THE NEAR-OPTIMALITY OF TWO-PART TARIFFS FOR NONLINEAR PRICING

By

Dirk Bergemann, Yang Cai, Jinzhao Wu

and Konstantin Zabarnyi

July 2026

COWLES FOUNDATION DISCUSSION PAPER NO. 2534

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS

YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.yale.edu/

The Near-Optimality of Two-Part Tariffs for Nonlinear Pricing*

Dirk Bergemann[†] Yang Cai[‡] Jinzhao Wu [§] Konstantin Zabarnyi[¶]

July 29, 2026

Abstract

The optimal mechanism for selling a divisible good under convex production costs can be arbitrarily complex, yet firms overwhelmingly use *two-part tariffs*: a fixed fee plus a constant markup over production cost. We quantify the profit this simplicity sacrifices. For regular value distributions, a two-part tariff guarantees a fraction of the optimal profit that depends only on a lower bound m on the elasticity of the marginal cost, independent of the value distribution. The guarantee approaches 1 as m grows, and is impossible without such a bound. Beyond regularity, no two-part tariff guarantees a constant fraction, but a menu of $K+1$ tariffs with a common markup does when the ironed virtual value has K ironing intervals; moreover, the menu size must scale with K. With multiple product lines, the better of separate sales and a single access fee granting purchases at production cost achieves a constant fraction of the optimal profit. These results provide a theoretical foundation for the ubiquity of simple cost-based pricing.

JEL CLASSIFICATION: D44, D47, D83, D84.

KEYWORDS: Nonlinear Pricing, Two-Part Tariff, Simple Contracts, Second-Degree Price Discrimination, Competitive Ratio, Approximation Guarantee.

*We gratefully acknowledge financial support from NSF SES 2519400 (DB), CCF 2342642 (YC and JW), and ONR MURI (DB). The authors are grateful to Tibor Heumann, Divyarthi Mohan, Stephen Morris, Haifeng Xu, and anonymous reviewers for their helpful comments and suggestions.

[†]Yale University | *E-mail*: dirk.bergemann@yale.edu.

[‡]Yale University | *E-mail*: yang.cai@yale.edu.

[§]Yale University | *E-mail*: jinzhao.wu@yale.edu.

[¶]Yale University | *E-mail*: konstantin.zabarnyi@yale.edu.

1 Introduction

1.1 Motivation

The theory of nonlinear pricing, going back to Mussa and Rosen (1978), instructs a monopolist facing privately informed buyers to offer an intricate price schedule—one whose shape tracks fine details of both the value distribution and the cost function. Practice tells a different story. Cloud providers, utilities, telecommunication carriers, and membership retailers overwhelmingly post *two-part tariffs*: a fixed fee plus a constant markup over production cost.[1] How much profit does this simplicity sacrifice? Our central finding is that the loss is modest and, as production costs become more convex, vanishes: the profit of the best two-part tariff converges to that of the fully optimal mechanism, uniformly over all regular value distributions. Simplicity is nearly free precisely in the markets where two-part tariffs are most widely used.

Those markets share the relevant feature of steeply increasing marginal costs. Cloud infrastructure costs rise sharply at high utilization; network congestion makes bandwidth increasingly costly at the margin; electricity, water, and gas display the same pattern. The consumed "amount" may also be quality rather than quantity: reducing a machine learning model's error rate from 5% to 2% requires far more effort than reducing it from 10% to 5%, and each additional "nine" of cloud uptime demands disproportionately greater investment. In all of these settings, our results imply that a fixed fee combined with a constant markup over cost captures nearly all of the attainable profit—and where a single tariff is not enough, we show that a small menu of them is.

A rich literature examines the tradeoff between simplicity and optimality in mechanism design (e.g., Hartline and Roughgarden, 2009; Chawla et al., 2010; Cai et al., 2016; Cai and Zhao, 2017; Babaioff et al., 2020, 2022), yet the fundamental setting of nonlinear pricing has remained largely unexplored since Wilson (1993), who studied how a finite menu of quantity–payment pairs can approximate the optimal mechanism. We revisit Wilson's question through the lens of the mechanism that firms actually use. Cost-indexed pricing of this form also has a long tradition: Bergemann et al. (2021a) show that pure markup mechanisms are robustly optimal under the welfare benchmark, and the fixed-fee-plus-cost-share contracts of Laffont and Tirole (1993), McAfee and McMillan (1986), and Bajari and Tadelis (2001) are exactly the class we analyze.

1.2 Results

We study a single-seller, single-buyer problem in which the seller produces a divisible good at a cost that increases convexly with quantity, and the buyer's value is private and drawn from a known distribution. In a *two-part tariff*, the buyer pays a *fixed fee* plus $(1 + \text{markup})$ times the production cost of any quantity he purchases, and participates if and only if his surplus exceeds the fee. The fixed fee corresponds to subscription, access, or connection charges; the markup is the per-unit margin over cost. The key parameter in our analysis is the *elasticity of the marginal cost function*, which quantifies the convexity of the cost structure; for *monomial* costs—total cost a power function of quantity—it is constant and

[1] See Chapter 3 of Tirole (1988) for the theory and many applications of two-part tariffs.

equals the exponent minus one. As a leading example, we show in Section 3 that for Pareto value distributions with monomial costs, two-part tariffs are exactly *profit optimal*, and that both components—fee and markup—are essential.

Regular distributions. For regular value distributions and any cost function whose marginal cost elasticity is bounded below by $m > 0$, setting the markup equal to m and choosing the fixed fee optimally guarantees a fraction of the optimal profit that depends only on m—independent of the value distribution—and improves further if the elasticity is also bounded from above (Theorem 1). In particular, for monomial cost functions, the approximation guarantee is always at least $1/e$. As $m \to \infty$, the guarantee approaches 1 even for sophisticated cost functions: two-part tariffs become nearly optimal exactly when production costs are highly convex. The argument decomposes the optimal profit into the production cost of the optimal allocation, which the markup covers, and a residual term that is linear in the allocation rule, which the fixed fee extracts. The guarantee is tight—for monomial costs no two-part tariff can do better (Proposition 4.1)—and the elasticity bound is necessary: without it, no two-part tariff attains a non-negligible fraction of the optimal profit (Proposition 4.2).

Non-regular distributions. Beyond regularity, simplicity has limits: for any $\varepsilon > 0$, there is a non-regular distribution and a monomial cost for which every two-part tariff achieves at most an ε-fraction of the optimal profit (Proposition 5.1). The reason is that a single fixed fee cannot price-discriminate across buyer types in different *ironing* intervals of the virtual value. A short *menu* restores the guarantee: when the ironed virtual value has K ironing intervals, a menu of $K+1$ two-part tariffs—sharing a common markup, with option-specific fixed fees and capacity limits—achieves a constant-factor approximation (Theorem 2 and Corollary 5.1). Each menu entry is tailored to one ironing interval, and the slack created by ironing allows fees and capacities to be set so that every type selects its intended entry. The required complexity thus scales with the number of ironing intervals—small or zero for many practical distributions—and this is unavoidable: any menu of size s can be forced to a negligible fraction of the optimal profit by a suitable non-regular distribution (Theorem 3).

Multiple product lines. With multiple product lines and independent values—a setting in which optimal mechanisms may require unboundedly large menus (Daskalakis et al., 2017)—we show that the optimal profit is bounded by a constant multiple of the sum of the profits from selling each line separately and from bundling all products under a single two-part tariff (Theorem 4). This generalizes the "selling separately or bundling is approximately optimal" theorem for indivisible goods (Babaioff et al., 2020) to divisible goods with convex costs; the proof adapts the duality framework of Cai et al. (2016) to accommodate production costs. When each product line satisfies the conditions above, the separate-sales mechanisms can themselves be two-part tariffs or menus thereof, so simple cost-based pricing captures essentially all achievable profit in multi-product settings as well.

Nonlinear buyer values. Finally, we extend the analysis beyond linear values. When the buyer's gross value is multiplicatively separable, $v\,U(q)$, a change of variables reduces the problem to the linear case with an induced cost function, and the guarantee of Theorem 1 carries over (Proposition 7.1). For a general supermodular value $u(v,q)$, the markup acts as an implicit *type contraction* rather than an exact rescaling; we formulate a conjecture on the resulting guarantee in which the elasticity floor m is replaced by an effective bound involving the elasticity of the marginal utility with respect to the type, and we explain why the general case is a challenging open problem.

1.3 Paper organization

Section 2 defines the framework. Section 3 analyzes the Pareto leading example. Sections 4 and 5 contain the results for regular and non-regular distributions, respectively, together with the underlying techniques. Section 6 treats the multi-product setting. Section 7 extends the analysis to nonlinear buyer values. Related work is discussed in Section 8, and Section 9 concludes.

2 Model

We consider a single-seller, single-buyer environment. The seller/producer/monopolist (she) can produce an unlimited quantity of a divisible good. Producing $q \in \mathbb{R}_{\geq 0}$ units incurs cost $C(q) \geq 0$. We focus on cost functions C that are strictly increasing, strictly convex, and in $C^1([0,\infty)) \cap C^2((0,\infty))$, with $C(0) = 0$ and $\lim_{q\to\infty} C'(q) = \infty$.[2]

The buyer (he) has a private *per-unit value* v, drawn from a known continuous distribution $F \in \Delta([\underline{v}, \overline{v}])$ with density $f(v)$, where the upper bound $\overline{v}$ may be infinite. The buyer's value is linear in quantity; that is, his value for receiving quantity q is vq.

If the buyer makes a payment $p \geq 0$ in exchange for quantity $q \geq 0$, his *utility* is $vq - p$, while the seller's *profit* is $p - C(q)$.

A *mechanism* is specified by a pair (x,p), where $x(v)$ denotes the quantity allocated to a buyer who reports value v, and $p(v)$ denotes the corresponding payment. We restrict attention to mechanisms that are *truthful (incentive compatible)* and *ex post individually rational.* That is, the following constraints must hold:

$$x(v)\,v - p(v) \geq x(v')\,v - p(v') \qquad \forall\, v, v' \in \mathrm{Supp}(F), \tag{IC}$$

$$x(v)\,v - p(v) \geq 0 \qquad \forall\, v \in \mathrm{Supp}(F). \tag{IR}$$

Our model assumes linear buyer values and convex production costs. This framework can be extended to more general environments via an appropriate change of variables, which is done in Section 7.

Simple mechanisms (a.k.a. two-part tariffs). For a given buyer's value v, define the maximum *social welfare* as

$$W(v) := \max_{q\geq 0}\{vq - C(q)\},$$

[2] Whenever C is not twice differentiable at 0, we shall refer to the right limit of C'' at 0, if it exists, as $C''(0)$.

and let $q(v)$ denote the corresponding welfare-maximizing quantity, i.e.,

$$q(v) \in \operatorname*{argmax}_{q \geq 0} \{vq - C(q)\}.$$

We study a class of simple mechanisms that charge the buyer through a fixed fee and a cost-based surcharge. Each mechanism is parameterized by a pair (T, β), where $T \geq 0$ is a *fixed fee* and $\beta \geq 0$ is a *markup* applied to the production cost. Upon participation, if the buyer purchases q units, his total payment is

$$T + (1+\beta) \cdot C(q).$$

The buyer is free to opt out of the mechanism; in that case, he purchases nothing and pays no fixed fee. If he chooses to participate, he selects a quantity

$$x_{T,\beta}(v) \in \operatorname*{argmax}_{q \geq 0} \left\{ vq - T - (1+\beta) \cdot C(q) \right\},$$

provided that the resulting utility is nonnegative.

Given $T, \beta \geq 0$, we use $\pi(T, \beta)$ to denote the expected profit under the two-part tariff with the fixed fee T and the markup β.

Our research focus. For a cost function C and a distribution F, let π^* be the expected optimal profit. In this paper, given a cost function C, we aim to find, for every distribution F, a two-part tariff $(T, \beta) = (T(F), \beta(F))$ for which:

$$\pi(T, \beta) \geq \rho \cdot \pi^*,$$

for some constant ρ independent of F (note that ρ can still depend on C, as the cost function is considered to be part of the setting). If this task is impossible, we broaden the class of simple mechanisms by considering *menus* of two-part tariffs (see Section 5).

For further detailed discussions about the properties of our model, see Appendix A.

3 A Leading Example: Values Drawn from a Pareto Distribution

We begin the analysis with the case in which the buyer's value is drawn from a Pareto distribution with a parameter $\lambda \in [1, \infty)$. That is, we consider distributions F satisfying

$$F(v) = 1 - v^{-\lambda}, \qquad \text{for } v \geq 1. \tag{1}$$

We further assume that the cost function is a monomial:

$$C(q) = \frac{1}{k} q^k, \qquad k > 1. \tag{2}$$

Before we describe the results of this section, we need to present some general properties of the social welfare, truthful mechanisms, and profit-maximizing mechanisms.

Social welfare. For a given buyer's value v, recall that the maximum social welfare is

$$W(v) := \max_{q \geq 0} \{vq - C(q)\},$$

and the corresponding welfare-maximizing quantity is denoted by $q(v)$. To ensure that the social welfare is finite, we impose the condition:

$$\lambda > \frac{k}{k-1} > 1 \Leftrightarrow k > \frac{\lambda}{\lambda - 1} > 1. \tag{3}$$

Note that since $\lambda > 1$, the distribution F is regular.

Properties of truthful mechanisms. Consider any truthful mechanism (x, p) with allocation rule $x(\cdot)$ and payment rule $p(\cdot)$. Observe that the incentive compatibility and individual rationality constraints (IC)–(IR) are identical to those in the unit-supply setting of Myerson (1981). By the same argument, any truthful mechanism must have a nondecreasing allocation rule $x(\cdot)$, and the associated payment rule is given by:[3]

$$p(v) = v\,x(v) - \int_{\underline{v}}^{v} x(t)\,dt. \tag{4}$$

For any mechanism (x, p), its expected profit can be expressed as:

$$\pi = \int_{\underline{v}}^{\overline{v}} f(v)\Big(x(v)\,\varphi(v) - C(x(v))\Big)\,dv,$$

where $\varphi(v) := v - \frac{1-F(v)}{f(v)}$ denotes the *virtual value*.

Profit maximization. We initially assume that the distribution F is regular – i.e., $\varphi(v)$ is nondecreasing in v. Mussa and Rosen (1978) characterize the optimal expected profit when selling a divisible good with nonzero production cost: for each $v \in [\underline{v}, \overline{v}]$, the seller should choose $x(v)$ to maximize

$$x(v)\,\varphi(v) - C(x(v)),$$

which, by definition, equals $W(\varphi(v))$. Hence, the optimal expected profit can be written as follows.

Fact 3.1 (Follows from Mussa and Rosen (1978))**.** *Given any regular distribution F, the optimal expected profit is given by*

$$\pi^* = \int_{\underline{v}}^{\overline{v}} f(v)\,W\big(\varphi(v)^+\big)\,dv,$$

where $\varphi(v)^+ = \max\{\varphi(v), 0\}$. For a general distribution F, an analogous relation holds, with $\varphi(v)$ replaced by the ironed virtual value of v.

[3] Here we normalize the buyer's utility for the lowest possible type $\underline{v}$ to be 0; this normalization is without loss of generality in our analysis, since it cannot decrease the profit of the mechanism.

A first-order condition. Fix any value v. Since $C(q)$ is convex and twice differentiable, the welfare-maximizing quantity $q(v)$ is characterized by the first-order condition that marginal value equals marginal cost. In particular, whenever $v \geq C'(0)$, we have:

$$C'(q(v)) = v. \tag{5}$$

Similarly, we get for $\frac{v}{1+\beta} \geq C'(0)$ (conditional on the buyer participating in the mechanism given the value v):

$$C'\big(x_{T,\beta}(v)\big) = \frac{v}{1+\beta}. \tag{6}$$

Back to the leading example. The results in this section illustrate that *both* components of a two-part tariff are essential. Specifically, we show that two-part tariffs exactly recover the profit-maximizing mechanism for this family of distributions. In contrast, if either the fixed fee or the markup is removed, the resulting class of mechanisms cannot achieve any constant-factor approximation to the optimal profit.

Our first result shows that, in the special case of Pareto distributions and monomial cost functions, two-part tariffs are profit-maximizing. The key point here is the linearity of the virtual value.

Proposition 3.1. *The two-part tariff (T,β) with*

$$\beta = \frac{1}{\lambda - 1} \quad \textit{and} \quad T = \frac{k-1}{k}\frac{1}{(1+\beta)^{1/(k-1)}} \tag{7}$$

is a profit-maximizing mechanism for the Pareto distribution F with parameter $\lambda > \frac{k}{k-1}$ and cost function $C(q) = \frac{1}{k}q^k$ $(k > 1)$.

Proof of Proposition 3.1. Let (x^*, p^*) be the (truthful) profit-maximizing mechanism. The first-order condition 5 and Fact 3.1 yield:

$$x^{*k-1}(v) = C'(x^*(v)) = \varphi(v) = (1 - 1/\lambda)v,$$

where $\varphi(\cdot)$ is the virtual valuation function.

Note that setting $\beta = \frac{1}{\lambda-1}$ and $T = \frac{k-1}{k}\frac{1}{(1+\beta)^{1/(k-1)}}$ yields exactly the above expression. Indeed, by the first-order condition 6, the allocation $x(\cdot)$ of the two-part tariff satisfies:

$$x^{k-1}(v) = C'(x(v)) = \frac{v}{1+\beta} = (1 - 1/\lambda)v.$$

We stress that this computation holds for every v because setting $T = \frac{k-1}{k}\frac{1}{(1+\beta)^{1/(k-1)}} = x(v)v - (1+\beta)C(x(v))|_{v=1}$ ensures that the buyer is willing to pay at least the fixed fee T for every v.

Since the payment is fully specified by the allocation through Equation 4, the conclusion follows. □

Our second result implies that to achieve constant-factor approximation guarantees, both the fixed fee and the markup are necessary. That is, setting $T = 0$ or $\beta = 0$ leads to an unbounded approximation ratio. Moreover, using a posted-price mechanism for selling a certain quantity x also cannot provide a constant approximation ratio. The proof is deferred to Appendix B.

Proposition 3.2. *Every one-part tariff yields vanishing small profits as either $k \to \infty$ or $k \to \frac{\lambda}{\lambda-1}^{+}$ with a Pareto distribution with parameter $\lambda > \frac{k}{k-1}$ and a monomial cost function $C(q) = \frac{1}{k}q^k$ with $k > 1$:*

1. *When $k \to \infty$, setting $T = 0$ yields $\pi(T = 0, \beta(k)) = o(\pi^*)$, where $\beta(k)$ is an arbitrary non-negative function of k.*
2. *When $k \to \frac{\lambda}{\lambda-1}^{+}$, setting $\beta = 0$ yields $\pi(T(k), \beta = 0) = o(\pi^*)$, where $T(k)$ is an arbitrary non-negative function of k.*
3. *When $k \to \frac{\lambda}{\lambda-1}^{+}$, for every mechanism that sells a fixed number $x \geq 0$ of units for a posted price-per-unit $p \geq 0$, the profit is $o(\pi^*)$.*

4 A General Environment with Regular Value Distributions

In this section, we demonstrate that when the cost function satisfies certain structural properties, a simple mechanism that combines a markup with a fixed fee suffices to approximate the optimal profit under a regular distribution.

Cost function and elasticity. A helpful object is the notion of *cost elasticity*. Elasticity is a widely used notion in economics that measures the responsiveness of one economic variable to change of another variable in a unit-free representation (see Chapter 7, Sydsæter and Hammond (2008)). Throughout most of the paper, we focus on cost functions whose marginal cost exhibits *bounded elasticity*. In our context, it measures the percentage change in the marginal cost $C'(q)$ induced by a percentage change in the produced quantity q.

Definition 1. *The elasticity of the marginal cost function $C'(q)$ with respect to quantity q is defined as*

$$\eta(q) := \frac{qC''(q)}{C'(q)}.$$

Throughout the paper (apart from Subsection 4.1), we assume that the elasticity of the marginal cost is uniformly bounded away from zero: there exists a constant $m > 0$ such that

$$\eta(q) \geq m, \qquad \forall q \geq 0.$$

Note that it implies $C'(0) = 0$, and thus the first-order conditions 5 and 6 hold for every v.

To improve our results, we shall sometimes further assume that the elasticity is also uniformly bounded away from infinity: there exist $M > 0$ such that

$$\eta(q) \leq M, \qquad \forall q \geq 0.$$

A particularly natural special case is given by monomial cost functions of the form

$$C(q) = \frac{1}{k}q^k,$$

for some constant $k > 1$. For this class, the elasticity of the marginal cost is constant: for all $q \geq 0$, we have $\eta(q) = k - 1$.

Our main result establishes that for any cost function with bounded elasticity and any regular value distribution, a simple two-part tariff mechanism, comprising a fixed fee T and a markup β, is sufficient to approximate the optimal profit.

