J. Econ. Theory 231 (2026) 106112

Contents lists available at ScienceDirect

Journal of Economic Theory

journal homepage: www.elsevier.com/locate/jet

Robust product design and pricing

Kyungmin Kim [a], Nenad Kos [b,*]

[a] Emory University, Atlanta, Georgia, USA
[b] Bocconi University, CEPR, IGIER, and PERICLES, Italy

ARTICLE INFO

JEL classification:
D42
D80
L12

Keywords:
Hotelling
Product design; robust pricing
Profit guarantee
Circular city
Product variety

ABSTRACT

We study design and pricing by a monopolist who has no information about the distribution of consumers' tastes and maximizes her profit under the worst-case scenario. We show that her optimal strategy takes a 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 and establishing strong duality. The profit-minimizing distributions exhibit a uniformity property, assigning equal probability mass to a finite number of partition cells of equal width. Through the dual, we also determine the seller's lowest profit in the Bayesian setting, establish that it is strictly positive, and derive the set of achievable profits.

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 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.

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.[2]

* Corresponding author.
E-mail addresses: kyungmin.kim@emory.edu (K. Kim), nenad.kos@unibocconi.it (N. Kos).

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

[2] For examples of recent use of the Hotelling model in information economics, see Hidir and Vellodi (2021), Callander et al. (2021), Callander et al. (2022), Ali et al. (2023) and Rhodes and Zhou (2024).

https://doi.org/10.1016/j.jet.2025.106112
Received 4 April 2025; Received in revised form 1 October 2025; Accepted 12 November 2025
Available online 19 November 2025

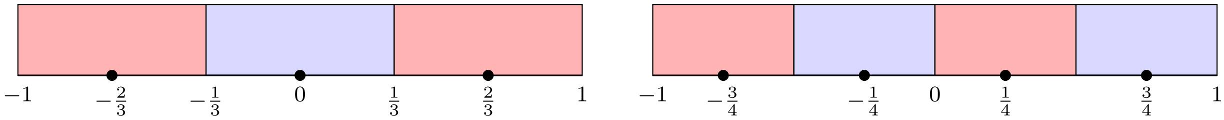

Fig. 1. The seller's robustly optimal strategy when $n^* = 3$ (left) or $n^* = 4$.

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 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 Fig. 1.[3] 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. The randomization over a finite number of designs distinguishes our results from results in models without product design, where the seller typically randomizes over a (continuous) range of prices; see, for example, Bergemann and Schlag (2011), Auster (2018), Carrasco et al. (2018), and Du (2018).

To derive the robust solution, we begin by tackling the dual problem, where nature moves first and selects a distribution of tastes with the intention of minimizing the seller's profit.[4] 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 through the max-min inequality; see, e.g., Osborne and Rubinstein (1994). In addition, it is of economic interest in and of itself, as it establishes the lowest profit the seller achieves under any distribution in the Bayesian setting, and enables us to characterize the set of profits that can be obtained across all distributions of tastes.[5] Our characterization of the set of profits provides two key insights. First, it highlights how the seller's profit depends on the distribution of consumers' tastes within the Hotelling model. This is particularly relevant since most of the existing literature focuses on the case of a uniform distribution; for recent examples, see Hidir and Vellodi (2021) and Bar-Isaac et al. (2023). Second, the lower bound on profit we establish applies universally across all distributions, meaning it remains valid even in richer settings where, for example, information is initially revealed or acquired.

The result uses the following observation: when the seller knows the distribution of tastes, for any integer n, 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.

There are two notable by-products of our analysis of the dual problem. First, the seller's profit is bounded away from 0. This is in stark contrast to the monopoly model without product design where the seller's profit can be arbitrarily small. Second, the uniform distribution, which is commonly adopted in applications, is close to seller-worst. 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). In Section 6, we show that in the circular city environment, the uniform distribution is always seller-worst.

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 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

[3] At the early stages of their work, online content creators typically operate with limited knowledge of consumer demand. This lack of clarity prompts them to pursue ideas at random, with outcomes largely contingent upon chance. For example, MrBeast, one of YouTube's biggest creators, started with Minecraft videos and random gaming content, to eventually find fame with challenge and stunt videos where contestants can win large sums of money; contributors (2025). We would like to thank an anonymous referee for suggesting this application.

[4] The duality approach (equivalently, zero-sum game representation) is common in the robust pricing literature. See, e.g., Bergemann and Schlag (2008, 2011), Carrasco et al. (2018), Du (2018), Che and Zhong (2021).

[5] The corresponding upper bound on profit is 1, which is achieved whenever the distribution of tastes is degenerate. See Kim and Kos (2025) for a characterization of the maximal consumer surplus in monopoly with product design.

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

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 bounds on the support are 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 (2008) consider the case where the seller has no information about the buyer's valuation and minimizes regret. Bergemann and Schlag (2011) study a standard monopoly model where the seller only knows the neighborhood the distribution of consumers' values belongs to; robustness to such local uncertainty is further examined in Li and Wang (2024). Carrasco et al. (2018) explore a model where the seller knows some moments of the distribution of consumers' values. Carroll (2017), Deb and Roesler (2021) and Che and Zhong (2021) study robustness in multi-dimensional settings, while Auster (2018) and Du (2018) characterize robust mechanisms in the common value settings. Che (2022), He and Li (2022), and Suzdaltsev (2022) study robust auctions. Informationally robust mechanism design is investigated by Bergemann et al. (2016) and Brooks and Du (2021, 2024, 2025). See Carroll (2019) for a comprehensive survey of the literature.[7] 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. Section 5 and 6 consider the optimal product line design problem and the circular city model, respectively. Section 7 concludes.

