Differentially Private Market Segmentation*

Chengcheng Wang †

November 5, 2026

Click here for the latest version

Abstract

Access to consumer data enables firms to segment markets and charge different prices to different consumers, but the use of such data raises privacy concerns. This paper studies monopoly price discrimination under differential privacy, which limits how much segmentation can reveal about consumers' valuations. We characterize the attainable combinations of consumer and producer surplus, which form a convex polygon when valuations are finite. Relative to unrestricted disclosure, differential privacy lowers the firm's maximum profit and raises minimum consumer surplus while reducing its maximum. We show that every segmentation that maximizes consumer surplus leaves the firm with its uniform monopoly profit. Under an upper-tail demand bound, minimum consumer surplus can also be attained at this profit level. Our analysis clarifies how privacy regulation shapes the benefits and harms of market segmentation and informs debates over the use of consumer data in pricing.

Keywords: Differential privacy; market segmentation; price discrimination; information design; welfare.

JEL classification: D42, D82, D83, L51.

*I am grateful to Nageeb Ali, Andreas Blume, Inga Deimen, Laurent Mathevet, Ina Taneva, and Matthijs Wildenbeest for their insightful comments and suggestions. I appreciate Yiwei Qu, Zexin Ye, and all participants at University of Arizona theory seminar, the 2026 Midwest Economic Theory Conference, and the 2026 Stony Brook International Conference on Game Theory for helpful discussions and feedback. ChatGPT and Claude provided excellent research assistance. All errors are my own.

†Department of Economics, University of Arizona. Email: cwang111@arizona.edu

1. Introduction

Firms use consumer data to decide which customers to charge more and which to offer a discount. Some consumers end up paying more, while others receive lower prices or buy a product they would otherwise have forgone. Shiller (2020) finds that web-browsing histories substantially increase the potential profit gains from price discrimination relative to demographics alone. In a field experiment, Dubé and Misra (2023) find that personalization raises expected profits by 19% relative to optimized uniform pricing. More than 60% of consumers benefit, yet aggregate consumer surplus falls. These findings suggest that the welfare consequences of consumer information depend on how it changes the prices offered to different consumers.

Privacy regulation constrains the information on which these prices can depend. The European Union's General Data Protection Regulation (GDPR) and the California Privacy Rights Act (CPRA) are prominent efforts to limit businesses' use of personal information and strengthen consumers' control over their data.[1] Restricting what a firm learns can prevent it from identifying consumers who would pay a higher price. It can also prevent the firm from identifying groups to whom a discount would be profitable. How much can consumers gain or lose from price discrimination when their privacy are protected? Can privacy regulation allow consumers to benefit from price discrimination while protecting them from its harms?

We study these questions in a monopoly market in which disclosure must satisfy differential privacy (Dwork et al., 2006). An information designer observes each consumer's valuation and commits to a disclosure rule that sends a report to the firm, which then sets a price for each report. Differential privacy limits how much these reports can reveal about individual valuations: the probability of any given report can differ across valuations by at most a factor of e^{ϵ}.[2] We characterize the consumer and producer surplus attainable under all disclosure rules that meet a given level of privacy regulation; because reports that induce the same price can be pooled,

[1] Article 5(1)(b)–(c) of Regulation (EU) 2016/679 (General Data Protection Regulation) establishes purpose limitation and data minimization. The CCPA, as amended by the CPRA, includes rights to correct inaccurate personal information and to limit the use and disclosure of sensitive personal information; see the California Attorney General's explanation.

[2] Applications of differential privacy include consumer-data collection by Google, Apple, and Microsoft (Erlingsson, Pihur, and Korolova, 2014; Apple, 2017; Ding, Kulkarni, and Yekhanin, 2017), as well as the publication of tabulations from the 2020 U.S. Census (Abowd et al., 2022).

these surplus pairs form a polygon. The segmentation that is best for consumers never benefits the firm, even though privacy typically makes it inefficient. Under a bound on upper-tail demand, the worst outcome for consumers can also be reached without raising the firm's profit. Whether privacy regulation separates helpful from harmful price discrimination depends on the order in which relaxing privacy regulation makes discounts and price increases available.