Theorem 1. *Suppose the cost function $C(q)$ satisfies the elasticity lower bound $\eta(q) \geq m$ for all $q \geq 0$. Then the universal markup choice $\beta := m$ has the following guarantee: for every regular buyer distribution F, there exists a fixed fee T such that*

$$\pi(T,\beta) \geq \frac{m}{(m+1)^{1+1/m}}\pi^*.$$

Moreover, if C also satisfies the elasticity upper bound $\eta(q) \leq M$ for all $q \geq 0$, then the guarantee improves to

$$\pi(T,\beta) \geq \frac{m}{(m+1)^{1+1/m}} \cdot \frac{M+1}{M}\pi^*.$$

The approximation ratio that only depends on the elasticity lower bound m, i.e. $\frac{m}{(m+1)^{1+1/m}}$, is increasing in m. It tends to 0 as $m \to 0$, while it converges to 1 as $m \to \infty$. Thus, a constant approximation is guaranteed whenever the elasticity lower bound is bounded away from zero. In Proposition 4.2, we show that this condition is necessary: there exists an instance with regular distributions and cost functions whose minimum elasticity is 0, for which no two-part tariff achieves any constant approximation.

As an important special case, the monomial cost function $C(q) = \frac{1}{k}q^k$ belongs to the family of bounded elasticity marginal cost functions with $m = M = k-1$. By substituting this expression into Theorem 1, we obtain an approximation ratio of $\frac{1}{k^{1/(k-1)}}$. In Proposition 4.1, we further show that this ratio is tight.

Proposition 4.1. *Let the cost function be $C(q) = \frac{1}{k}q^k$ for some $k > 1$. Fix the markup $\beta := k-1$. Then for every regular value distribution F, there exists a fixed fee T such that the two-part tariff mechanism (T,β) achieves at least a fraction of $\frac{1}{k^{1/(k-1)}}$ of the optimal profit. Formally,*

$$\pi(T,\beta) \geq \frac{1}{k^{1/(k-1)}}\pi^*.$$

Moreover, this ratio is tight: for every $\varepsilon > 0$, there exist $\ell \in (0,1)$ and a regular distribution F supported on $[\ell,1]$ such that F is continuous on $[\ell,1]$, admits a strictly positive density on $[\ell,1]$, and every two-part tariff (T,β) satisfies

$$\pi(T,\beta) \leq \left(\frac{1}{k^{1/(k-1)}} + \varepsilon\right)\pi^*.$$

Remark 4.1. *The approximation ratio $\frac{1}{k^{1/(k-1)}}$ is increasing in k, with $\lim_{k\to1^+} \frac{1}{k^{1/(k-1)}} = \frac{1}{e}$ and $\lim_{k\to\infty} \frac{1}{k^{1/(k-1)}} = 1$. The real line in Figure 1 illustrates the dependence of the ratio on k.*

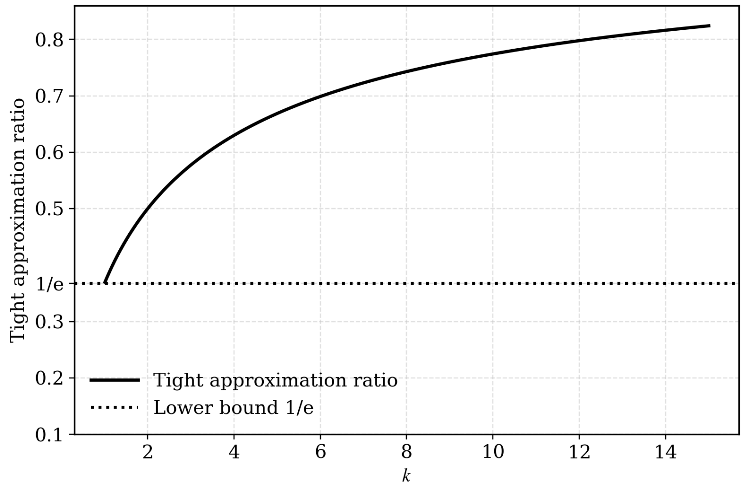

Figure 1: Approximation ratio as a function of $k > 1$ for $C(q) = \frac{1}{k}q^k$.

Our proof of the positive result follows a two-step approach. We first derive an analytically tractable upper bound on the optimal profit, and then show that the profit achieved by simple mechanisms matches this benchmark up to a constant factor.

We begin by sketching the main ideas of the proof; the complete argument is deferred to Appendix C. In the following, we assume that the cost function is monomial, namely,

$$C(q) = \frac{1}{k}q^k$$

for some constant $k > 1$. Under this cost function, the welfare-maximizing quantity and the corresponding social welfare admit closed-form expressions:

$$q(v) = \operatorname*{argmax}_{q \geq 0} \left\{ vq - \frac{1}{k}q^k \right\} = v^{\frac{1}{k-1}}, \qquad W(v) = v\,q(v) - C(q(v)) = \left(1 - \frac{1}{k}\right) v^{\frac{k}{k-1}}.$$

A benchmark for the optimal profit. As our first step, we begin by deriving an analytically convenient upper bound on the optimal profit. The following upper bound works for any cost function $C(q)$.

By Fact 3.1, the optimal expected profit can be written as

$$\pi^* = \int_{v_0}^{\infty} f(v)\,W(\varphi(v))\,dv,$$

where

$$v_0 := \inf\{v \mid \varphi(v) \geq 0\}.$$

Recall that under monomial costs $C(q) = \frac{1}{k}q^k$, the welfare function satisfies:

$$W(\zeta) = \zeta\,q(\zeta) - C(q(\zeta)) = \left(1 - \frac{1}{k}\right)\zeta\,q(\zeta).$$

Substituting this expression into the above integral yields

$$\begin{aligned}\pi^* &= \int_{v_0}^{\infty} f(v)\,(\varphi(v)\,q(\varphi(v)) - C(q(\varphi(v))))\,dv \\ &= \left(1-\frac{1}{k}\right)\int_{v_0}^{\infty} f(v)\,\varphi(v)\,q(\varphi(v))\,dv.\end{aligned}$$

We further upper bound this expression by observing that $\varphi(v) \le v$ for all $v \ge v_0$, and that $q(\cdot)$ is nondecreasing. It follows that

$$\pi^* \le \left(1-\frac{1}{k}\right)\int_{v_0}^{\infty} f(v)\,\varphi(v)\,q(v)\,dv.$$

Recall that the virtual value satisfies

$$\varphi(v) = v - \frac{1-F(v)}{f(v)}.$$

Substituting this into the upper bound derived above, we obtain

$$\begin{aligned}\pi^* &\le \left(1-\frac{1}{k}\right)\int_{v_0}^{\infty} f(v)\left(v - \frac{1-F(v)}{f(v)}\right)q(v)\,dv \\ &= \left(1-\frac{1}{k}\right)\int_{v_0}^{\infty} (f(v)v - (1-F(v)))\,q(v)\,dv.\end{aligned}$$

By swapping the order of integration,

$$\begin{aligned}\int_{v_0}^{\infty}(1-F(v))\,q(v)\,dv &= \int_{v_0}^{\infty}\left(\int_{v}^{\infty} f(t)\,dt\right)q(v)\,dv \\ &= \int_{v_0}^{\infty} f(t)\left(\int_{v_0}^{t} q(v)\,dv\right)dt.\end{aligned}$$

Plugging this back yields

$$\pi^* \le \left(1-\frac{1}{k}\right)\int_{v_0}^{\infty} f(v)\left(v\,q(v) - \int_{v_0}^{v} q(t)\,dt\right)dv.$$

Finally, note that

$$W(v) + C\big(q(v)\big) = v\,q(v) - C\big(q(v)\big) + C\big(q(v)\big) = v\,q(v).$$

Therefore,

$$\pi^* \le \left(1-\frac{1}{k}\right)\int_{v_0}^{\infty} f(v)\left(W(v) + C\big(q(v)\big) - \int_{v_0}^{v} q(t)\,dt\right)dv.$$

To further simplify this, we introduce the following lemma.

Lemma 4.1. *For any cost function C and any $v \in Supp(F)$, it holds that*

$$W(v) = \int_{0}^{v} q(t)\,dt.$$

Proof. Notice that the following holds for any v:

$$W'(v) = q'(v)v + q(v) - C'(q(v))q'(v) = q(v).$$

where the last equality follows from the fact that $C'(q(v)) = v$ by the first-order condition 5. Combining this with $W(0) = 0$ finishes the proof. $\square$

Plugging this expression back yields our upper bound on the optimal profit. This bound decomposes naturally into two components: the first term exhibits like a revenue from some fixed fee, while the second term can be covered as profit generated by the markup:

$$\pi^* \leq \left(1 - \frac{1}{k}\right) \int_{v_0}^{\infty} f(v) \left(\underbrace{W(v_0)}_{\text{fixed fee}} + \underbrace{C\big(q(v)\big)}_{\text{markup}} \right) dv.$$

Approximating the benchmark. We now construct a simple mechanism that approximates this benchmark. Consider a two-part tariff with markup $\beta = k - 1$ and fixed fee $k\,W\left(\frac{v_0}{k}\right)$.

Fix a buyer with value v, and first ignore the fixed fee. The buyer then chooses a quantity t that maximizes

$$tv - kC(t) = k\left(t\frac{v}{k} - C(t)\right),$$

which yields the optimal quantity $q\left(\frac{v}{k}\right)$. The buyer's resulting utility (before paying the fixed fee) is therefore

$$q\left(\tfrac{v}{k}\right)v - kC\left(q\left(\tfrac{v}{k}\right)\right) = k\,W\left(\tfrac{v}{k}\right).$$

The seller additionally charges a fixed fee of $k\,W\left(\frac{v_0}{k}\right)$. Since $W(\cdot)$ is monotone, the buyer participates if and only if $v \geq v_0$. Hence, the expected profit of this two-part tariff is

$$\pi = \int_{v_0}^{\infty} f(v) \left(\underbrace{k\,W\left(\tfrac{v_0}{k}\right)}_{\text{profit from fixed fee}} + \underbrace{(k-1)C\left(q\left(\tfrac{v}{k}\right)\right)}_{\text{profit from markup}} \right) dv.$$

To compare π with π^*, it therefore suffices to compare the two corresponding components term by term. A direct computation shows that

$$k\,W\left(\tfrac{v_0}{k}\right) \geq \frac{1}{k^{\frac{1}{k-1}}}\left(1 - \frac{1}{k}\right)W(v_0),$$
$$(k-1)C\left(q\left(\tfrac{v}{k}\right)\right) \geq \frac{1}{k^{\frac{1}{k-1}}}\left(1 - \frac{1}{k}\right)C(q(v)).$$

Combining these bounds, we conclude that the two-part tariff achieves an approximation ratio of $\frac{1}{k^{\frac{1}{k-1}}}$ for monomial cost functions.

To extend the argument beyond monomial costs, observe that the benchmark derived above remains essentially unchanged: its derivation does not quite rely on the specific functional form $C(q) = \frac{1}{k} q^k$. For a general cost function and any $\beta > 0$, we have

$$(1+\beta)\, W\Big(\tfrac{v_0}{1+\beta}\Big) = (1+\beta) \int_0^{\frac{v_0}{1+\beta}} q(t)\, dt = \int_0^{v_0} q\Big(\tfrac{t}{1+\beta}\Big)\, dt.$$

Consequently, to show that the profit achieved by a two-part tariff is comparable to the optimal profit, it suffices to identify a choice of β such that both

$$\frac{q\Big(\frac{v}{\beta+1}\Big)}{q(v)} \quad \text{and} \quad \frac{\beta\, C\Big(q\Big(\frac{v}{\beta+1}\Big)\Big)}{C(q(v))}$$

are bounded by constants independent of v. Thus, the main technical challenge is to obtain sufficient control over the function $q(\cdot)$ to ensure that these ratios remain constant.

The following lemma shows that such control is guaranteed when the marginal cost has bounded elasticity. The scaling bound controls how allocations change under a markup, while the cost-share bounds are used to tighten the benchmark upper bound on the optimal profit.

Lemma 4.2. *Fix a cost function C whose marginal cost C' satisfies*

$$m \le \eta_{C'}(q) \le M, \qquad \forall q > 0.$$

Recall that $q(v)$ denotes the welfare-maximizing production quantity for a buyer of value v under cost function C, i.e.,

$$q(v) := \operatorname*{argmax}_{q \ge 0} \{vq - C(q)\}.$$

Then the following statements hold for every v:

- *For every $s \ge 1$, $q\big(\frac{v}{s}\big) \ \ge\ s^{-1/m}\, q(v)$.*
- $\frac{1}{M+1}\, v\, q(v) \le C(q(v)) \le \frac{1}{m+1}\, v\, q(v)$.

4.1 Necessity of the lower bound on the elasticity

In this subsection, we prove that if the elasticity of the marginal cost is not bounded away from 0, then two-part tariffs cannot achieve any constant-factor approximation for the optimal profit. That is, the assumption $m > 0$ is necessary to derive positive results.

Proposition 4.2. *Fix the cost function $C(q) = e^q - 1$. Then for every $\varepsilon > 0$, there exists a regular F with Supp$(F) \subseteq [1,2]$ for which every two-part tariff (T, β) satisfies:*

$$\pi(T,\beta) \ \le\ \varepsilon \pi^*.$$

Note that $C'(q) = C''(q) = e^q$, and thus $\eta_{C'}(q) = q$. In particular, $\eta_{C'}(q = 0) = 0$, which corresponds by the first-order condition 5 to $v = 1$. In contrast, $\eta_{C'}(q)$ is bounded from above by $q(v = 2) = \log 2$.

The proof is based on looking at a variant of the truncated equal-revenue distribution over $W(v)$ in which the support size depends on ε. As a thought experiment, consider first a weaker version of the statement in which the distribution F is allowed to depend on the markup β. A natural candidate for the "hard" distribution F is a variant of the truncated equal-revenue distribution over $z_\beta(v) := v\, q\left(\frac{v}{1+\beta}\right) - (1+\beta)C\left(q\left(\frac{v}{1+\beta}\right)\right)$ with the support depending on ε, since such a distribution will make the expected profit $\pi(T,\beta)$ independent of T (if we ignore "unreasonable" choices of T for which the expected profit is clearly suboptimal). Since in reality F must be independent of β, an educated guess is to check a truncated equal-revenue distribution over $W(v)$, with the support depending on ε.

A technical challenge arises since the profit-maximizing mechanism is very complex. To circumvent this issue, we guess another mechanism which outperforms two-part tariffs by an unbounded factor. A possible guess is the welfare-maximizing mechanism. Unfortunately, in this mechanism the buyer extracts the full surplus. However, multiplying the allocation $x(v)$ and the price $p(v)$ in the welfare-maximizing mechanism by an appropriate constant (which turns out to be 0.5) provides the desired benchmark mechanism. For the full proof, see Appendix D.

5 General Environment with Non-Regular Distributions

In this section, we extend our analysis from regular to general distributions. As a natural starting point, one may ask whether a constant-factor approximation can still be achieved using the same class of mechanisms, which impose a markup together with a fixed fee, as in the regular case. The following result shows that even under the canonical monomial cost function $C(q) = \frac{1}{k}q^k$, such a guarantee is impossible: no choice of a markup and a fixed fee can ensure a constant-factor approximation for non-regular distributions.

Proposition 5.1. *For any $\varepsilon > 0$, there exists an instance with a non-regular buyer's value distribution F and a monomial cost function $C(q) = \frac{1}{k}q^k$ for some $k > 1$, such that for every choice of fixed fee T and markup parameter β, the (T,β)-mechanism attains at most an ε-fraction of the optimal profit. Moreover, the exponent k of the cost function goes to infinity as $\varepsilon \to 0$.*

Proposition 5.1 demonstrates that a single combination of a markup and a fixed fee is insufficient to recover any constant fraction of the optimal profit when the buyer's value distribution is non-regular, even under the simplest nontrivial form of cost functions. Specifically, the result holds for cost functions of the form $C(q) = \frac{1}{k}q^k$ even when $k \to \infty$. This finding is somewhat counterintuitive: as $k \to \infty$, the monomial cost converges to the unit-supply case, for which the result of Myerson (1981) implies that a posted pricing mechanism, which is a strict subclass of two-part tariffs, is revenue-optimal for any value distribution. Indeed, this intuition is valid in the regular case, as established by Proposition 4.1, which shows that the two-part tariff mechanism becomes revenue-optimal as $k \to \infty$. However, this correspondence breaks down in the non-regular case. The technical reason is that the worst-case distribution may also change as k increases, preventing the environment from converging to the unit-supply setting. This serves as an indication of the hardness introduced by non-linear environments.

This limitation of two-part tariff mechanisms arises because the seller cannot engage in effective price discrimination across buyer types associated with different ironing intervals when constrained to a single fixed fee. To address this issue, we introduce a broader class of mechanisms parameterized by a size s, and show that for any non-regular distribution with s ironing intervals, a menu with at most $s+1$ entries achieves a constant fraction of the optimal profit.

Definition 2. *[Menu of Two-Part Tariffs] A mechanism with a menu of size s consists of s entries. Each entry is specified by a triplet (T_i, β_i, L_i), where T_i is the fixed fee, β_i represents the markup factor, and L_i is the maximum purchasable quantity under that option. Upon choosing entry i, the buyer pays T_i to gain the right to purchase up to L_i units, each priced at $\beta_i + 1$ times the production cost. The buyer then selects the entry that maximizes his utility, or opts out of the mechanism if no option yields non-negative utility.*

It is clear that this family of mechanisms strictly generalizes the mechanisms considered earlier. The following theorem demonstrates that we can achieve the same approximation ratio with this stronger family of mechanisms.

Theorem 2. *Suppose the cost function $C(q)$ satisfies $m \le \eta_{C'}(q)$ for all $q \ge 0$. Then for any value distribution F with s ironing intervals, there exists a menu of two-part tariffs of size at most $s+1$ that applies a uniform markup $\beta = m$ across all menu entries while assigning distinct fixed fees corresponding to different capacity levels that guarantees at least a $\frac{m}{(m+1)^{1+1/m}}$ fraction of the optimal profit – i.e.,*

$$\text{MenuProfit} \ge \frac{m}{(m+1)^{1+1/m}} \cdot \pi^*.$$

Moreover, if C also satisfies the elasticity upper bound $\eta(q) \le M$ for all $q \ge 0$, then the guarantee improves to

$$\text{MenuProfit} \ge \frac{m}{(m+1)^{1+1/m}} \frac{M+1}{M} \cdot \pi^*.$$

Remark 5.1. *The mechanism from Theorem 2 employs a uniform markup across all menu entries. Hence, it can be interpreted as follows: the mechanism offers multiple levels of capacity constraints. Whenever the buyer reaches the capacity limit of the current level, he may either pay an additional fixed fee to unlock the next level or remain at the current level, while always paying the same fixed markup β.*

Similarly, we can apply Theorem 2 to the monomial cost functions. For this family, we obtain the same approximation ratio as in the regular case, albeit requiring a larger menu size.

Corollary 5.1. *Let the cost function be $C(q) = \frac{1}{k} q^k$ for some $k > 1$. For any value distribution F with s ironing intervals, there exists a menu of two-part tariffs of size at most $s+1$ that applies a uniform markup $\beta = k-1$ across all menu entries while assigning distinct fixed fees corresponding to different capacity levels, that achieves at least a fraction of $\frac{1}{k^{1/(k-1)}}$ of the optimal profit. That is to say,*

$$\text{MenuProfit} \ge \frac{1}{k^{\frac{1}{k-1}}} \cdot \pi^*.$$

We now sketch the main ideas behind the proof of Theorem 2. The argument follows the same overall structure as in the regular case. Again we consider the monomial cost functions where $C(q) = \frac{1}{k}q^k$ for some $k > 1$. Given any non-regular distribution F, suppose its ironed virtual value function $\tilde{\varphi}(v)$ contains s ironing intervals $\{[l_i, r_i]\}_{i=1}^s$. We first derive an upper bound on the optimal profit, and then show that a menu of two-part tariffs can approximate this benchmark up to a constant factor.

As in the regular case, the first step is to derive a suitable upper bound on the optimal profit for non-regular distributions. Using analogous arguments, and carefully accounting for the presence of ironed (flat) regions of the virtual value function, we obtain the following bound:

$$\pi^* \le \left(1 - \frac{1}{k}\right) \sum_{i=1}^{s} \int_{r_{i-1}}^{r_i} f(v) \Bigg(\underbrace{W(r_{i-1})}_{\text{fixed fee}} + \underbrace{C\big(q(\min\{v, l_i\})\big)}_{\text{markup}} - \underbrace{\sum_{j=1}^{i-1} \int_{r_{j-1}}^{r_j} q(\min\{t, l_j\})\, dt}_{\text{information rent for IC}} \Bigg) dv.$$

The value space is thus partitioned into $s+1$ intervals. Within each interval, the contribution to the optimal profit decomposes into components that can be interpreted as revenue from a fixed fee and from a markup, mirroring the structure of the regular case. This decomposition naturally suggests designing a menu of size approximately s, with menu options tailored to the intervals.

The main technical challenge lies in ensuring incentive compatibility across intervals, namely, guaranteeing that types in each interval select the menu option intended for them. The additional term relative to the regular case plays a crucial role in addressing this issue. It can be interpreted as an *information rent* that must be given up to prevent types from deviating to options designed for other intervals. By reducing the fixed fee by this amount and imposing a capacity constraint at $q(l_i)$ for each menu option, we are able to construct a menu of two-part tariffs that approximates the benchmark profit.

In the following, we address the necessary menu size for non-regular distributions, demonstrating that a menu of size $\Omega(s)$ in Theorem 2 is, in fact, unavoidable. The following theorem establishes that to obtain any constant approximation of the optimal profit, a menu of size $\Omega(s)$ is required, even with a simple monomial cost function $C(q) = \frac{1}{k}q^k$.