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 a 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 $x \in [-1, 1]$ 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.[8]

Given (ℓ, p), a consumer at x purchases the product if and only if $1 - c(|x - \ell|) \geq p$.[9] 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. Δ_p is continuous and strictly decreasing in p, so the seller can be interpreted as choosing reach Δ (around ℓ) instead of p. We make use of the following notation: For any $\ell \leq \ell'$, $F([\ell, \ell']) := F(\ell') - F_-(\ell)$, where $F_-(x) := \lim_{x' \uparrow x} F(x)$, represents the measure of consumers that lie on $[\ell, \ell']$. If the underlying distribution of tastes is F and the seller chooses (ℓ, p), the quantity demanded is

$$D(\ell, p; F) := F\big([\ell - \Delta_p, \ell + \Delta_p]\big),$$

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 $[-1, 1]$.[10] The seller, therefore, entertains any distribution F over $[-1, 1]$ as plausible and evaluates each strategy according to its worst-case scenario regarding F.[11] 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

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

[7] For the related literature on robust contracting, see Carroll (2015) and Carroll and Meng (2016). Robustness in games of incomplete information is explored in Brooks et al. (2025).

[8] The problem simplifies considerably if the seller is restricted to a pure strategy: If $1 - c(1) \leq 0$, then the seller's profit guarantee is equal to 0 because for any strategy that does not cover the whole interval, the seller fears that the whole mass is on the portion that is not covered (where consumers do not buy). If $1 - c(1) > 0$ then the seller cannot do better than selecting $\ell = 0$ and $p = 1 - c(1)$, which ensures the profit $1 - c(1)$ regardless of the distribution F.

[9] We assume that consumers who are indifferent purchase the seller's product. This ensures that the seller's best response to any F is well defined. Dropping the assumption would not alter the results but would require cumbersome use of ϵ-best responses and limits.

[10] Our formulation allows for uncertainty about the bounds of the support, by interpreting -1 as the smallest possible lower bound and 1 as the highest possible upper bound.

[11] 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).

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*.

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, and the seller designs and prices the product after having observed the chosen distribution. 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 the 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*.

Discussion. We study the canonical model of Hotelling in which all consumers have the same valuation for their preferred design; see for example Tirole (1988). Our results extend to the case where the seller does not know the distribution over valuations and the lowest valuation is 1.

Kim and Kos (2025) show that the version of the here-studied model where the seller knows the distribution of tastes nests the monopoly model without design as the special case corresponding to the seller's fixed design. Namely, given the fixed design, the distribution of tastes induces a distribution over willingness to pay. Conversely, for a fixed design and any distribution over willingness to pay, there exists a distribution of tastes that induces the distribution over willingness to pay.

The feature that distinguishes our model from the monopoly model without design is a lack of natural rankings over locations. In that 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_*.

3.1. 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.[12] 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 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 1. *The seller's lowest profit π_* must be at least $\underline{\pi}$, where*

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

Restricting attention to an integer number of intervals—equivalently, requiring $1/\Delta$ to be 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: Consider the binary distribution that assigns equal probability to -1 and 1. In this case, an interval can include at most mass $1/2$ of consumers, unless it encompasses the whole space $[-1, 1]$. This implies that for $\Delta \in (1/2, 1)$, there does not exist any ℓ such that $F([\ell - \Delta, \ell + \Delta]) \geq \Delta$.[13]

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

[13] More generally, if there is probability mass of equal size $(1/n)$ at locations $\{-1, -1 + 2/(n - 1), \ldots, 1\}$, the seller has to cover at least portion $1/(n - 1)$ of the interval to serve more than $1/n$ consumers.

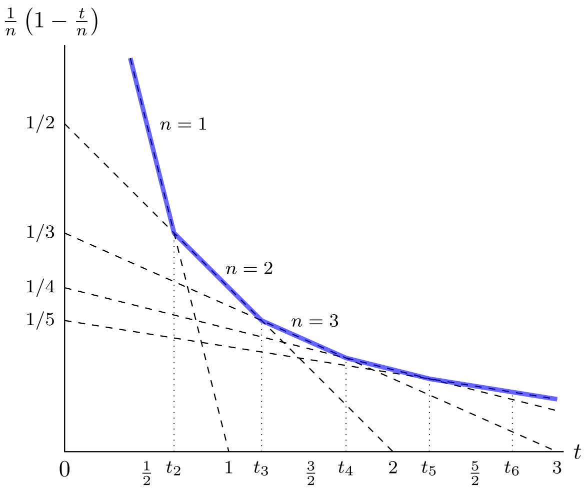

Fig. 2. The optimal value of n that yields $\underline{\pi}$ depending on t in the linear case where $c(y) = ty$. Each dashed line depicts $\frac{1}{n}\left(1-\frac{t}{n}\right)$ for some n, and the blue translucent curve shows its upper envelope, which coincides with $\underline{\pi}$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

If Δ can take any value in $[0,1]$, then the resulting profit corresponds to the profit obtained under the uniform distribution: if F is uniform over $[-1,1]$ then it is always optimal to choose $\ell = 0$, in which case the seller's problem reduces to

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

Let Δ^U denote a solution to this maximization problem under the uniform distribution. Together with concavity of $\Delta(1-c(\Delta))$, the relationship between $\underline{\pi}$ and π^U leads to the subsequent result.

Lemma 2. *The following statements hold:*

(a) If $\Delta^U = 1/n$ for some $n \in \mathbb{N}$ then $\underline{\pi} = \pi_* = \pi^U$.
(b) Let $\hat{n} = \lceil 1/\Delta^U \rceil$.[14] *Then,*

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

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

Lemma 2 provides a tractable way to solve the discrete maximization problem in Lemma 1. One can first consider the unrestricted problem (3) where Δ is permitted to take any value in $[0,1]$ or, equivalently, find the optimal reach under the uniform distribution. Since the problem is concave, it suffices to consider the two reaches of the form $1/n$ that are closest to the solution of the unrestricted problem. In the following paragraph we use this approach to fully solve the problem with linear disutility.

