellers have heterogeneous quality distributions (Proposition 6). In Section 6.2, we introduce buyer-side private information. We first show that the general mechanism space reduces to *garbled contingent pricing*, in which each seller garbles his quality and prices the induced posterior mean (Proposition 7), and we then characterize when full revelation survives as an equilibrium through the convexity of the posterior-mean value function (Proposition 8).

6.1 Ex-ante Asymmetric Sellers

We now relax the ex-ante symmetry assumption among sellers while maintaining the benchmark assumption that the buyer has no private information, $\underline{\theta} = \bar{\theta} = 1$. Specifically, seller i ’s quality is independently drawn from a possibly different distribution F_i on $[0, 1]$.

The previous analysis relies heavily on symmetry to obtain a closed-form pricing rule. The next result shows that symmetry is not essential for the equilibrium disclosure structure: even when sellers have heterogeneous quality distributions, full revelation remains optimal in equilibrium.

Proposition 6 (TCP equilibrium under asymmetric sellers). *Suppose the buyer has no private information. For any collection of independent quality distributions $(F_1, \dots, F_n)$, where each F_i is atomless and admits a continuous strictly positive density on $[0, 1]$, there exists an equilibrium in transparent contingent pricing mechanisms defined in Definition 1.*

This proposition shows that ex-ante asymmetry affects equilibrium pricing but not the equilibrium disclosure structure. Because the buyer's type is degenerate, all sellers' offers are ranked according to a common score, $b_i = m_i - p_i$, which is independent of sellers' identities and prior distributions. Hence, the pricing problem is equivalent to an asymmetric first-price auction in which sellers compete by choosing the net utility left to the buyer. Therefore, the existence of an equilibrium in the corresponding asymmetric auction directly implies the existence of a transparent contingent pricing equilibrium, in which each seller fully reveals quality and adjusts his price accordingly.

However, unlike the i.i.d. benchmark, asymmetric first-price auctions generally do not admit closed-form solutions: equilibrium bidding strategies are characterized by a system of coupled differential equations that must be solved jointly, except in special cases such as $n = 2$ (see [Maskin and Riley \(2000a\)](#)).

6.2 Privately Informed Buyer

We now allow the buyer to have a nondegenerate private type $\theta \sim G \in \Delta([\underline{\theta}, \bar{\theta}])$, independently of $(Q_1, \dots, Q_n)$. This extension captures heterogeneity in the marginal value buyers assign to product quality. A buyer with a larger θ is willing to sacrifice more income for a quality improvement, whereas a buyer with a smaller θ places greater relative weight on price.

The buyer's private information changes the competitive problem in a fundamental way. With a degenerate buyer type, every offer can be ranked by the type-independent score, i.e., the net utility left for the buyer $m_i - p_i$. With a nondegenerate type, the difference between two offers is $\theta(m_i - m_j) - (p_i - p_j)$, whose sign may vary with θ . The same pair of offers can therefore be ranked differently by different buyer types, and the seller's problem no longer collapses to a single one-dimensional bid. However, we can perform the following equivalent reduction.

Definition 2 (Garbled contingent pricing (GCP)). *A mechanism is a garbled contingent pricing mechanism if the seller chooses an information structure $\pi_i : [0, 1] \rightarrow \Delta(\mathcal{S}_i)$ and a deterministic pricing rule $p_i : \mathcal{S}_i \rightarrow \mathcal{P}$. Given quality q_i , the seller draws $s_i \sim \pi_i(\cdot \mid q_i)$ and posts the price $p_i(s_i)$.*

Let $\text{MPC}(F)$ denote the set of all mean-preserving contraction distributions for prior F . A canonical GCP mechanism is then summarized by a feasible posterior-mean distribution $H \in \text{MPC}(F)$ and a deterministic price function $p(m)$ based on posterior mean m .

Proposition 7 (Reduction to garbled contingent pricing). *The following statements hold.*

- (i) *Each mechanism in $\mathcal{M}_{\text{gen}}$ admits an outcome-equivalent representation as a GCP mechanism.*
- (ii) *Fix seller i and an arbitrary profile κ_{-i} of rivals. Let $\Pi_i(m, p; \kappa_{-i})$ denote seller i 's expected revenue from inducing posterior mean m and posting price p . Define the pointwise optimized revenue $V_i(m; \kappa_{-i}) := \sup_{p \in \mathcal{P}} \Pi_i(m, p; \kappa_{-i})$. Restricting attention to GCP mechanisms entails*

no loss of achievable revenue:

$$\sup_{\kappa_i \in \mathcal{M}_{\text{gen}}} R_i(\kappa_i, \kappa_{-i}) = \sup_{H \in \text{MPC}(F)} \int_{[0,1]} V_i(m; \kappa_{-i}) dH(m).$$

Although GCP substantially narrows the form of an optimized deviation, a full equilibrium characterization remains difficult. We therefore turn to a natural benchmark that isolates the remaining role of information design. To determine whether transparent contingent pricing continues to arise with a privately informed buyer, we introduce an auxiliary *full-information pricing game* in which every seller is compelled to fully reveal realized quality and can choose only a deterministic quality-contingent price function $\chi_i : [0, 1] \rightarrow \mathcal{P}$.

This auxiliary game provides a necessary benchmark for any TCP equilibrium of the general mechanism game. If a profile of transparent contingent pricing mechanisms constitutes an equilibrium of the general game, the associated price functions must form an equilibrium of the auxiliary game. The converse need not hold. We therefore proceed in two stages: first characterize equilibria of the full-information pricing game, and then determine which of them remain stable when sellers can also design information. If the auxiliary game has no equilibrium, or if none of its equilibria satisfies the survival condition developed below, then we can conclude that the general mechanism game admits no TCP equilibrium.

Proposition 8 (When full-information equilibrium survives). *Suppose the full-information pricing game admits a symmetric equilibrium χ^* . Assume that F is atomless with full support on $[0, 1]$. Fix seller i , hold all rivals at full revelation and χ^* , and define the posterior-mean value function*

$$V(m) := \sup_{p \geq 0} p \cdot \Pr_{\theta \sim G, Q_{-i} \sim F^{n-1}} \left(\theta m - p \geq \max \left\{ 0, \max_{j \neq i} (\theta Q_j - \chi^*(Q_j)) \right\} \right).$$

Then the full revelation strategy together with χ^ constitutes a symmetric equilibrium of the general mechanism game if and only if the function V is convex on $[0, 1]$.*

The characterization above nests the baseline model considered in Section 4 as a special case. When the buyer's type is degenerate, the posterior mean quality plays the same role as the realized quality in the full-information pricing game. In particular, when $\theta \equiv 1$, the posterior-mean value function under the TCP equilibrium characterized in Theorem 2 reduces to $V(m) = \int_0^m F(t)^{n-1} dt$, which is convex almost everywhere. Therefore, the convexity condition in Proposition 8 is always satisfied, and the TCP equilibrium survives in the general mechanism space.

With private buyer types, the posterior-mean value function $V(m)$ need not be convex. Consequently, full revelation and the associated transparent contingent pricing equilibrium may fail to survive in some instances. The next example first identifies a structured environment in which it does, thereby showing that private buyer information does not by itself eliminate transparent contingent pricing.

Example 3 (Full revelation under a binary private buyer type). *Consider a duopoly in which sellers' qualities are independently and identically distributed according to the uniform distribution on $[0, 1]$. The buyer has a binary private type, $\theta \in \{\theta_l, \theta_r\}$, where $0 < \theta_r/2 \leq \theta_l < \theta_r$ and $\Pr\{\theta = \theta_r\} = \lambda \in (0, 1)$. Let $a^* \in (0, \theta_l)$ be the unique solution to*

$$1 = a^* \left(\frac{\lambda}{\theta_r - a^*} + \frac{1 - \lambda}{\theta_l - a^*} \right).$$

Then the full-information pricing game admits a symmetric equilibrium with a linear pricing rule $\chi^(q) = a^*q$. Moreover, the associated posterior-mean value function is*

$$V(m) = a^*m^2, \quad \forall m \in [0, 1].$$

Since V is convex, Proposition 8 implies that full revelation together with the linear pricing rule χ^ remains a symmetric equilibrium of the general mechanism game.*

The uniform-prior example shows how full revelation can survive private buyer information. This property is not generic. The following example constructs a regular i.i.d. environment in which the posterior-mean value function is locally concave and pooling is strictly profitable.

Example 4 (An example where full revelation fails in equilibrium). *There are two ex-ante symmetric sellers whose qualities are independently and identically distributed according to*

$$F(q) = q \left(\frac{20 - 9q^2}{11} \right)^{\frac{3}{26}} \left(\frac{q^2 + 5}{6} \right)^{-\frac{55}{26}}, \quad q \in [0, 1].$$

The buyer's private type is independent of product qualities and satisfies $\Pr\{\theta = \frac{5}{9}\} = \Pr\{\theta = 1\} = \frac{1}{2}$. In the full-information pricing game, the symmetric pricing rule $\chi^(q) = \frac{q}{3} + \frac{q^3}{15}$ constitutes an equilibrium. Nevertheless, this equilibrium does not survive in the general mechanism space: a seller can profitably pool all qualities in $[8/9, 1]$ and post a single price contingent on the posterior mean induced by the pooled signal.*

The positive and negative examples isolate the role of the posterior-mean value function. Under the uniform prior, the posterior-mean value function $V(m) = a^*m^2$ is convex, so the seller prefers the dispersion generated by full revelation. In Example 4, the function $V(m)$ is locally concave over $(8/9, 1)$, so pooling these high qualities raises the seller's optimized payoff.

Private buyer information therefore changes the equilibrium role of disclosure. Because the buyer's ranking of offers depends on her type, a seller may use garbling to reshape which buyer types are attracted by a given price. Full revelation is not generally without loss, even under ex-ante symmetry, i.i.d. qualities, and a regular full-support prior. At the same time, private buyer information does not invariably eliminate transparent contingent pricing.

7 Conclusion

We study markets in which privately informed sellers commit to mechanisms that jointly govern what buyers learn and what they pay. The analysis delivers three results. For a monopolist, information design is payoff-irrelevant: any disclosure policy attains the optimal revenue once the price moves proportionally with the induced posterior mean. Under competition, the vast mechanism space collapses to transparent contingent pricing: whatever rivals do, a seller loses nothing by fully revealing his quality and pricing deterministically, and in the resulting symmetric equilibrium the buyer purchases the best available product at first-price-auction markups. Besides, when the prior is log-concave, the price tag alone carries all quality information. Comparing decision spaces, pricing flexibility raises social welfare and, in thick markets, buyer surplus, yet leaves every seller with less revenue than a fixed-price regime would: an *over-competition* effect whereby individually optimal flexibility is collectively self-defeating for sellers. Once pricing is flexible, disclosure mandates are redundant because competition itself makes sellers transparent; rules that freeze prices instead restore the strategic value of obfuscation. A platform's incentive to permit contingent pricing depends on whether it profits from transaction volume or from seller revenue.