Theorem 3. *For any $K > 1$, there exists an instance with a non-regular buyer's value distribution F with K ironing intervals and a monomial cost function $C(q) = \frac{1}{K}q^K$, such that any menu of two-part tariffs of size s can achieve at most an $\frac{O(s)}{K}$ fraction of the optimal profit.*

The complete proofs of the results in this section are deferred to Appendix E.

6 Approximation Guarantees in the Multi-Product Case

In this section, we consider a more general framework with $n \ge 1$ product lines. We assume that the products have a production cost function C for each, and the buyer's values for different product lines $v^1, \ldots, v^n$ are drawn independently from $F_1, \ldots, F_n$, respectively, where each of the above distributions has a finite support and a corresponding probability mass function f_i (see also Remark 6.1). Define $F := F_1 \times \ldots F_n$ and let $T := \mathrm{Supp}(F)$. For $v = \left(v^1, \ldots, v^n\right) \in T$, define $f(v) := \Pi_{i=1}^n f_i\left(v^i\right)$.

A mechanism $M(v) = (x(v), p(v))$ consists of: (i) an *allocation rule* $x(v) = (x_1(v), \dots, x_n(v))$, which maps the buyer's type $v \in T$ to a vector $(x_1(v), \dots, x_n(v)) \in \mathbb{R}^n_{\geq 0}$, where $x_i(v)$ is the number of allocated units of product i; and (ii) a *pricing rule* $p(v)$, which maps $v \in T$ to a non-negative price. The buyer's utility under the mechanism M for a true type v and a reported type v' is $\sum_{i=1}^n v^i x_i(v') - p(v')$. We focus on truthful mechanisms. The *incentive-compatibility* constraints require that:

$$\sum_{i=1}^n v^i x_i(v) - p(v) \geq \sum_{i=1}^n v^i x_i(v') - p(v') \quad \forall v, v' \in T.$$

The *(ex post) individual rationality* constraints require that:

$$\sum_{i=1}^n v^i x_i(v) - p(v) \geq 0 \forall v \in T.$$

The following result provides approximation guarantees for the multi-product setting. The proof follows an outline similar to Cai et al. (2016). Namely, we use a partial Lagrangian corresponding to the seller's profit maximization problem and analyze the solutions of its dual via a flow framework. However, the presence of a cost function imposes additional technical challenges. Another challenge arises since the approach of Cai et al. (2016) involves approximation by a mechanism that sells the product lines separately and a mechanism that bundles the product lines. Although the former has a clear analogue in our setting, it is not immediately clear what the analogue should be for selling a bundle of product lines. In fact, the pricing of the bundle in our setting turns out to be a two-part tariff with zero markup over the production cost of each product.

The approximation guarantee with a continuum of products in every product line are more modest than those of Cai et al. (2016) where every product line contains only one product, but they still offer positive constants.

We define the following benchmarks. Let SREV be the optimal profit for selling the n product lines separately and BREV be the profit for selling the grand bundle consisting of $q\left(v^i\right)$ units of product line i (for every $1 \leq i \leq n$) for the total price of $\sum_{i=1}^n C\left(q\left(v^i\right)\right) + \frac{\mathrm{Med}\left(\sum_{i=1}^n W\left(v^i\right)\right)}{2}$, where Med is the median.

Theorem 4. *A simple menu of separate product lines and grand bundles guarantees:*

$$\pi^* \leq 4\mathrm{SREV} + 4\mathrm{BREV}. \tag{8}$$

Note that the mechanism for selling the grand bundle described in Theorem 4 is incentive-compatible, since after paying an access fee of $\frac{\mathrm{Med}\left(\sum_{i=1}^n W\left(v^i\right)\right)}{2}$, the quantities and prices of the products correspond to selling the product lines separately in an incentive-compatible mechanism. In other words, upon paying the access fee, the buyer gains the opportunity to purchase all the products for their production costs with no additional markup.

Remark 6.1. *In this paper, we focus on a discrete setting. However, Theorem 4 also holds for continuous distributions $F_1, \dots, F_n$ (see Cai et al. (2016) for a connection between the discrete and the continuous frameworks).*

Remark 6.2. *If the cost function C and each of the distributions $F_1, \dots, F_n$ satisfy the conditions from either of Theorem 1 or Proposition 4.1, then Theorem 4 immediately implies that we can get a constant-factor approximation of the optimal profit using two-part tariffs for any number of goods $n \geq 1$. Similarly, if the conditions of either of Theorem 2 or Corollary 5.1 are satisfied, then a menu of two-part tariffs achieves constant-factor approximation guarantees for any number of goods.*

7 Extensions to Nonlinear Buyer Values

The model of Section 2 assumes that the buyer's gross value for quantity q is linear in q, namely $v\,q$. In this section, we extend Theorem 1 to buyers with multiplicatively separable values, $v\,U(q)$, for which the problem reduces to the linear case by a change of variables and the guarantee carries over essentially unchanged. We then turn to general supermodular values $u(v,q)$: we describe what is known, formulate a conjecture, and explain why the general case is a challenging open problem. Proofs are deferred to Appendix G.

Suppose first that the buyer with type v derives gross value $v\,U(q)$ from quantity q, so that his utility from receiving q at payment p is $v\,U(q) - p$.

Assumption 1. *The function $U : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is strictly increasing, unbounded from above, continuous, concave, in $C^2((0,\infty))$, with $U(0) = 0$.*

Assumption 2. *The function $C : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is strictly increasing, convex, in $C \in C^1([0,\infty)) \cap C^2((0,\infty))$, with $C(0) = 0$.*

This environment reduces to the linear one by measuring quantity in units of the buyer's utility. Setting $\tilde{q} := U(q)$ and defining the *induced cost function* $\tilde{C}(\tilde{q}) := C\big(U^{-1}(\tilde{q})\big)$, the buyer's gross value becomes the linear expression $v\,\tilde{q}$, and a two-part tariff in the original units is a two-part tariff in the new units. Since U is strictly increasing and concave, $\tilde{C}$ is strictly increasing and convex; under the elasticity condition (Eq. (9) below), it is *strictly* convex, so the mechanism design problem is exactly the one studied in Section 2, with $\tilde{C}$ in place of C. The only remaining ingredient is to express the bounded-elasticity condition (cf. Definition 1) for the induced cost in terms of the primitives (C,U); a direct computation (Lemma G.1 in the appendix) shows that the marginal-cost elasticity of $\tilde{C}$ equals the ratio appearing in the following statement.

Proposition 7.1 (Generalization of Theorem 1 to $v\,U(q)$)**.** *Let U satisfy Assumption 1 and let C satisfy Assumption 2. Suppose there exist constants $M \geq m > 0$ such that:*[4]

$$m \leq \frac{\eta_{C'}(q) - \eta_{U'}(q)}{\eta_U(q)} \leq M \qquad \forall\, q \geq 0, \tag{9}$$

where $\eta_U(q) := qU'(q)/U(q)$ and $\eta_{U'}(q) := qU''(q)/U'(q)$. Then for every regular value distribution F,

[4] Here we allow $M = \infty$; if this is the case, we interpret $(M+1)/M$ as 1.

the universal markup $\beta := m$ admits a fixed fee $T = T(F)$ such that

$$\pi(T,\beta) \geq \frac{m}{(m+1)^{1+1/m}} \cdot \frac{M+1}{M} \cdot \pi^*.$$

Two special cases illustrate the condition. For linear values, $U(q) = q$, the ratio in Eq. (9) collapses to $\eta_{C'}(q)$ and the proposition recovers Theorem 1. At the other extreme, when the buyer's value is strictly concave in q and the seller's cost is *linear*—the case alluded to informally in Section 2—the ratio becomes $-\eta_{U'}/\eta_U$, so the proposition applies whenever the curvature of the value function itself is suitably bounded: the concavity of demand plays the role that the convexity of cost plays in the baseline model. When both U and C are monomial, the condition holds with equality throughout and the guarantee takes a closed form.

Corollary 7.1 (Monomial values and monomial costs)**.** *Suppose $U(q) = q^\alpha$ for some $\alpha \in (0,1]$ and $C(q) = \frac{1}{k}q^k$ for some $k > 1$, with $k > \alpha$. Then the universal markup $\beta := (k-\alpha)/\alpha$ achieves, for every regular value distribution F, a fraction $(k/\alpha)^{-\alpha/(k-\alpha)}$ of the optimal profit.*

Beyond separability. For a general supermodular gross value $u(v,q)$, concave in q, the relevant new primitive is the elasticity of the marginal utility with respect to the type, $\sigma(v,q) := v\, u_{qv}(v,q)/u_q(v,q)$, which is identically 1 for separable values; a *constant* elasticity σ is equivalent to $u(v,q) = v^\sigma U(q)$ and reduces to the separable case by reparametrizing the type, so Proposition 7.1 effectively covers that situation. The open case is a σ that varies between bounds $\underline{\sigma} < \overline{\sigma}$. We conjecture that two-part tariffs with a universal markup continue to guarantee a distribution-free constant fraction of the optimal profit, with the elasticity floor m of Theorem 1 deflated to $m\underline{\sigma}/\overline{\sigma}$—so that, as in the rest of the paper, the guarantee approaches 1 as costs become highly convex, at a rate dampened by the heterogeneity of the type response. Much of the machinery supports this: the markup still acts as a type contraction confined to an explicit band by the σ-bounds, the envelope decomposition of the tariff profit carries over with a generalized virtual utility in place of $\varphi(v)\,q$, and the optimal profit admits an upper bound of the same form.

The difficulty is that the two sides of the comparison are governed by *different* ends of the elasticity band: the benchmark is largest when the buyer's information rent is small (σ near $\underline{\sigma}$), while the tariff guarantee must withstand rents that are large (σ near $\overline{\sigma}$). Near the participation threshold the two resulting virtual weights come apart, and because $\sigma(v,q)$ may vary with the quantity, an adversary can make the rent small at the quantities served by the optimal mechanism yet large at the quantities induced by the tariff—no single value of σ governs both sides. The entry fee collects additional profit from precisely the threshold types, but combining it with the markup profit uniformly across distributions appears to require either a structural link between the values of σ at different quantities or a genuinely different benchmark. We leave this as an open problem; the fact that the conjectured guarantee is exact at $\underline{\sigma} = \overline{\sigma}$ suggests it should hold at least when $\overline{\sigma}/\underline{\sigma}$ is close to one.

8 Related Work

Our work relates to several streams of literature. The foundational works of Mussa and Rosen (1978) and Maskin and Riley (1984) study monopoly pricing with quality discrimination, establishing the framework for analyzing nonlinear pricing under private information.

Simple mechanisms. There is an extensive literature on simple, approximately optimal mechanisms for settings with indivisible goods, including single-item auctions (e.g., Hartline and Roughgarden, 2009; Jin et al., 2019), multi-item monopolist pricing (e.g., Hart and Nisan, 2012; Li and Yao, 2013; Rubinstein and Weinberg, 2015; Babaioff et al., 2020), and more general multi-item auction environments (e.g., Chawla et al., 2010; Cai and Huang, 2013; Yao, 2014; Cai et al., 2016; Chawla and Miller, 2016; Cai and Zhao, 2017; Babaioff et al., 2022; Cai et al., 2022, 2023). Related results for nonlinear agents consider different environments (Feng et al., 2023): while their model considers nonlinear agents with *unit-demand* allocations and non-linear utility of the form $U(vq - p)$, we study agents with utility $U(v, q) - p$ under general allocations. Consequently, the two models are not directly comparable. Caragiannis et al. (2020) study a related multi-buyer setting with concave valuation functions; under shortsighted buyers and regularity of the distributions of marginal values, they show that linear pricing achieves a logarithmic approximation to the optimal revenue. In contrast, we establish constant-factor guarantees for two-part tariffs in a single-buyer setting with convex production costs. Our guarantees hold for regular value distributions; moreover, when distributions are non-regular, we show that menus of two-part tariffs recover constant-factor approximations. To our knowledge, this provides the first approximation guarantees for nonlinear pricing with general value distributions and convex production costs.

Two-part tariffs have been extensively studied in industrial organization and applied microeconomics. Wilson's 1993 monograph on nonlinear pricing (Wilson, 1993) provides a comprehensive treatment, with applications spanning utilities, telecommunications, and service industries. Oi (1971) establishes early theoretical foundations, and Rochet and Stole (2002) analyze nonlinear pricing in competitive settings. In contrast to this classical literature, our work provides the first approximation guarantees for broad classes of value distributions and cost functions, bridging the gap between theoretical optimality and practical simplicity. Related approximation results for revenue maximization also appear in multi-item monopolist pricing and auction settings with indivisible goods, where two-part tariffs and their sequential variants are shown to approximate optimal revenue (Cai and Zhao, 2017). Beyond mechanism design, Iyer et al. (2025) study a principal–agent model in which the agent can commit to an action and the principal may reject the realized outcome; they show that a menu of two-part tariffs is optimal for the principal.

Optimal multi-item mechanisms can require unbounded menu sizes even with independent values, as shown in (Daskalakis et al., 2017). Our work contributes to understanding when simple menus suffice. The duality approach we build upon was developed by Cai et al. (2016), providing a unified framework for Bayesian mechanism design that accommodates multiple agents and arbitrary objectives.

Nonlinear pricing with finite or limited information. A complementary line of work examines nonlinear pricing under explicit *informational* restrictions on the menu rather than parametric restrictions on its functional form. Bergemann et al. (2012) introduce a quantization framework for multi-dimensional nonlinear pricing in which the seller is constrained to offer a finite menu of n bundles, and connect the resulting welfare loss to classical results from the theory of vector quantization. Bergemann et al. (2021b) substantially extend this framework: for a buyer with preferences over a d-dimensional variety of goods and a seller restricted to a menu of n bundles, they characterize necessary conditions for the optimal menu under both welfare- and revenue-maximizing objectives, and apply the Lloyd–Max algorithm from information theory to compute the optimal quantization of the menu. Their work and ours pose the same conceptual question of *how should a seller restrict the complexity of the price schedule when the optimal mechanism is intractable*, but along orthogonal axes of simplicity. They restrict the cardinality of the menu while allowing each bundle to be priced arbitrarily; we restrict the menu to a two-part tariff (and, for non-regular distributions, to small menus of two-part tariffs) while allowing the buyer to choose any quantity from a continuum.

Cost-based and prior-free nonlinear pricing. A closely related strand of recent work studies cost-based nonlinear pricing from a prior-free perspective. Bergemann et al. (2023) analyze a monopolist offering a quantity- or quality-differentiated product when the demand distribution is unknown and show that the worst-case ratio of expected profit to expected social surplus is maximized by setting the markup over marginal cost equal to the elasticity of the marginal cost function. Bergemann et al. (2025) develop and substantially extend this analysis in a procurement environment: when the buyer has no prior over the supplier's cost, mechanisms that share a constant fraction of the buyer's utility with the seller guarantee a constant fraction of the efficient social surplus across all cost realizations, and the optimal sharing rule can be characterized in closed form. Bergemann et al. (2026), with a different focus but a closely related toolkit, consider a monopolist selling multiple products who can engage in both second- and third-degree price discrimination via market segmentation, and characterize consumer-optimal segmentations from a consumer-surplus perspective.

The methodological connection to our results is striking. The prior-free benchmark of Bergemann et al. (2023) prescribes a markup equal to the marginal cost elasticity, which for a monomial cost function $C(q) = q^k/k$ coincides with the universal markup $\beta = k-1$ that we identify in Proposition 4.1. Indeed, both analyses isolate the elasticity of the marginal cost function as the primitive that pins down the optimal markup. The two approaches differ along three substantive dimensions. *First*, the benchmark: Bergemann et al. (2023) and Bergemann et al. (2025) measure the seller's (or buyer's) profit against the efficient social surplus and seek a guarantee that holds for every demand (or cost) distribution, whereas we measure expected profit against the optimal mechanism designed for a known distribution, and seek a guarantee that holds for every distribution within a regularity class. *Second*, the role of the fixed fee: the prior-free benchmark in Bergemann et al. (2023) is achieved with a pure cost-plus markup and no fixed fee, while in our Bayesian environment the fixed fee is essential to extract the information rent of the marginal type, and its calibration depends explicitly on the distribution. *Third*, the scope: the prior-free analysis applies to general convex costs but delivers a surplus-share guarantee, while our Bayesian analysis applies to costs with bounded marginal-cost elasticity but delivers a profit guarantee against the Bayesian optimum and extends to non-regular distributions, multiple products, and generalized nonlinear values (see Section 7 above). Taken together, the two literatures suggest that the cost-elasticity-based markup is a robust object that arises as the optimal simple instrument under markedly different informational assumptions about the demand side.

Connection to bilateral trade and intermediary pricing. We consider the problem of a seller who faces uncertain demand for a variety of products. In this setting, there is only one-sided private information. We believe that the tools developed here could also apply to bilateral trade with two-sided asymmetric information Myerson and Satterthwaite (1983); Rochet (1985); Green and Honkapohja (1983). Hagerty and Rogerson (1987) provided the first version of a robust trading mechanism in the bilateral trade environment. Recent work on constant approximation involves (Blumrosen and Mizrahi, 2016; Brustle et al., 2017; Deng et al., 2022; Fei, 2022; Hartline and Wang, 2025) for the gains from trade and (Blumrosen and Dobzinski, 2016; Kang et al., 2022; Cai and Wu, 2023; Liu et al., 2023) for the welfare. The work of Zhang (2022) makes a distinction between bilateral trade with and without an intermediary. In the framework of bilateral trade with an intermediary and unit demand, Loertscher and Niedermayer (2007) and Loertscher and Niedermayer (2012) investigate the shape of the optimal, possibly nonlinear, fee of the transaction price. We leave the applicability of our methodology for studying simple bilateral trade pricing mechanisms as an interesting direction for future research.

9 Conclusion

We have studied the approximation guarantees of simple two-part tariffs for nonlinear pricing with convex production costs. Our main contributions establish both the power and the limitations of simplicity in mechanism design. While most of the paper focuses on convex marginal costs and a buyer's utility that is linear in the type, Proposition 7.1 extends the guarantee to multiplicatively separable value functions, and we conjecture in Section 7 that a similar guarantee extends to general supermodular values.

For the special case of Pareto distributions with monomial costs, we derived explicit optimal parameters and showed that two-part tariffs are indeed optimal.

For general regular distributions with cost functions of bounded elasticity, we proved that two-part tariffs with appropriately chosen markups achieve approximation ratios that depend only on the elasticity bounds. For monomial cost functions, these ratios approach one as cost convexity increases; for general convex cost functions with elasticity uniformly bounded away from zero, a similar conclusion holds when the lower bound on the elasticity grows large. Our analysis shows that when production cost is highly convex—characteristic of capacity-constrained or congested systems—simple two-part tariffs capture nearly all of the profit attainable by complex optimal mechanisms.

For non-regular distributions, we established that simple two-part tariffs can fail catastrophically, extracting only an arbitrarily small fraction of the optimal profit in carefully constructed examples. However, we proved that menus of two-part tariffs with a number of entries equal to the number of ironing intervals for the virtual value function plus one achieve constant-factor approximations, where the constant depends only on the elasticity bounds. This result shows that the required complexity grows only with the intrinsic non-regularity of the distribution—measured by the amount of ironing needed—rather than with the full complexity of the optimal mechanism.

Our multi-product extension bounded the optimal profit by a constant multiple of the profit from selling the products separately and from bundling them together, showing that simple mechanism formats capture essentially all attainable profit even in multi-dimensional settings in which optimal mechanisms may be exponentially complex. Under additional assumptions on each item's value distribution, as discussed above for the single-item framework, the mechanism can be further simplified to either (i) selling each product separately using two-part tariffs or a menu of them, or (ii) selling the bundle using a two-part tariff; the better of the two achieves a constant fraction of the optimal profit.

These results provide a rigorous justification for the widespread use of simple pricing schemes in practice. Although optimal mechanisms may in principle yield higher profit, we show that the gap between simple and optimal mechanisms shrinks under natural economic and informational conditions. Complexity matters less when production costs are highly convex, when value distributions are regular, or when non-regularities are limited.

The gap between simple and optimal mechanisms is both quantifiable and often small in practical regimes. This helps explain why firms across diverse industries employ two-part tariffs despite the availability of sophisticated mechanism design theory: the additional profit from fully optimal mechanisms may not justify the costs of implementation, communication, or potential behavioral frictions faced by consumers under complex pricing schedules.

For practitioners, our results suggest that two-part tariffs are near-optimal when cost functions exhibit high elasticity, corresponding to large values of the power parameter in monomial cost structures. The markup should be chosen based on the elasticity bounds, while the fixed fee should be set to extract surplus from marginal buyers without inducing excessive exclusion. For non-regular demand distributions, a small menu of options suffices—rather than a fully continuous pricing schedule.

References

Moshe Babaioff, Yannai A. Gonczarowski, and Noam Nisan. 2022. The menu-size complexity of revenue approximation. *Games and Economic Behavior* 134 (July 2022), 281–307. doi:10.1016/j.geb.2021.03.001

Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg. 2020. A Simple and Approximately Optimal Mechanism for an Additive Buyer. *J. ACM* 67, 4 (June 2020), 24:1–24:40. doi:10.1145/3398745

Patrick Bajari and Steven Tadelis. 2001. Incentives versus transaction costs: A theory of procurement contracts. *Rand journal of Economics* (2001), 387–407.