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

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

See Fig. 2 for a graphical illustration of this argument.

3.2. Achieving the profit lower bound

The main result of this section is that $\underline{\pi}$ in Lemma 1 is the tight lower bound for the seller's profit. In other words, $\underline{\pi}$ is the seller's lowest possible profit in the monopoly model with product design.

[14] W use $\lceil x \rceil$ to denote the smallest integer greater or equal to x.

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

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

We explicitly construct a distribution F that delivers Theorem 1. The distribution we present below has a particularly simple structure: its density function is a step function that takes only two values. The distribution function we construct may not be the only distribution function that minimizes the seller's payoff. However, any such distribution should satisfy certain properties we derive. In what follows, we let n^* refer to the natural number such that $\underline{\pi} = 1/n^*(1 - c(1/n^*))$.

A *necessary* condition for F to satisfy $\pi(F) = \underline{\pi}$ is that it never assigns strictly more probability than $1/n^*$ to any interval with length $2/n^*$, that is,

$$F\left(\left[\ell - \frac{1}{n^*}, \ell + \frac{1}{n^*}\right]\right) \leq \frac{1}{n^*} \text{ for all } \ell \in [-1, 1]. \tag{5}$$

Otherwise, the seller could obtain strictly more than $\underline{\pi}$ by choosing that particular ℓ along with $\Delta = 1/n^*$. On the other hand, since the total measure of consumers is 1, it must be that

$$\sum_{k=1}^{n^*} F\left(\left[-1 + \frac{2(k-1)}{n^*}, -1 + \frac{2k}{n^*}\right]\right) \geq F([-1, 1]) = 1.$$

The inequality (instead of equality) is because of potential probability mass at $-1 + 2k/n^*$ for $k = 1, \dots, n^* - 1$. Combining these inequalities, it follows that F should assign probability mass $1/n^*$ to all $\left[-1, -1 + \frac{2}{n^*}\right], \dots, \left[1 - \frac{2}{n^*}, 1\right]$, that is,

$$F\left(\left[-1 + \frac{2(k-1)}{n^*}, -1 + \frac{2k}{n^*}\right]\right) = \frac{1}{n^*} \text{ for all } k = 1, \dots, n^*. \tag{6}$$

Given the condition (5) and the requirement that $F([-1, 1]) = 1$, one might be led to believe that only the uniform distribution satisfies the above necessary properties. This is not the case. In fact, Lemma 2 demonstrates that the uniform distribution is usually not profit-minimizing.

Another necessary property for F to yield $\underline{\pi}$ is that serving the consumers on $[-1 + 2(k-1)/n^*, -1 + 2k/n^*]$—by 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 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 $(1 - c(\Delta))F([-1, -1 + 2\Delta])$. Let f denote the density function of F. Evaluating the first-order condition at $\Delta = 1/n^*$ and invoking (6), we obtain

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

Applying the same argument to all other subintervals results in:

$$f\left(-1 + \frac{2}{n^*}\right) = \dots = f\left(1 - \frac{2}{n^*}\right) = \frac{1}{2n^*} \frac{c'(1/n^*)}{1 - c(1/n^*)}. \tag{7}$$

The following distribution combines Lemma 2.(b) with the two necessary conditions, (6) and (7), in a simple manner.

Definition 1. Let $\hat{n}$ be the value defined in Lemma 2 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

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

where $\kappa := \frac{1-2f_{\hat{n}}}{\hat{n}(f_{\hat{n}-1} - f_{\hat{n}})}$.[15]

Fig. 3 shows a representative structure of the density f^*. It alternates between density levels $f_{\hat{n}-1}$ (high) and $f_{\hat{n}}$ (low). By construction, f coincides with $f_{\hat{n}}$ around $-1 + \frac{2}{\hat{n}}, \dots, 1 - \frac{2}{\hat{n}}$ and with $f_{\hat{n}-1}$ around $-1 + \frac{2}{\hat{n}-1}, \dots, 1 - \frac{2}{\hat{n}-1}$, thus ensuring that f^* satisfies the necessary condition (7), whether the optimal n is $\hat{n} - 1$ or $\hat{n}$. In addition, the lengths of the subintervals are chosen so that F assigns probability $1/\hat{n}$ to all intervals $\left[-1, -1 + \frac{2}{\hat{n}}\right], \dots, \left[1 - \frac{2}{\hat{n}}, 1\right]$ and $1/(\hat{n} - 1)$ to all intervals $\left[-1, -1 + \frac{2}{\hat{n}-1}\right], \dots, \left[1 - \frac{2}{\hat{n}-1}, 1\right]$; the properties guarantee condition (6), whether the optimal n is $\hat{n} - 1$ or $\hat{n}$. Note that κ is the width of each high-density interval.

Given F^*, the seller can achieve $\underline{\pi}$ by serving $\left[-1, -1 + \frac{2}{\hat{n}}\right]$ or $\left[-1, -1 + \frac{2}{\hat{n}-1}\right]$, depending on which of the two intervals leads to a higher profit. It remains to show that the seller's profit *cannot* exceed $\underline{\pi}$ under F^*. We establish this result in two steps. First, we show that for any $\Delta > 0$ and $\ell \in [-1, 1]$,

$$F^*([-1, -1 + 2\Delta]) \geq F^*([\ell - \Delta, \ell + \Delta]).$$

[15] The value κ is such that $\hat{n}\kappa f_{\hat{n}-1} + (2 - \hat{n}\kappa)f_{\hat{n}} = 1$. 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 $\kappa \in [0, 1/(2\hat{n})]$.

