Games and Economic Behavior 131 (2022) 275–291

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Note

Third-degree price discrimination versus uniform pricing ☆

Dirk Bergemann [a], Francisco Castro [b], Gabriel Weintraub [c]

[a] *Department of Economics, Yale University, New Haven, USA*
[b] *Anderson School of Management, UCLA, Los Angeles, USA*
[c] *Graduate School of Business, Stanford University, Stanford, USA*

A R T I C L E   I N F O

Article history:
Received 1 July 2020
Available online 30 November 2021

JEL classification:
C72
D82
D83

Keywords:
Third-degree price discrimination
Uniform price
Approximation
Concave profit function
Market segmentation

A B S T R A C T

We compare the profit of the optimal third-degree price discrimination policy against a uniform pricing policy. A uniform pricing policy offers the same price to all segments of the market. Our main result establishes that for a broad class of third-degree price discrimination problems with concave profit functions (in the price space) and common support, a uniform price is guaranteed to achieve one half of the optimal monopoly profits. This profit bound holds for any number of segments and prices that the seller might use under third-degree price discrimination. We establish that these conditions are tight and that weakening either common support or concavity can lead to arbitrarily poor profit comparisons even for regular or monotone hazard rate distributions.

© 2021 Elsevier Inc. All rights reserved.

1. Introduction

1.1. Motivation and results

An important use of information about demand is to engage in price discrimination. A large body of literature, starting with the classic work of Pigou (1920), examines what happens to prices, quantities, and various measures of welfare as the market is segmented. A seller engages in third-degree price discrimination if he uses information about consumer characteristics to offer different prices to different segments of the market. As every segment is offered a different price, there is scope for the producer to extract more surplus from the consumer. Additional information about consumer demand increases flexibility for market segmentation.[1]

Our main contribution is to compare the profit performance of third-degree price discrimination against a uniform pricing policy. A uniform pricing policy offers the same price to all segments of the market. Theorem 1 establishes that for a broad class of third-degree price discrimination problems with concave profit functions (in the price space) and

☆ The first author would like to acknowledge financial support through NSF grant SES 1459899 and NSF grant SES 1948336. We thank the editor, Tim Roughgarden, and two anonymous referees for productive comments. We thank John Vickers for helpful conversations that lead to an alternative strategy to establish the bound in the main result. We thank Miho Hong for excellent research assistance.

E-mail addresses: dirk.bergemann@yale.edu (D. Bergemann), francisco.castro@anderson.ucla.edu (F. Castro), gweintra@stanford.edu (G. Weintraub).

[1] Pigou (1920) suggested a classification of different forms of price discrimination. First-degree (or perfect) price discrimination is given when the monopolist charges each unit with a price that is equal to the consumer's maximum willingness to pay for that unit. Second-degree price discrimination arises when the price depends on the quantity (or quality) purchased. Third-degree price discrimination occurs when different market segments are offered different prices, e.g., due to temporal or geographical differentiation.

https://doi.org/10.1016/j.geb.2021.11.008
0899-8256/© 2021 Elsevier Inc. All rights reserved.

common support, a uniform price is guaranteed to achieve one half of the optimal monopoly profits. The profit bound in Theorem 1 is independent of the number of segments and holds for any arbitrarily large number of segments. Interestingly, the performance guarantee of Theorem 1 can be established with different choices regarding the uniform price, each of which uses different sources of information regarding the market demand.

We investigate the limits of this result by weakening the assumptions of concavity and common support. First, Proposition 4 shows the significance of the common support assumption by studying a setting with concave profit functions that have finite but different supports. We display a sequence of segments under which the profit ratio of uniform price to third-degree price discrimination goes to zero.[2] Second, Proposition 5 and Proposition 6 note that the approximation result does not hold for the commonly studied class of regular distributions. More specifically, we can weaken the concavity of the profit function to merely assume regular environments while maintaining common support. In other words, we assume that the profit function is only concave in the space of quantiles, rather than prices. Proposition 5 establishes that for some regular distributions, uniform pricing can perform arbitrarily poorly compared to optimal third-degree price discrimination. Proposition 6 establishes that when we consider the even more selective sub-family of monotone hazard rate distributions, the poor performance of uniform pricing still holds. That is, in the third-degree price discrimination setting we need more stringent conditions, such as concavity, beyond the most commonly used notion of regularity to attain good approximations. The importance of the aforementioned results is that they establish that if any of these assumptions is dropped (not necessarily at the same time), then the profit ratio of uniform price to third-degree price discrimination can be small.

As an application of our main result, we consider the dynamic mechanism design problem of sequential screening with ex-post participation constraints investigated by Krähmer and Strausz (2015) and Bergemann et al. (2020). In Section 6, we establish the connection between the aforementioned problem and our setting and show that Theorem 1 implies that the static mechanism in sequential screening can lead to a half approximation of the optimal dynamic mechanism.

1.2. Related literature

Our work builds on the classic literature on third-degree price discrimination, see e.g., Pigou (1920), Robinson (1933), and Schmalensee (1981). In more recent work, Aguirre et al. (2010) identify conditions on the shape of the demand function for price discrimination to increase welfare and output compared to the non-discriminating price case.

Bergemann et al. (2015) analyze the limits of price discrimination. They show that the segmentation and pricing induced by the additional information can achieve every combination of consumer and producer surplus such that: (*i*) consumer surplus is nonnegative, (*ii*) producer surplus is at least as high as profits under the uniform monopoly price, and (*iii*) total surplus does not exceed the surplus generated by the efficient trade. Building on this work, Cummings et al. (2020) provide approximate guarantees to segment the market when an intermediary has only partial information about the buyer's values.

In contrast, in this paper we analyze the *profit* implications of uniform pricing versus third-degree price discrimination. We are particularly interested in understanding the approximation guarantees that a uniform price can deliver. Closest to our work is a paper by Malueg and Snyder (2006) which examines the profit effects of third-price discrimination compared to uniform pricing. They consider a setting similar to ours in which the monopolist experiences a total cost function for serving different segments. They show that when the demand is continuous and the total cost is superadditive, the ratio of third-degree price discrimination profit to uniform price profit is bounded above by the number of segments that are served under price discrimination. They provide an example under which this bound is tight and the bound for the ratio equals the total number of segments. In contrast, in the present paper we identify a key condition which leads to a bound that is not contingent on the number of segments in the market.

Their Proposition 2, adjusted to our setting, implies our Proposition 4. They also provide an example that attains the worst-case performance for distributions with different support (and linear demand). While their proof is inductive, we provide an alternative and constructive argument.

Since the seminal work of Myerson (1981), there has been great interest in simple and approximate mechanisms design. In general, characterizing optimal selling mechanisms is a difficult task, see e.g., Daskalakis et al. (2014) and Papadimitriou et al. (2016). Hence, deriving simple-practical mechanisms is of utmost importance. Chawla et al. (2007), Hartline and Roughgarden (2009), Alaei et al. (2019), and Jin et al. (2018), among others, have made remarkable progress toward establishing performance guarantees of simple mechanisms in a variety of settings. One of the key observations is that regular environments—non-decreasing virtual value—consistently lead to good bounds. In particular, triangular instances—instances for which the revenue functions in the *quantile space* are triangle-shaped (see e.g., Alaei et al. (2019))—are the worst-case in terms of performance guarantees. In contrast to this stream of literature, we consider the problem faced by a monopolist selling to many distinct segments of the market. Nevertheless, in line with these earlier papers, we aim to obtain performance guarantees when comparing the best possible pricing for the monopolist (third-degree price discrimination) to the simple pricing scheme (uniform pricing). In terms of techniques, as we discuss in the next paragraph, we resort to related triangular instances as worst-case performance settings, but we also establish arbitrarily poor performance in the regular case.

[2] In the related literature section we discuss the relation between this result and Malueg and Snyder (2006).

Our work also shares some similarities with the approach taken by Dhangwatnotai et al. (2015) (see also Hartline (2020) chapter 5) to study the prior-independent single sample mechanism. In this mechanism, bidders are allocated an object according to the VCG mechanism with reserves randomly computed from other bidders' bids. In a setting with n bidders, the authors establish that random pricing achieves half of the optimal profit. The authors expand this result to more complex settings and formalize it by using an intuitive geometric approach similar to the one we present in Section 3. In particular, under the assumption of regular distributions, the profit function in the quantile space turns out to be concave. Consequentially, the profit function is bounded below by a triangle with height equal to the maximum profit. This implies that the expected profit from uniformly selecting a quantile is bounded below by the area of the triangle or, equivalently, by half the maximum profit. One of the proofs we give for our result in Theorem 1 uses this observation to prove a different result, namely, that uniform pricing can deliver at least half the value of optimal third-degree price discrimination. However, we must assume that the profit functions are *concave in the price space*, otherwise our half approximation result might not hold. Indeed, in Proposition 6 we show that, in contrast to the aforementioned papers in our third-degree price discrimination setting, for regular distributions, simple pricing leads to arbitrarily poor guarantees.