Dirk Bergemann, Tibor Heumann, and Stephen Morris. 2021a. Selling Impressions: Efficiency vs. Competition. (2021).

Dirk Bergemann, Tibor Heumann, and Stephen Morris. 2023. Cost Based Nonlinear Pricing. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC '23)*. Extended abstract; full version: Cowles Foundation Discussion Paper No. 2353.

Dirk Bergemann, Tibor Heumann, and Stephen Morris. 2025. *Procurement without Priors: A Simple Mechanism and Its Notable Performance*. Technical Report Discussion Paper No. 2479. Cowles Foundation for Research in Economics, Yale University.

Dirk Bergemann, Tibor Heumann, and Michael Wang. 2026. *Screening and Segmenting: A Consumer Surplus Perspective*. Technical Report Discussion Paper No. 2498. Cowles Foundation for Research in Economics, Yale University.

Dirk Bergemann, Ji Shen, Yun Xu, and Edmund M. Yeh. 2012. Mechanism Design with Limited Information: The Case of Nonlinear Pricing. In *Game Theory for Networks (GameNets 2011)*. Springer, 1–10.

Dirk Bergemann, Edmund M. Yeh, and Jinkun Zhang. 2021b. Nonlinear Pricing with Finite Information. *Games and Economic Behavior* 130 (2021), 62–84.

Liad Blumrosen and Shahar Dobzinski. 2016. (Almost) Efficient Mechanisms for Bilateral Trading. *CoRR* abs/1604.04876 (2016). `http://arxiv.org/abs/1604.04876`

Liad Blumrosen and Yehonatan Mizrahi. 2016. Approximating Gains-from-Trade in Bilateral Trading. In *Web and Internet Economics - 12th International Conference, WINE 2016, Montreal, Canada, December 11-14, 2016, Proceedings*. 400–413. doi:`10.1007/978-3-662-54110-4_28`

Johannes Brustle, Yang Cai, Fa Wu, and Mingfei Zhao. 2017. Approximating Gains from Trade in Two-sided Markets via Simple Mechanisms. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC '17)*. Association for Computing Machinery, New York, NY, USA, 589–590. doi:`10.1145/3033274.3085148`

Yang Cai, Ziyun Chen, and Jinzhao Wu. 2023. Simultaneous Auctions are Approximately Revenue-Optimal for Subadditive Bidders. In *Proceedings of IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS 2023)*. IEEE Computer Society, 134–147. doi:`10.1109/FOCS57990.2023.00017`

Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. 2016. A duality based unified approach to bayesian mechanism design. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 926–939.

Yang Cai and Zhiyi Huang. 2013. Simple and Nearly Optimal Multi-Item Auctions. In *Proceedings of the 2013 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) (Proceedings)*. Society for Industrial and Applied Mathematics, 564–577. doi:`10.1137/1.9781611973105.41`

Yang Cai, Argyris Oikonomou, and Mingfei Zhao. 2022. Computing simple mechanisms: Lift-and-round over marginal reduced forms. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2022)*. Association for Computing Machinery, New York, NY, USA, 704–717. doi:10.1145/3519935.3520029

Yang Cai and Jinzhao Wu. 2023. On the optimal fixed-price mechanism in bilateral trade. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 737–750.

Yang Cai and Mingfei Zhao. 2017. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2017)*. Association for Computing Machinery, New York, NY, USA, 170–183. doi:10.1145/3055399.3055465

Ioannis Caragiannis, Zhile Jiang, and Apostolis Kerentzis. 2020. Bounds on the revenue gap of linear posted pricing for selling a divisible item. *arXiv e-prints* (2020), arXiv–2007.

Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing (STOC '10)*. Association for Computing Machinery, New York, NY, USA, 311–320. doi:10.1145/1806689.1806733

Shuchi Chawla and J. Benjamin Miller. 2016. Mechanism Design for Subadditive Agents via an Ex-Ante Relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC '16, Maastricht, The Netherlands, July 24-28, 2016*. 579–596. doi:10.1145/2940716.2940756 Journal Abbreviation: CoRR.

Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. 2017. Strong Duality for a Multiple-Good Monopolist. *Econometrica* 85, 3 (2017), 735–767. Publisher: Wiley Online Library.

Yuan Deng, Jieming Mao, Balasubramanian Sivan, and Kangning Wang. 2022. Approximately efficient bilateral trade. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC 2022)*. Association for Computing Machinery, New York, NY, USA, 718–721. doi:10.1145/3519935.3520054

Yumou Fei. 2022. Improved approximation to first-best gains-from-trade. In *International Conference on Web and Internet Economics*. Springer, 204–218.

Yiding Feng, Jason D. Hartline, and Yingkai Li. 2023. Simple Mechanisms for Non-linear Agents. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Society for Industrial and Applied Mathematics, 3802–3816. doi:10.1137/1.9781611977554.ch148

Jerry Green and Seppo Honkapohja. 1983. Bilateral contracts. *Journal of Mathematical Economics* 11, 2 (1983), 171–187.

Kathleen M Hagerty and William P Rogerson. 1987. Robust trading mechanisms. *Journal of Economic Theory* 42, 1 (1987), 94–107.

Sergiu Hart and Noam Nisan. 2012. Approximate Revenue Maximization with Multiple Items. In *the 13th ACM Conference on Electronic Commerce (EC)*.

Jason D. Hartline and Tim Roughgarden. 2009. Simple versus optimal mechanisms. In *ACM Conference on Electronic Commerce*. 225–234.

Jason D. Hartline and Kangning Wang. 2025. A Geometric Analysis of Gains from Trade. *arXiv preprint arXiv:2508.06469* (2025).

Krishnamurthy Iyer, Alec Sun, Haifeng Xu, and You Zu. 2025. How to Sell a Service with Uncertain Outcomes. *arXiv preprint arXiv:2502.13334* (2025).

Yaonan Jin, Pinyan Lu, Qi Qi, Zhihao Gavin Tang, and Tao Xiao. 2019. Tight approximation ratio of anonymous pricing. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC 2019)*. Association for Computing Machinery, New York, NY, USA, 674–685. doi:10.1145/3313276.3316331

Zi Yang Kang, Francisco Pernice, and Jan Vondrák. 2022. Fixed-price approximations in bilateral trade. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2964–2985.

Jean-Jacques Laffont and Jean Tirole. 1993. *A theory of incentives in procurement and regulation*. MIT press.

Xinye Li and Andrew Chi-Chih Yao. 2013. On Revenue Maximization for Selling Multiple Independently Distributed Items. *Proceedings of the National Academy of Sciences* 110, 28 (2013), 11232–11237.

Zhengyang Liu, Zeyu Ren, and Zihe Wang. 2023. Improved approximation ratios of fixed-price mechanisms in bilateral trades. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. 751–760.

Simon Loertscher and Andras Niedermayer. 2007. *When is seller price setting with linear fees optimal for intermediaries?* Technical Report. Discussion Papers.

Simon Loertscher and Andras F Niedermayer. 2012. Fee-setting mechanisms: On optimal pricing by intermediaries and indirect taxation. *Available at SSRN 2172386* (2012).

Eric Maskin and John Riley. 1984. Monopoly with incomplete information. *The RAND Journal of Economics* 15, 2 (1984), 171–196.

R Preston McAfee and John McMillan. 1986. Bidding for contracts: a principal-agent analysis. *The RAND Journal of Economics* (1986), 326–338.

Michael Mussa and Sherwin Rosen. 1978. Monopoly and product quality. *Journal of Economic Theory* 18, 2 (Aug. 1978), 301–317. doi:10.1016/0022-0531(78)90085-6

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

Roger B Myerson and Mark A Satterthwaite. 1983. Efficient mechanisms for bilateral trading. *Journal of economic theory* 29, 2 (1983), 265–281. Publisher: Elsevier.

Walter Y. Oi. 1971. A Disneyland Dilemma: Two-Part Tariffs for a Mickey Mouse Monopoly. *The Quarterly Journal of Economics* 85, 1 (Feb. 1971), 77–96. doi:10.2307/1881841 Publisher: Oxford Academic.

Jean-Charles Rochet. 1985. Bilateral monopoly with imperfect information. *Journal of Economic Theory* 36, 2 (1985), 214–236.

Jean-Charles Rochet and Lars A. Stole. 2002. Nonlinear Pricing with Random Participation. *The Review of Economic Studies* 69, 1 (2002), 277–311. https://www.jstor.org/stable/2695961 Publisher: [Oxford University Press, Review of Economic Studies, Ltd.].

Aviad Rubinstein and S. Matthew Weinberg. 2015. Simple Mechanisms for a Subadditive Buyer and Applications to Revenue Monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*. 377–394. doi:10.1145/2764468.2764510

Knut Sydsæter and Peter J Hammond. 2008. *Essential mathematics for economic analysis*. Pearson Education.

Jean Tirole. 1988. *The Theory of Industrial Organization*. MIT Press, Cambridge.

Robert Wilson. 1993. *Nonlinear Pricing*. New York: Oxford University Press.

Andrew Chi-Chih Yao. 2014. An n-to-1 Bidder Reduction for Multi-item Auctions and its Applications. In *Proceedings of the 2015 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. Society for Industrial and Applied Mathematics, 92–109. doi:10.1137/1.9781611973730.8

Wanchang Zhang. 2022. Random double auction: A robust bilateral trading mechanism. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 182–183.

A Additional Discussion About The Model

Further properties of simple mechanisms. Consider a two-part tariff specified by a fixed fee T and a markup β, and let $x_{T,\beta}(v)$ denote its allocation rule. When the fixed fee or the markup is clear from context, we may omit the corresponding subscript.

First, suppose there is no fixed fee. The buyer then solves:

$$\max_{q\geq 0}\left(vq-(1+\beta)C(q)\right)=(1+\beta)\cdot\max_{q\geq 0}\left(\frac{v}{1+\beta}q-C(q)\right).$$

By the definition of $q(\cdot)$, the optimal purchase quantity is, therefore:

$$q\left(\frac{v}{1+\beta}\right).$$

When a fixed fee T is present, the buyer participates if and only if his utility at least T. Define

$$z(v)=z_\beta(v):=v\,q\left(\frac{v}{1+\beta}\right)-(1+\beta)C\left(q\left(\frac{v}{1+\beta}\right)\right)$$

to be the buyer's utility in the absence of the fixed fee. Notice that $z(v)$ is non-decreasing, which can be checked by simple derivation. The allocation rule induced by the two-part tariff (T,β) can then be written as:

$$x_{T,\beta}(v)=\begin{cases}0 & \text{if } z(v)<T,\\ q\left(\frac{v}{1+\beta}\right) & \text{if } z(v)\geq T.\end{cases}$$

Finally, by the first-order condition (5), for every simple mechanism (T,β) there exists a cutoff $\tau_2=\tau_2(T,\beta)$ such that $x_{T,\beta}(v)=0$ for all $v<\tau_2$, while for all $v\geq\tau_2$ the allocation satisfies

$$C'\left(x_{T,\beta}(v)\right)=\frac{v}{1+\beta}. \tag{10}$$

Therefore, the expected profit of a simple mechanism (T,β) is:

$$\pi(T,\beta)=\Pr[z(v)\geq T]\cdot T+\beta\cdot\mathbb{E}\left[C(x(v))\cdot\mathbb{1}_{z(v)\geq T}\right]. \tag{11}$$

Additional discussion about elasticity. Consider the *elasticity of the cost* $C(q)$ with respect to the quantity q:

$$\eta_C=\frac{\frac{dC(q)}{dq}}{\frac{C(q)}{q}}=\frac{\frac{dC(q)}{C(q)}}{\frac{dq}{q}}=\frac{d\log(C(q))}{d\log(q)}$$

Formally, it is the ratio of the percentage change of cost $C(q)$ induced by a percentage change in the quantity q. The notion of elasticity is used in economics as it is unit-free. That is, the elasticity of the cost function is independent of how input and output are measured. It is also the (double) logarithmic derivative of $C(q)$ with respect to q. The cost elasticity measures returns to scale, as it can be rewritten as the ratio of marginal cost to average cost:

$$\eta_C = \frac{\frac{dC(q)}{dq}}{\frac{C(q)}{q}} = \frac{\text{Marginal Cost}}{\text{Average Cost}}$$

Similarly, we can define the *elasticity of the marginal cost function $C'(q)$*:

$$\eta_{C'} = \frac{\frac{dC'(q)}{dq}}{\frac{C'(q)}{q}} = \frac{C''(q)\,q}{C'(q)}.$$

With a monomial cost function $C(q) = \frac{1}{k}q^k$ $(k > 1)$, the cost elasticity and the marginal cost elasticity are tightly connected:

$$\eta_{C'} = \eta_C - 1.$$

The notion of marginal cost elasticity is closely related to the notion of *supply elasticity*. A competitive firm (i.e., a firm supply in a competitive equilibrium) expands its supply until the marginal cost is equal to the competitive equilibrium (unit) price p of the product:

$$C'(q) = p. \tag{12}$$

Thus, its supply function is the inverse of the marginal cost function:

$$q = (C')^{-1}(p)$$

Differentiating the above first-order condition (12) with respect to p yields:

$$1 = C''(q)\frac{dq}{dp} \Rightarrow \frac{dq}{dp} = \frac{1}{C''(q)},$$

and thus the elasticity of supply, denoted by ε_S, can be written as:

$$\varepsilon_S = \frac{dq}{dp}\frac{p}{q} = \frac{1}{C''(q)}\frac{C'(q)}{q} = \frac{1}{\eta_{C'}}.$$

Therefore, *the elasticity of supply with respect to price is the reciprocal of the marginal cost elasticity.*

Consequently, the assumption $\eta_{C'}(q) \in [m, M]$ for all $q > 0$ is equivalent to the assumption that

$$\varepsilon_S(q) \in \left[\frac{1}{M}, \frac{1}{m}\right] \qquad \forall q > 0,$$

i.e., that the firm's supply elasticity is uniformly bounded away from 0 and ∞. Intuitively, a large marginal-cost elasticity $\eta_{C'}$ corresponds to a sharply rising marginal-cost curve and hence to a relatively inelastic supply, while a small $\eta_{C'}$ corresponds to a gently rising marginal cost and a relatively elastic supply.

For the monomial cost function $C(q) = \frac{1}{k}q^k$, the elasticity of the marginal cost is the constant $\eta_{C'} \equiv$

$k-1$, so the corresponding supply elasticity is the constant $\varepsilon_S \equiv 1/(k-1)$. The quadratic-cost case $k=2$, often used as a benchmark in the applied literature, corresponds to a unit-elastic supply, $\varepsilon_S \equiv 1$. The limiting case $k \to 1^+$ corresponds to perfectly elastic supply ($\varepsilon_S \to \infty$), in which the firm is essentially price-taking at a constant marginal cost; the limiting case $k \to \infty$ corresponds to perfectly inelastic supply ($\varepsilon_S \to 0$), in which the firm can produce only at a single quantity.

This reciprocal relationship has two consequences worth noting. *First*, the approximation guarantee of Theorem 1 can be restated entirely in terms of supply elasticities: writing $\underline{\varepsilon} := 1/M$ and $\overline{\varepsilon} := 1/m$ for the bounds on $\varepsilon_S(q)$, the universal markup is $\beta = 1/\overline{\varepsilon}$ and the approximation ratio is:

$$\frac{1+\underline{\varepsilon}}{\overline{\varepsilon}} \cdot (1+1/\overline{\varepsilon})^{-(1+\overline{\varepsilon})}.$$

For a monomial cost function, this reduces to $(1+1/\varepsilon_S)^{-\varepsilon_S}$, a particularly transparent function of the supply elasticity alone. *Second*, since supply elasticities (rather than marginal-cost elasticities) are the object most commonly estimated in the empirical industrial-organization literature, this restatement allows the quantitative bounds in our main theorem to be anchored directly to empirically-measured primitives.

B Proofs from Section 3

We begin this section with the following observation.

Observation B.1. *For a cost function $C(q) = \frac{1}{k}q^k$ ($k>1$) and F the Pareto distribution with parameter $\lambda > \frac{k}{k-1}$ supported on $[1,\infty)$, we have:*

$$\pi^* = \left(\frac{k-1}{k}\right)\left(\frac{\lambda-1}{\lambda}\right)^{k/(k-1)} \frac{\lambda}{\lambda - \frac{k}{k-1}}.$$

Proof of Observation B.1. We have:

$$\begin{aligned}
\pi^* &= \mathbb{E}_{v\sim F}\big[W\big(\varphi(v)^+\big)\big] = \frac{k-1}{k}\mathbb{E}_{v\sim F}\left[\left(\frac{\lambda-1}{\lambda}v\right)^{\frac{k}{k-1}}\right] \\
&= \left(\frac{k-1}{k}\right)\left(\frac{\lambda-1}{\lambda}\right)^{k/(k-1)} \lambda \int_{v=1}^{\infty} v^{\frac{k}{k-1}-\lambda-1}\,dv \\
&= \left(\frac{k-1}{k}\right)\left(\frac{\lambda-1}{\lambda}\right)^{k/(k-1)} \frac{\lambda}{\frac{k}{k-1}-\lambda} v^{\frac{k}{k-1}-\lambda}\Big|_{v=1}^{v=\infty} = \left(\frac{k-1}{k}\right)\left(\frac{\lambda-1}{\lambda}\right)^{k/(k-1)} \frac{\lambda}{\lambda-\frac{k}{k-1}},
\end{aligned}$$

where the integration step uses the assumption $\lambda > \frac{k}{k-1}$. □

Proof of Proposition 3.2. Define $\zeta^{k-1} := \frac{1}{1+\beta}$ and $z(v) := x(v)v - (1+\beta)C(x(v))$. Assume first $T=0$. Let us compute the expected profit in a two-part tariff with a markup β:

$$\begin{aligned}\pi &= \beta\mathbb{E}[C(x(v))] = \beta\mathbb{E}\left[\frac{1}{k}\zeta^k v^{\frac{k}{k-1}}\right]\\ &= \frac{\lambda\beta\zeta^k}{k}\int_{v=1}^{\infty} v^{\frac{k}{k-1}-\lambda-1}dv = \frac{\beta\zeta}{k(1+\beta)}\cdot\frac{\lambda}{\frac{k}{k-1}-\lambda}v^{\frac{k}{k-1}-\lambda}\Big|_{v=1}^{v=\infty}\\ &= \zeta\frac{1-\zeta^{k-1}}{k}\frac{\lambda}{\lambda-\frac{k}{k-1}}.\end{aligned}$$

Derivation shows that it is optimal for the seller to take $\zeta^{k-1} = 1/k$ (i.e., $\beta = k-1$). This choice yields, together with Observation B.1:

$$\frac{\pi(T=0,\beta=k-1)}{\pi^*} = \frac{\zeta}{((\lambda-1)/\lambda)^{k/(k-1)}}\frac{1-\zeta^{k-1}}{k-1} = \frac{1}{k^{1+1/(k-1)}}\cdot\frac{1}{((\lambda-1)/\lambda)^{k/(k-1)}} \to_{k\to\infty} 0,$$

which proves the first part.

For a general T, note that we have (see Equation 6) $x(v) = \zeta v^{\frac{1}{k-1}}$, and therefore:

$$z(v) = x(v)v-(1+\beta)C(x(v)) = x(v)v-\zeta^{-(k-1)}C(x(v)) = \frac{k-1}{k}\zeta v^{\frac{k}{k-1}}.$$

Hence, we have:

$$v = \left(\frac{k/\zeta}{k-1}z(v)\right)^{\frac{k-1}{k}}.$$

Note that it is unreasonable to choose $T < z(v=1)$, as it is possible to increase the pointwise profit by increasing such T. For $T \geq z(v=1)$, we get:

$$\Pr[z(v)\geq T]\cdot T = (1-F(v(z=T)))T = \left(\frac{k}{\zeta(k-1)}\right)^{-\lambda\frac{k-1}{k}}T^{1-\lambda\frac{k-1}{k}}.$$

Since $\lambda > \frac{k}{k-1}$, the optimal fixed fee for every fixed β is $T = z(v=1) = \frac{k-1}{k}\zeta(\beta)$.[5]

In particular, when $\beta = 0$, the resultant expected profit is $\frac{k-1}{k}$. Together with Observation B.1, we deduce:

$$\frac{\pi\left(T=\frac{k-1}{k},\beta=0\right)}{\pi^*} = \frac{\frac{k-1}{k}}{\frac{k-1}{k}\left(\frac{\lambda-1}{\lambda}\right)^{k/(k-1)}\frac{\lambda}{\lambda-\frac{k}{k-1}}} = \left(\frac{\lambda}{\lambda-1}\right)^{k/(k-1)}\frac{\lambda-\frac{k}{k-1}}{\lambda} \to_{k\to\frac{\lambda}{\lambda-1}^+} 0,$$

as needed.

For the last bullet, consider a posted-price mechanism offering x units for a posted price-per-unit of p. The expected profit is:

$$\Pr[v\geq p](px-C(x)) \leq (1-F(p))W(p) = \left(1-\frac{1}{k}\right)(1-F(p))p^{k/(k-1)}.$$

[5]Note also that the component of the expected profit coming from the markup is weakly decreasing in T.

