Robust Product Design and Pricing*

Kyungmin Kim† Nenad Kos‡

October 30, 2023

Abstract

We study the problem of product design and pricing by a monopolist who has no information about the distribution of consumers' horizontal tastes and maximizes her profit under the worst-case scenario. We show that her optimal strategy takes a strikingly simple form of dividing the taste space into a finite number of equal-length intervals and serving consumers on a randomly chosen interval. We obtain this result by studying the dual problem of finding a distribution of consumers' tastes that minimizes the seller's profit. Our analysis extends to the case where the seller can design several varieties of the product.

JEL Classification Numbers: D42, D80, L12.

Keywords: Hotelling; product design; robust pricing; profit guarantee; circular city; product variety.

1 Introduction

Anticipating the launch of the Game Boy handheld console in 1989, Nintendo faced a crucial decision: whether to bundle it with Super Mario Land, appealing mainly to children, or Tetris, which had the potential to attract a more mature audience. The decision was particularly challenging because the market for handheld consoles was fairly nascent, so information about consumers' tastes was scarce. Eventually, Nintendo chose the latter option for most markets (excluding Japan) and

*We thank Christoph Carnehl, Jay Pil Choi, Nima Haghpanah, Vijay Krishna, Ali Shourideh, and seminar audiences at Bocconi, CMU/Pittsburgh, Emory, and Penn State for many helpful comments. Part of this project was completed while Nenad was visiting EconTribute center at Bonn University. He is very grateful for their hospitality. Kyungmin Kim is supported by the Ministry of Education of the Republic of Korea and the National Research Foundation of Korea (NRF-2020S1A5A2A03043516).

†Department of Economics, Emory University, Email: kyungmin.kim@emory.edu

‡Bocconi University, Department of Economics, IGIER and CEPR, nenad.kos@unibocconi.it.

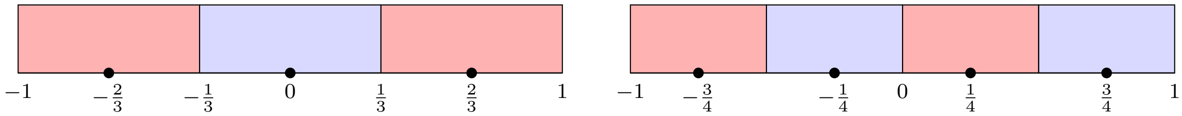

Figure 1: The seller's robustly optimal strategy when $n^* = 3$ (left) or $n^* = 4$.

experienced a resounding success. Companies developing new products commonly face such a design problem with little information about consumers' preferences. Opting for the wrong design can result in the product having a narrow appeal, a predicament that is not easily rectified even by a meticulous choice of the pricing strategy.

Motivated by the above, we study the problem of a monopolist who decides on the design of a new product and its price while lacking information about consumers' preferences. We use the Hotelling model to capture consumers' heterogeneous preferences as well as the seller's product design (choice).[1] Consumers are located along the interval $[-1, 1]$, with their location denoting their favorite design. The willingness to pay for their favorite product is normalized to 1. The products further away from consumers' favored designs are valued less. That is, consumers incur disutility from a mismatch between the design of the product and their preferred version. The disutility is an increasing and convex function of the distance between the two. Meanwhile, the seller chooses a location on the interval (the design) and the price to charge for the product.

To ascertain the limits of how the seller's lack of information affects the product design, pricing, and profit, we adopt the assumption that the seller has no information about the distribution of consumers' tastes. She considers any distribution over the interval $[-1, 1]$ plausible and evaluates her strategy according to the worst-case scenario: for each strategy prescribing (a distribution over) location-price pairs, she computes the lowest profit across all distributions. The seller's goal is to maximize such lowest profit. We refer to the resulting max-min profit as the *profit guarantee*. The model lends itself to the standard interpretation of the seller facing adversarial nature that wishes to minimize her profit.

Our main result shows that the seller's robustly optimal strategy takes a strikingly simple form: the seller divides the interval $[-1, 1]$ into a finite number of subintervals of equal width and targets (all) consumers on a randomly selected subinterval; visualized in Figure 1. The seller serves a subinterval by locating at its center and offering a price that makes the consumers at the edges of the subinterval indifferent between buying and not buying. A crucial property of the robustly optimal strategy is that the seller obtains the same payoff—profit guarantee—for *almost all* distributions,

[1] Hotelling (1929) proposes this product design interpretation of his model, using *sweetness of cider* as an example.

and never less than that.[2] The robust strategy fully insures the seller against the lack of information about the distribution of tastes.

To derive the robust outcome, we begin by tackling the dual problem, where nature first selects a distribution of tastes with the intention of minimizing the seller's profit. The seller observes the distribution and chooses the design/price combination to maximize the profit. We term the resulting min-max profit as the *lowest profit*. The dual problem serves two purposes. It provides an upper bound for the seller's profit guarantee.[3] In addition, it is of economic interest in and of itself, establishing the lowest profit the seller can be held to in the Bayesian setting.[4] The solution to the problem reveals crucial features of distributions yielding low profits to the seller.

When the seller knows the distribution of tastes, she can guarantee to sell to at least mass $1/n$ of consumers by splitting the interval into n subintervals of equal width and selling on the one that has the most customers. The seller's profit is therefore at least $1/n$ times the price required to cover the subinterval. Taking the supremum over n yields a more precise lower bound on the profit. Our main characterization result for the dual problem is that this bound is tight, that is, there exists a distribution of tastes under which the seller's maximized profit coincides with the bound. The seller-worst distribution we construct has a particularly simple structure: its density is a step function with only two values. The distribution and the seller's strategy of randomly covering one of the intervals constitute a saddle point, establishing strong duality in our environment.

A by-product of our analysis is the finding that the uniform distribution, which is commonly adopted in applications, is close to seller-worst.[5] Specifically, we show that the robustness problem corresponds to an integer version of the problem the seller solves when facing the uniform distribution. In the case when the disutility from the mismatch between the actual and preferred design of the product is linear, the seller's profit under the uniform distribution is never more than 12.5% above the minimal profit (the seller's profit guarantee).

The seller's robust strategy calls for the seller to cover the whole market, at least in probability. This is achieved by randomizing over a finite number of designs while keeping the price constant. A natural question to ask is, how the seller's strategy and the resulting profit would be affected by the possibility to design more than one variety of the product. Let n^* be the optimal number of subintervals when the seller is selling a single product. We show that the seller who can produce m

[2] The seller's strategy covers almost every location with probability $1/n^*$, where n^* is the number of subintervals, with the same price. The exceptions are the locations at the edges of the subintervals, which would be served twice as likely. Nature is, thus, indifferent between all the distributions that do not assign probability mass to those points and strictly prefers them to distributions that do.

[3] This is the standard max-min inequality; see, e.g., Osborne and Rubinstein (1994).

[4] The corresponding upper bound is trivially 1, which is achievable whenever the distribution of tastes is degenerate.

[5] In Section 6, we show that in the circular city environment, the uniform distribution is always seller-worst.

varieties of the product optimally divides the market into $\max\{m, n^*\}$ equal intervals and serves m randomly selected ones.[6] This preserves the seller's insurance against her lack of information, while simultaneously enabling her to reach a wider consumer base. The seller-worst distribution for the single-product case remains seller-worst as long as $m \leq n^*$, while if $m > n^*$, the uniform distribution minimizes the seller's profit.

Literature Review. Our paper lies at the intersection of the literatures on product design and robust pricing. The former literature dates back to Hotelling (1929) and is too large for us to survey meaningfully here. In broad strokes, product design has been studied in environments with multiple firms and both vertical and horizontal differentiation (e.g., Moorthy, 1988; Kuksov, 2004; Lauga and Ofek, 2011), optimal dynamic (re)positioning (e.g., Sweeting, 2013; Villas-Boas, 2018), and portfolio design (e.g., Villas-Boas, 2004; Orhun, 2009; Ke et al., 2022). The novelty of our work lies in studying the seller's product design problem when she has little information about consumers' tastes. Jovanovic (1981) and Meagher and Zauner (2004) incorporate uncertainty about the distribution of consumers' tastes into the Hotelling framework. They consider an environment where consumers are uniformly distributed over an interval of a fixed length, but the sellers do not know where the midpoint of the interval is. In our model, the support is fixed, but the seller is ignorant of the distribution.