For the consumer optimum, the answer matches the unrestricted benchmark, but for a different reason. Without privacy regulation, Bergemann, Brooks, and Morris (2015) (BBM hereafter) show that consumer-optimal segmentation leaves the firm with exactly its uniform monopoly profit. In BBM, this is a consequence of efficiency: the segmentation is efficient, so consumers receive all surplus above the profit the firm can secure by ignoring information. Privacy regulation breaks this argument. Every report must remain possible under every valuation, so every segment contains every valuation, and unless the uniform monopoly price is the lowest valuation, no private segmentation achieves efficient trade. A segmentation that lets the firm gain could then, in principle, expand trade enough to leave consumers better off. We show that it cannot: at every privacy level, every segmentation that maximizes consumer surplus leaves the firm with exactly its uniform monopoly profit.

The firm's pricing incentives explain why. A discount can strictly beat the uniform monopoly price only in a segment with more consumers who buy only because of the discount, per consumer who would buy anyway, than in the aggregate market; otherwise the uniform monopoly price would still be more profitable. Another segment must then have relatively fewer of them. At a consumer optimum, exchanging a small mass of consumers between the two segments, those who would buy anyway into the discounted segment and those who buy only because of the discount out, keeps every price optimal and never lowers consumer surplus. Privacy regulation never blocks this exchange, because it makes the two segments more alike: within each, the affected likelihood ratios move toward each other and stay within the segment's original range. The exchange either raises consumer surplus or leaves both surpluses unchanged while balancing the segments, so the most balanced consumer optimum at any profit level contains no profitable discount. Every consumer optimum therefore leaves the firm exactly its uniform monopoly profit, and any segmentation that raises profit leaves consumers strictly short of their

optimum.

The worst segmentation for consumers behaves differently. Raising some consumers' prices while keeping the firm indifferent requires enough excluded consumers in each high-price segment: consumers who would buy at the uniform monopoly price but not at the higher one. With several higher prices, such consumers can make an intermediate price more profitable than the higher one, and privacy regulation requires every other segment to retain some of them. Under an upper-tail demand bound, which limits the revenue that high-valuation demand can support relative to uniform monopoly profit, some segmentation still minimizes consumer surplus without raising profit, so consumers lose exactly the value of trades that no longer take place. The bound holds automatically with at most three valuations. Without it, both obstacles can bind at once, and minimizing consumer surplus can require rewarding the firm.

The level of privacy regulation determines which price changes segmentation can support. A price's activation threshold is the smallest privacy budget at which some private report makes that price optimal; a linear program computes it. Below every threshold, every segment leaves the uniform monopoly price uniquely optimal, so segmentation has no effect. Discounts and price increases have their own thresholds, and the attainable range of consumer surplus jumps once the relevant one is reached: a tie suffices for the firm to charge the new price, so a positive mass of consumers can already face it while the firm earns only its uniform monopoly profit. When discounts activate first, every budget between the two thresholds permits segmentation that raises consumer surplus but none that lowers it. No price above the uniform monopoly price is then optimal after any report, so every private segmentation in this range weakly benefits the firm and every consumer. When price increases activate first, the same range admits only segmentation that weakly harms every consumer.

The binary model shows these forces in their simplest form. With a single alternative price, segmentation can only help consumers when the uniform monopoly price is high and only hurt them when it is low. With two valuations, privacy regulation also changes the efficiency of price discrimination. Without privacy regulation, the extreme market segmentations for the firm

and for consumers are both efficient; with privacy regulation, the firm's preferred segmentation excludes the low-valuation consumers that every high-price segment must contain, while consumer-optimal segmentation still maximizes total surplus among private disclosures.