2. Model

We consider a monopolist selling to K different customer segments. Each segment k is in proportion α_k in the market where $\alpha_k \geq 0$ for all $k \in \{1,\dots,K\}$ and $\sum_{k=1}^{K}\alpha_k = 1$. If the monopolist offers price p_k to segment k and has constant marginal cost $c \geq 0$, then the monopolist receives an associated profit of:

$$\pi_k(p_k) \triangleq (p_k - c)\cdot(1 - F_k(p_k)),$$

where $F_k(\cdot)$ is the cumulative distribution function of a distribution with support in $\Theta_k \subset \mathbb{R}_+$. We assume that $c \leq \sup\{\theta : \theta \in \Theta_k\}, \forall k$, i.e., the efficient allocation would generate sales with positive probability in every segment. The total profit the monopolist receives from the different segments by pricing according to a vector of prices $\boldsymbol{p} = (p_1,\dots,p_K)$ is

$$\Pi(\boldsymbol{p}) = \sum_{k=1}^{K}\alpha_k\pi_k(p_k).$$

The monopolist wishes to choose $\boldsymbol{p}$ to maximize $\Pi(\boldsymbol{p})$.

The monopolist can choose prices in different manners. First, for each segment k, the monopolist can set the price $p_k^\star$ where

$$p_k^\star \in \operatorname*{argmax}_{p\in\Theta_k}\pi_k(p).$$

Let $\boldsymbol{p}^\star$ be the vector of prices $\{p_k^\star\}_{k=1}^K$; we refer to these prices as the *per-segment optimal prices*. Note that $\boldsymbol{p}^\star$ corresponds to the case of third-degree price discrimination. We use $\Pi^\star$ to denote $\Pi(\boldsymbol{p}^\star)$. Another way of setting prices is to simply use a uniform price for all segments. In this case, the monopolist solves the problem

$$\Pi^U \triangleq \max_{p\in\cup_{k=1}^K\Theta_k}\sum_{k=1}^{K}\alpha_k\pi_k(p). \tag{1}$$

We use p_u to denote the optimal price in the above problem, which we refer to as the *optimal uniform price*. With some abuse of notation we sometimes use $\Pi^U(p)$ to denote $\Pi(\boldsymbol{p})$ when all the components of $\boldsymbol{p}$ are equal to p. We call the ratio between the best third-degree price discrimination scheme and the best uniform price scheme the *profit ratio*:

$$\frac{\Pi^U}{\Pi^\star}. \tag{2}$$

We use $\Pi^\star(\alpha,F)$ and $\Pi^U(\alpha,F)$ to make explicit the dependence of the monopolist profit on the model parameters (α,F): the segmentation, $\alpha = \{\alpha_k\}_{k=1}^K$, and the demand, $F = \{F_k\}_{k=1}^K$. Our main objective in this paper is to study how this ratio performs across a wide rage of parameter environments:

$$\inf_{\alpha,F}\frac{\Pi^U(\alpha,F)}{\Pi^\star(\alpha,F)}. \tag{$\mathcal{P}$}$$

3. Concave profit functions

In this section, we assume that the profit functions, $\pi_k(\cdot)$, are concave. We will further assume that the segments' supports, Θ_k, are identical across segments, that is, $\Theta_k = \Theta$ for all k where Θ is a closed and bounded interval $[0,\overline{\theta}]$ of $\mathbb{R}_+$. In later sections, we analyze ($\mathcal{P}$) under relaxed assumptions.

We now establish that a particularly simple uniform price, formed as the midpoint between the marginal cost c and the largest possible value $\overline{\theta}$,

$$p_s = \frac{c+\overline{\theta}}{2}, \tag{3}$$

can achieve half of the monopolist's profit.

Theorem 1 (***Uniform price is a half approximation***).
Suppose that the profit functions $\pi_k(p)$ are concave and defined in the same bounded interval $\Theta \subset \mathbb{R}_+$ for all $k \in \{1, \dots, K\}$. Then the uniform price p_s delivers a 1/2-approximation for the monopolist's profits.

Proof. We show that by setting $p_s = (c+\overline{\theta})/2$ the monopolist can obtain at least half the profit of third-degree price discrimination. Indeed, the concavity of $\pi_k(\cdot)$ ensures that:

$$\frac{\pi_k(p_s)}{p_s - c} \geq \frac{\pi_k(p_k^\star)}{p_k^\star - c} \text{ if } p_s \leq p_k^\star, \quad \text{and} \quad \frac{\pi_k(p_s)}{\overline{\theta} - p_s} \geq \frac{\pi_k(p_k^\star)}{\overline{\theta} - p_k^\star} \text{ if } p_s > p_k^\star. \tag{4}$$

By noticing that $(p_s - c)/(p_k^\star - c) \geq 1/2$, $(\overline{\theta} - p_s)/(\overline{\theta} - p_k^\star) \geq 1/2$, and $p_k^* \in [c, \overline{\theta}]$, we deduce that in either case $\pi_k(p_s) \geq \pi_k(p_k^\star)/2$ for all $k \in \{1, \dots, K\}$. Multiplying this inequality by α_k, adding it up over k, and observing that $\Pi^U \geq \Pi^U(p_S)$, we obtain the desired result:

$$\inf_{\alpha, F} \frac{\Pi^U(\alpha, F)}{\Pi^\star(\alpha, F)} \geq 1/2. \quad \square$$

Theorem 1 provides a fundamental guarantee of uniform pricing compared to the optimal third-degree price discrimination. In particular, the monopolist can simply use a judiciously chosen price across all customer segments to ensure half of the best possible profit from perfectly discriminating across the different segments in the market. Interestingly, it is possible to achieve this profit guarantee by setting a price that only uses the upper bound of the support and the marginal cost.[3] An implication of this is that half of the optimal profit can be secured by using only information about the upper end of segments' support, $\overline{\theta}$. While the simplicity of this pricing policy is appealing, the informational requirements might be too stringent. In practice, finding the upper bound of the support can be difficult because it can entail experimenting with high prices. We next investigate what other simple pricing policies can achieve the approximation guarantee in Theorem 1 and discuss their informational requirements.

In what follows, we provide a simple geometric argument as an alternative to the proof of Theorem 1. This alternative proof not only sheds light on different informational requirements of simple pricing policies that can achieve at least 1/2 performance, but also shows how ideas used in approximate mechanism design translate to our setting. Indeed, the next argument is similar to the one presented in Dhangwatnotai et al. (2015) and Hartline (2020) for concave profit functions in quantile space, i.e., for regular distributions.

Let

$$r_k \triangleq \alpha_k \pi_k(p_k^\star),$$

that is, r_k corresponds to the maximum profit the seller can obtain from the fraction α_k of customers in segment k. Note that $\Pi^\star$ equals $\sum_{k=1}^K r_k$. Since for each segment k the profit function is concave on Θ, we can lower bound it by a triangular-shaped function that we denote by $L_k(p)$ as depicted in Fig. 1 (a).

More precisely, we define the lower bound functions

$$L_k(p) \triangleq \begin{cases} \frac{r_k}{p_k^\star - c} \cdot (p - c), & \text{if } p \in [c, p_k^\star]; \\ \frac{r_k}{\overline{\theta} - p_k^\star} \cdot (\overline{\theta} - p), & \text{if } p \in [p_k^\star, \overline{\theta}]. \end{cases}$$

Observe that $\sum_{k=1}^K L_k(p)$ is a concave piecewise linear function that achieves its maximum at some $p \in \{p_k^\star\}_{k=1}^K$. We use Π^L to denote its maximum value. Then, it is easy to see that $\Pi^U \geq \Pi^L$ because $\sum_{k=1}^K L_k(p)$ lower bounds $\sum_{k=1}^K \alpha_k \pi_k(p)$, see Fig. 1 (b). Next, we argue that

$$\Pi^L = \max_{p \in \{p_1^\star, \dots, p_K^\star\}} \left\{ \sum_{k=1}^K L_k(p) \right\} \geq \frac{1}{2} \sum_{k=1}^K r_k = \frac{1}{2} \Pi^\star. \tag{5}$$

[3] We thank John Vickers for this suggestion.

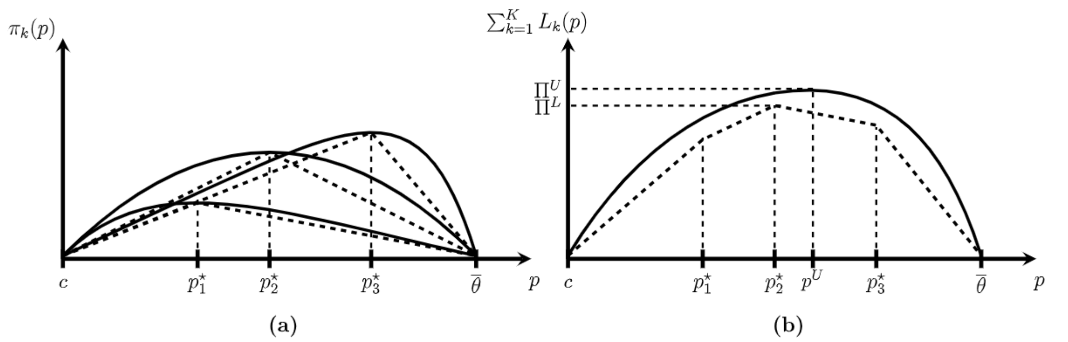

Fig. 1. **(a)** The solid curves depict the concave profit function of each segment, $\alpha_k \pi_k(p)$. The dashed lines depict the lower bounds $L_k(p)$ for each segment. **(b)** The solid curve shows the sum of the profit functions over segments, $\sum_{k=1}^K \alpha_k \pi_k(p)$. The dashed curve shows the sum of the lower bound over segments, $\sum_{k=1}^K L_k(p)$.