Robust pricing (mechanism design) has been extensively studied in the environment without product design. Bergemann and Schlag (2011) study a standard monopoly model where the seller only knows the neighborhood the distribution of consumers' values belongs to. Carrasco et al. (2018) explore a model where the seller knows some moments of the distribution of consumers' values. Carroll (2017) studies robustness in a multi-dimensional setting, while Auster (2018) and Du (2018) characterize robust mechanisms in the common value settings. We benefit from basic and common insights in that literature but, to our knowledge, our analysis of the sellers' robustly optimal strategy and worst distributions is novel.

The remainder of this paper is organized as follows. Section 2 introduces our formal model. Section 3 studies the dual problem, explicitly constructing a seller-worst distribution and characterizing the seller's lowest profit. Section 4 analyzes the primal problem, presenting the seller's robustly optimal strategy and identifying the seller's profit guarantee. Sections 5 and 6 consider the optimal product line design problem and the circular city model, respectively. Section 7 concludes.

[6] If $m \geq n^*$, the seller covers the whole market with varieties.

2 The Model

A seller is facing a unit mass of consumers with heterogeneous tastes, modelled as in Hotelling (1929). Specifically, consumers are distributed over $[-1, 1]$ according to some distribution F. The seller designs her product as well as chooses its price p. We model product design as the seller's choice of location ℓ in $[-1, 1]$: Given ℓ, the willingness to pay of a consumer with taste $x \in [0, 1]$ is $1 - c(|x - \ell|)$, where $c : \mathbb{R}_+ \to \mathbb{R}_+$ is a differentiable, strictly increasing, and weakly convex function with $c(0) = 0$. The (normalized) value of the product that perfectly matches a consumer's taste is 1, and $c(|x - \ell|)$ represents disutility from preference misalignment between a consumer's taste x and the seller's product ℓ. We use σ to denote the seller's mixed strategy over $(\ell, p) \in [-1, 1] \times [0, 1]$ and Σ to denote the set of all mixed strategies.

Given (ℓ, p), a consumer at x purchases the product if and only if $1 - c(|x - \ell|) \geq p$. Let $\Delta(p) := c^{-1}(1 - p)$ denote the *reach* at price p—the largest distance from the seller at which the consumer is willing to purchase the product. $\Delta(\cdot)$ is continuous and strictly decreasing in p, so the seller can be interpreted as choosing reach Δ (around ℓ) instead of p. If the underlying distribution of tastes is F and the seller chooses (ℓ, p), the quantity demanded is

$$D(\ell, p; F) := F(\ell + \Delta(p)) - F_-(\ell - \Delta(p)),$$

where $F_-(x) := \lim_{x' \uparrow x} F(x)$, and the seller's profit is given by $\pi(\ell, p; F) := pD(\ell, p; F)$.

The seller has no information about the distribution F other than that its support lies in $[-1, 1]$. The seller, therefore, entertains any distribution F over $[-1, 1]$ as plausible and evaluates each strategy according to its worst-case scenario regarding F.[7] To be precise, let $\mathcal{F}$ denote the set of all distributions over $[-1, 1]$. The seller's payoff from a strategy $\sigma \in \Sigma$ is given by

$$\inf_{F \in \mathcal{F}} \pi(\sigma; F) = \inf_{F \in \mathcal{F}} \mathbb{E}_\sigma[\pi(\ell, p; F)].$$

Her problem is to find a σ that maximizes $\inf_{F \in \mathcal{F}} \pi(\sigma; F)$:

$$\pi^* := \sup_{\sigma \in \Sigma} \inf_{F \in \mathcal{F}} \pi(\sigma; F) = \sup_{\sigma \in \Sigma} \inf_{F \in \mathcal{F}} \mathbb{E}_\sigma[\pi(\ell, p; F)]. \tag{1}$$

Following the literature on robust pricing and mechanism design (e.g., Carrasco et al., 2018; Bergemann et al., 2019; Hinnosaar and Kawai, 2020), we refer to π^* as the seller's *profit guarantee*.

[7]The assumption that the seller knows the disutility function $c(\cdot)$ but not the distribution of tastes is in line with the common assumption in the robust pricing literature, namely, that the seller knows buyers' risk preferences (usually risk-neutral) but not their willingness to pay (e.g., Bergemann and Schlag, 2011; Du, 2018).

The dual problem. To identify π^*, we begin by studying the dual problem:

$$\pi_* := \inf_{F\in\mathcal{F}} \sup_{\sigma\in\Sigma} \pi(\sigma; F) = \inf_{F\in\mathcal{F}} \sup_{\sigma\in\Sigma} \mathbb{E}_\sigma[\pi(\ell, p; F)]. \tag{2}$$

In the dual problem nature chooses a distribution of tastes so as to minimize the seller's profit, the seller observes the distribution and designs and prices the product. Besides being a stepping stone in solving the robustness problem, the dual problem is of economic interest on its own: It establishes the lowest profit the seller can be held to in a Bayesian setting. For this reason, we refer to the solution to (2) as a *seller-worst distribution* and π_* as the seller's *lowest profit*.[8]

Discussion. The feature that starkly distinguishes our model from the standard monopoly model is the lack of natural rankings over locations. In the standard monopoly model, malevolent nature can assign all the probability mass to the lowest valuation, in which case a pessimistic seller with no information is powerless and obtains the minimal payoff. The seller can secure a payoff above the lowest valuation only if she is endowed with some information about the distribution. Bergemann and Schlag (2011), for example, consider a seller who knows the neighborhood the distribution belongs to, while Carrasco et al. (2018) bestow the seller with knowledge of some moments of the distribution. In our model where the seller's location is endogenous, no particular location corresponds to the lowest valuation. This limits the extent to which nature can reduce the seller's profit.

3 Seller-Worst Distributions and Payoffs

This section analyzes the dual problem (2) and determines π_*. To understand the working of our model, as well as facilitate the subsequent analysis, we first examine two benchmarks.

3.1 Benchmarks

The first benchmark characterizes the seller's optimal strategy when she is facing the uniform distribution of tastes. The second benchmark explores the worst distribution for the seller when the design of the product is fixed.

Uniform distribution. Consider the case where consumers are uniformly distributed over $[-1, 1]$ with constant density $1/2$; $F = U[-1, 1]$. Under the distribution, $\ell = 0$ is *an* optimal location for

[8] More precisely, F is a seller-worst distribution if there exists a $\hat{\sigma}$ such that $\pi(\hat{\sigma}, F) = \pi_*$.

the seller. Hence,

$$\pi(0,p;F) = pD(0,p;F) = p(F(\Delta(p)) - F(-\Delta(p))) = p\Delta(p),$$

for all p, and the seller's problem can be rewritten in terms of Δ as

$$\pi^U := \max_{\Delta\in[0,1]} p(\Delta)\Delta = \max_{\Delta\in[0,1]} \Delta(1-c(\Delta)),$$

leading to the following result.

Lemma 1 *If F is uniform over $[-1,1]$, then it is optimal for the seller to choose $\ell = 0$ and $p = 1 - c(\Delta^U)$, where*

$$\Delta^U := \arg\max_{\Delta\in[0,1]} \Delta(1-c(\Delta)).$$

Closed form solutions can be obtained for specific cost functions c. For instance, suppose $c(y) = ty$ for some $t > 0$. Then, $\Delta(1-c(\Delta)) = \Delta(1-t\Delta)$ is maximized at $\Delta = 1/(2t)$, implying that the seller optimally covers the whole market if $2t \geq 1$ and only $[-1/(2t), 1/(2t)]$ otherwise. Therefore, the seller's optimal strategy and profit are given by

$$\Delta^U = \min\left\{\frac{1}{2t}, 1\right\} \text{ and } \pi(U[-1,1]) = \begin{cases} \frac{1}{4t} & \text{if } t > \frac{1}{2} \\ 1-t & \text{if } t \leq \frac{1}{2}. \end{cases}$$

Seller-worst distributions given $\ell = 0$. Consider the case where the product design is exogenously fixed at $\ell = 0$.[9] The further away consumers are from the seller, the lower is their willingness to pay for the given design. Therefore, any distribution that puts all probability mass on -1 or 1 minimizes the seller's profit. The seller serves those customers if the price necessary to cover them, $1 - c(1)$, is non-negative; the resulting profit is $\max\{1 - c(1), 0\}$.

Due to a high concentration of consumers at -1 and 1, the above binary distributions might not minimize the seller's profit if she can change the design of the product. Given a binary distribution on -1 and 1, the seller's optimal strategy would be to either locate at the centre and serve all consumers (i.e., $(\ell, p) = (0, 1 - c(1))$) or to locate at -1 or 1, and extract full surplus from consumers at that location (i.e., $(\ell, p) = (-1, 1)$ or $(\ell, p) = (1, 1)$). In the latter case, the seller is worst off if consumers are equally divided between -1 and 1. Therefore, among the binary

[9] The subsequent analysis can be modified in a straightforward fashion for the case where $\ell \neq 0$. It requires heavy investment into notation and exposition, but does not produce any qualitatively different results.

distributions with support $\{-1, 1\}$, the seller-worst distribution (when she can design the product) is the symmetric one, and the resulting profit is

$$\underline{\pi}^0 := \max\left\{\frac{1}{2}, 1 - c(1)\right\}.$$

The seller optimally locates in the middle and covers the whole market if $c(1) < 1/2$ and chooses one of the corners otherwise.

