Pricing for Coordination

Marina Halac Elliot Lipnowski Daniel Rappoport*

May 1, 2024

Preliminary and Incomplete

Abstract

A seller prices a good with network externalities. Purchasing decisions being complementary, a pricing policy can yield equilibrium multiplicity. We study how personalized pricing can be used to mitigate this strategic uncertainty, guaranteeing a high revenue. An optimal policy offers personalized discounts to successively insulate against low-demand equilibria, and posts a high price to extract revenue from the induced higher demand. The result is price dispersion and a higher quantity of trade than would occur if the seller could choose her preferred equilibrium. We examine how the optimal policy changes with the strength of externalities and heterogeneity across buyer groups.

*Halac: Department of Economics, Yale University, 30 Hillhouse Avenue, New Haven, CT 06511, and CEPR (email: marina.halac@yale.edu); Lipnowski: Department of Economics, Columbia University, 420 West 118th Street, New York, NY 10027 (email: e.lipnowski@columbia.edu); Rappoport: University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637 (email: daniel.rappoport@chicagobooth.edu). Acknowledgments to be added.

1. Introduction

We study a seller who sells a good to a population of buyers, with two key features. First, there is incomplete information: a buyer's value from purchasing the good depends on some privately known characteristic, namely his type. Second, there are network externalities: a buyer's value from purchasing the good increases with the number of other buyers who also purchase it.

There are many applications for this canonical setting. Incomplete information is the quintessential feature of the monopoly problem, and network externalities are prevalent across industries. For a classic example, take a seller of a communications system (Rohlfs, 1974), say a messaging app. The utility a buyer derives from the app depends on his privately known propensity to communicate via messages, and also increases with the number of users with whom he can exchange messages. Similar considerations apply to sellers of file-sharing services, online social media platforms, and multiplayer game websites, among others. In finance, these features are central to a firm raising capital. An investor's incentive to invest with the firm depends on his other planned investments, which are his private information, and is higher if more other investors invest, as the firm is then more likely to be successful.

The seller's problem is to offer each buyer a price to maximize revenue. Prices can be personalized—e.g., via discounts and promotional deals directed to different buyers—but they cannot be conditioned on buyers' types, which are hidden. Buyers decide whether to purchase given the seller's price offers and their types, and given their expectations of other buyers' purchasing decisions. Due to the externalities in consumption, a pricing policy can yield multiple outcomes, with a high total quantity of trade if buyers anticipate that many others will purchase, or a low total quantity if buyers are less optimistic about others' purchases. Low-quantity outcomes are naturally bad for revenue.

Our main result characterizes the optimal pricing policy that guarantees the seller a high revenue, i.e. that maximizes revenue in her worst-case outcome. This policy takes the form of a posted price with dispersed discounts. The seller offers personalized discounts to (some) buyers to successively insulate against low-demand outcomes, and posts a high price to extract revenue from the induced higher demand. Targeted discounts on a list price are common in applications. Our analysis provides a rationale for these practices and sheds light on their comparative statics.

To illustrate the seller's problem and our main results, we next describe an example that is a special case of our model. Suppose there is a unit mass of ex-ante identical buyers with types drawn uniformly from $\Theta = [0, 1]$. The seller offers a price $p_i \in \mathbb{R}_+$ to each buyer i, and then buyers simultaneously decide whether to purchase. If a buyer of type $\theta_i \in \Theta$ purchases at a price p_i and the total quantity demanded (i.e., the total mass of buyers who purchase) is $q \in [0, 1]$, then the buyer's payoff is $\theta_i q - p_i$. The buyer purchases if, given the total quantity he anticipates, this payoff is weakly greater than his payoff from not purchasing, which is 0. Summarizing the seller's price offers by their distribution $\Pi \in \Delta(\mathbb{R}_+)$, a total quantity q is an equilibrium quantity given Π if it is the quantity demanded when buyers anticipate it.[1]

Suppose first that the seller sets only one price. That is, Π is degenerate, taking the form of a posted price. We show in Proposition 1 that this policy is optimal under best-case selection, namely if the seller were able to pick the equilibrium that buyers play whenever multiple equilibria arise. In this case, the seller would post a price $p^{\mathrm{B}} \approx 0.22$, yielding a best-case equilibrium with total quantity $q^{\mathrm{B}} \approx 0.66$ and revenue $R^{\mathrm{B}} \approx 0.148$. However, given this or any other strictly positive posted price, there is also an equilibrium with zero total quantity—no buyer is willing to purchase at such a price if they anticipate that no other buyer will purchase. This policy thus performs poorly under worst-case selection: the seller's guaranteed revenue is zero.

It is clear that the seller must offer some buyers a price of 0 to ensure a positive demand. How about then just setting two prices? The seller can offer a price 0 to a share $\pi \in (0, 1)$ of the population and some price p^* to the remaining $1 - \pi$ share. Buyers anticipate that at least π buyers will purchase,

[1] Given Π, if buyers anticipate a total quantity q, then the quantity demanded is $D_q(\Pi) := \int (1 - p/q)\, \mathrm{d}\Pi(p)$. The quantity q is an equilibrium quantity if $D_q(\Pi) = q$.

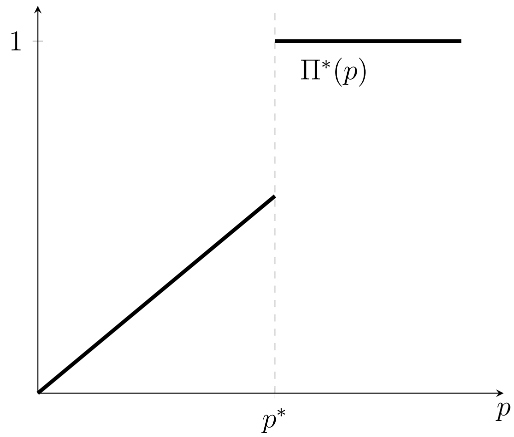

Figure 1: Optimal price distribution in the example described in the Introduction.

so setting $p^* \in (0, \pi)$ guarantees a positive revenue. In fact, we can verify that an optimal two-price distribution has $\pi \approx p^* \approx 0.25$, yielding a worst-case equilibrium quantity $q \approx 0.75$ and revenue $R \approx 0.125$. But why two prices and not more? For example, the seller could choose a uniform price distribution, $\Pi(p) = p/p^*$ for $p \in [0, p^*]$ and some $p^* > 0$ (perturbed to add some small mass at price 0). The optimal such distribution has $p^* \approx 0.41$, yielding a worst-case equilibrium quantity $q \approx 0.71$ and revenue $R \approx 0.126$.

Our results show that the seller's optimal price distribution in this example is indeed uniform, but only up to a mass point at the top. This distribution is given by $\Pi^*(p) = 2p$ for $p \in [0, p^*)$ and $\Pi^*(p) = 1$ for $p > p^*$, where $p^* \approx 0.28$. We provide an illustration in Figure 1. The worst-case equilibrium under Π^* has total quantity $q^* \approx 0.72$ and revenue $R^* \approx 0.133$.

The shape of the optimal price distribution reflects two goals of the seller. On the one hand, the seller wishes to ensure a high demand. We show that the optimal way to do so is by using a greedy function that builds the demand from the bottom, placing as little mass on low prices as is needed to iteratively rule out low-demand outcomes. This greedy function is uniform in the example above. On the other hand, the seller also wishes to extract revenue given the induced demand. The optimal way to do so is with a posted price, hence the mass point at the highest supported price p^*.

Theorem 1 provides a characterization for our general model. The main primitive of our model is the distribution over buyers' willingness to pay given an anticipated total quantity. We identify concavity conditions on this primitive under which the seller's optimal price distribution is greedy up to its highest supported price, with a mass point at that price. A key point in our proof is that contractions of the price distribution that preserve demand given an anticipated quantity increase both demand and revenue given any higher anticipated quantity. This is why a posted price obtains under best-case selection, and why greediness pins down price dispersion under worst-case selection. We prove that the greedy function—which corresponds to the solution to an integral equation—is continuous and strictly increasing. Thus, the seller's optimal policy can be interpreted as a posted price with (fully) dispersed discounts.

We use our characterization of the optimal policy to study the effects of externalities. We show that the seller's concern for strategic uncertainty in the presence of externalities results in not only price dispersion but also a higher total quantity of trade compared to the benchmark of best-case selection. The example described above, where $q^* > q^{\mathrm{B}}$, provides an illustration.[2] Turning to comparative statics, we find that the stronger the externalities between buyers, the less weight the seller puts on low prices, and the higher the total quantity that she induces. Our characterization also applies to a population where buyers belong to (observable) groups of heterogeneous strength of externalities. The seller's optimal policy in this case offers larger discounts to weak-externality buyers in order to build demand and extract higher revenue from strong-externality buyers.

We conclude with a discussion of variants of our model and potential avenues for future research. In this paper, we focus on a simple model that introduces externalities into an otherwise standard monopoly setting. We believe this framework can be enriched in a number of directions—for example,

[2] We also show that the worst-case solution has a higher maximum price—observe $p^* > p^{\mathrm{B}}$ in the example—yet it yields higher consumer surplus than the best-case benchmark.

incorporating congestion, dynamics, and two-sided markets—to shed further light on the use of pricing for coordination.

Literature. Our paper relates to three main literatures. First, there is a sizable literature on monopoly pricing under incomplete information but in the absence of externalities. Most useful to our analysis is Bulow and Roberts (1989), which relates concepts from optimal auction design (as in Myerson, 1981) to the problem of third-degree price discrimination under capacity constraints. We build on their insights to solve our benchmark problem of best-case selection in Section 3.

Second, there is also a large literature on markets with network externalities. Classic references include Rohlfs (1974), which highlights the possibility of multiple equilibria under a posted price, and Katz and Shapiro (1985, 1986) and Ellison and Fudenberg (2000), which study models of technology adoption with potentially incompatible products/upgrades. Oren, Smith and Wilson (1982), Csorba (2008), Aoyagi (2013), and Veiga (2018) consider settings more similar to ours, but the former two focus on second-degree price discrimination, and the latter two allow for schemes that condition on the number of buyers. We are not aware of work in this literature that studies optimal personalized pricing (or more general bilateral contracts; see Section 6) under incomplete information—neither with best-case nor with worst-case selection.

Third, our paper belongs to a growing literature on contracting with externalities that focuses on worst-case selection (or unique implementation). Following respectively the seminal contributions of Segal (2003) and Winter (2004), one strand of this literature studies settings where agents' actions are bilaterally contractible, as in our model, while another strand examines moral hazard problems with unobservable actions. Within the first strand, Halac, Kremer and Winter (2020) consider agents with heterogeneous but observable attributes.[3,4] Our main departure from this literature is that we study a

[3] Bernstein and Winter (2012) and Sákovics and Steiner (2012) examine related models with observable heterogeneity. More broadly, there is related work on exclusionary contracts; see, e.g., Rasmusen, Ramseyer and Wiley (1991); Innes and Sexton (1994); Segal and Whinston (2000); Spiegler (2000); Genicot and Ray (2006).

[4] Within the second strand, see, e.g., Eliaz and Spiegler (2015); Moriya and Yamashita

monopoly setting in which agents' attributes are hidden. We discuss how our analysis would change under complete information in Section 6.

Finally, in addition to these literatures, we relate to papers that predict pricing policies similar to the ones we characterize but in quite different environments. For example, Perry (1984) studies an incumbent firm that seeks to prevent entry and can post different prices for different units of its total supply. The firm uses a continuum of prices, with unlimited supply at the top and just enough supply at each lower price to make entry unattractive. In Heidhues and Kőszegi (2014), a monopolist sells to a loss-averse consumer who forms expectations prior to purchasing based on the monopolist's announced price distribution. To lure the consumer and exploit his attachment, an optimal distribution combines a continuum of sale prices with an atom at a high price. Our paper provides a complementary theory that emphasizes the role of externalities in consumption. These externalities and the strategic uncertainty they generate determine the optimal form of price dispersion in our model.

2. Model

Our model introduces strategic complementarities into a canonical monopoly setting. Below we describe the setup, the seller's problem, and our assumptions. We also provide examples of special cases.

2.1. Setup

We study a seller who sells a good to a population of buyers, each with a unit demand. Buyers' identities $i \in I := [0, 1]$ are uniformly distributed and independent of their payoff types $\theta \in \Theta$. The seller makes a price offer $p_i \in \mathbb{R}_+$ to each buyer $(i, \theta) \in I \times \Theta$. Prices are personalized, namely they can depend on a buyer's identity i. The offered price however cannot condition on a buyer's type θ, which is the buyer's private information.[5]

(2020); Halac, Lipnowski and Rappoport (2021, 2022); Cusumano, Gan and Pieroth (2023); Camboni and Porcellacchia (2024); Halac, Kremer and Winter (forthcoming).

[5] Since a buyer's type is independent of his identity, it is also independent of his price offer. Using this fact, we show in Section 6 that our focus on personalized price offers is without loss of generality within the class of public bilateral contracts.

Given the price offers, the buyers simultaneously decide whether or not to purchase from the seller. Denote the total quantity of purchases—i.e., the total mass of buyers who purchase—by $q \in [0,1]$. If a buyer of type $\theta \in \Theta$ purchases at a price p_i and the total purchased quantity is q, the buyer gets a payoff of

$$u(\theta, q) - p_i. \tag{1}$$

The measurable function $u : \Theta \times [0,1] \to \mathbb{R}_+$ is increasing in its second argument, reflecting that buyers' purchasing decisions are complementary. A buyer who does not purchase gets a payoff of 0.

The random variable $u(\cdot, q)$ represents a buyer's willingness to pay given an anticipated total quantity q. Let $F_q : \mathbb{R} \to [0,1]$ denote its cumulative distribution function. We assume F_q has support $[0, \overline{v}(q)] \subset \mathbb{R}_+$ for $q \in [0,1]$, where $\overline{v}$ is continuously differentiable with $\overline{v}' > 0$. This says that the lowest value is zero, whereas the highest value is strictly increasing in anticipated quantity. We further make the "cold-start" assumption that $\overline{v}(0) = 0$, so a buyer's value is almost surely zero if he anticipates no other buyer will purchase.[6] For strictly positive anticipated quantity $q \in (0,1]$, we suppose F_q admits a density f_q which is strictly positive on $(0, \overline{v}(q)]$, and that $f_q(v)$ is continuous in (q, v) where $0 \leq v \leq \overline{v}(q)$, having a partial derivative with respect to q that is also continuous in (q, v) over this domain.

Given a fixed anticipated quantity $q \in [0,1]$, the quantity that buyers demand and the seller's revenue can be easily computed. Assume that a buyer who is indifferent over purchasing chooses to purchase.[7] Then, given anticipated quantity q, the **quantity demanded** from a price p is equal to the mass of buyers whose willingness to pay is weakly greater than p, denoted $D_q(p) := 1 - F_q(p^-)$, and the quantity demanded from a price distribution $\Pi \in \Delta(\mathbb{R}_+)$ is $D_q(\Pi) := \int D_q(p)\, \mathrm{d}\Pi(p)$. Similarly, the seller's **revenue** from a price p is $R_q(p) := pD_q(p)$, and her revenue from a price distribution Π is

[6] We relax the zero-lowest-value and cold-start assumptions in Section 6.

[7] Without this tie-breaking assumption, our results would remain unchanged if we allow the seller to use (slightly) negative prices for a small fraction of the buyer population.

$R_q(\Pi) := \int R_q(p)\ \mathrm{d}\Pi(p)$.[8] By our assumptions on F_q, the demand function is continuous in quantity, and demand and revenue are jointly continuous (see Lemma 1 and Lemma 2 in the Appendix).

2.2. Seller's problem

The seller's price offers $(p_i)_{i \in I}$ induce a game between the buyers. In this game, each buyer (i, θ) simultaneously makes a decision of whether to purchase, with his payoff from purchasing given by (1). Since a buyer's identity conveys no information about his type, we can summarize the seller's price offers by their distribution $\Pi \in \Delta(\mathbb{R}_+)$. Given such a price distribution Π, if all buyers anticipate a total quantity q, the total quantity demanded is $D_q(\Pi)$. Thus, a total quantity q is an **equilibrium quantity** given Π if it is the quantity demanded when buyers anticipate it: $D_q(\Pi) = q$.

Due to the complementarity in buyers' purchasing decisions, multiple equilibrium quantities may arise given a price distribution Π. The seller wishes to guarantee a high revenue, and is therefore concerned with maximizing revenue in the worst-case equilibrium. Formally, her optimal value is given by

$$\begin{aligned} \sup_{\Pi \in \Delta(\mathbb{R}_+)} \min_{q \in [0,1]} \ & R_q(\Pi) \\ \text{subject to} \ \ & D_q(\Pi) = q. \end{aligned} \tag{P}$$

By the continuity of revenue and demand, the objective is continuous and the set of equilibrium quantities is closed and nonempty (see Lemma 3 in the Appendix). We say that (Π^*, q^*) is **optimal** if there exists a sequence $(\Pi_k, q_k)_k$ that converges to (Π^*, q^*) such that quantity q_k is the worst-case equilibrium quantity given price distribution Π_k for every k and $R_{q_k}(\Pi_k)$ converges to the seller's optimal value in (P).

Remark 1. The complementarity in buyers' purchasing decisions implies that the seller's revenue R_q is increasing in q. Hence, a worst-case equilibrium

[8] In a slight abuse of notation, we let $D_q(\Pi)$ and $R_q(\Pi)$ be similarly defined by such integrals for any function $\Pi : \mathbb{R}_+ \to \mathbb{R}$ such that the integral is well defined.

and a best-case equilibrium for the seller are respectively a lowest-quantity equilibrium and a highest-quantity equilibrium, and these equilibria exist for a given price distribution (see Lemma 3 in the Appendix).

Remark 2. While we have stated the seller's problem as maximizing revenue in the worst-case equilibrium, in our setting this is equivalent to maximizing revenue in the worst-case rationalizable outcome. The reason is that buyers' purchasing decisions are complementary, and thus the game they play under any price distribution is supermodular. The equivalence then follows from Guesnerie and Jara-Moroni (2011), who extend results of Milgrom and Roberts (1990) to games with a continuum of players. Further building on this observation and given Assumption 2 below, it will also turn out that the seller's worst-case problem is essentially equivalent to a more constrained one that maximizes revenue subject to inducing a unique equilibrium.

Remark 3. We have set up the model as one with positive externalities (i.e., with buyers' payoffs increasing in q). This implies that the worst-case equilibrium for the seller is also the worst-case equilibrium for the buyers. However, virtually nothing in our analysis changes if we assume that a buyer's payoff from purchasing is 0 while that from not purchasing is $-u(\theta, q)$.[9] In this case, there are negative externalities on nontraders (cf. Segal, 1999), and the seller's worst-case equilibrium is the best-case equilibrium for the buyers.

2.3. Concavity assumptions

We make three assumptions that we maintain throughout our analysis. Observe that while it is natural to describe our model in terms of the buyers' willingness-to-pay function $u(\theta, q)$ and the distribution of buyers' types $\theta \in \Theta$ (as we will do when providing examples in Section 2.4), there is a sense in which this is over-specified. In fact, different pairs of function $u(\theta, q)$ and distribution of θ map to the same distribution F_q over buyers' willingness to pay and therefore yield the same equilibrium conditions. Below, we thus express our model assumptions in terms of our model primitive F_q.

[9] The only result that changes in this case is the consumer surplus claim in Proposition 3.

Our first two assumptions concern the shape of externalities in buyers' purchasing decisions. Our model is one in which externalities are positive (but see Remark 3) and increasing (i.e., purchasing decisions are strategic complements).[10] Our first assumption strengthens the sense in which externalities are increasing by requiring a monotone likelihood ratio property (MLRP):

Assumption 1 (MLRP). *For any $0 < q \leq \tilde{q} \leq 1$, the likelihood ratio $f_{\tilde{q}}(v)/f_q(v)$ is weakly increasing in v over $(0, \overline{v}(q)]$.*

Our model assumption that $u(\cdot, q)$ is increasing in q says that the distribution of willingness to pay under an anticipated quantity $\tilde{q}$ first-order stochastically dominates that under any lower anticipated quantity $q \leq \tilde{q}$, and MLRP requires such dominance even when conditioning on any set of values. Using the Arrow-Pratt equivalence, this implies that for any price distribution $\Pi \in \Delta(\mathbb{R}_+)$ and quantities $0 < q \leq \tilde{q} \leq 1$, the demand $D_{\tilde{q}}(\Pi)$ is a concave transformation of the demand $D_q(\Pi)$ over the common support $[0, \overline{v}(q)]$.

Our second assumption requires the demand function to be strictly concave in anticipated quantity.

Assumption 2 (Concave externalities). *Wherever $q \in [0, 1]$ and $p \in \mathbb{R}_{++}$ have $p < \overline{v}(q)$, the demand function $D_q(p)$ is strictly concave in q.*

This assumption can be interpreted as saying that externalities in our model are concave: a buyer's payoff from purchasing relative to not purchasing increases with the anticipated total quantity demanded at a decreasing rate.

Finally, for our third assumption, we define the **cross virtual value** associated with a buyer's willingness to pay v. For any $0 < q \leq \tilde{q} \leq 1$, the cross virtual value function $\varphi_{q,\tilde{q}} : (0, \overline{v}(q)] \to \mathbb{R}$ is given by

$$\varphi_{q,\tilde{q}}(v) := \frac{f_{\tilde{q}}(v)}{f_q(v)} \left[v - \frac{1 - F_{\tilde{q}}(v)}{f_{\tilde{q}}(v)} \right].$$

This function is exactly the Myerson virtual value function under total quantity $\tilde{q}$ in the special case that $q = \tilde{q}$, and is otherwise the Myerson virtual value

[10] This terminology follows Segal (2003).

function under $\tilde{q}$ normalized by the likelihood ratio $f_{\tilde{q}}(v)/f_q(v)$. Recall that Myerson regularity says that the virtual value function $\varphi_{\tilde{q},\tilde{q}}$ is increasing. We make an analogous assumption on the cross virtual value function:

Assumption 3 (Cross regularity). *For any $0 < q \leq \tilde{q} \leq 1$, the cross virtual value function $\varphi_{q,\tilde{q}}(v)$ is strictly increasing in v over $(0, \overline{v}(q)]$.*

For intuition, fix an actual total quantity $\tilde{q}$. As noted by Bulow and Roberts (1989), the Myerson virtual value function corresponds to the seller's marginal revenue, and thus Myerson regularity implies that the seller's revenue is concave in the quantity demanded. Cross regularity serves an analogous role, but applies across different hypothetical anticipated quantities $q \leq \tilde{q}$. In particular, it implies that for any price distribution $\Pi \in \Delta(\mathbb{R}_+)$ and anticipated and actual total quantities $0 < q \leq \tilde{q} \leq 1$, the revenue $R_{\tilde{q}}(\Pi)$ is a concave transformation of the demand $D_q(\Pi)$ over the common support $[0, \overline{v}(q)]$.

2.4. Examples

We will illustrate our results with the following special cases of our model.

Linear demand. Suppose a buyer's willingness to pay given type $\theta \in \Theta$ and anticipated quantity $q \in [0,1]$ is $u(\theta, q) = \theta\overline{v}(q)$ for $\overline{v}$ satisfying $1/\overline{v}(q)$ strictly convex in q over $(0,1]$ (as well as $\overline{v}(0) = 0$ and $\overline{v}' > 0$), and let buyers' types be drawn uniformly from $\Theta = [0,1]$. The condition on $\overline{v}$ is equivalent to our concave externalities assumption; it holds, for example, if $\overline{v}$ is log-concave. The demand function takes the linear form $D_q(p) = 1 - p/\overline{v}(q)$, and one can verify that all of our model assumptions are satisfied. Our comparative-static results in Section 5 will focus on this environment.

Proportional values. Suppose a buyer's willingness to pay given type $\theta \in \Theta$ and anticipated quantity $q \in [0,1]$ is $u(\theta, q) = \theta q$. Denoting by G the distribution of buyers' types, with positive density g, we can then rewrite our assumptions as follows. MLRP and cross regularity say, respectively, that for all $\lambda > 1$, the ratio $g(\theta)/g(\lambda\theta)$ is weakly increasing in θ and the cross virtual value function

$$\frac{g(\theta)}{g(\lambda\theta)}\left[\theta - \frac{1 - G(\theta)}{g(\theta)}\right]$$

is strictly increasing in θ wherever $g(\lambda\theta)$ is strictly positive. Concave externalities says that $G(p/q)$ is convex in q. An example that satisfies these conditions is the power distribution with $g(\theta) = \kappa\theta^{\kappa-1}$ for $\kappa \geq 1$ and $\Theta = [0, 1]$.