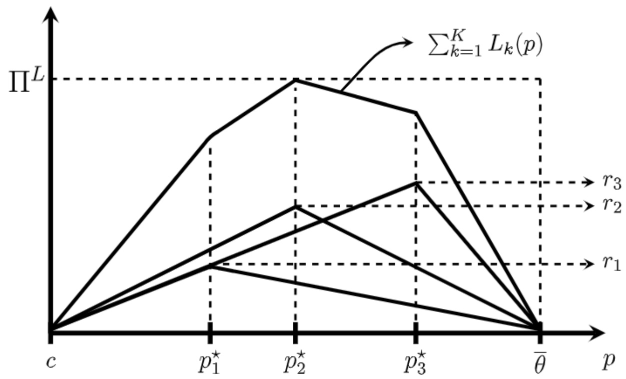

Fig. 2. Alternative geometric proof of Theorem 1: the area below $\sum_{k=1}^K L_k(p)$ equals the sum of the areas below $L_k(p)$ for all k.

Consider Fig. 2 and note that $\Pi^L \cdot (\overline{\theta} - c)$ is equal to the area of the smallest rectangle that contains the graph of $\sum_{k=1}^K L_k(p)$. As a consequence, $\Pi^L \cdot (\overline{\theta} - c)$ is an upper bound for the area below the curve $\sum_{k=1}^K L_k(p)$. That is,

$$\Pi^L \cdot (\overline{\theta} - c) \geq \int_c^{\overline{\theta}} \sum_{k=1}^K L_k(p) dp = \sum_{k=1}^K \int_c^{\overline{\theta}} L_k(p) dp = \sum_{k=1}^K \frac{r_k \cdot (\overline{\theta} - c)}{2},$$

where in the last equality we have used the fact that $L_k(p)$ is triangle-shaped and, therefore, the area below its curve equals $r_k \cdot (\overline{\theta} - c)/2$. Dividing both sides in the expression above by $(\overline{\theta} - c)$ yields (5), completing the proof. In particular,

$$\Pi^U = \max_{p \in \Theta} \left\{ \sum_{k=1}^K \alpha_k \pi_k(p) \right\} \geq \max_{p \in \{p_1^\star, \dots, p_K^\star\}} \left\{ \sum_{k=1}^K \alpha_k \pi_k(p) \right\} \geq \frac{1}{2} \Pi^\star. \tag{6}$$

This argument suggests two distinct yet simple ways of selecting the uniform price. First, the monopolist can optimize against the mixture of customer segments to derive the optimal uniform price. This is advantageous for situations in which the monopolist possesses aggregate market information but discriminating across segments is not an available option. When the monopolist has more granular market information, for example, the monopolist knows the prices $\{p_k^\star\}_{k=1}^K$, then it is not necessary for the monopolist to optimize over the full range of prices; he can simply choose one of the K prices at hand.

In sum, the two distinct arguments presented for Theorem 1 complement each other and point to different ways the monopolist has of achieving the performance guarantee of 1/2, depending on the information available. If the monopolist knows the upper bound of the support, setting $p_s = (c + \overline{\theta})/2$ is a simple choice. However, if such information is not available then there are two more options that can be directly inferred from the inequalities (6): *(i)* the monopolist can find the optimal uniform price for the aggregate demand or *(ii)* identify among the best per segment prices the uniform price that maximizes the revenue from the aggregate demand.

As mentioned in the introduction, our result and our second approach share some similarities with Dhangwatnotai et al. (2015). For the case of one buyer—and regular distributions—they show that the expected profit of randomly selecting a price achieves half of the optimal profit. Their approach uses the fact that the profit function in the quantile space for regular distributions is concave, and then proposes a uniform randomization over quantities. We can use a similar argument

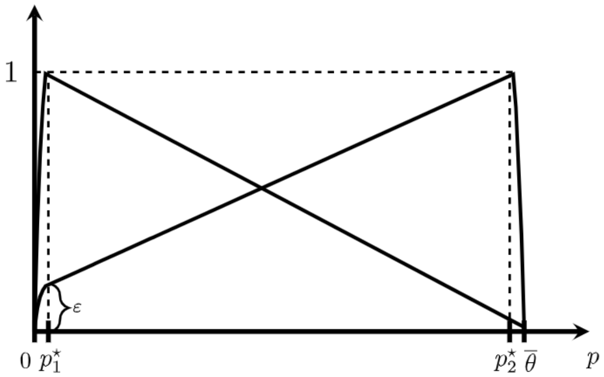

Fig. 3. Illustration of the construction in Proposition 2. Profit functions for both segments.

to show that the expected profit of uniformly choosing prices achieves half the profit of third-degree price discrimination. Indeed, suppose we set a price p at random such that $p \sim U[c, \overline{\theta}]$. Then, the expected profit is

$$\mathsf{E}[\Pi^U(p)] = \int_c^{\overline{\theta}} \sum_{k=1}^K \alpha_k \pi_k(p) \cdot \frac{1}{\overline{\theta} - c}\, dp \geq \frac{1}{\overline{\theta} - c} \cdot \int_c^{\overline{\theta}} \sum_{k=1}^K L_k(p)\, dp = \frac{1}{\overline{\theta} - c} \cdot \sum_{k=1}^K \frac{r_k \cdot (\overline{\theta} - c)}{2} = \frac{1}{2}\Pi^\star. \tag{7}$$

We summarize this discussion in the following proposition.

Proposition 1 *(**Uniformly at random pricing**).*
Suppose that the profit functions $\pi_k(p)$ are concave and have common and compact support. Then for $p \sim U[c, \overline{\theta}]$ we have that $\mathsf{E}_p[\Pi^U(p)]$ is at least half as large as $\Pi^\star$.

To conclude this section, we note that the above pricing policies all correspond to specific instances of simple pricing. Hence, it is still possible that the optimal uniform pricing achieves a better performance than the one established in Theorem 1. In the next proposition, we show that the latter is not possible by establishing that the profit guarantee in Theorem 1 is tight. To see why this is true, consider Fig. 3. There are two segments in the same proportion with maximum profit equal to 1. Assume that the profit of the first segment is very low at the price of the second segment and vice-versa. Then uniform pricing will achieve only the maximum profit of one of the segments, but very little of the profit from the other segment. In the figure, the best uniform pricing is $p_1^\star$ and it achieves $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \varepsilon$; while perfect price discrimination achieves $\frac{1}{2} \cdot 1 + \frac{1}{2} \cdot 1 = 1$. As ε becomes small, the ratio of the profits of optimal uniform pricing to third-degree price discrimination approaches $1/2$. In the proof of Proposition 2, we carefully construct the cumulative distribution functions to mimic the behavior in the above illustration.

Proposition 2 *(**Tightness**).*
Under the assumptions of Theorem 1

$$\underset{\alpha, F}{inf} \frac{\Pi^U(\alpha, F)}{\Pi^\star(\alpha, F)} = \frac{1}{2}.$$

Proof. Without loss of generality we take $c = 0$. Consider a case with two segments in the same proportion and the following distributions:

$$1 - F_1(p) \triangleq \begin{cases} \frac{1}{\lambda_1 p}\left(1 - e^{-\lambda_1 p}\right), & \text{if } p \leq a; \\ \frac{1}{p}\frac{(M-p)}{(M-a)}, & \text{if } p > a, \end{cases}$$

and

$$1 - F_2(p) \triangleq \begin{cases} \frac{1}{\lambda_2 p}\left(1 - e^{-\lambda_2 p}\right), & \text{if } p \leq a; \\ \frac{1}{p}\left\{\frac{1-\left(\frac{a}{M-a}+\varepsilon\right)}{(M-2a)}(p-a) + \left(\frac{a}{M-a}+\varepsilon\right)\right\}, & \text{if } p \in (a, M-a); \\ \frac{1}{p}\left\{-\frac{1}{a}(p - M + a) + 1\right\}, & \text{if } p \in [M-a, M], \end{cases}$$

with support $\Theta = [0, M]$ for some $M > 2a$ to be determined, and we choose $\lambda_k > 0$ such that $F_k(\cdot)$ is continuous at $a > 1$, and $\varepsilon > 0$ is a small parameter. Below, we will show how to choose the parameters to ensure that $F_k(\cdot)$ is a well-defined continuous distribution for $k \in \{1, 2\}$. First, we verify that π_k is concave for $k \in \{1, 2\}$. Indeed, for $k = 1$ we have

$$\pi_1(p) \triangleq \begin{cases} \frac{1}{\lambda_1}\left(1 - e^{-\lambda_1 p}\right), & \text{if } p \le a; \\ \frac{(M-p)}{(M-a)}, & \text{if } p > a. \end{cases}$$

The first piece of $\pi_1(\cdot)$ is increasing and concave while the second piece is decreasing and linear. The continuity of F_1 requires that $\pi_1(a) = 1$, equivalently, $(1 - e^{-a\lambda_1})/(a\lambda_1) = 1/a$. The latter always has a solution because the function $(1 - e^{-x})/x$ maps to the entire interval $(0, 1]$ and we are assuming $a > 1$. We conclude that $\pi_1(\cdot)$ is concave on Θ. Additionally, by taking the derivative of $1 - F_1(\cdot)$, it is possible to verify that $1 - F_1(\cdot)$ is decreasing as long as $M > a$.