Notice that $\underline{\pi}^0$ may be strictly *larger* than the seller's profit under $U[-1, 1]$. Specifically, if $c(y) = ty$ and $t > \frac{1}{2}$ then

$$\pi(U[-1, 1]) = \frac{1}{4t} < \underline{\pi}^0 = \frac{1}{2}.$$

Thus, a seller-worst distribution given $\ell = 0$ is not necessarily a seller-worst distribution once the seller can choose a design.

The above suggests that a seller-worst distribution should position consumers as far away from the seller as possible while inducing her to choose the given design. An intuitive way to limit demand irrespective of the design the seller chooses is to spread consumers evenly across $[-1, 1]$, thereby guaranteeing that there are never too many consumers nearby. Indeed, the uniform distribution will play a prominent part in our analysis.

3.2 A Profit Lower Bound

We start by deriving a lower bound on π_*. A simple lower bound is achieved by the strategy where the seller positions at the center and offers a price that covers the whole market. This strategy yields the profit $1 - c(1)$ irrespective of the distribution of tastes. The seller is guaranteed to sell to all consumers but at the cost of selling at a low price.[10] The fact that the seller can observe the distribution allows her to refine the strategy and target consumers in either $[-1, 0]$ or $[0, 1]$, depending on which of the two intervals is more populous. This is achieved by positioning at either $-1/2$ or $1/2$ and charging the price $1 - c(1/2)$. One of the two intervals must have at least half of the consumers, thus the seller is guaranteed a profit of at least $1/2(1 - c(1/2))$.

The above strategies can be extended to an arbitrary integer n: Partition the interval $[-1, 1]$ into n subintervals of equal width and serve only those consumers on the most densely-populated subinterval (by positioning at its center and charging $p = 1 - c(1/n)$). Since total consumer mass is 1, some subiniterval must be inhabited by at least $1/n$ of consumers. Therefore, the strategy

[10]Indeed, $1-c(1)$ could be negative. If the consumers' disutility from the object is too large, the far away consumers may need to be subsidized to consume the product.

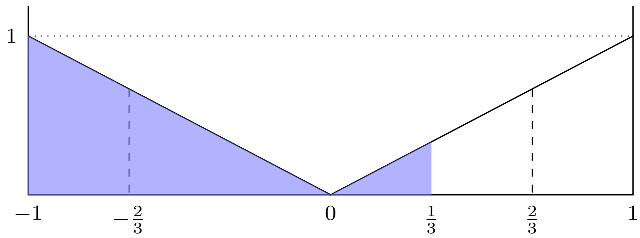

Figure 2: The density function for Example 1.

ensures profit of at least $1/n(1-c(1/n))$. This lower bound holds for any $n \in \mathbb{N}$ and *regardless of* F, leading to the following result.

Lemma 2 *The seller's lowest profit π_* must be at least $\underline{\pi}$, where*

$$\underline{\pi} := \sup_{\Delta \in \left\{1, \frac{1}{2}, \frac{1}{3} \dots\right\}} \Delta(1 - c(\Delta)).$$

Restricting attention to an integer number of intervals—equivalently, requiring that $1/\Delta$ is an integer—might seem arbitrary. A natural method to extend the technique would be to fix an arbitrary reach $\Delta \leq 1$ and serve the interval $[\ell - \Delta, \ell + \Delta]$ with the most customers. One could think that some such interval contains at least Δ consumers. Yet, this is not the case unless $1/\Delta$ is an integer. The following example demonstrates the idea.

Example 1 *Consider F with density $f(x) = |x|$, depicted in Figure 2. If $\Delta = 2/3$ then for any $\ell \in [-1/3, 1/3]$,*

$$F(\ell + \Delta) - F(\ell - \Delta) \leq F(1/3) - F(-1) = \frac{5}{9} < \Delta = \frac{2}{3},$$

where the first inequality holds because f is increasing towards the two edges and so $F(\ell + \Delta) - F(\ell - \Delta)$ rises as ℓ moves in either direction; for example, as shown in Figure 2, $F(1/3) - F(-1) > F(2/3) - F(-2/3)$. This means that any interval of length $4/3$ (covering $2/3$ of the whole interval) has at most $5/9$ consumers. Note that there is nothing particular about $\Delta = 2/3$; the same argument extends to any $\Delta \in (1/2, 1)$.

If Δ can take any value in $[0, 1]$, then the resulting profit corresponds to π^U, that is,

$$\pi^U = \max_{\Delta \in [0,1]} \Delta \left(1 - c(\Delta)\right). \tag{3}$$

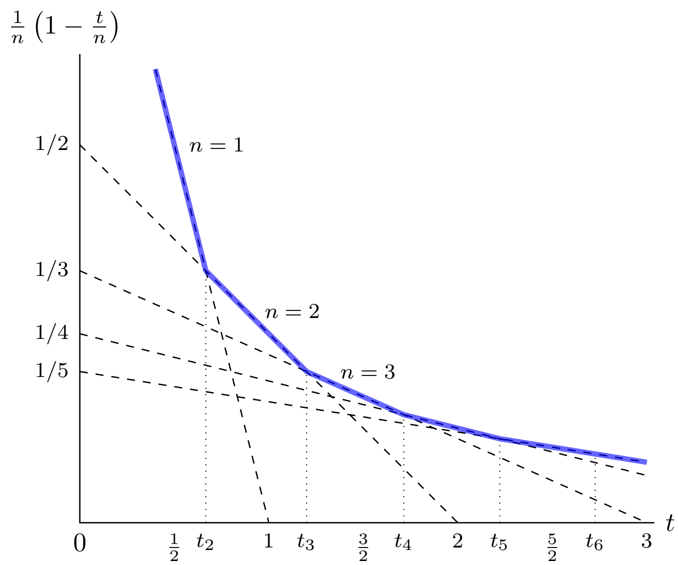

Figure 3: The optimal value of n that yields $\underline{\pi}$ depending on t in the linear case where $c(y) = ty$. Each dashed line depicts $\frac{1}{n}\left(1 - \frac{t}{n}\right)$ for some n, and the blue translucent curve shows its upper envelope, which coincides with $\underline{\pi}$.

Together with concavity of $\Delta(1-c(\Delta))$, the relationship between $\underline{\pi}$ and π^U leads to the subsequent result.

Lemma 3 *The following statements hold:*

(a) If $\Delta^U = 1/n$ *for some* $n \in \mathbb{N}$ *then* $\underline{\pi} = \pi_* = \pi^U$.

(b) Let $\widehat{n} = \lceil 1/\Delta^U \rceil$.[11] *Then,*

$$\underline{\pi} = \max\left\{ \frac{1}{\widehat{n}-1}\left(1 - c\left(\frac{1}{\widehat{n}-1}\right)\right), \frac{1}{\widehat{n}}\left(1 - c\left(\frac{1}{\widehat{n}}\right)\right) \right\}.$$

Proof. Part (a) holds because if $\Delta^U = 1/n$ then $\underline{\pi} \leq \pi_* \leq \pi^U = \underline{\pi}$. Part (b) is because, by the definition of $\widehat{n}$ and concavity of $\Delta(1 - c(\Delta))$, $1/n(1 - c(1/n))$ is strictly increasing if $n < \widehat{n} - 1$ and decreasing if $n > \widehat{n}$. ■

Linear Disutility. Suppose $c(y) = ty$ for some $t > 0$. Then, $\Delta^U = \min\{1/(2t), 1\}$. For any t, $\widehat{n}$ in Lemma 3.(b) is the smallest integer such that $\widehat{n} \geq 2t$. Let t_n be the value of t such that the

[11] Where $\lceil x \rceil$ is the smallest integer greater or equal to x.

seller is indifferent between $\Delta = 1/(n-1)$ and $\Delta = 1/n$:

$$t_n = \frac{n(n-1)}{2n-1} \in \left(\frac{n-1}{2}, \frac{n}{2}\right).$$

As a consequence, if $t \in [t_n, t_{n+1}]$ then

$$\underline{\pi} = \frac{1}{n}\left(1 - c\left(\frac{1}{n}\right)\right).$$

3.3 Achieving the Profit Lower Bound

The main result of this section is that $\underline{\pi}$ in Lemma 2 is the tight lower bound for the seller's profit.