The last expression is $\left(1-\frac{1}{k}\right)p^{k/(k-1)-\lambda}$ for $p \geq 1$ and is $\left(1-\frac{1}{k}\right)p^{k/(k-1)}$ for $p < 1$. In any case, it is upper-bounded by $1-\frac{1}{k}$. Thus, the expected profit of a posted-price mechanism is upper-bounded by $1-\frac{1}{k}$.[6] Observation B.1 completes the proof. □

C Proofs from Section 4

C.1 Positive results: Proof of Theorem 1

To start with, we first complete the proof of Lemma 4.2.

Proof of Lemma 4.2. First fix $x \geq y > 0$. The lower elasticity bound gives

$$\log \frac{C'(x)}{C'(y)} = \int_y^x \frac{C''(t)}{C'(t)}\, dt \geq m \int_y^x \frac{dt}{t} = m \log \frac{x}{y}.$$

Thus

$$\frac{C'(x)}{C'(y)} \geq \left(\frac{x}{y}\right)^m.$$

Taking $x = q(v)$ and $y = q(v/s)$, and using $C'(q(z)) = z$, gives

$$s = \frac{C'(q(v))}{C'(q(v/s))} \geq \left(\frac{q(v)}{q(v/s)}\right)^m.$$

Hence $q(v/s) \geq s^{-1/m} q(v)$.

It remains to prove the cost-share bounds. Let $q = q(v)$, so $C'(q) = v$. For $0 < t \leq q$, the upper elasticity bound gives

$$\log \frac{C'(q)}{C'(t)} = \int_t^q \frac{C''(u)}{C'(u)}\, du \leq M \log \frac{q}{t},$$

and therefore $C'(t) \geq C'(q)(t/q)^M$. Integrating,

$$C(q) = \int_0^q C'(t)\, dt \geq C'(q) \int_0^q \left(\frac{t}{q}\right)^M dt = \frac{vq}{M+1}.$$

Similarly, the lower elasticity bound gives

$$\log \frac{C'(q)}{C'(t)} \geq m \log \frac{q}{t},$$

so $C'(t) \leq C'(q)(t/q)^m$. Integrating again,

$$C(q) \leq C'(q) \int_0^q \left(\frac{t}{q}\right)^m dt = \frac{vq}{m+1}.$$

Substituting $q = q(v)$ proves the desired cost-share inequalities.

[6] Note that setting $x = p = 1$ achieves the upper bound.

$\square$

We now complete the proof of Theorem 1.

Proof of Theorem 1. Assume throughout this proof that:[7]

$$m \leq \eta_{C'}(q) \leq M, \qquad \forall q > 0.$$

We prove a slightly more general statement. Fix any markup $\beta \geq 0$, set $\tau := 1 + \beta$, and choose the fixed fee

$$T := \tau\, W\left(\frac{v_0}{\tau}\right), \qquad v_0 := \inf\{v : \varphi(v) \geq 0\},$$

where v_0 is the lowest type with nonnegative virtual value. We show that:

$$\pi(T, \beta) \geq \frac{M+1}{M} \cdot \frac{\beta}{(1+\beta)^{1+1/m}} \cdot \pi^*.$$

The claimed guarantee follows by setting $\beta = m$.

The optimal profit benchmark. Regularity implies that $\varphi(v) \geq 0$ precisely on the upper tail $v \geq v_0$. By Fact 3.1,

$$\pi^* = \int_{v_0}^{\overline{v}} f(v)\, W(\varphi(v))\, dv.$$

The cost-share statement in Lemma 4.2 implies:

$$C(q(x)) \geq \frac{1}{M+1} x q(x), \qquad \forall x.$$

Hence:

$$W(x) = xq(x) - C(q(x)) \leq \frac{M}{M+1} x q(x).$$

Applying this with $x = \varphi(v)$, and using $\varphi(v) \leq v$ together with monotonicity of $q(\cdot)$, gives:

$$\pi^* \leq \frac{M}{M+1} \int_{v_0}^{\overline{v}} f(v)\, \varphi(v)\, q(v)\, dv. \tag{13}$$

It is useful to name the remaining integral:

$$B := \int_{v_0}^{\overline{v}} f(v)\, \varphi(v)\, q(v)\, dv$$

This is the virtual-value benchmark against which we compare the tariff profit. The preceding display says

$$\pi^* \leq \frac{M}{M+1} B. \tag{14}$$

[7] If only the lower bound is available, we set $M := +\infty$, $M/(M+1) := 1$, and the proof remains valid.

Rewriting tariff profit as virtual surplus. First ignore the entry fee. A type-v buyer then chooses a quantity solving

$$\max_{q\geq 0}\{vq-\tau C(q)\}=\tau\max_{q\geq 0}\left\{\frac{v}{\tau}q-C(q)\right\},$$

and hence chooses $q(v/\tau)$. The corresponding utility before paying the entry fee is $\tau W(v/\tau)$. Since $T=\tau W(v_0/\tau)$ and $W(\cdot)$ is monotone, exactly the types $v\geq v_0$ participate.

The seller's expected profit is therefore:

$$\pi(T,\beta)=(1-F(v_0))\tau W\left(\frac{v_0}{\tau}\right)+\beta\int_{v_0}^{\overline{v}}f(v)C\left(q\left(\frac{v}{\tau}\right)\right)dv. \tag{15}$$

Using

$$C\left(q\left(\frac{v}{\tau}\right)\right)=\frac{v}{\tau}q\left(\frac{v}{\tau}\right)-W\left(\frac{v}{\tau}\right),$$

we rewrite (15) as:

$$\begin{aligned}\pi(T,\beta)=(1-F(v_0))\tau W\left(\frac{v_0}{\tau}\right)&+\frac{\beta}{\tau}\int_{v_0}^{\overline{v}}f(v)vq\left(\frac{v}{\tau}\right)dv\\ &-\beta\int_{v_0}^{\overline{v}}f(v)W\left(\frac{v}{\tau}\right)dv.\end{aligned} \tag{16}$$

The last integral is the only term that still needs to be rewritten. By Lemma 4.1, followed by the change of variables $u=t/\tau$,

$$W\left(\frac{v}{\tau}\right)=W\left(\frac{v_0}{\tau}\right)+\frac{1}{\tau}\int_{v_0}^{v}q\left(\frac{t}{\tau}\right)dt.$$

Integrating both sides with respect to $f(v)dv$ over the participating types gives:

$$\begin{aligned}\int_{v_0}^{\overline{v}}f(v)W\left(\frac{v}{\tau}\right)dv&=(1-F(v_0))W\left(\frac{v_0}{\tau}\right)+\frac{1}{\tau}\int_{v_0}^{\overline{v}}f(v)\int_{v_0}^{v}q\left(\frac{t}{\tau}\right)dt\,dv\\ &=(1-F(v_0))W\left(\frac{v_0}{\tau}\right)+\frac{1}{\tau}\int_{v_0}^{\overline{v}}q\left(\frac{t}{\tau}\right)\int_{t}^{\overline{v}}f(v)\,dv\,dt\\ &=(1-F(v_0))W\left(\frac{v_0}{\tau}\right)+\frac{1}{\tau}\int_{v_0}^{\overline{v}}(1-F(t))q\left(\frac{t}{\tau}\right)dt.\end{aligned} \tag{17}$$

Here the second equality swaps the order of integration over the region $\{(t,v):v_0\leq t\leq v\leq\overline{v}\}$, and the last equality uses $\int_t^{\overline{v}}f(v)\,dv=1-F(t)$.

Substituting (17) into (16), and then using $\tau=1+\beta$ and $f(v)\varphi(v)=f(v)v-(1-F(v))$ gives:

$$\begin{aligned}\pi(T,\beta)&=(1-F(v_0))\tau W\left(\frac{v_0}{\tau}\right)+\frac{\beta}{\tau}\int_{v_0}^{\overline{v}}f(v)vq\left(\frac{v}{\tau}\right)dv\\ &\quad-\beta(1-F(v_0))W\left(\frac{v_0}{\tau}\right)-\frac{\beta}{\tau}\int_{v_0}^{\overline{v}}(1-F(v))q\left(\frac{v}{\tau}\right)dv\\ &=(\tau-\beta)(1-F(v_0))W\left(\frac{v_0}{\tau}\right)+\frac{\beta}{\tau}\int_{v_0}^{\overline{v}}\big(f(v)v-(1-F(v))\big)q\left(\frac{v}{\tau}\right)dv\\ &=(1-F(v_0))W\left(\frac{v_0}{\tau}\right)+\frac{\beta}{\tau}\int_{v_0}^{\overline{v}}f(v)\varphi(v)\,q\left(\frac{v}{\tau}\right)dv.\end{aligned} \tag{18}$$

We now compare (18) to the benchmark B. Dropping the nonnegative first term, and using the scaling statement in Lemma 4.2, we obtain:

$$\pi(T,\beta) \geq \frac{\beta}{\tau}\int_{v_0}^{\overline{v}} f(v)\varphi(v)\, q\left(\frac{v}{\tau}\right)dv \geq \frac{\beta}{\tau^{1+1/m}}\int_{v_0}^{\overline{v}} f(v)\varphi(v)\, q(v)\, dv = \frac{\beta}{\tau^{1+1/m}}B,$$

where the second inequality uses $\varphi(v) \geq 0$ for $v \geq v_0$. Combining this lower bound with (14) gives:

$$\pi(T,\beta) \geq \frac{M+1}{M}\cdot\frac{\beta}{(1+\beta)^{1+1/m}}\cdot\pi^*.$$

It remains to choose the markup. Set

$$g(\beta) := \frac{\beta}{(1+\beta)^{1+1/m}}.$$

Then:

$$g'(\beta) = (1+\beta)^{-2-1/m}\left(1-\frac{\beta}{m}\right),$$

so g is maximized at $\beta = m$, with

$$g(m) = \frac{m}{(m+1)^{1+1/m}}.$$

Therefore the universal choice $\beta = m$ gives:

$$\pi(T,m) \geq \frac{M+1}{M}\cdot\frac{m}{(m+1)^{1+1/m}}\cdot\pi^*$$

and completes the proof. □

C.2 Tightness Results: Proof of Proposition 4.1

The positive result of $\frac{1}{k^{1/(k-1)}}$ is straighforward. By substituting $m = M = k-1$ into Theorem 1, we then get the ratio.

In the following, we show this ratio is tight for regular distributions and power functions. Fixing any $\eta > 0$, the proof proceeds in three steps. First, we construct a family of hard instances in quantile space. Second, we verify that the induced distribution has the claimed support and regularity properties. Finally, we show that, for a suitable choice of the construction parameters, every two-part tariff extracts at most a

$$k^{-1/(k-1)} + \eta$$

fraction of the optimal profit.

C.2.1 Construction of the hard instance

Let

$$p := \frac{k}{k-1}.$$

For the monomial cost $C(q) = q^k/k$, the welfare-maximizing quantity and the corresponding welfare are

$$q(v) := \operatorname*{argmax}_{q\geq 0}\{vq - C(q)\} = v^{1/(k-1)} \qquad \text{and} \qquad W(v) := \max_{q\geq 0}\{vq - C(q)\} = \frac{k-1}{k}v^p.$$

We specify the hard instance through its upper-quantile revenue curve. Fix an integer $n \geq 2$ and a parameter $\varepsilon \in (0, 1/2)$. Define upper-quantile breakpoints

$$q_0 := 0, \qquad q_i := \varepsilon^{p(n-i)} \qquad (i = 1, \dots, n),$$

so that

$$0 = q_0 < q_1 < \cdots < q_n = 1.$$

Define a step function $\phi : [0, 1] \to \mathbb{R}_+$ by

$$\phi(q) = \varepsilon^{i-1} \qquad \text{for } q \in (q_{i-1}, q_i], \quad i = 1, \dots, n.$$

The function ϕ represents the virtual value as a function of the upper quantile.

Let

$$R(q) := \int_0^q \phi(t)\,dt, \qquad v(q) := \frac{R(q)}{q} \quad (q > 0),$$

and let $F = F_{n,\varepsilon}$ be the distribution whose upper-quantile function is v; equivalently, $V = v(U)$ for $U \sim \mathrm{Unif}[0, 1]$. By construction, $R(q) = qv(q)$ is the revenue curve of F.

C.2.2 Regularity

We first verify that the distribution constructed above has the claimed support and regularity properties. Since ϕ is nonincreasing, the revenue curve R is concave. Hence the induced distribution F is regular.

It remains to describe the support and density of F. On the first block $(0, q_1]$, we have $R(q) = q$, and hence $v(q) \equiv 1$. Thus F has an atom of mass q_1 at $v = 1$. For every $i \geq 2$, the function R is affine on $(q_{i-1}, q_i]$ with slope ε^{i-1}. Therefore, there exists B_i such that

$$R(q) = \varepsilon^{i-1}q + B_i \qquad \text{for } q \in (q_{i-1}, q_i].$$

Equivalently,

$$v(q) = \varepsilon^{i-1} + \frac{B_i}{q} \qquad \text{for } q \in (q_{i-1}, q_i].$$

Moreover,

$$B_i = R(q_{i-1}) - \varepsilon^{i-1}q_{i-1} = q_{i-1}\big(v(q_{i-1}) - \varepsilon^{i-1}\big) > 0,$$

because $v(q_{i-1})$ is the average of slopes of R on $[0, q_{i-1}]$, all of which are at least $\varepsilon^{i-2} > \varepsilon^{i-1}$.

It follows that v is continuous on $(0,1]$, strictly decreasing on $(q_1,1]$, and constant on $(0,q_1]$. Write

$$v_i := v(q_i), \qquad i = 1,\dots,n, \qquad \ell := v_n = v(1).$$

Since $0 < \phi \le 1$, and $\phi \not\equiv 1$ for $n \ge 2$, we have

$$0 < \ell = R(1) < 1.$$

Thus F is supported on the interval $[\ell,1]$, has one atom at $v = 1$, and is continuous on $[\ell,1)$.

For $i \ge 2$ and $v \in [v_i, v_{i-1})$, the inverse map is

$$q = \frac{B_i}{v - \varepsilon^{i-1}}.$$

Hence, on this interval,

$$F(v) = 1 - \frac{B_i}{v - \varepsilon^{i-1}}, \qquad f(v) = \frac{B_i}{(v - \varepsilon^{i-1})^2} > 0.$$

The virtual value on this interval is

$$\varphi(v) = v - \frac{1 - F(v)}{f(v)} = v - \frac{B_i}{v - \varepsilon^{i-1}} \cdot \frac{(v - \varepsilon^{i-1})^2}{B_i} = \varepsilon^{i-1}.$$

As v increases, the corresponding block index i weakly decreases, and therefore ε^{i-1} weakly increases. Thus the virtual value is monotone on the continuous part of the support.

C.2.3 Two-part tariff upper bound

We first compute the optimal profit for this instance. Since $\phi(q) \ge 0$ is the virtual value at upper quantile q, Fact 3.1 gives

$$\pi^* = \int_0^1 W(\phi(q))\,dq = \frac{k-1}{k} \int_0^1 \phi(q)^p\,dq.$$

Let

$$S := \int_0^1 \phi(q)^p\,dq.$$

Then

$$\pi^* = \frac{k-1}{k} S.$$

We next upper-bound the profit of an arbitrary two-part tariff (T,β). The argument proceeds by parameterizing any tariff by its markup β and the upper-quantile cutoff q of participating types. For each fixed cutoff q, we first optimize over the markup. We then show that, for every fixed markup, the best cutoff is attained at one of the breakpoints $q_1,\dots,q_n$. It remains only to bound the profit of tariffs whose cutoffs are among these breakpoints.

For a type with value v, the consumption choice under markup β is

$$q\left(\frac{v}{1+\beta}\right).$$

The buyer's gross utility before paying the entry fee is

$$(1+\beta)W\left(\frac{v}{1+\beta}\right).$$

This quantity is increasing in v, so any fixed fee induces a cutoff rule: the participating types are the top quantiles of the distribution. Let $q \in (0,1]$ denote the upper-quantile cutoff, so that $v(q)$ is the lowest participating value. For fixed β and cutoff q, the largest fixed fee consistent with participation of the cutoff type is

$$T = (1+\beta)W\left(\frac{v(q)}{1+\beta}\right).$$

Therefore every tariff with markup β and participation cutoff q has profit at most the profit obtained with this maximal entry fee.

Define

$$U(q) := q\,v(q)^p, \qquad H(q) := \int_0^q v(t)^p\,dt.$$

Using

$$W(z) = \frac{k-1}{k}z^p \qquad \text{and} \qquad C(q(z)) = \frac{1}{k}z^p,$$

the corresponding upper bound on profit is

$$\Pi(\beta,q) = \frac{k-1}{k}(1+\beta)^{1-p}U(q) + \frac{\beta}{k(1+\beta)^p}H(q). \tag{19}$$

Indeed, the first term is the entry-fee revenue, and the second term is the usage markup revenue net of production cost.

Let

$$z := (1+\beta)^{-1} \in (0,1], \qquad A(q) := (k-1)U(q) + H(q).$$

For a fixed cutoff q, optimizing over the markup $\beta \geq 0$ is equivalent to optimizing over $z \in (0,1]$. In terms of z, (19) becomes

$$\Pi(\beta,q) = \frac{1}{k}z^{p-1}\big(A(q) - zH(q)\big).$$

Notice that $A(q)$ and $H(q)$ are constants when q is fixed. The first-order condition with respect to z is

$$(p-1)A(q) - pzH(q) = 0.$$

Since $p = k/(k-1)$, this gives the unconstrained maximizer

$$z^*(q) = \frac{A(q)}{kH(q)}.$$

To check that this maximizer is feasible, note that $v(t) \geq v(q)$ for all $t \leq q$. Hence

$$H(q) = \int_0^q v(t)^p\,dt \geq q\,v(q)^p = U(q).$$

Therefore

$$A(q) = (k-1)U(q) + H(q) \le kH(q),$$

so $z^*(q) \in (0,1]$ and therefore corresponds to a feasible markup $\beta^*(q) = 1/z^*(q) - 1$. Let $G(q)$ denote the maximum approximation ratio among two-part tariffs with upper-quantile cutoff q. Substituting $z^*(q)$ gives

$$G(q) := \sup_{\beta \ge 0} \frac{\Pi(\beta, q)}{\pi^*} = \frac{A(q)^p}{k^p H(q)^{p-1} S} = \frac{\big((k-1)U(q) + H(q)\big)^p}{k^p H(q)^{p-1} S}. \tag{20}$$

We next show that, for every fixed markup, the best cutoff is attained at one of the breakpoints.

Claim C.1. *For every fixed $\beta \ge 0$, the function $q \mapsto \Pi(\beta, q)$ attains its maximum over $(0,1]$ at one of the points $q_1, \ldots, q_n$.*

Proof. On the first block $(0, q_1]$, we have $v(q) \equiv 1$. Hence

$$U(q) = H(q) = q,$$

and therefore (19) reduces to

$$\Pi(\beta, q) = q\left[\frac{k-1}{k}(1+\beta)^{1-p} + \frac{\beta}{k(1+\beta)^p}\right].$$

The bracketed term is nonnegative and does not depend on q. Thus $\Pi(\beta, q)$ is increasing in q on this block, and the maximum over $(0, q_1]$ is attained at q_1.

Now fix $i \ge 2$ and $q \in (q_{i-1}, q_i]$. We have

$$v(q) = \varepsilon^{i-1} + \frac{B_i}{q}, \qquad v'(q) = -\frac{B_i}{q^2}.$$

Therefore

$$\begin{aligned} U'(q) &= \frac{d}{dq}\big(qv(q)^p\big) \\ &= v(q)^p + qpv(q)^{p-1}v'(q) \\ &= v(q)^{p-1}\left(\varepsilon^{i-1} - \frac{B_i}{(k-1)q}\right), \end{aligned}$$

where we used $p = k/(k-1)$. Also,

$$H'(q) = v(q)^p.$$

Differentiating (19) gives

$$k(1+\beta)^p \partial_q \Pi(\beta, q) = v(q)^{p-1}\left((k-1+k\beta)\varepsilon^{i-1} - \frac{B_i}{q}\right).$$

The prefactor $v(q)^{p-1}$ is positive, and the term in parentheses is strictly increasing in q. Hence $\partial_q \Pi(\beta, q)$ changes sign at most once on $(q_{i-1}, q_i]$, from negative to positive. Thus $\Pi(\beta, \cdot)$ has no strict interior maximum on this block, and its maximum over the block is attained at an endpoint.

Combining all blocks proves the claim. □

Now any two-part tariff is bounded above by some $\Pi(\beta, q)$. Therefore, by Claim C.1,

$$\sup_{T,\beta} \frac{\pi(T,\beta)}{\pi^*} \le \sup_{\beta\ge 0} \max_{i\in[n]} \frac{\Pi(\beta,q_i)}{\pi^*}$$
$$\le \max_{i\in[n]} G(q_i).$$

It remains to estimate $G(q_i)$.

Let

$$c := q_1 = \varepsilon^{p(n-1)}.$$

Throughout, $o_n(1)$ denotes a quantity that tends to 0 as $\varepsilon \to 0$ for fixed n, uniformly over all relevant $i \in \{1,\dots,n\}$.