References

- Shipra Agrawal, Yiding Feng, and Wei Tang. Dynamic pricing and advertising with demand learning. *arXiv e-prints*, pages arXiv–2304, 2023.
- Mark Armstrong and Jidong Zhou. Consumer information and the limits to competition. *American Economic Review*, 112(2):534–577, 2022.
- Pak Hung Au and Keiichi Kawai. Competitive information disclosure by multiple senders. *Games and Economic Behavior*, 119:56–78, 2020.
- Kyle Bagwell and Michael H Riordan. High and declining prices signal product quality. *The American Economic Review*, pages 224–239, 1991.
- Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.
- Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–957, 2015.
- David Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, volume 2, pages 93–103, 1951.
- Raphael Boleslavsky and Christopher Cotton. Grading standards and education quality. *American Economic Journal: Microeconomics*, 7(2):248–279, 2015.
- Ozan Candogan. Information design in operations. In *Pushing the Boundaries: Frontiers in Impactful OR/OM Research*, pages 176–201. INFORMS, 2020.
- Zhicheng Du, Wei Tang, Zihe Wang, and Shuo Zhang. Competitive information design with asymmetric senders. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 1203–1203, 2024.
- Matthew Gentzkow and Emir Kamenica. Competition in persuasion. *The Review of Economic Studies*, 84(1):300–322, 2016.
- Matthew Gentzkow and Emir Kamenica. Bayesian persuasion with multiple senders and rich signal spaces. *Games and Economic Behavior*, 104:411–429, 2017.
- Safwan Hossain, Tonghan Wang, Tao Lin, Yiling Chen, David C Parkes, and Haifeng Xu. Multi-sender persuasion: A computational perspective. In *International Conference on Machine Learning*, pages 18944–18971. PMLR, 2024.
- Harold Hotelling. Stability in competition. *The economic journal*, 39(153):41–57, 1929.
- Ilwoo Hwang, Kyungmin Kim, and Raphael Boleslavsky. Competitive advertising and pricing. *The Review of Economic Studies*, page rdag062, 2026. doi: 10.1093/restud/rdag062.
- Maxim Ivanov. Information revelation in competitive markets. *Economic Theory*, 52(1):337–365, 2013.
- Vasudha Jain and Mark Whitmeyer. Competing to persuade a rationally inattentive agent. *arXiv preprint arXiv:1907.09255*, 2019.
- Emir Kamenica. Bayesian persuasion and information design. *Annual Review of Economics*, 11(1):249–272, 2019.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Ming Li, Binyan Pu, and Renkun Yang. The design of quality disclosure policy and the limits to competition.

- In *Proceedings of the 26th ACM Conference on Economics and Computation*, pages 640–640, 2025.
- Zi Ling, Chris Ryan, and Mengzhenyu Zhang. On the interplay between pricing and information policies under social learning. *Available at SSRN 5047585*, 2024.
- David Lingenbrink and Krishnamurthy Iyer. Optimal signaling mechanisms in unobservable queues. *Operations research*, 67(5):1397–1416, 2019.
- Eric Maskin and John Riley. Monopoly with incomplete information. *The RAND Journal of Economics*, 15(2):171–196, 1984.
- Eric Maskin and John Riley. Asymmetric auctions. *The review of economic studies*, 67(3):413–438, 2000a.
- Eric Maskin and John Riley. Equilibrium in sealed high bid auctions. *The Review of Economic Studies*, 67(3):439–454, 2000b.
- Paul Milgrom and John Roberts. Price and Advertising Signals of Product Quality. *Journal of Political Economy*, 94(4):796–821, August 1986.
- Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- Jeffrey M Perloff and Steven C Salop. Equilibrium with product differentiation. *The Review of Economic Studies*, 52(1):107–120, 1985.
- Dilip Ravindran and Zhihan Cui. Competing persuaders in zero-sum games. *arXiv preprint arXiv:2008.08517*, 2020.
- Daniel F Spulber. Bertrand competition when rivals’ costs are unknown. *The Journal of Industrial Economics*, 43(1):1–11, 1995.
- Yanwei Sun, Niloofar Zamani-Foroushani, Zhe Liu, and Jiahua Wu. Information design and pricing in lead-selling platforms with supply competition. *Available at SSRN 4872681*, 2024.
- Jean Tirole. *The theory of industrial organization*. MIT press, 1988.
- Asher Wolinsky. Prices as signals of product quality. *The review of economic studies*, 50(4):647–658, 1983.

A Proofs of Section 3

Proof of Theorem 1. The proof first derives an upper bound on the seller's revenue conditional on each realized message. We then aggregate this bound using Bayes plausibility and show that it is attained by pairing any feasible information structure with proportional posterior-mean pricing.

Step 1: Derive a messagewise upper bound. This step bounds the seller's revenue after each realized signal–price pair. The key is to express the conditional revenue as the posterior mean multiplied by a one-dimensional cutoff-revenue term.

Fix an arbitrary mechanism $\kappa \in \mathcal{M}$. Under the joint distribution induced by κ , quality is drawn according to $q \sim F$, and the signal–price pair is subsequently drawn according to $(s, p) \sim \kappa(\cdot \mid q)$. Fix a realized pair (s, p) and write $m := m(s, p)$ for the induced posterior mean.

If $m > 0$, the buyer purchases if and only if her type exceeds the cutoff $\tau := p/m$. Because θ is independent of both quality and the mechanism's randomization, observing (s, p) does not change its distribution. Hence, letting $\text{Rev}(s, p)$ denote the seller's expected revenue conditional on the realized message (s, p) , we have

$$\text{Rev}(s, p) = \mathbb{E}_{\theta \sim G}[p \cdot \mathbf{1}\{\theta m \geq p\}] = p \cdot \bar{G}(\tau) = m \cdot \tau \bar{G}(\tau).$$

The maximum defining θ^* is attained: $\theta \mapsto \theta \bar{G}(\theta)$ is upper semicontinuous on the compact interval $[\underline{\theta}, \bar{\theta}]$, because $\bar{G}(\theta) = \Pr_G\{\theta' \geq \theta\}$ is nonincreasing and left-continuous. By the definition of θ^* , $\tau \bar{G}(\tau) \leq \theta^* \bar{G}(\theta^*)$ for every $\tau \geq 0$. This comparison also covers cutoffs outside the support of G : $\bar{G}(\tau) = 1$ for $\tau < \underline{\theta}$, whereas $\bar{G}(\tau) = 0$ for $\tau > \bar{\theta}$. If $m = 0$, the conditional revenue is zero, because a positive price induces no purchase and a zero price generates no revenue. Therefore, every realized message satisfies

$$\text{Rev}(s, p) \leq m(s, p) \cdot \theta^* \bar{G}(\theta^*).$$

Step 2: Aggregate the upper bound. This step converts the messagewise bound into an ex-ante revenue bound. We take expectations under the joint distribution induced by F and κ and then apply the law of iterated expectations.

For any $\kappa \in \mathcal{M}$, we obtain the following upper bound on expected revenue

$$\begin{aligned} R(\kappa) &= \mathbb{E}_{q \sim F, (s, p) \sim \kappa(\cdot \mid q)}[\text{Rev}(s, p)] \\ &\leq \theta^* \bar{G}(\theta^*) \cdot \mathbb{E}_{q \sim F, (s, p) \sim \kappa(\cdot \mid q)}[m(s, p)] \\ &= \theta^* \bar{G}(\theta^*) \cdot \mathbb{E}_{q \sim F, (s, p) \sim \kappa(\cdot \mid q)}[\mathbb{E}[q \mid s, p]] \\ &= \theta^* \bar{G}(\theta^*) \cdot \mathbb{E}_{q \sim F}[q] \\ &= m_0 \theta^* \bar{G}(\theta^*), \end{aligned}$$

where $m_0 := \mathbb{E}_{q \sim F}[q]$. Because $\kappa \in \mathcal{M}$ was arbitrary, no joint information–pricing mechanism can

generate expected revenue above $m_0\theta^*\bar{G}(\theta^*)$.

Step 3: Attain the bound under any information structure. We next establish that the upper bound is tight for every feasible information structure. We pair an arbitrary information structure with a deterministic price proportional to its induced posterior mean.

Fix any feasible information structure $(\mathcal{S}, \pi)$, where $q \sim F$ and $s \sim \pi(\cdot | q)$. Let $m(s) := \mathbb{E}[q | s]$ denote the posterior mean under the joint distribution induced by F and π . Construct a joint mechanism κ^π that draws s according to $\pi(\cdot | q)$ and posts the deterministic price $p^*(m(s)) := \theta^*m(s)$. Because the posted price is a deterministic function of s , observing the pair $(s, p^*(m(s)))$ leaves the posterior mean equal to $m(s)$.

Whenever $m(s) > 0$, the pricing rule induces the purchase cutoff $p^*(m(s))/m(s) = \theta^*$. When $m(s) = 0$, the posted price is also zero, so this realization contributes no revenue. Using the independence of θ from (q, s) , the induced joint mechanism earns

$$\begin{aligned} R(\kappa^\pi) &= \mathbb{E}_{q \sim F, s \sim \pi(\cdot | q), \theta \sim G}[\theta^*m(s) \cdot \mathbf{1}\{\theta \geq \theta^*\}] \\ &= \theta^* \cdot \mathbb{E}_{\theta \sim G}[\mathbf{1}\{\theta \geq \theta^*\}] \cdot \mathbb{E}_{q \sim F, s \sim \pi(\cdot | q)}[m(s)] \\ &= \theta^*\bar{G}(\theta^*) \cdot \mathbb{E}_{q \sim F, s \sim \pi(\cdot | q)}[\mathbb{E}[q | s]] \\ &= \theta^*\bar{G}(\theta^*) \cdot \mathbb{E}_{q \sim F}[q] \\ &= m_0\theta^*\bar{G}(\theta^*). \end{aligned}$$

Thus, the proportional pricing rule attains the upper bound under every feasible information structure. Consequently,

$$\max_{\kappa \in \mathcal{M}} R(\kappa) = m_0\theta^*\bar{G}(\theta^*),$$

and the seller solves Problem $\mathcal{P}_{\text{mono}}$ by charging $p^*(m) = \theta^*m$ after every posterior mean m . $\square$

B Proofs of Section 4

Proof of Theorem 2. The proof has two parts. First, fixing the rivals' mechanisms, we reduce seller i 's problem to a statewise choice of the buyer's net utility and show that the resulting upper bound is attained by transparent contingent pricing, proving part (i). Second, we derive and verify a symmetric pricing rule whose induced score is strictly increasing, so the highest-quality seller wins almost surely and the outside option is never chosen, proving part (ii).

Step 1: Reduce the joint mechanisms to scores. The first step is to show that, holding the rivals' mechanisms fixed, an arbitrary joint information–pricing mechanism affects seller i 's revenue only through the net utility it induces for the buyer.

Fix seller i and an arbitrary strategy profile κ_{-i} of rival mechanisms. Let (s_i, p_i) denote seller i 's realized signal–price pair, let $m_i := \mathbb{E}[q_i | s_i, p_i]$ be the induced posterior mean, and define

$B_i := m_i - p_i$ as the buyer's net utility from seller i , which we also refer to as seller i 's score. Recall that $H(b)$ is the probability that a nonnegative score b weakly exceeds all rival scores.