Theorem 1 *There exists a distribution F such that $\pi(F) = \underline{\pi}$. Therefore,*

$$\pi_* = \underline{\pi} = \sup_{\Delta \in \left\{1, \frac{1}{2}, \frac{1}{3} \ldots\right\}} \Delta(1 - c(\Delta)). \tag{4}$$

We explicitly construct a distribution F that delivers Theorem 1.[12] Suppose $\underline{\pi} = \frac{1}{n}\left(1 - c\left(\frac{1}{n}\right)\right)$, for some $n \in \mathbb{N}$. A *necessary* condition for F is that it assigns probability mass $1/n$ to all $\left[-1, -1 + \frac{2}{n}\right]$,...,$\left[1 - \frac{2}{n}, 1\right]$:

$$F\left(-1 + \frac{2k}{n}\right) - F\left(-1 + \frac{2(k-1)}{n}\right) = \frac{1}{n}, \tag{5}$$

for every $k = 1, ..., n$. If this condition did not hold, the seller would obtain strictly more than $\underline{\pi}$ from the most densely-populated subinterval.

Another property that the distribution needs to satisfy is that for each subinterval, positioning at its center and charging the price $1 - c(1/n)$ must be locally optimal for the seller; global optimality is addressed later. Consider the lowest subinterval $\left[-1, -1 + \frac{2}{n}\right]$. For the seller to obtain (no more than) $\underline{\pi}$, $1/n$ should be the reach maximizing $(1 - c(\Delta))F(-1 + 2\Delta)$.[13] Let f denote the density function of F. Evaluating the first-order condition at $\Delta = 1/n$ and invoking (5), yields

$$f\left(-1 + \frac{2}{n}\right) = \frac{1}{2n}\frac{c'(1/n)}{1 - c(1/n)}.$$

[12] The construction implies that there may be multiple seller-worst distributions. The distribution we present below has a particularly simple structure.

[13] Here, the seller is maximizing over price/location pairs with $\ell = -1 + \Delta$ and $p = 1 - c(\Delta)$.

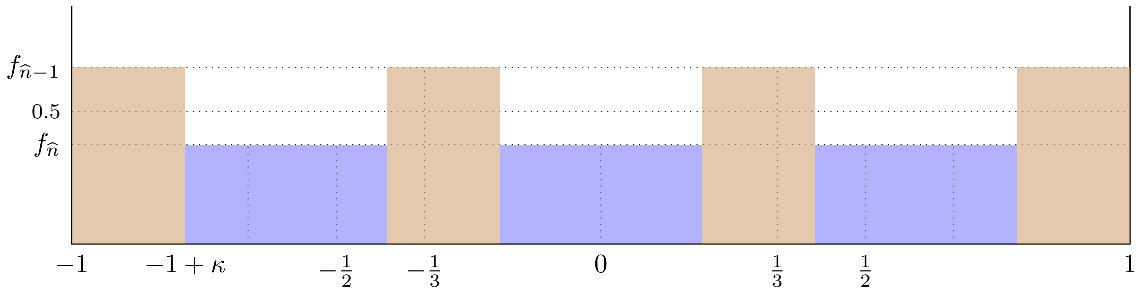

Figure 4: The density of distribution F_* defined in Definition 1. The cost function used for this figure is $c(y) = \frac{12}{7}y$. In this case, $\widehat{n} = 4$, and $\pi_* = \frac{1}{3}\left(1 - c\left(\frac{1}{3}\right)\right) = \frac{1}{4}\left(1 - c\left(\frac{1}{4}\right)\right)$.

Applying the same argument to all other subintervals results in:

$$f\left(-1 + \frac{2}{n}\right) = ... = f\left(1 - \frac{2}{n}\right) = \frac{1}{2n}\frac{c'(1/n)}{1 - c(1/n)}. \tag{6}$$

The following distribution combines Lemma 3.(b) with the two necessary conditions, (5) and (6), in a simple manner.

Definition 1 *Let $\widehat{n}$ be the value defined in Lemma 3 and $f_n := \frac{1}{2n}\frac{c'(1/n)}{1-c(1/n)}$ for each $n \in \mathbb{N}$. We define F_* to be a piecewise linear distribution function with density*

$$f_*(x) = \begin{cases} f_{\widehat{n}-1} & \textit{if } x \in \left[-1 + \frac{2(k-1)}{\widehat{n}-1} - \kappa\frac{k-1}{\widehat{n}-1}, -1 + \frac{2(k-1)}{\widehat{n}-1} + \kappa\frac{n-k}{\widehat{n}-1}\right) \\ f_{\widehat{n}} & \textit{otherwise,} \end{cases}$$

where $\kappa := \frac{1-2f_{\widehat{n}}}{\widehat{n}(f_{\widehat{n}-1}-f_{\widehat{n}})}$.[14]

Figure 4 shows a representative structure of the density f_*. It alternates between density levels $f_{\widehat{n}-1}$ (high) and $f_{\widehat{n}}$ (low). By construction, f coincides with $f_{\widehat{n}}$ around $-1 + \frac{2}{\widehat{n}}, ..., 1 - \frac{2}{\widehat{n}}$ and with $f_{\widehat{n}-1}$ around $-1 + \frac{2}{\widehat{n}-1}, ..., 1 - \frac{2}{\widehat{n}-1}$, thus ensuring that f_* satisfies the necessary condition (6), whether the optimal n is $\widehat{n}-1$ or $\widehat{n}$. In addition, the lengths of the subintervals are chosen so that F assigns probability $1/\widehat{n}$ to all intervals $\left[-1, -1 + \frac{2}{\widehat{n}}\right], ..., \left[1 - \frac{2}{\widehat{n}}, 1\right]$ and $1/(\widehat{n}-1)$ to all intervals $\left[-1, -1 + \frac{2}{\widehat{n}-1}\right], ..., \left[1 - \frac{2}{\widehat{n}-1}, 1\right]$; the properties guarantee condition (5), whether the optimal n is $\widehat{n} - 1$ or $\widehat{n}$. Note that κ is the width of each high-density interval.

It remains to show that the seller has no other profitable deviations, that is, given F_*, the

[14] The value κ is such that $\widehat{n}\kappa f_{\widehat{n}-1} + (2 - \widehat{n}\kappa)f_{\widehat{n}} = 1$. In other words, κ is defined to be the common width of high-density regions that makes total probability equal to 1. For any convex $c(\cdot)$, f_n is strictly decreasing in n, and $\kappa \in [0, 1/(2\widehat{n})]$.

seller's maximized profit coincides with $\underline{\pi}$. By construction, the seller can achieve $\underline{\pi}$ by serving, for example, $\left[-1, -1+\frac{2}{\widehat{n}}\right]$ or $\left[-1, -1+\frac{2}{\widehat{n}-1}\right]$. To show that the seller's profit *cannot* exceed π_*, first consider the profit maximization problem $\max_\Delta \widehat{\pi}(\Delta) := (1-c(\Delta))F_*(-1+2\Delta)$ where the seller maximizes profit over $(\ell, p) = (-1+\Delta, 1-c(\Delta))$ (i.e., by choosing the interval $[-1, -1+2\Delta]$). By definition of F^*:

$$\widehat{\pi}'\left(\frac{1}{\widehat{n}}\right) = \widehat{\pi}'\left(\frac{1}{\widehat{n}-1}\right) = 0.$$

In the proof of Theorem 1, we show that $\frac{1}{\widehat{n}}$ and $\frac{1}{\widehat{n}-1}$ are the only local maximizers of $\widehat{\pi}(\Delta)$, which suggests that $\pi_* = \max_\Delta \widehat{\pi}(\Delta)$. Given this, it suffices to show that the seller cannot increase the profit by moving her location ℓ. Indeed, we show that for any Δ, the payoff from locating at $-1+\Delta$ is at least as large as that from locating anywhere else: $\pi(-1+\Delta, \Delta) \geq \pi(\ell, \Delta)$, for every Δ and ℓ.[15] The last inequality holds because, by its construction, f_* is periodic and has the higher density $f_{\widehat{n}-1}$ over $[-1, -1+\kappa]$, and thus if ℓ increases from $-1+\Delta$ then $F_*(\ell+\Delta) - F_*(\ell-\Delta)$ may become smaller, but cannot exceed, $F_*(-1+2\Delta)$; see the proof in the appendix for the details of the argument.