Other examples. The examples described above take a willingness-to-pay function of the form $u(\theta, q) = \theta\overline{v}(q)$ for a buyer type $\theta \in \Theta$ and anticipated quantity $q \in [0, 1]$. Our model can also accommodate other formulations; for example, $u(\theta, q) = e^{\theta q} - 1$ paired with a uniform distribution of types over $\Theta = [0, 1]$ would satisfy all of our assumptions.

A natural setting that is outside our model as stated is one in which a buyer's willingness to pay takes an additive form, $u(\theta, q) = \theta + q$. This formulation does not satisfy our zero-lowest-value and cold-start assumptions. Both of these assumptions however can be relaxed—see Section 6 for details—and our main takeaways remain valid in settings like the additive one.

3. Benchmark: best-case selection

Before we solve the seller's problem in (P), we consider a benchmark setting in which the seller has no concern for strategic uncertainty. Suppose that for any price distribution $\Pi \in \Delta(\mathbb{R}_+)$ that the seller chooses, she can select the equilibrium that buyers play in the induced game if multiple equilibria arise. Rather than being concerned with the worst case as in (P), such a seller maximizes revenue in the best-case equilibrium:

$$\sup_{\Pi \in \Delta(\mathbb{R}_+)} \max_{q \in [0,1]} R_q(\Pi)$$
$$\text{subject to} \quad D_q(\Pi) = q.$$

We find that the seller's solution under best-case selection takes the form of a posted price.

Proposition 1. *Under best-case selection, some optimum exists, and any optimum has strictly positive revenue and degenerate price distribution.*

The argument builds on Bulow and Roberts (1989). Recall that by cross regularity (Assumption 3)—in fact, by Myerson regularity, which is implied by cross-regularity—the seller's revenue is concave in the quantity demanded. This means that revenue increases if price dispersion is reduced in a way that keeps quantity unchanged. Hence, given any nondegenerate price distribution Π and best-case equilibrium quantity $q > 0$ that it induces, the seller can improve upon Π with a q-preserving posted price, namely by offering each buyer a price $p = D_q^{-1}(q) \in [0, \overline{v}(q)]$.

The posted-price mechanism is a familiar result in the monopoly setting, as it is the one that obtains in the absence of externalities. If, for a fixed quantity $q \in (0, 1]$, our seller faced an exogenous distribution F_q of buyer values and thus an exogenous demand curve D_q, she would maximize revenue by offering the same price, call it $p^M(q)$, to each buyer. No equilibrium multiplicity would arise in such a setting with no externalities, and by (cross) regularity, the optimal posted price would be the unique $p^M \in (0, \overline{v}(q))$ with $\varphi_{q,q}(p^M) = 0$.

Proposition 1 tells us that the presence of externalities does not alter the nature of the seller's optimal mechanism provided that the seller can select her preferred equilibrium. The seller uses a posted price as in a standard monopoly setting, although naturally the externalities do affect the price she chooses. The complementarity in buyers' purchasing decisions implies that the total quantity demanded is more responsive to price changes. Thus, letting p^B and q^B be respectively an optimal price and equilibrium quantity that solve the best-case problem above, one can show that $p^B \leq p^M(q^B)$, with strict inequality under sufficient smoothness conditions.

4. Optimal price distribution

We now return to the problem in (P), where the seller is concerned with worst-case outcomes. The seller chooses a price distribution to maximize revenue in her least preferred equilibrium of the induced game between the buyers. In Section 4.1, we present an auxiliary program that clarifies the key constraints that the worst-case focus introduces. We then use this auxiliary program in Section 4.2 to derive our main result on the seller's optimal price distribution. We provide intuition for the proof of this result in Section 4.3.

4.1. Which constraints matter?

Recall that (Π^*, q^*) is optimal if it is the limit of a sequence $(\Pi_k, q_k)_k$ of price distributions and corresponding worst-case equilibrium quantities whose revenue $R_{q_k}(\Pi_k)$ converges to the seller's optimal value in (P). The next proposition establishes that (Π^*, q^*) can be computed as the solution to an auxiliary program.

Proposition 2. (Π^*, q^*) *is optimal if and only if it solves*

$$\max_{\Pi \in \Delta(\mathbb{R}_+),\ q \in [0,1]} R_q(\Pi) \tag{P*}$$

$$\textit{subject to} \quad D_{\hat{q}}(\Pi) \geq \hat{q} \quad \forall \hat{q} \in (0, q).$$

Moreover, this program has a maximizer, generating strictly positive revenue.

Proposition 2 elucidates the constraints that are introduced by the seller having a concern for strategic uncertainty. Observe that as in the best-case benchmark of Section 3, program (P*) maximizes over both a price distribution $\Pi \in \Delta(\mathbb{R}_+)$ and an equilibrium quantity $q \in [0, 1]$, and any optimum (Π^*, q^*) in (P*) satisfies the equilibrium condition $D_{q^*}(\Pi^*) = q^*$ (for otherwise raising q^* would yield a strict improvement). However, there are additional constraints that (P*) imposes to guarantee that (Π^*, q^*) is optimal under worst-case selection. Plainly, the seller's price distribution cannot admit any lower-quantity equilibrium, and therefore the demand $D_{\hat{q}}(\Pi)$ at any anticipated quantity $\hat{q} < q^*$ must exceed $\hat{q}$. Program (P*) imposes these demand constraints as weak inequalities, with the solution being the limit of a sequence $(\Pi_k, q_k)_k$ that satisfies the constraints strictly for every k.

To prove Proposition 2, we first show that the auxiliary program (P*) is a relaxation of the original program (P). In fact, any price distribution $\Pi \in \Delta(\mathbb{R}_+)$ and its corresponding worst-case equilibrium quantity q are feasible in (P*): if $q = 0$, the program imposes no constraints, and if $q > 0$, then this being the lowest equilibrium quantity under Π implies that the constraints in (P*) hold strictly for all $\hat{q} \in [0, q)$.[11] Next, in the other direction, we show that (P*) cannot yield strictly higher revenue than (P). Given (Π, q) feasible in (P*), we construct a perturbed price distribution Π_ε which coincides with Π except for a small fraction ε of buyers who are offered a zero price. For every $\varepsilon > 0$, the price distribution Π_ε generates a worst-case equilibrium quantity $q_\varepsilon \geq q$. Hence, since revenue is increasing in the quantity demanded and Π_ε converges to Π as $\varepsilon \to 0$, we obtain $R_{q_\varepsilon}(\Pi_\varepsilon) \geq R_q(\Pi)$ in this limit.

While program (P*) clarifies which constraints are the relevant ones for the seller under worst-case selection, it is not immediate what the solution to this program looks like. The seller chooses a continuum of prices which must satisfy a continuum of demand constraints. At each of these constraints, she faces tradeoffs between increasing one price versus lowering another one to preserve demand, and the tightness of the constraints depends on the relative slopes of the demand function at different anticipated total quantities. It might be intuitive to think that the solution to (P*) should satisfy all the demand constraints with equality—but this may not be feasible for a given target quantity, and even when feasible, we will see that it is not optimal.

In the next two sections, we show that the principle that guides the solution to program (P*) is essentially the same principle behind the seller's solution in the benchmark of best-case selection. This principle is that price dispersion is bad for revenue, so quantity-preserving contractions of the price distribution benefit the seller. Of course, unlike under best-case selection, the result will not be a degenerate price distribution; as discussed in the Introduction, price dispersion is needed to generate strictly positive revenue under worst-case selection. Instead, our analysis will show how this fact can be used to pin down the optimal form of price dispersion.

[11] In particular, note that $D_0(\Pi) \geq 0$ and the demand function is continuous in anticipated quantity. Thus, if any $\hat{q} \in (0, q)$ had $D_q(\Pi) \leq q$, then some $\tilde{q} \in [0, \hat{q}]$ would be an equilibrium quantity, meaning q is not the worst equilibrium.

4.2. Posted price with dispersed discounts

We define a class of functions that we will use in our characterization of the seller's optimal pricing policy.

Definition 1. *Let $\Gamma : \mathbb{R}_+ \to \mathbb{R}_+$ be right-continuous and nondecreasing. Given $p \in \mathbb{R}_+$, say Γ is* ***greedy up to p*** *if it satisfies*

$$D_{\hat{q}}(\Gamma) = \hat{q}$$

for all $\hat{q} \in (0,1)$ with $\overline{v}(\hat{q}) \leq p$. Say Γ is ***greedy*** *if it is greedy up to every $p \geq 0$.*

A function Γ that is greedy up to p satisfies the demand constraints in program (P^*) with equality for all anticipated quantities for which the highest willingness to pay is no greater than p. This means that Γ iteratively sets to zero the demand-constraint difference $D_{\hat{q}}(\Gamma) - \hat{q}$ starting from the lowest anticipated quantity up to $\underline{q}(p) := \overline{v}^{-1}(p)$. Intuitively, a greedy function follows a greedy procedure: for each anticipated quantity $\hat{q}$ starting from 0, given a measure over prices $[0, \overline{v}(\hat{q}))$, the seller pushes up the next prices as much as possible subject to satisfying the demand constraint at $\hat{q}$. Following this greedy procedure up to q is equivalent to solving the Volterra integral equation $\int_0^{\overline{v}(\hat{q})} \Pi(p) f_{\hat{q}}(p)\, \mathrm{d}p = \hat{q}$ for all $\hat{q} \in (0, q)$.

The next theorem presents our main result.

Theorem 1. *Suppose (Π^*, q^*) is optimal, and let p^* be the highest price in the support of Π^*. Then $p^* \leq p^M(q^*)$, and Π^* is greedy up to p^*, with a mass point at p^*.*

A seller's optimal price distribution balances two goals. On the one hand, being concerned with worst-case outcomes, the seller wishes to insulate against low-quantity equilibria. She does so by using a greedy function that seeds the demand from the bottom, placing as little mass on low prices as is needed to iteratively rule out low-demand outcomes.[12] On the other hand, the seller also wishes to extract revenue from the induced higher demand. This is achieved with the mass point at the highest offered price p^*—such an extraction point plays the same role as the seller's posted price $p^M(q^*)$ in the standard monopoly problem with exogenous demand D_{q^*}. The resulting price distribution maximizes the seller's worst-case revenue by minimizing the demand constraints up to anticipated quantity $\underline{q}^* := \overline{v}^{-1}(p^*)$ and satisfying with slack the demand constraints for quantities in $(\underline{q}^*, q^*)$.

Theorem 1 suggests an appealing interpretation for the seller's optimal pricing policy: the seller posts a high price and simultaneously offers personalized discounts to some buyers to build a high demand. The use of list prices together with promotions and special deals that vary across buyers is common in applications. The shape of Π^* tells us precisely how these personalized discounts are optimally distributed in the population. We show in the Appendix that that any greedy function must be continuous and strictly increasing. Hence, it follows from Theorem 1 that the seller's optimal price distribution has only one mass point, and personalized discounts are (fully) dispersed across buyers.

Corollary 1. *Any optimal price distribution is continuous and strictly increasing up to a mass point at the top of its support. Said differently, the seller's policy is a posted price with dispersed discounts.*

In many environments, one can verify directly that the seller's problem admits a unique greedy function Γ^* over $[0, \overline{v}(1))$. In such cases, Theorem 1 reduces the seller's problem to a one-parameter optimization over q^*, as any optimal (Π^*, q^*) must then have Π^* coincide with the unique greedy function Γ^* up to its highest supported price $p^* \in (0, \overline{v}(q^*))$. This is the case in the examples that we describe next, and more generally we show it is always true in the linear demand environment—see Lemma 9 in the Appendix.

[12] Recall that while the demand constraints in program (P*) are weak inequalities, the solution (Π^*, q^*) is the limit of a sequence $(\Pi_k, q_k)_k$ which, for every k, satisfies the demand constraints strictly and thus rules out equilibrium quantities $\hat{q} < q_k$.

Figure 2 illustrates our results with three examples. Each graph depicts the unique greedy function $\Gamma^*(p)$ (gray dotted line) and the seller's optimal price distribution $\Pi^*(p)$ (black solid line), which is also unique. The first two examples, in the top panel, belong to the linear demand environment—with a willingness-to-pay function $u(\theta, q) = \theta\overline{v}(q)$ and a uniform distribution of types. The first example on the left is the one discussed in the Introduction, with $\overline{v}(q) = q$. The second example on the right takes $\overline{v}(q) = q + q^2$. Finally, the third example in the bottom panel belongs to the proportional values environment, with $u(\theta, q) = \theta q$ and a power distribution of types.

Observe that the first and third examples in Figure 2 both have proportional values, and while they assume different distributions of types, in both cases the unique greedy function is uniform (and given by $\Gamma^*(p) = p/\mathbb{E}[\theta]$). This is not a coincidence, as we report in the next corollary.

Corollary 2. *In the proportional values environment, the seller's policy is a posted price with uniform discounts.*

4.3. Intuition for proof of Theorem 1

We next provide intuition for the proof of Theorem 1. To highlight the main ideas, we focus on the linear demand environment. We comment on the differences with respect to our general proof in the Appendix at the end of this section.

As a preliminary step, observe that under a linear demand, we can rewrite the demand constraints in the auxiliary program (P*) as follows:

$$\int_0^{\overline{v}(\hat{q})} \Pi(p)\mathrm{d}p \geq \hat{q}\overline{v}(\hat{q}) \quad \forall \hat{q} \in (0, q). \tag{2}$$

One can readily verify that all of these constraints are satisfied with equality if Π agrees with the unique greedy function Γ^* up until at least $\overline{v}(q)$, where

$$\Gamma^*(p) = \underline{q}(p) + p\underline{q}'(p). \tag{3}$$

Suppose by contradiction that (Π^*, q^*) is optimal and Π^* is not greedy up to

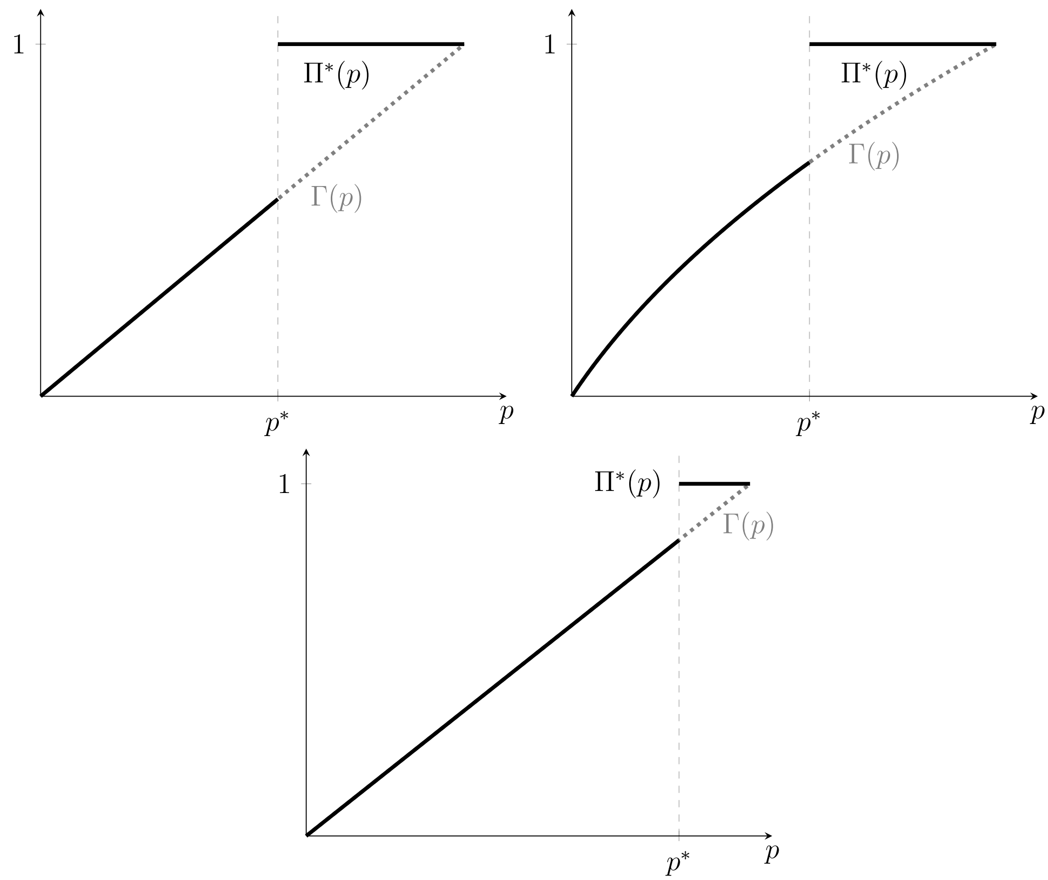

Figure 2: Greedy function (gray dotted line) and optimal price distribution (black solid line). The top-panel examples take a linear demand environment. The left example takes $\overline{v}(q) = q$, and has $\Gamma^*(p) = 2p$ with $p^* \approx 0.28$ and $q^* \approx 0.72$. The right example takes $\overline{v}(q) = q + q^2$, and has $\Gamma^*(p) = (\sqrt{1+4p} - 1)/2 + p/\sqrt{1+4p}$ with $p^* \approx 0.51$ and $q^* \approx 0.76$. The bottom-panel example takes proportional values with $g(\theta) = 2\theta$ over $\Theta = [0, 1]$, and has $\Gamma^*(p) = (3/2)p$ with $p^* \approx 0.56$ and $q^* \approx 0.76$.

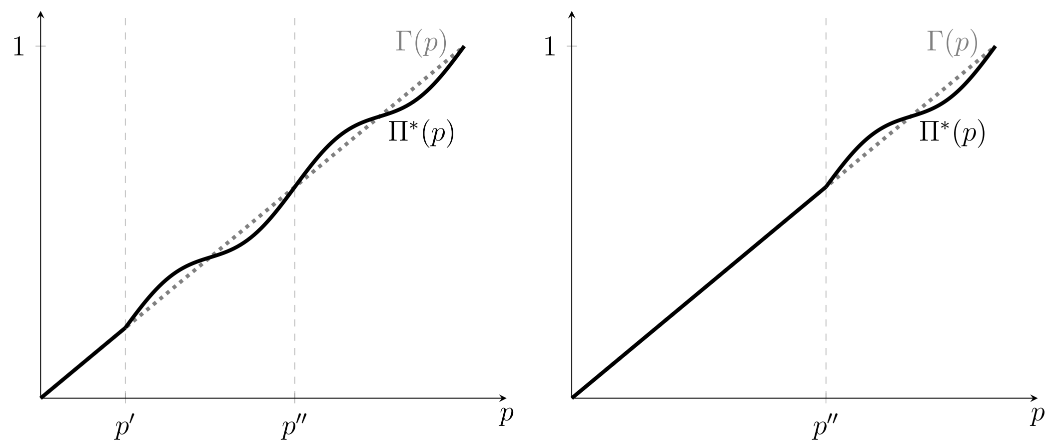

Figure 3: Illustration of arguments in Section 4.3. See the text for details.

its highest supported price p^*. Since we have shown in Proposition 2 that the seller's optimal value is strictly positive, we take $q^*, p^* > 0$. By definition of the greedy function Γ^*, and assuming here that $\Pi^* - \Gamma^*$ is piecewise monotone, it follows that there exists some price $p' \in [0, \overline{v}(q^*))$ such that $\Pi^*(p) = \Gamma^*(p)$ for $p \leq p'$ and $\Pi^*(p) > \Gamma^*(p)$ right above p'. An illustration is provided in the left panel of Figure 3, where we have drawn $\Gamma^*(p)$ (gray dotted line) for the same environment as in the top left example of Figure 2.

There are two scenarios to consider. First, suppose that there exists a price $p'' \in (p', \overline{v}(q^*))$ such that

$$\int_{p'}^{p''} [\Pi^*(p) - \Gamma^*(p)] \, \mathrm{d}p = 0. \tag{4}$$

Then we can take the lowest such price p'', in which case

$$\int_{p'}^{\hat{p}} [\Pi^*(p) - \Gamma^*(p)] \, \mathrm{d}p > 0 \quad \forall \hat{p} \in (p', p''), \tag{5}$$

as illustrated in Figure 3.

Now let us define a new price distribution $\tilde{\Pi}$ which coincides with the greedy function Γ^* up to p'' and is otherwise equal to Π^*:[13]

$$\tilde{\Pi}(p) = \begin{cases} \Gamma^*(p) & \text{for } p < p'' \\ \Pi^*(p) & \text{otherwise.} \end{cases}$$

The right panel of Figure 3 provides an illustration. By definition of Γ^*, the price distribution $\tilde{\Pi}$ satisfies the demand constraints for all anticipated quantities $\hat{q} \in (0, \underline{q}(p''))$. Moreover, observe that by (4) and (5), $\tilde{\Pi}$ is a mean-preserving contraction of Π^* below p''. MLRP (Assumption 1) therefore implies $D_{\hat{q}}(\tilde{\Pi}) \geq D_{\hat{q}}(\Pi^*)$ for all $\hat{q} \in [\underline{q}(p''), 1]$, which means that $\tilde{\Pi}$ also satisfies the demand constraints for all anticipated quantities $\hat{q} \in [\underline{q}(p''), q^*)$. Furthermore, cross regularity (Assumption 3) implies $R_{\hat{q}}(\tilde{\Pi}) > R_{\hat{q}}(\Pi^*)$ for all $\hat{q} \in [\underline{q}(p''), 1]$. It follows that $\tilde{\Pi}$ yields strictly higher revenue than Π^*, contradicting the assumption that Π^* is optimal.

We are then left with the second scenario, in which no $p'' \in (p', \overline{v}(q^*))$ exists that satisfies equation (4). In this case, each $\hat{p} \in (p', \overline{v}(q^*))$ has

$$\int_{p'}^{\hat{p}} [\Pi^*(p) - \Gamma^*(p)] \, \mathrm{d}p > 0.$$

By (2), it follows that the corresponding demand constraints are satisfied with slack; that is, $D_{\hat{q}}(\Pi^*) > \hat{q}$ for all anticipated quantities $\hat{q} \in (\underline{q}(p'), q^*)$. However, this means that if Π^* puts any mass (strictly) above p', then again a strict improvement is feasible. Specifically, if Π^* is supported on more than one price above p', we show that a small mean-preserving contraction above p' preserves the demand constraints (by them being slack) and increases revenue (by cross regularity). If Π^* has only one mass point above p', then satisfaction of the demand constraints for quantities right above $\underline{q}(p')$ requires a mass point also at p', and we show that revenue can be increased with a small contraction that takes mass from these two points. We thus conclude that Π^* cannot have support above p', i.e., $p' =: p^*$. This contradicts the assumption that Π^* is not

[13] We can verify that the price p'' satisfies $p'' < p^*$ and $\Gamma^*(p'') \leq \Pi^*(p'')$, so $\tilde{\Pi}$ is a distribution function.

greedy up to p^*.

The steps above yield that any optimal price distribution Π^* coincides with the greedy function Γ^* up to its highest supported price p^*. Theorem 1 also states that p^* is no greater than the monopoly price $p^M(q^*)$. This is intuitive: by definition of $p^M(q^*)$, any price distribution with highest price $p^* > p^M(q^*)$ can be improved upon by lowering all prices above $p^M(q^*)$ to this level. Finally, to prove that Π^* has a mass point at p^*, observe that since $p^M(q^*) < \overline{v}(q^*)$, we have $p^* < \overline{v}(q^*)$. Hence, while the greedy function Γ^* up to p^* satisfies the demand constraints up to $\underline{q}^* = \underline{q}(p^*)$, there must be a mass point at p^* to satisfy the demand constraints over $(\underline{q}^*, q^*)$.

The proof of Theorem 1 in the Appendix proceeds via perturbations as we did here: taking a price distribution Π^* that is not greedy up to the top of its support, and showing how it can be improved while preserving the demand constraints. However, we do not build on a fixed greedy function Γ^*, nor do we rely on $\Pi^* - \Gamma^*$ being well-behaved. Instead, we show that given Π^*, we can locate an interval of anticipated quantities where the demand constraints are slack, and where we can apply contraction arguments analogous to those used in the second scenario above. For locating such an interval, concave externalities (Assumption 2) is important. For arguing that contractions improve revenue while satisfying the demand constraints, a difficulty is that these constraints do not take the form of majorization as in (2) outside the linear demand environment. While this means that we cannot use off-the-shelf comparative statics on mean-preserving contractions as we did above, we show that similar comparative statics can be derived for our general model. This step extends results from Rappoport (2024); see Lemma 4 in the Appendix.

5. The effects of externalities

We use our characterization of the seller's optimal pricing policy to study the effects of externalities. In Section 5.1, we compare the seller's solution under worst-case selection to the best-case benchmark of Section 3. In Section 5.2, we examine how the solution changes as the externalities in buyers' purchasing decisions become stronger. Finally, in Section 5.3, we consider a setting where buyers belong to groups with heterogeneous strength of externalities, and the seller's price offers can condition on buyer group. Throughout this section, we focus on the linear demand environment.