We now verify the concavity of $\pi_2(\cdot)$. We have

$$\pi_2(p) \triangleq \begin{cases} \frac{1}{\lambda_2}\left(1 - e^{-\lambda_2 p}\right), & \text{if } p \le a; \\ \frac{1-\left(\frac{a}{M-a}+\varepsilon\right)}{(M-2a)}(p-a) + \left(\frac{a}{M-a} + \varepsilon\right), & \text{if } p \in (a, M-a); \\ -\frac{1}{a}(p - M + a) + 1, & \text{if } p \in [M-a, M]. \end{cases}$$

Note that as in the previous case, each piece of $\pi_2(\cdot)$ is concave. We can also verify that it is continuous. This is clear at $p = M - a$, and for $p = a$ we must choose λ_2 to ensure continuity. In this case, we must have $\pi_2(a) = \varepsilon + a/(M - a)$; equivalently, $(1 - e^{-a\lambda_2})/(a\lambda_2) = \varepsilon/a + 1/(M - a)$. Assuming that

$$\varepsilon/a + 1/(M - a) < 1, \tag{8}$$

we can always find λ_2 that makes $\pi_2(\cdot)$ continuous at $p = a$. In turn, $\pi_2(\cdot)$ linearly decreases in $[M - a, M]$, linearly increases in $(a, M - a]$ (because of (8)), and increases as a strictly concave function in $[0, a]$. Thus to ensure concavity we need to verify that the slope of $\pi_2(p)$ from the left of $p = a$ is larger than the slope from the right at $p = a$, that is,

$$e^{-\lambda_2 a} \ge \frac{1 - \left(\frac{a}{M-a} + \varepsilon\right)}{(M - 2a)}. \tag{9}$$

To see why (9) holds, note first that λ_2, which solves $(1 - e^{-a\lambda_2})/(a\lambda_2) = \varepsilon/a + 1/(M - a)$, is bounded above by the solution to $1/(a\tilde{\lambda}_2) = \varepsilon/a + 1/(M - a)$, that is, $\tilde{\lambda}_2 \ge \lambda_2$. Hence to verify (9) it suffices to show that

$$\exp\left(-\frac{1}{\varepsilon/a + 1/(M - a)}\right) \ge \frac{1 - \left(\frac{a}{M-a} + \varepsilon\right)}{(M - 2a)}. \tag{10}$$

Let us choose M such that $\varepsilon = 1/(M - a)^\kappa$ with $\kappa > 0$ to be determined, that is, $M - a = \varepsilon^{-1/\kappa}$. Then (10) becomes

$$\exp\left(-\frac{1}{\varepsilon/a + \varepsilon^{1/\kappa}}\right) \ge \frac{1 - \left(a\varepsilon^{1/\kappa} + \varepsilon\right)}{(\varepsilon^{-1/\kappa} - a)}.$$

Note that for $\varepsilon < 1$, as $\kappa \downarrow 0$ the left-hand side converges to $e^{-a/\varepsilon}$ while the right-hand side converges to 0. Hence we can always choose $\kappa \in (0, 1)$ such that (9) holds. Moreover, since $\varepsilon^{1/\kappa} < \varepsilon$ (because $\kappa < 1$), we have that (8) is always satisfied for ϵ small enough. In conclusion, we can always choose the parameter of our instance such that $\pi_2(\cdot)$ is concave. Additionally, by taking the derivative of $F_2(\cdot)$, it is possible to verify that it is decreasing as long as $M > 2a$ (which can also be achieved for $\kappa \in (0, 1)$ small).

Under perfect price discrimination we have that $\pi_k(p_k^\star) = 1$ for $k \in \{1, 2\}$, so that $\Pi^\star = 1$. Now, because both $\pi_1(\cdot)$ and $\pi_2(\cdot)$ increase in $[0, a]$ and decrease in $[M - a, a]$, the optimal uniform price must lie in $[a, M - a]$. In this interval both functions are linear and the derivative of their sum is $-\varepsilon/(M - 2a) < 0$. Hence the optimal uniform price is $p_u = a$, which yields

$$\Pi^U = \frac{1}{2} \cdot 1 + \frac{1}{2} \cdot \left(a\varepsilon^{1/\kappa} + \varepsilon\right) \le \frac{1}{2} + \frac{1}{2}\varepsilon(a + 1).$$

This implies that

$$\inf_{\alpha, F} \frac{\Pi^U(\alpha, F)}{\Pi^\star(\alpha, F)} \le \frac{1}{2} + \frac{1}{2}\varepsilon(a + 1).$$

Because $\varepsilon > 0$ is arbitrary, we conclude that the lower bound performance of 1/2 in Theorem 1 is tight. □

We note that Proposition 2 and Theorem 1 imply that the ratio in ($\mathcal{P}$) does not degrade too fast in the number of segments. Indeed, the worst performance can be achieved for the case $K = 2$, but as K increases the performance does not continue to degrade. This is a consequence of the concavity assumption. If we did not assume concavity we could have profit functions that take values close to zero around the optimal prices for other segments and a value of, for example, 1 at their own optimal per-segment prices. In this case, the ratio in ($\mathcal{P}$) would degrade at rate $1/K$ (see also Section 5).

4. Profit performance in general environments

In this section we examine weaker conditions relative to the environment studied in Section 3. In particular, we aim to understand how the profit ratio behaves when we relax the assumptions in Theorem 1. We first consider the assumption of compact support and then consider different supports across customer segments. Then, we study non-concave environments. In the latter, we are especially interested in common well-behaved environments such as regular and monotone hazard rate (MHR) value distributions. For ease of exposition and without loss of generality in this section and the following we will assume $c = 0$, unless otherwise stated.

4.1. Concave with unbounded support

In Theorem 1, we considered concave profit functions supported on some common finite interval Θ. In the next proposition, we relax the finite support assumption while keeping a common support and concave profit functions across customer segments.

Proposition 3 (***Zero profit gap with unbounded support***).
Suppose that the profit functions for all segments are concave with common and unbounded support $\Theta = \mathbb{R}_+$. Then $\Pi^U = \Pi^\star$.

Proof. Without loss of generality, assume that $p_1^\star \leq \cdots \leq p_K^\star$. Note that the concavity of the profit functions together with the unbounded support assumption causes each $\pi_k(p)$ to be increasing up to $p_k^\star$ and then constant and equal to $\pi_k(p_k^\star)$ for any price p larger than $p_k^\star$ for all $k \geq 1$. This leads to the following distributions:

$$1 - F_k(p) = \begin{cases} h_k(p)/p, & \text{if } p \leq p_k^\star; \\ \pi_k(p_k^\star)/p, & \text{if } p \geq p_k^\star, \end{cases}$$

for some increasing and concave function $h_k(\cdot)$ such that $\lim_{p\to 0} h_k(p)/p = 1$, $h_k(p)/p$ is decreasing, and $h_k(0) = 0$. In turn, by setting p_u equal to $p_K^\star$, the profit Π^U becomes $\sum_{k=1}^K \alpha_k \pi_k(p_k^\star) = \Pi^\star$. □

The proposition establishes that in the concave case with unbounded support there is no gap in the profit between no price discrimination and full price discrimination. The intuition behind Proposition 3 is simple. Concavity, together with the unbounded support assumption, implies that the marginal profit for each segment must equal zero for sufficiently large prices. As a consequence, setting an equal and sufficiently large price for every segment achieves the optimal third-degree price discrimination outcome.

4.2. Significance of common support

Here we consider concave profit functions supported on some finite interval Θ_k for each segment $k \geq 1$. In contrast to the previous section, we will not assume that $\Theta_k = \Theta$ for all segments. In order to gain intuition, note that, for an arbitrary distribution, the optimal revenue can be arbitrarily small compared to the expected surplus. Indeed, consider for example the distribution $F(v) = 1 - 1/(1 + v)$ for $v \geq 0$.[4] For this distribution, the optimal revenue $\sup_v\{v \cdot (1 - F(v))\}$ is 1 while the expected surplus is ∞. Hence, if we consider segments with point-mass distributions for every value v, the monopolist profit under uniform pricing (which would correspond to the optimal revenue for F) can be arbitrarily bad compared to optimal third-degree price discrimination—which would correspond to the expected surplus for F. In the next result, we make this intuition precise and show that there is a discrete collection of non-point-mass distributions with non-common support for which uniform pricing delivers arbitrarily small profit as the number of segments increases.

Proposition 4 (***Significance of common support***).
If the segments can have distinct supports then the optimal uniform price may yield an arbitrarily small profit ratio as the number of segments increases.

[4] Hartline and Roughgarden (2009) use this distribution to provide a lower bound on the worst revenue performance ratio of Vickrey with duplicated bidders and the optimal auction.