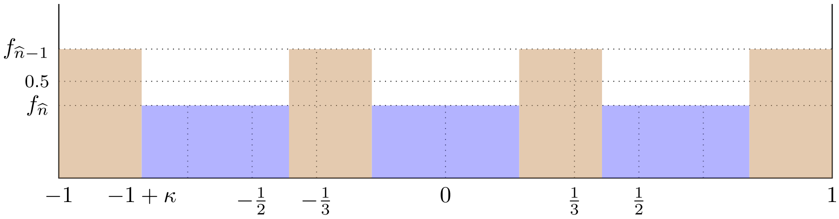

Fig. 3. The density of distribution F^* defined in Definition 1. The cost function used for this figure is $c(y) = \frac{12}{7}y$. In this case, $\hat{n} = 4$, and $\pi_* = \frac{1}{3}\left(1 - c\left(\frac{1}{3}\right)\right) = \frac{1}{4}\left(1 - c\left(\frac{1}{4}\right)\right)$.

In other words, given F^*, for any reach Δ, it is weakly better for the seller to serve $[-1, -1 + 2\Delta]$ than any other $[\ell - \Delta, \ell + \Delta]$. This is because f^* has a periodic structure and each subinterval with high density $f_{\hat{n}-1}$ has the same length; if ℓ increases from $-1 + \Delta$ then $F([\ell - \Delta, \ell + \Delta])$ may become smaller than, but cannot exceed, $F^*([-1, -1 + 2\Delta])$. Second, restricting attention to the intervals of the form $[-1, -1 + 2\Delta]$, the seller's profit $(1 - c(\Delta))F^*([-1, -1 + 2\Delta])$ has only two local maximizers, $\frac{1}{\hat{n}}$ and $\frac{1}{\hat{n}-1}$. These together imply that the seller cannot do strictly better than serving either $\left[-1, -1 + \frac{2}{\hat{n}}\right]$ or $\left[-1, -1 + \frac{2}{\hat{n}-1}\right]$. The remaining details can be found in the appendix.

3.3. Discussion

We conclude this section by providing three implications of Theorem 1 that are of economic interest.

The Set of Attainable Profits. The highest profit the seller can obtain in the Hotelling model is 1, which is available if (and only if) the distribution of tastes is degenerate. The following result shows that any profit between π_* and 1 is attainable.

Proposition 1. *There exists a distribution $F \in \Delta([-1, 1])$ under which the seller obtains π (i.e., $\pi = \sup_{\sigma \in \Sigma} \pi(\sigma; F)$) if and only if $\pi \in [\pi_*, 1]$.*

Proof. We provide a sketch of the proof here, relegating a comprehensive proof to the appendix. For each $\eta \in [0, 1]$, we let $\mathcal{F}^\eta$ denote the set of all distributions over $[-\eta, \eta]$ and then find a seller-worst distribution in $\mathcal{F}^\eta$. If F coincides on its support with some distribution in $\mathcal{F}^\eta$, the seller has no incentive to cover $[-1, -\eta)$ and $(\eta, 1]$. Combining this with the logic used to derive the profit lower bound in Section 3.1, it follows that the seller's profit cannot be lower than

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

Then, following the same steps as in Section 3.2, one can construct a distribution under which the seller's maximized profit is $\underline{\pi}^\eta$. The desired result follows from the fact that $\underline{\pi}^\eta$ continuously increases from $\underline{\pi} = \underline{\pi}^1$ to $1 = \underline{\pi}^0$ as η decreases from 1 to 0. □

Comparison to the classical monopoly model without design. This result is in stark contrast to the corresponding result in the monopoly model without product design where the seller can obtain any profit in $[0, 1]$. The classical monopoly model, thus, underestimates the seller's profit at the lower end by neglecting the fact that the seller designs the product to appeal to a wider public.

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 latter is an integer version of the former. Lemma 2 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 2. *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]

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^*.

[16] While it may seem that generically the uniform distribution is not a worst distribution, this is not quite the case, because the boundary solution with $\Delta^U = 1$ arises for a non-negligible 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.

Observing 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., $\pi^*\leq\pi_*$).

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, if n^* is the (unique) value such that $\pi_*=1/n^*(1-c(1/n^*))$ then the seller's (uniquely) robust strategy is to set $p=1-c(1/n^*)$ and uniformly randomize the design ℓ over $\left\{-1+\frac{1}{n^*},-1+\frac{3}{n^*},\dots,1-\frac{1}{n^*}\right\}$.[18]

Proof. As argued above, the inequality $\pi^*\leq\pi_*$ holds invariably. We prove 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

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

where the inequality holds with equality whenever F has no atom on $\{-1+2/n^*,\dots,1-2/n^*\}$. Since this holds for any $F\in\mathcal{F}$, we arrive at the desired inequality $\pi^*\geq\pi_*$. See the appendix for a proof that σ^* is the unique robust strategy of the seller whenever n^* is unique. □

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, $\{-1+2/n^*,\dots,1-2/n^*\}$) and thus almost all distributions.[19] To see why this property is essential, notice that given the 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 $\left\{-1+\frac{1}{n^*},-1+\frac{3}{n^*},\dots,1-\frac{1}{n^*}\right\}$ 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 $1/n^*$). Second, all submarkets have the same size (length), or equivalently, all the consumers are offered the same price. Both of these features differentiate our model from existing models of robust pricing, where the robustly optimal strategy typically involves price randomization (see, e.g., Bergemann and Schlag, 2011; Auster, 2018; Carrasco et al., 2018). The differences are driven by the fact that in our model, there is a horizontal taste dimension, which can be effectively hedged only through a randomization over designs.