5.1. Worst-case versus best-case

The seller's optimal pricing policy in Theorem 1 is shaped by her concern for strategic uncertainty. Recall from Proposition 1 that under best-case selection—i.e., if the seller could choose the equilibrium that buyers play given her price offers—a posted price mechanism would be optimal. Instead, when concerned with worst-case outcomes, Theorem 1 says that the seller uses a posted price together with personalized discounts. An immediate consequence of the seller's worst-case focus is thus price dispersion. But, what does this imply for price levels and for the induced total quantity of trade? And how does the resulting consumer surplus compare under worst-case versus best-case selection? The next proposition provides answers to these questions. We denote the consumer surplus associated with anticipated quantity $q \in [0, 1]$ and price $p \in \mathbb{R}_+$ by

$$\mathrm{CS}_q(p) := \int_p^{\overline{v}(q)} D_q(v) \, \mathrm{d}v,$$

and let $\mathrm{CS}_q(\Pi) := \int \mathrm{CS}_q(p) \, \mathrm{d}\Pi(p)$ for any price distribution $\Pi \in \Delta(\mathbb{R}_+)$.

Proposition 3. *Take the linear demand environment. Relative to the best-case benchmark, the seller's worst-case solution has a higher maximum offered price $p^* > p^B$, induces a higher total quantity $q^* > q^B$, and yields a higher consumer surplus $CS_{q^*}(\Pi^*) > CS_{q^B}(p^B)$.*

This result reveals that not all buyers benefit from lower prices when the seller is concerned with worst-case outcomes. Some buyers receive generous discounts as the seller seeks to ensure a high demand, but others receive price offers strictly higher than the seller's best-case posted price. At the same time, Proposition 3 tells us that buyers on average do purchase at a lower price in the worst-case solution, and thus the total quantity demanded is higher than in the best-case benchmark. Interestingly, while the seller is concerned with ruling out low-quantity equilibria—and must therefore offer discounts and sacrifice revenue to ensure any quantity as a worst-case equilibrium—she ends up inducing a higher quantity of trade than in the absence of this concern. Proposition 3 further establishes that, as a consequence, consumer surplus increases due to the seller's worst-case focus.

The linear-demand example discussed in the Introduction (with $\overline{v}(q) = q$) provides an illustration of the comparisons reported in Proposition 3. As noted, the seller's worst-case and best-case solutions in that setting have maximum prices $p^* \approx 0.28 > 0.22 \approx p^B$ and total quantities $q^* \approx 0.72 > 0.66 \approx q^B$. The resulting consumer surpluses are $\mathrm{CS}_{q^*}(\Pi^*) \approx 0.19 > 0.15 \approx \mathrm{CS}_{q^B}(p^B)$.

To prove Proposition 3, we use our characterizations of the seller's worst-case and best-case optima. By Theorem 1, any optimal worst-case price distribution is a function $\Pi(\cdot|\hat{p})$ that coincides with the greedy function—unique in the linear demand environment—up to a highest supported price $\hat{p}$. Define $\mathcal{R}(\hat{p}, \hat{q})$ as the seller's worst-case revenue given such a price distribution $\Pi(\cdot|\hat{p})$ and a buyers' anticipated quantity $\hat{q}$. In a worst-case equilibrium, $\hat{q}$ is equal to the lowest quantity demanded given that $\Pi(\cdot|\hat{p})$ is the limit worst-case price distribution; call it $\mathcal{Q}(\hat{p})$. Then $\mathcal{R}(\hat{p}, \mathcal{Q}(\hat{p}))$ gives the seller's worst-case revenue parametrized by $\hat{p}$. We define analogous objects for the best-case problem, with $\mathcal{R}^B(\hat{p}, \hat{q})$ being the seller's best-case revenue given a posted price $\hat{p}$ and anticipated quantity $\hat{q}$, and $\mathcal{R}^B(\hat{p}, \mathcal{Q}^B(\hat{p}))$ taking $\hat{q}$ to equal the highest equilibrium quantity $\mathcal{Q}^B(\hat{p})$ under $\hat{p}$.

Our analysis is facilitated by the fact that, in the linear demand environment, these worst-case and best-case revenue functions are strictly quasiconcave, with unique interior optima p^* and p^B given by

$$\left.\frac{\mathrm{d}\mathcal{R}(\hat{p}, \mathcal{Q}(\hat{p}))}{\mathrm{d}\hat{p}}\right|_{\hat{p}=p^*} = 0 \quad \text{and} \quad \left.\frac{\mathrm{d}\mathcal{R}^B(\hat{p}, \mathcal{Q}^B(\hat{p}))}{\mathrm{d}\hat{p}}\right|_{\hat{p}=p^B} = 0. \tag{6}$$

Hence, to establish the ranking between p^* and p^B, it suffices to sign

$$\frac{\mathrm{d}\mathcal{R}}{\mathrm{d}\hat{p}} = \underbrace{\frac{\partial \mathcal{R}}{\partial \hat{p}}}_{\substack{\text{monopoly}\\ \text{effect}}} + \underbrace{\frac{\partial \mathcal{R}}{\partial \hat{q}} \frac{\mathrm{d}\mathcal{Q}}{\mathrm{d}\hat{p}}}_{\substack{\text{externality}\\ \text{effect}}} \tag{7}$$

at $\hat{p} = p^B$. We call the first term on the right-hand side the **monopoly effect**. This effect tells us how revenue changes with $\hat{p}$ while keeping the anticipated quantity $\hat{q}$, and thus the demand function $D_{\hat{q}}$, fixed. As is familiar, raising $\hat{p}$ increases revenue from inframarginal buyers via a higher price, but reduces revenue from marginal buyers via a lower quantity. Note that if $\hat{p} = p^B$, the quantity demanded at the worst-case highest price $\hat{p}$ is larger than the quantity demanded at the best-case posted price p^B. A comparison of the monopoly effects then follows from (cross) regularity (Assumption 3): conditional on pricing at $\hat{p}$, the worst-case monopoly effect of raising $\hat{p}$ starting from $\hat{p} = p^B$ is higher (i.e., more positive) than the analog best-case monopoly effect.

The second term on the right-hand side of (7) is the **externality effect**. This effect tells us how revenue changes as the demand function $D_{\hat{q}}$ shifts towards the new equilibrium—that is, given that the anticipated quantity $\hat{q}$ must adjust to match the quantity demanded $\mathcal{Q}(\hat{p})$ following an increase in $\hat{p}$. We show that conditional on pricing at $\hat{p}$, the worst-case externality effect of raising $\hat{p}$ starting from $\hat{p} = p^B$ is higher than the analog best-case externality effect. Hence, given the definition of p^B in (6), the monopoly and externality effects imply $\mathrm{d}\mathcal{R}(\hat{p}, \mathcal{Q}(\hat{p}))/\mathrm{d}\hat{p} > 0$ at $\hat{p} = p^B$. We conclude that the worst-case highest price p^* is strictly higher than the best-case posted price p^B.

The idea behind the ranking of the worst-case and best-case quantities, q^* and q^B, is similar. We show that $\mathcal{R}(\hat{p}, \mathcal{Q}(\hat{p}))$ increases as $\hat{p}$ is reduced from a level that makes the worst-case equilibrium quantity equal to q^B, and therefore the optimal such quantity must satisfy $q^* > q^B$. Finally, this result allows us to prove the ranking on consumer surplus. In particular, using $q^* > q^B$, we show that the average price offered under a price distribution $\Pi(\cdot|\overline{v}(q^B))$ is no greater than p^B.[14] Since consumer surplus is decreasing in average price, and

[14] This follows from the fact that, given $q^* > q^B$, the demand constraint in program (P*)

is strictly increasing in price dispersion due to buyer option value, it follows that $\mathrm{CS}_{q^B}(\Pi(\cdot|\overline{v}(q^B))) > \mathrm{CS}_{q^B}(p^B)$. Moreover, since consumer surplus is also increasing in quantity, we obtain

$$\mathrm{CS}_{q^*}(\Pi(\cdot|p^*)) \geq \mathrm{CS}_{q^B}(\Pi(\cdot|p^*)) = \mathrm{CS}_{q^B}(\Pi(\cdot|\overline{v}(q^B))) > \mathrm{CS}_{q^B}(p^B).$$

5.2. Strength of externalities

The externalities in buyers' purchasing decisions are a key novel ingredient of our seller's problem. We next study how the seller's solution in Theorem 1 changes as these externalities become stronger. Recall that in the linear demand environment, a buyer's willingness to pay given type $\theta \in \Theta$ and anticipated quantity $q \in [0,1]$ is $u(\theta, q) = \theta\overline{v}(q)$. The strength of externalities is reflected in the function $\overline{v}$.

Definition 2. *In the linear demand environment, say $\overline{v}_1$ has* **stronger externalities** *than $\overline{v}_0$ if, for all $q \in (0,1]$,*

(i) $\overline{v}_1(q) > \overline{v}_0(q)$, and

(ii) $\overline{v}_1(q)/\overline{v}_0(q)$ is strictly increasing in q.

Intuitively, buyers' purchasing decisions are more complementary if their willingness to pay grows with the anticipated total quantity demanded at a higher rate, as captured by (ii).[15] Since our model assumes $\overline{v}(0) = 0$, it is then also natural that buyers' willingness to pay under any given anticipated total quantity will be higher when externalities are stronger, as captured by (i).

Proposition 4. *Take the linear demand environment. Suppose $\overline{v}_1$ has stronger externalities than $\overline{v}_0$, with corresponding optimal price distributions Π_1^* and Π_0^*. Relative to Π_0^*, then Π_1^* induces a higher total quantity $q_1^* > q_0^*$. Moreover, Π_1^* puts lower weight on low prices: $\Pi_1^*(p) < \Pi_0^*(p)$ for all $p < \min\{p_1^*, p_0^*\}$.*

---

is satisfied at anticipated quantity $\hat{q} = q^B$.

[15] Observe that taking $\overline{v}_1'(\cdot) > \overline{v}_0'(\cdot)$ would not yield the desired externality order: multiplying $\overline{v}$ by a constant $\kappa > 0$ has no effect on the seller's solution up to a change of numeraire.

An increase in the strength of externalities makes it less costly for the seller to insulate against low-demand equilibria. Specifically, take any anticipated quantity $\hat{q} \in (0, 1)$ and greedy prices over $[0, \overline{v}(\hat{q}))$ that satisfy the demand constraints in program (P*) up to $\hat{q}$. Under stronger externalities, because the demand is more responsive to anticipated quantity, the seller can then satisfy the demand constraint at $\hat{q}$ without the need to offer price $\overline{v}(\hat{q})$ to such a large mass of buyers. As a result, the optimal price distribution places relatively less weight on discounted prices below a given posted price. Moreover, because the seller can guarantee a given equilibrium quantity while charging higher prices, it is optimal for her to induce a higher quantity when externalities are stronger. These price and quantity effects combined explain why Proposition 4 does not pin down how the posted price itself changes with the externalities.

For illustration, we can compare the linear-demand examples shown in the top panel of Figure 2. The second example on the right (with $\overline{v}(q) = q + q^2$) has stronger externalities than the first example on the left (with $\overline{v}(q) = q$), and accordingly induces a higher total quantity of trade (as reported in the figure caption). The second example also exhibits a higher posted price and lower weight on discounted prices below the first-example posted price.

The proof of the comparative static concerning the seller's optimal price distribution follows directly from equation (3), which defines the unique greedy function in the linear demand environment. If $\overline{v}_1$ has stronger externalities than $\overline{v}_0$, then the greedy function under $\overline{v}_1$ is lower than that under $\overline{v}_0$ in the first-order stochastic dominance (FOSD) sense.

To prove the comparative static on the optimal total quantity, we use arguments similar to those described in Section 5.1. Given $\overline{v}_1$ and $\overline{v}_0$, let $\mathcal{R}_1(\hat{p}, \mathcal{Q}_1(\hat{p}))$ and $\mathcal{R}_0(\hat{p}, \mathcal{Q}_0(\hat{p}))$ be the respective revenue functions parametrized by the highest offered price $\hat{p}$. We study how the strong-externality revenue $\mathcal{R}_1$ changes as we increase the highest price $\hat{p}$, starting from a level that makes the induced strong-externality quantity equal to the optimal weak-externality quantity q_0^*. By the FOSD ranking of the greedy functions, such a starting level for $\hat{p}$ is strictly higher than p_0^*. We then show that increasing $\hat{p}$ from that level causes strong-externality monopoly and externality effects, as defined in equation (7), which are both lower (i.e., more negative) than the corresponding weak-externality effects caused by increasing $\hat{p}$ from p_0^*.[16] Since the latter weak-externality effects add to zero by definition of p_0^*, this means that the strong-externality revenue $\mathcal{R}_1$ can be increased by lowering $\hat{p}$. Thus, we obtain that the optimal total quantities satisfy $q_1^* > q_0^*$, as stated in Proposition 4.

5.3. Heterogeneity

In the previous section, we studied how the seller's pricing policy changes as the externalities in buyers' purchasing decisions become stronger. A related but distinct question is how the seller's policy changes if the strength of externalities varies across buyers. For example, take a seller of file sharing services. Because these services are more heavily used in the corporate sector, corporate buyers' willingness to pay would tend to be higher and to grow at a faster rate with the total number of subscribers compared to that of retail buyers. How should the seller's price offers take this heterogeneity into account?

We consider $N > 1$ buyer groups indexed by $n \in \{1, \ldots, N\}$, each making up a proportion λ_n of the population, with $\sum_n \lambda_n = 1$. A buyer's willingness to pay is increasing in the anticipated quantity q demanded by all buyers, but this externality is stronger on buyers in higher-indexed groups. Specifically, in the linear demand environment with $u(\theta, q) = \theta \overline{v}(q)$, and consistent with Definition 2, we assume that for all $q \in [0, 1]$ and all $n \in \{1, \ldots, N-1\}$: (*i*) $\overline{v}_{n+1}(q) > \overline{v}_n(q)$, and (*ii*) $\overline{v}_{n+1}(q)/\overline{v}_n(q)$ is strictly increasing in q.

The seller's price offers can condition on both a buyer's identity i and the group n to which the buyer belongs (but not on the buyer's private type θ). The seller's problem thus amounts to choosing a price distribution $\Pi_n \in \Delta(\mathbb{R}_+)$ for each buyer group $n \in \{1, \ldots, N\}$, with the objective of maximizing her total worst-case revenue. Given an anticipated total quantity $q \in [0, 1]$ and a price $p \in \mathbb{R}_+$, denote the quantity demanded by group-n buyers by $D_{n,q}(p) := 1 - p/\overline{v}_n(q)$, and let $D_{n,q}(\Pi_n) := \int D_{n,q}(p)\, \mathrm{d}\Pi_n(p)$ and $R_{n,q}(\Pi_n) :=$

[16] The monopoly effect is lower under stronger externalities because marginal buyers are located at a higher price point, so the loss in revenue from their quantity going down is more pronounced. The externality effect is lower because the feedback effects of a lower quantity are larger under stronger externalities.

$\int p D_{n,q}(p)\, \mathrm{d}\Pi_n(p)$ for any $\Pi_n \in \Delta(\mathbb{R}_+)$. Applying the logic of Proposition 2, we can write the seller's problem analogously as we did in program (P^*) for our baseline model:

$$\max_{\{\Pi_n \in \Delta(\mathbb{R}_+)\}_n, q\in[0,1]} \sum_n \lambda_n R_{n,q}(\Pi_n) \tag{P^*_N}$$
$$\text{subject to} \quad \sum_n \lambda_n D_{n,\hat{q}}(\Pi_n) \geq \hat{q} \quad \forall \hat{q} \in (0,q).$$

As in (P^*), the demand constraints in (P^*_N) say that to implement an equilibrium total quantity $q \in [0,1]$, the seller must rule out any lower quantity as an equilibrium. That requires that for each anticipated quantity $\hat{q} \in (0,q)$, the total quantity demanded exceed $\hat{q}$.[17] Now, in this setting, the total quantity demanded is the sum of the demands from each of the N buyer groups. The seller thus makes a choice on how to use the different groups to build the demand up to q. The following definition introduces a class of price distributions that build the demand in an ordered manner. We let $\underline{q}_n := \overline{v}_n^{-1}$.

Definition 3. *Given prices $p_1 < \cdots < p_N$, say price distributions $(\Pi_n)_{n=1}^N$ are **residual greedy up to $(p_n)_{n=1}^N$** if each $n \in \{1,\ldots,N\}$ has*

$$\mathrm{Supp}(\Pi_n) \subseteq [\min\{\overline{v}_n(q_{n-1}), p_n\}, p_n]\,,$$
$$\sum_{m=1}^{n} \lambda_m D_{m,\hat{q}}(\Pi_m) = \hat{q} \quad \forall \hat{q} \in (q_{n-1}, \underline{q}_n(p_n)),$$

where $q_n := \max\{q \in [0,1] : \sum_{m=1}^n \lambda_m D_{m,q}(\Pi_m) = q\}$.

Price distributions $(\Pi_n)_n$ that are residual greedy up to $(p_n)_n$ have two key properties. First, since the quantities $(q_n)_n$ as defined must satisfy $p_n \leq \overline{v}_n(q_n)$ and $q_1 < q_2 < \ldots < q_N$, the supports of the price distributions are ordered.[18] This means that all buyers in group n are offered lower prices than any buyer

[17] As in (P^*), we impose the demand constraints as weak inequalities, with the solution being the limit of a sequence $(\{\Pi_{nk}\}_n, q_k)_k$ that satisfies them strictly for every k.

[18] Observe that there is no circularity in Definition 3 since Π_n depends on $(q_1,\ldots,q_{n-1})$ while q_n depends on Π_n.

in group $n+1$, and the seller uses buyers only from groups $\{1,\dots,n\}$ to satisfy the demand constraints up to anticipated quantity q_n. Second, given the price distributions for groups $\{1,\dots,n-1\}$, the price distribution for group n makes the demand constraints for anticipated quantities $\hat{q} \in (q_{n-1}, \underline{q}_n(p_n))$ hold with equality. Intuitively, the seller follows a greedy procedure as in our main model, but because these demand constraints aggregate the quantity demanded by buyers in all groups $(1,\dots,n)$, the prices are greedy in a residual sense: the seller offers discounts to group-n buyers only as much as is needed to build the residual demand not fulfilled by lower-index-group buyers.

We show that the seller's optimal pricing policy consists of price distributions that are residual greedy up to their highest supported prices.

Proposition 5. *Take the linear demand environment with buyer groups $n \in \{1,\dots,N\}$ indexed by increasing strength of externalities. Suppose $((\Pi_n^*)_{n=1}^N, q^*)$ is optimal, and let p_n^* be the highest price in the support of Π_n^*. Then the price distributions $(\Pi_n^*)_{n=1}^N$ are residual greedy up to $(p_n^*)_{n=1}^N$, and Π_N^* has a mass point at $p_N^* < \overline{v}_N(q^*)$. Therefore, for each $n \in \{1,\dots,N-1\}$,*

$$\max \operatorname{Supp}(\Pi_n^*) < \min \operatorname{Supp}(\Pi_{n+1}^*).$$

This result sheds light on how the seller optimally builds the demand towards an equilibrium total quantity. Buyers from strong-externality groups are more responsive to the anticipated quantity of trade than those from weak-externality groups. Hence, the seller benefits from offering lower prices to weak-externality buyers in order to provide assurance of a higher total quantity to strong-externality buyers; this allows her to extract higher revenue from the latter. Going back to the example of corporate and retail buyers of file sharing services, Proposition 5 says that all retail buyers will enjoy larger discounts than any corporate buyer.

The proposition further shows that the methodology from our main model extends to the setting with heterogeneous buyer groups. Once we establish that the optimal price distributions $(\Pi_n^*)_n$ have ordered supports—more specifically, that any prices $(p_n)_n$ respectively in the supports of $(\Pi_n^*)_n$ have

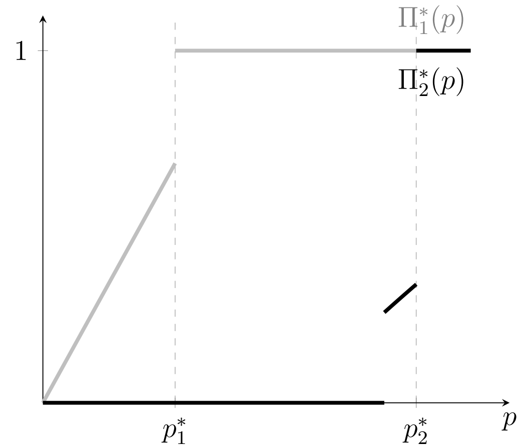

Figure 4: Optimal price distributions for a population with 2 buyer groups with equal weights, where group-1 and group-2 buyers have willingness to pay as in the first and second examples of Figure 2 respectively. For group 1 (light gray line), we obtain $\Pi_1^*(p) = 4p$ for $p < p_1^*$ and $\Pi_1^*(p) = 1$ for $p \geq p_1^*$, with $p_1^* \approx 0.17$ and $q_1 \approx 0.33$. For group 2 (black line), we obtain $\Pi_2^*(p) = 0$ for $p < \overline{v}_2(q_1)$, $\Pi_2^*(p) = (1 + 6p - 2\sqrt{1+4p} + p_1^*(1 - 2p_1^*))/\sqrt{1+4p}$ for $\overline{v}_2(q_1) \leq p < p_2^*$, and $\Pi_2^*(p) = 1$ for $p \geq p_2^*$, where $p_2^* \approx 0.48$ and $q_2 = q^* \approx 0.74$.

$\underline{q}_n(p_n) \leq \underline{q}_{n+1}(p_{n+1})$—then we are able to apply the arguments of Theorem 1 to each of the N buyer groups. This yields the characterization in Proposition 5, with price distributions that are residual greedy up to the top of their supports, and a mass point at p_N^*, as well as possibly mass at other points in $(p_1^*, \dots, p_{N-1}^*)$.[19] The interpretation is intuitive: the seller sets a high posted price together with group-exclusive discounts and personalized discounts within each group.

Figure 4 provides an illustration. We take a population with $N = 2$ buyer groups with equal weights. Group-1 and group-2 buyers have willingness to pay as given respectively in the first and second examples of Figure 2. We can interpret the seller's solution as posting a price of p_2^* and offering all buyers in the weak-externality group 1 a group-exclusive discount of $p_2^* - p_1^*$, in addition

[19] Observe that residual greediness can be consistent with a price distribution Π_n^* also having a mass point at the lower limit of its support. An example is shown in Figure 4.

to offering personalized discounts to some buyers in this group and some buyers in the strong-externality group 2. In this way, the seller builds the demand with group-1 buyers up to a quantity q_1, and then extracts higher revenue from group-2 buyers as she continues to grow the demand with buyers from both groups up to the equilibrium quantity $q^* > q_1$.

6. Discussion

In this section, we describe different variants of our model, discuss how our analysis and results would (or not) change, and offer some concluding remarks.

Complete information. Our seller's problem has two key features: externalities in consumption and unobservable buyer types. In Section 3, we studied a benchmark describing what happens when either the externalities are absent or the strategic uncertainty they generate is not a concern for the seller. That benchmark placed our analysis within the literature on monopoly pricing. We next consider the other natural benchmark for our problem, in which the externalities and the concern for strategic uncertainty are present, but buyer types are observable. This benchmark connects our analysis to the literature on contracting with externalities, which until now had focused on complete-information settings.

Take our baseline model but suppose the seller can make price offers that condition on both a buyer's identity i and his type θ. Since the seller knows exactly how much each buyer (i, θ) is willing to pay for each anticipated total quantity, her problem simplifies significantly. Given the increasing externalities, it is easy to see that the seller will want to ensure that all buyers purchase. This means that all buyers purchasing must be the unique equilibrium, and thus the unique rationalizable outcome. Therefore, the seller's problem reduces to choosing an order in which buyers iteratively delete the no-purchase action as being dominated, together with revenue-maximizing price offers that implement this iterated deletion. The solution prescribes a permutation of buyers, such that each buyer in the permutation is indifferent over purchasing if all buyers preceding him purchase and the rest do not.[20] If $u(\theta, q)$ is supermodular (as in our linear-demand and proportional-values environments), then an optimal permutation orders buyers in increasing type order.[21]

This approach is the same as used in other papers on contracting with externalities. However, this methodology is not available to us in our model with incomplete information. Plainly, the fact that types are unobservable means that the seller cannot control the order in which buyers iteratively delete the no-purchase action. Our analysis develops a new methodology that excludes low-revenue outcomes by working not through the buyer types but through the anticipated quantities of trade. The seller's solution iteratively deletes anticipated quantities as candidates for equilibrium quantities. Observe that, in this solution, the order in which buyers delete the no-purchase action in not necessarily monotonic, neither in their types nor in their price offers.