Related literature. Our analysis builds on information design (Kamenica and Gentzkow, 2011; Bergemann and Morris, 2019; Taneva, 2019) and the market-segmentation framework of Bergemann, Brooks, and Morris (2015). They characterize the welfare triangle attainable through unrestricted segmentation of a single-product monopoly. Subsequent work extends this analysis to multiple products (Haghpanah and Siegel, 2022) and competition (Elliott et al., 2026; Rhodes and Zhou, 2024). We keep the single-product monopoly setting and restrict what the firm can learn about consumers through segmentation instead. Privacy rules out the segments that omit some valuations, on which BBM's efficient extremes rely; once efficiency is out of reach, the consumer optimum still leaves the firm its uniform monopoly profit, and under the upper-tail demand bound so can the worst outcome for consumers.

The closest paper is Fallah et al. (2024), who also introduce privacy into the environment of BBM. They model privacy guarantee using a market-level randomized response mechanism: with some probability, the mechanism outputs a market drawn uniformly from the set of all markets instead of the true market. In our paper, the designer can choose any information structure satisfying differential privacy, without imposing a particular form of noise. Because reports never distort the market itself, the firm can always secure its uniform monopoly profit by ignoring them, so every welfare effect we identify comes from what reports reveal about individual consumers.

A related literature studies how the incentives of consumers and intermediaries determine the information that reaches firms. Taylor (2004) examines the sale of customer information; Ichihashi (2020), Hidir and Vellodi (2021), and Ali, Lewis, and Vasserman (2023) give consumers control over disclosure; and Bonatti and Cisternas (2020) study the scores intermediaries construct from consumer histories. We instead let a committed designer choose any disclosure that satisfies an exogenous privacy constraint, which separates what the guarantee permits from the

incentives that select a particular disclosure.[3]

Our formulation of the privacy constraint is based on Schmutte and Yoder (2024), who express differential privacy as a support restriction on Bayes-plausible posteriors for statistical publication. Pan et al. (2025) study Bayesian persuasion under several privacy notions. We use the posterior approach to characterize consumer and producer surplus jointly. The all-pairs likelihood-ratio constraint generates two-level extreme points related to the staircase mechanisms of Kairouz, Oh, and Viswanath (2016).[4]

Sections 2 and 3 establish the settings and the uniform monopoly price benchmarks. Section 4 illustrates the pricing and welfare effects in the binary model. Section 5 characterizes the welfare polygon, consumer-surplus extrema, and the activation thresholds for finitely many valuations. Section 6 concludes. Appendix A contains the supporting results and proofs.

2. Model

2.1 Environment and Market Segmentation

A monopolistic firm produces a single product at zero marginal cost. Information about consumers allows the firm to offer different prices to different groups; the privacy constraint will limit how those groups can be formed. There is a continuum of consumers with total mass one; each demands at most one unit and knows her own valuation. Valuations lie in the finite set

$$V = \{v_1, v_2, \ldots, v_K\}, \qquad K \geq 2,$$

where

$$0 < v_1 < v_2 < \ldots < v_K < \infty.$$

[3] Work on differentially private dynamic pricing studies the revenue cost of privacy for learning algorithms (e.g., Chen, Miao, and Wang, 2023). Acquisti, Taylor, and Wagman (2016) survey the broader economics of privacy.

[4] Other approaches include the social choice treatment of accuracy and privacy in Abowd and Schmutte (2019), the posterior-based privacy measure in Eilat, Eliaz, and Mu (2021), and the exact non-disclosure of designated events in Strack and Yang (2024).

A consumer who buys at price p obtains utility $v - p$; one who does not buy obtains zero.[5]

A *market* is a probability distribution x over V, with x_k the share of consumers whose valuation is v_k. The set of all markets is

$$X := \Delta(V) = \left\{ x \in \mathbb{R}_+^K : \sum_{k=1}^{K} x_k = 1 \right\}.$$

An exogenously given aggregate market x^* has full support, $x_k^* > 0$ for every k, so every valuation occurs before any information is disclosed. This also keeps the posterior-to-prior likelihood ratios used below well-defined.