Uniform vs Seller-Worst Distributions. The only difference between maximization when the seller is facing the uniform distribution, (3), and when the seller is facing a worst distribution, (4), is that the first is an integer version of the second. Lemma 3 established that the uniform distribution is a seller-worst distribution whenever $\Delta^U = 1/n$ for some $n \in \mathbb{N}$.[16] The following result suggests that the uniform distribution generally gives low profits to the seller when the disutility function is linear.

Proposition 1 *Suppose $c(y) = ty$ for some $t > 0$. Then $\pi^U/\pi_* \leq 9/8$.*

Proof. See the appendix. ∎

When costs are linear, the seller's profit under the uniform distribution is at most 12.5% above her lowest profit. In fact, if $t > 1$ then the ratio reduces to $25/24 \approx 1.0417$, so the maximum difference becomes around 4%. The standard practice of assuming the uniform distribution in the Hotelling environment imposes fairly low profits on the monopolist.[17]

[15] The same inequality holds if $\pi(-1+\Delta, \Delta)$ is replaced with any $\pi(-1+(2k-1)\Delta, \Delta)$, for $k \in \{1, ..., n^*\}$.

[16] While it may seem that generically the uniform distribution is not a worst distribution, this is not quite the case, due to the boundary solution $\Delta_u = 1$. The latter can obtain for a non-trivial set of cost functions.

[17] The conclusion is even stronger in the circular city model where the uniform distribution is *always* a seller-worst distribution; see Section 6.

4 Robust Pricing and Design

This section returns to the primal problem (1) and determines the seller's profit guarantee π^* as well as the corresponding optimal strategy σ^*.

Observing the malevolent nature's choice of distribution of tastes before making the strategic choice can not make the seller worse off than having to choose the strategy knowing that nature will respond maliciously. This is but a restatement of the max-min inequality (see, e.g., Osborne and Rubinstein, 1994):

$$\sup_{\sigma\in\Sigma}\inf_{F\in\mathcal{F}}\pi(\sigma;F)\leq\inf_{F\in\mathcal{F}}\sup_{\sigma\in\Sigma}\pi(\sigma;F).$$

The seller's lowest profit characterized in Section 3, therefore, provides an upper bound for the seller's profit guarantee (i.e., $\pi_*\geq\pi^*$).

The following result establishes the converse inequality (thus strong duality) and, moreover, pinpoints the seller's robustly optimal strategy.

Theorem 2 *The following holds:*

$$\pi^*=\pi_*=\underline{\pi}.$$

In addition, let n^* be a value such that $\pi_*=1/n^*(1-c(1/n^*))$. The seller's optimal strategy is to set $p=1-c(1/n^*)$ and uniformly randomizes her location ℓ over $\left\{-1+\frac{1}{n^*},-1+\frac{3}{n^*},...,1-\frac{1}{n^*}\right\}$.

Proof. As argued above, the inequality $\pi^*\leq\pi_*$ holds invariably. Therefore, it suffices to show that the indicated strategy, denoted by σ^*, yields at least π_* regardless of $F\in\mathcal{F}$. This implies that $\pi^*\geq\pi_*$ and so $\pi^*=\pi_*$.

The seller's expected profit under σ^* and $F\in\mathcal{F}$ is

$$\begin{aligned}\pi(\sigma^*;F)&=\sum_{k=1}^{n^*}\frac{1}{n^*}\left(F\left(-1+\frac{2k}{n^*}\right)-F_-\left(-1+\frac{2(k-1)}{n^*}\right)\right)\left(1-c\left(\frac{1}{n^*}\right)\right)\\&=\frac{1}{n^*}\left(1-c\left(\frac{1}{n^*}\right)\right)\sum_{k=1}^{n^*}\left(F\left(-1+\frac{2k}{n^*}\right)-F_-\left(-1+\frac{2(k-1)}{n^*}\right)\right)\\&\geq\frac{1}{n^*}\left(1-c\left(\frac{1}{n^*}\right)\right)=\pi_*,\end{aligned}$$

where the inequality holds with equality whenever F has no atom on $\{-1+2/n^*,...,1-2/n^*\}$. Since this holds for any $F\in\mathcal{F}$, we arrive at the desired inequality $\pi^*\geq\pi_*$. ■

The seller divides $[-1, 1]$ into n^* subintervals of equal length and serves a randomly selected subinterval. A crucial property of this strategy is to render nature indifferent over all degenerate distributions (with exceptions of those on the borderline points, $\{-1 + 2/n^*, ..., 1 - 2/n^*\}$) and thus almost all distributions.[18] To see why this property is essential, notice that given seller's strategy, nature's best response is to choose a degenerate (Dirac) distribution that assigns all the mass to the location that keeps the seller to the lowest profit. A strategy that equates the profit across degenerate distributions insures the seller against nature's malevolence and makes it inconsequential whether the seller first observes nature's choice or not. Meanwhile, the distribution F_* given in Definition 1 (and used to prove Theorem 1) renders the seller's choice of design in $\left\{-1 + \frac{1}{n^*}, -1 + \frac{3}{n^*}, ..., 1 - \frac{1}{n^*}\right\}$ and the price $1 - c(1/n^*)$ optimal. The strategy and the distribution, therefore, form a saddle point.

Two other properties of the seller's optimal strategy are worth noting. First, it divides the relevant space $[-1, 1]$ into n^* *disjoint* intervals. Each location therefore belongs to only one "submarket" and is served only when the submarket is chosen (with probability $1/n^*$). If an interval were to belong to multiple submarkets, nature would simply avoid assigning consumers to that interval. This, in turn, would eradicate the seller's incentive to create such an overlap. Second, all submarkets have the same size (length), or equivalently, all the consumers are offered the same price.

Full market coverage. Under severe uncertainty about the distribution of tastes, the "safest" strategy the seller can take is to locate at the center $\ell = 0$ and charge a sufficiently low price $1-c(1)$ that can be accepted by all consumers. Theorem 2 shows that such a full coverage strategy can be indeed optimal, but not always. The following result provides a necessary and sufficient condition for its optimality.

Corollary 1 *Choosing the design $\ell = 0$ and price $1 - c(1)$ is optimal for the seller if and only if $c(1) \leq 1/2 + 1/2c(1/2)$.*

Proof. The result follows from Theorem 2 and concavity of $\Delta(1 - c(\Delta))$. The former shows that equally dividing $[-1, 1]$ into n^* subintervals for some $n^* \in \mathbb{N}$ is optimal, while the latter ensures that $1/n(1 - c(1/n))$ is decreasing in $n \geq n^*$. $n^* = 1$ is thus equivalent to requiring that $1 - c(1) \geq 1/2(1 - c(1/2))$. ∎

To understand this result better, consider the case where $c(y) = ty^\beta$ for some $t > 0$ and $\beta \geq 1$.

[18] Mass points on the borderline points benefit the seller.

Corollary 1 implies that the full coverage strategy is optimal if and only if

$$t \leq \underline{t}(\beta) := \frac{1}{2 - 1/2^{\beta}}.$$

It is intuitive that the strategy is optimal when t is sufficiently small; the seller's profit from the strategy, $1 - t$, approximates 0 if t tends to 1 and 1—the maximal possible value—if t tends to 0. The cutoff $\underline{t}(\beta)$ is decreasing in β. This is because as β rises, c becomes more convex, in which case the opportunity cost of using the strategy, $1/2(1 - c(1/2))$, rises.

5 Optimal Product Line Design

Thus far, we have studied the canonical environment where the seller designs and produces one product. Unless the condition in Corollary 1 holds, this setup requires the seller to employ uniform randomization across multiple locations to ensure comprehensive coverage of the entire space $[-1, 1]$, while avoiding excessive price reductions. An alternative response of the seller could be to diversify and produce multiple products, thereby directly broadening her market reach. In this section, we demonstrate how the above-developed techniques can be used to study the robustness problem of a seller who can design and sell several varieties of a product.[19]

Setup. Let $m \in \mathbb{N}$ denote the number of varieties of the product the seller can produce; our main model is the special case where $m = 1$. For each product variety $i = 1, ..., m$, the seller chooses its design $\ell_i \in [-1, 1]$ and price $p_i \in \mathbb{R}_+$; the seller's pure strategy is $(\ell_1, p_1; ...; \ell_m, p_m) \in ([-1, 1] \times \mathbb{R}_+)^m$. We use σ_m to represent the seller's mixed strategy and Σ_m the set of all mixed strategies. The seller's profit guarantee is

$$\pi_m^* := \sup_{\sigma_m \in \Sigma_m} \inf_F \pi(\sigma_m; F),$$

where $\pi(\sigma_m; F)$ is the seller's expected profit from the strategy σ_m when consumers' tastes are distributed according to the distribution F.