In addition to requiring a new methodology, our incomplete-information problem yields results that are qualitatively different from those obtained under complete information. As noted, when types are observable, the seller induces the whole population of buyers to purchase. Moreover, except in special cases, no two buyers receive the same price offer.[22] Instead, our model with unobservable types yields exclusion and comparative statics on the total quantity of trade, as well as the result that any optimal price distribution is continuous and strictly increasing up to a mass point at the top. The latter allows us to interpret the seller's solution as a posted price with dispersed discounts, and thus to relate our findings to pricing policies used in applications.

Screening menus. We have phrased our model with the seller choosing personalized price offers. Since buyers have private information about their payoff types, it is natural to ask whether the seller could do better with more

[20] Recall that we have assumed that buyers purchase under indifference.

[21] The intuition for this order is similar to that in Section 5.3: the seller benefits from building the demand with lower types so that she can extract more value from higher types.

[22] This is always true if $u(\theta, q)$ is supermodular. In our Introduction example, where $u(\theta, q) = \theta q$ and types are drawn uniformly from $\Theta = [0, 1]$, the optimal price distribution under complete information is $\Pi^C(p) = \sqrt{p}$ for $p \in [0, 1]$.

sophisticated mechanisms. In this section, we argue that our focus on price offers is without loss of generality within the class of public bilateral contracts.

Let $\mathcal{M}$ denote the set of all compact subsets of $[0,1]\times\mathbb{R}_+$ that contain $(0,0)$. We consider a general contracting environment in which the seller offers a menu $M_i \in \mathcal{M}$ to each buyer $(i,\theta) \in I \times \Theta$, and buyers then simultaneously choose an option from their offered menus. Each menu option specifies a probability of trade $x \in [0,1]$ and a transfer $t \in \mathbb{R}_+$ from the buyer to the seller, with $(0,0)$ corresponding to a buyer's option of not purchasing the good and not making any transfer. Clearly, this is a generalization of our main model, as menus in $\mathcal{M}^P := \{\{(0,0),(1,p)\} : \ p \in \mathbb{R}_+\}$ correspond exactly to price offers.

For any menu $M \in \mathcal{M}$ and willingness to pay $v \in \mathbb{R}_+$, let $(x_M(v), t_M(v))$ be the element of $\arg\max_{(x,t)\in M}(xv - t)$ with highest x. If a buyer anticipates total quantity of trade q and faces menu offer M, his expected quantity demanded is $D_q(M) := \int_0^{\overline{v}(q)} x_M(v) f_q(v) \ \mathrm{d}v$, and the expected revenue he generates is $R_q(M) := \int_0^{\overline{v}(q)} t_M(v) f_q(v) \ \mathrm{d}v$.[23] Analogous to our main model, we can summarize the seller's mechanism choice via a distribution—here, a distribution $\mu \in \Delta\mathcal{M}$ over menu offers. Given such a μ, a total quantity q is an equilibrium quantity if and only if $q = D_q(\mu) := \int D_q(M) \ \mathrm{d}\mu(M)$, and the resulting revenue is $R_q(\mu) := \int R_q(M) \ \mathrm{d}\mu(M)$.

We argue that any menu distribution $\mu \in \Delta\mathcal{M}$ admits some price distribution $\Pi_\mu \in \Delta(\mathbb{R}_+)$ with the same set of equilibrium quantities $q \in [0,1]$ and generating the same revenue for every equilibrium quantity. The idea is simple. First, it follows by standard arguments (Myerson, 1981) that any menu $M \in \mathcal{M}$ can be replaced by a revenue-equivalent random posted price. That is, given M, we can define a distribution Π_M such that a buyer who has willingness to pay $v \in [0, \overline{v}(1)]$ and faces a random posted price with distribution Π_M (and purchases whenever doing so is weakly optimal) would then purchase with probability $x_M(v)$ and generate an expected transfer of $t_M(v)$.[24] Thus, for any $q \in [0,1]$, we obtain $D_q(\Pi_M) = D_q(M)$ and $R_q(\Pi_M) = R_q(M)$.

[23] The assumption that buyers choose the highest-x option among their preferred menu options corresponds to our main model's assumption that buyers purchase when indifferent.

[24] Taking $\nu > 0$, we can let $\Pi_M(v) = x_M(v)$ for $v < \overline{v}(1) + \nu$, and $\Pi_M(v) = 1$ otherwise.

Next, because there is a continuum of buyers, we can take the distribution of prices that the individual random posted prices generate in the population and implement it directly as a distribution of price offers. That is, we can define Π_μ to be the barycenter $\int \Pi_M \,\mathrm{d}\mu(M)$, yielding $D_q(\Pi_\mu) = D_q(\mu)$ and $R_q(\Pi_\mu) = R_q(\mu)$ for every $q \in [0, 1]$.

The implication is that our focus on price offers rather than menu offers is without loss. Instead, what matters for our analysis is our maintained assumption that contracts are bilateral and public. Bilateral contracts means that the contract offered to a buyer cannot directly condition on the purchasing decisions of other buyers. If such multilateral contract offers were feasible, the seller's concern for strategic uncertainty would be mute.[25] Multilateral contracts are difficult to verify and enforce in practice, and for this reason they are commonly ruled out in the contracting-with-externalities literature.[26] Finally, public contracts means that buyers know the realized distribution of prices offered in the population. Whether revenue can be improved in a setting with private contracts—as is the case in the moral-hazard problem of Halac et al. (2021)—is an open question.

Warm start. Our model assumes a cold-start problem: no buyer is willing to purchase at a strictly positive price if he anticipates that no other buyer will purchase. Formally, we assumed that the highest willingness to pay as a function of the anticipated quantity of trade, $\overline{v}(q)$, satisfies $\overline{v}(0) = 0$. In this section, we discuss how our results change if we relax this assumption and consider a "warm-start" model where $\overline{v}(0) > 0$. We maintain all of our other assumptions, including that $\overline{v}$ is continuously differentiable with $\overline{v}' > 0$.

Conceptually, our analysis can be extended to the warm-start model with little modification. Both our restatement of the seller's problem in Proposition 2 and our characterization of the seller's solution in Theorem 1 continue to apply. The latter in particular says that any optimal (Π^*, q^*) has Π^* greedy

[25] The seller would be able to guarantee the best-case outcome by offering each buyer a contract that specifies the best-case price p^B conditional on total quantity $q \geq q^B$, and a zero price otherwise.

[26] See, e.g., Innes and Sexton (1994), Segal (2003), and Halac et al. (2020).

up to its highest supported price p^*, and thus that our definition of greediness remains useful for describing the seller's optimal price distribution. A key difference, however, is that greediness now implies zero mass on prices strictly below $\overline{v}(0)$, so we must have $\Pi^*(p) = 0$ for all $p < \min\{p^*, \overline{v}(0)\}$. Therefore, Theorem 1 in the warm-start model says that Π^* takes one of two forms: either Π^* has no supported prices strictly below $\overline{v}(0)$—in which case it takes the form of a posted price with dispersed discounts, as in our cold-start model—or Π^* is degenerate on $p^* < \overline{v}(0)$—in which case it is simply a posted price.

The intuition for why a degenerate price distribution could be optimal for the seller under warm start can be seen immediately by taking $\overline{v}(0)$ to be high enough. Indeed, observe that by (the application of) Theorem 1, any optimal (Π^*, q^*) has Π^* with highest supported price $p^* \leq p^M(q^*)$, where $p^M(q^*)$ is the monopoly price that obtains under an exogenous demand D_{q^*}. Hence, a sufficient condition for the seller to choose a degenerate price distribution in the warm-start model is $\overline{v}(0) > p^M(1)$. In this case, the externalities in consumption operate in a region of highest values that is above the highest price the seller could ever want to offer. The seller cannot gain from setting prices above $\overline{v}(0)$, and thus she cannot gain from price dispersion.

Low-value externalities. We have assumed that the lowest buyer value is zero for all anticipated quantities; i.e., that F_q has support $[0, \overline{v}(q)]$ for all $q \in [0, 1]$. Suppose instead that the support of F_q is $[\underline{v}(q), \overline{v}(q)]$, where $\underline{v}$ is continuously differentiable with $\underline{v}' \geq 0$. Adapting our concavity assumptions to this more general setting, we can show that our main results in Proposition 2 and Theorem 1 go through essentially unchanged. The proof of Theorem 1 would combine demand-preserving contractions as the ones we use in our baseline model together with some price increases below $\underline{v}(q^*)$, namely in a price range where all buyers are willing to purchase in equilibrium.

Concluding remarks. We have presented a framework for studying personalized pricing in markets with network externalities. Our analysis provides an explanation for the use of posted prices together with discounts that are dispersed across buyers. We showed how the seller's solution is shaped by her concern for strategic uncertainty, and how it changes with the strength of externalities and with heterogeneity across buyer groups.

We believe there are several potentially fruitful directions for future research. For example, one could build on our model to examine the possibility of congestion in consumption; this could be introduced by assuming that buyers' highest-value function $\overline{v}(q)$ is inverse-U-shaped in the anticipated quantity of trade q. Another interesting direction would be to extend our analysis to a two-sided platform, say with sellers on one side and buyers on the other.[27] Unlike in our heterogeneous-groups setting of Section 5.3, here participants on each side would have a value of participating that is increasing in the number of participants on the other side but (weakly) decreasing in the number of participants on their same side.

Finally, another possible direction would be to introduce dynamics. Suppose buyers receive offers from the seller and can hold onto them, so they decide not only whether to purchase but also when to purchase. Taking others' decisions to be independent of his own, assume a buyer purchases at a time $t \geq 0$ if and only if it is dominant for him to do so given the publicly observed quantity of purchases up until time $t-1$.[28] If the seller offers a constant price to each buyer, then her solution coincides with that in our static model. In fact, this dynamic setting offers a transparent dynamic implementation of our seller's solution that requires buyers to know neither the seller's price distribution nor the distribution of other buyers' types. In future work, we are interested in studying the conditions under which this solution remains optimal even when the seller can commit to prices that change over time.

References

Aoyagi, Masaki, "Coordinating Adoption Decisions under Externalities and Incomplete Information," *Games and Economic Behavior*, 2013, *77*, 77–89.

[27] See Jullien, Pavan and Rysman (2023) for a recent survey of the literature on two-sided markets with network effects.

[28] See Hartline, Mirrokni and Sundararajan (2008) for a related model of dynamic pricing.

Bernstein, Shai and Eyal Winter, "Contracting with Heterogeneous Externalities," *American Economic Journal: Microeconomics*, 2012, *4*, 50–76.

Bulow, Jeremy and John Roberts, "The Simple Economics of Optimal Auctions," *Journal of Political Economy*, 1989, *97* (5), 1060–1090.

Camboni, Matteo and Michael Porcellacchia, "Monitoring Team Members: Information Waste and the Transparency Trap," 2024.

Csorba, Gergely, "Screening contracts in the presence of positive network effects," *International Journal of Industrial Organization*, 2008, *26* (1), 213–226.

Cusumano, Carlo M., Tan Gan, and Ferdinand Pieroth, "Misaligning Incentives in Teams," 2023. Working paper.

Eliaz, Kfir and Ran Spiegler, "X-games," *Games and Economic Behavior*, 2015, *89*, 93–100.

Ellison, Glenn and Drew Fudenberg, "The Neo-Luddite's Lament: Excessive Upgrades in the Software Industry," *The RAND Journal of Economics*, 2000, *31* (2), 253–272.

Genicot, Garance and Debraj Ray, "Contracts and externalities: How things fall apart," *Journal of Economic Theory*, 2006, *131* (1), 71–100.

Guesnerie, Roger and Pedro Jara-Moroni, "Expectational Coordination in Simple Economic Contexts," *Economic Theory*, 2011, *47*, 205–246.

Halac, Marina, Elliot Lipnowski, and Daniel Rappoport, "Rank Uncertainty in Organizations," *American Economic Review*, 2021, *111* (3), 757–786.

_ , _ , and _ , "Addressing Strategic Uncertainty with Incentives and Information," *AEA Papers and Proceedings*, 2022, *112*, 431–437.

_ , Ilan Kremer, and Eyal Winter, "Raising Capital from Heterogeneous Investors," *American Economic Review*, 2020, *110* (3), 889–921.

_ , _ , and _ , "Monitoring Teams," *American Economic Journal: Microeconomics*, forthcoming.

Hartline, Jason, Vahab Mirrokni, and Mukund Sundararajan, "Optimal marketing strategies over social networks," in "Proceedings of the 17th international conference on World Wide Web" 2008, pp. 189–198.

Heidhues, Paul and Botond Kőszegi, "Regular Prices and Sales," *Theoretical Economics*, 2014, *9*, 217–251.

Innes, Robert and Richard J. Sexton, "Strategic Buyers and Exclusionary Contracts," *American Economic Review*, 1994, *84* (3), 566–584.

Jullien, Bruno, Alessandro Pavan, and Marc Rysman, "Two-sided Markets, Pricing, and Network Effects," *Handbook of Industrial Organization*, 2023, *4*, 485–592.

Katz, Michael L. and Carl Shapiro, "Network Externalities, Competition, and Compatibility," *The American Economic Review*, 1985, *75* (3), 424–440.

_ and _ , "Technology Adoption in the Presence of Network Externalities," *Journal of Political Economy*, 1986, *94* (4), 822–841.

Milgrom, Paul and John Roberts, "Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities," *Econometrica*, 1990, pp. 1255–1277.

Moriya, Fumitoshi and Takuro Yamashita, "Asymmetric-Information Allocation to Avoid Coordination Failure," *Journal of Economics & Management Strategy*, 2020, *29* (1), 173–186.

Myerson, Roger B., "Optimal Auction Design," *Mathematics of Operations Research*, 1981, *6* (1), 58–73.

Oren, Shmuel S, Stephen A Smith, and Robert B Wilson, "Nonlinear pricing in markets with interdependent demand," *Marketing Science*, 1982, *1* (3), 287–313.

Perry, Motty, "Sustainable positive profit multiple-price strategies in contestable markets," *Journal of Economic Theory*, 1984, *32* (2), 246–265.

Rappoport, Daniel, "Evidence and Skepticism in Verifiable Disclosure Games," 2024. Working Paper.

Rasmusen, Eric B., J. Mark Ramseyer, and John S. Wiley Jr., "Naked Exclusion," *American Economic Review*, 1991, *81*, 1137–45.

Rohlfs, Jeffrey, "A Theory of Interdependent Demand for a Communications Service," *The Bell Journal of Economics and Management Science*, 1974, *5* (1), 16–37.

Sákovics, József and Jakub Steiner, "Who Matters in Coordination Problems?," *American Economic Review*, 2012, *102* (7), 3439–3461.

Segal, Ilya R., "Contracting with Externalities," *Quarterly Journal of Economics*, 1999, *114*, 337–388.

_ , "Coordination and Discrimination in Contracting with Externalities: Divide and Conquer?," *Journal of Economic Theory*, 2003, *113*, 147–81.

_ **and Michael D. Whinston**, "Naked Exclusion: Comment," *American Economic Review*, 2000, *90*, 296–309.

Spiegler, Ran, "Extracting Interaction-Created Surplus," *Games and Economic Behavior*, 2000, *30*, 142–162.

Veiga, André, "A note on how to sell a network good," *International Journal of Industrial Organization*, 2018, *59*, 114–126.

Winter, Eyal, "Incentives and Discrimination," *American Economic Review*, 2004, *94*, 764–773.

A. Preliminaries

Lemma 1. *The map taking anticipated quantity $q \in (0,1]$ to its demand function D_q is continuous (with respect to the supremum norm).*

Proof. Fix $q \in (0,1]$ and consider $\tilde{q} \in (0,1]$. We want to show $D_{\tilde{q}} \xrightarrow{\|\cdot\|_\infty} D_q$ as $\tilde{q} \to q$. To that end, define

$$\beta(\tilde{q}) := \left| \|f_q\|_\infty - \|f_{\tilde{q}}\|_\infty \right| \text{ and } \gamma(\tilde{q}) := \left\| (f_q - f_{\tilde{q}})|_{[0,\ \min\{\overline{v}(q), \overline{v}(\tilde{q})\}]} \right\|_\infty .$$

Now, every $p \in \mathbb{R}_+$ has

$$\begin{aligned} |D_q(p) - D_{\tilde{q}}(p)| &= \left| \left(\int_{[p,\infty)\cap[0,\overline{v}(q)]\cap[0,\overline{v}(\tilde{q})]} + \int_{[p,\infty)\cap \mathrm{co}\{\overline{v}(q),\overline{v}(\tilde{q})\}} \right) (f_q - f_{\tilde{q}}) \right| \\ &\leq \gamma(\tilde{q})\overline{v}(1) + |\overline{v}(q) - \overline{v}(\tilde{q})| \left[\|f_q\|_\infty + \beta(\tilde{q}) \right]. \end{aligned}$$

Because $\overline{v}$ is continuous and the right-hand side of the above inequality is independent of p, the lemma will follow if we establish that $\beta(\tilde{q})$ and $\gamma(\tilde{q})$ both converge to zero as $\tilde{q} \to q$. So given any $\varepsilon > 0$, we want to show $\tilde{q} \in (0,1]$ close enough to q has $\beta(\tilde{q}), \gamma(\tilde{q}) < \varepsilon$.

Let $Q := [\frac{1}{2}q, 1]$, a compact neighborhood of q in $(0,1]$. Because a continuous function on a compact space is uniformly continuous, some $\delta > 0$ exists such that any $q_1, q_2 \in Q$ and $v_1 \in [0, \overline{v}(q_1)]$, $v_2 \in [0, \overline{v}(q_2)]$ such that $|q_1 - q_2|, |v_1 - v_2| < \delta$ have $|f_{q_1}(v_1) - f_{q_2}(v_2)| < \frac{\varepsilon}{2}$. But then, consider any $\tilde{q} \in Q$ with $|\tilde{q} - q|, |\overline{v}(\tilde{q}) - \overline{v}(q)| < \delta$—close enough $\tilde{q}$ satisfies these inequalities because $\overline{v}$ is continuous. Clearly, $\gamma(\tilde{q}) \leq \frac{\varepsilon}{2} < \varepsilon$, so it remains to show $\beta(\tilde{q}) < \varepsilon$.

Take any $\{q_1, q_2\} = \{q, \tilde{q}\}$. Some $v_1 \in [0, \overline{v}(q_1)]$ exists such that $f_{q_1}(v_1) = \|f_{q_1}\|_\infty$. But then, because $v_2 := \min\{v_1, \overline{v}(q_2)\}$ has $|v_1 - v_2| < \delta$, we have

$$\|f_{q_1}\|_\infty = f_{q_1}(v_1) \leq |f_{q_1}(v_1) - f_{q_1}(v_2)| + |f_{q_1}(v_2) - f_{q_2}(v_2)| + f_{q_2}(v_2) < 2\tfrac{\varepsilon}{2} + \|f_{q_2}\|_\infty .$$

Hence, $\beta(\tilde{q}) < \varepsilon$, as required. *Q.E.D.*

Lemma 2. *The functions $\Delta[0, \overline{v}(1)] \times (0,1] \to \mathbb{R}$ taking a pair (Π, q) to its* *quantity demanded $D_q(\Pi)$ and revenue $R_q(\Pi)$ are continuous (when $\Delta[0,\overline{v}(1)]$ is endowed with its weak* topology).*

Proof. Given a continuous function $\psi : [0,\overline{v}(1)] \to \mathbb{R}_+$, we will show $(\Pi, q) \mapsto \int \psi D_q \, \mathrm{d}\Pi$ is continuous. The demand and revenue results will follow by, respectively, setting $\psi(p) := 1$ and $\psi(p) := p$. Now, fix any $(\Pi, q) \in \Delta[0,\overline{v}(1)] \times (0,1]$. Then, any $(\tilde{\Pi}, \tilde{q}) \in \Delta[0,\overline{v}(1)] \times (0,1]$ has

$$\begin{aligned}\left|\int \psi D_q \, \mathrm{d}\Pi - \int \psi D_{\tilde{q}} \, \mathrm{d}\tilde{\Pi}\right| &\leq \left|\int \psi D_q \, \mathrm{d}(\Pi - \tilde{\Pi})\right| + \int \psi \left|D_q - D_{\tilde{q}}\right| \, \mathrm{d}\tilde{\Pi} \\ &\leq \left|\int \psi D_q \, \mathrm{d}(\Pi - \tilde{\Pi})\right| + \|\psi\|_\infty \|D_q - D_{\tilde{q}}\|_\infty,\end{aligned}$$

which converges to zero as $\tilde{\Pi} \to \Pi$ and $\tilde{q} \to q$ by Lemma 1 and because ψD_q is continuous and bounded on $[0,\overline{v}(1)]$. *Q.E.D.*

Lemma 3. *For any given price distribution $\Pi \in \Delta(\mathbb{R}_+)$, the seller's revenue is weakly increasing in the anticipated quantity, strictly so wherever the revenue is strictly positive. Moreover, the induced set of equilibrium quantities is closed and nonempty. Hence, a least-quantity equilibrium exists and is the unique worst-case equilibrium.*

Proof. By Lemma 2, the set of equilibrium quantities in $(0,1]$ is closed in $(0,1]$. To see it is closed in $\mathbb{R}$, then, it remains to show 0 is an equilibrium quantity if it is a limit of them. To that end, suppose some sequence $\{q_n\}_{n=1}^\infty \subset (0,1]$ of equilibrium quantities converges to zero. Because $\Pi(0) \leq D_{q_n}(\Pi) = q_n$ for each $n \in \mathbb{N}$, we have $\Pi(0) \leq \lim_{n\to\infty} q_n = 0$. Thus 0 is an equilibrium quantity in this case, as desired.

Let us establish that an equilibrium exists. If $D_0(\Pi) = 0$, then 0 is an equilibrium quantity, so focus on the case in which $D_0(\Pi) > 0$. For sufficiently small q_0 (e.g., those strictly below $D_0(\Pi)$), we then have $D_{q_0}(\Pi) \geq q_0$. Meanwhile, $D_1(\Pi) \leq 1$. It follows from Lemma 2 and the intermediate value theorem that some equilibrium quantity in $[q_0, 1]$ exists.

Toward the payoff ranking, note an anticipated quantity q generates revenue

$$R_q(\Pi) = \int_0^\infty pD_q(p)\ \mathrm{d}\Pi(p).$$

Because the integrand weakly increases with $D_q(p)$ at every $p \geq 0$, it weakly increases (given monotonicity of u) with q.

Now we pursue the strict revenue ranking. Suppose two quantities $\tilde{q}, q \in [0,1]$ have $\tilde{q} < q$ and $R_q(\Pi) > 0$. We want to show $R_q(\Pi) > R_{\tilde{q}}(\Pi)$. The claim holds if $\tilde{q} = 0$ because then $R_{\tilde{q}}(\Pi) = 0$; so focus on the alternative case. In this case, we can pair Assumption 1 with the fact that (given $\overline{v}' > 0$) the distributions F_q and $F_{\tilde{q}}$ are not identical, to deduce $D_q(p) > D_{\tilde{q}}(p)$ for every $p \in (0, \overline{v}(q))$. That $R_q(\Pi) > 0$ implies Π puts positive mass on such prices then implies $\int_0^\infty pD_q(p)\ \mathrm{d}\Pi(p) > \int_0^\infty pD_{\tilde{q}}(p)\ \mathrm{d}\Pi(p)$, as desired.

Having shown the set of equilibrium quantities is closed in the compact set $[0,1]$, a lowest equilibrium quantity q exists. We also know q is a worst-case equilibrium quantity, uniquely so if $R_q(\Pi) > 0$. Finally, if q is a zero-revenue equilibrium quantity, then our tie-breaking assumption implies $q = \Pi(0)$, and so q is the unique zero-revenue equilibrium quantity. The lemma follows. *Q.E.D.*

B. Proofs for Section 3 and Section 4

B.1. Proof of Proposition 1

First, let us show that a best degenerate-price equilibrium exists and generates strictly positive revenue. To that end, consider the program

$$\begin{aligned} \max_{(p,q)\in[0,\overline{v}(1)]\times[0,1]} \quad & pq \\ \text{s.t.} \quad & q\,[D_q(p) - q] = 0. \end{aligned}$$

First, let us observe the program admits some optimal solution (p^B, q^B) with strictly positive value. Indeed, notice the constraint function $(p,q) \mapsto q\,[D_q(p) - q]$ is continuous wherever the quantity is strictly positive by Lemma 2, and it is continuous at zero quantity because $(p, q) \mapsto D_q(p)$ is bounded. Therefore, the program has continuous objective on a compact domain and so admits an optimal solution (p^B, q^B). Moreover, because $(D_q^{-1}(q), q)$ is feasible and yields strictly positive value in the program for $q \in (0, 1)$, it follows that $p^B q^B > 0$.