Worst Distributions. To prove Theorem 1, we constructed a particular class of distributions. However, the saddle point might not be unique. Nevertheless, if there are two saddle points (σ_1,F_1) and (σ_2,F_2), the combination of the two (σ_1,F_2) is also a saddle point; see Section 2.5 in Osborne and Rubinstein (1994). Since the value of the inf sup problem is unique, this implies that any distribution that solves the inf sup problem has to assign mass $1/n^*$ to all the intervals $\left[-1,-1+\frac{2}{n^*}\right],\dots,\left[1-\frac{2}{n^*},1\right]$.

Intuitively, when minimizing the seller's profit, nature strives for the balance between making the seller indifferent over a large set of designs and incentivizing her to offer a price that will lead to a low profit. Uniform distribution ensures the seller's indifference over an interval of designs, but the price the seller charges might not minimize the profit across all the distributions. The class of distributions provided in Definition 1 strikes the balances between keeping the seller indifferent between a set of designs and offering a price that minimizes her profit.

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 complete coverage strategy can indeed be optimal but is not optimal always. The following result provides a necessary and sufficient condition for its optimality.

[18] Recall that n^* such that $\pi_*=1/n^*(1-c(1/n^*))$ is typically unique, but can be two consecutive integers (e.g., whenever $t=t_n=\frac{n(n-1)}{2n-1}$ with linear disutility $c(y)=ty$). In the latter case, there are many equilibria, because the seller can not only divide $[-1,1]$ in two different ways, but also mix between those two cases.

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

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 $\beta \geq 1$. 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. Robust 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.[20]

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 $i = 1, \ldots, m$, the seller chooses its design $\ell_i \in [-1, 1]$ and price $p_i \in \mathbb{R}_+$; the seller's pure strategy is $(\ell_1, p_1; \ldots; \ell_m, p_m) \in \left([-1, 1] \times \mathbb{R}_+\right)^m$. 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 $F \in \mathcal{F}$. We use π_{*m} to denote the corresponding inf-sup value and, as in the single-product case, refer to it as the seller's lowest profit. 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 smaller than

$$\underline{\pi}_m := \max_{n \geq m} \frac{m}{n}\left(1 - c\left(\frac{1}{n}\right)\right). \tag{8}$$

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 3 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

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

Lemma 3. *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 $i = 1, \ldots, m$.*

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

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

[20] 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).

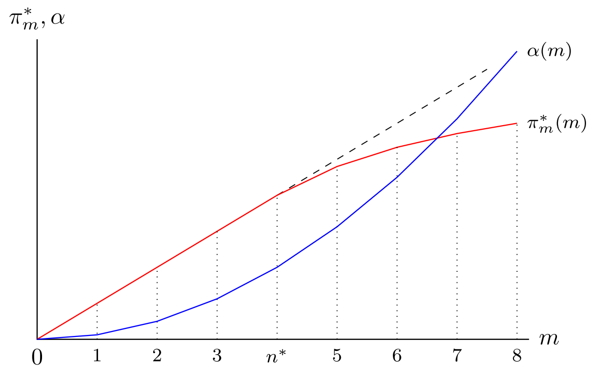

Fig. 4. The red curve shows the seller's profit guarantee π^*_m as a function of m for the case when $c(y) = 2y$ (in which case $n^* = 4$). The blue curve depicts an example of the convex cost function $\alpha(m)$. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

If $m \le n^*$, 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.[21] It follows that $\pi^*_m \ge \underline{\pi}_m$, which together with the above inequality yields 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*

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

Let $n^\dagger := \max\{m, n^*\}$. The seller's robust strategy is to divide $[-1, 1]$ into $n^\dagger$ subintervals of equal length and serve m randomly selected subinterrvals.[22]

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 Fig. 4, 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 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.[23]

Recall that in the dual problem of Section 3.2, the seller's lowest profit π_* is given by

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

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

[22] As in Theorem 2, the seller's optimal strategy is unique whenever n^* such that $\pi_* = 1/n^*(1 - c(1/n^*))$ is unique. Uniqueness also holds whenever $m \ge n^*$.

[23] 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.

which is necessarily smaller than her profit under the uniform distribution

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

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 3.1, 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 4. *Given any probability measure μ over a circle with circumference 2, for any $\Delta \le 1$, there exists ℓ such that $\mu([\ell - \Delta, \ell + \Delta]) \ge \Delta$.*

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. $\square$

Lemma 4 implies that the seller's lowest profit satisfies

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

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

Proposition 3. *In the circular city, the uniform distribution is always a seller-worst distribution. If its circumference is 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 $p = 1 - c(\Delta^U)$ at every location.*[24]

The above result sheds light on the distribution of surplus under the common practice of assuming that the distribution of tastes in the circular city model is uniform.

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 nature's choice of distribution can be held to (inf-sup).

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. This is, of course, the extreme case. 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. Future work could address the question of optimal design and pricing if the seller has partial information about the distributions of tastes (for example, moments or quantiles of the distribution), or how much information the seller wants to acquire prior to designing a product.

Consistent with standard practice in Hotelling models, this paper examines optimal product design and pricing. In future work, we intend to investigate general screening mechanisms. Our conjecture is that restricting the seller to designing and pricing a product in an environment with a continuum of consumers is essentially without loss of generality, since in that setting the seller effectively knows the distribution of realized consumer preferences across parameters. This equivalence, however, does not hold when the number of consumers is finite.

[24] If $\Delta^U = 1/n^*$ for some $n^* \in \mathcal{N}$ then the seller can employ a simpler strategy of uniformly randomizing over n^* (or, any multiple of n^*) equally dispersed locations.

CRediT authorship contribution statement

Kyungmin Kim: Writing – original draft, Formal analysis; **Nenad Kos:** Writing – original draft, Formal analysis.

Data availability

No data was used for the research described in the article.