Conditional on (s_i, p_i) , both p_i and B_i are fixed. Moreover, seller i 's message is independent of the rivals' realized offers. For expositional convenience and to ensure that a best response is attained even when the distribution of rival scores has atoms, we use the seller-favorable tie-breaking convention of Section 4.1. Specifically, when evaluating seller i 's payoff, every tie involving seller i is resolved in his favor. Under any other fixed tie-breaking rule, the same reduction guarantees a transparent contingent pricing ε -best response for every $\varepsilon > 0$; an exact transparent contingent pricing best response continues to exist whenever the distribution of B_{-i} is atomless. Under the maintained convention, seller i is chosen with conditional probability $H(B_i)$ when $B_i \geq 0$ and with probability zero otherwise. By the law of iterated expectations,

$$\begin{aligned} R_i(\kappa_i, \kappa_{-i}) &= \mathbb{E}_{q_i \sim F, (s_i, p_i) \sim \kappa_i(\cdot | q_i)} [p_i \cdot H(B_i) \cdot \mathbf{1}\{B_i \geq 0\}] \\ &= \mathbb{E}_{q_i \sim F, (s_i, p_i) \sim \kappa_i(\cdot | q_i)} [(q_i - B_i) \cdot H(B_i) \cdot \mathbf{1}\{B_i \geq 0\}]. \end{aligned}$$

To establish the second equality, define $g_i := H(B_i)\mathbf{1}\{B_i \geq 0\}$. Because B_i is determined by (s_i, p_i) , the variable g_i is measurable with respect to (s_i, p_i) . Using $p_i = m_i - B_i$ and $m_i = \mathbb{E}[q_i | s_i, p_i]$, the tower property implies

$$\begin{aligned} \mathbb{E}_{q_i \sim F, (s_i, p_i) \sim \kappa_i(\cdot | q_i)} [m_i g_i] &= \mathbb{E}_{q_i \sim F, (s_i, p_i) \sim \kappa_i(\cdot | q_i)} [\mathbb{E}[q_i | s_i, p_i] g_i] \\ &= \mathbb{E}_{q_i \sim F, (s_i, p_i) \sim \kappa_i(\cdot | q_i)} [q_i g_i]. \end{aligned}$$

Substituting this identity into the term $\mathbb{E}_{q_i \sim F, (s_i, p_i) \sim \kappa_i(\cdot | q_i)} [(m_i - B_i) g_i]$ yields the second equality. This representation replaces the posterior mean with realized quality and reduces the seller's problem to the problem of choosing the score B_i given any realized quality q_i .

For each realized quality $q_i \in [0, 1]$, define the statewise value

$$V(q_i) := \max_{b \in [0, q_i]} (q_i - b)H(b).$$

This value is the highest revenue seller i can obtain by directly choosing the buyer's score after observing q_i . Restricting the choice to $b \in [0, q_i]$ is without loss: a negative score cannot beat the outside option and therefore generates no revenue, while any $b > q_i$ yields $(q_i - b)H(b) \leq 0$ and is weakly dominated by $b = q_i$. Moreover, because H is a CDF, it is nondecreasing and right-continuous. At every $b \in [0, q_i]$, the objective $(q_i - b)H(b)$ is therefore continuous from the right, while its left limit is $(q_i - b)H(b^-) \leq (q_i - b)H(b)$. Hence, the objective is upper semicontinuous on the compact interval $[0, q_i]$, so a maximizer exists. Applying this statewise bound to the score representation above and conditioning on q_i gives the following revenue upper bound

$$\begin{aligned} R_i(\kappa_i, \kappa_{-i}) &= \mathbb{E}_{q_i \sim F} \left[\mathbb{E}_{(s_i, p_i) \sim \kappa_i(\cdot | q_i)} [(q_i - B_i) \cdot H(B_i) \cdot \mathbf{1}\{B_i \geq 0\}] \right] \\ &\leq \mathbb{E}_{q_i \sim F} [V(q_i)]. \end{aligned}$$

Step 2: Construct a transparent best response. We show that the statewise upper bound derived in Step 1 is attainable through transparent contingent pricing. The construction fully reveals realized quality and uses the price to implement a revenue-maximizing score in each state.

For every $q_i \in [0, 1]$, define the score function

$$b^\dagger(q_i) := \min_{b \in [0, q_i]} \arg\max (q_i - b)H(b).$$

This selection is well defined because Step 1 shows that the argmax set is nonempty and compact. We next verify that $b^\dagger$ is a measurable function. In fact, it is nondecreasing. To see this, fix $0 \leq q < q' \leq 1$ and write $x := b^\dagger(q)$ and $y := b^\dagger(q')$. Suppose, toward a contradiction, that $y < x$. Then both x and y are feasible at both qualities, and optimality gives

$$(q - x)H(x) \geq (q - y)H(y) \quad \text{and} \quad (q' - y)H(y) \geq (q' - x)H(x).$$

Adding these inequalities yields

$$(q' - q)(H(y) - H(x)) \geq 0.$$

Because H is nondecreasing and $y < x$, the left-hand side is also nonpositive. Hence both optimality inequalities must bind. In particular, y is also a maximizer at quality q , contradicting $y < x = b^\dagger(q)$ and the definition of $b^\dagger(q)$ as the smallest maximizer. Therefore, $b^\dagger$ is Borel measurable.

Consider the mechanism that fully reveals q_i and posts the deterministic price $p_i^\dagger(q_i) := q_i - b^\dagger(q_i)$. This price is nonnegative because $b^\dagger(q_i) \in [0, q_i]$. Under this mechanism, full revelation implies $m_i = q_i$, so the buyer's score is $B_i = m_i - p_i^\dagger(q_i) = b^\dagger(q_i)$. Conditional on q_i , seller i is therefore chosen with probability $H(b^\dagger(q_i))$ and earns

$$p_i^\dagger(q_i)H(b^\dagger(q_i)) = (q_i - b^\dagger(q_i))H(b^\dagger(q_i)) = V(q_i).$$

Taking expectations, the mechanism yields $\mathbb{E}[V(q_i)]$, attaining the upper bound established in Step 1. Hence, for any profile κ_{-i} of rival mechanisms, seller i has a best response within the family of transparent contingent pricing mechanisms. This proves part (i).

Step 3: Derive the symmetric candidate. We derive a candidate symmetric pricing rule within the family of transparent contingent pricing mechanisms. We parameterize each score choice by the quality type whose equilibrium score it mimics and then solve the resulting first-order condition.

Under transparent contingent pricing, seller i with realized quality q_i chooses a score $b \in [0, q_i]$ for the buyer and charges the residual price $p = q_i - b$. Suppose all rivals use the same strictly increasing score schedule $b(\cdot)$. For any pseudo-quality $x \in [0, 1]$, consider a deviation in which seller i offers the score $b(x)$. Because $b(\cdot)$ is strictly increasing, this score beats rival j 's offer if and only if $q_j \leq x$. Independence across qualities then implies that seller i wins with probability $F(x)^{n-1}$.

Conditional on quality q_i , the revenue from mimicking type x is therefore

$$W_{q_i}(x) := (q_i - b(x)) \cdot F(x)^{n-1}.$$

To construct a symmetric candidate, impose that each seller optimally chooses the score prescribed for its own quality, so that $x = q_i$ maximizes $W_{q_i}(x)$. At every point of differentiability, the associated first-order condition is

$$0 = W'_{q_i}(q_i) = -b'(q_i) \cdot F(q_i)^{n-1} + (n-1) \cdot (q_i - b(q_i)) \cdot F(q_i)^{n-2} \cdot f(q_i).$$

Let $p(q_i) := q_i - b(q_i)$ denote the corresponding price schedule. Since $b'(q_i) = 1 - p'(q_i)$, the first-order condition reduces to the linear differential equation

$$p'(q_i) + \frac{(n-1)f(q_i)}{F(q_i)} \cdot p(q_i) = 1 \quad \text{for a.e. } q_i \in (0, 1].$$

Multiplying both sides by the integrating factor $F(q_i)^{n-1}$ gives $\frac{d}{dq_i} [p(q_i)F(q_i)^{n-1}] = F(q_i)^{n-1}$. Integrating from 0 to q_i yields

$$p(q_i) \cdot F(q_i)^{n-1} = \int_0^{q_i} F(t)^{n-1} dt + C.$$

Feasibility under transparent contingent pricing requires $0 \leq p(q_i) \leq q_i$. Because feasibility bounds $0 \leq p(q_i) \leq q_i \leq 1$ and $F(q_i) \rightarrow 0$ as $q_i \downarrow 0$, it follows that $p(q_i)F(q_i)^{n-1} \rightarrow 0$. The integral on the right-hand side also converges to zero, so the boundary condition implies $C = 0$. The candidate pricing rule is therefore

$$p(q_i) = \frac{\int_0^{q_i} F(t)^{n-1} dt}{F(q_i)^{n-1}} = p^*(q_i), \quad q_i \in (0, 1],$$

which coincides with (2). The next step verifies that the induced score schedule is strictly increasing and that p^* is globally optimal against itself.

Step 4: Verify global optimality. We next verify that the candidate pricing rule derived in Step 3 is a global best response to itself. We first establish that the induced score schedule is feasible and strictly increasing, and then use this monotonicity to parameterize every score deviation by a pseudo-quality.

Define $p^*(0) := 0$ and $b^*(q_i) := q_i - p^*(q_i)$ for all $q_i \in [0, 1]$. For every $q_i > 0$, the numerator in (2) is strictly positive, so $p^*(q_i) > 0$. Moreover, because F is strictly increasing, $\int_0^{q_i} F(t)^{n-1} dt < q_i \cdot F(q_i)^{n-1}$, which implies $p^*(q_i) < q_i$. Hence, $b^*(q_i) \in (0, q_i)$ for every $q_i > 0$.

We next establish that b^* is strictly increasing. Differentiating the identity $p^*(q_i) \cdot F(q_i)^{n-1} = \int_0^{q_i} F(t)^{n-1} dt$ at points of differentiability yields

$$b^{*'}(q_i) = \frac{(n-1)f(q_i)}{F(q_i)} \cdot p^*(q_i) > 0 \quad \text{for a.e. } q_i \in (0, 1].$$

Because b^* is locally absolutely continuous on $(0, 1]$, this derivative condition implies that it is

strictly increasing on every compact subinterval of $(0, 1]$. Moreover, $0 < b^*(q_i) < q_i$ implies $b^*(q_i) \rightarrow 0$ as $q_i \downarrow 0$. Thus, b^* is strictly increasing on $(0, 1]$.

Now suppose all rivals use the pricing rule p^* , or equivalently the score schedule b^* . Fix seller i with realized quality $q_i > 0$ and consider an arbitrary deterministic score $y \in [0, q_i]$. It is sufficient to consider deterministic scores because the seller's expected revenue is linear in the distribution over score realizations.

Consider first a score $y \in [0, b^*(1)]$. Because b^* is strictly increasing, there exists a unique pseudo-quality $x \in [0, 1]$ such that $y = b^*(x)$. Seller i then beats each rival if and only if that rival's quality is at most x , so the probability of winning is $F(x)^{n-1}$. The resulting revenue is $W_{q_i}(x) := (q_i - b^*(x)) \cdot F(x)^{n-1}$. Using the derivative formula for b^* and the identity $p^*(x) = x - b^*(x)$, we obtain

$$\begin{aligned} W'_{q_i}(x) &= -b'^*(x) \cdot F(x)^{n-1} + (q_i - b^*(x)) \cdot (n-1)F(x)^{n-2}f(x) \\ &= (n-1)f(x)F(x)^{n-2} \cdot (q_i - x) \quad \text{for a.e. } x \in (0, 1]. \end{aligned}$$

Therefore, W_{q_i} is increasing for $x < q_i$ and decreasing for $x > q_i$. Since W_{q_i} is absolutely continuous, it is globally maximized at $x = q_i$. The optimal score within the range of b^* is consequently $y = b^*(q_i)$.