Let us now see (p^B, q^B) is a best degenerate-price equilibrium. First, because $q^B > 0$, we know $D_{q^B}(p^B) - q^B = 0$, so q^B is an equilibrium quantity for the degenerate price distribution on p^B. Next, any alternative degenerate-price equilibrium (p, q) would either have $p > \overline{v}(1)$ and hence generate zero revenue, or would be feasible in the above program and so generate a weakly lower revenue.

It remains to show any nondegenerate price distribution $\Pi \in \Delta(\mathbb{R}_+)$, with any equilibrium quantity q it generates, does strictly worse than some degenerate-price equilibrium. If q is zero (and so too is revenue), then we have nothing to show because we have already shown a degenerate-price equilibrium can yield strictly positive revenue. So focus on the case of $q \in (0, 1]$. In this case, some uniform price—specifically $p = D_q^{-1}(q) \in [0, \overline{v}(q)]$—exists for which q is an equilibrium quantity. Moreover, because $\varphi_{q,q}$ is strictly increasing (given Assumption 3), the degenerate price yields a strictly higher revenue. *Q.E.D.*

B.2. Proof of Proposition 2

Toward showing this program's solutions are exactly the optimal pairs (Π^*, q^*), let us invest in some terminology. Say a pair $(\Pi, q) \in \Delta(\mathbb{R}_+) \times [0, 1]$ is *worst-feasible* if q is a worst equilibrium for the principal given price distribution Π. Say a pair $(\Pi^*, q^*) \in \Delta(\mathbb{R}_+) \times [0, 1]$ is *limit-worst-feasible (LWF)* if it is a limit of a sequence of worst-feasible pairs. Finally, let $R^* := \sup_{(\Pi,q) \text{ worst-feasible}} R_q(\Pi)$ denote the principal's optimal value.

Let us make three preliminary observations. First, any convergent sequence $(\Pi_n, q_n)_{n=1}^{\infty}$ of worst-feasible pairs has

$$\lim_{n\to\infty} R_{q_n}(\Pi_n) = R_{q_\infty}(\Pi_\infty), \text{ where } (\Pi_\infty, q_\infty) = \lim_{n\to\infty}(\Pi_n, q_n). \tag{8}$$

This result follows immediately from Lemma 2 for the case of $q_\infty > 0$, and in the case of $q_\infty = 0$ it follows from $0 \le R_{q_n}(\Pi_n) \le \overline{v}(1) D_{q_n}(\Pi_n) = \overline{v}(1) q_n \to 0$. Second, observe that some LWF pair $(\hat{\Pi}, \hat{q})$ has $R_{\hat{q}}(\hat{\Pi}) = R^*$. Indeed, to find such a pair, take some sequence of worst-feasible pairs $(q_n, \Pi_n)_{n=1}^\infty$ with $\lim_{n\to\infty} R_{q_n}(\Pi_n) = R^*$, which exists by definition of R^*. Because all prices in $[\overline{v}(1), \infty)$ yield the same revenue (zero), we can assume without loss that $\Pi_n \in \Delta[0, \overline{v}(1)]$. Then, by compactness, we can (dropping to a subsequence) assume without loss that $(\Pi_n, q_n)_{n=1}^\infty$ converges to some $(\hat{\Pi}, \hat{q})$—which is then as desired by (8). Third, some worst-feasible pair (hence some LWF pair) generates strictly positive revenue. Indeed, given $p \in (0, \overline{v}_{1/2})$, Lemma 3 implies one can pair $\Pi = \frac{1}{2}\mathbf{1}_{[0,\infty)} + \frac{1}{2}\mathbf{1}_{[p,\infty)}$ with its (strictly positive) lowest equilibrium quantity.

Let us now establish, given $(\Pi^*, q^*) \in [0,1] \times \Delta(\mathbb{R}_+)$, a four-way equivalence:

(i) The pair (Π^*, q^*) is optimal, in the sense defined in the main text.

(ii) The pair (Π^*, q^*) is LWF and has $R_{q^*}(\Pi^*) = R^*$.

(iii) The pair (Π^*, q^*) solves the program $\max_{(\Pi,q) \text{ LWF}} R_q(\Pi)$.

(iv) The pair (Π^*, q^*) solves program (P*).

Because we have noted above that some LWF pair $(\hat{\Pi}, \hat{q})$ has $R_{\hat{q}}(\hat{\Pi}) = R^*$, and because we have noted that some LWF pair generates strictly positive revenue, proving this four-way equivalence will prove the proposition. We will prove that (i) $\iff$ (ii) $\iff$ (iii) $\iff$ (iv). First, note that (i) $\iff$ (ii) follows immediately from (8).

Now, let us see that (ii) $\iff$ (iii). Recall that some some LWF pair $(\hat{q}, \hat{\Pi})$ has $R_{\hat{q}}(\hat{\Pi}) = R^*$. This equivalence will therefore follow if every LWF (Π, q) has $R_q(\Pi) \le R^*$. And indeed, taking some sequence $(\Pi_n, q_n)_{n=1}^\infty$ of worst-feasible pairs converging to it, every n has $R_{q_n}(\Pi_n) \le R^*$ by definition of R^*—but then $R_q(\Pi) \le R^*$ by (8).

Finally, toward showing (iii) $\iff$ (iv), note that the two programs have the same objective, but different constraint sets. We will first show that any LWF (Π, q) satisfies the constraints of program (P*). Then, we will show that any (Π, q) satisfying the constraints of program (P*) is either LWF or admits an alternative LWF pair $(\tilde{\Pi}, \tilde{q})$ that generates strictly higher revenue. This will imply the equivalence.

Take any LWF (Π, q), as witnessed by $(\Pi_n, q_n)_n$. Toward showing (Π, q) satisfies the constraints of program (P*), suppose $\hat{q} \in (0, q)$. For sufficiently large n, we have $q_n > \hat{q}$; let us argue that $D_{\hat{q}}(\Pi_n) > \hat{q}$ for such n, which will then imply $D_{\hat{q}}(\Pi) \geq \hat{q}$ by Lemma 2. To show it, assume $D_{\hat{q}}(\Pi_n) \leq \hat{q}$ for a contradiction. Now, a worst equilibrium for price distribution Π_n is $q_n > \hat{q} > 0$. Hence, either no zero-quantity equilibrium exists, or the equilibrium quantity $q_n > 0$ also generates zero revenue: in either case, $\Pi_n(0) > 0$. But then, some small enough $\tilde{q}_n \in (0, q_n)$—for instance, any one below $\Pi_n(0)$—has $D_{\tilde{q}}(\Pi_n) > \tilde{q}_n$, so that the intermediate value theorem delivers (given Lemma 2) some equilibrium quantity in $[\tilde{q}_n, q_n]$ for Π_n, a contradiction.

Finally, consider any (Π, q) satisfying the constraints of program (P*). We want to show either that (Π, q) is LWF or that an alternative LWF pair $(\tilde{\Pi}, \tilde{q})$ is a LWF generating strictly higher revenue. Because we know some LWF pair generates strictly positive revenue, the conclusion follows immediately if $R_q(\Pi) = 0$; so focus on the case of $R_q(\Pi) > 0$ from now on. Now, for any $\varepsilon \in (0, 1)$, define the price distribution $\Pi_\varepsilon := (1 - \varepsilon)\Pi + \varepsilon \mathbf{1}_{[0,\infty)}$. Then, every quantity $\hat{q} \in (0, q)$ has $D_{\hat{q}}(\Pi_\varepsilon) = (1 - \varepsilon)D_{\hat{q}}(\Pi) + \varepsilon \geq (1 - \varepsilon)\hat{q} + \varepsilon > \hat{q}$, and $D_0(\Pi_\varepsilon) \geq \varepsilon > 0$. In particular, the worst equilibrium for price distribution Π_ε is at least q. Now, considering some sequence $(\varepsilon_n)_n$ from $(0, 1)$ converging to zero, the sequence $(\Pi_{\varepsilon_n})_n$ of price distributions converges to Π, and has the property that the worst equilibrium quantity q_n for each price distribution Π_{ε_n} has $q_n \geq q$. Dropping to a subsequence if necessary, we may without loss assume q_n converges to some $\tilde{q} \in [q, 1]$ as $n \to \infty$. By construction, the pair $(\Pi, \tilde{q})$ is a LWF. If $\tilde{q} = q$ then (Π, q) is a LWF, and if $\tilde{q} > q$ then the LWF $(\Pi, \tilde{q})$ generates strictly higher revenue than (Π, q) by Lemma 3. The proposition follows. *Q.E.D.*

B.3. Inputs for the proof of Theorem 1

The following lemma records a useful technical result that generalizes Proposition 4 of Rappoport (2024).

Lemma 4. *Suppose* $[\underline{v}, \overline{v}] \subset \mathbb{R}$ *is a nondegenerate interval;* $f, g : [\underline{v}, \overline{v}] \to \mathbb{R}$ *are absolutely integrable functions; and* $\psi : [\underline{v}, \overline{v}] \to \mathbb{R}$ *is a function of bounded variation.*[29]

(i) *Suppose* g *is zero wherever* f *is zero on* $[\underline{v}, \overline{v}]$*, and the ratio* $\frac{g}{f}$ *is weakly increasing where its denominator is nonzero. If* $\int_{\underline{v}}^{v} \psi f \geq 0$ *for every* $v \in [\underline{v}, \overline{v}]$*, with equality at* $v = \overline{v}$*, then* $\int_{\underline{v}}^{\overline{v}} \psi g \leq 0$*.*

(ii) *Suppose* g *is zero wherever* f *is zero on* $[\underline{v}, \overline{v}]$*, and the ratio* $\frac{g}{f}$ *is weakly increasing where its denominator is nonzero. If* $\int_{\underline{v}}^{v} \psi f \geq 0$ *for every* $v \in [\underline{v}, \overline{v}]$*, with equality at* $v = \overline{v}$*, and some* $v \in [\underline{v}, \overline{v}]$ *exists such that* $\int_{\underline{v}}^{v} \psi f > 0$ *and* $\frac{g}{f}$ *is not constant on any neighborhood of* v*, then* $\int_{\underline{v}}^{\overline{v}} \psi g < 0$*.*

(iii) *Suppose* f *is zero wherever* g *is zero on* $[\underline{v}, \overline{v}]$*, and the ratio* $\frac{f}{g}$ *is weakly decreasing where its denominator is nonzero. If* $f(\overline{v}), g(\overline{v}) \geq 0$ *and* $\int_{\underline{v}}^{v} \psi g \geq 0$ *for every* $v \in [\underline{v}, \overline{v}]$*, then* $\int_{\underline{v}}^{\overline{v}} \psi f \geq 0$*.*

Proof. First, we prove parts (i) and (ii). To that end, suppose the hypotheses of part (i) are satisfied. In what follows, we interpret $\frac{g}{f}$ as an arbitrary nondecreasing function $[\underline{v}, \overline{v}] \to \mathbb{R}$ that agrees with $\frac{g}{f}$ wherever f is nonzero. Now, define the absolutely continuous function $\Psi : [\underline{v}, \overline{v}] \to \mathbb{R}$ by letting $\Psi(v) := \int_{\underline{v}}^{v} \psi f$. Then,

$$
\begin{aligned}
\int_{\underline{v}}^{\overline{v}} \psi g &= \int_{\underline{v}}^{\overline{v}} \tfrac{g}{f} \Psi' \\
&= \left[\Psi \tfrac{g}{f}\right]_{\underline{v}}^{\overline{v}} - \int_{\underline{v}}^{\overline{v}} \Psi \, \mathrm{d}\tfrac{g}{f} \text{ (by integration by parts)} \\
&= 0 - \int_{\underline{v}}^{\overline{v}} \Psi \, \mathrm{d}\tfrac{g}{f} \text{ (since } \Psi(\underline{v}) = \Psi(\overline{v}) = 0\text{)} \\
&\leq 0 \text{ (since } \Psi \geq 0 \text{ and } \tfrac{g}{f} \text{ is weakly increasing),}
\end{aligned}
$$

establishing part (i). Now, suppose in addition that some $v \in [\underline{v}, \overline{v}]$ exists such that $\Psi(v) > 0$ and $\frac{g}{f}$ is not constant on any neighborhood of v. By continuity,

[29] Throughout, for any interval $[\underline{v}, \overline{v}] \subseteq \mathbb{R}$ and any Lebesgue integrable function $h : [\underline{v}, \overline{v}] \to \mathbb{R}$, we let $\int_{\underline{v}}^{\overline{v}} h$ denote the Lebesgue integral $\int_{\underline{v}}^{\overline{v}} h(v) \, \mathrm{d}v$.

Ψ is strictly positive on some some nondegenerate interval of v. Because $\frac{g}{f}$ is not constant on this interval, it follows that $\int_{\underline{v}}^{\overline{v}} \psi g = -\int_{\underline{v}}^{\overline{v}} \Psi \, \mathrm{d}\frac{g}{f} < 0$, delivering (ii).

Next, we prove part (iii); suppose its hypotheses are satisfied. In what follows, we interpret $\frac{f}{g}$ as an arbitrary nonincreasing function $[\underline{v}, \overline{v}] \to \mathbb{R}$ that agrees with $\frac{f}{g}$ wherever g is nonzero.

Now, define the absolutely continuous functions $\Phi : [\underline{v}, \overline{v}] \to \mathbb{R}$ by letting $\Phi(v) := \int_{\underline{v}}^{v} \psi g$. Then,

$$\begin{aligned}
\int_{\underline{v}}^{\overline{v}} \psi f &= \int_{\underline{v}}^{\overline{v}} \tfrac{f}{g} \Phi' \\
&= \left[\Phi \tfrac{f}{g}\right]_{\underline{v}}^{\overline{v}} - \int_{\underline{v}}^{\overline{v}} \Phi \, \mathrm{d}\tfrac{f}{g} \text{ (by integration by parts)} \\
&= \Phi(\overline{v}) \tfrac{f(\overline{v})}{g(\overline{v})} - \int_{\underline{v}}^{\overline{v}} \Phi \, \mathrm{d}\tfrac{f}{g} \text{ (since } \Phi(\underline{v}) = 0) \\
&\geq \Phi(\overline{v}) \tfrac{f(\overline{v})}{g(\overline{v})} \text{ (since } \Phi \geq 0 \text{ and } \tfrac{f}{g} \text{ is weakly decreasing)} \\
&\geq 0 \text{ (since } \Phi(\overline{v}), f(\overline{v}), g(\overline{v}) \geq 0),
\end{aligned}$$

as required. *Q.E.D.*

The following lemma is a comparative statics result for comparing different price distributions: if a reduction in price dispersion preserves aggregate demand under a low anticipated quantity (and the only modified prices are those that will sometimes be exercised), then the reduction increases both demand and revenue when the anticipated quantity is higher.

Lemma 5. *Given $q \in (0, 1]$, suppose distinct price distributions $\Pi, \tilde{\Pi} \in \Delta(\mathbb{R}_+)$ are such that $\Pi|_{(\overline{v}(q),\infty)} = \tilde{\Pi}|_{(\overline{v}(q),\infty)}$, and*

$$\int_0^v (\Pi - \tilde{\Pi}) f_q \geq 0$$

for every $v \in [0, \overline{v}(q)]$*, with equality at* $v = \overline{v}(q)$.[30] *Then, any* $\tilde{q} \in [q, 1]$ *has*

$$D_{\tilde{q}}(\tilde{\Pi}) \geq D_{\tilde{q}}(\Pi) \text{ and } R_{\tilde{q}}(\tilde{\Pi}) > R_{\tilde{q}}(\Pi).$$

Proof. Both rankings can be derived as applications of Lemma 4, with $(\underline{v}, \overline{v}, f, \psi) = (0, \overline{v}_q, f_q, \Pi - \tilde{\Pi})$ and different choices of g.

First, consider $g := f_{\tilde{q}}|_{[0,\overline{v}(q)]}$, and apply Assumption 1. By Lemma 4(i),[31]

$$\begin{aligned} 0 &\leq \int_0^{\overline{v}(q)} (\tilde{\Pi} - \Pi) f_{\tilde{q}} \\ &= \int_0^{\overline{v}(\tilde{q})} (\tilde{\Pi} - \Pi) f_{\tilde{q}} \\ &= D_{\tilde{q}}(\tilde{\Pi}) - D_{\tilde{q}}(\Pi). \end{aligned}$$

Next, consider $g := \varphi_{q,\tilde{q}} f_q$. As Assumption 3 holds, Lemma 4(ii) tells us

$$\begin{aligned} 0 &< \int_0^{\overline{v}(q)} (\tilde{\Pi} - \Pi) \varphi_{q,\tilde{q}} f_q \\ &= \int_0^{\overline{v}(\tilde{q})} (\tilde{\Pi} - \Pi) \varphi_{\tilde{q},\tilde{q}} \, \mathrm{d}F_{\tilde{q}} \\ &= \int_0^{\overline{v}(\tilde{q})} (\Pi - \tilde{\Pi}) \, \mathrm{d}R_{\tilde{q}} \\ &= 0 - \int R_{\tilde{q}} \, \mathrm{d}(\Pi - \tilde{\Pi}) \\ &= R_{\tilde{q}}(\tilde{\Pi}) - R_{\tilde{q}}(\Pi). \end{aligned}$$

Q.E.D.

The following lemma extends our concave externalities assumption to price distributions rather than just prices.

Lemma 6. *Suppose* $\Pi \in \Delta(\mathbb{R}_+)$ *and* $0 \leq q_0 < q_1 \leq 1$ *have* $\Pi(\overline{v}(q_0)) = \Pi(\overline{v}(q_1)^-)$*. Then* $q \mapsto D_q(\Pi)$ *is concave on* $[q_0, q_1]$*, strictly so if* $\Pi(\overline{v}(q_0)) > 0$.

[30] Note, the equality at $\overline{v}(q)$ says exactly that $D_q(\tilde{\Pi}) = D_q(\Pi)$.

[31] One can alternatively prove this ranking by using the fact that $F_{\tilde{q}} \circ F_q^{-1}$ is convex under Assumption 1.

Proof. For any price $p \leq \overline{v}(q_0)$, the function $q \mapsto D_q(p)$ is strictly concave on (q_0, q_1) by Assumption 2, hence on $[q_0, q_1]$ by Lemma 2. For any price $p \geq \overline{v}(q_1)$, the function $q \mapsto D_q(p)$ is zero on $[q_0, q_1]$. Because a pointwise weighted average of concave functions is concave, strictly so if this average puts strictly positive weight on strictly concave functions, the lemma follows. *Q.E.D.*

To state the next lemma, we invest in some notation.

Notation 1.

- *Let $\dot{f}_q(v)$ denote the partial derivative of $f_q(v)$ with respect to q, which exists wherever $q \in (0,1]$ and $0 \leq v \leq \overline{v}(q)$.*
- *Let $\partial D_q(\Pi)$ [resp. $\partial^- D_q(\Pi)$ or $\partial^+ D_q(\Pi)$] denote the partial derivative [resp. left derivative or right derivative] of $D_q(\Pi)$ with respect to q, if it exists.*

The following lemma establishes that one-sided derivatives of demand with respect to anticipated quantity are finite, and that the demand function is kinked if and only if the price distribution has a mass point.

Lemma 7. *Suppose $\Gamma : [0, \overline{v}(1)) \to \mathbb{R}_+$ is increasing and right continuous, and $q \in (0,1]$. Then:*

- $\partial^- D_q(\Gamma) = \int_0^{\overline{v}(q)} \Gamma \dot{f}_q + \overline{v}'(q)\Gamma\left(\overline{v}(q)^-\right) f_q\left(\overline{v}(q)\right) \in \mathbb{R}$.
- *If $q < 1$, then $\partial^+ D_q(\Pi) = \int_0^{\overline{v}(q)} \Gamma \dot{f}_q + \overline{v}'(q)\Gamma\left(\overline{v}(q)\right) f_q\left(\overline{v}(q)\right) \in \mathbb{R}$.*
- *If Γ is continuous at $\overline{v}(q)$, then $\tilde{q} \mapsto D_{\tilde{q}}(\Gamma)$ is differentiable at q.*
- *If $q < 1$ and Γ is discontinuous at $\overline{v}(q)$, then $\tilde{q} \mapsto D_{\tilde{q}}(\Gamma)$ has a convex kink at q.*

Proof. Whenever $0 \leq q_0 < q_1 \leq 1$, we have

$$\begin{aligned} \tfrac{D_{q_1}(\Gamma) - D_{q_0}(\Gamma)}{q_1 - q_0} &= \tfrac{1}{q_1 - q_0}\left[\int_0^{\overline{v}(q_1)} \Gamma f_{q_1} - \int_0^{\overline{v}(q_0)} \Gamma f_{q_0}\right] \\ &= \int_0^{\overline{v}(q_0)} \Gamma \tfrac{f_{q_1} - f_{q_0}}{q_1 - q_0} + \tfrac{\overline{v}(q_1) - \overline{v}(q_0)}{q_1 - q_0} \tfrac{1}{\overline{v}(q_1) - \overline{v}(q_0)} \int_{\overline{v}(q_0)}^{\overline{v}(q_1)} \Gamma f_{q_1}. \end{aligned}$$

Given the Lebesgue dominated convergence theorem, the first two points come from applying this expression as $q_0 \nearrow q = q_1$ and as $q_1 \searrow q = q_0$, respectively. Then, combine the first two points for $q \in (0,1)$ to obtain

$$\partial^+ D_q(\Gamma) - \partial^- D_q(\Gamma) = \overline{v}'(q) f_q(\overline{v}(q)) \left[\Gamma\left(\overline{v}(q)\right) - \Gamma\left(\overline{v}(q)^-\right)\right],$$

directly implying the last two points. *Q.E.D.*

B.4. Proof of Theorem 1

We begin with some useful terminology.

Definition 4. *Consider any price distribution* Π. *Given* $q \in [0,1]$*:*

- *Say* Π ***has mass at*** q^{++} *if* $\Pi(p) > \Pi\left(\overline{v}(q)\right)$ *for every* $p > \overline{v}(q)$.
- *Say* Π ***has mass at*** q^{--} *if* $\Pi(p) < \Pi\left(\overline{v}(q)^-\right)$ *for every* $p < \overline{v}(q)$.
- *Say* Π ***has mass at*** q^+ *[resp.* q^-*] if it has a mass at* q^{++} *[resp. at* q^{--}*] or has a mass point at* $\overline{v}(q)$.

Given $q_0, q_1 \in [0,1]$ *with* $q_0 < q_1$, *say* Π *is* ***degenerate on*** $[q_0, q_1]$ *if some* $p \in [\overline{v}(q_0), \overline{v}(q_1)]$ *exists such that* $\Pi\left(p^-\right) = \Pi\left(\overline{v}(q_0)^-\right)$ *and* $\Pi\left(p\right) = \Pi\left(\overline{v}(q_1)\right)$.

The following claim shows any optimal price distribution in the subproblem associated with any targeted quantity uses only prices below the monopoly price for that anticipated quantity's demand curve.

Claim 1. *Suppose* $\Pi \in \Delta(\mathbb{R}_+)$ *and* $\hat{q} \in (0,1]$ *have* $\Pi\left(p^M(\hat{q})\right) < 1$. *Then, some* $\tilde{\Pi} \in \Delta(\mathbb{R}_+)$ *exists such that* $D_q(\tilde{\Pi}) \geq D_q(\Pi)$ *for every* $q \in [0,1]$, *and* $R_{\hat{q}}(\tilde{\Pi}) > R_{\hat{q}}(\Pi)$.

Proof. Let $p^* := p^M(\hat{q})$, and let $\tilde{\Pi} := \Pi|_{[0,p^*)} \cup \mathbf{1}|_{[p^*,\infty)}$. The distribution $\tilde{\Pi}$ is below Π in the sense of first-order stochastic dominance, so that $D_q(\tilde{\Pi}) \geq D_q(\Pi)$ for every $q \in [0,1]$. Moreover, Assumption 3 implies any price $p \neq p^*$ has $R_{\hat{q}}(p) < R_{\hat{q}}(p^*)$. Therefore, given that $\Pi(p^*) < 1$, we have

$$R_{\hat{q}}(\tilde{\Pi}) - R_{\hat{q}}(\Pi) = \int_{p^*}^{\infty} \left[R(p^*) - R(p)\right] \, \mathrm{d}\Pi(p) > 0,$$

as desired. *Q.E.D.*

The following claim uses concave externalities to establish that the slack on the demand constraint is first-order wherever the price distribution has a gap at the edge of a slack region.