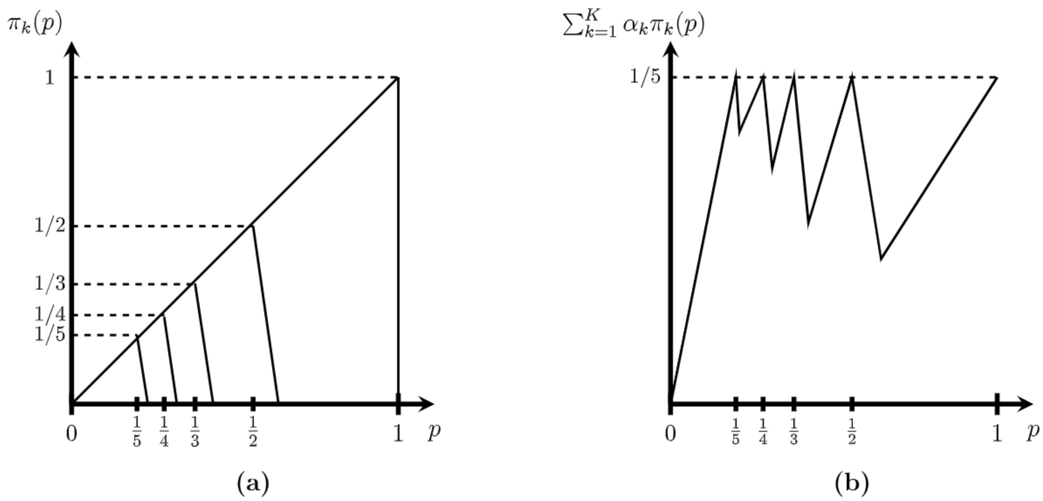

Fig. 4. Example for the construction of concave profit functions with different supports, as in the proof of Proposition 4 for $K = 5$. In **(a)** we illustrate the profit functions for each segment where the value and optimal per-segment price decay as $1/k$. In **(b)** we show $\sum_{k=1}^K \alpha_k \pi_k(p)$ which is maximized at any of the per-segment optimal prices and is bounded above by $1/K$.

Proof. We construct concave profit functions with finite support such that $\Theta_k \neq \Theta_j$ for all $k \neq j$. Let the distributions $\{F_k(\cdot)\}_k$ be defined by

$$F_k(p) = \begin{cases} 0, & \text{if } p \in [0, v_k]; \\ \frac{(p-v_k)(v_k+\varepsilon_k)}{\varepsilon_k p}, & \text{if } p \in (v_k, v_k + \varepsilon_k], \end{cases}$$

with

$$v_k = \frac{1}{(K-k+1)}, \quad \alpha_k = \frac{1}{K}, \quad \varepsilon_k \in (0, v_{k+1} - v_k), \quad \forall k \in \{1, \dots, K\}.$$

This leads to the following profit functions:

$$\pi_k(p) = \begin{cases} p, & \text{if } p \in [0, v_k]; \\ \frac{v_k}{\varepsilon_k}(v_k + \varepsilon_k - p), & \text{if } p \in (v_k, v_k + \varepsilon_k]. \end{cases}$$

The perfect price discrimination profit is

$$\Pi^\star = \sum_{k=1}^K \alpha_k v_k = \frac{1}{K} \sum_{k=1}^K \frac{1}{K-k+1} = \frac{1}{K} \sum_{k=1}^K \frac{1}{k}.$$

The optimal uniform price must be achieved at one of the v_k, hence

$$\Pi^U = \frac{1}{K} \max_{k=1,\dots,K} \left\{ \sum_{j=k}^K v_k \right\} = \frac{1}{K} \max_{k=1,\dots,K} \left\{ \frac{1}{K-k+1} \cdot (K-k+1) \right\} = \frac{1}{K}.$$

Hence,

$$\frac{\Pi^U}{\Pi^\star} = \frac{\frac{1}{K}}{\frac{1}{K}\sum_{k=1}^K \frac{1}{k}} = \frac{1}{\sum_{k=1}^K \frac{1}{k}} \approx \frac{1}{\log(K)} \to 0 \quad \text{as } K \uparrow \infty. \quad \square$$

Proposition 4 shows the importance of the common support assumption. Above we construct concave profit functions that have finite support, but the endpoints of the supports are increasing, see Fig. 4 **(a)**. In the construction, all segments are given the same weight and the profit functions are triangle-shaped. All of them start at zero, go up along the 45 degree line, peak at $1/(K-k+1)$, and then go down sharply such that the upper end of the support of segment k is strictly between $1/(K-k+1)$ and $1/(K-k)$ (solid lines in Fig. 4 **(a)**). In turn, $\Pi^\star$ (normalized by the per-segment proportions) grows logarithmically with K. Since the upper bound of the supports are strictly increasing and non-overlapping, the uniform price profit at the per-segment optimal prices, $\Pi^U(1/(K-k+1))$, is constant and equal to 1 (normalized by the per-segment proportions), see Fig. 4 **(b)**. For example, in the case of $K = 5$, consider $p_3^\star = 1/3$. At this price, $\pi_1(1/3) = \pi_2(1/3) = 0$

and $\pi_3(1/3) = \pi_4(1/3) = \pi_5(1/3) = 1/3$. Hence $5 \cdot \Pi^U = 0 + 0 + 3 \cdot \frac{1}{3} = 1$. As a result, the ratio in ($\mathcal{P}$) goes to zero as the number of segments increases. Note that this only works because the profit functions do not have common support. The non-common support allows for the possibility of having profit functions such that at the per-segment optimal prices some of them have zero profit. We note that a similar version of this result was stated in Malueg and Snyder (2006), Proposition 2. Their proof is inductive whereas our proof is constructive and provides a transparent view regarding the significance of the common support. Finally, we note that in the context of simple optimal auctions, Hartline and Roughgarden (2009) developed a related construction to argue the necessity of a single-item setting and regular distributions to obtain revenue guarantees for anonymous reserve price versus the optimal auction that do not scale with the number of bidders.

4.3. Non-concave environments

One of the most commonly analyzed families of distributions in the mechanism design literature is regular distributions. These are distributions such that the virtual value function is non-decreasing. Formally, we say that distribution F is regular if and only if $\phi(p)$ is non-decreasing:

$$\phi(p) \triangleq p - \frac{1 - F(p)}{f(p)}.$$

As pointed out in Section 1.2, several approximation guarantees have been obtained for these distributions in diverse settings. One of the main insights used in the literature is that the profit function associated with this family of distributions is concave in the quantile space. Indeed, let $\pi(p) = p \cdot (1 - F(p))$ and consider the change of variables $q = 1 - F(p)$. Define the profit function in the quantile space as $\widehat{\pi}(q) = q \cdot F^{-1}(1-q)$. Then

$$\frac{d}{dq}\widehat{\pi}(q) = F^{-1}(1-q) - \frac{q}{f(F^{-1}(1-q))} = \phi(F^{-1}(1-q)),$$

and since $\phi(\cdot)$ is non-decreasing, we can conclude that $\widehat{\pi}(q)$ is concave. The concavity of $\widehat{\pi}(q)$ allows arguments similar to the ones employed in the triangular proof of Theorem 1. For example, Dhangwatnotai et al. (2015) use this property to show that with one bidder, the expected profit from random pricing (uniformly selecting a quantile) is half the profit of the optimal monopoly price. This suggests that a similar approach may work in our framework. In particular, we ask whether it is possible with regular distributions to exploit the concavity of the profit functions to obtain good approximation guarantees.

Recall that for each segment $k \geq 1$, the profit function comes from a cdf F_k, for which we assume its pdf f_k is well-defined. To switch to the quantile space, for any $p \in \Theta$, we would need to define

$$q_k = 1 - F_k(p) \quad \text{and} \quad \widehat{\pi}_k(q) = q \cdot F_k^{-1}(q).$$

Let $\boldsymbol{q} = \{q_k\}_{k=1}^K$. The optimal uniform price profit, Π^U, is given by

$$\begin{aligned} \Pi^U = \max_{0 \leq \boldsymbol{q} \leq 1} & \sum_{k=1}^{K} \alpha_k \cdot \widehat{\pi}_k(q_k) \\ \text{s.t.} \quad & F_k^{-1}(1-q_k) = F_j^{-1}(1-q_j) \quad \forall k, j. \end{aligned}$$

Note that in this formulation the objective function is the sum of concave functions. However, we have additional constraints compared to the original formulation of Π^U in Section 2. These constraints stem from the fact that under uniform pricing each segment receives the uniform price p, and since $q_k = 1 - F_k(p)$ we must have that $F^{-1}(1-q_k) = F^{-1}(1-q_j)$ for all segments $k, j \geq 1$. At this point, the natural approach would be to lower bound each $\widehat{\pi}_k$ by a triangle-shaped function—similar to Fig. 1 **(a)**, but in the quantile space. We would then solve the resulting optimization problem and, hopefully, obtain a good approximation guarantee.

Unfortunately, for regular distributions, in general, the former approach fails. Consider the case of triangular instances in quantile space. These are instances for which the profit functions in the quantile space are triangle-shaped—they have corresponding distributions that are regular. They are widely used in the literature of approximate mechanism design as a bridge to provide good profit guarantees. However, in our setting they can perform arbitrarily poorly.