It remains to consider a feasible score $y > b^*(1)$. Such a score wins with probability one, but yields revenue $q_i - y < q_i - b^*(1) = W_{q_i}(1) \leq W_{q_i}(q_i)$. Thus, no score above the equilibrium range is profitable. The case $q_i = 0$ is immediate, because the only feasible nonnegative price and score generate zero revenue.

It follows that, for every realized quality q_i , the score $b^*(q_i)$ maximizes the statewise objective $(q_i - b) \cdot H(b)$ over $b \in [0, q_i]$. The corresponding transparent contingent pricing rule $p^*(q_i) = q_i - b^*(q_i)$ therefore attains the upper bound derived in Step 1. Hence, p^* is a best response to itself not only within transparent contingent pricing mechanisms, but within the full mechanism space $\mathcal{M}$. The resulting symmetric strategy profile is therefore an equilibrium.

Step 5: Characterize the equilibrium allocation. This step derives the allocation induced by the symmetric equilibrium. The argument uses the strict monotonicity and positivity of the equilibrium score schedule established in Step 4.

Because $b^*(\cdot)$ is strictly increasing, equilibrium scores preserve the ranking of realized qualities: a seller with a higher quality offers the buyer a strictly higher net utility. Since F is atomless, the highest realized quality is unique almost surely. The seller with the highest quality therefore also offers the highest score, so the buyer purchases the highest-quality product almost surely.

It remains to rule out the outside option. Step 4 establishes that $b^*(q_i) > 0$ for every $q_i > 0$. Moreover, atomlessness of F implies $\Pr\{q_i = 0\} = 0$, so at least one seller offers a strictly positive score almost surely. Because the outside option yields zero utility, the buyer never selects it. This proves part (ii) and completes the proof. $\square$

Proof of Proposition 1. The proof uses the observation that the equilibrium price is the reciprocal of the slope of the log cumulative payoff function. This identity yields the exact monotonicity condition, while a tangent-line argument establishes the stated sufficient conditions.

Step 1: Characterize price monotonicity. This step links the monotonicity of p^* directly to the log-concavity condition in the proposition. The argument rewrites the equilibrium price in terms of the logarithmic derivative of a single auxiliary function.

Define $A(q) := \int_0^q F(t)^{n-1} dt$. Because $F(q) > 0$ for every $q > 0$, we have $A(q) > 0$ and $A'(q) = F(q)^{n-1} > 0$ on $(0, 1]$. By Equation (2), we have

$$p^*(q) = \frac{A(q)}{A'(q)}, \quad (\log A(q))' = \frac{A'(q)}{A(q)} = \frac{1}{p^*(q)}.$$

Because $p^*(q) > 0$, the price schedule p^* is strictly increasing if and only if $1/p^*$ is strictly decreasing. The identity above shows that this is equivalent to $(\log A)'$ being strictly decreasing, which, since $\log A$ is differentiable, is equivalent to $\log A$ being strictly concave. Therefore, p^* is strictly increasing on $(0, 1]$ if and only if $A(q) = \int_0^q F(t)^{n-1} dt$ is strictly log-concave in q .

Step 2: Verify the sufficient conditions. We next show that strict log-concavity of F guarantees strict monotonicity of the equilibrium price. We use the supporting-tangent inequality for strictly log-concave functions and then verify that strict log-concavity of f implies the same property for F .

Suppose first that F is strictly log-concave on $(0, 1]$. Let $g(q) := F(q)^{n-1}$ and define $\lambda(q) := (\log g(q))' = (n-1)f(q)/F(q)$ at points of differentiability. The function g is strictly log-concave, and $\lambda(q) > 0$ under the maintained assumption that f is strictly positive. Thus, for every $t \in (0, q)$,

$$g(t) < g(q) \cdot e^{(\lambda(q)(t-q))}.$$

Integrating this strict tangent bound gives

$$p^*(q) = \frac{\int_0^q g(t) dt}{g(q)} < \frac{1 - e^{-\lambda(q)q}}{\lambda(q)} < \frac{1}{\lambda(q)}.$$

Differentiating $p^*(q) = A(q)/g(q)$ at points of differentiability yields $p^{*'}(q) = 1 - \lambda(q)p^*(q)$. The preceding bound therefore implies $p^{*'}(q) > 0$ for almost every $q \in (0, 1]$. Because p^* is locally absolutely continuous, it follows that p^* is strictly increasing on $(0, 1]$.

It remains to show that strict log-concavity of f implies strict log-concavity of F . Suppose that f is strictly log-concave, and define the reverse hazard rate

$$\rho(q) := f(q)/F(q) = (\log F(q))'.$$

At a point where $\log f$ is differentiable, let $\eta(q) := (\log f(q))'$. If $\eta(q) \leq 0$, then $\rho(q) > 0 \geq \eta(q)$. If

$\eta(q) > 0$, strict log-concavity of f gives $f(t) < f(q) \exp(\eta(q)(t - q))$ for every $t < q$, and hence

$$F(q) < f(q)(1 - e^{-\eta(q)q})/\eta(q) < f(q)/\eta(q).$$

Thus, $\rho(q) > \eta(q)$ in either case. Consequently,

$$\rho'(q) = \rho(q)(\eta(q) - \rho(q)) < 0 \quad \text{for a.e. } q \in (0, 1].$$

Since $\rho = (\log F)'$ is strictly decreasing, $\log F$ is strictly concave. Strict log-concavity of f therefore implies strict log-concavity of F , which in turn implies strict monotonicity of p^* .

Step 3: Recover quality from prices. When p^* is strictly increasing, it is one-to-one on $(0, 1]$ and therefore admits an inverse on its range. Replace the explicit quality message in the TCP mechanism with a degenerate signal while retaining the pricing rule $p^*(q_i)$. Upon observing a price $p^*(q_i)$, the buyer recovers the realized quality as $q_i = (p^*)^{-1}(p^*(q_i))$. Her posterior mean, net utility, and purchase decision are therefore unchanged. The price-only mechanism induces the same allocation and equilibrium payoffs as the original TCP mechanism, establishing the claimed implementation result. $\square$

Proof of Proposition 2. We first represent equilibrium revenue, buyer utility, and social welfare through the two highest order statistics $Q_{(1;n)}$ and $Q_{(2;n)}$. We then establish the uniform asymptotic behavior of the pricing rule using a quantile transformation and derive the relevant order-statistic expansions from the concentration of their quantile kernels near the upper endpoint.

Step 1: Express equilibrium outcomes through order statistics. This first step reduces the equilibrium outcome measures to the highest and second-highest realized qualities. Let $Q_{(1;n)}$ and $Q_{(2;n)}$ denote the highest and second-highest values among $Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F$, respectively. Because the equilibrium score $b^*(\cdot)$ is strictly increasing, the buyer purchases from the highest-quality seller almost surely. Hence,

$$SW_{\text{gen}}(n) = \mathbb{E}_{Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F}[Q_{(1;n)}].$$

The density of $Q_{(1;n)}$ is $nF(y)^{n-1}f(y)$. Using $p_n^*(y)F(y)^{n-1} = \int_0^y F(t)^{n-1} dt$ and changing the order of integration, the total expected seller revenue satisfies

$$\begin{aligned} n \cdot R_{\text{gen}}(n) &= \mathbb{E}_{Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F}[p_n^*(Q_{(1;n)})] \\ &= n \int_0^1 \left(\int_0^y F(t)^{n-1} dt \right) f(y) dy \\ &= n \int_0^1 F(t)^{n-1} (1 - F(t)) dt \\ &= \mathbb{E}_{Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F}[Q_{(1;n)}] - \mathbb{E}_{Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F}[Q_{(2;n)}]. \end{aligned}$$

The last equality follows from $\Pr\{Q_{(1;n)} \leq t\} = F(t)^n$ and $\Pr\{Q_{(2;n)} \leq t\} = F(t)^n + nF(t)^{n-1}(1 -$

$F(t)$). By symmetry, each seller receives a $1/n$ share of the total expected revenue. The buyer's expected utility is consequently

$$U_{\text{gen}}(n) = SW_{\text{gen}}(n) - n \cdot R_{\text{gen}}(n) = \mathbb{E}_{Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F}[Q_{(2;n)}].$$

Step 2: Establish uniform price convergence. Then we derive a uniform first-order approximation for $p_n^*(\cdot)$. Because f is continuous and strictly positive on $[0, 1]$, F is strictly increasing and admits a continuous inverse. Define $h(u) := 1/f(F^{-1}(u))$ for $u \in [0, 1]$. The function h is continuous and bounded. For $q > 0$, let $a := F(q)$ and apply the changes of variables $z = F(t)$ and $v = z/a$ to obtain

$$n \cdot p_n^*(q) = a \int_0^1 n v^{n-1} h(av) dv, \quad \frac{F(q)}{f(q)} = ah(a).$$

The same identities extend to $q = 0$ after setting $p_n^*(0) = 0$.

To make the concentration of the kernel nv^{n-1} explicit, introduce an auxiliary random variable $V_n \sim \text{Beta}(n, 1)$ with density nv^{n-1} on $[0, 1]$. Its role is purely notational. In particular, for $a := F(q)$, we have

$$n \cdot p_n^*(q) = a \cdot \mathbb{E}_{V_n \sim \text{Beta}(n, 1)}[h(aV_n)], \quad \frac{F(q)}{f(q)} = a \cdot h(a).$$

Let $\omega_h(\delta) := \sup\{|h(u) - h(v)| : u, v \in [0, 1], |u - v| \leq \delta\}$ denote the modulus of continuity of h . Since h is uniformly continuous on $[0, 1]$, $\omega_h(\delta) \rightarrow 0$ as $\delta \downarrow 0$. For any $\delta \in (0, 1)$, splitting according to whether V_n lies within δ of 1 gives

$$\begin{aligned} \sup_{q \in [0, 1]} \left| n \cdot p_n^*(q) - \frac{F(q)}{f(q)} \right| &\leq \sup_{a \in [0, 1]} \mathbb{E}_{V_n \sim \text{Beta}(n, 1)}[|h(aV_n) - h(a)|] \\ &\leq \omega_h(\delta) + 2\|h\|_\infty \cdot (1 - \delta)^n. \end{aligned}$$

Indeed, on the event $\{V_n \geq 1 - \delta\}$, we have $|aV_n - a| \leq \delta$ uniformly over $a \in [0, 1]$. On the complementary event, the difference is bounded by $2\|h\|_\infty$, while $\Pr_{V_n \sim \text{Beta}(n, 1)}\{V_n < 1 - \delta\} = (1 - \delta)^n$. Letting $n \rightarrow \infty$ and then $\delta \downarrow 0$ yields

$$\sup_{q \in [0, 1]} \left| n \cdot p_n^*(q) - \frac{F(q)}{f(q)} \right| \rightarrow 0.$$

Therefore, for every $q \in (0, 1]$, $p_n^*(q) \sim F(q)/(nf(q))$. Moreover, F/f is continuous and bounded on $[0, 1]$, with $(F/f)(1) = 1/f(1) > 0$. The uniform approximation consequently implies $n\|p_n^*\|_\infty \rightarrow \|F/f\|_\infty \in (0, \infty)$. Hence, p_n^* converges uniformly to zero at rate $\Theta(1/n)$, proving part (i).