Claim 2. *Suppose* $\Pi \in \Delta(\mathbb{R}_+)$ *and* $q \in [0,1]$ *have* $D_q(\Pi) = q$*.*

- *If* $q < 1$*, every* $\tilde{q} > q$ *close enough to* q *has* $D_{\tilde{q}}(\Pi) > \tilde{q}$*, and* Π *has no mass at* q^{++}*, then* $\partial^+ D_q(\Pi) > 1$*.*
- *If* $q > 0$*, every* $\tilde{q} < q$ *close enough to* q *has* $D_{\tilde{q}}(\Pi) > \tilde{q}$*, and* Π *has no mass at* q^{--}*, then* $\partial^- D_q(\Pi) < 1$*.*

Proof. Define the function $\psi : [0,1] \to \mathbb{R}$ via $\psi(\tilde{q}) := D_{\tilde{q}}(\Pi) - \tilde{q}$, which is continuous by Lemma 2. By Lemma 6, we know ψ is concave in an interval to the right [left] of q if $q < 1$ [resp. $q > 0$] and Π has no mass at q^{++} [resp. q^{--}].

Now, if ψ is zero at q and concave and strictly positive in a right [resp. left] neighborhood of q, it follows that its right [resp. left] derivative at q is strictly positive [resp. strictly negative], delivering the claim. *Q.E.D.*

The following claim shows that a feasible price distribution is always nondegenerate over (the closure of) any slack region in the range of its support.

Claim 3. *Suppose* $\Pi \in \Delta(\mathbb{R}_+)$ *and* $p^* := \max \operatorname{Supp} \Pi$ *has* $D_q(\Pi) \geq q$ *for every* $q \in (0, \hat{q})$*, for some* $\hat{q} \in (0,1]$ *with* $\overline{v}(\hat{q}) > p^*$*. If* (q_0, q_1) *is a connected component of*

$$\left\{ q \in \big(0, \underline{q}(p^*)\big) : \; D_q(\Pi) > q \right\},$$

then Π *is nondegenerate on* $[q_0, q_1]$*.*

Proof. The claim holds vacuously if $p^* = 0$, so focus on the case in which $p^* > 0$.

If Π has mass at q_0^{++} or at q_1^{--}, it is clearly nondegenerate on $[q_0, q_1]$. So now, focus on the case in which Π has mass neither at q_0^{++} nor at q_1^{--}. The claim will now follow if we establish that Π has mass points both at $\overline{v}(q_0)$ and at $\overline{v}(q_1)$.

Observe first that $\min \operatorname{Supp} \Pi = 0$, for otherwise small enough $q \in \big(0, \underline{q}(p^*)\big)$ will have $D_q(\Pi) = 0 < q$. Then, by definition of the support (and the hypothesis that Π has mass neither at q_0^{++} nor at q_1^{--}), we know that Π has a mass point at $0 = \overline{v}(q_0)$ if $q_0 = 0$, and has a mass point at $p^* = \overline{v}(q_1)$ if $q_1 = \underline{q}(p^*)$.

It remains now to show that Π has a mass point at $\overline{v}(q_0)$ if $q_0 > 0$, and has a mass point at $\overline{v}(q_1)$ if $q_1 < \underline{q}(p^*)$. So suppose $q_0 > 0$ [resp. $q_1 < \underline{q}(p^*)$]. By definition of (q_0, q_1), no $\tilde{q}_0 < q_0$ [resp. $\tilde{q}_1 > q_1$] exists such that every $q \in [\tilde{q}_0, q_0]$ [resp. every $q \in [q_1, \tilde{q}_1]$] has $D_q(\Pi) > q$. But then, by Lemma 2—which applies given that $D_q \geq q$ for every q in a neighborhood of $(0, \underline{q}(p^*)]$—we in fact have that $D_{q_0}(\Pi) = q_0$ [resp. $D_{q_1}(\Pi) = q_1$]. Claim 2 thus implies $\partial^+ D_{q_0}(\Pi) > 1$ [resp. $\partial^- D_{q_1}(\Pi) < 1$]. Meanwhile, that $q \mapsto D_q(\Pi) - q$ is zero at q_0 [resp. q_1] and nonnegative just to the left [resp. right] of it implies $\partial^- D_{q_0}(\Pi) \leq 1$ [resp. $\partial^+ D_{q_1}(\Pi) \geq 1$]. Thus, $q \mapsto D_q(\Pi)$ has a convex kink at q_0 [resp. q_1], and so Lemma 7 tells us Π has a mass point at $\overline{v}(q_0)$ [resp. $\overline{v}(q_1)$] as desired. *Q.E.D.*

The following claim says that whenever the price distribution is nondegenerate over some interval, a smaller such interval can be found on which the price distribution is also well-behaved.

Claim 4. *Suppose $\Pi \in \Delta(\mathbb{R}_+)$ and $0 \leq q_0 < q_1 \leq 1$ are such that Π is nondegenerate on $[q_0, q_1]$. Then some $\tilde{q}_0, \tilde{q}_1 \in [q_0, q_1]$ with $\tilde{q}_0 < \tilde{q}_1$ such that:*

- *Π is nondegenerate on $[\tilde{q}_0, \tilde{q}_1]$;*
- *either $\tilde{q}_0 \in (q_0, q_1)$ or Π has no mass at q_0^{++};*
- *either $\tilde{q}_1 \in (q_0, q_1)$ or Π has no mass at q_1^{--}.*

Proof. If Π has mass at q_1^{--}, then any $\tilde{q}_0 \in (q_0, q_1)$, paired with any $\tilde{q}_1 \in (\tilde{q}_0, q_1)$ close enough to q_1, is as desired. If Π has mass at q_0^{++}, then any $\tilde{q}_1 \in (q_0, q_1)$, paired with any $\tilde{q}_0 \in (q_0, \tilde{q}_1)$ close enough to q_0, is as desired. If Π has no mass at q_1^{--} or at q_0^{++}, then $\tilde{q}_0 = q_0$ and $\tilde{q}_1 = q_1$ are as desired. *Q.E.D.*

The following claim shows that small enough perturbations preserve the demand constraint on any interval where it is slack (with first-order slack at the edges).

Claim 5. *Suppose* $\Pi, \tilde{\Pi} \in \Delta(\mathbb{R}_+)$ *and* $0 \le q_0 < q_1 \le 1$ *are such that:*

- *Every* $q \in (q_0, q_1)$ *has* $D_q(\Pi) > q$, *and each* $q \in \{q_0, q_1\}$ *has* $D_q(\tilde{\Pi}) \ge q$;
- *Either* $D_{q_0}(\Pi) > q_0$ *or* $\partial^+ D_{q_0}(\Pi) > 1$, *with the former case if* $q_0 = 0$; *and*
- *Either* $D_{q_1}(\Pi) > q_1$ *or* $\partial^- D_{q_1}(\Pi) < 1$.

Then, letting $\Pi_\varepsilon := (1-\varepsilon)\Pi + \varepsilon\tilde{\Pi}$, *any small enough* $\varepsilon \in (0,1)$ *has*

$$D_q(\Pi_\varepsilon) \ge q, \ \forall q \in [q_0, q_1].$$

Proof. First, for either $q \in \{q_0, q_1\}$, if $D_q(\Pi) > q$, then (given that $D_q(\tilde{\Pi}) \ge q$) every $\varepsilon \in (0,1)$ has $D_q(\Pi_\varepsilon) > q$. Next, if either $q \in \{q_0, q_1\}$ has $D_q(\Pi) = q$ (which in particular means $q > 0$ given our hypotheses), then Lemma 7 tells us one-sided derivatives of $\tilde{q} \mapsto D_{\tilde{q}}(\tilde{\Pi})$ are finite there. So, for small enough $\overline{\varepsilon} \in (0,1)$:

- Each $q \in \{q_0, q_1\}$ has $D_q(\Pi_{\overline{\varepsilon}}) \ge q$;
- Either $D_{q_0}(\Pi_{\overline{\varepsilon}}) > q_0$ or $\partial^+ D_{q_0}(\Pi_{\overline{\varepsilon}}) > 1$; and
- Either $D_{q_1}(\Pi_{\overline{\varepsilon}}) > q_1$ or $\partial^- D_{q_1}(\Pi_{\overline{\varepsilon}}) < 1$.

Fixing such an $\overline{\varepsilon}$, some $\tilde{q}_0, \tilde{q}_1 \in (q_0, q_1)$ exist such that every $q \in (q_0, \tilde{q}_0] \cup [\tilde{q}_1, q_1)$ has $D_q(\Pi_{\overline{\varepsilon}}) > q$. Hence, because $\varepsilon \mapsto D_q(\Pi_\varepsilon)$ is affine for every q, it follows that every $q \in (q_0, \tilde{q}_0] \cup [\tilde{q}_1, q_1)$ and $\varepsilon \in (0, \overline{\varepsilon}]$ have $D_q(\Pi_{\overline{\varepsilon}}) \ge q$.

Hence, all that remains is to see (focusing on the nontrivial case that $q_0 < q_1$) that sufficiently small $\varepsilon \in (0, \overline{\varepsilon}]$ have $D_q(\Pi_\varepsilon) \ge q$ for every $(\tilde{q}_0, \tilde{q}_1)$. And indeed, given Lemma 2, Berge's theorem tells us the function $[0, \overline{\varepsilon}] \to \mathbb{R}$ given by $\varepsilon \mapsto \min_{q \in [\tilde{q}_0, \tilde{q}_1]} [D_q(\Pi_\varepsilon)]$ is well-defined and continuous. Because $[\tilde{q}_0, \tilde{q}_1] \subset (q_0, q_1)$, this function is strictly positive at $\varepsilon = 0$, and so is strictly positive for small enough $\varepsilon \in (0, \overline{\varepsilon}]$, delivering the claim. *Q.E.D.*

Now, with these claims in hand, we pursue the proof of the theorem.

Proof of Theorem 1. First, given Claim 1, any optimal (Π^*, q^*) must have $\max \operatorname{Supp} \Pi^* \le p^M(q^*)$.

Now, we show that any optimal (Π^*, q^*) has Π^* greedy up to the top of its support. To that end, consider $q^* \in [0,1]$ and $\Pi \in \Delta(\mathbb{R}_+)$ such that $D_q(\Pi) \geq q$ for every $q \in (0, q^*)$, and Π is not greedy up to $p^* := \max \operatorname{Supp} \Pi$. We want to show (Π, q^*) cannot be optimal. We have nothing to show (given the previous paragraph) if $p^* \geq \overline{v}(q^*)$, so without loss say $p^* < \overline{v}(q^*)$. Now, by hypothesis, the set

$$\left\{q \in \left(0, \underline{q}(p^*)\right) : \ D_q(\Pi) > q\right\}$$

is nonempty. Meanwhile, Lemma 2 implies this set is open in $\mathbb{R}$, and so every connected component of it is an open interval. Let (q_0, q_1) be such a connected component. Claim 3 (which applies because $p^* < \overline{v}(q^*)$) tells us Π is nondegenerate on $[q_0, q_1]$. Hence, Claim 4 delivers some $\tilde{q}_0, \tilde{q}_1 \in [q_0, q_1]$ with $\tilde{q}_0 < \tilde{q}_1$ such that Π is nondegenerate on $[\tilde{q}_0, \tilde{q}_1]$; either $\tilde{q}_0 \in (q_0, q_1)$ or Π has no mass at q_0^{++}; and either $\tilde{q}_1 \in (q_0, q_1)$ or Π has no mass at q_1^{--}. Moreover, by Lemma 2, we know $D_{q_0}(\Pi) \geq q_0$ and $D_{q_1}(\Pi) \geq q_1$. Hence, applying Claim 2, we therefore have that either $D_{\tilde{q}_0}(\Pi) > \tilde{q}_0$ or $\partial^+ D_{\tilde{q}_0}(\Pi) > 1$; and either $D_{\tilde{q}_1}(\Pi) > \tilde{q}_1$ or $\partial^- D_{\tilde{q}_1}(\Pi) < 1$. So given any $\tilde{\Pi} \in \Delta(\mathbb{R}_+)$ with $D_{\tilde{q}_0}(\tilde{\Pi}) \geq \tilde{q}_0$ and $D_{\tilde{q}_1}(\tilde{\Pi}) \geq \tilde{q}_1$, Claim 5 tells us sufficiently small $\varepsilon \in (0,1)$ has $D_q\left((1-\varepsilon)\Pi + \varepsilon\tilde{\Pi}\right) \geq q$ for every $q \in [\tilde{q}_0, \tilde{q}_1]$.

We are now equipped to show (Π, q^*) is suboptimal. For any $p \in [\overline{v}(\tilde{q}_0), \overline{v}(\tilde{q}_1)]$, consider the price distribution Π^p which coincides with Π on $[0, \overline{v}(\tilde{q}_0)) \cup [\overline{v}(\tilde{q}_1), \infty)$, takes value $\Pi\left(\overline{v}(\tilde{q}_0)^-\right)$ on $[\overline{v}(\tilde{q}_0), p)$, and takes value $\Pi\left(\overline{v}(\tilde{q}_1)\right)$ on $[p, \overline{v}(\tilde{q}_1))$. Observe that $p \mapsto D_{q_1}(p)$ is decreasing, Π lies between $\Pi^{\overline{v}(\tilde{q}_0)}$ and $\Pi^{\overline{v}(\tilde{q}_1)}$ (in the sense of first-order stochastic dominance), and $p \mapsto D_{q_1}(\Pi^p)$ is continuous by Lemma 2. Hence, the intermediate value function yields some $p \in [\overline{v}(\tilde{q}_0), \overline{v}(\tilde{q}_1)]$ such that $D_{q_1}(\Pi^p) = D_{q_1}(\Pi)$. For any $\varepsilon \in (0,1)$, let $\Pi_\varepsilon := (1-\varepsilon)\Pi + \varepsilon\Pi^p$. By construction, every $q \in (0, q_0]$ has $D_q(\Pi^p) = D_q(\Pi) \geq q$. Meanwhile Lemma 5 tells us $R_{q^*}(\Pi^p) > R_{q^*}(\Pi)$ and every $q \in [q_1, q^*)$ has $D_q(\Pi^p) \geq D_q(\Pi) \geq q$. Therefore, for any $\varepsilon \in (0,1)$, we have $R_{q^*}(\Pi_\varepsilon) > R_{q^*}(\Pi)$ and $D_q(\Pi^p) \geq q$ for every $q \in (0, q_0] \cup [q_1, q^*)$. Finally, as noted in the previous paragraph, sufficiently small $\varepsilon \in (0,1)$ has $D_q(\Pi^p) \geq q$ for every $q \in [\tilde{q}_0, \tilde{q}_1]$. So (Π^ε, q^*) witnesses that (Π, q^*) is suboptimal, as claimed above.

Now, letting (Π^*, q^*) be optimal and $p^* := \max \operatorname{Supp} \Pi^*$, we have established that Π^* is greedy up to p^* and $p^* \leq p^M(q^*)$. All that remains is to see Π^* has a mass point at p^*. To that end, note that $p^* \leq p^M(q^*) < \overline{v}(q^*)$ implies $\underline{q}(p^*) < q^*$. We can therefore apply Claim 2 to learn $\partial^+ D_{\underline{q}(p^*)}(\Pi^*) > 1$. But greediness up to p^* directly tells us $\partial^- D_{\underline{q}(p^*)}(\Pi^*) = 1$, and so Lemma 7 implies Π^* has a mass point at $\underline{q}(p^*)$. *Q.E.D.*

B.5. Proof of Corollary 1

Corollary 1 follows directly from Theorem 1 and the next Lemma 8, which shows that greediness rules out mass points and gaps in a price distribution.

Lemma 8. *Suppose $\hat{q} \in (0,1]$ and $\Gamma : [0, \overline{v}(\hat{q})) \to \mathbb{R}_+$ is increasing and right continuous with $D_q(\Gamma) = q$ for every $q \in (0, \hat{q})$. Then Γ is strictly increasing and continuous on $[0, \overline{v}(\hat{q}))$ with $\Gamma(0) = 0$.*

Proof. By hypothesis, $q \mapsto D_q(\Gamma)$ is differentiable on $(0, \hat{q})$, and so Lemma 7 tells us Γ has no discontinuities in $(0, \overline{v}(\hat{q}))$. Moreover, $\Gamma(0) = D_0(\Gamma) = 0$.

To show Γ is strictly increasing on $[0, \overline{v}(\hat{q}))$, it suffices to show (given that it is weakly increasing by definition) that it is not constant over any interval. So suppose $0 < q_0 < q_1 < \overline{v}(\hat{q})$. Because $\Gamma(q_0) = q_0 > 0$, Lemma 6 would imply $q \mapsto D_q(\Gamma)$ is strictly concave if $\Gamma \circ \overline{v}$ were constant on (q_0, q_1). But this function is linear by hypothesis, hence not strictly concave. It follows that $\Gamma \circ \overline{v}$ is constant on (q_0, q_1), delivering the claim. *Q.E.D.*

C. Proofs for Section 5

C.1. Preliminaries

Lemma 9. *In the linear demand environment, define $\Gamma^* : [0, \overline{v}(1)] \to \mathbb{R}$ by*

$$\Gamma^*(v) := \underline{q}(v) + \frac{v}{\overline{v}'(\underline{q}(v))}.$$

(i) The function Γ^* is continuous and strictly increasing, and every $q \in [0,1]$ has

$$\overline{v}(q) D_q(\Gamma^*) = \int_0^{\overline{v}(q)} \Gamma^* = q\overline{v}(q).$$

(ii) *The function* Γ^* *is greedy.*[32] *Conversely, if* $\hat{q} \in [0,1]$ *and* Γ *is greedy up to* $\overline{v}(\hat{q})$*, then* Γ *agrees with* Γ^* *on* $[0, \overline{v}(\hat{q}))$*.*

(iii) *A unique* $\overline{p}^* \in (0, \overline{v}(1))$ *exists with* $\Gamma^*(\overline{p}^*) = 1$*.*

(iv) *Every* $\hat{p} \in [0, \overline{p}^*]$ *admits a unique* $\hat{q} \in [\underline{q}(\overline{p}^*), 1]$ *such that* $\int_{\hat{p}}^{\overline{v}(\hat{q})}(1-\Gamma^*) = 0$*, and this* $\hat{q}$ *strictly decreases as* $\hat{p}$ *increases.*

Proof. First, let us observe that Γ^* is continuous and weakly increasing. It is continuous because $\overline{v}$ is continuously differentiable and $\overline{v}'$ is strictly positive. To see it is weakly increasing (or equivalently, that $\Gamma^* \circ \overline{v}$ is) we apply convexity of $\overline{v}$. First, consider the case in which $\overline{v}$ is twice differentiable. In this case, on $(0,1]$ (where $\overline{v}, \overline{v}'$ are both strictly positive), we have[33]

$$\begin{aligned}
(\Gamma \circ \overline{v})' &= 1 + \left(\tfrac{\overline{v}}{\overline{v}'}\right)' = 1 + \tfrac{(\overline{v}')^2 - \overline{v}\ \overline{v}''}{(\overline{v}')^2} = \tfrac{2(\overline{v}')^2 - \overline{v}\ \overline{v}''}{(\overline{v}')^2} \\
&= \tfrac{\overline{v}^3}{(\overline{v}')^2} \tfrac{2(\overline{v}')^2 - \overline{v}\ \overline{v}''}{\overline{v}^3} = \tfrac{\overline{v}^3}{(\overline{v}')^2} \tfrac{2\overline{v}\ \overline{v}'\overline{v}' - \overline{v}^2 \overline{v}''}{\overline{v}^4} = \tfrac{\overline{v}^3}{(\overline{v}')^2} \left(\tfrac{-\overline{v}'}{\overline{v}^2}\right)' = \tfrac{\overline{v}^3}{(\overline{v}')^2} \left(\tfrac{1}{\overline{v}}\right)'' \\
&\geq 0,
\end{aligned}$$

where the last inequality holds because $\frac{1}{\overline{v}}$ is convex. The general case—in which $\overline{v} : [0,1] \to \mathbb{R}$ is an arbitrary continuously differentiable function that is zero at zero, has strictly positive derivative, and has $\frac{1}{\overline{v}}$ convex on $(0,1]$—follows from an approximation argument.[34]

Now, let $\Gamma : \mathbb{R}_+ \to \mathbb{R}_+$ be any increasing and right-continuous function, and note that $\overline{v}(q) D_q(\Gamma) = \int_0^{\overline{v}(q)} \Gamma$ for every $q \in [0,1]$. Given $\hat{q} \in [0,1]$, let us now show Γ is greedy up to $\overline{v}(\hat{q})$—or equivalently, has $\int_0^{\overline{v}(q)} \Gamma = q\overline{v}(q)$ for every $q \in [0, \hat{q}]$—if and only if Γ agrees with Γ^* on $[0, \overline{v}(\hat{q}))$. To that end, note Lemma 8 tells us Γ can be greedy up to $\overline{v}(\hat{q})$ only if it is continuous on $[0, \overline{v}(\hat{q}))$ with $\Gamma(0) = 0$, so we can focus on such Γ. Because the equality $\int_0^{\overline{v}(q)} \Gamma = q\overline{v}(q)$ holds for $q = 0$ and both sides are differentiable in q, it holds for every $q \in (0, \hat{q})$

[32] Our main text defines greediness only for (increasing and right-continuous) functions $\mathbb{R}_+ \to \mathbb{R}_+$, but the definition can be applied verbatim to a function defined on $[0, \overline{v}(1)]$.

[33] This argument is substantively the same as the observation (**?**, footnote 11) that a type distribution is regular if and only if the inverse of its survival function is convex.

[34] Any such function is easily seen to be a limit in $C^1[0,1]$ of such functions that are also twice differentiable, and the induced $\Gamma^* \circ \overline{v}$ is then a limit of those for the approximating models, hence is weakly increasing.

if and only if the derivatives coincide at every $q \in (0, \hat{q})$—that is

$$q\overline{v}'(q) + \overline{v}(q) = \overline{v}'(q)\Gamma\left(\overline{v}(q)\right).$$

Rearranging, Γ is greedy if and only if it agrees with Γ^* on $[0, \overline{v}(\hat{q}))$.

In particular, the equivalence of the previous paragraph tells us Γ^* is greedy, and Lemma 8 says it is strictly increasing on $[0, \overline{v}(1))$. Now, because Γ^* is strictly increasing, at most one $\overline{p}^* \in (0, \overline{v}(1))$ can exist with $\Gamma^*(\overline{p}^*) = 1$. Because Γ^* is continuous and

$$\Gamma^*(0) = 0 < 1 < 1 + \tfrac{\overline{v}(1)}{\overline{v}'(1)} = \Gamma^*(\overline{v}(1)),$$

the intermediate value theorem tells us some such $\overline{p}^*$ exists.

Observe next, because Γ^* is strictly increasing, it follows that the function $[0, \overline{v}(1)] \to \mathbb{R}$ given by $p \mapsto \int_0^p (1 - \Gamma^*)$ is continuous and strictly concave and is maximized at $\overline{p}^*$. Moreover, its value at the right endpoint of its domain is $\int_0^{\overline{v}(1)} (1 - \Gamma^*) = \overline{v}(1) - \int_0^{\overline{v}(1)} \Gamma^* = \overline{v}(1) - 1\overline{v}(1) = 0$, the same as its value at the left endpoint. Therefore, every $\hat{p} \in [0, \overline{p}^*]$ admits a unique $\hat{p}' \in [\overline{p}^*, \overline{v}(1)]$ such that $\int_0^{\hat{p}'} (1 - \Gamma^*) = \int_0^{\hat{p}} (1 - \Gamma^*)$—hence a unique $\hat{q} \in [\underline{q}(\overline{p}^*), 1]$ such that $\int_{\hat{p}}^{\overline{v}(\hat{q})} (1 - \Gamma^*) = 0$. Moreover, this $\hat{p}'$ continuously strictly decreases as $\hat{p}$ increases, and so too does $\hat{q}$. *Q.E.D.*

In line with the previous lemma, we can introduce the following notations:

Notation 2. *In the linear demand environment:*

(i) Define $\mathcal{Q} : [0, \overline{p}^*] \to [\underline{q}(\overline{p}^*), 1]$ *to be the unique function mapping any* $\hat{p} \in [0, \overline{p}^*]$ *to the unique* $\hat{q} \in [\underline{q}(\overline{p}^*), 1]$ *such that* $\int_{\hat{p}}^{\overline{v}(\hat{q})} (1 - \Gamma^*) = 0$.

(ii) Define $\mathcal{P} := \mathcal{Q}^{-1} : [\underline{q}(\overline{p}^*), 1] \to [0, \overline{p}^*]$ *and* $\mathcal{V} := v \circ \mathcal{Q} : [0, \overline{p}^*] \to [\overline{p}^*, \overline{v}(1)]$.