Proposition 5 (***Triangular instances and failure of regular distributions***).
For triangular instances defined by

$$F_k(p) = \begin{cases} 1, & \text{if } p \geq v_k; \\ \frac{p\cdot(1-q_k)}{p\cdot(1-q_k)+v_k\cdot q_k}, & \text{if } p < v_k, \end{cases}$$

there exists a choice of $\{\alpha_k\}_{k=1}^K \in (0,1)^K$, $\{v_k\}_{k=1}^K \in \mathbb{R}_+^K$ *and* $\{q_k\}_{k=1}^K \in (0,1)^K$ *that delivers an arbitrarily small profit ratio as the number of segments increases.*

Proof. Let us start by considering triangular instances. The profit functions are

$$\pi_k(p) = p \cdot (1 - F_k(p)) = \begin{cases} 0, & \text{if } p \geq v_k; \\ \frac{p \cdot v_k \cdot q_k}{p \cdot (1-q_k) + v_k \cdot q_k}, & \text{if } p < v_k. \end{cases}$$

Note that $\pi_k(p)$ is increasing and concave up to v_k and then is constant and equal to zero for $p \geq v_k$. For each curve the optimal price is v_k (minus small $\varepsilon > 0$), and thus

$$\Pi^\star = \sum_{k=1}^{K} \alpha_k \cdot v_k \cdot q_k. \tag{11}$$

For the uniform price, the optimal price must be achieved at one of the $v_1, \ldots, v_K$. Therefore,

$$\Pi^U = \max_{i \in \{1,\ldots,K\}} \left\{ \sum_{k=i}^{K} \alpha_k \cdot \frac{v_i \cdot v_k \cdot q_k}{v_i \cdot (1 - q_k) + v_k \cdot q_k} \right\}.$$

Next we establish that $\Pi^U / \Pi^\star \to 0$ as $K \to \infty$. Consider the instance

$$v_k = \frac{1}{(K - k + 1)}, \quad q_k = 0.5, \quad \text{and} \quad \alpha_k = \frac{1}{K}, \quad \forall k \in \{1, \ldots, K\}.$$

Hence,

$$\sum_{k=i}^{K} \alpha_k \cdot \frac{v_i \cdot v_k \cdot q_k}{v_i \cdot (1 - q_k) + v_k \cdot q_k} = \frac{1}{K} \sum_{k=i}^{K} \frac{\frac{1}{(K-i+1)} \cdot \frac{1}{(K-k+1)}}{\frac{1}{(K-i+1)} + \frac{1}{(K-k+1)}} = \frac{1}{K} \sum_{k=i}^{K} \frac{1}{2(K+1) - (k+i)}.$$

The last term above is decreasing in i, and therefore

$$\Pi^U = \frac{1}{K} \sum_{k=1}^{K} \frac{1}{2K + 1 - k} = \frac{1}{K} \sum_{k=1}^{K} \frac{1}{K + k} \approx \frac{1}{K} \int_1^K \frac{1}{K + x} dx = \frac{1}{K} \log\left(\frac{2K}{K + 1}\right).$$

We also have that

$$\Pi^\star = \sum_{k=1}^{K} \alpha_k \cdot v_k \cdot q_k = \frac{1}{2K} \sum_{k=1}^{K} \frac{1}{(K - k + 1)} = \frac{1}{2K} \sum_{k=1}^{K} \frac{1}{k} \approx \frac{1}{2K} \log(K).$$

Thus,

$$\frac{\Pi^U}{\Pi^\star} \approx \frac{\frac{1}{K} \log\left(\frac{2K}{K+1}\right)}{\frac{1}{2K} \log(K)} = 2 \frac{\log\left(\frac{2K}{K+1}\right)}{\log(K)} \to 0 \quad (\approx 2 \cdot \log(2)/\log(K)). \quad \square$$

Given that for general regular distributions we may obtain arbitrarily poor guarantees, we next investigate if such results can be improved upon by considering a commonly used sub-family of distributions with more structure. We consider distributions with monotone (non-increasing) inverse hazard rate $(1 - F_k(p))/f_k(p)$. Note that triangular distributions do not belong to this family as they an have increasing hazard rate. Interestingly, even distributions with monotone hazard rate can deliver only an arbitrarily small profit guarantee.

Proposition 6 *(**Failure of monotone hazard rate distributions**).*
There exist distributions with monotone inverse hazard rate and common and bounded support for which the optimal uniform price delivers an arbitrarily small profit ratio as the number of segments increases.

Proof. We construct regular distributions $\{F_k\}_{k=1}^K$ such that $\Pi^U / \Pi^\star \to 0$ as $K \uparrow \infty$. For $L > 0$ large, define

$$F_k(p) = \frac{1 - e^{-(K-k+1)p}}{1 - e^{-(K-k+1)L}} \quad \forall p \geq 0, \quad \text{and} \quad \alpha_k = 1/K \quad \forall k \in \{1, \ldots, K\}.$$

Thus, we consider truncated exponential distributions with support in $[0, L]$. Note that these distributions have monotone inverse hazard rate because

$$\frac{1 - F_k(p)}{f_k(p)} = \frac{1}{K - k + 1} - \frac{1}{(K - k + 1)} \cdot \frac{e^{-(K-k+1)L}}{e^{-(K-k+1)p}}$$

is non-increasing. For all $k \geq 1$, the profit functions are

$$\pi_k(p) = p \cdot (1 - F_k(p)) = p \cdot \frac{e^{-(K-k+1)p} - e^{-(K-k+1)L}}{1 - e^{-(K-k+1)L}},$$

whereas the per-segment optimal prices satisfy

$$p_k^\star = \frac{1 - e^{-(K-k+1)(L-p_k^\star)}}{K-k+1}, \quad \text{and} \quad \pi_k(p_k^\star) = \frac{1 - e^{-(K-k+1)(L-p_k^\star)}}{K-k+1} \cdot \frac{e^{-(K-k+1)p_k^\star} - e^{-(K-k+1)L}}{1 - e^{-(K-k+1)L}}.$$

First, notice that $p_k^\star \leq 1/(K-k+1)$. Then the third-degree price discrimination profit can be bounded as follows:

$$\begin{aligned}
\Pi^\star &= \sum_{k=1}^{K} \alpha_k \pi_k(p_k^\star) \\
&= \sum_{k=1}^{K} \frac{1}{K} \frac{e^{(K-k+1)p_k^\star}\left(e^{-(K-k+1)p_k^\star} - e^{-(K-k+1)L}\right)^2}{(K-k+1)(1 - e^{-(K-k+1)L})} \\
&\overset{(a)}{\geq} \sum_{k=1}^{K} \frac{1}{K} \frac{e^{1}\left(e^{-1} - e^{-(K-k+1)L}\right)^2}{K-k+1} \\
&\overset{(b)}{\geq} \frac{e\left(e^{-1} - e^{-L}\right)^2}{K} \sum_{k=1}^{K} \frac{1}{K-k+1} \\
&\approx \frac{e\left(e^{-1} - e^{-L}\right)^2}{K} \cdot \log(K),
\end{aligned}$$

where in (a) we used the fact that the function $e^{\lambda x}\left(e^{-\lambda x} - e^{\lambda L}\right)^2$ is decreasing for $x \in [0, L]$ and that $p_k^\star < L$. In (b) we used L large enough such that $L > 1$. The uniform price profit for some price p is

$$\begin{aligned}
\sum_{k=1}^{K} \alpha_k \pi_k(p) &= \frac{1}{K} \sum_{k=1}^{K} p \cdot \frac{e^{-(K-k+1)p} - e^{-(K-k+1)L}}{1 - e^{-(K-k+1)L}} \\
&\leq \frac{p}{K} \sum_{k=1}^{K} e^{-(K-k+1)p} \\
&= \frac{p}{K} \sum_{k=1}^{K} e^{-kp} \\
&= \frac{p}{K} \cdot \frac{1 - e^{-Kp}}{e^{p} - 1} \\
&\leq \frac{1 - e^{-Kp}}{K} \\
&\leq \frac{1}{K},
\end{aligned}$$

where the second to last inequality holds because we always have that $p + 1 \leq e^{p}$. With this, we can conclude that

$$\frac{\Pi^U}{\Pi^\star} = \frac{\max\limits_{p \geq 0}\left\{\sum_{k=1}^{K} \alpha_k \pi_k(p)\right\}}{\Pi^\star} \leq \frac{\frac{1}{K}}{\frac{e\left(e^{-1} - e^{-L}\right)^2}{K} \cdot \log(K)} = \frac{1}{e\left(e^{-1} - e^{-L}\right)^2 \cdot \log(K)} \to 0, \quad K \uparrow \infty. \quad \square$$

Proposition 6 establishes that for some monotone hazard rate distributions, uniform pricing can perform arbitrarily poorly compared to optimal third-degree price discrimination. The intuition behind this result is similar to that of Proposition 4. We consider exponential distributions such that at the optimal uniform price, most of the per-segment profits will be low, and therefore they will not contribute much to Π^U, see Fig. 5 (a). Since for exponentials, the associated profit functions decay quickly after they peak, they behave in a similar manner as the case of non-common support distributions where the upper end of the support is increasing. Indeed, in the proof of Proposition 6 we obtain a profit guarantee similar to that found in the proof of Proposition 4, namely, $O(1/\log(K))$. Finally, we note that to prove the proposition, we use truncated

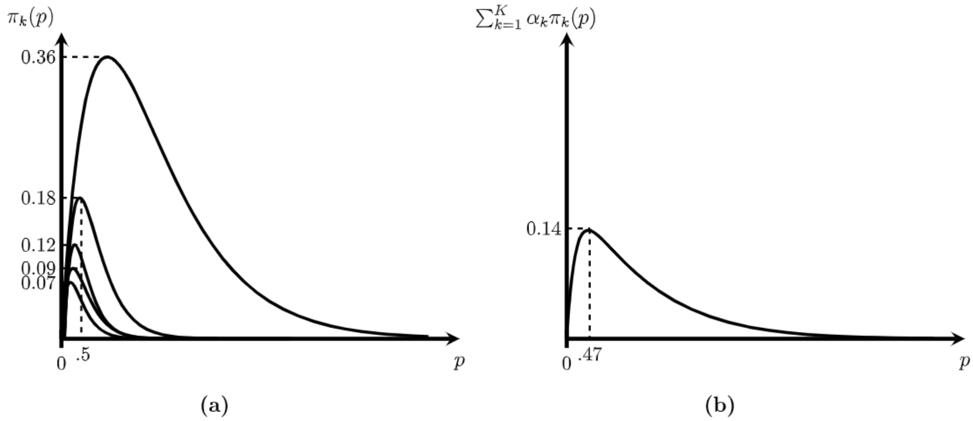

Fig. 5. Example for the construction of profit functions from regular distributions in Proposition 6 with $K = 5$. In **(a)** we illustrate the profit functions for each segment where the per-segment optimal profit is e^{-1}/k and the per-segment optimal price $1/k$. In **(b)** we show $\sum_{k=1}^{K} \alpha_k \pi_k(p)$ with maximum value 0.14 (bounded above by $1/K$).

exponential distributions. In turn, we only relax the concavity of the profit functions, but we keep the common and finite support assumptions intact.