Step 3: Derive the order-statistic asymptotics. This step obtains the first-order behavior of the two highest qualities. Using the tail-integral representation and the change of variables

$u = F(t)$, we can rewrite these two quantities as follows

$$\begin{aligned} 1 - \mathbb{E}_{Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F} [Q_{(1;n)}] &= \int_0^1 u^n h(u) du, \\ 1 - \mathbb{E}_{Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F} [Q_{(2;n)}] &= \int_0^1 [u^n + nu^{n-1}(1-u)] h(u) du. \end{aligned}$$

The kernels u^n and $u^n + nu^{n-1}(1-u)$ have total masses $1/(n+1)$ and $2/(n+1)$, respectively. After normalization, both concentrate at $u = 1$: for every $\varepsilon > 0$, their mass on $[0, 1 - \varepsilon]$ converges to zero exponentially. Since h is bounded and continuous at 1, it follows that

$$\begin{aligned} 1 - \mathbb{E}_{Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F} [Q_{(1;n)}] &= \frac{1}{nf(1)} + o\left(\frac{1}{n}\right), \\ 1 - \mathbb{E}_{Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F} [Q_{(2;n)}] &= \frac{2}{nf(1)} + o\left(\frac{1}{n}\right). \end{aligned}$$

Step 4: Conclude the payoff asymptotics. We finally combine the outcome identities from Step 1 with the order-statistic expansions from Step 3. In particular, we obtain the following asymptotic expansions for the equilibrium outcomes

$$\begin{aligned} U_{\text{gen}}(n) &= \mathbb{E}_{Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F} [Q_{(2;n)}] \\ &= 1 - \frac{2}{nf(1)} + o\left(\frac{1}{n}\right) \sim 1 - \frac{1}{n} \cdot \frac{2}{f(1)}, \\ R_{\text{gen}}(n) &= \frac{\mathbb{E}_{Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F} [Q_{(1;n)}] - \mathbb{E}_{Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F} [Q_{(2;n)}]}{n} \\ &= \frac{1}{n^2 f(1)} + o\left(\frac{1}{n^2}\right) \sim \frac{1}{n^2} \cdot \frac{1}{f(1)}, \\ SW_{\text{gen}}(n) &= \mathbb{E}_{Q_1, \dots, Q_n \stackrel{\text{i.i.d.}}{\sim} F} [Q_{(1;n)}] \\ &= 1 - \frac{1}{nf(1)} + o\left(\frac{1}{n}\right) \sim 1 - \frac{1}{n} \cdot \frac{1}{f(1)}. \end{aligned}$$

These expansions imply the asymptotic relations stated in parts (ii)–(iv). $\square$

C Proofs of Section 5

Proof of Proposition 3. Let $p_n := p_{\text{ff}}^*(n)$ denote the common equilibrium price in the full-information, fixed-price benchmark, and let $Q_{(1;n)} := \max\{Q_1, \dots, Q_n\}$, where the qualities $Q_1, \dots, Q_n$ are i.i.d. draws from the prior distribution F .

Step 1: Compare finite-market efficiency. We first compare the three decision spaces along the extensive margin of trade and the allocative margin of product selection.

Under $\mathcal{M}$, Theorem 2 implies that $b^*(q) > 0$ for every $q > 0$. Because F is atomless, at least one seller offers a strictly positive score almost surely, and the buyer always purchases. Moreover, the strict monotonicity of b^* implies that she selects the highest-quality seller. Hence, $\text{MC}_{\text{gen}}(n) = 1$

and $\text{SW}_{\text{gen}}(n) = \mathbb{E}[Q_{(1;n)}]$.

Under $\mathcal{M}_{\text{nf}}$, Bertrand competition yields the equilibrium price $p_{\text{nf}}^* = 0$. The buyer therefore always purchases, so $\text{MC}_{\text{nf}}(n) = 1$. Because no information about realized qualities is disclosed, the selected product has expected quality $\mathbb{E}[Q]$, and thus $\text{SW}_{\text{nf}}(n) = \mathbb{E}[Q]$.

Under $\mathcal{M}_{\text{ff}}$, the equilibrium price satisfies $p_n \in (0, 1)$. Indeed, a price of zero can be profitably increased while retaining positive demand, whereas a price of at least one generates no demand and is dominated by a sufficiently small positive price. Given full disclosure and the common price p_n , the buyer purchases if and only if $Q_{(1;n)} \geq p_n$ and, conditional on purchasing, selects the highest-quality seller. Therefore,

$$\text{MC}_{\text{ff}}(n) = 1 - F(p_n)^n < 1, \quad \text{SW}_{\text{ff}}(n) = \mathbb{E}\left[Q_{(1;n)} \cdot \mathbf{1}\{Q_{(1;n)} \geq p_n\}\right].$$

This establishes part (i).

To compare social welfare, observe that $Q_{(1;n)} \geq Q_1$ almost surely, with strict inequality with positive probability because F is nondegenerate and $n \geq 2$. Hence, $\text{SW}_{\text{gen}}(n) = \mathbb{E}[Q_{(1;n)}] > \mathbb{E}[Q] = \text{SW}_{\text{nf}}(n)$. Moreover,

$$\text{SW}_{\text{gen}}(n) - \text{SW}_{\text{ff}}(n) = \mathbb{E}\left[Q_{(1;n)} \cdot \mathbf{1}\{Q_{(1;n)} < p_n\}\right] > 0,$$

where strictness follows from $p_n > 0$ and the strictly positive density of F . Thus, $\text{SW}_{\text{gen}}(n) > \max\{\text{SW}_{\text{nf}}(n), \text{SW}_{\text{ff}}(n)\}$, proving part (ii).

Step 2: Establish the large-market ranking. We next show that market exclusion under $\mathcal{M}_{\text{ff}}$ vanishes as the number of sellers grows. A deviation argument first bounds the equilibrium price away from the upper endpoint of the quality support.

At the symmetric fixed-price equilibrium, each seller receives an equal share of equilibrium demand. His expected revenue is therefore

$$\frac{p_n}{n} (1 - F(p_n)^n) \leq \frac{1}{n}.$$

Suppose that $p_n \geq 3/4$. A seller could deviate to the price $1/2$. Whenever his realized quality satisfies $q_i > 3/4$, his offer gives the buyer net utility $q_i - 1/2 > 1/4$, whereas every rival offers net utility at most $1 - p_n \leq 1/4$. The deviating seller therefore wins with probability at least $1 - F(3/4)$ and earns expected revenue of at least $\frac{1}{2}(1 - F(3/4)) > 0$. For all sufficiently large n , this deviation payoff exceeds $1/n$, contradicting equilibrium. Consequently, there exists N_0 such that $p_n < 3/4$ for every $n > N_0$.

It follows that the probability of market exclusion satisfies

$$\Pr\{Q_{(1;n)} < p_n\} = F(p_n)^n \leq F(3/4)^n \rightarrow 0.$$

At the same time, $Q_{(1;n)} \rightarrow 1$ almost surely and, because $Q_{(1;n)} \in [0, 1]$, $\mathbb{E}[Q_{(1;n)}] \rightarrow 1$. Moreover,

$$\mathbb{E}[Q_{(1;n)}] - F(p_n)^n \leq \text{SW}_{\text{ff}}(n) \leq \mathbb{E}[Q_{(1;n)}],$$

which implies that $\text{SW}_{\text{ff}}(n) \rightarrow 1$.

Finally, $\text{SW}_{\text{nf}}(n) = \mathbb{E}[Q] < 1$ is independent of n . Hence, there exists $N \in \mathbb{Z}^+$ such that $\text{SW}_{\text{ff}}(n) > \text{SW}_{\text{nf}}(n)$ for every $n > N$, which proves part (iii). $\square$

Proof of Proposition 4. We proceed with the proof in six steps.

Step 1: Establish the finite-market ranking. We first prove part (i) by making the source of strict positivity explicit in each decision space. Under $\mathcal{M}_{\text{nf}}$, Bertrand competition yields $p_{\text{nf}}^* = 0$, and hence $R_{\text{nf}}(n) = 0$. At the TCP equilibrium, strict monotonicity of the buyer's score implies that a seller with quality q wins with probability $F(q)^{n-1}$. Substituting the price in (2) and applying Fubini's theorem therefore gives

$$R_{\text{gen}}(n) = \int_0^1 p^*(q) F(q)^{n-1} f(q) dq = \int_0^1 (1 - F(x)) F(x)^{n-1} dx.$$

For every realization of the i.i.d. qualities, $Q_{(1;n)} - Q_{(2;n)} = \int_0^1 \mathbf{1}\{Q_{(1;n)} > x, Q_{(2;n)} \leq x\} dx$. Taking expectations and using independence,

$$\mathbb{E}[Q_{(1;n)} - Q_{(2;n)}] = n \int_0^1 (1 - F(x)) F(x)^{n-1} dx.$$

Hence $R_{\text{gen}}(n) = \mathbb{E}[Q_{(1;n)} - Q_{(2;n)}]/n$. The strictly positive density implies $F(x) \in (0, 1)$ for every $x \in (0, 1)$, so this revenue is strictly positive.

For the full-disclosure fixed-price benchmark, let $p_n := p_{\text{ff}}^*(n)$ denote the symmetric equilibrium price. Under Assumption 1, (FOC_{ff}) can be written as

$$p_n = \frac{1 - F(p_n)^n}{n D_n(p_n)}, \quad D_n(p) := F(p)^{n-1} f(p) + (n-1) \int_p^1 F(x)^{n-2} f(x)^2 dx.$$

The denominator is strictly positive. Moreover, neither $p_n = 0$ nor $F(p_n) = 1$ can satisfy this equation; therefore $p_n > 0$ and $1 - F(p_n)^n > 0$. It follows that $R_{\text{ff}}(n) = p_n(1 - F(p_n)^n)/n > 0$. This proves part (i).

Step 2: Expand TCP revenue at the upper endpoint. We next derive the third-order expansion of $R_{\text{gen}}(n)$. The argument rewrites revenue in quantile space, where the beta kernel concentrates at the upper endpoint.