Declaration of competing interest

Kyungmin Kim, an author of the paper titled 'Robust Product Design and Pricing,' is supported by the Ministry of Education of the Republic of Korea and the National Research Foundation of Korea (NRF-2024S1A5A2A03038509). He declares that he has no relevant or material financial interests that relate to the research described in this paper. I, an author of the paper titled 'Robust Product Design and Pricing', declare that I have no relevant or material financial interests that relate to the research described in this paper.

Funding

We thank Christoph Carnehl, Jay Pil Choi, Hülya Eraslan, Nima Haghpanah, Vijay Krishna, Ali Shourideh, and seminar audiences at Bocconi, CERGE-EI, CMU/Pittsburgh, Emory, HEC Lausanne, 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-2024S1A5A2A03038509).

Appendix A. Omitted Proofs

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

$$\pi(\ell, p; F^*) \leq \hat{\pi}(\Delta) := \pi(-1+\Delta, p; F^*),$$

which is equivalent to

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

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 $\hat{\pi}(\Delta) = (1 - c(\Delta))F^*(-1+2\Delta)$ over Δ, we show $\hat{\pi}(\Delta)$ 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) $\hat{\pi}(\Delta)$ is increasing if $\Delta < \Delta_1$.

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

$$\begin{aligned}\hat{\pi}'(\Delta) &= -c'(\Delta)F^*(-1+2\Delta) + 2(1-c(\Delta))f^*(-1+2\Delta)\\ &> -c'(\Delta_1)F^*(-1+2\Delta_1) + 2(1-c(\Delta_1))f^*(-1+2\Delta_1)\\ &= \hat{\pi}'(\Delta_1)\\ &= 0.\end{aligned}$$

(ii) $\hat{\pi}(\Delta)$ 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 $c'(\Delta) \ge c'(\Delta_1)$) leads to

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

(iii) $\hat{\pi}(\Delta)$ 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),

$$\begin{aligned}\hat{\pi}'(\Delta) &= -c'(\Delta)F^*(-1+2\Delta) + 2(1-c(\Delta))f^*(-1+2\Delta)\\ &> -c'(\Delta_2)F^*(-1+2\Delta_2) + 2(1-c(\Delta_2))f^*(-1+2\Delta_2)\\ &= \hat{\pi}'(\Delta_2)\\ &= 0.\end{aligned}$$

(iv) $\hat{\pi}(\Delta)$ is decreasing if $\Delta > \Delta_2$.
Since $f^*(-1+2\Delta_2) = f_{n^*-1}$ (high density), similarly to (ii),

$$\begin{aligned}\hat{\pi}'(\Delta) &= -c'(\Delta)F^*(-1+2\Delta) + 2(1-c(\Delta))f^*(-1+2\Delta)\\ &\le -c'(\Delta)F^*(-1+2\Delta) + 2(1-c(\Delta))f_{n^*-1}\\ &< -c'(\Delta_2)F^*(-1+2\Delta_2) + 2(1-c(\Delta_2))f^*(-1+2\Delta_2)\\ &= \hat{\pi}'(\Delta_2)\\ &= 0.\end{aligned}$$

□

Proof. (Proof of Proposition 1) Given Theorem 1, it suffices to show that for each $\pi \in [\pi_*, 1]$, there exists a distribution that gives π to the seller. Recall that we defined $\mathcal{F}^\eta$ as the set of all distributions over $[-\eta, \eta]$ and $\underline{\pi}^\eta$ as

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

We first show that for any $\eta \in (0, 1]$, there exists a distribution in $\mathcal{F}^\eta$ under which the seller's maximized profit is $\underline{\pi}^\eta$. Let $\hat{c}$ denote the function such that $\hat{c}(y) = c(\eta y)$ for any $y \ge 0$. By Theorem 1 and the definition of $\hat{c}$, there exists a distribution F under which the seller's maximized profit is equal to

$$\underline{\pi}^\eta = \max_{\Delta \in \left\{1, \frac{1}{2}, \frac{1}{3} \dots\right\}} \Delta(1 - \hat{c}(\Delta)) = \max_{\Delta \in \left\{1, \frac{1}{2}, \frac{1}{3} \dots\right\}} \Delta(1 - c(\eta\Delta)).$$

Consider the distribution F^η such that $F^\eta(x) = F(x/\eta)$. By construction, the support of F^η belongs to $[-\eta, \eta]$, so $F^\eta \in \mathcal{F}^\eta$. It suffices to show that the seller's profit cannot exceed $\underline{\pi}^\eta$ under F^η. Consider any strategy $(\ell, p) = (\ell, 1 - c(\Delta))$, and let $\ell' := \ell/\eta$ and $\Delta' := \Delta/\eta$. Then, we have

$$\begin{aligned}\pi(\ell, 1-c(\Delta); F^\eta) &= F^\eta([\ell-\Delta, \ell+\Delta])(1-c(\Delta))\\ &= F\left(\left[\frac{\ell-\Delta}{\eta}, \frac{\ell+\Delta}{\eta}\right]\right)(1-\hat{c}(\Delta/\eta))\\ &= F\big([\ell'-\Delta', \ell'+\Delta']\big)(1-\hat{c}(\Delta'))\\ &= \pi(\ell', 1-\hat{c}(\Delta'); F)\\ &\le \underline{\pi}^\eta,\end{aligned}$$

where the second equality is because of the definitions of F^η and $\hat{c}$, the third equality is because of the definitions of ℓ' and Δ', and the inequality is due to the fact that $\underline{\pi}^\eta$ is the seller's maximized profit under F.