Dual problem. We first consider the dual inf-sup problem in which the seller moves after having observed nature's choice $F \in \mathcal{F}$. For any $F \in \mathcal{F}$, the seller can divide $[-1, 1]$ into $n(\geq m)$ equal-length regions and serve the best m of them. Therefore, her lowest profit, π_{*m}, cannot be

[19]Product line design has been studied in the context of oligopoly. Schmalensee (1978) introduces the idea that firms can use brand proliferation as a deterrent to entry. See Dasci and Laporte (2005) and Janssen et al. (2005) for more recent studies on competition among multi-product firms in spatial models. For a monopoly model incorporating product line design, but no strategic pricing, see Ke et al. (2022).

smaller than

$$\underline{\pi}_m := \max_{n \geq m} \frac{m}{n} \left(1 - c\left(\frac{1}{n}\right)\right). \tag{7}$$

Suppose $m \leq n^*$ and consider the distribution F_* in Definition 1. Due to Theorem 1, the seller's profit from a single product cannot exceed $1/n^*(1 - c(1/n^*))$. The seller's overall profit, therefore, cannot exceed m times the single-product profit: $\pi(F_*) \leq m/n^*(1 - c(1/n^*))$. Together with $\pi_* \leq \pi(F_*)$, this implies

$$\underline{\pi}_m \leq \pi_{*m} \leq \pi(F_*) \leq \frac{m}{n^*} \left(1 - c\left(\frac{1}{n^*}\right)\right) \leq \underline{\pi}_m.$$

Next, suppose $m > n^*$ and consider the uniform distribution over $[-1, 1]$. The seller obtains $\underline{\pi}_m = 1 - c(1/m)$ if she equally divides $[-1, 1]$ into m regions and serves all consumers with the price $1 - c(1/m)$. The following Lemma 4 argues that the seller cannot obtain strictly more than $\underline{\pi}_m$, so it is the seller's maximized profit under the uniform distribution. Consequently, $\pi(U[-1, 1]) = \underline{\pi}_m$, and thus

$$\underline{\pi}_m \leq \pi_{*m} \leq \pi(U[-1, 1]) = \underline{\pi}_m.$$

Lemma 4 *If the seller can produce m varieties of a product for some $m > n^*$ and consumers' tastes are uniformly distributed over $[-1, 1]$ then the seller's optimal strategy is to set $p_i = 1 - c(1/m)$ and $\ell_i = -1 + (2i - 1)/m$ for all $i = 1, ..., m$.*

Proof. See the appendix. ■

This result holds because $\Delta(1 - c(\Delta))$ is concave and strictly increasing while $\Delta \leq 1/m \leq 1/(n^* + 1)$; the former (concavity) implies that all submarkets must have the same size (reach), while the latter suggests that it is more profitable for the seller to serve the whole market (by setting $\Delta = 1/m$ in each submarket) than to lose some consumers in order to charge a higher price (by setting $\Delta < 1/m$).

Primal problem. The above analysis of the dual problem enables us to apply a saddle-point characterization to the primal problem, just as in Section 4. Since the inequality $\sup_{\sigma_m \in \Sigma_m} \inf_F \pi(\sigma_m; F) \leq \inf_F \sup_{\sigma_m \in \Sigma_m} \pi(\sigma_m; F)$ always holds, we have

$$\pi_m^* \leq \pi_{*m} = \underline{\pi}_m.$$

If $m \leq n^*$, the seller can divide $[-1,1]$ into n^* subintervals of equal length and choose m of them to serve at random (uniformly). This strategy guarantees the payoff of at least $\underline{\pi}_m$ regardless of the distribution of tastes.[20] It follows that $\pi^*_m \geq \underline{\pi}_m$, which together with the above inequality yields the strong duality: $\pi^*_m = \pi_{*m} = \underline{\pi}_m$. If $m > n^*$, the same logic holds except that the seller divides the interval into m subintervals of equal length and serves all of them; one with each variety of the product.

Theorem 3 *If the seller can produce m different varieties of the product, then her profit guarantee is*

$$\pi^*_m = \pi_{*m} = \begin{cases} \frac{m}{n^*}\left(1 - c\left(\frac{1}{n^*}\right)\right), & \text{if } m \leq n^* \\ 1 - c\left(\frac{1}{m}\right)), & \text{if } m > n^*. \end{cases}$$

Let $n^\dagger := \max\{m, n^*\}$. The seller's optimal strategy is to divide $[-1,1]$ into $n^\dagger$ subintervals of equal length and serve m randomly selected subinterrvals.

When $m = 1$, the seller divides the market into n^* subintervals of equal length and serves a randomly chosen one. As m increases but remains below n^*, the seller continues to partition the market into the same number of submarkets but serves a greater number of them. The fact that all the submarkets are of the same size and served at the same price implies that the marginal benefit of adding a variety is constant at $1/n^*(1-c(1/n^*))$. Once the seller fully covers the market ($m = n^*$), increasing m enables the seller to extract more surplus by creating more varieties and charging a higher price for each of them. The resulting profit is $1 - c(1/m)$, which is concave in m. Producing more varieties raises the seller's profit, but its marginal returns are decreasing. As depicted in Figure 5, π^*_m is linear in m below n^* and concave above n^*.

Optimal number of varieties. Theorem 3 presents the seller's optimal strategy when she produces m varieties of the product. The optimal number of varieties can easily be determined if one is given a convex cost function of producing varieties, $\alpha(\cdot)$. The optimal number of varieties maximizies the seller's profit $\pi^*(m) - \alpha(m)$ over $m \in \mathbb{N}$.

6 Circular City

We explore how our results extend to the circular city model by Salop (1979), where consumers are distributed on a circle. To make the model directly comparable to our main model, we normalize

[20]This is a generalization of the argument in the proof of Theorem 2.

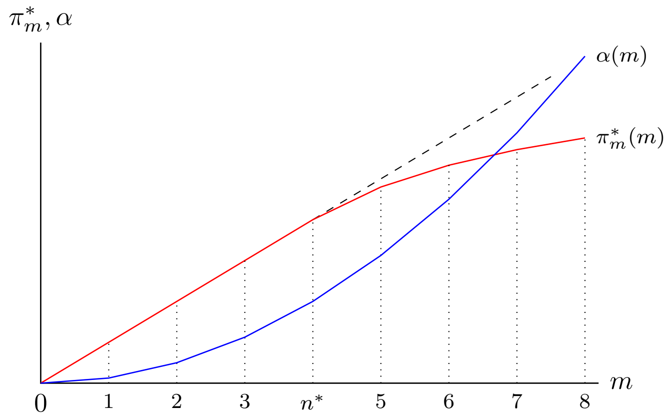

Figure 5: The red curve shows the seller's profit guarantee π_m^* as a function of m for the case when $c(y) = 2y$ (in which case $n^* = 4$). The blue curve depicts an example of the convex cost function $\alpha(m)$.

the circumference of the circle to 2. Although we find this model less appropriate to address product-design questions, it provides additional insights into the results obtained for the linear city model.[21]

Recall that in the dual problem of Section 3.3, the seller's lowest profit π_* is given by

$$\pi_* = \sup_{\Delta \in \left\{1, \frac{1}{2}, \frac{1}{3} \ldots\right\}} \Delta(1 - c(\Delta)),$$

which is necessarily smaller than her profit under the uniform distribution

$$\pi^U = \max_{\Delta \in [0,1]} \Delta \left(1 - c(\Delta)\right).$$

The difference between the two profits arises because given any distribution F, the seller can always secure at least $1/n(1 - c(1/n))$ for some $n \in \mathbb{N}$ but, as illustrated in Section 3.2, she cannot guarantee $\Delta(1 - c(\Delta))$ for Δ that is not of the form $1/n$ for some $n \in \mathbb{N}$. Unlike in the linear city model, the same is not true in the circular city model.

Lemma 5 *Given any probability measure μ over a circle with circumference 2, for any $\Delta \leq 1$, there exists ℓ such that $\mu\left([\ell - \Delta, \ell + \Delta]\right) \geq \Delta$.*

[21] Hotelling (1929) provides an example in which a position on an interval represents a particular level of sweetness of cider, with one end of the interval representing extremely sweet and the other completely bitter cider. More generally, the linear city model is the more appropriate model whenever there is a natural order over designs (locations). Bar-Isaac et al. (2023), however, propose an interesting model of product design where consumers are distributed on a circle, while firms choose a design inside of the circle.