Recall that $S = \int_0^1 \phi(q)^p \, dq$. Using the definition of ϕ on each quantile block, we get

$$\begin{aligned} S &= q_1 + \sum_{i=2}^{n} \varepsilon^{p(i-1)}(q_i - q_{i-1}) \\ &= c + \sum_{i=2}^{n} \varepsilon^{p(i-1)} \varepsilon^{p(n-i)}(1-\varepsilon^p) \\ &= c\big(n - (n-1)\varepsilon^p\big). \end{aligned} \tag{21}$$

Next define

$$\lambda_i := \frac{v_i}{\varepsilon^{i-1}}, \qquad i = 1,\dots,n.$$

Since $v_1 = 1$, we have $\lambda_1 = 1$. For $i = 1,\dots,n-1$, the slope of R on $(q_i, q_{i+1}]$ is ε^i. Therefore

$$R(q_{i+1}) = R(q_i) + \varepsilon^i (q_{i+1} - q_i).$$

Using $v_i = R(q_i)/q_i$ and $q_i/q_{i+1} = \varepsilon^p$, we get

$$\begin{aligned} v_{i+1} &= \frac{R(q_{i+1})}{q_{i+1}} \\ &= \frac{R(q_i) + \varepsilon^i (q_{i+1} - q_i)}{q_{i+1}} \\ &= \frac{q_i}{q_{i+1}} v_i + \varepsilon^i \left(1 - \frac{q_i}{q_{i+1}}\right) \\ &= \varepsilon^p v_i + \varepsilon^i (1-\varepsilon^p). \end{aligned}$$

Dividing by ε^i, we obtain

$$\lambda_{i+1} = 1 + \varepsilon^{p-1}(\lambda_i - \varepsilon). \tag{22}$$

For fixed n, this recursion implies

$$\lambda_i = 1 + o_n(1)$$

uniformly over $i = 1, \dots, n$. Hence

$$U(q_i) = q_i v_i^p = \varepsilon^{p(n-i)} \big(\varepsilon^{i-1} \lambda_i\big)^p = c\,\lambda_i^p = c\big(1 + o_n(1)\big) \tag{23}$$

uniformly in i.

We now estimate $H(q_i)$. For $i \geq 2$, write

$$\Delta H_i := H(q_i) - H(q_{i-1}).$$

For $q \in [q_{i-1}, q_i]$, we have

$$v(q) = \varepsilon^{i-1} + \frac{B_i}{q}, \qquad B_i = q_{i-1}(v_{i-1} - \varepsilon^{i-1}).$$

Substitute $q = q_i y$, where $y \in [\varepsilon^p, 1]$. Using $q_{i-1}/q_i = \varepsilon^p$ and $v_{i-1} = \varepsilon^{i-2} \lambda_{i-1}$, we obtain

$$\begin{aligned} v(q_i y) &= \varepsilon^{i-1} + \frac{q_{i-1}(v_{i-1} - \varepsilon^{i-1})}{q_i y} \\ &= \varepsilon^{i-1} + \frac{\varepsilon^p(\varepsilon^{i-2}\lambda_{i-1} - \varepsilon^{i-1})}{y} \\ &= \varepsilon^{i-1}\left(1 + \frac{\varepsilon^{p-1}(\lambda_{i-1} - \varepsilon)}{y}\right). \end{aligned}$$

Therefore

$$\Delta H_i = \int_{q_{i-1}}^{q_i} v(q)^p \, dq = c\, J_{A_i}(\varepsilon), \qquad A_i := \lambda_{i-1} - \varepsilon,$$

where

$$J_A(\varepsilon) := \int_{\varepsilon^p}^1 \left(1 + \frac{A\varepsilon^{p-1}}{y}\right)^p dy.$$

Claim C.2. *Uniformly for A in any compact subset of $(0, \infty)$,*

$$J_A(\varepsilon) = 1 + (k-1)A^p + o(1) \qquad \text{as } \varepsilon \to 0.$$

Proof. Set

$$\theta := A\varepsilon^{p-1}.$$

With the substitution $y = \theta/t$, we obtain

$$J_A(\varepsilon) = \theta \int_\theta^{A/\varepsilon} \frac{(1+t)^p}{t^2} \, dt.$$

Integrating by parts,

$$\begin{aligned} J_A(\varepsilon) = (1+\theta)^p - \varepsilon^p \left(1 + \frac{A}{\varepsilon}\right)^p \\ + p\theta \int_\theta^{A/\varepsilon} \frac{(1+t)^{p-1}}{t} \, dt. \end{aligned}$$

The first term is $1+o(1)$, uniformly for A in compact subsets of $(0,\infty)$. The second term is

$$\varepsilon^p\left(1+\frac{A}{\varepsilon}\right)^p = (A+\varepsilon)^p = A^p + o(1),$$

again uniformly on compact subsets.

It remains to estimate the integral term. On $[\theta,1]$, the factor $(1+t)^{p-1}$ is uniformly bounded above by 2^{p-1}. Hence

$$\theta\int_\theta^1 \frac{(1+t)^{p-1}}{t}\,dt \le 2^{p-1}\theta\int_\theta^1 \frac{dt}{t} = 2^{p-1}\theta\log\frac{1}{\theta} = o(1),$$

where the last equality follows from $\theta = A\varepsilon^{p-1}\to 0$. On $[1, A/\varepsilon]$,

$$(1+t)^{p-1} = t^{p-1} + O(t^{p-2}),$$

uniformly. Hence

$$\begin{aligned}\theta\int_1^{A/\varepsilon}\frac{(1+t)^{p-1}}{t}\,dt &= \theta\int_1^{A/\varepsilon} t^{p-2}\,dt + \theta\int_1^{A/\varepsilon} O(t^{p-3})\,dt\\ &= \frac{A^p}{p-1} + o(1).\end{aligned}$$

The error term is $o(1)$: if $p<2$ the integral of t^{p-3} is bounded; if $p=2$ it contributes $O(\varepsilon\log(1/\varepsilon))$; and if $p>2$ it contributes $O(\varepsilon)$.

Therefore

$$p\theta\int_\theta^{A/\varepsilon}\frac{(1+t)^{p-1}}{t}\,dt = \frac{p}{p-1}A^p + o(1) = kA^p + o(1).$$

Combining the estimates gives

$$J_A(\varepsilon) = 1 - A^p + kA^p + o(1) = 1 + (k-1)A^p + o(1),$$

as claimed. □

Since $A_i = \lambda_{i-1} - \varepsilon = 1 + o_n(1)$ uniformly in i, all A_i's lie in a fixed compact subset of $(0,\infty)$ for sufficiently small ε. Thus Claim C.2 implies

$$\Delta H_i = c\big(k + o_n(1)\big) \qquad \text{uniformly for } i = 2,\dots,n.$$

Since $H(q_1) = q_1 = c$, we get

$$H(q_i) = c + \sum_{j=2}^{i}\Delta H_j = c\big(ki - (k-1) + o_n(1)\big) \tag{24}$$

uniformly over $i = 1,\dots,n$.

Endpoint ratios. Combining (23) and (24), we have

$$(k-1)U(q_i) + H(q_i) = c\big(ki + o_n(1)\big),$$

Recall from (20) that

$$G(q_i) = \frac{\big((k-1)U(q_i) + H(q_i)\big)^p}{k^p H(q_i)^{p-1} S}.$$

Also, by (21), $S = c(n + o_n(1))$. Therefore, uniformly over $i = 1, \dots, n$,

$$G(q_i) = \frac{i^p}{n\big(ki - (k-1)\big)^{p-1}} + o_n(1). \tag{25}$$

Define

$$g(x) := \frac{x^p}{\big(kx - (k-1)\big)^{p-1}}, \qquad x \geq 1.$$

Then

$$\begin{aligned} \frac{d}{dx} \log g(x) &= \frac{p}{x} - \frac{k(p-1)}{kx - (k-1)} \\ &= \frac{p(k-1)(x-1)}{x\big(kx - (k-1)\big)} \geq 0. \end{aligned}$$

Thus g is increasing on $[1, \infty)$. By the uniformity of (25),

$$\begin{aligned} \max_{1 \leq i \leq n} G(q_i) &= \frac{1}{n} \max_{1 \leq i \leq n} \frac{i^p}{\big(ki - (k-1)\big)^{p-1}} + o_n(1) \\ &= \frac{1}{n} \frac{n^p}{\big(kn - (k-1)\big)^{p-1}} + o_n(1) \\ &= \left(\frac{n}{kn - (k-1)}\right)^{p-1} + o_n(1) \\ &= \left(\frac{n}{kn - (k-1)}\right)^{1/(k-1)} + o_n(1). \end{aligned}$$

Consequently,

$$\sup_{T, \beta} \frac{\pi(T, \beta)}{\pi^*} \leq \left(\frac{n}{kn - (k-1)}\right)^{1/(k-1)} + o_n(1).$$

Finally,

$$r_{k,n} := \left(\frac{n}{kn - (k-1)}\right)^{1/(k-1)}$$

decreases to $k^{-1/(k-1)}$ as $n \to \infty$. Choose n large enough so that

$$r_{k,n} \leq k^{-1/(k-1)} + \frac{\eta}{2}.$$

Then choose $\varepsilon > 0$ small enough that the $o_n(1)$ term above is at most $\eta/2$. For this choice of n and ε, the constructed distribution satisfies

$$\sup_{T,\beta} \frac{\pi(T,\beta)}{\pi^*} \leq k^{-1/(k-1)} + \eta.$$

Removing the top atom. Theorem 1 remains valid for a regular distribution that is absolutely continuous below its top value but has an atom at that value. Thus the construction above already establishes tightness in this slightly enlarged class. To obtain an instance satisfying our maintained absolute-continuity assumption, replace the first segment $R_0(q) = q$, $q \in [0, q_1]$, by

$$R_\xi(q) = (1+\xi)q - \frac{\xi}{q_1} q^2,$$

and leave the rest of the revenue curve unchanged. For all sufficiently small $\xi > 0$, the new curve remains concave and simply spreads the top atom over $[1, 1+\xi]$ with a positive density. As $\xi \to 0$, the optimal profit and the maximum profit achievable by a two-part tariff converge to those in the construction above. Hence the same bound holds with an arbitrarily small additional loss.

This proves the theorem.

D The Role of the Lower Bound on the Elasticity: Proof of Proposition 4.2

Proof of Proposition 4.2. Fix $\varepsilon > 0$. Note that we have by the first-order condition 5 that $q(v) = \log v$ for $v \geq 1$. Thus, $W(v) = v \log v - v + 1$ for $v \geq 1$. In particular, $W(v)$ is continuous and strictly increasing for $v \geq 1$.

Fix a sufficiently small $\delta > 0$ (we shall later let it tend to 0). Define a distribution $F(v) := F_\delta(v) := G(W(v)) := G_\delta(W(v))$ by $G(W) = 0$ for $W < \delta$ and $G(W) = \min\left\{1, \frac{1}{1-\sqrt{\delta}}\left(1 - \frac{\delta}{W}\right)\right\}$ for $W \geq \delta$. Then $\mathrm{Supp}(F) = \left\{v : W(v) \in \left[\delta, \sqrt{\delta}\right]\right\} \subseteq [1, 2]$.

We have for $v = 1 + o(1)$:

$$0.25(v-1)^2 \leq W(v) \leq v(v-1) - (v-1) = (v-1)^2 \implies \sqrt{W(v)} + 1 \leq v \leq 2\sqrt{W(v)} + 1. \tag{26}$$

In particular, note that:

$$\mathrm{Supp}(F) = \left\{v : W(v) \in \left[\delta, \sqrt{\delta}\right]\right\} \subseteq \left[1 + \delta^{0.5}, 1 + 2\delta^{0.25}\right]. \tag{27}$$

Moreover:

$$\varphi(v) = v - \frac{1 - F(v)}{f(v)} = v - \frac{1 - \frac{1}{1-\sqrt{\delta}}\left(1 - \frac{\delta}{W(v)}\right)}{\frac{\delta}{1-\sqrt{\delta}} \frac{W'(v)}{W^2(v)}} = v - W(v) \frac{\left(1 - \sqrt{\delta}\right)W(v) - \left(W(v) - \delta\right)}{\delta \log v}$$
$$= v - \frac{W(v)\left(\sqrt{\delta} - W(v)\right)}{\sqrt{\delta} \log v}.$$

The derivative with respect to v inside $\mathrm{Supp}(F)$ is (note that $0 \le W(v) \le \sqrt{\delta}$ and $q(v) = \log v$):

$$1 - \frac{\left(\sqrt{\delta} q(v) - 2W(v) q(v)\right)\log v - W(v)\left(\sqrt{\delta} - W(v)\right)(1/v)}{\sqrt{\delta}\log^2 v} \ge 1 - \frac{\left(\sqrt{\delta} - 2W(v)\right)\log^2 v}{\sqrt{\delta}\log^2 v} = \frac{2W(v)}{\sqrt{\delta}} \ge 0,$$

which implies that F is regular.

We shall show that for every $\varepsilon > 0$ there exists $\delta > 0$ s.t. the distribution $F = F_\delta$ yields $\frac{\pi(T,\beta)}{\pi^*} \le \varepsilon$ for every T, β.

A simple derivation shows that the expression $z_\beta(v) := v\, q\left(\frac{v}{1+\beta}\right) - (1+\beta) C\left(q\left(\frac{v}{1+\beta}\right)\right)$, for every fixed v, is decreasing with respect to β. Therefore, the expected profit attributed to the fixed fee T satisfies for $T \in \left[\delta, \sqrt{\delta}\right]$:

$$T\Pr[z(v) \ge T] \le T\Pr[W(v) \ge T] = \left(1 - \frac{1}{1-\sqrt{\delta}}\right)T + \frac{\delta}{1-\sqrt{\delta}} \le \frac{\delta}{1-\sqrt{\delta}}. \tag{28}$$

Note that the above inequality actually holds for any $T \ge 0$, since $T\Pr[z(v) \ge T] \le \delta \le \frac{\delta}{1-\sqrt{\delta}}$ for $0 \le T < \delta$, and $T\Pr[z(v) \ge T] = 0$ for $T > \sqrt{\delta}$.

Moreover, the expected profit attributed to the markup is:

$$\beta \mathbb{E}\left[C\left(q\left(\frac{v}{1+\beta}\right)\right) \cdot 1_{z(v) \ge T}\right] \le \beta \mathbb{E}\left[C\left(q\left(\frac{v}{1+\beta}\right)\right)\right] = \frac{\beta\delta}{1-\sqrt{\delta}} \int_{W=\max\{W(v=1+\beta),\delta\}}^{W=\sqrt{\delta}} \frac{C\left(q\left(\frac{v(W)}{1+\beta}\right)\right)}{W^2} dW. \tag{29}$$

Note that by (26), unless $\beta = \beta(\delta) \le 2\delta^{0.25} = o(1)$, the two-part tariff yields expected profit of 0. We get from (26) that for $1+\beta \le v \le W^{-1}\left(\sqrt{\delta}\right)$:

$$\frac{C\left(q\left(\frac{v}{1+\beta}\right)\right)}{W^2(v)} = \frac{v-1-\beta}{(1+\beta)W^2(v)} \le \frac{v-1}{W^2(v)} \le \frac{2}{W^{1.5}(v)}. \tag{30}$$

From (26), (28), (29), and (30), we get that for every $\beta > 0$:

$$\pi \le \frac{\delta}{1-\sqrt{\delta}}\left(1 + 2\beta \int_{W=W(v=1+\beta)}^{W=\sqrt{\delta}} W^{-1.5} dW\right) \le \frac{\delta}{1-\sqrt{\delta}}\left(1 + \frac{4\beta}{\sqrt{W(v=1+\beta)}}\right) \le \frac{\delta}{1-\sqrt{\delta}}\left(1 + \frac{8\beta}{\beta}\right) = O(\delta), \tag{31}$$

while for $\beta = 0$:

$$\pi \le \frac{\delta}{1-\sqrt{\delta}} = O(\delta). \tag{32}$$

To easily lower-bound π^*, consider the direct revelation mechanism $(x(v), p(v))$, with $x(v) = p(v) = 0$ for $0 \le v < 1$, and $x(v) = 0.5\log v$, $p(v) = 0.5(v-1)$ for $v \ge 1$. This mechanism is IC and IR since for any $t \ge 1$ and $v \ge 1$: $v x(t) - p(t) = 0.5(v\log t - t + 1)$, which attains its global maximum for a fixed v at $t = v$, and the maximum value is non-negative when $v \ge 1$. The profit of this mechanism for any $v \in \mathrm{Supp}(F)$ is:

$$0.5 + 0.5v - v^{0.5} \geq \frac{(v-1)^2}{8} - \frac{(v-1)^3}{16} \geq \frac{(v-1)^2}{16}. \tag{33}$$

From (26) and (33):

$$\pi^* \geq \frac{\delta}{1-\sqrt{\delta}} \int_{W=\delta}^{W=\sqrt{\delta}} \frac{0.5 + 0.5v(W) - v(W)^{0.5}}{W^2} dW \geq \frac{\delta}{1-\sqrt{\delta}} \int_{W=\delta}^{W=\sqrt{\delta}} \frac{(v-1)^2}{16(v-1)^4} dW$$
$$= \frac{\delta}{16\left(1-\sqrt{\delta}\right)} \int_{W=\delta}^{W=\sqrt{\delta}} \frac{1}{(v-1)^2} dW \geq \frac{\delta}{64\left(1-\sqrt{\delta}\right)} \int_{W=\delta}^{W=\sqrt{\delta}} \frac{1}{W} dW = \frac{\delta}{64\left(1-\sqrt{\delta}\right)} \log \frac{1}{\sqrt{\delta}} = \Omega(\delta|\log\delta|).$$

Since $\pi = O(\delta) = o(\delta|\log\delta|)$ when $\delta \to 0^+$ for any $\beta = \beta(\delta)$, we are done.

□

E Proofs from Section 5

E.1 Negative results: The necessity of a menu

In this section, we complete the proofs of Proposition 5.1 and Theorem 3.

Proof of Proposition 5.1. Let $s = k$ be a sufficiently large number, and suppose the cost function is given by $C(q) = \frac{1}{k} q^k$. We begin by constructing a discrete distribution, which we will later extend to a continuous one.

Consider the discrete distribution F supported on the set $\{v_i\}_{i=1}^k$, where each value is defined as $v_i = \left(\frac{1}{\varepsilon}\right)^{(k-1)i}$. The probability mass function of F is given by

$$\Pr_{v \sim F}[v = v_i] = \frac{\varepsilon^{ki}}{C},$$

where C is the appropriate normalization constant.

We begin by computing the optimal profit for this instance. To establish a lower bound on the optimal profit, we construct a specific mechanism and analyze its performance. In this mechanism, the buyer may choose an index $i \in [k]$, receiving a quantity of $\left(\frac{1}{\varepsilon}\right)^i$ while paying $\frac{1}{2}\left(\frac{1}{\varepsilon}\right)^{ki}$ to the seller, or alternatively, choose to opt out.

Suppose the buyer's value is v_i. The utility from selecting an index j is given by

$$\left(\frac{1}{\varepsilon}\right)^{(k-1)i+j} - \frac{1}{2}\left(\frac{1}{\varepsilon}\right)^{kj}.$$

It is straightforward to verify that this expression is maximized when $j = i$, yielding a positive utility. Consequently, a buyer with value v_i will optimally choose index i. The seller's total profit in this mechanism is therefore

$$\sum_{i=1}^k \left(\frac{1}{2}\left(\frac{1}{\varepsilon}\right)^{ki} - \frac{1}{k}\left(\frac{1}{\varepsilon}\right)^{ki}\right) \cdot \frac{\varepsilon^{ki}}{C} = \Theta\left(\frac{k}{C}\right).$$

This implies that the optimal profit is at least of the same order. In what follows, we derive an upper bound on the profit achievable by any two-part tariff mechanism. Consider a mechanism that charges a fixed fee T and a markup β. It is evident that its total profit cannot exceed the sum of two components: the profit of a two-part tariff mechanism that charges a fixed fee T but no markup, and the profit of a mechanism that charges no fixed fee but applies a markup of β. Therefore, we can analyze these two components separately and compute the maximum profit extractable either from the fixed fee or from the markup.

We first analyze the profit generated from the surcharge. Suppose the mechanism charges a markup of 1+β. In this case, the profit when the buyer's value is v is given by

$$\beta C\left(q\left(\frac{v}{\beta+1}\right)\right) = \frac{\beta}{k}\left(\frac{v}{\beta+1}\right)^{\frac{k}{k-1}}.$$

Taking the derivative with respect to β, we obtain

$$\left(\frac{v}{\beta+1}\right)^{\frac{k}{k-1}}\left(1-\frac{\beta}{\beta+1}\cdot\frac{k}{k-1}\right).$$

It follows that the profit is maximized when $\beta = k-1$. In this case, the optimal profit from the surcharge is:

$$\frac{k-1}{k}\left(\frac{v}{k}\right)^{\frac{k}{k-1}}.$$

Therefore, the optimal profit from the surcharge can be computed as follows:

$$\begin{aligned}\frac{k-1}{k}\sum_{i=1}^{k}\left(\frac{1}{k\varepsilon^{(k-1)i}}\right)^{\frac{k}{k-1}}\cdot\Pr_{v\sim F}[v=v_i] &= \frac{k-1}{k}\left(\frac{1}{k}\right)^{\frac{k}{k-1}}\sum_{i=1}^{k}\left(\frac{1}{\varepsilon}\right)^{ki}\cdot\frac{\varepsilon^{ki}}{C}\\ &= \frac{k-1}{k}\left(\frac{1}{k}\right)^{\frac{1}{k-1}}\frac{1}{C} = \Theta\left(\frac{1}{C}\right).\end{aligned}$$

which is at most $\frac{O(1)}{k}$ fraction of the optimal profit, as $\frac{k-1}{k}(\frac{1}{k})^{\frac{1}{k-1}} \to 1$ when k goes to infinity.