It remains to show that $\underline{\pi}^\eta$ continuously decreases in η; since $\underline{\pi}^1 = \pi_*$ and $\underline{\pi}^0 = 1$, the intermediate value theorem ensures that for any $\pi \in [\pi_*, 1]$ there exists $\eta \in [0,1]$ such that $\underline{\pi}^\eta = \pi$. Consider the following relaxed problem where Δ is allowed to be 0 as well:

$$\underline{\pi}_0^\eta := \max_{\Delta \in \left\{1, \frac{1}{2}, \frac{1}{3} \ldots\right\} \cup \{0\}} \Delta(1 - c(\eta\Delta)).$$

For each Δ, the function $\Delta(1 - c(\eta\Delta))$ is continuous in η. In addition, the domain $\left\{1, \frac{1}{2}, \frac{1}{3} \ldots\right\} \cup \{0\}$ is compact. The theorem of the maximum then implies that $\underline{\pi}_0^\eta$ is always well-defined and continuous in η. In fact, since $\Delta(1 - c(\eta\Delta))$ is decreasing in η for any Δ, the maximum $\underline{\pi}_0^\eta$ is also decreasing in η. The desired result then follows from the fact that $\Delta = 0$ can never be the solution to the above maximization problem, so $\underline{\pi}^\eta = \underline{\pi}_0^\eta$ always holds. □

Proof. (Proof of Proposition 2) If $t \le \frac{1}{2}$, 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,

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

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. (Remaining Proof of Theorem 2) We now establish that if n^* such that $1/n^*(1 - c(1/n^*)) = \pi^*$ is unique then σ^* is the unique optimal strategy. Let σ' denote an optimal strategy of the seller.

First, $\pi(\sigma'; F^*) = \pi^*$, that is, the seller earns π^* when she plays σ' against the seller-worst distribution F^* defined in Definition 1. This follows from the fact that σ' should yield at least π^* regardless of F (i.e., $\pi(\sigma'; F^*) \ge \pi^*$), while, as shown in Section 3.2, the seller's profit can never exceed π_* against F^* (i.e., $\pi(\sigma'; F^*) \le \pi(F^*) \le \pi_* = \pi^*$). Alternatively, this directly follows from a general observation that if both (σ_1, F_1) and (σ_2, F_2) are saddle points then (σ_1, F_2) is also a saddle point.

Then, given σ', nature should be indifferent over almost all degenerate distributions, that is,

$$\pi(\sigma'; \delta_x) = \pi^* \text{ for almost all } x \in [-1, 1], \tag{A.1}$$

where δ_x denotes the Dirac distribution on design x. Indeed, since $\pi(\sigma'; F) \ge \pi^*$ for all distributions, $\pi(\sigma'; \delta_x) \ge \pi^*$ for all $x \in [-1, 1]$. If the inequality were not to hold with equality for almost all x's (i.e., it holds strictly for a non-negligible set of x's) then

$$\pi(\sigma'; F^*) = \int \pi(\sigma'; \delta_x) dF^*(x) > \pi^*,$$

which would contradict the above finding.

The analysis in Section 3.2 then implies that given F^*, the seller can obtain the highest profit π_* only when $p = 1 - c(1/n^*)$, that is, her reach is $1/n^*$. From now on, we take the optimal reach $1/n^*$ as given and focus on the seller's optimal design choice. Denote by G the distribution that represents the seller's design choice under σ'; that is, $G(\ell)$ denotes the probability that the chosen design is weakly below ℓ. Given the unique optimal reach $1/n^*$, the seller has no reason to choose $\ell < -1 + 1/n^*$ or $\ell > 1 - 1/n^*$, so $G(\ell) = 0$ for all $\ell < -1 + 1/n^*$ and $G(\ell) = 1$ for all $\ell > 1 - 1/n^*$.

We show that $G(\ell) = 1/n^*$ for $\ell \in [-1 + 1/n^*, -1 + 3/n^*)$, that is, the seller chooses design $-1 + 1/n^*$ with probability $1/n^*$ and does not choose any designs in $(-1 + 1/n^*, -1 + 3/n^*)$ with positive probability. On the one hand, $G(-1 + 1/n^*) \ge 1/n^*$ is necessary for $\pi(\sigma'; \delta_x) = \pi^*$ to hold for almost all x close to -1; otherwise, there exists $\varepsilon > 0$ such that $\pi(\sigma'; \delta_x) < \pi^*$ whenever $x \in [-1, -1 + \varepsilon)$, which contradicts (A.1). On the other hand, if $G(\ell) > 1/n^*$ for some $\ell \in [-1 + 1/n^*, -1 + 3/n^*)$ then $\pi(\sigma'; \delta_x) > \pi^*$ whenever x is close to $(-1 + 1/n^* + \ell)/2$, which again contradicts (A.1).

The logic in the above paragraph can be recursively applied to the other subintervals, leading to the conclusion that G has a discrete jump of $1/n^*$ at the center of every subinterval. This distribution corresponds to σ^*, and thus $\sigma' = \sigma^*$.

□

Proof. (Proof of Lemma 3) 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 $i = 1, \dots, m$, let q_i denote the expected mass of consumers variety i is sold to. By definition, $\sum_i q_i \le 1$. 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 $q_i < 1/m \le 1/(n^* + 1)$. 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 $\Delta_1, \Delta_2, \dots, \Delta_m$,

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

where the inequality is due to 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. □