Proof. If $\mu([\ell-\Delta,\ell+\Delta]) < \Delta$ for all $\ell \in [0,2]$ then the following contradiction emerges:

$$\begin{aligned}\int_0^2 \mu([\ell-\Delta,\ell+\Delta])\,d\ell &< \int_0^2 \Delta d\ell \\ &= 2\Delta,\end{aligned}$$

but also

$$\begin{aligned}\int_0^2 \mu([\ell-\Delta,\ell+\Delta])\,d\ell &= \int_0^2 \left[\int_{-\Delta}^{\Delta} d\mu(\ell+x)\right] d\ell \\ &= \int_{-\Delta}^{\Delta} \left[\int_0^2 d\mu(\ell+x)d\ell\right] dx \\ &= \int_{-\Delta}^{\Delta} dx \\ &= 2\Delta,\end{aligned}$$

where the second equality is through changing the order of integration and the third one holds because $\int_0^2 d\mu(\ell+x)$ is the sum of probability over the circle, which is 1 regardless of the value x. ■

Lemma 5 implies that the seller's lowest profit satisfies

$$\pi_* \geq \pi^U = \max_{\Delta\in[0,1]} \Delta\,(1-c(\Delta))\,.$$

But, since $\pi_* \leq \pi^U$ necessarily holds, we have $\pi_* = \pi^U$. Applying a duality argument analogous to the one in Section 4 yields the following result.

Proposition 2 *In the circular city, the uniform distribution is always a seller-worst distribution. When its circumference is equal to 2, then the seller's profit guarantee is equal to π^U, which she can achieve by uniformly randomizing her location over the circle and charging $p = 1 - c\left(\Delta^U\right)$ at every location.*

7 Conclusion

We study robust product design and pricing by a monopolist in the Hotelling framework. Specifically, we consider a seller who has no information about the distribution of consumers' tastes and wishes to maximize her worst-case payoff.

We show that the seller's optimal strategy is to partition the taste space into a finite number of subintervals of the same length and randomize uniformly over which of these subintervals to serve. The seller charges the same price, regardless of the subinterval selected. We also construct a distribution of tastes that, together with the seller's strategy, forms a saddle point. By implication, strong duality holds and, therefore, the seller's profit guarantee (sup-inf) coincides with the lowest profit a seller who first observes the choice of nature's distribution can be held to (inf-sup). We extend our analysis to the case where the seller can produce multiple varieties of the same product. Producing more varieties enables the seller to insure against nature more effectively by covering several subintervals simultaneously.

Information about consumers' tastes plays a pivotal role in product design, allowing sellers to create products that better resonate with consumers. We consider a fully agnostic seller who has no information about the distribution of consumers' tastes. Much is still to be explored regarding the broader impact of information on design and pricing, as well as the value of information in this context.

References

Auster, Sarah, "Robust contracting under common value uncertainty," *Theoretical Economics*, 2018, *13* (1), 175–204.

Bar-Isaac, Heski, Guillermo Caruana, and Vicente Cuñat, "Targeted product design," *American Economic Journal: Microeconomics*, 2023, *15* (2), 157–186.

Bergemann, Dirk and Karl Schlag, "Robust monopoly pricing," *Journal of Economic Theory*, 2011, *146* (6), 2527–2543.

_ , Benjamin Brooks, and Stephen Morris, "Revenue guarantee equivalence," *American Economic Review*, 2019, *109* (5), 1911–29.

Carrasco, Vinicius, Vitor Farinha Luz, Nenad Kos, Matthias Messner, Paulo Monteiro, and Humberto Moreira, "Optimal selling mechanisms under moment conditions," *Journal of Economic Theory*, 2018, *177*, 245–279.

Carroll, Gabriel, "Robustness and separation in multidimensional screening," *Econometrica*, 2017, *85* (2), 453–488.

Dasci, Abdullah and Gilbert Laporte, "A continuous model for multistore competitive location," *Operations Research*, 2005, *53* (2), 263–280.

Du, Songzi, "Robust mechanisms under common valuation," *Econometrica*, 2018, *86* (5), 1569–1588.

Hinnosaar, Toomas and Keiichi Kawai, "Robust pricing with refunds," *The RAND Journal of Economics*, 2020, *51* (4), 1014–1036.

Hotelling, Harold, "Stability in competition," *Economic Journal*, 1929, *39* (153), 41–57.

Janssen, Maarten CW, Vladimir A Karamychev, and Peran Van Reeven, "Multi-store competition: Market segmentation or interlacing?," *Regional Science and Urban Economics*, 2005, *35* (6), 700–714.

Jovanovic, Boyan, "Entry with private information," *The Bell Journal of Economics*, 1981, pp. 649–660.

Ke, T Tony, Jiwoong Shin, and Jungju Yu, "A Model of Product Portfolio Design: Guiding Consumer Search Through Brand Positioning," *Marketing Science*, 2022.

Kuksov, Dmitri, "Buyer search costs and endogenous product design," *Marketing Science*, 2004, *23* (4), 490–499.

Lauga, Dominique Olié and Elie Ofek, "Product positioning in a two-dimensional vertical differentiation model: The role of quality costs," *Marketing Science*, 2011, *30* (5), 903–923.

Meagher, Kieron J and Klaus G Zauner, "Product differentiation and location decisions under demand uncertainty," *Journal of Economic Theory*, 2004, *117* (2), 201–216.

Moorthy, K Sridhar, "Product and price competition in a duopoly," *Marketing science*, 1988, *7* (2), 141–168.

Orhun, A Yeşim, "Optimal product line design when consumers exhibit choice set-dependent preferences," *Marketing Science*, 2009, *28* (5), 868–886.

Osborne, Martin J and Ariel Rubinstein, *A course in game theory*, MIT press, 1994.

Salop, Steven C, "Monopolistic competition with outside goods," *The Bell Journal of Economics*, 1979, pp. 141–156.

Schmalensee, Richard, "Entry deterrence in the ready-to-eat breakfast cereal industry," *The Bell Journal of Economics*, 1978, pp. 305–327.

Sweeting, Andrew, "Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry," *Econometrica*, 2013, *81* (5), 1763–1803.

Villas-Boas, J Miguel, "Communication strategies and product line design," *Marketing Science*, 2004, *23* (3), 304–316.

_ , "A dynamic model of repositioning," *Marketing Science*, 2018, *37* (2), 279–293.

Appendix: Omitted Proofs

Proof of Theorem 1. Let F_* be as in Definition 1. We first show that for any $p = 1 - c(\Delta)$ and $\ell \in [-1+\Delta, 1-\Delta]$,

$$\pi(\ell, p; F_*) \leq \pi(-1+\Delta, p; F_*) =: \widehat{\pi}(\Delta),$$

which is equivalent to

$$\begin{aligned} D(\ell, p; F_*) &= F_*(\ell+\Delta) - F_*(\ell-\Delta) \\ &\leq D(-1+\Delta, p; F_*) \\ &= F_*(-1+2\Delta). \end{aligned}$$

First, consider the case where $f_*(\ell - \Delta) = f_{n^*-1}$ (higher density). Let ℓ' be the left-most location such that $f_*(x) = f_{n^*-1}$ for all $x \in [\ell' - \Delta, \ell - \Delta]$. In this case, the seller can increase her demand (profit) by moving to the left by $\ell - \ell'$, that is,

$$\begin{aligned} D(\ell', p; F_*) - D(\ell, p; F_*) &= (\ell - \ell') f_{n^*-1} - (\ell - \ell') \int_{\ell'+\Delta}^{\ell+\Delta} f_*(x) dx \\ &= (\ell - \ell') \int_{\ell'+\Delta}^{\ell+\Delta} \left(f_{n^*-1} - f_*(x)\right) dx \geq 0, \end{aligned}$$

where the inequality holds because f_* cannot strictly exceed f_{n^*-1} anywhere. The desired result then follows from the fact that f_* is periodic, so $D(-1+\Delta, p; F_*) = D(\ell', p; F_*)$.

Next, consider the case where $f_*(\ell + \Delta) = f_{n^*-1}$. In this case, for the same reason as in the previous case, the seller can increase her demand by moving to the *right* (and then to $1 - \Delta$), which yields $D(1 - \Delta, p; F_*) \geq D(\ell, p; F_*)$. Combining this with the symmetry of F_*, we have $D(-1 + \Delta, p; F_*) = D(1 - \Delta, p; F_*) \geq D(\ell, p; F_*)$.