Now consider the profit generated from the fixed fee. To maximize this component, the optimal mechanism sets the markup to 0 and selects the fixed fee that maximizes total profit. Observe that

$$W(v_i) = \left(1-\frac{1}{k}\right)\left(\frac{1}{\varepsilon}\right)^{ki}.$$

Hence,

$$W(v_i)\left(\sum_{j=i}^{k}\Pr_{v\sim F}[v=v_j]\right) = O\left(\frac{1}{C}\right)$$

for any $i \in [k]$. This implies that the maximum profit extractable from setting the fixed fee at $W(v_i)$ is at most an $\frac{O(1)}{k}$ fraction of the optimal profit. Since the optimal fixed fee, when the markup is 0, must correspond to one of the values $W(v_i)$ for some $i \in [k]$, the claim follows by taking the limit as $k \to \infty$.

To convert this into a continuous distribution, we simply mix a small probability of a uniform distribution into the current discrete distribution. Specifically, consider a new distribution where

$$v \sim F, \quad v = \begin{cases} U[v_i, v_i + \Delta] & \text{with probability } (1-\Delta)\cdot \frac{\varepsilon^{ki}}{C}, \\ U[v_1, v_k] & \text{with probability } \Delta. \end{cases}$$

By taking $\Delta \to 0$, the argument continues to hold.

□

Proof of Theorem 3. We adopt the same construction as in the proof of Proposition 5.1. Specifically, we consider the following distribution:

$$v \sim F, \quad v = \begin{cases} U[v_i, v_i + \varepsilon] & \text{with probability } (1-\varepsilon^{K^3})\cdot \dfrac{\varepsilon^{Ki}}{C}, \\ U[v_1, v_K] & \text{with probability } \varepsilon^{K^3}, \end{cases}$$

where $v_i = \left(\frac{1}{\varepsilon}\right)^{(K-1)i}$, and $\varepsilon \to 0$ is sufficiently small. Our first goal is to show that this distribution admits at least $K-1$ ironing intervals.

Observe that for any $i \in [K]$, the revenue curve R satisfies

$$\begin{aligned} R(1-F(v_i+\varepsilon)) &= \left(v_i \pm \Theta(\varepsilon)\right)\left(\sum_{j=i+1}^{K} \frac{\varepsilon^{Kj}}{C} \pm \Theta(\varepsilon^{K^3})\right) \\ &= \frac{\varepsilon^{K+i} + o(\varepsilon^{K+i})}{C}, \end{aligned}$$

whereas

$$R(1-F(v_i)) = v_i\left(\sum_{j=i}^{K} \frac{\varepsilon^{Kj}}{C} \pm \Theta(\varepsilon^{K^3})\right) \tag{34}$$

$$= \frac{\varepsilon^{i} + o(\varepsilon^{i})}{C}. \tag{35}$$

It follows that the set $\{(1-F(v_i), R(1-F(v_i)))\}_{i=1}^{K}$ corresponds to the local maxima of the revenue curve R, where the point $(1-F(v_1), R(1-F(v_1)))$ is both the global maximum and the rightmost among them. Between each consecutive pair of these points, R increases almost linearly from $1-F(v_i+\varepsilon)$ to $1-F(v_i)$, and then gradually decreases toward $R(1-F(v_{i-1}))$ as it approaches $1-F(v_{i-1})$.

Notice that

$$\frac{R(1-F(v_i))}{1-F(v_i)} = v_i$$

represents the slope of the line connecting the point $(1-F(v_i), R(1-F(v_i)))$ to the origin $(0,0)$. Since this slope decreases as the x-coordinate $1-F(v_i)$ increases, it follows that, when constructing the convex hull of R, one must connect these adjacent points in order of increasing i. Consequently, the convex hull consists of $K-1$ linear segments, implying that the ironed virtual value function of this distribution contains at least $K-1$ ironing intervals.

As established in the proof of Proposition 5.1, the optimal profit attainable under this distribution satisfies:

$$\pi^* = \Omega\left(\frac{K}{C}\right).$$

In the following, we show that any menu of two-part tariffs of size k can achieve a revenue of at most $\frac{O(k)}{K}$. By the same argument as before, the profit derived from the markup component is at most $O\left(\frac{1}{C}\right)$.

Now consider each entry $j \in [k]$ in the menu. Let S_j denote the set of buyer values that select entry j, and let T_j denote the corresponding fixed fee. Observe that, in the absence of a fixed fee, the buyer's maximum achievable utility is upper bounded by:

$$W(v_i) = \left(1 - \frac{1}{K}\right)\left(\frac{1}{\varepsilon}\right)^{K\cdot i}.$$

For each buyer value $v_i \in S_j$, if the buyer chooses entry j, his utility without the fixed fee must exceed T_j in order to satisfy the incentive compatibility constraint. Let $i_j = \min S_j$ be the smallest index in the set S_j. It then follows that

$$T_j \le \left(1 - \frac{1}{K}\right)\left(\frac{1}{\varepsilon}\right)^{K\cdot i_j}.$$

Therefore, the total profit from the fixed fee corresponding to entry j is upper bounded by:

$$\begin{aligned}
\sum_{i\in S_j} \Pr_{v\sim F}[v \in [v_i, v_i+\varepsilon]]\, T_j &\le \sum_{i\in S_j} \Pr_{v\sim F}[v \in [v_i, v_i+\varepsilon]] \left(1 - \frac{1}{K}\right)\left(\frac{1}{\varepsilon}\right)^{K\cdot i_j} \\
&= \left(1 - \frac{1}{K}\right) \sum_{i\in S_j} \frac{\varepsilon^{K\cdot i}}{C}\left(\frac{1}{\varepsilon}\right)^{K\cdot i_j} \\
&= \left(1 - \frac{1}{K}\right) \frac{1 + O(\varepsilon^K)}{C} \\
&= \frac{O(1)}{C}.
\end{aligned}$$

The second equality follows from the fact that the indices in each set S_j are integers and mutually disjoint, with i_j being the smallest among them. Summing over all entries in the menu, we obtain that the total profit of any mechanism with menu size s is upper bounded by $\frac{O(s)}{C}$, whereas the optimal profit is at least $\frac{\Omega(K)}{C}$. This completes the proof.

□

E.2 Positive results: Proof of Theorem 2

In the following, we present the proof of Theorem 2. We start with the following lemma that will be used in the proof.

Lemma E.1. *Given any distribution F, define $\tilde{\varphi}(v)$ to be the ironed virtual value. Then:*

$$\tilde{\varphi}(v) \le v.$$

Proof. Consider the revenue curve $R(q)$. For any value v, let its corresponding (inverse) quantile be $q = 1 - F(v)$, and conversely, let $v(q)$ denote the value associated with quantile q. The revenue from posting a price $v(q)$ is then given by

$$R(q) = q \cdot v(q) = v(1 - F(v)).$$

Define $\tilde{R}$ to be the convex hull of function R. It is clear that

$$\tilde{R}'(q) = \varphi(v(q)).$$

For any x, let $q = 1 - F(x)$ be its corresponding (inverse) quantile. If x does not belong to any iron interval, it follows that

$$\tilde{\varphi}(x) = \varphi(x) = x - \frac{1 - F(x)}{f(x)} \leq x.$$

Otherwise, suppose q is in the ironed interval $[q_1, q_2]$. It follows that:

$$\begin{aligned}
\tilde{\varphi}(x) &= \tilde{R}'(q) \\
&= \frac{\tilde{R}(q_2) - \tilde{R}(q_1)}{q_2 - q_1} \\
&= \frac{R(q_2) - R(q_1)}{q_2 - q_1} \\
&= \frac{q_2 \cdot v(q_2) - q_1 \cdot v(q_1)}{q_2 - q_1} \\
&\leq \frac{q_2 \cdot x - q_1 \cdot x}{q_2 - q_1} = x.
\end{aligned}$$

where the inequality follows form the fact that $v(q)$ is decreasing in the interval $[q_1, q_2]$ and $v(q) = x$.

□

We are now ready to prove Theorem 2.[8]

Proof of Theorem 2. Fix a markup $\beta > 0$, and write

$$\tau := 1 + \beta.$$

Let

$$r_0 := \inf\{v : \tilde{\varphi}(v) \geq 0\}.$$

We only need to consider the region $v \geq r_0$, where the ironed virtual value is non-negative. Let

$$\{[\underline{\gamma}_h, \overline{\gamma}_h]\}_{h=1}^{s}$$

be the ironing intervals of $\tilde{\varphi}$. Discard the ironing intervals that lie entirely below r_0. On the remaining

[8] Again, if only the lower bound on the elasticity is available, we set $M := +\infty$ and $M/(M+1) := 1$.

region, form consecutive blocks by pairing each remaining ironing interval with the non-ironed segment immediately before it, and by adding a final non-ironed tail if one is present. Thus we obtain intervals

$$[r_{i-1}, r_i], \qquad i = 1, \dots, N,$$

with $N \le s+1$. For each block, choose $\ell_i \in [r_{i-1}, r_i] \cup \{\infty\}$ so that $[r_{i-1}, \ell_i]$ is non-ironed, while, if $\ell_i < r_i$, the terminal segment $[\ell_i, r_i]$ is one of the remaining ironing intervals. If the final block is a non-ironed tail, we set $r_N = \ell_N = \infty$. Define

$$a_i(v) := \min\{v, \ell_i\},$$

with the convention that $\min\{v, \infty\} = v$. Thus $v \mapsto q(a_i(v))$ is increasing on the non-ironed part of block i and constant on the ironing interval, if such an interval is present.

The optimal profit benchmark. By the ironed-virtual-surplus characterization of the optimal mechanism,

$$\pi^* = \int_{r_0}^{\infty} f(v)\, W(\tilde{\varphi}(v))\, dv.$$

The cost-share statement in Lemma 4.2 gives

$$C(q(z)) \ge \frac{1}{M+1} z q(z), \qquad W(z) = zq(z) - C(q(z)) \le \frac{M}{M+1} z q(z).$$

Therefore,

$$\pi^* \le \frac{M}{M+1} \int_{r_0}^{\infty} f(v)\, \tilde{\varphi}(v)\, q(\tilde{\varphi}(v))\, dv.$$

For $v \in [r_{i-1}, \ell_i]$, we have $\tilde{\varphi}(v) = \varphi(v) \le v = a_i(v)$. When $\ell_i < r_i$ and $v \in [\ell_i, r_i]$, the ironed virtual value is constant on the ironing interval, and Lemma E.1 gives $\tilde{\varphi}(v) = \tilde{\varphi}(\ell_i) \le \ell_i = a_i(v)$. Since q is monotone,

$$\pi^* \le \frac{M}{M+1} \sum_{i=1}^{N} \int_{r_{i-1}}^{r_i} f(v)\, \tilde{\varphi}(v)\, q(a_i(v))\, dv. \tag{36}$$

Define

$$B := \sum_{i=1}^{N} \int_{r_{i-1}}^{r_i} f(v)\, \tilde{\varphi}(v)\, q(a_i(v))\, dv.$$

Then (36) says

$$\pi^* \le \frac{M}{M+1} B. \tag{37}$$

We are only left to approximate B.

The menu mechanism and its implemented allocation. We now construct the menu and identify the allocation it induces. For each block i, create one menu entry (T_i, β, L_i). All entries use the same markup β; the capacity of entry i is

$$L_i := q\left(\frac{\ell_i}{\tau}\right),$$

with the convention that $L_i = \infty$ when $\ell_i = \infty$. Its fixed fee is

$$T_i := \tau W\left(\frac{r_{i-1}}{\tau}\right) - \sum_{j=1}^{i-1} \int_{r_{j-1}}^{r_j} q\left(\frac{a_j(t)}{\tau}\right) dt. \tag{38}$$

This menu has $N \le s+1$ entries. The fee is non-negative because, by Lemma 4.1,

$$\begin{aligned} T_i &= \int_0^{r_{i-1}} q\left(\frac{t}{\tau}\right) dt - \sum_{j=1}^{i-1} \int_{r_{j-1}}^{r_j} q\left(\frac{a_j(t)}{\tau}\right) dt \\ &= \int_0^{r_0} q\left(\frac{t}{\tau}\right) dt + \sum_{j=1}^{i-1} \int_{r_{j-1}}^{r_j} \left(q\left(\frac{t}{\tau}\right) - q\left(\frac{a_j(t)}{\tau}\right)\right) dt \ge 0. \end{aligned}$$

The first equality uses $\tau W(r_{i-1}/\tau) = \int_0^{r_{i-1}} q(t/\tau)dt$, which follows from Lemma 4.1 after the change of variables $t = \tau u$. The final inequality follows because $a_j(t) = \min\{t, \ell_j\} \le t$ and q is monotone.

Ignoring the fixed fee T_i, a buyer of value v who selects entry i solves

$$\max_{0 \le q \le L_i} \{vq - \tau C(q)\},$$

The unconstrained maximizer is $q(v/\tau)$. Since $L_i = q(\ell_i/\tau)$ and q is increasing, the quantity purchased from entry i is

$$x_i(v) = \min\left\{q\left(\frac{v}{\tau}\right), q\left(\frac{\ell_i}{\tau}\right)\right\} = q\left(\frac{\min\{v, \ell_i\}}{\tau}\right) = q\left(\frac{a_i(v)}{\tau}\right).$$

Let $U_i(v)$ be the utility from choosing entry i. For type r_{i-1}, the capacity constraint of entry i does not bind, since $r_{i-1} \le \ell_i$. Hence the gross utility before paying the fixed fee is

$$\max_{q \ge 0}\{r_{i-1}q - \tau C(q)\} = \tau W\left(\frac{r_{i-1}}{\tau}\right).$$

Recalling the definition of the fixed fee,

$$T_i = \tau W\left(\frac{r_{i-1}}{\tau}\right) - \sum_{j=1}^{i-1} \int_{r_{j-1}}^{r_j} q\left(\frac{a_j(t)}{\tau}\right) dt,$$

we have

$$\begin{aligned} U_i(r_{i-1}) &= \tau W\left(\frac{r_{i-1}}{\tau}\right) - T_i \\ &= \sum_{j=1}^{i-1} \int_{r_{j-1}}^{r_j} q\left(\frac{a_j(t)}{\tau}\right) dt. \end{aligned} \tag{39}$$

Moreover, within a fixed entry the fee and capacity do not vary with the type. Thus the envelope theorem gives

$$U_i'(v) = x_i(v) = q\left(\frac{a_i(v)}{\tau}\right)$$

wherever the derivative is defined. Hence, for any type v,

$$\begin{aligned}U_i(v) &= U_i(r_{i-1}) + \int_{r_{i-1}}^{v} U_i'(t)\,dt \\ &= \sum_{j=1}^{i-1} \int_{r_{j-1}}^{r_j} q\left(\frac{a_j(t)}{\tau}\right) dt + \int_{r_{i-1}}^{v} q\left(\frac{a_i(t)}{\tau}\right) dt.\end{aligned}$$

This lets us compare adjacent entries globally. For $i < N$, entry i gives type r_i utility

$$U_i(r_i) = \sum_{j=1}^{i} \int_{r_{j-1}}^{r_j} q\left(\frac{a_j(t)}{\tau}\right) dt.$$

Applying (39) to entry $i+1$, whose left endpoint is r_i gives

$$\begin{aligned}U_{i+1}(r_i) &= \tau W\left(\frac{r_i}{\tau}\right) - T_{i+1} \\ &= \sum_{j=1}^{i} \int_{r_{j-1}}^{r_j} q\left(\frac{a_j(t)}{\tau}\right) dt.\end{aligned}$$

Thus adjacent entries cross at the block boundary r_i.

Moreover, the block thresholds are ordered, so $\ell_i \le \ell_{i+1}$. Hence

$$a_i(v) = \min\{v, \ell_i\} \le \min\{v, \ell_{i+1}\} = a_{i+1}(v),$$

and the monotonicity of q implies $x_i(v) \le x_{i+1}(v)$ for all v. Therefore $U_{i+1}(v) - U_i(v)$ is increasing in v. Since this difference is zero at r_i, types below r_i weakly prefer entry i to entry $i+1$, while types above r_i weakly prefer entry $i+1$ to entry i.

Iterating this adjacent-crossing argument, every type $v \in [r_{i-1}, r_i]$ weakly prefers entry i to all other entries. For types below r_0, the same comparison shows that entry 1 weakly dominates all later entries. Moreover, $U_1(r_0) = 0$, and for $v < r_0$,

$$U_1(v) = \int_{r_0}^{v} q\left(\frac{a_1(t)}{\tau}\right) dt \le 0.$$

Thus types below r_0 opt out. Overall, the menu implements the allocation

$$x^{\beta}(v) = q\left(\frac{a_i(v)}{\tau}\right) \qquad \text{for } v \in [r_{i-1}, r_i].$$

Expected profit from the menu mechanism. Let P_β denote the expected payment collected by the menu. Since the implemented allocation x^β is monotone and types below r_0 opt out, revenue equivalence gives

$$P_\beta = \int_{r_0}^{\infty} f(v)\varphi(v)\, x^{\beta}(v)\, dv.$$

Using the block representation of x^β, this becomes

$$P_\beta = \sum_{i=1}^{N} \int_{r_{i-1}}^{r_i} f(v)\,\varphi(v)\,q\!\left(\frac{a_i(v)}{\tau}\right) dv.$$

We now replace the virtual value by the ironed virtual value. On each non-ironed segment, $\varphi(v) = \tilde{\varphi}(v)$. On an ironing interval $[\ell_i, r_i]$, notice that $a_i(v) = \ell_i$, so the implemented allocation is constant:

$$q\!\left(\frac{a_i(v)}{\tau}\right) = q\!\left(\frac{\ell_i}{\tau}\right).$$

Hence

$$\begin{aligned}&\int_{\ell_i}^{r_i} f(v)\varphi(v)\,q\!\left(\frac{a_i(v)}{\tau}\right) dv\\ &\quad = q\!\left(\frac{\ell_i}{\tau}\right) \int_{\ell_i}^{r_i} f(v)\,\varphi(v)\,dv.\end{aligned}$$

By the definition of the ironed virtual value, ironing preserves total virtual surplus on each ironing interval:

$$\int_{\ell_i}^{r_i} f(v)\,\varphi(v)\,dv = \int_{\ell_i}^{r_i} f(v)\,\tilde{\varphi}(v)\,dv.$$

Multiplying by the constant allocation on the interval gives

$$\int_{\ell_i}^{r_i} f(v)\,\varphi(v)\,q\!\left(\frac{a_i(v)}{\tau}\right) dv = \int_{\ell_i}^{r_i} f(v)\,\tilde{\varphi}(v)\,q\!\left(\frac{a_i(v)}{\tau}\right) dv.$$

Combining the non-ironed and ironed parts of all blocks yields

$$P_\beta = \sum_{i=1}^{N} \int_{r_{i-1}}^{r_i} f(v)\,\tilde{\varphi}(v)\,q\!\left(\frac{a_i(v)}{\tau}\right) dv. \tag{40}$$

The general scaling statement in Lemma 4.2 gives

$$q\!\left(\frac{z}{\tau}\right) \geq \tau^{-1/m} q(z) \qquad \forall z \geq 0.$$

Using this inequality and $\tilde{\varphi}(v) \geq 0$ on the active region,

$$P_\beta \geq \tau^{-1/m} \sum_{i=1}^{N} \int_{r_{i-1}}^{r_i} f(v)\,\tilde{\varphi}(v)\,q(a_i(v))\,dv = \tau^{-1/m} B. \tag{41}$$

We next compare, pointwise in the buyer's type, the payment collected by the menu and the profit kept by the seller. Fix $v \in [r_{i-1}, r_i]$, and write

$$x_i(v) = q\!\left(\frac{a_i(v)}{\tau}\right)$$

for the quantity purchased from entry i. The payment collected from this type is

$$p_i(v) := T_i + \tau C(x_i(v)),$$

because the buyer pays the fixed fee and then pays markup $\tau = 1+\beta$ on production cost. The seller's profit from the same type is this payment minus the production cost:

$$\pi_i(v) := p_i(v) - C(x_i(v)) = T_i + \beta C(x_i(v)).$$

We claim that $\pi_i(v) \geq (\beta/\tau)p_i(v)$ pointwise. Indeed,

$$\begin{aligned}\pi_i(v) - \frac{\beta}{\tau}p_i(v) &= T_i + \beta C(x_i(v)) - \frac{\beta}{\tau}\big(T_i + \tau C(x_i(v))\big)\\ &= \left(1-\frac{\beta}{\tau}\right)T_i = \frac{T_i}{\tau} \geq 0,\end{aligned}$$

where the inequality uses $T_i \geq 0$.