(iii) For each $\hat{p} \in [0, \overline{p}^*]$*, define* $\Pi(\cdot|\hat{p}) \in \Delta(\mathbb{R}_+)$ *via*

$$\Pi(p|\hat{p}) := \begin{cases} \Gamma^*(p) & : \ p < \hat{p} \\ 1 & : \ p \geq \hat{p}. \end{cases}$$

(iv) Define $\mathcal{R} : [0, \overline{p}^*] \times [\underline{q}(\overline{p}^*), 1] \to \mathbb{R}$ *by* $\mathcal{R}(\hat{p}, \hat{q}) := R_{\hat{q}}(\Pi_{\hat{p}})$.

Now, let us record some useful computations about these objects.

Lemma 10. *In the linear demand environment:*

(i) The functions $\mathcal{V}$ *and* $\mathcal{Q}$ *are continuously differentiable on* $[0, \overline{p}^*)$*, and* $\mathcal{P}$ *is continuously differentiable on* $(\underline{q}(\overline{p}^*), 1]$*. Any* $\hat{p} \in [0, \overline{p}^*)$ *and* $\hat{q} = \mathcal{Q}(\hat{p}) \in (\underline{q}(\overline{p}^*), 1]$ *have*

$$\mathcal{V}'(\hat{p}) = -\frac{1 - \Gamma^*(\hat{p})}{\Gamma^*(\overline{v}(\hat{q})) - 1}, \ \mathcal{Q}'(\hat{p}) = \frac{\mathcal{V}'(\hat{p})}{\overline{v}'(\hat{q})}, \ \textit{and } \mathcal{P}'(\hat{q}) = \frac{1}{\mathcal{Q}'(\hat{p})},$$

which are all strictly negative.

(ii) Any $\hat{p} \in [0, \overline{p}^*]$ *has*

$$\int_0^\infty p^2 \ \mathrm{d}\Pi(p|\hat{p}) = 2 \int_0^{\hat{p}} p[1 - \Gamma^*(p)] \ \mathrm{d}p = \hat{p}^2 \left[1 - \underline{q}(\hat{p})\right] - \int_0^{\underline{q}(\hat{p})} \overline{v}^2.$$

(iii) Any $\hat{p} \in [0, \overline{p}^*]$ *and* $\hat{q} \in [\underline{q}(\overline{p}^*), 1]$ *have*

$$\mathcal{R}(\hat{p}, \hat{q}) = \int_0^{\hat{p}} \left[1 - \tfrac{2p}{\overline{v}(\hat{q})}\right] [1 - \Gamma^*(p)] \ \mathrm{d}p.$$

(iv) Any $\hat{p} \in [0, \overline{p}^*)$ *has* $\frac{\mathrm{d}}{\mathrm{d}\hat{p}} \mathcal{R}\left(\hat{p}, \mathcal{Q}(\hat{p})\right) = \frac{1-\Gamma^*(\hat{p})}{\mathcal{V}(\hat{p})} r(\hat{p})$*, where*

$$r(\hat{p}) := [\mathcal{V}(\hat{p}) - 2\hat{p}] - \frac{2}{\mathcal{V}(\hat{p}) \left[\Gamma^*(\mathcal{V}(\hat{p})) - 1\right]} \int_0^{\hat{p}} p \left[1 - \Gamma^*(p)\right] \ \mathrm{d}p.$$

(v) The function $r : [0, \overline{p}^*) \to \mathbb{R}$ *is continuously differentiable with strictly negative derivative.*

(vi) The function $[0, \overline{p}^*] \to \mathbb{R}$ *given by* $\hat{p} \mapsto \mathcal{R}\left(\hat{p}, \mathcal{Q}(\hat{p})\right)$ *is strictly quasiconcave with interior maximum.*

Proof. We first establish the derivative computations for $\mathcal{V}$, $\mathcal{Q}$, and $\mathcal{P}$. We need only show the given properties for $\mathcal{V}$, and then those for $\mathcal{Q}$ and $\mathcal{P}$ follow directly from the chain rule. At any $\hat{p} \in [0, \overline{p}^*)$, that $\mathcal{Q}(\hat{p}) > \underline{q}(\overline{p}^*)$ implies the second partial derivative of the continuously differentiable function $(p, v) \mapsto \int_p^v (1 - \Gamma^*)$ is nonzero at $(\hat{p}, \mathcal{V}(\hat{p}))$. The implicit function theorem therefore implies $\mathcal{V}$ is differentiable at $\hat{p}$ with

$$0 = \frac{\mathrm{d}}{\mathrm{d}\hat{p}} \int_{\hat{p}}^{\mathcal{V}(\hat{p})} (1 - \Gamma^*) = \mathcal{V}'(\hat{p}) \left[1 - \Gamma^*(\mathcal{V}(\hat{p}))\right] - \left[1 - \Gamma^*(\hat{p})\right].$$

Thus, $\mathcal{V}'$ is as desired.

Next, observe that the expectation of the squared price given $\Pi(\cdot|\hat{p})$ is

$$\begin{aligned}
\int_0^\infty p^2 \ \mathrm{d}\Pi(p|\hat{p}) &= [1 - \Gamma^*(\hat{p})]\hat{p}^2 + \int_0^{\hat{p}} p^2 \ \mathrm{d}\Gamma^*(p) \\
&= \hat{p}^2 - \hat{p}^2\Gamma^*(\hat{p}) + \left[p^2\Gamma^*(p)\right]_{p=0}^{\hat{p}} - \int_0^{\hat{p}} 2p\Gamma^*(p) \ \mathrm{d}p \\
&= \hat{p}^2 - 2\int_0^{\hat{p}} p\Gamma^*(p) \ \mathrm{d}p \\
&= 2\int_0^{\hat{p}} p \left[1 - \Gamma^*(p)\right] \ \mathrm{d}p,
\end{aligned}$$

which is in turn equal to $\hat{p}^2 \left[1 - \underline{q}(\hat{p})\right] - \int_0^{\underline{q}(\hat{p})} \overline{v}^2$ because the two expressions are both zero for $\hat{p} = 0$ and have the same derivative with respect to $\hat{p}$.

Toward computing $\mathcal{R}(\hat{p}, \hat{q})$, note that

$$\int_0^\infty p \ \mathrm{d}\Pi(p|\hat{p}) = \int_0^\infty [1 - \Pi(\cdot|\hat{p})] = \int_0^{\hat{p}} (1 - \Gamma^*).$$

Moreover, that $\hat{p} \in [0, \overline{p}^*]$ and $\hat{q} \in [\underline{q}(\overline{p}^*), 1]$ implies $\hat{p} \leq \overline{v}(\hat{q})$. Thus,

$$\begin{aligned}
\overline{v}(\hat{q})\mathcal{R}(\hat{p}, \hat{q}) &= \overline{v}(\hat{q}) \int_0^\infty p \left[1 - \tfrac{p}{\overline{v}(\hat{q})}\right] \mathrm{d}\Pi(p|\hat{p}) \\
&= \overline{v}(\hat{q}) \int_0^\infty p \ \mathrm{d}\Pi(p|\hat{p}) - \int_0^\infty p^2 \ \mathrm{d}\Pi(p|\hat{p}) \\
&= \overline{v}(\hat{q}) \int_0^{\hat{p}} (1 - \Gamma^*) - 2\int_0^{\hat{p}} p \left[1 - \Gamma^*(p)\right] \mathrm{d}p.
\end{aligned}$$

Hence, $\mathcal{R}(\hat{p}, \hat{q}) = \int_0^{\hat{p}} \left[1 - \frac{2p}{\overline{v}(\hat{q})}\right] \left[1 - \Gamma^*(p)\right] \mathrm{d}p$.

Now, because

$$\begin{aligned}
\left.\frac{\partial}{\partial \hat{q}}\right|_{\hat{q}=\mathcal{Q}(\hat{p})} \mathcal{R}\left(\hat{p}, \hat{q}\right) &= \left\{\int_0^{\hat{p}} 2p\left[1-\Gamma^*(p)\right] \,\mathrm{d}p\right\} \left.\frac{\partial}{\partial \hat{q}}\right|_{\hat{q}=\mathcal{Q}(\hat{p})} \left[\frac{-1}{\overline{v}(\hat{q})}\right] \\
&= \frac{2\overline{v}'(\hat{q})}{\overline{v}(\hat{q})^2} \int_0^{\hat{p}} p\left[1-\Gamma^*(p)\right] \,\mathrm{d}p,
\end{aligned}$$

the chain rule yields

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}\hat{p}}\mathcal{R}\left(\hat{p}, \mathcal{Q}(\hat{p})\right) &= \frac{\partial}{\partial \hat{p}}\mathcal{R}\left(\hat{p}, \mathcal{Q}(\hat{p})\right) + \mathcal{Q}'(\hat{p}) \left.\frac{\partial}{\partial \hat{q}}\right|_{\hat{q}=\mathcal{Q}(\hat{p})} \mathcal{R}\left(\hat{p}, \hat{q}\right) \\
&= \left[1-\frac{2\hat{p}}{\overline{v}(\mathcal{Q}(\hat{p}))}\right]\left[1-\Gamma^*(\hat{p})\right] + \frac{\mathcal{V}'(\hat{p})}{\overline{v}'(\mathcal{Q}(\hat{p}))}\frac{2\overline{v}'(\mathcal{Q}(\hat{p}))}{\overline{v}(\mathcal{Q}(\hat{p}))^2}\int_0^{\hat{p}} p\left[1-\Gamma^*(p)\right] \,\mathrm{d}p \\
&= \frac{1-\Gamma^*(\hat{p})}{\overline{v}(\mathcal{Q}(\hat{p}))}\left\{\left[\overline{v}(\mathcal{Q}(\hat{p}))-2\hat{p}\right] + \frac{2\mathcal{V}'(\hat{p})}{\overline{v}(\mathcal{Q}(\hat{p}))\left[1-\Gamma^*(\hat{p})\right]}\int_0^{\hat{p}} p\left[1-\Gamma^*(p)\right] \,\mathrm{d}p\right\} \\
&= \frac{1-\Gamma^*(\hat{p})}{\overline{v}(\mathcal{Q}(\hat{p}))} r(\hat{p}).
\end{aligned}$$

Next, that r is continuously differentiable on $[0, \overline{p}^*)$ follows directly from $\mathcal{V}$ being so and Γ^* being continuous. To see r has strictly negative derivative on $[0, \overline{p}^*)$, it suffices to see that $r(\hat{p}) + 2\hat{p}$ is decreasing on this range. And indeed, because $\mathcal{V}$ is decreasing there, it follows that $\hat{p} \mapsto r(\hat{p}) + 2\hat{p}$ is a decreasing function minus the ratio of a positive increasing function to a positive decreasing function—and hence is decreasing as desired.

Finally, because $\frac{\mathrm{d}}{\mathrm{d}\hat{p}}\mathcal{R}\left(\hat{p}, \mathcal{Q}(\hat{p})\right)$ is a strictly positive multiple of $r(\hat{p})$, which is strictly decreasing in $\hat{p} \in [0, \overline{p}^*)$, it follows that $\hat{p} \mapsto \mathcal{R}\left(\hat{p}, \mathcal{Q}(\hat{p})\right)$ is strictly quasiconcave on $[0, \overline{p}^*)$—hence on $[0, \overline{p}^*]$ by continuity. Moreover, $\hat{p} \mapsto \mathcal{R}\left(\hat{p}, \mathcal{Q}(\hat{p})\right)$ is maximized on the interior of its domain if r has an interior root. And indeed, $r(0) = \overline{v}(1) > 0$, whereas any $\hat{p} \in [0, \overline{p}^*)$ close enough to $\overline{p}^*$ has $\mathcal{V}(\hat{p}) < 2\hat{p}$ and so $r(\hat{p}) < 0$. Therefore, r has an interior root by the intermediate value theorem. *Q.E.D.*

Lemma 11. *In the linear demand environment,* (Π^*, q^*) *is optimal if and only if* $\Pi^* = \Pi(\cdot|p^*)$ *and* $q^* = \mathcal{Q}(p^*)$ *for the unique* $p^* \in (0, \overline{p}^*)$ *satisfying* $r(p^*) = 0$.

Proof. First, we observe any optimal (Π^*, q^*) must have $\Pi^* = \Pi(\cdot|\hat{p})$ for some

$\hat{p} \in [0, \overline{p}^*]$. To see this, note that Theorem 1 tells us Π^* is greedy up to the top of its support p^*. But then Lemma 9 tells us Π^* agrees with Γ^* on $[0, p^*)$, and so fact that Π^* is in $\Delta(\mathbb{R}_+)$ tells us $\Pi^* = \Pi(\cdot|p^*)$ and $p^* \leq \overline{p}^*$.

Now we argue that, given $\hat{p} \in [0, \overline{p}^*]$, the set of all $q \in (0,1]$ with $D_q(\Pi(\cdot|\hat{p})) \geq q$ is equal to $[0, \mathcal{Q}(\hat{p})]$. Toward this characterization, first note (given Lemma 9) any $q \in [0, \underline{q}(\hat{p})]$ has $D_q(\Pi(\cdot|\hat{p})) = \int_0^{\overline{v}(q)} \Pi(\cdot|\hat{p}) f_q = \int_0^{\overline{v}(q)} \Gamma^* f_q = D_q(\Gamma^*) = q$. Next observe, any $q \in [\underline{q}(\hat{p}), 1]$ has (again by Lemma 9)

$$\overline{v}(q) D_q(\Pi(\cdot|\hat{p})) = \overline{v}(q) \int_0^{\overline{v}(q)} \Pi(\cdot|\hat{p}) f_q = \int_0^{\overline{v}(q)} \Pi(\cdot|\hat{p}) = \int_0^{\hat{p}} \Gamma^* + \int_{\hat{p}}^{\overline{v}(q)} 1,$$

and so $\overline{v}(q) \left[D_q(\Pi(\cdot|\hat{p})) - q\right] = \int_0^{\hat{p}} \Gamma^* + \int_{\hat{p}}^{\overline{v}(q)} 1 - \int_0^{\overline{v}(q)} \Gamma^* = \int_{\hat{p}}^{\overline{v}(q)} (1 - \Gamma^*)$. Therefore, because Lemma 9 says Γ^* is strictly increasing, the function $q \mapsto \overline{v}(q) \left[D_q(\Pi(\cdot|\hat{p})) - q\right]$ is strictly quasiconcave on $[\underline{q}(\hat{p}), 1]$ and zero at $\underline{q}(\hat{p})$. Because the function also takes value zero at $\mathcal{Q}(\hat{p})$, it is then nonnegative up to $\mathcal{Q}(\hat{p})$ and strictly negative to the right. Thus, $\{q \in (0,1] : \ D_q(\Pi(\cdot|\hat{p})) \geq q\} = [0, \mathcal{Q}(\hat{p})]$, as desired.

By the previous two paragraphs, we can write the seller's problem (P*) as

$$\max_{\hat{p} \in [0, \overline{p}^*], \ \hat{q} \in [0,1]} R_{\hat{q}}(\Pi(\cdot|\hat{p})) \text{ s.t. } \hat{q} \leq \mathcal{Q}(\hat{p}).$$

Because Lemma 3 tells us the objective is strictly increasing (wherever strictly positive, as the optimal revenue is) in the quantity, the seller optimally sets $\hat{q} = \mathcal{Q}(\hat{p})$, and so her problem can be written as

$$\max_{\hat{p} \in [0, \overline{p}^*]} \mathcal{R}(\hat{p}, \mathcal{Q}(\hat{p})).$$

By Lemma 10, this objective is strictly quasiconcave with interior optimum, and the optimum p^* is characterized by $r(p^*) = 0$. *Q.E.D.*

C.2. Proof of Proposition 3

We begin by proving the following Lemma 12, which we will then use in the proof of Proposition 3.

Lemma 12. *Take the linear demand environment.*

(i) If the seller posts a price strictly greater than $\overline{p}^B := \int_0^{\overline{p}^*}(1-\Gamma^*) \in (0, \overline{p}^*)$*, then the highest equilibrium quantity is zero.*

(ii) If the seller posts a price $\hat{p} \in [0, \overline{p}^B]$*, then a given* $\hat{q}$ *is an equilibrium quantity if and only if* $\hat{p} = \mathcal{P}^B(\hat{q})$*, where* $\mathcal{P}^B(\hat{q}) := \int_0^{\overline{v}(\hat{q})}(1-\Gamma^*)$*. In particular, the highest such quantity is the unique* $\mathcal{Q}^B(\hat{p}) \in [\underline{q}(\overline{p}^*), 1]$ *such that* $\mathcal{P}^B(\mathcal{Q}^B(\hat{p})) = \hat{p}$*.*

(iii) The functions $\mathcal{P}^B$ *and* $\mathcal{Q}^B$ *are continuously differentiable on* $[0, \overline{p}^B)$ *and* $(\underline{q}(\overline{p}^*), 1]$*, respectively. Any* $\hat{p} \in [0, \overline{p}^*)$ *and* $\hat{q} = \mathcal{Q}(\hat{p}) \in (\underline{q}(\overline{p}^*), 1]$ *have*

$$\frac{\mathrm{d}}{\mathrm{d}\hat{q}}\mathcal{P}^B(\hat{q}) = -\overline{v}'(\hat{q})\,[\Gamma^*(\overline{v}(\hat{q})) - 1] \text{ and } \frac{\mathrm{d}}{\mathrm{d}\hat{p}}\mathcal{Q}^B(\hat{p}) = \frac{1}{\frac{\mathrm{d}}{\mathrm{d}\hat{q}}\big|_{\hat{q}=\mathcal{Q}^B(\hat{p})}\mathcal{P}^B(\hat{q})},$$

which are both strictly negative.

(iv) Letting $\mathcal{R}^B(\hat{p}, \hat{q}) := R_{\hat{q}}(\hat{p})$*, any* $\hat{p} \in [0, \overline{p}^B)$ *has* $\frac{\mathrm{d}}{\mathrm{d}\hat{p}}\mathcal{R}^B\left(\hat{p}, \mathcal{Q}^B(\hat{p})\right) = \frac{1}{\overline{v}(\mathcal{Q}^B(\hat{p}))} r^B(\hat{p})$*, where*

$$r^B(\hat{p}) := \left[\overline{v}(\mathcal{Q}^B(\hat{p})) - 2\hat{p}\right] - \frac{1}{\overline{v}(\mathcal{Q}^B(\hat{p}))\,[\Gamma^*(\overline{v}(\mathcal{Q}^B(\hat{p}))) - 1]}\;\hat{p}^2.$$

(v) The function $r^B : [0, \overline{p}^B) \to \mathbb{R}$ *is continuously differentiable with strictly negative derivative.*

(vi) The function $[0, \overline{p}^B] \to \mathbb{R}$ *given by* $\hat{p} \mapsto \mathcal{R}^B\left(\hat{p}, \mathcal{Q}^B(\hat{p})\right)$ *is strictly quasi-concave with interior maximum.*

(vii) Under best-case equilibrium selection, the unique optimal price distribution is degenerate at the unique price $p^B \in (0, \overline{p}^*)$ *with* $r^B(p^B) = 0$*, and the unique best equilibrium quantity at that price is* $q^B := \mathcal{Q}^B(p^B)$*.*

Proof. By Proposition 1, the seller optimally chooses a deterministic price and so solves

$$\max_{\hat{p}\in\mathbb{R}_+,\ \hat{q}\in[0,1]} \mathcal{R}^B(\hat{p}, \hat{q}) \text{ s.t. } D_{\hat{q}}(\hat{p}) = \hat{q}.$$

By the same proposition, an optimum exists and has both price and quantity being strictly positive. Now, let us rewrite the equilibrium constraint. Any quantity $\hat{q} \in (0,1]$ and price $\hat{p} \in \mathbb{R}_+$ have

$$D_{\hat{q}}(\hat{p}) = \hat{q} \iff 1 - \tfrac{\hat{p}}{\overline{v}(\hat{q})} = \hat{q} \iff \hat{p} = (1-\hat{q})\overline{v}(\hat{q}) \iff \hat{p} = \int_0^{\overline{v}(\hat{q})} (1-\Gamma^*),$$

where the last equivalence follows from Lemma 9. Now, defining $\mathcal{P}^B : [0,1] \to \mathbb{R}$ given by $\mathcal{P}^B(\hat{q}) := \int_0^{\overline{v}(\hat{q})}(1-\Gamma^*)$, Lemma 9 tells us $\mathcal{P}^B$ is continuous and strictly quasiconcave with $\mathcal{P}^B(0) = \mathcal{P}^B(1) = 0$ and maximizer $\underline{q}(\overline{p}^*)$. Therefore, the range of $\mathcal{P}^B$ is $[0, \overline{p}^B]$ for $\overline{p}^B := \int_0^{\overline{p}^*}(1-\Gamma^*) \in (0, \overline{p}^*)$, and every $\hat{p} \in [0, \overline{p}^B]$ has one solution in $[0, \underline{q}(\overline{p}^*)]$ and one solution in $[\underline{q}(\overline{p}^*), 1]$ to $\mathcal{P}^B(\cdot) = \hat{p}$. So let $\mathcal{Q}^B : [0, \overline{p}^B] \to [\underline{q}(\overline{p}^*), 1]$ be such that $\mathcal{P}^B(\mathcal{Q}^B(\hat{p})) = \hat{p}$ for every $\hat{p} \in [0, \overline{p}^B]$; Lemma 3 tells us seller revenue is strictly increasing (wherever strictly positive, as the optimal revenue is) in the quantity, and so the best equilibrium quantity if the seller set price $\hat{p}$ is $\hat{q} = \mathcal{Q}^B(\hat{p})$. We can thus write the seller's problem under best case selection as

$$\max_{\hat{p} \in [0, \overline{p}^B]} \mathcal{R}^B(\hat{p}, \mathcal{Q}^B(\hat{p})).$$

Now, the function $\mathcal{Q}^B$ is continuous and strictly decreasing by construction. Moreover, because $\frac{\mathrm{d}}{\mathrm{d}\hat{q}}\mathcal{P}^B(\hat{q}) = -\overline{v}'(\hat{q})\,[\Gamma^*(\overline{v}(\hat{q})) - 1]$, which is continuous and strictly negative for $\hat{q} \in [0, \underline{q}(\overline{p}^B))$, the inverse function theorem tells us $\mathcal{Q}^B$ is continuously differentiable on $(0, \overline{p}^B]$ with derivative $\frac{\mathrm{d}}{\mathrm{d}\hat{p}}\mathcal{Q}^B(\hat{p}) = \frac{1}{\overline{v}'(\mathcal{Q}^B(\hat{p}))[\Gamma^* \circ \overline{v} \circ \mathcal{Q}^B(\hat{p}) - 1]}$ there. We are now equipped to compute the seller's first-order condition under best-case selection. For any $\hat{p} \in (0, \overline{p}^B)$, at $\hat{q} = \mathcal{Q}^B(\hat{p})$

and $\hat{v} = \overline{v}(\hat{q})$ we have

$$\begin{aligned}
\overline{v}(\mathcal{Q}^B(\hat{p}))\ \frac{\mathrm{d}}{\mathrm{d}\hat{p}}\mathcal{R}^B(\hat{p},\mathcal{Q}^B(\hat{p})) &= \hat{v}\left[\frac{\partial}{\partial\hat{p}}\mathcal{R}^B(\hat{p},\hat{q}) + (\mathcal{Q}^B)'(\hat{p})\ \frac{\partial}{\partial\hat{q}}\mathcal{R}^B(\hat{p},\hat{q})\right] \\
&= \hat{v}\left\{\frac{\partial}{\partial\hat{p}}\left[\hat{p}\left(1-\frac{\hat{p}}{\overline{v}(\hat{q})}\right)\right] + (\mathcal{Q}^B)'(\hat{p})\ \frac{\partial}{\partial\hat{q}}\left[\hat{p}\left(1-\frac{\hat{p}}{\overline{v}(\hat{q})}\right)\right]\right\} \\
&= \hat{v}\left\{\left[1-\frac{2\hat{p}}{\hat{v}}\right] + \frac{-1}{\overline{v}'(\hat{q})\,[\Gamma^*(\hat{v})-1]}\ \frac{\hat{p}^2}{\hat{v}^2}\overline{v}'(\hat{q})\right\} \\
&= [\hat{v}-2\hat{p}] - \frac{1}{\hat{v}\,[\Gamma^*(\hat{v})-1]}\ \hat{p}^2 \\
&= \left[\overline{v}(\mathcal{Q}^B(\hat{p}))-2\hat{p}\right] - \frac{1}{\overline{v}(\mathcal{Q}^B(\hat{p}))\,[\Gamma^*(\overline{v}(\mathcal{Q}^B(\hat{p})))-1]}\ \hat{p}^2,
\end{aligned}$$

which is $r^B(\hat{p})$. Because the denominator $\overline{v}(\mathcal{Q}^B(\hat{p}))\left[\Gamma^*(\overline{v}(\mathcal{Q}^B(\hat{p})))-1\right]$ is positive and decreasing in $\hat{p}\in(0,\overline{p}^B)$, it follows that the decreasing r^B has strictly negative derivative on $(0,\overline{p}^B)$—hence the seller's problem under best-case selection is strictly quasiconcave in the choice of posted price. To see an interior root of r^B exists, note that $r(0)=\overline{v}(1)>0$, whereas any $\hat{p}\in(0,\overline{p}^B)$ close enough to $\overline{p}^B$ has $v(\mathcal{Q}^B(\hat{p}))<2\hat{p}$ and so $r^B(\hat{p})<0$; thus the intermediate value theorem applies. *Q.E.D.*

Proof of Proposition 3. Given Lemma 12, we know the strictly positive price $p^B<\overline{p}^B<\overline{p}^*$ is the unique price such that $r^B(p^B)=0$, and $q^B=\mathcal{Q}^B(p^B)$. We will use these facts to compare with worst-case selection.