Write $b := f(1) > 0$ and $c := f'(1)$. Under condition (a), there exist constants $0 < m \leq M < \infty$ such that $m \leq f(x) \leq M$ on $[0, 1]$. Because F is strictly increasing, the function $g(u) := f(F^{-1}(u))$ is well defined on $(0, 1)$ and also satisfies $m \leq g(u) \leq M$. The preceding order-statistic identity

and the substitutions $u = F(x)$ and $t = 1 - u$ yield

$$R_{\text{gen}}(n) = \int_0^1 \frac{(1-u)u^{n-1}}{g(u)} du = \int_0^1 \frac{t(1-t)^{n-1}}{g(1-t)} dt.$$

To expand the last integrand, let $y := 1 - F^{-1}(1-t)$. Continuous differentiability at the upper endpoint gives

$$1 - F(1-y) = by - \frac{c}{2}y^2 + o(y^2), \quad f(1-y) = b - cy + o(y).$$

Thus $y = t/b + O(t^2)$ as $t \downarrow 0$, and therefore

$$g(1-t) = b - \frac{c}{b}t + o(t), \quad \frac{1}{g(1-t)} = \frac{1}{b} + \frac{c}{b^3}t + o(t).$$

Using these two equations

$$\int_0^1 t(1-t)^{n-1} dt = \frac{1}{n(n+1)} \quad \text{and} \quad \int_0^1 t^2(1-t)^{n-1} dt = \frac{2}{n(n+1)(n+2)},$$

we obtain the final expression

$$R_{\text{gen}}(n) = \frac{1}{bn^2} + \left(-\frac{1}{b} + \frac{2c}{b^3}\right) \frac{1}{n^3} + o(n^{-3}).$$

For completeness, write the endpoint remainder as $r(t) = o(t)$. Given $\varepsilon > 0$, choose $\eta > 0$ such that $|r(t)| \leq \varepsilon t$ for $t \leq \eta$. The contribution from $(0, \eta]$ is bounded by $2\varepsilon/[n(n+1)(n+2)]$, whereas boundedness of g^{-1} makes the contribution from $[\eta, 1]$ exponentially small. After multiplication by n^3 and then letting $\varepsilon \downarrow 0$, these bounds establish the stated $o(n^{-3})$ remainder.

Step 3: Control the fixed-price equilibrium. We now obtain a uniform bound on p_n and show that the lower endpoint is asymptotically irrelevant. These bounds justify replacing the truncated denominator in the fixed-price first-order condition by an integral over the full support.

Changing variables from x to $u = F(x)$ gives

$$D_n(p) = F(p)^{n-1}f(p) + (n-1) \int_{F(p)}^1 u^{n-2}g(u) du.$$

The two terms form a weighted average: their weights are $F(p)^{n-1}$ and $1 - F(p)^{n-1}$, respectively. Since both $f(p)$ and $g(u)$ lie in $[m, M]$, $m \leq D_n(p) \leq M$ for every $p \in [0, 1]$. Combining the lower bound with the first-order condition yields

$$0 < p_n \leq \frac{1}{mn}. \tag{C-1}$$

Because $F(x) \leq Mx$ on $[0, 1]$, (C-1) further implies

$$F(p_n)^n \leq \left(\frac{M}{mn}\right)^n = o(n^{-k}) \quad \text{for every fixed } k > 0.$$

Define

$$\overline{D}_n := (n-1) \int_0^1 F(x)^{n-2} f(x)^2 dx.$$

The omitted boundary term is bounded by $MF(p_n)^{n-1}$, while the lower truncation is bounded by $(n-1)M^2p_nF(p_n)^{n-2}$. The preceding super-polynomial bound therefore gives $D_n(p_n) = \bar{D}_n + o(n^{-1})$.

Step 4: Expand fixed-price revenue. This step uses the same endpoint expansion as Step 2 to evaluate $\bar{D}_n$, and then substitutes the result into the equilibrium first-order condition.

The changes of variables $u = F(x)$ and $t = 1 - u$ imply

$$\bar{D}_n = (n-1) \int_0^1 (1-t)^{n-2} g(1-t) dt.$$

Substituting $g(1-t) = b - (c/b)t + o(t)$ and using

$$(n-1) \int_0^1 (1-t)^{n-2} dt = 1, \quad (n-1) \int_0^1 t(1-t)^{n-2} dt = \frac{1}{n},$$

gives

$$D_n(p_n) = b - \frac{c}{bn} + o(n^{-1}).$$

The remainder is $o(n^{-1})$ by the same split-and-bound argument used in Step 2. Hence,

$$\frac{1}{D_n(p_n)} = \frac{1}{b} + \frac{c}{b^3n} + o(n^{-1}).$$

Since $F(p_n)^n$ is smaller than every polynomial order, the first-order condition yields

$$p_n = \frac{1}{bn} + \frac{c}{b^3n^2} + o(n^{-2}).$$

Substituting this expansion into $R_{\text{ff}}(n) = p_n(1 - F(p_n)^n)/n$ gives

$$R_{\text{ff}}(n) = \frac{1}{bn^2} + \frac{c}{b^3n^3} + o(n^{-3}).$$

Step 5: Compare the two revenue expansions. The leading n^{-2} terms coincide, so the revenue ranking is determined at order n^{-3} . Subtracting the TCP expansion in Step 2 from the fixed-price expansion in Step 4 gives

$$R_{\text{ff}}(n) - R_{\text{gen}}(n) = \frac{b^2 - c}{b^3} \frac{1}{n^3} + o(n^{-3}).$$

Condition (c) makes the leading coefficient strictly positive. Therefore, there exists $N \in \mathbb{Z}^+$ such that $R_{\text{ff}}(n) > R_{\text{gen}}(n)$ for every $n \geq N$, which is stronger than the weak inequality stated in part (ii).

Step 6: Verify the log-concavity implication. It remains to justify the sufficient condition stated in part (ii). Suppose that f is log-concave and let $\phi := \log f$. If $\phi'(1) \leq 0$, then $c \leq 0 < b^2$. If instead $\lambda := \phi'(1) = c/b > 0$, concavity implies $f(q) \leq b \exp\{\lambda(q-1)\}$ for every $q \in [0, 1]$. Integrating and using $\int_0^1 f(q) dq = 1$ gives

$$1 \leq b \int_0^1 e^{\lambda(q-1)} dq = \frac{b}{\lambda} (1 - e^{-\lambda}).$$

Thus $\lambda \leq b(1 - e^{-\lambda}) < b$, or equivalently $c < b^2$. Hence log-concavity implies condition (c), completing the proof. $\square$

Proof of Proposition 5. Let $n \geq 2$ and set $m := \min_{q \in [0,1]} f(q) > 0$, which exists because f is continuous and strictly positive on $[0, 1]$. We establish the three comparisons separately.

Step 1: Compare no disclosure with TCP. We first prove part (i) by exploiting the monotonicity of the second-highest quality as the market expands. The equilibrium characterizations give $U_{\text{nf}}(n) = \mathbb{E}[Q]$ and $U_{\text{gen}}(n) = \mathbb{E}[Q_{(2;n)}]$. Define $\Delta(n) := \mathbb{E}[Q_{(2;n)}] - \mathbb{E}[Q]$.

Couple the markets with n and $n+1$ sellers by using the same first n quality draws and adding one independent draw. Then $Q_{(2;n+1)} \geq Q_{(2;n)}$ almost surely, so $\Delta(n)$ is weakly increasing. At $n = 2$, $Q_{(2;2)} = \min\{Q_1, Q_2\} \leq Q_1$ almost surely, with strict inequality with positive probability because F has a strictly positive density. Hence $\Delta(2) < 0$.

It remains to verify the limiting sign. For every $\varepsilon \in (0, 1)$, the event $Q_{(2;n)} \leq 1 - \varepsilon$ occurs precisely when at most one of the n qualities exceeds $1 - \varepsilon$. Therefore,

$$\Pr\{Q_{(2;n)} \leq 1 - \varepsilon\} = F(1 - \varepsilon)^n + n(1 - F(1 - \varepsilon))F(1 - \varepsilon)^{n-1} \rightarrow 0.$$

Thus $Q_{(2;n)} \rightarrow 1$ in probability. More explicitly, $0 \leq 1 - \mathbb{E}[Q_{(2;n)}] \leq \varepsilon + \Pr\{Q_{(2;n)} \leq 1 - \varepsilon\}$, so first letting $n \rightarrow \infty$ and then $\varepsilon \downarrow 0$ gives $\mathbb{E}[Q_{(2;n)}] \rightarrow 1$. Moreover, absolute continuity on $[0, 1]$ gives $\mathbb{E}[Q] < 1$, and hence $\Delta(n) \rightarrow 1 - \mathbb{E}[Q] > 0$.

Consequently, the set $\{n \geq 2 : \Delta(n) > 0\}$ is nonempty. Define $N_1 := \min\{n \geq 2 : \Delta(n) > 0\}$. Monotonicity implies $\Delta(n) \leq 0$ for every $n < N_1$ and $\Delta(n) > 0$ for every $n \geq N_1$. In particular, $U_{\text{nf}}(n) \geq U_{\text{gen}}(n)$ for every $n < N_1$, whereas $U_{\text{nf}}(n) \leq U_{\text{gen}}(n)$ for every $n > N_1$, as claimed.

Step 2: Compare no disclosure with fixed-price full disclosure. We next establish part (ii) by bounding both the fixed equilibrium price and the gap between the highest quality and the top of the support. Let $p_n := p_{\text{ff}}^*(n)$. The denominator bound in (C-1), whose derivation uses only continuity and strict positivity of f , gives $0 < p_n \leq 1/(mn)$.

For every $t \in [0, 1]$, the lower density bound yields $1 - F(1 - t) \geq mt$ and therefore $F(1 - t)^n \leq (1 - mt)^n \leq e^{-mnt}$. The tail-integral identity then implies

$$1 - \mathbb{E}[Q_{(1;n)}] = \int_0^1 \Pr\{Q_{(1;n)} \leq 1 - t\} dt \leq \int_0^1 e^{-mnt} dt \leq \frac{1}{mn}.$$

Because $(q - p)_+ \geq q - p$ for every $q \in [0, 1]$ and $p \geq 0$, we obtain

$$U_{\text{ff}}(n) = \mathbb{E}[(Q_{(1;n)} - p_n)_+] \geq \mathbb{E}[Q_{(1;n)}] - p_n \geq 1 - \frac{2}{mn}.$$

Set $\delta := 1 - \mathbb{E}[Q] > 0$ and choose $N_2 := \lceil 2/(m\delta) \rceil$. For every $n > N_2$, we have $2/(mn) < \delta$, and hence

$$U_{\text{ff}}(n) > 1 - \delta = \mathbb{E}[Q] = U_{\text{nf}}(n).$$

This strict inequality is stronger than part (ii).

Step 3: Compare the two full-information regimes. We impose the additional conditions in part (iii). By part (ii) of Proposition 4, there exists $N_3 \in \mathbb{Z}^+$ such that $R_{\text{ff}}(n) \geq R_{\text{gen}}(n)$ for every $n \geq N_3$. In addition, part (ii) of Proposition 3 gives $SW_{\text{gen}}(n) > SW_{\text{ff}}(n)$ for every finite n .

Since R is per-seller revenue, the accounting identity $U = SW - nR$ gives, for every $n \geq N_3$,

$$U_{\text{gen}}(n) - U_{\text{ff}}(n) = SW_{\text{gen}}(n) - SW_{\text{ff}}(n) + n(R_{\text{ff}}(n) - R_{\text{gen}}(n)) > 0.$$

Therefore, $U_{\text{gen}}(n) > U_{\text{ff}}(n)$ for every $n \geq N_3$, which proves our part (iii). $\square$

D Proofs of Section 6

Proof of Proposition 6. The proof establishes the result by mapping the asymmetric pricing problem into an asymmetric first-price auction. The key observation is that, without buyer-side private information, every seller's offer can be summarized by the net utility left to the buyer.