5. Worst-case performance

In this brief section, our objective is to investigate how the worst case performance of uniform pricing depends on the number of segments. For this purpose, suppose there are K segments and without loss of generality let us assume that

$$0 \le \alpha_1 \pi_1(p_1^\star) \le \alpha_2 \pi_2(p_2^\star) \le \cdots \le \alpha_K \pi_K(p_K^\star). \tag{12}$$

Using this condition we can verify that the ratio $\Pi^U/\Pi^\star$ is always bounded below by $1/K$. Indeed, note that $\Pi^U \ge \Pi^U(p_K^\star)$ and

$$\Pi^U(p_K^\star) = \sum_{k=1}^{K} \alpha_k \pi_k(p_K^\star) \ge \alpha_K \pi_K(p_K^\star) = \frac{\overbrace{\alpha_K \pi_K(p_K^\star) + \cdots + \alpha_K \pi_K(p_K^\star)}^{K \text{ times}}}{K} \ge \frac{1}{K} \sum_{k=1}^{K} \alpha_k \pi_k(p_k^\star),$$

where in the last inequality we use (12). This proves that the worst-case performance of uniform pricing with respect to third-degree price discrimination is $1/K$. In what follows, we argue that this lower bound performance can indeed be achieved.[5]

We now suppose that the demand in every segment k is described by a Dirac distribution with an atom at value v_k. We denote by $\mathcal{D}$ the set of the segmentations and per segment demand functions that are generated by Dirac distributions, thus $\{\alpha_k, v_k\}_{k=1}^{K} \in \mathcal{D}$.

Under perfect price discrimination, the monopolist can charge the price $p_k = v_k$ to the buyer and extract full surplus: $\sum_{k=1}^{K} \alpha_k v_k$. Under uniform pricing, the monopolist charges a fixed price p across all segments and collects profit only from those segments whose value, v_k, is larger than p: $\sum_{k=1}^{K} \alpha_k p \mathbf{1}_{v_k \ge p}$. In turn, the optimization problem we would like to solve to assess the performance of uniform pricing becomes

$$\inf_{\alpha_k, v_k} \left\{ \frac{\max_{p \ge 0} \sum_{k=1}^{K} \alpha_k p \mathbf{1}_{v_k \ge p}}{\sum_{k=1}^{K} \alpha_k v_k}, \quad \text{s.t.} \quad \sum_{k=1}^{K} \alpha_k = 1, \;\; \alpha_k \ge 0 \;\; \forall k \right\}.$$

Observe that in the above problem, without loss of generality, we can assume that the values v_k are ordered. This allows us to simplify the numerator in the objective above. Note that the maximum in the numerator must be achieved at some price $p_j = v_j$ for $j \in \{1, \ldots, K\}$ and for any p_j we have

$$\sum_{k=1}^{K} \alpha_k p_j \mathbf{1}_{p_k \ge p_j} = p_j \sum_{k=j}^{K} \alpha_k.$$

[5] Proposition 2 in Malueg and Snyder (2006) provides an inductive argument for this result with linear demand functions. Here we provide a constructive argument with atomic distributions.

In turn, this enables us to reformulate the problem as

$$\inf_{\alpha_k, p_k} \left\{ \frac{\max\limits_{j\in\{1,\ldots,K\}} p_j \sum_{k=j}^{K} \alpha_k}{\sum_{k=1}^{K} \alpha_k p_k}, \quad \text{s.t.} \quad p_1 \leq \cdots \leq p_K, \quad \sum_{k=1}^{K} \alpha_k = 1, \ \alpha_k \geq 0 \ \forall k \right\}. \tag{13}$$

Next, we exhibit values of p_k and α_k such that the ratio in the problem above is arbitrarily close to the lower bound $1/K$. Let $\varepsilon > 0$ be small, and define the prices

$$p_k = \frac{\varepsilon}{K}\left(\frac{1+\varepsilon}{\varepsilon}\right)^k, \quad k \in \{1,\ldots,K-1\}, \quad \text{and} \quad p_K = \frac{1}{K}\left(\frac{1+\varepsilon}{\varepsilon}\right)^{K-1},$$

and let the per-segment proportions be

$$\alpha_k = \frac{1}{K}\frac{1}{p_k}, \quad k \in \{1,\ldots,K\}.$$

Next, we verify that the above prices and proportions are feasible. Indeed, it is easy to verify that p_k is increasing in k, while for the per-segment proportions we have that $\alpha_k > 0$ and

$$\sum_{k=1}^{K} \alpha_k = \sum_{k=1}^{K-1} \frac{1}{\varepsilon}\left(\frac{\varepsilon}{1+\varepsilon}\right)^k + \left(\frac{\varepsilon}{1+\varepsilon}\right)^{K-1} = 1.$$

Now let us look at the objective in problem (13). Given our choice of prices and proportions, we have that $\alpha_k p_k = 1/K$, and therefore the denominator in (13) equals 1. For the numerator consider $j \in \{1,\ldots,K-1\}$. Then

$$\begin{aligned} p_j \sum_{k=j}^{K} \alpha_k &= \frac{\varepsilon}{K}\left(\frac{1+\varepsilon}{\varepsilon}\right)^j \left(\sum_{k=j}^{K-1} \frac{1}{\varepsilon}\left(\frac{\varepsilon}{1+\varepsilon}\right)^k + \left(\frac{\varepsilon}{1+\varepsilon}\right)^{K-1}\right) \\ &= \frac{\varepsilon}{K}\left(\frac{1+\varepsilon}{\varepsilon}\right)^j \left(\frac{\varepsilon}{1+\varepsilon}\right)^{j-1} \\ &= \frac{1+\varepsilon}{K}, \end{aligned}$$

and $p_K \alpha_K = 1/K$. Therefore the objective in (13) evaluated at our current choice of prices and proportions equals

$$\max_{j\in\{1,\ldots,K\}} \left\{ p_j \sum_{k=j}^{K} \alpha_k \right\} = \max\left\{\frac{1}{K}, \frac{1+\varepsilon}{K}\right\} = \frac{1+\varepsilon}{K}.$$

In conclusion, we have exhibited an instance for which the performance of the optimal uniform price is arbitrarily close to the worst performance guarantee $1/K$. We summarize this result in the following proposition.

Proposition 7 (***Worst performance achieved***).
The worst case profit ratio in the class of Dirac demand functions is given by

$$\inf_{\{\alpha_k, v_k\}_{k=1}^{K} \in \mathcal{D}} \frac{\Pi^U}{\Pi^\star} = \frac{1}{K}.$$

6. Connection to sequential screening

The result in Theorem 1 is intimately related to the problem of ex-post individually rational sequential screening, in which the seller must optimally design a menu of contracts that incentivize buyers of different types to self-select, see e.g., Krähmer and Strausz (2015) and Bergemann et al. (2020). The connection between the two settings comes from considering the types in sequential screening as the segments of our paper and observing that optimal static pricing in the screening setting is the same as our optimal uniform pricing. Additionally, the optimal screening profit is bounded above by the optimal perfect price discrimination profit. As we show next, this allows us to prove a half approximation result for the sequential screening setting.

In what follows we formally introduce the sequential screening setting and then establish the connection with our third-degree price discrimination setting. There is a seller selling one unit of an object at zero cost to a buyer with an outside option of zero. The buyer is of type $k \in \{1,\ldots,K\}$, with probability α_k, $\alpha_k > 0$ and $\sum_{k=1}^{K} \alpha_k = 1$. Both parties are risk-neutral and have quasilinear utility functions. There are two periods. In the first period, the buyer privately learns her type k—which provides information about her true willingness-to-pay for the object—and then the parties contract. The contract specifies allocation and payment as a function of reported interim type and ex-post value. In the second period, the buyer privately learns her value θ—drawn from a distribution function $F_k(\theta)$ with density function $f_k(\theta)$—and allocations and transfers are realized.