References

Ali, S.N., Lewis, G., Vasserman, S., 2023. Voluntary disclosure and personalized pricing. Rev. Econ. Stud. 90 (2), 538–571.
Auster, S., 2018. Robust contracting under common value uncertainty. Theor. Econ. 13 (1), 175–204.
Bar-Isaac, H., Caruana, G., Cuñat, V., 2023. Targeted product design. Am. Econ. J.: Microeconom. 15 (2), 157–186.
Bergemann, D., Brooks, B., Morris, S., 2019. Revenue guarantee equivalence. Am. Econ. Rev. 109 (5), 1911–29.
Bergemann, D., Brooks, B.A., Morris, S., 2016. Informationally robust optimal auction design .
Bergemann, D., Schlag, K., 2011. Robust monopoly pricing. J. Econ. Theory 146 (6), 2527–2543.
Bergemann, D., Schlag, K.H., 2008. Pricing without priors. J. Eur. Econ. Assoc. 6 (2–3), 560–569.
Brooks, B., Du, S., 2021. Optimal auction design with common values: an informationally robust approach. Econometrica 89 (3), 1313–1360.
Brooks, B., Du, S., 2024. On the structure of informationally robust optimal mechanisms. Econometrica 92 (5), 1391–1438.
Brooks, B., Du, S., 2025. Informational robustness in mechanism design a case for (and against) the worst case .
Brooks, B., Du, S., Haberman, A., 2025. Robust predictions with bounded information. Technical Report. Working Paper.
Callander, S., Lambert, N., Matouschek, N., 2021. Arrow meets hotelling: modeling spatial innovation. In: AEA Papers and Proceedings. Vol. 111. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 538–543.
Callander, S., Lambert, N., Matouschek, N., 2022. Innovation and competition on a rugged technological landscape. Technical Report. Working paper Northwestern University.
Carrasco, V., Luz, V.F., Kos, N., Messner, M., Monteiro, P., Moreira, H., 2018. Optimal selling mechanisms under moment conditions. J. Econ. Theory 177, 245–279.
Carroll, G., 2015. Robustness and linear contracts. Am. Econ. Rev. 105 (2), 536–563.
Carroll, G., 2017. Robustness and separation in multidimensional screening. Econometrica 85 (2), 453–488.
Carroll, G., 2019. Robustness in mechanism design and contracting. Annu. Rev. Econom. 11, 139–166.
Carroll, G., Meng, D., 2016. Robust contracting with additive noise. J. Econ. Theory 166, 586–604.
Che, E., 2022. Robustly optimal auction design under mean constraints. In: Proceedings of the 23Rd ACM Conference on Economics and Computation, pp. 153–181.
Che, Y.-K., Zhong, W., 2021. Robustly-optimal mechanism for selling multiple goods. In: Proceedings of the 22Nd ACM Conference on Economics and Computation, pp. 314–315.
contributors, W., 2025. Mrbeast. https://en.wikipedia.org/wiki/MrBeast. Accessed: 2025-09-12.
Dasci, A., Laporte, G., 2005. A continuous model for multistore competitive location. Oper. Res. 53 (2), 263–280.
Deb, R., Roesler, A.-K., 2021. Multi-dimensional screening: buyer-optimal learning and informational robustness. In: Proceedings of the 22Nd ACM Conference on Economics and Computation, pp. 343–344.
Du, S., 2018. Robust mechanisms under common valuation. Econometrica 86 (5), 1569–1588.
He, W., Li, J., 2022. Correlation-robust auction design. J. Econ. Theory 200, 105403.
Hidir, S., Vellodi, N., 2021. Privacy, personalization, and price discrimination. J. Eur. Econ. Assoc. 19 (2), 1342–1363.
Hinnosaar, T., Kawai, K., 2020. Robust pricing with refunds. Rand. J. Econ. 51 (4), 1014–1036.
Hotelling, H., 1929. Stability in competition. Econ. J. 39 (153), 41–57.
Janssen, M. C.W., Karamychev, V.A., Van Reeven, P., 2005. Multi-store competition: market segmentation or interlacing? Reg. Sci. Urban Econ. 35 (6), 700–714.
Jovanovic, B., 1981. Entry with private information. Bell J. Econ., 649–660.
Ke, T.T., Shin, J., Yu, J., 2022. A model of product portfolio design: guiding consumer search through brand positioning. Market. Sci.
Kim, K., Kos, N., 2025. Pricing and consumer surplus in monopoly with product design. mimeo.
Kuksov, D., 2004. Buyer search costs and endogenous product design. Market. Sci. 23 (4), 490–499.
Lauga, D.O., Ofek, E., 2011. Product positioning in a two-dimensional vertical differentiation model: the role of quality costs. Market. Sci. 30 (5), 903–923.
Li, J., Wang, K., 2024. A robust optimization approach to mechanism design .
Meagher, K.J., Zauner, K.G., 2004. Product differentiation and location decisions under demand uncertainty. J. Econ. Theory 117 (2), 201–216.
Moorthy, K.S., 1988. Product and price competition in a duopoly. Market. Sci. 7 (2), 141–168.
Orhun, A.Y., 2009. Optimal product line design when consumers exhibit choice set-dependent preferences. Market. Sci. 28 (5), 868–886.
Osborne, M.J., Rubinstein, A., 1994. A course in game theory. MIT press.
Rhodes, A., Zhou, J., 2024. Personalized pricing and competition. Am. Econ. Rev. 114 (7), 2141–2170.
Salop, S.C., 1979. Monopolistic competition with outside goods. Bell J. Econ., 141–156.
Schmalensee, R., 1978. Entry deterrence in the ready-to-eat breakfast cereal industry. Bell J. Econ., 305–327.
Suzdaltsev, A., 2022. Distributionally robust pricing in independent private value auctions. J. Econ. Theory 206, 105555.
Sweeting, A., 2013. Dynamic product positioning in differentiated product markets: the effect of fees for musical performance rights on the commercial radio industry. Econometrica 81 (5), 1763–1803.
Tirole, J., 1988. The theory of industrial organization. MIT press.
Villas-Boas, J.M., 2004. Communication strategies and product line design. Market. Sci. 23 (3), 304–316.
Villas-Boas, J.M., 2018. A dynamic model of repositioning. Market. Sci. 37 (2), 279–293.