Step 1: Reduce transparent contingent pricing to bidding. Fix seller i with realized quality q_i . Under a transparent contingent pricing mechanism, define the buyer's net utility from seller i as $b_i(q_i) := q_i - p_i(q_i)$. Because the buyer's valuation is common across all sellers and normalized to one, she purchases from the seller offering the highest nonnegative net utility. Therefore, seller i wins if and only if

$$b_i(q_i) \geq \max\{0, b_j(q_j), j \neq i\}.$$

Hence, choosing a price function $p_i(\cdot)$ is equivalent to choosing a bid function $b_i(\cdot)$ in which the seller's value is his realized quality q_i and his payment is the residual $p_i(q_i) = q_i - b_i(q_i)$.

The seller's expected payoff becomes

$$(q_i - b_i(q_i)) \Pr \left\{ b_i(q_i) \geq \max_{j \neq i} b_j(Q_j) \right\},$$

which is exactly the payoff function of seller i in a first-price auction with independent but non-identically distributed bidder values.

Step 2: Construct the equilibrium pricing rule. Because each F_i has a continuous, strictly positive density on the common support $[0, 1]$, the associated asymmetric first-price auction admits a pure-strategy equilibrium $(b_1^*, \dots, b_n^*)$ in strictly increasing bidding functions (Maskin and Riley, 2000b). Strict monotonicity and the atomlessness of F_i imply that the induced bid $b_i^*(Q_i)$ is atomless for every seller i : the preimage of any bid under b_i^* contains at most one quality type. Hence, ties among equilibrium bids occur with probability zero. Moreover, the highest rival bid has a continuous distribution and therefore ties with any unilateral bid that is independent of the rivals with probability zero. Thus, the tie-breaking rule is immaterial in this construction.

For each seller i , define the corresponding quality-contingent price rule as

$$p_i^*(q_i) := q_i - b_i^*(q_i).$$

Consider the mechanism in which seller i fully reveals q_i and posts the price $p_i^*(q_i)$. Under this mechanism, the buyer's net utility from seller i is exactly $q_i - p_i^*(q_i) = b_i^*(q_i)$. Therefore, the induced allocation and seller payoffs coincide with those in the asymmetric first-price auction equilibrium.

Because no seller can profitably deviate in the auction problem, no seller can profitably deviate within transparent contingent pricing mechanisms. Thus, the constructed pricing profile constitutes an equilibrium. Finally, because the buyer has no private information, any signal-price pair is evaluated only through the induced posterior mean quality and price. Under the constructed equilibrium, each seller fully reveals quality, so the posterior mean equals the realized quality.

Therefore, the equilibrium outcome is implemented by transparent contingent pricing, even though sellers are ex-ante asymmetric. $\square$

Proof of Proposition 7. We first establish the claim by enlarging the signal space without changing any observable outcome. Start from an arbitrary mechanism that generates (s_i, p_i) conditional on q_i . Define the augmented signal $\hat{s}_i := (s_i, p_i)$ and the deterministic pricing rule $\hat{p}_i(\hat{s}_i) := p_i$. The joint distribution of (Q_i, s_i, p_i) is unchanged, so the buyer's posterior belief, choice probabilities, and all sellers' payoffs are unchanged. Applying this transformation seller by seller preserves every mechanism profile and every unilateral deviation payoff. Hence, an equilibrium, if one exists, has an outcome-equivalent GCP representation.

Then we derive the substantive reduction. Let (S_i, P_i) be the realized offer under an arbitrary deviation and define $M_i := \mathbb{E}[Q_i \mid S_i, P_i]$. The distribution of M_i belongs to $\text{MPC}(F)$. Conditional on $(M_i, P_i) = (m, p)$, seller i earns $\Pi_i(m, p; \kappa_{-i}) \leq V_i(m; \kappa_{-i})$. Therefore,

$$R_i(\kappa_i, \kappa_{-i}) \leq \mathbb{E}[V_i(M_i; \kappa_{-i})] \leq \sup_{H \in \text{MPC}(F)} \int V_i(m; \kappa_{-i}) dH(m).$$

Conversely, fix $H \in \text{MPC}(F)$. By the definition of $\text{MPC}(F)$, there exists an information structure that induces a posterior mean $M_i \sim H$. It remains to construct a measurable near-optimal price rule. Define $W_i := \max\{0, \max_{j \neq i}(\theta M_j - P_j)\}$. Under the seller-favorable tie-breaking,

$$\Pi_i(m, p; \kappa_{-i}) = p \Pr_{\theta, \kappa_{-i}} \{\theta m - p \geq W_i\}.$$

The map $(m, p, \omega) \mapsto p \mathbf{1}\{\theta(\omega)m - p \geq W_i(\omega)\}$ is Borel measurable and bounded. Hence, $\Pi_i(\cdot, \cdot; \kappa_{-i})$ is Borel measurable. For fixed m , if $p_k \uparrow p$, then

$$\{\theta m - p_k \geq W_i\} \downarrow \{\theta m - p \geq W_i\}.$$

Continuity of probability from above therefore implies that $p \mapsto \Pi_i(m, p; \kappa_{-i})$ is left-continuous. Then we define $D := (\mathbb{Q} \cap [0, \bar{p}]) \cup \{\bar{p}\} = \{r_1, r_2, \dots\}$. Because every positive price can be approxi-

mated from below by prices in D , left-continuity gives

$$V_i(m; \kappa_{-i}) = \sup_{r \in D} \Pi_i(m, r; \kappa_{-i}).$$

Thus, $V_i(\cdot; \kappa_{-i})$ is Borel measurable as a countable supremum of Borel functions. For every $\varepsilon > 0$, define

$$k_\varepsilon(m) := \min \{k : \Pi_i(m, r_k; \kappa_{-i}) > V_i(m; \kappa_{-i}) - \varepsilon\}, \quad p_i^\varepsilon(m) := r_{k_\varepsilon(m)}.$$

For every k , the set $\{m : k_\varepsilon(m) = k\}$ is Borel, so p_i^ε is measurable. Using M_i as the signal and posting $p_i^\varepsilon(M_i)$ therefore defines a GCP mechanism whose expected revenue is at least

$$\int_{[0,1]} V_i(m; \kappa_{-i}) dH(m) - \varepsilon.$$

Taking the supremum over $H \in \text{MPC}(F)$ and then letting $\varepsilon \downarrow 0$ proves the result. $\square$

Proof of Proposition 8. We divide the proof into three steps.

Step 1: Reduce survival to posterior-mean values. This step expresses the seller's full-revelation payoff and his maximal information-design payoff through V . We also establish that V is automatically continuous under the maintained bounded-type and flexible-price assumptions.

Let $\Pi(m, p)$ denote the seller's continuation payoff at posterior mean m and price p , so that $V(m) = \sup_{p \geq 0} \Pi(m, p)$. Because the seller's pricing problem in the auxiliary game separates across realized qualities, the equilibrium condition in that game implies $\chi^*(q) \in \arg\max_{p \geq 0} \Pi(q, p)$ for F -almost every q . The seller's expected payoff under full revelation is therefore $\mathbb{E}_{Q \sim F}[V(Q)]$.

By the posterior-mean reduction established above, an arbitrary joint information-pricing deviation induces a feasible posterior-mean random variable M and yields at most $\mathbb{E}[V(M)]$. Conversely, any feasible distribution of posterior means can be paired with posterior-mean-contingent prices that attain, or approximate arbitrarily closely, the corresponding value $\mathbb{E}[V(M)]$. Hence, full revelation survives if and only if

$$\mathbb{E}_{Q \sim F}[V(Q)] \geq \mathbb{E}_{M \sim H}[V(M)] \quad \text{for every } H \in \text{MPC}(F). \quad (\text{FI})$$

We next show that V is continuous. Fix $m_2 \geq m_1$ and let $\delta := m_2 - m_1$. At any fixed price, a higher posterior mean weakly raises the seller's utility for every buyer type, so $V(m_2) \geq V(m_1)$.

For the reverse bound, fix any price p at posterior mean m_2 . If $p < \bar{\theta}\delta$, then $\Pi(m_2, p) \leq p < \bar{\theta}\delta$. If $p \geq \bar{\theta}\delta$, define $p' := p - \bar{\theta}\delta$. For every buyer type $\theta \leq \bar{\theta}$,

$$\theta m_1 - p' = \theta m_2 - p + (\bar{\theta} - \theta)\delta \geq \theta m_2 - p.$$

Thus, every realization under which the seller wins at (m_2, p) also leads him to win at (m_1, p') . Consequently, $\Pi(m_2, p) \leq \Pi(m_1, p') + \bar{\theta}\delta$. Taking the supremum over p gives

$$0 \leq V(m_2) - V(m_1) \leq \bar{\theta}(m_2 - m_1).$$

Therefore, V is $\bar{\theta}$ -Lipschitz continuous on $[0, 1]$.

Step 2: Prove sufficiency. This step applies the martingale property of posterior means. Convexity makes the seller prefer the dispersion generated by full revelation to every Bayes-plausible garbling.

Let M be any feasible posterior mean. Bayes plausibility implies $M = \mathbb{E}[Q \mid M]$ almost surely. If V is convex, conditional Jensen's inequality gives

$$\mathbb{E}[V(M)] = \mathbb{E}[V(\mathbb{E}[Q \mid M])] \leq \mathbb{E}[V(Q)].$$

Thus, condition (FI) holds for every feasible M , and no joint information-pricing deviation is profitable. Full revelation together with χ^* therefore constitutes an equilibrium of the general mechanism game.

Step 3: Prove necessity. We next show that any violation of convexity generates a profitable information deviation. Full support and atomlessness allow the seller to pool small sets of qualities around the points involved in the convexity violation.

Suppose that V is not convex. Then there exist $q_L < q_H$, $\lambda \in (0, 1)$, and $\eta > 0$ such that, defining $\bar{q} := \lambda q_L + (1 - \lambda)q_H$,

$$V(\bar{q}) \geq \lambda V(q_L) + (1 - \lambda)V(q_H) + 4\eta. \quad (\text{NC})$$

Choose disjoint neighborhoods of q_L and q_H . Because F is atomless and has full support, these neighborhoods contain measurable subsets A_L and A_H satisfying

$$\Pr\{Q \in A_L\} = \lambda\alpha, \quad \Pr\{Q \in A_H\} = (1 - \lambda)\alpha$$

for some $\alpha > 0$. The seller pools $A_L \cup A_H$ into a common signal and fully reveals all qualities outside this set.

Let $\mu_L := \mathbb{E}[Q \mid Q \in A_L]$ and $\mu_H := \mathbb{E}[Q \mid Q \in A_H]$. The posterior mean following the pooled signal is $m_A = \lambda\mu_L + (1 - \lambda)\mu_H$. As the two neighborhoods shrink to q_L and q_H , respectively, we have $m_A \rightarrow \bar{q}$ and

$$\mathbb{E}[V(Q) \mid Q \in A_L \cup A_H] \longrightarrow \lambda V(q_L) + (1 - \lambda)V(q_H).$$

Because V is continuous by Step 1, condition (NC) implies that, for sufficiently small neighborhoods,

$$V(m_A) > \mathbb{E}[V(Q) \mid Q \in A_L \cup A_H] + 2\eta.$$

After the pooled signal, choose a price that yields a continuation payoff within η of $V(m_A)$. The seller then earns strictly more than under full revelation conditional on $A_L \cup A_H$, while his behavior is unchanged outside this set. The pooling deviation is therefore strictly profitable, contradicting condition (FI).