We consider direct revelation mechanisms, with allocations $x_k : \Theta \to [0,1]$ and transfers $t_k : \Theta \to \mathbb{R}$, that depend on reported interim type k' and ex-post value θ'. Then the ex-post utility of a buyer who reported k in the first period and v' in the second period while her true value θ is given by:

$$u_k(\theta;\theta') \triangleq \theta \cdot x_k(\theta') - t_k(\theta').$$

Similarly, the interim expected utility of a buyer whose true interim type is k, but reported to the mechanism k' and is truthful in the second period, is given by:

$$U_{kk'} \triangleq \int_{\Theta} u_{k'}(z;z) \cdot f_k(z)dz.$$

There are two kinds of incentive compatibility constraints that must be satisfied. The first is the ex-post incentive compatibility constraint which requires that for any report in the first period, truth-telling is optimal in the second period:

$$u_k(\theta;\theta) \geq u_k(\theta;\theta') \quad \forall k \in \{1,\ldots,K\}, \forall \theta,\theta' \in \Theta.$$

The second is the interim incentive compatibility constraint which requires that truth-telling is optimal in the first period:

$$U_{kk} \geq U_{kk'} \quad \forall k,k' \in \{1,\ldots,K\}.$$

Finally, the key constraint in the sequential screening problem with ex-post participation constraints is that the buyer must be willing to participate after having learned her type and value:

$$u_k(\theta;\theta) \geq 0, \quad \forall k \in \{1,\ldots,K\}, \quad \forall \theta \in \Theta.$$

The seller aims to maximize the expected transfers from a mechanism that satisfies both incentive compatibility constraints and the ex-post participation constraint. Lemma 1 in Bergemann et al. (2020) implies the following reformulation of the seller's problem in which we only need to solve for the allocations and the utility of the lowest ex-post type. $\theta = 0$, of each interim type k, which is denoted by u_k, thus $u_k = u_k(0;0)$:

$$\begin{aligned} \Pi^{seq} \triangleq \max_{0\leq x_k \leq 1, u_k} \quad & -\sum_{k=1}^{K} \alpha_k u_k + \sum_{k=1}^{K} \alpha_k \int_{\Theta} x_k(z)\left(z - \frac{1-F_k(z)}{f_k(z)}\right) f_k(z)dz \\ \text{s.t} \quad & x_k(\cdot) \quad \text{is non-decreasing}, \quad \forall k \in \{1,\ldots,K\} \\ & u_k \geq 0, \quad \forall k \in \{1,\ldots,K\} \\ & u_k + \int_{\Theta} x_k(z)(1-F_k(z))dz \geq u_{k'} + \int_{\Theta} x_{k'}(z)(1-F_k(z))dz, \quad \forall k,k' \in \{1,\ldots,K\}. \end{aligned}$$

The first set of constraints comes from the ex-post incentive compatibility constraints in the original formulation. The second set of constraints ensures ex-post participation. The final constraints come from the interim incentive compatibility constraints.

To see the connection between the setting above and our third-degree price discrimination setting, first consider Π^{seq} without the interim incentive compatibility constraints. In this case the problem decouples across types, and for each type it reduces to:

$$\begin{aligned} \max_{0\leq x \leq 1} \quad & \int_{\Theta} x(z) \cdot \alpha_k \left(z - \frac{1-F_k(z)}{f_k(z)}\right) f_k(z)dz \\ \text{s.t} \quad & x(\cdot) \quad \text{is non-decreasing}. \end{aligned}$$

Note that $u_k = 0$ for all k is optimal in all decoupled problems. It is well known (see e.g., Riley and Zeckhauser (1983)) that the problem above has a bang-bang solution, say p_k. Therefore,

$$\Pi^{seq} \leq \sum_{k=1}^{K} \int_{p_k}^{\overline{\theta}} \alpha_k \left(z - \frac{1-F_k(z)}{f_k(z)}\right) f_k(z)dz = \sum_{k=1}^{K} \alpha_k \pi_k(p_k) \leq \sum_{k=1}^{K} \alpha_k \pi_k(p_k^{\star}) = \Pi^{\star}. \tag{14}$$

That is, the optimal solution in the sequential screening problem is bounded above by the optimal third-degree price discrimination solution.

Additionally, we can consider Π^{seq} in the case where the seller uses a static price—a price that is the same regardless of the buyer's interim type. In other words, we set $x_k(\cdot) = x(\cdot)$ for all k. In this case, after setting $u_k = 0$ for all k, the interim incentive compatibility constraints are directly satisfied, and the problem becomes:

$$\Pi^{static} \triangleq \max_{0 \le x \le 1} \int_{\Theta} x(z) \cdot \left(\sum_{k=1}^{K} \alpha_k \left(z - \frac{1 - F_k(z)}{f_k(z)} \right) f_k(z) \right) dz$$
$$\text{s.t} \quad x(\cdot) \quad \text{non-decreasing.}$$

As before, the optimal solution is bang-bang:

$$\Pi^{static} = \max_{p \in \Theta} \int_{p}^{\overline{\theta}} \sum_{k=1}^{K} \alpha_k \left(z - \frac{1 - F_k(z)}{f_k(z)} \right) f_k(z) dz = \max_{p \in \Theta} \sum_{k=1}^{K} \alpha_k \pi_k(p) = \Pi^U. \tag{15}$$

That is, the static optimal solution in the sequential screening problem coincides with the optimal uniform price solution.

Under the conditions of Theorem 1, conditions (14) and (15) imply that:

$$\frac{\Pi^{static}}{\Pi^{seq}} = \frac{\Pi^U}{\Pi^{seq}} \ge \frac{\Pi^U}{\Pi^{\star}} \ge 1/2.$$

We have thus established the following corollary of Theorem 1.

Corollary 1 (***Half approximation in sequential screening***).
Suppose that the assumptions of Theorem 1 hold. Then in the ex-post individually rational screening setting of Krähmer and Strausz (2015) and Bergemann et al. (2020), the optimal static contract delivers a 1/2-approximation for the seller's profits.

7. Conclusion

We consider the profit performance of uniform pricing in settings where a monopolist may engage in third-degree price discrimination. We establish that, for concave profit functions with common support, using a single price can achieve half of the optimal profit the monopolist could potentially garner by engaging in third-degree price discrimination. Our profit guarantee does not depend on the number of market segments or prices the seller might use. The different arguments we provide for Theorem 1 highlight that the monopolist can achieve the performance guarantee of 1/2 in different informational settings by using a simple uniform price.

We then investigate the scope of our profit guarantee. We establish that by relaxing either the concavity or the common support assumption, uniform pricing can lead to arbitrarily poor profit guarantees as the number of market segments increases. Interestingly, for regular distributions and triangular instances—leading cases in the literature of approximate mechanism design—we show that uniform pricing can again secure only very poor profit guarantees that scale with the number of market segments.

Depending on the nature of a market, different types of price discrimination are possible. A plausible direction for future work is to consider an environment in which the seller can exercise second-degree price discrimination by creating a menu of prices and quantities. For example, consider a setting with K different markets, each of which is characterized by a different distribution of valuations. Within each market the seller may offer an optimal menu of prices and quantities. However, due to legal or business constraints such powerful price discrimination may not be implementable. Instead, the seller might only be able to offer the same menu across all markets. In turn, it becomes a natural question to explore the performance of this limited menu versus the full discriminating one.

References

Aguirre, I., Cowan, S., Vickers, J., 2010. Monopoly price discrimination and demand curvature. Am. Econ. Rev. 100, 1601–1615.

Alaei, S., Hartline, J., Niazadeh, R., Pountourakis, E., Yuan, Y., 2019. Optimal auctions vs. anonymous pricing. Games Econ. Behav. 118, 494–510.

Bergemann, D., Brooks, B., Morris, S., 2015. The limits of price discrimination. Am. Econ. Rev. 105, 921–957.

Bergemann, D., Castro, F., Weintraub, G.Y., 2020. The scope of sequential screening with ex post participation constraints. J. Econ. Theory 188, 105055.

Chawla, S., Hartline, J.D., Kleinberg, R., 2007. Algorithmic pricing via virtual valuations. In: Proceedings of the 8th ACM Conference on Electronic Commerce. ACM, pp. 243–251.

Cummings, R., Devanur, N., Huang, Z., Wang, X., 2020. Algorithmic price discrimination. In: Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM, pp. 2432–2451.

Daskalakis, C., Deckelbaum, A., Tzamos, C., 2014. The complexity of optimal mechanism design. In: Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms. SIAM, pp. 1302–1318.

Dhangwatnotai, P., Roughgarden, T., Yan, Q., 2015. Revenue maximization with a single sample. Games Econ. Behav. 91, 318–333.

Hartline, J.D., 2020. Mechanism Design and Approximation. Book draft.

Hartline, J.D., Roughgarden, T., 2009. Simple versus optimal mechanisms. In: Proceedings of the 10th ACM Conference on Electronic Commerce. ACM, pp. 225–234.
Jin, Y., Lu, P., Qi, Q., Tang, Z.G., Xiao, T., 2018. Tight approximation ratio of anonymous pricing. arXiv preprint. arXiv:1811.00763.
Krähmer, D., Strausz, R., 2015. Optimal sales contracts with withdrawal rights. Rev. Econ. Stud. 82, 762–790.
Malueg, D.A., Snyder, C.M., 2006. Bounding the relative profitability of price discrimination. Int. J. Ind. Organ. 24, 995–1011.
Myerson, R., 1981. Optimal auction design. Math. Oper. Res. 6, 58–73.
Papadimitriou, C., Pierrakos, G., Psomas, C.-A., Rubinstein, A., 2016. On the complexity of dynamic mechanism design. In: Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms. Society for Industrial and Applied Mathematics, pp. 1458–1475.
Pigou, A., 1920. The Economics of Welfare. Macmillan, London.
Riley, J., Zeckhauser, R., 1983. Optimal selling strategies: when to haggle, when to hold firm. Q. J. Econ. 98, 267–290.
Robinson, J., 1933. The Economics of Imperfect Competition. Macmillan, London.
Schmalensee, R., 1981. Output and welfare implications of monopolistic third-degree price discrimination. Am. Econ. Rev. 71, 242–247.