Let p^* and q^* be the highest supported price and equilibrium quantity as described in Lemma 11. We want to show that $p^*>p^B$ and $q^*>q^B$. Because $\mathcal{Q}$ is strictly decreasing, we can equivalently show that $\mathcal{P}(q^B)>p^*>p^B$. By Lemma 10, we can rewrite this condition as the requirement that $r(\mathcal{P}(q^B))<0<r(p^B)$. Toward both of these rankings, observe that any $\hat{p}\in(0,\overline{p}^*)$ has

$$\begin{aligned}
r(\hat{p}) &= r(\hat{p}) - r^B(p^B) \\
&= [\mathcal{V}(\hat{p})-2\hat{p}] - \left[\overline{v}(q^B)-2p^B\right] \\
&\quad -\tfrac{1}{\mathcal{V}(\hat{p})[\Gamma^*(\mathcal{V}(\hat{p}))-1]}\int_0^\infty p^2\ \mathrm{d}\Pi(p|\hat{p}) + \tfrac{1}{\overline{v}(q^B)[\Gamma^*(\overline{v}(q^B))-1]}\ (p^B)^2
\end{aligned}$$

Toward the price ranking, note that specializing the above calculation yields

$$r(p^B) = \mathcal{V}(p^B) - \overline{v}(q^B) - \tfrac{1}{\mathcal{V}(p^B)[\Gamma^*(\mathcal{V}(p^B))-1]} \int_0^\infty p^2 \, \mathrm{d}\Pi(p|p^B) + \tfrac{1}{\overline{v}(q^B)[\Gamma^*(\overline{v}(q^B))-1]} \, (p^B)^2.$$

That $\Pi(p^B|p^B) = 1$ implies $\int_0^\infty p^2 \, \mathrm{d}\Pi(p|p^B) \le (p^B)^2$; and that Γ^* is increasing (Lemma 9) and $\mathcal{V}$ decreasing implies $\frac{1}{\mathcal{V}\cdot[\Gamma^*\circ\mathcal{V}-1]}$ is increasing on $[0, \overline{p}^*)$. Hence, $r(p^B) > 0$ will follow if we show $\mathcal{V}(p^B) > \overline{v}(q^B)$. And indeed,

$$\begin{aligned}
\int_{\overline{v}(q^B)}^{\mathcal{V}(p^B)} (\Gamma^* - 1) &= \left(\int_0^{\overline{v}(q^B)} - \int_{p^B}^{\mathcal{V}(p^B)} - \int_0^{p^B} \right) (1 - \Gamma^*) \\
&= p^B - 0 - \int_0^{p^B} (1 - \Gamma^*) \\
&= \int_0^{p^B} \Gamma^* > 0,
\end{aligned}$$

delivering $\mathcal{V}(p^B) > \overline{v}(q^B)$ because $\Gamma^* > 1$ between them. The price ranking follows.

Toward the quantity ranking, observe that

$$\begin{aligned}
r(\mathcal{P}(q^B)) &= \left[\mathcal{V}(\mathcal{P}(q^B)) - 2\mathcal{P}(q^B)\right] - \left[\overline{v}(q^B) - 2p^B\right] \\
&\quad - \tfrac{1}{\mathcal{V}(\mathcal{P}(q^B))[\Gamma^*(\mathcal{V}(\mathcal{P}(q^B)))-1]} \int_0^\infty p^2 \, \mathrm{d}\Pi(p|\mathcal{P}(q^B)) + \tfrac{1}{\overline{v}(q^B)[\Gamma^*(\overline{v}(q^B))-1]} \, (p^B)^2 \\
&= -2\left[\mathcal{P}(q^B) - p^B\right] - \tfrac{1}{\overline{v}(q^B)[\Gamma^*(\overline{v}(q^B))-1]} \left[\int_0^\infty p^2 \, \mathrm{d}\Pi(p|\mathcal{P}(q^B)) - (p^B)^2 \right].
\end{aligned}$$

So $r(\mathcal{P}(q^B)) < 0$ would follow if we knew $\mathcal{P}(q^B) > p^B$ and $\int_0^\infty p^2 \, \mathrm{d}\Pi(p|\mathcal{P}(q^B)) \ge (p^B)^2$. To establish both inequalities, observe that

$$p^B = p^B - 0 = \left(\int_0^{\overline{v}(q^B)} - \int_{\mathcal{P}(q^B)}^{\mathcal{V}(\mathcal{P}(q^B))} \right) (1 - \Gamma^*) = \int_0^{\mathcal{P}(q^B)} (1 - \Gamma^*).$$

This identity first implies $\mathcal{P}(q^B) > p^B$ because $\int_0^{\mathcal{P}(q^B)} (1 - \Gamma^*) < \mathcal{P}(q^B)$. Then, to show $\int_0^\infty p^2 \, \mathrm{d}\Pi(p|\mathcal{P}(q^B)) \ge (p^B)^2$, it suffices to show $\int_0^\infty p^2 \, \mathrm{d}\Pi(p|\hat{p}) - \left[\int_0^{\hat{p}} (1 - \Gamma^*) \right]^2$ is nonnegative for any $\hat{p} \in [0, \overline{p}^*]$. And indeed, the expression is obviously zero for $\hat{p}=0$, and it satisfies

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}\hat{p}}\left\{\int_0^\infty p^2\ \mathrm{d}\Pi(p|\hat{p}) - \left[\int_0^{\hat{p}}(1-\Gamma^*)\right]^2\right\} &= \frac{\mathrm{d}}{\mathrm{d}\hat{p}}\left\{2\int_0^{\hat{p}} p\,[1-\Gamma^*(p)]\ \mathrm{d}p - \left[\int_0^{\hat{p}}(1-\Gamma^*)\right]^2\right\} \\
&= 2\hat{p}\,[1-\Gamma^*(\hat{p})] - 2\left[\int_0^{\hat{p}}(1-\Gamma^*)\right][1-\Gamma^*(\hat{p})] \\
&= 2\,[1-\Gamma^*(\hat{p})]\int_0^{\hat{p}}\Gamma^* \geq 0.
\end{aligned}$$

The quantity ranking follows.

Finally, we turn to the consumer surplus ranking. Let $\Pi^* := \Pi(\cdot|p^*)$ be the optimal (under worst-case selection) price distribution, and let Π be the modified price distribution given by capping the price at $\overline{v}(q^B)$—that is $\Pi(p)$ is equal to $\Pi^*(p)$ for $p < \overline{v}(q^B)$, and is equal to 1 for $p \geq \overline{v}(q^B)$. Observe, any $p \in \mathbb{R}_+$ has

$$\begin{aligned}
\tfrac{1}{\overline{v}(q^B)}\left[p^B - \min\{p,\ \overline{v}(q^B)\}\right] &= D_{q^B}(\min\{p,\ \overline{v}(q^B)\}) - D_{q^B}(p^B) \\
&= D_{q^B}(p) - q^B,
\end{aligned}$$

and so $\frac{1}{\overline{v}(q^B)}\left[p^B - \int_0^\infty p\ \mathrm{d}\Pi(p)\right] = D_{q^B}(\Pi^*) - q^B$, which is nonnegative because $q^B < q^*$ and (Π^*, q^*) satisfies the demand constraints. Having established $p^B \geq \int_0^\infty p\ \mathrm{d}\Pi(p)$, we now pursue the surplus ranking. To that end, define

$$\mathrm{CS}_q(p) := \int (v-p)_+\ \mathrm{d}F_q(v) = \int_p^\infty D_q$$

the consumer surplus associated with anticipated quantity q (hence demand curve D_q) and a price offer of p; and let $\mathrm{CS}_q(\hat{\Pi}) := \int \mathrm{CS}_q(p)\ \mathrm{d}\hat{\Pi}(p)$ for any price distribution $\hat{\Pi}$. Observe that $\mathrm{CS}_q(p)$ is decreasing in p, strictly convex in p (because D_q is strictly decreasing) wherever $0 \leq p \leq \overline{v}(q)$, and (because $u(\theta,\cdot)$ is increasing) increasing in q. Moreover, the price distribution Π is nondegenerate because $p^B > 0$ and Π^* has nondegenerate convex support including zero. Therefore,

$$\mathrm{CS}_{q^B}(p^B) \leq \mathrm{CS}_{q^B}\left(\int p \; \mathrm{d}\Pi(p)\right) < \mathrm{CS}_{q^B}(\Pi) = \mathrm{CS}_{q^B}(\Pi^*) \leq \mathrm{CS}_{q^*}(\Pi^*),$$

where the last inequality holds because $q^* \geq q^B$. *Q.E.D.*

C.3. Proof of Proposition 4

For any $\omega \in [0,1]$, define $\overline{v}_\omega : [0,1] \to \mathbb{R}$ by letting $\overline{v}_\omega(q) := \frac{1}{(1-\omega)\frac{1}{\overline{v}_0(q)} + \omega \frac{1}{\overline{v}_1(q)}}$ for $q \in (0,1]$, and $\overline{v}_\omega(0) := 0$; this $\overline{v}_\omega$ is also an instance of the linear demand environment. In particular, $\frac{1}{\overline{v}_\omega}$ inherits strict convexity from $\frac{1}{\overline{v}_0}$ and $\frac{1}{\overline{v}_1}$. Observe that any $q \in (0,1]$ has

$$\frac{\partial}{\partial \omega} \log \overline{v}_\omega(q) = -\frac{\partial}{\partial \omega} \log \frac{1}{\overline{v}_\omega(q)} = -\frac{\frac{1}{\overline{v}_1(q)} - \frac{1}{\overline{v}_0(q)}}{\frac{1}{\overline{v}_\omega(q)}} = \frac{\frac{\overline{v}_1(q)}{\overline{v}_0(q)} - 1}{(1-\omega)\frac{\overline{v}_1(q)}{\overline{v}_0(q)} + \omega},$$

which is strictly increasing in q because $\frac{\overline{v}_1(q)}{\overline{v}_0(q)}$ is. Equivalently, whenever $0 < q < \tilde{q} \leq 1$, we have $\frac{\partial}{\partial \omega}\left[\frac{\overline{v}_\omega(\tilde{q})}{\overline{v}_\omega(q)}\right] > 0$, a log-supermodularity property that will be useful in establishing the quantity ranking.

First, we pursue the price distribution ranking. That $\frac{\overline{v}_1}{\overline{v}_0}$ has nonnegative derivative on $(0,1]$ means $\frac{\overline{v}_1}{\overline{v}_1'} \leq \frac{\overline{v}_0}{\overline{v}_0'}$, and so $\Gamma_1^* \circ \overline{v}_1 \leq \Gamma_0^* \circ \overline{v}_0$. Using this fact, let us see that $\Gamma_1^*(p) < \Gamma_0^*(p)$ for any price p with $0 < p \leq \min\{\overline{v}_0(1), \overline{v}_1(1)\}$. To see it, let $q_\omega := \overline{v}_\omega^{-1}(p) \in (0,1]$ for each $\omega \in \{0,1\}$. That $\overline{v}_1$ exhibits stronger externalities than $\overline{v}_0$ implies $q_1 < q_0$—since the strictly increasing function $\frac{\overline{v}_1}{\overline{v}_0}$ is above 1 on $(0,1]$, it is in fact strictly above 1 there.[35] It follows that

$$\Gamma_0^*(p) = \Gamma_0^*(\overline{v}_0(q_0)) \geq \Gamma_1^*(\overline{v}_1(q_0)) > \Gamma_1^*(\overline{v}_1(q_1)) = \Gamma_1^*(p).$$

Therefore, given Lemma 11, any p with $0 < p < \min\{p_1^*, p_0^*\}$ has $\Pi_0^*(p) = \Gamma_0^*(p) < \Gamma_1^*(p) = \Pi_1^*(p)$.

Next, we turn to the quantity ranking. Define $\tilde{r}(\hat{q}, \omega) := \frac{r_\omega(\mathcal{P}_\omega(\hat{q}))}{\mathcal{P}_\omega(\hat{q})}$ for any

[35] This observation is the only part of the proposition that uses the condition that $\overline{v}_1 \geq \overline{v}_0$. In particular, the quantity ranking follows only from $\frac{\overline{v}_1}{\overline{v}_0}$ being strictly increasing on $(0,1]$.

$(\hat{q},\omega) \in [0,1) \times [0,1]$ with $\overline{v}_\omega(\hat{q}) > \overline{p}^*_\omega$. Below, we will show that the function $\tilde{r}$ has strictly negative partial derivative with respect to its second argument at (q^*_ω, ω) for any $\omega \in [0,1]$; let us now see that doing so would establish the result. First, an application of the implicit function theorem tells us $(\hat{q},\omega) \mapsto \mathcal{P}_\omega(\hat{q})$ is continuously differentiable on the range of $(\hat{q},\omega) \in [0,1) \times [0,1]$ with $\overline{v}_\omega(\hat{q}) > \overline{p}^*_\omega$, and that $\mathcal{P}_\omega$ is strictly decreasing (Lemma 10) tells us it is strictly positive there. Because $(\hat{p},\omega) \mapsto r_\omega(\hat{p})$ is continuously differentiable on the range of $(\hat{p},\omega) \in \mathbb{R}_+ \times [0,1]$ with $\hat{p} < \overline{p}^*_\omega$, it follows that $\tilde{r}$ is continuously differentiable on its domain. Next, observe, any $\omega \in [0,1]$ has

$$\frac{\partial}{\partial \hat{q}}\bigg|_{\hat{q}=q^*_\omega} \tilde{r}(\hat{q},\omega) = \frac{p^*_\omega r'_\omega(p^*_\omega) - 1 r_\omega(p^*_\omega)}{(p^*_\omega)^2}\mathcal{P}'_\omega(q^*_\omega) = \frac{\mathcal{P}'_\omega(q^*_\omega)}{p^*_\omega} r'_\omega(p^*_\omega) > 0,$$

where the second equality holds by Lemma 10 and Lemma 11 and the strict inequality follows from Lemma 10. Because q^*_ω is the unique solution $\hat{q}$ to $\tilde{r}(\hat{q},\omega) = 0$ for each $\omega \in [0,1]$, it follows that $\omega \mapsto q^*_\omega$ is continuously differentiable. Therefore, at any $\omega \in [0,1]$ and $\hat{q} = q^*_\omega$, we have

$$0 = \frac{\mathrm{d}}{\mathrm{d}\omega} 0 = \frac{\mathrm{d}}{\mathrm{d}\omega}\tilde{r}(q^*_\omega,\omega) = \left[\frac{\partial}{\partial\omega}\tilde{r}(\hat{q},\omega)\right] + \left[\frac{\partial}{\partial\hat{q}}\tilde{r}(\hat{q},\omega)\right]\left[\frac{\partial}{\partial\omega}q^*_\omega\right].$$

Because we have shown $\frac{\partial}{\partial\hat{q}}\big|_{\hat{q}=q^*_\omega}\tilde{r}(\hat{q},\omega) > 0$, the hypothesis that $\frac{\partial}{\partial\omega}\big|_{\hat{q}=q^*_\omega}\tilde{r}(\hat{q},\omega) < 0$ therefore implies $\frac{\partial}{\partial\omega}q^*_\omega > 0$; hence, $\omega \mapsto q^*_\omega$ is strictly increasing, and so $q^*_1 > q^*_0$.

Thus, all that remains is to show is that $\frac{\partial}{\partial\omega}\tilde{r}(\hat{q},\omega) < 0$ wherever $\tilde{r}(\hat{q},\omega)$ is zero. To that end, fix any $q^* \in (0,1)$ for the remainder of our analysis. Define now the continuously differentiable (by Lemma 9 and Lemma 10) functions

$$\begin{aligned}
r^* : (0,1) \times [0,1] &\to \mathbb{R} \\
(q,\omega) &\mapsto \frac{1-q}{1-q^*} - 2 - \frac{2(1-q^*)}{\overline{v}_\omega(q)^2(1-q)\left[\Gamma^*_\omega(\overline{v}_\omega(q^*)) - 1\right]}\int_0^{\overline{v}_\omega(q)} p\left[1-\Gamma^*_\omega(p)\right]\mathrm{d}p \\
&= \frac{1-q}{1-q^*} - 2 - \frac{1-q^*}{\Gamma^*_\omega(\overline{v}_\omega(q^*)) - 1}\left[1 - \frac{1}{\overline{v}_\omega(q)^2(1-q)}\int_0^q \overline{v}^2_\omega\right]
\end{aligned}$$

and

$$\begin{array}{rcl} q_* : \{\omega \in [0,1] : \ \overline{v}_\omega(q^*) > \overline{p}^*_\omega\} & \to & (0,1) \\ \omega & \mapsto & \overline{v}_\omega^{-1}\left(\mathcal{P}_\omega(q^*)\right). \end{array}$$

Now, for any ω in the domain of q_*, the definition of $\mathcal{P}_\omega$ implies

$$0 = \int_{\overline{v}_\omega(q_*(\omega))}^{\overline{v}_\omega(q^*)} (1-\Gamma^*) = \int_0^{\overline{v}_\omega(q^*)} (1-\Gamma^*) - \int_0^{\overline{v}_\omega(q_*(\omega))} (1-\Gamma^*) = (1-q^*)\overline{v}_\omega(q^*) - [1-q_*(\omega)]\,\overline{v}_\omega(q_*(\omega)),$$

where the last identity follows from Lemma 9. It follows that every such ω has $\tilde{r}(q^*,\omega) = r^*(q_*(\omega),\omega)$. Therefore,

$$\frac{\partial}{\partial\omega}\tilde{r}(q^*,\omega) = \left[\frac{\partial}{\partial q}\bigg|_{q=q_*(\omega)} r^*(q,\omega)\right] q_*'(\omega) + \left[\frac{\partial}{\partial\omega}\bigg|_{q=q_*(\omega)} r^*(q,\omega)\right].$$

To show $\frac{\partial}{\partial\omega}\tilde{r}(q^*,\omega) < 0$, it thus suffices to show $\frac{\partial}{\partial q}\big|_{q=q_*(\omega)} r^*(q,\omega)$ and $\frac{\partial}{\partial\omega}\big|_{q=q_*(\omega)} r^*(q,\omega)$ are both strictly negative and $q_*'(\omega)$ is strictly positive. We pursue each of these three inequalities.

Toward signing $q_*'(\omega)$, observe the definition of $q_*(\cdot)$ and Lemma 9 imply

$$0 = \int_{\overline{v}_\omega(q_*(\omega))}^{\overline{v}_\omega(q^*)} (1-\Gamma^*) = (1-q^*)\overline{v}_\omega(q^*) - [1-q_*(\omega)]\,\overline{v}_\omega(q_*(\omega)),$$

which rearranges to

$$1 - q^* = [1-q_*(\omega)]\,\frac{\overline{v}_\omega(q_*(\omega))}{\overline{v}_\omega(q^*)}.$$

Differentiating the above equation tells us

$$\begin{aligned}
0 &= [1-q_*(\omega)]\left.\frac{\partial}{\partial\omega}\right|_{q=q_*(\omega)}\left[\frac{\overline{v}_\omega(q)}{\overline{v}_\omega(q^*)}\right]+q_*'(\omega)\left.\frac{\partial}{\partial q}\right|_{q=q_*(\omega)}\left[(1-q)\,\frac{\overline{v}_\omega(q)}{\overline{v}_\omega(q^*)}\right]\\
&< q_*'(\omega)\left.\frac{\partial}{\partial q}\right|_{q=q_*(\omega)}\left[(1-q)\,\frac{\overline{v}_\omega(q)}{\overline{v}_\omega(q^*)}\right]\\
&= q_*'(\omega)\frac{\overline{v}_\omega'(q_*(\omega))}{\overline{v}_\omega(q^*)}\left[1-\Gamma^*(\overline{v}_\omega(q_*(\omega)))\right].
\end{aligned}$$

Hence, $q_*'(\omega)$ is strictly positive.

Now, to sign $\frac{\partial}{\partial\omega}r^*(q,\omega)$, it suffices to show

$$\frac{1-\frac{1}{\overline{v}_\omega(q)^2(1-q)}\int_0^q\overline{v}_\omega^2}{\Gamma_\omega^*(\overline{v}_\omega(q^*))-1}$$

has strictly positive partial derivative with respect to ω at $q=q_*(\omega)$. Because both the numerator and denominator are strictly positive there, it suffices to show the numerator is has positive partial derivative and denominator has negative partial derivative with respect to ω, at least one of them strictly so. First, the numerator's partial derivative is a strictly negative multiple of

$$\frac{\partial}{\partial\omega}\left[\frac{\int_0^q\overline{v}_\omega^2}{\overline{v}_\omega(q)^2}\right]=\frac{\partial}{\partial\omega}\int_0^q\left[\frac{\overline{v}_\omega(\hat{q})}{\overline{v}_\omega(q)}\right]^2\,\mathrm{d}\hat{q}=2\int_0^q\frac{\overline{v}_\omega(\hat{q})}{\overline{v}_\omega(q)}\,\frac{\partial}{\partial\omega}\left[\frac{\overline{v}_\omega(\hat{q})}{\overline{v}_\omega(q)}\right]\,\mathrm{d}\hat{q},$$

which is strictly negative. Second, the denominator's partial derivative is

$$\frac{\partial}{\partial\omega}\left[\Gamma_\omega^*(\overline{v}_\omega(q^*))-1\right]=\frac{\partial}{\partial\omega}\left[\frac{\overline{v}_\omega(q^*)}{\overline{v}_\omega'(q^*)}\right]=\frac{\partial}{\partial\omega}\left\{\frac{1}{\left.\frac{\partial}{\partial q}\right|_{q=q^*}\log\overline{v}_\omega(q)}\right\},$$

which is nonpositive by log-supermodularity. So $\frac{\partial}{\partial\omega}r^*(q,\omega)<0$.

Finally, we turn to signing $\frac{\partial}{\partial q}r^*(q,\omega)$. To that end, let $q:=q_*(\omega)$, let $v:=\overline{v}_\omega(q)$, and let $v':=\overline{v}_\omega'(q)$. Then, that $\Gamma_\omega^*(v)<1$ rearranges to $(1-q)v'>v$.

Hence,

$$\begin{aligned}
\frac{\partial}{\partial q}\left[\frac{\int_0^q \overline{v}_\omega^2}{\overline{v}_\omega(q)^2(1-q)}\right] &= \frac{v^2(1-q)v^2 - [2(1-q)vv' - v^2]\int_0^q \overline{v}_\omega^2}{v^4(1-q)^2} \\
&< \frac{v^2(1-q)v^2 - (2v^2 - v^2)\int_0^q \overline{v}_\omega^2}{v^4(1-q)^2} \\
&= \frac{1}{1-q}\left[1 - \frac{1}{(1-q)v^2}\int_0^q \overline{v}_\omega^2\right].
\end{aligned}$$

Therefore, at such q we have

$$\begin{aligned}
\frac{\partial}{\partial q} r^*(q,\omega) &= \frac{-1}{1-q^*} + \frac{1-q^*}{\Gamma_\omega^*(\overline{v}_\omega(q^*)) - 1}\cdot\frac{\partial}{\partial q}\left[\frac{\int_0^q \overline{v}_\omega^2}{\overline{v}_\omega(q)^2(1-q)}\right] \\
&< \frac{-1}{1-q^*} + \frac{1-q^*}{\Gamma_\omega^*(\overline{v}_\omega(q^*)) - 1}\left\{\frac{1}{1-q}\left[1 - \frac{1}{(1-q)v^2}\int_0^q \overline{v}_\omega^2\right]\right\} \\
&= \frac{-1}{1-q}\left[r^*(q,\omega) + 2\right] \\
&= \frac{-2}{1-q} < 0.
\end{aligned}$$

The quantity ranking follows. *Q.E.D.*

C.4. Proof of Proposition 5

To be added.