Hence, if full revelation together with χ^* survives in the general mechanism game, V must be convex on $[0, 1]$. $\square$

Proof of Example 3. We first establish existence and uniqueness of the candidate coefficient. Define $h(a) := a(\lambda/(\theta_r - a) + (1 - \lambda)/(\theta_l - a))$ on $[0, \theta_l)$. The function h is continuous and strictly increasing, with $h(0) = 0$ and $h(a) \rightarrow \infty$ as $a \uparrow \theta_l$. Hence, the equation $h(a) = 1$ admits a unique solution $a^* \in (0, \theta_l)$.

Then, we check every price region and simultaneously compute the posterior-mean value function. Suppose the rival fully reveals quality and charges a^*q_j . Fix a posterior mean $m > 0$ and write a deviation price as $p = ym$. For a buyer of type $\theta \in \{\theta_l, \theta_r\}$, the seller's demand is

$$D_\theta(m, y) = \min \left\{ \frac{m(\theta - y)_+}{\theta - a^*}, 1 \right\}.$$

The deviation payoff is $ym[\lambda D_{\theta_r}(m, y) + (1 - \lambda)D_{\theta_l}(m, y)]$. Using $\min\{z, 1\} \leq z$, this payoff is bounded above by $m^2\Phi(y)$, where

$$\Phi(y) := y \left[\lambda \frac{(\theta_r - y)_+}{\theta_r - a^*} + (1 - \lambda) \frac{(\theta_l - y)_+}{\theta_l - a^*} \right].$$

On $[0, \theta_l]$, Φ is a strictly concave quadratic. The defining equation for a^* gives $\Phi'(a^*) = 0$, while $\Phi(a^*) = a^*$. Hence, a^* is the unique maximizer on this interval.

On $[\theta_l, \theta_r]$, only the high type is active and $\Phi(y) = \lambda y(\theta_r - y)/(\theta_r - a^*)$. Because $\theta_l \geq \theta_r/2$, this expression is decreasing on the interval. Moreover, the defining equation for a^* implies

$$a^* - \lambda \frac{\theta_l(\theta_r - \theta_l)}{\theta_r - a^*} = \frac{(\theta_l - a^*)^2}{a^*} \geq 0.$$

Thus, no high-type-only price yields more than a^* . Prices with $y \geq \theta_r$ generate zero demand.

At $y = a^*$, neither demand term is truncated: $D_{\theta_l}(m, a^*) = D_{\theta_r}(m, a^*) = m$. The upper bound is therefore attained, and the optimized payoff is $V(m) = a^*m^2$. The case $m = 0$ is immediate. Setting $m = q$ proves that $\chi^*(q) = a^*q$ is a global best response in the full-information pricing game. Since V is convex, Proposition 8 implies that this equilibrium survives in the general mechanism space. $\square$

Proof of Example 4. The proof first verifies global optimality of χ^* in the auxiliary game. It then identifies a strictly concave region of the posterior-mean value function and applies strict Jensen's inequality to an explicit pooling deviation.

Step 1: Verify the prior and price schedule. This step establishes that the constructed prior is regular and records the reverse-hazard identity that implements the equilibrium first-order condition. Direct differentiation gives, for $q \in (0, 1]$,

$$\frac{f(q)}{F(q)} = \frac{(10 - 3q^2)(10 - 9q^2)}{q(q^2 + 5)(20 - 9q^2)} > 0.$$

Thus, F is strictly increasing, with $F(0) = 0$ and $F(1) = 1$. Moreover, as $q \downarrow 0$, $F(q) = c_0 q + O(q^3)$, where $c_0 = \left(\frac{20}{11}\right)^{\frac{3}{26}} \left(\frac{5}{6}\right)^{-\frac{55}{26}} \in (0, \infty)$. The density therefore extends continuously to $q = 0$ by setting $f(0) = c_0$, and is strictly positive, bounded, and continuous on $[0, 1]$.

The candidate price satisfies $\chi^{*'}(q) = 1/3 + q^2/5 \leq 8/15 < 5/9$. Hence, for either buyer type $\theta \in \{5/9, 1\}$, the utility schedule $u_\theta(q) := \theta q - \chi^*(q)$ is strictly increasing.

Define

$$\Gamma(z) := \frac{1}{2} \frac{1}{\frac{5}{9} - z} + \frac{1}{2} \frac{1}{1 - z}.$$

A direct calculation yields

$$\frac{f(q)}{F(q)} = \frac{1}{\chi^*(q)\Gamma(\chi^{*'}(q))}. \quad (\text{D-2})$$

Step 2: Establish the full-information equilibrium. Here we ~~proves~~**prove** global optimality of χ^* . We establish concavity over prices that attract both buyer types and then rule out prices that target only the high type.

Fix a perceived quality $m \in [0, 1]$. For type θ , let $D_\theta(m, p)$ denote the probability that the seller wins at price p when the rival fully reveals quality and charges χ^* . Whenever the relevant rival-quality threshold r is interior, it solves $u_\theta(r) = \theta m - p$. Implicit differentiation gives

$$\frac{\partial^2 D_\theta(m, p)}{\partial p^2} = \frac{f(r)}{(\theta - \chi^{*'}(r))^3} [(\log f)'(r)(\theta - \chi^{*'}(r)) + x^{*''}(r)]. \quad (\text{D-3})$$

For the two buyer types,

$$\begin{aligned} (\log f)'(q) \left(\frac{5}{9} - \chi^{*'}(q) \right) + x^{*''}(q) &= -\frac{q(10 - 9q^2)(54q^4 - 335q^2 + 350)}{15(q^2 + 5)(10 - 3q^2)(20 - 9q^2)}, \\ (\log f)'(q) (1 - \chi^{*'}(q)) + x^{*''}(q) &= -\frac{3q(10 - 3q^2)(54q^4 - 175q^2 + 250)}{5(q^2 + 5)(20 - 9q^2)(10 - 9q^2)}. \end{aligned}$$

Both expressions are strictly negative on $(0, 1]$. The two polynomials $54y^2 - 335y + 350$ and $54y^2 - 175y + 250$ are decreasing on $[0, 1]$ and equal 69 and 129, respectively, at $y = 1$.

It follows from (D-3) that each $D_\theta(m, p)$ is concave in p over the region in which the corresponding threshold is interior. The truncation at demand one preserves concavity because the derivative drops from zero to a negative value when the threshold enters the support.

For $p \in [0, 5m/9]$, both buyer types obtain nonnegative utility. The aggregate demand $D(m, p) := \frac{1}{2}D_{5/9}(m, p) + \frac{1}{2}D_1(m, p)$ is nonincreasing and concave. Therefore, $\Pi(m, p) := pD(m, p)$ is concave on this interval because $\Pi_{pp} = 2D_p + pD_{pp} \leq 0$ on every smooth segment, with downward jumps in slope at the remaining kinks.

At $p = \chi^*(m)$, both types rank sellers by quality, so $D_{5/9}(m, \chi^*(m)) = D_1(m, \chi^*(m)) = F(m)$.

Moreover,

$$\left. \frac{\partial D(m, p)}{\partial p} \right|_{p=\chi^*(m)} = -f(m)\Gamma(\chi^{*'}(m)).$$

Using (D-2),

$$\left. \frac{\partial \Pi(m, p)}{\partial p} \right|_{p=\chi^*(m)} = F(m) - \chi^*(m)f(m)\Gamma(\chi^{*'}(m)) = 0.$$

Since $\chi^*(m)/m = (m^2 + 5)/15 \leq 2/5 < 5/9$, the candidate lies in the region in which both types are active. Concavity therefore makes $\chi^*(m)$ optimal over $[0, 5m/9]$. At $m = 1$, the preceding derivative is the right derivative at the demand-saturation kink; the left derivative is positive, so the same conclusion follows.

Now consider $p \in [5m/9, m]$, where only the high type can purchase. Let $d := D_1(m, p)$. If $d > 0$, let r solve $r - \chi^*(r) = m - p$. At the lower price $p/2$, high-type demand is strictly greater than d . For the low type,

$$\frac{5}{9}m - \frac{p}{2} - \left(\frac{5}{9}r - \chi^*(r) \right) = \frac{p}{18} + \frac{4}{9}\chi^*(r) > 0.$$

Low-type demand at $p/2$ is therefore also strictly greater than d , and

$$\Pi\left(m, \frac{p}{2}\right) > \frac{p}{2}d = \Pi(m, p).$$

Thus, every positive-demand price that targets only the high type is strictly dominated. Prices above m generate no demand. We conclude that $\chi^*(m)$ is a global maximizer for every $m \in [0, 1]$. Setting $m = q$ establishes the symmetric equilibrium of the full-information pricing game.

Step 3: Identify a concave value region. Next we compute the posterior-mean value function and identify an interval on which pooling is valuable. Since Step 2 applies to every $m \in [0, 1]$,

$$V(m) = \chi^*(m)F(m).$$

Differentiating twice gives

$$V''(m) = -F(m) \frac{567m^6 - 3480m^4 + 7500m^2 - 4000}{3m(m^2 + 5)(20 - 9m^2)^2}, \quad m \in (0, 1]. \quad (\text{D-4})$$

Let $\psi(y) := 567y^3 - 3480y^2 + 7500y - 4000$. Since $\psi'(y) = 3(567y^2 - 2320y + 2500) > 0$ on $[0, 1]$ and $\psi(64/81) = 216928/6561 > 0$, (D-4) implies $V''(m) < 0$ for every $m \in [8/9, 1]$.

Step 4: Construct a profitable pooling deviation. This step converts local concavity into a strict information-design gain. Define

$$\bar{m} := \mathbb{E}\left[Q_i \mid Q_i \geq \frac{8}{9}\right] = \frac{\int_{8/9}^1 qf(q) dq}{1 - F(8/9)}.$$

The seller reveals every $q_i < 8/9$ and continues to charge $\chi^*(q_i)$, but pools all $q_i \in [8/9, 1]$ into one

signal and charges $\chi^*(\bar{m})$ following that signal. The pooled signal induces posterior mean $\bar{m}$, and Step 2 implies that its optimized conditional payoff is $V(\bar{m})$.

Under full revelation, the conditional payoff over the same interval is $\mathbb{E}[V(Q_i) \mid Q_i \geq 8/9]$. Strict concavity of V and nondegeneracy of the conditional distribution give

$$V(\bar{m}) > \mathbb{E}\left[V(Q_i) \mid Q_i \geq \frac{8}{9}\right].$$

The ex-ante gain is therefore

$$\left(1 - F\left(\frac{8}{9}\right)\right) V(\bar{m}) - \int_{8/9}^1 V(q)f(q) dq > 0.$$

Thus, the displayed full-information equilibrium does not survive in the general mechanism space.

Finally, F is atomless and each $u_\theta(\cdot)$ is strictly increasing, so payoff-relevant seller ties occur with probability zero under both the full-information profile and the pooling deviation. The result is independent of the seller tie-breaking rule. $\square$