Finally, consider the case where $f_*(\ell - \Delta) = f_*(\ell + \Delta) = f_{n^*}$ (lower density). Let ℓ' be the largest location below ℓ such that $f_*(\ell' - \Delta) = f_{n^*-1}$ or $f_*(\ell' + \Delta) = f_{n^*-1}$ (i.e., ℓ^* is the location at which one of the edges of $[\ell' - \Delta, \ell' + \Delta]$ meets the higher density region). By construction, $D(\ell', p; F_*) = D(\ell, p; F_*)$. But, now $f_*(\ell' - \Delta) = f_{n^*-1}$ or $f_*(\ell' + \Delta) = f_{n^*-1}$, so one of the above cases applies to ℓ'. This leads to $D(-1 + \Delta, p; F_*) \geq D(\ell', p; F_*) = D(\ell, p; F_*)$.

Having established that it is without loss of generality to maximize $\widehat{\pi}(\Delta) = (1 - c(\Delta))F_*(-1 + 2\Delta)$ over Δ, we show $\widehat{\pi}(\Delta)$ has only two local maximizers, $\Delta_1 := \frac{1}{n^*}$ and $\Delta_2 := \frac{1}{n^*-1}$. Define $\Delta_3 := \frac{2-\kappa}{n^*-1}$, so that

$$f_*(x) = \begin{cases} f_{n^*-1} & \text{if } x \in [-1, -1 + \kappa) \text{ or } x \in [-1 + \Delta_3, -1 + \Delta_3 + \kappa) \\ f_{n^*} & \text{if } x \in [-1 + \kappa, -1 + \Delta_3). \end{cases}$$

The desired result follows from the following four results.

(i) $\widehat{\pi}(\Delta)$ is increasing if $\Delta < \Delta_1$.

Since $c(\cdot)$ is increasing and convex, $c(\Delta) < c(\Delta_1)$ and $c'(\Delta) \leq c'(\Delta_1)$. Combining this with the fact that $f_*(-1 + 2\Delta_1) = f_{n^*} \leq f_*(-1 + 2\Delta)$ for any Δ, we get

$$\begin{aligned} \widehat{\pi}'(\Delta) &= -c'(\Delta)F_*(-1 + 2\Delta) + 2(1 - c(\Delta))f_*(-1 + 2\Delta) \\ &> -c'(\Delta_1)F_*(-1 + 2\Delta_1) + 2(1 - c(\Delta_1))f_*(-1 + 2\Delta_1) \\ &= \widehat{\pi}'(\Delta_1) \\ &= 0. \end{aligned}$$

(ii) $\widehat{\pi}(\Delta)$ is decreasing if $\Delta \in (\Delta_1, \Delta_3]$.

In this case, $f_*(-1+2\Delta) = f_{n^*}$ (low density). Combining this with $\Delta > \Delta_1$ (so $c(\Delta) > c(\Delta_1)$

and $c'(\Delta) \geq c'(\Delta_1)$) leads to

$$
\begin{aligned}
\widehat{\pi}'(\Delta) &= -c'(\Delta)F_*(-1+2\Delta) + 2(1-c(\Delta))f_*(-1+2\Delta) \\
&= -c'(\Delta)F_*(-1+2\Delta) + 2(1-c(\Delta))f_{n^*} \\
&< -c'(\Delta_1)F_*(-1+2\Delta_1) + 2(1-c(\Delta_1))f_*(-1+2\Delta_1) \\
&= \widehat{\pi}'(\Delta_1) \\
&= 0.
\end{aligned}
$$

(iii) $\widehat{\pi}(\Delta)$ is increasing if $\Delta \in (\Delta_3, \Delta_2)$.

In this case, $f_*(-1+2\Delta) = f_{n^*-1}$ (high density). Then, similarly to (i),

$$
\begin{aligned}
\widehat{\pi}'(\Delta) &= -c'(\Delta)F_*(-1+2\Delta) + 2(1-c(\Delta))f_*(-1+2\Delta) \\
&> -c'(\Delta_2)F_*(-1+2\Delta_2) + 2(1-c(\Delta_2))f_*(-1+2\Delta_2) \\
&= \widehat{\pi}'(\Delta_2) \\
&= 0.
\end{aligned}
$$

(iv) $\widehat{\pi}(\Delta)$ is decreasing if $\Delta > \Delta_2$.

Since $f_*(-1+2\Delta_2) = f_{n^*-1}$ (high density), similarly to (ii),

$$
\begin{aligned}
\widehat{\pi}'(\Delta) &= -c'(\Delta)F_*(-1+2\Delta) + 2(1-c(\Delta))f_*(-1+2\Delta) \\
&\leq -c'(\Delta)F_*(-1+2\Delta) + 2(1-c(\Delta))f_{n^*-1} \\
&< -c'(\Delta_2)F_*(-1+2\Delta_2) + 2(1-c(\Delta_2))f_*(-1+2\Delta_2) \\
&= \widehat{\pi}'(\Delta_2) \\
&= 0.
\end{aligned}
$$

■

Proof of Proposition 1. If $t \leq \frac{1}{2}$, then the result is immediate, because $\pi^U = \pi_* = \underline{\pi} = 1-t$. From now on, we restrict attention to $t > 1/2$ by redefining $t_1 = \frac{1}{2}$. Note that for any $t > t_1$, the seller's profit under the uniform distribution is given by $\pi^U = \frac{1}{4t}$.

Take any $t \in (t_n, t_{n+1}]$. Then,

$$\begin{aligned}\frac{\pi^U}{\pi_*} &= \frac{\frac{1}{4t}}{\frac{1}{n}\left(1-\frac{t}{n}\right)} \\ &= \frac{1}{4}\frac{1}{\frac{t}{n}\left(1-\frac{t}{n}\right)}.\end{aligned}$$

Define a function $\phi : [t_n, t_{n+1}] \to \mathcal{R}_+$ as $\phi_n(t) := \frac{t}{n}\left(1-\frac{t}{n}\right)$. If $n = 1$ (i.e., $t \in \left[\frac{1}{2}, \frac{2}{3}\right]$) then its maximum is given by $\phi_n(1/2) = 1/4$, while its minimum is given by $\phi_n(t_2) = \phi_n(2/3) = 2/9$. It follows that

$$\frac{\pi^U}{\pi_*} \in \left[\frac{1}{4\phi_1(1/2)}, \frac{1}{4\phi_1(2/3)}\right] = \left[1, \frac{9}{8}\right] \text{ for } t \in [t_1, t_2].$$

For $n > 1$, the maximum of ϕ_n is given by $\phi_n(n/2) = 1/4$, while its minimum is given by $\phi_n(t_n) = \frac{n(n-1)}{(2n-1)^2}$. This implies that

$$\frac{\pi^U}{\pi_*} \in \left[\frac{1}{4\phi_n(n/2)}, \frac{1}{4\phi_n(t_n)}\right] = \left[1, \frac{(2n-1)^2}{4n(n-1)}\right] \text{ for } t \in [t_n, t_{n+1}].$$

Since $\phi_n(t_n) = \frac{n(n-1)}{(2n-1)^2} = \frac{1}{1+1/(n(n-1))}$ is strictly increasing in n, the global maximum of π^U/π_* is given by $9/8$, which is achieved when $t = \frac{2}{3}$. ■

Proof of Lemma 4. We show that if $m > n^*$ and $F = U[-1, 1]$ then the seller can never obtain strictly more than $\underline{\pi}_m = 1 - c(1/m)$. Consider any strategy by the seller. For each $i = 1, ..., m$, let q_i denote the expected mass of consumers variety i is sold to. By definition, $\sum_i q_i \leq 1$. Now, notice that, since $1 - c(q_i)$ is the highest price the seller can set to serve q_i mass of consumers, the seller's expected profit from variety i cannot exceed $q_i(1 - c(q_i))$. If $\sum_i q_i < 1$, there exists i such that $q_i < 1/m \leq 1/(n^* + 1)$. Since the seller's profit is increasing in Δ on $[0, 1/(n^* + 1)]$, she can do strictly better by expanding the partition element and so increasing q_i. This implies that the seller's profit is maximized when she fully utilizes $[-1, 1]$, that is, $\sum_i q_i = 1$.

Given any partition with reaches $\Delta_1, \Delta_2, ..., \Delta_m$,

$$\begin{aligned}\sum_{i=1}^m \Delta_i(1 - c(\Delta_i)) &= m\sum_{i=1}^m \frac{1}{m}\Delta_i(1 - c(\Delta_i)) \\ &\leq m\left(\frac{1}{m}\left(1 - c\left(\frac{1}{m}\right)\right)\right) \\ &= 1 - c\left(\frac{1}{m}\right).\end{aligned}$$

where the inequality is due to to Jensen's inequality and the fact that $\sum_i \Delta_i = 1$. Consequently, $\pi(U[-1,1]) = m/m(1 - c(1/m)) = 1 - c(1/m)$, which is the desired result. ■