Integrating this pointwise inequality over the implemented allocation gives

$$\text{Profit} = \sum_{i=1}^{N}\int_{r_{i-1}}^{r_i} f(v)\pi_i(v)\,dv \geq \frac{\beta}{\tau}\sum_{i=1}^{N}\int_{r_{i-1}}^{r_i} f(v)\,p_i(v)\,dv = \frac{\beta}{\tau}P_\beta. \tag{42}$$

Combining (37), (41), and (42), we conclude that

$$\text{Profit} \geq \frac{\beta}{\tau}\cdot\tau^{-1/m}B \geq \frac{M+1}{M}\cdot\frac{\beta}{(1+\beta)^{1+1/m}}\cdot\pi^*.$$

This proves the one-parameter guarantee.

Finally, the function

$$g(\beta) = \frac{\beta}{(1+\beta)^{1+1/m}}$$

is maximized at $\beta = m$. This gives

$$\text{Profit} \geq \frac{M+1}{M}\cdot\frac{m}{(m+1)^{1+1/m}}\cdot\pi^*,$$

as needed. □

F Theorem 4 Proof

Proof. Define $T_+ := T \cup \{\varnothing\}$ and extend x, p, f to T_+ by $f(\varnothing) = p(\varnothing) = x_i(\varnothing) = 0$ for $1 \le i \le n$. The seller aims to solve the following maximization problem over x, p:[9]

$$\begin{aligned}
\max \quad & \sum_{v \in T} f(v) \cdot \left(p(v) - \sum_{i=1}^{n} C(x_i(v)) \right) \\
\text{s.t.} \quad & \sum_{i=1}^{n} v^i x_i(v) - p(v) \ge \sum_{i=1}^{n} v^i x_i(v') - p(v') \quad \forall v \in T, v' \in T_+ \\
& x_i(v) \ge 0 \quad \forall 1 \le i \le n, v \in T
\end{aligned}$$

Similarly to Cai et al. (2016), we write the partial Lagrangian corresponding to the seller's optimization problem:

$$\begin{aligned}
L(\lambda, x, p) := & \sum_{v \in T} f(v) \cdot \left(p(v) - \sum_{i=1}^{n} C(x_i(v)) \right) \\
& + \sum_{v \in T, v' \in T_+} \lambda(v, v') \cdot \left[\sum_{i=1}^{n} v^i x_i(v) - p(v) - \left(\sum_{i=1}^{n} v^i x_i(v') - p(v') \right) \right] \\
= & \sum_{v \in T} p(v) \cdot \left[f(v) + \sum_{v' \in T} \lambda(v', v) - \sum_{v' \in T_+} \lambda(v, v') \right] \\
& + \sum_{i=1}^{n} \sum_{v \in T} x_i(v) \cdot \left[\sum_{v' \in T_+} v^i \lambda(v, v') - \sum_{v' \in T} {v'}^i \lambda(v', v) \right] \\
& - \sum_{i=1}^{n} \sum_{v \in T} f(v) C(x_i(v)).
\end{aligned}$$

The dual of the seller's maximization problem is as follows:

$$\begin{aligned}
& \min_{\lambda} \max_{x \ge 0, p} L(\lambda, x, p) \\
& \text{s.t.} \quad \lambda(v, v') \ge 0 \quad \forall v \in T, v' \in T_+.
\end{aligned}$$

Similarly to Cai et al. (2016), we call a feasible solution λ of the dual problem *useful* if $\max_{x \ge 0, p} L(\lambda, x, p) < \infty$. Note that each λ that is not useful can be ignored when solving the dual minimization problem. From Cai et al. (2016), we have the following.

Lemma F.1. *A solution λ of the dual problem is useful if and only if we have for every $v \in T$: $f(v) + \sum_{v' \in T} \lambda(v', v) - \sum_{v' \in T_+} \lambda(v, v') = 0$.*

In other words, the useful solutions exactly correspond to valid flows in a network with a node per type $v \in T$, sink $\varnothing$ and source s s.t. there is a flow of weight $f(v)$ from s to every $v \in T$ and a flow of weight $\lambda(v, v')$ from every $v \in T$ to every $v' \in T_+$.

[9] Allowing $v' \in T_+$ simultaneously captures incentive compatibility and individual rationality.

For every $1 \le i \le n$, define the *virtual value* for item i by $\Phi_i(v) := v^i - \frac{1}{f(v)} \sum_{v' \in T} \lambda(v', v) \cdot \left(v'^i - v^i\right)$. The proof of the following lemma is similar to the proof of Theorem 2 from Cai et al. (2016). For completeness, we also present it here.

Lemma F.2. *For a useful solution λ of the dual problem and a truthful mechanism $M = (x, p)$, it holds that:*

$$\sum_{v \in T} f(v) p(v) \le \sum_{i=1}^{n} \sum_{v \in T} f(v) x_i(v) \Phi_i(v).$$

Moreover, for a profit-maximizing mechanism $M^* = (x^*, p^*)$ and an optimal solution λ^* of the dual problem, equality holds.

That is, the expected virtual welfare is always at least as high as the expected seller's income from the buyer's payments, and for the optimal M^* and λ^* we have an equality.

Proof of Lemma F.2. Let λ be a useful dual solution. We have by Lemma F.1:

$$\begin{aligned}
L(\lambda, x, p) &= \sum_{v \in T} p(v) \cdot \left[f(v) + \sum_{v' \in T} \lambda(v', v) - \sum_{v' \in T_+} \lambda(v, v') \right] \\
&\quad + \sum_{i=1}^{n} \sum_{v \in T} x_i(v) \cdot \left[\sum_{v' \in T_+} v^i \lambda(v, v') - \sum_{v' \in T} v'^i \lambda(v', v) \right] - \sum_{i=1}^{n} \sum_{v \in T} f(v) C(x_i(v)) \\
&= \sum_{i=1}^{n} \sum_{v \in T} x_i(v) \cdot \left[v^i \left(\sum_{v' \in T} \lambda(v', v) + f(v) \right) - \sum_{v' \in T} v'^i \lambda(v', v) \right] - \sum_{i=1}^{n} \sum_{v \in T} f(v) C(x_i(v)) \\
&= \sum_{i=1}^{n} \sum_{v \in T} x_i(v) \left[f(v) v^i - \sum_{v' \in T} \lambda(v', v) \cdot \left(v'^i - v^i \right) \right] - \sum_{i=1}^{n} \sum_{v \in T} f(v) C(x_i(v)) \\
&= \sum_{i=1}^{n} \sum_{v \in T} f(v) x_i(v) \Phi_i(v) - \sum_{i=1}^{n} \sum_{v \in T} f(v) C(x_i(v)).
\end{aligned}$$

That is, the partial Lagrangian equals the expected virtual welfare minus the expected production cost. By the condition $\lambda \ge 0$ and the incentive compatibility constraints, we immediately deduce that $L(\lambda, x, p) \ge \sum_{v \in T} f(v) \cdot (p(v) - \sum_{i=1}^{n} C(x_i(v)))$. By adding the expected production cost to both sides of the inequality, we exactly get the first part of our result. The second part follows from strong duality. $\square$

For $1 \le i \le n$, let $\tilde{\varphi}_i$ be the ironed virtual value for item i; let r_i be s.t. $W(r_i) = \sum_{v \in T : \tilde{\varphi}_i(v^i) > 0} f_i\left(v^i\right) W\left(\tilde{\varphi}_i\left(v^i\right)\right)$. Let $r \ge 0$ be s.t. $W(r) = W(r_1) + \ldots + W(r_n)$. Note that since W is continuous, increasing and unbounded from above, such $r_1, \ldots, r_n, r$ exist.

Lemma F.3. *There exists a useful λ s.t. for every $1 \le i \le n$,*

$$\sum_{v \in T} f(v)(x_i(v)\Phi_i(v) - C(x_i(v))) \le \sum_{v^i \in Supp(F_i)} f_i\left(v^i\right) \left(W\left(v^i\right) \Pr_{\tilde{v}_{-i} \sim F_{-i}} [\exists j \ne i : \tilde{v}^j > v^i] + W\left(\tilde{\varphi}_i\left(v^i\right)\right) \right).$$

Proof of Lemma F.3. Fix $1 \le i \le n$. We get from Cai et al. (2016) that there exists a useful λ s.t. for every $v \in T$ with $v^i = \max_j v^j$, we have: $\Phi_i(v) \le \tilde{\varphi}_i\left(v^i\right)$ and $\Phi_j(v) = v^j$ for every $j \ne i$.

Fix $v \in T$. If $v^i = \max_j v^j$, we get:

$$x_i(v)\Phi_i(v) - C(x_i(v)) \leq x_i(v)\tilde{\varphi}_i\left(v^i\right) - C(x_i(v)) \leq W\left(\tilde{\varphi}_i\left(v^i\right)\right),$$

where the last transition is by definition of W. If $v^i < \max_j v^j$, we get:

$$x_i(v)\Phi_i(v) - C(x_i(v)) = x_i(v)v^i - C(x_i(v)) \leq W\left(v^i\right).$$

Taking the expectation over $v \sim F$ completes the proof. $\square$

From Lemma F.3 and the definition of SREV, we get:

$$\begin{aligned}
&\sum_{i=1}^{n}\sum_{v\in T} f(v)x_i(v)\Phi_i(v) - \sum_{i=1}^{n}\sum_{v\in T} f(v)C(x_i(v)) \\
&\leq \sum_{i=1}^{n}\sum_{v^i\in \mathrm{Supp}(F_i)} f_i\left(v^i\right)\left[W\left(v^i\right)\Pr_{v_{-i}\sim F_{-i}}[\exists j\neq i : v^j > v^i] + W\left(\tilde{\varphi}_i\left(v^i\right)\right)\right] \\
&\leq \sum_{i=1}^{n}\sum_{v^i\in \mathrm{Supp}(F_i)} f_i\left(v^i\right)W\left(v^i\right)\Pr_{v_{-i}\sim F_{-i}}[\exists j\neq i : v^j > v^i] + \sum_{i=1}^{n}\sum_{v^i:\tilde{\varphi}_i(v^i)>0} f_i\left(v^i\right)W\left(\tilde{\varphi}_i\left(v^i\right)\right) \\
&= \sum_{i=1}^{n}\sum_{v^i\in \mathrm{Supp}(F_i)} f_i\left(v^i\right)W\left(v^i\right)\Pr_{v_{-i}\sim F_{-i}}[\exists j\neq i : v^j > v^i] + \mathrm{SREV} \\
&\leq \sum_{i=1}^{n}\sum_{v^i\in \mathrm{Supp}(F_i):v^i>r} f_i\left(v^i\right)W\left(v^i\right)\Pr_{v_{-i}\sim F_{-i}}[\exists j\neq i : v^j > v^i] + \sum_{i=1}^{n}\sum_{v^i\in \mathrm{Supp}(F_i):v^i\leq r} f_i\left(v^i\right)W\left(v^i\right) + \mathrm{SREV}.
\end{aligned}$$

Set:

$$\begin{aligned}
\mathrm{Core} &:= \sum_{i=1}^{n}\sum_{v^i\in \mathrm{Supp}(F_i):v^i\leq r} f_i\left(v^i\right)W\left(v^i\right), \\
\mathrm{Tail} &:= \sum_{i=1}^{n}\sum_{v^i\in \mathrm{Supp}(F_i):v^i>r} f_i\left(v^i\right)W\left(v^i\right)\Pr_{v_{-i}\sim F_{-i}}[\exists j\neq i : v^j > v^i].
\end{aligned}$$

Bounding the tail. Fix some $1 \leq i \leq n$. Note that if $v^j > v^i$ for some j, then selling the product lines separately for the price of v^i per unit when exactly $q\left(v^i\right)$ or 0 units of each product line must be purchased yields the profit of at least $W\left(v^i\right)$. On the other hand, the expected profit of this mechanism is at most the expected profit for optimally selling the product lines separately, which is $W(r)$. Thus,

$$\mathrm{Tail} = \sum_{i=1}^{n}\sum_{v^i\in \mathrm{Supp}(F_i):v^i>r} f_i\left(v^i\right)W\left(v^i\right)\Pr_{v_{-i}\sim F_{-i}}[\exists j\neq i : v^j > v^i] \leq \sum_{i=1}^{n}\sum_{v^i\in \mathrm{Supp}(F_i):v^i>r} f_i\left(v^i\right)W(r).$$

Note that the last expression is the profit for selling the product lines separately for price-per-unit of r when only $q(r)$ units of each product line can be purchased. Therefore, it is upper-bounded by the profit for optimally selling the product lines separately, namely $W(r)$. Hence, we get:

$$\text{Tail} \le W(r). \tag{43}$$

Bounding the core. For every $1 \le i \le n$, define the random variable $X_i := W\left(v^i\right) \cdot \mathbb{1}_{v^i \le r}$. That is, $X_i = W\left(v^i\right)$ if $v^i \le r$ and $X_i = 0$ otherwise. Define further $X := \sum_{i=1}^n X_i$. By linearity of expectation, we have $\mathbb{E}[X] = \text{Core}$. Let us bound the variance of X to prove a concentration result using Chebyshev's inequality. Since the X_is are independent random variables, we have:

$$\text{Var}[X] = \sum_{i=1}^n \text{Var}[X_i] = \sum_{i=1}^n \left(\mathbb{E}\left[X_i^2\right] - \mathbb{E}[X_i]^2\right) \le \sum_{i=1}^n \mathbb{E}\left[X_i^2\right]. \tag{44}$$

Fix $1 \le i \le n$. Note that for every t, selling $q(t)$ units of product line i for price-per-unit of t leads to expected profit of $\Pr\left[v^i \ge t\right] W(t)$, which is at most the optimal expected profit for product line i, namely $W(r_i)$. Thus:

$$\Pr\left[v^i \ge t\right] W(t) \le W(r_i). \tag{45}$$

We use a technique similar to Lemma 9 from Cai et al. (2016) to bound the second moment of X_i. Let $\text{Supp}(F_i) \cap [0, r] = \{a_1, \ldots, a_l\}$ with $a_0 := 0 \le a_1 < a_2 < \ldots < a_l \le r$. We get:

$$\begin{aligned}
\mathbb{E}\left[X_i^2\right] &= \sum_{j=0}^{l} \Pr_{v^i \sim F_i}[v^i = a_j] \cdot W(a_j)^2 = \sum_{j=1}^{l} \left(W(a_j)^2 - W(a_{j-1})^2\right) \cdot \sum_{h=j}^{l} \Pr_{v^i \sim F_i}[v^i = a_h] \\
&= \sum_{j=1}^{l} \left(W(a_j)^2 - W(a_{j-1})^2\right) \cdot \Pr_{v^i \sim F_i}[v^i \ge a_j] \le 2 \sum_{j=1}^{l} W(a_j)(W(a_j) - W(a_{j-1})) \Pr_{v^i \sim F_i}[v^i \ge a_j] \\
&\le_{(45)} 2W(r_i) \sum_{j=1}^{l} (W(a_j) - W(a_{j-1})) = 2W(r_i) W(a_l) \le 2W(r_i) W(r),
\end{aligned}$$

where the last transition is since W is increasing. Combining the above with Inequality 44 yields:

$$\text{Var}[X] \le 2W(r) \sum_{i=1}^n W(r_i) = 2W(r)^2, \tag{46}$$

where the last transition is by the definition of r. From Chebyshev's inequality and Inequality 46, we have:

$$\Pr_{v \sim F}\left[\sum_{i=1}^n W\left(v^i\right) < \text{Core} - 2W(r)\right] \le \Pr[X < \mathbb{E}[X] - 2W(r)] \le \frac{\text{Var}[X]}{4W(r)^2} \le \frac{1}{2}. \tag{47}$$

Consider the mechanism selling the bundle consisting of $q\left(v^i\right)$ units of product line i for every $1 \le i \le n$ for the price of $\sum_{i=1}^n C\left(q\left(v^i\right)\right) + \frac{\text{Med}\left(\sum_{i=1}^n W\left(v^i\right)\right)}{2}$. Let us condition on the event that $\sum_{i=1}^n W\left(v^i\right) \ge \text{Core} - 2W(r)$, which has probability ≥ 0.5 by Inequality 47. If for some realization of v the seller can extract at least half of the surplus, then she gets at least $\frac{\text{Core}-2W(r)}{2}$;[10] otherwise, the seller gets at least the fixed fee, which is at least $\frac{\text{Core}-2W(r)}{2}$. In either case, the expected seller's profit is at least $\frac{\text{Core}-2W(r)}{4}$.

[10]This case includes the scenario $q\left(v^1\right) = \ldots = q(v^n) = 0$ in which $\sum_{i=1}^n W\left(v^i\right) = 0 \ge \text{Core} - 2W(r)$.

Therefore: $\text{BREV} \geq \frac{\text{Core}-2W(r)}{4}$; i.e., $\text{Core} \leq 4\text{BREV} + 2W(r)$. Combining our bounds, we get:

$$\sum_{i=1}^{n}\sum_{v\in T} f(v)x_i(v)\Phi_i(v) - \sum_{i=1}^{n}\sum_{v\in T} f(v)C(x_i(v)) \leq \text{Tail} + \text{Core} + \text{SREV} \leq$$
$$\leq W(r) + (4\text{BREV} + 2W(r)) + \text{SREV} = 4\text{SREV} + 4\text{BREV}.$$

Together with Lemma F.2, our theorem follows. $\square$

G Omitted Proofs from Section 7

Recall the change of variables $\tilde{q} := U(q)$ and the induced cost function

$$\tilde{C}(\tilde{q}) := C\big(U^{-1}(\tilde{q})\big). \tag{48}$$

Lemma G.1. *Let $\tilde{C}$ be defined by Eq.* (48)*. Then for every $q > 0$,*

$$\eta_{\tilde{C}'}(U(q)) = \frac{\eta_{C'}(q) - \eta_{U'}(q)}{\eta_U(q)}, \tag{49}$$

where $\eta_U(q) := qU'(q)/U(q) \in (0,1]$ is the elasticity of U, and $\eta_{U'}(q) := qU''(q)/U'(q) \leq 0$ is the elasticity of U'.

Proof. Differentiating Eq. (48) once and twice gives $\tilde{C}'(\tilde{q})\,U'(q) = C'(q)$, so $\tilde{C}'(\tilde{q}) = C'(q)/U'(q)$, and

$$\tilde{C}''(\tilde{q}) = \frac{C''(q)\,U'(q) - C'(q)\,U''(q)}{U'(q)^3}.$$

Therefore

$$\begin{aligned}\eta_{\tilde{C}'}(U(q)) = \frac{\tilde{q}\,\tilde{C}''(\tilde{q})}{\tilde{C}'(\tilde{q})} &= \frac{U(q)\big[C''(q)\,U'(q) - C'(q)\,U''(q)\big]}{U'(q)^2\,C'(q)} \\ &= \frac{U(q)}{q\,U'(q)}\left[\frac{q\,C''(q)}{C'(q)} - \frac{q\,U''(q)}{U'(q)}\right] \\ &= \frac{\eta_{C'}(q) - \eta_{U'}(q)}{\eta_U(q)}.\end{aligned}$$

$\square$

Proof of Proposition 7.1. Because U is strictly increasing and concave, U^{-1} is strictly increasing and convex; since C is strictly increasing and convex, the composition $\tilde{C}$ is also strictly increasing and convex. Moreover, by Lemma G.1, $\eta_{\tilde{C}'} \geq m > 0$, hence $\tilde{C}'' > 0$ on $(0,\infty)$, which implies *strict* convexity of $\tilde{C}$. Moreover, note that $\tilde{C}(0) = 0$ and $\tilde{C} \in C^1([0,\infty)) \cap C^2((0,\infty))$. Furthermore, the assumption $\eta_{\tilde{C}'} \geq m$ implies through integrating that for a fixed $\tilde{q}_0 > 0$ and $\tilde{q} \to \infty$, $\tilde{C}'(\tilde{q}) \geq \tilde{C}'(\tilde{q}_0)(\tilde{q}/\tilde{q}_0)^m \to \infty$.

In $\tilde{q}$-coordinates, the buyer's gross value is the linear expression $v\,\tilde{q}$, and the two-part tariff $T + (1+\beta)C(q)$ becomes $T + (1+\beta)\tilde{C}(\tilde{q})$; the mechanism design problem is, therefore, the one studied in Section 2 with $\tilde{C}$ in place of C. since U is a continuous strictly increasing bijection of $[0,\infty)$ onto itself,

$q \mapsto \tilde{q}$ is a bijection between the two choice problems, so every two-part tariff, and every mechanism and its expected profit, coincide across the two formulations (in particular, π^* is unchanged, and the value distribution F – and hence its regularity – is untouched).

By Lemma G.1 and Eq. (9), the induced cost function satisfies $\eta_{\tilde{C}'}(\tilde{q}) \in [m, M]$ for every $\tilde{q} > 0$. Applying Theorem 1 to the linear-value environment with cost function $\tilde{C}$ completes the proof. □

Proof of Corollary 7.1. With $U(q) = q^\alpha$, $\eta_U \equiv \alpha$ and $\eta_{U'} \equiv \alpha - 1$. With $C(q) = q^k/k$, $\eta_{C'} \equiv k - 1$. Plugging into Eq. (49), $\eta_{\tilde{C}'} = ((k-1) - (\alpha - 1))/\alpha = (k-\alpha)/\alpha =: A - 1$, a constant, with $A = k/\alpha > 1$. From Proposition 7.1, we get an approximation guarantee of $A^{-1/(A-1)} = (k/\alpha)^{-\alpha/(k-\alpha)}$, as needed. □