A *segmentation* σ divides the aggregate market into finitely many submarkets. Formally, σ is a finite-support probability distribution over X, with $\sigma(x)$ the share of consumers assigned to market x, and the feasible segmentations are those that average back to the aggregate market:

$$\Sigma := \left\{ \sigma \in \Delta(X) : \sum_{x \in \mathrm{supp}(\sigma)} \sigma(x) \cdot x = x^*, |\mathrm{supp}(\sigma)| < \infty \right\}.$$

2.2 Privacy-Constrained Information

An information designer observes each consumer's valuation and commits to an information structure (M, ψ): a finite message space M, consistent with the finite-support segmentations in Σ, and a stochastic disclosure rule

$$\psi : V \to \Delta(M)$$

that maps each valuation to a distribution over messages. Conditional on valuation v_k, a message is drawn from $\psi(\cdot | v_k)$ and disclosed to the firm. Without loss of generality, every message is sent with positive probability under some valuation.[6]

The disclosure must protect the consumer's valuation. For an exogenously given privacy budget

[5] The setting can equivalently be formulated as a single-consumer problem in which the consumer's valuation is uncertain.

[6] A message that is sent with probability zero under every valuation has no effect on posterior beliefs or market outcomes and can therefore be deleted.

$\epsilon \geq 0$, the information structure must satisfy *ϵ-differential privacy*.

Definition 1 (Differential privacy (Dwork et al., 2006)). *An information structure is* ***ϵ-differentially private (ϵ-DP)*** *if for every $m \in M$ and every pair $v_k, v_{k'} \in V$ with $k \neq k'$,*

$$\left| \log \frac{\psi(m \mid v_{k'})}{\psi(m \mid v_k)} \right| \leq \epsilon.$$ [7]

Definition 1 requires every message to remain possible under every valuation. A message may make a high valuation more likely, but it cannot identify a consumer as a high type with certainty. The likelihood bound can be written as

$$\psi(m \mid v_{k'}) \leq e^{\epsilon} \psi(m \mid v_k) \qquad \text{for every } m, k, \text{ and } k'.$$

The factor e^{ϵ} limits what any report can teach the firm: whatever the message, the firm's odds between any two valuations move by at most a factor e^{ϵ} relative to its prior. With $\epsilon = \log 2$, for example, no message can more than double the odds that a consumer's valuation is high rather than low. The restriction is imposed on the information structure itself and does not depend on the aggregate market x^*.

The information designer acts as a trusted intermediary, observing each consumer's true valuation and releasing only a randomized report to the firm.[8] The guarantee applies to the report used to set the consumer's price and does not cover information revealed by her subsequent purchase decision.

Consumers who generate the same message form a segment. After observing that message, the firm's posterior distribution is the distribution of valuations within the segment. By Bayes' rule, its odds satisfy

$$\frac{x_{k'}}{x_k} = \frac{\psi(m|v_{k'})}{\psi(m|v_k)} \frac{x^*_{k'}}{x^*_k}.$$

A message that is more likely under $v_{k'}$ than under v_k raises the relative share of $v_{k'}$ consumers

[7] Throughout the paper, log denotes the natural logarithm.

[8] The information structure satisfies the standard likelihood-ratio bound of local differential privacy (Kasiviswanathan et al., 2011). In the usual local implementation, consumers randomize their own data before disclosure; here, the trusted designer performs this randomization on their behalf.

in its segment. Differential privacy bounds precisely this change.

Therefore, (M, ψ) is ϵ-DP if and only if every induced segment's log odds between any two valuations stay within ϵ of the aggregate market's:

$$\left| \log \frac{x_{k'}}{x_k} - \log \frac{x^*_{k'}}{x^*_k} \right| \leq \epsilon, \qquad \text{for every } k \neq k'. \tag{1}$$

For the fixed aggregate market x^*, let

$$X(\epsilon) := \{x \in X : x \text{ satisfies (1)}\}$$

denote the set of private posterior markets. Every such market has full support: even a message suggesting a high willingness to pay must sometimes come from a low-valuation consumer. At zero budget, messages reveal nothing and $X(0) = \{x^*\}$. A larger budget allows larger changes in the odds, so the feasible set expands (Lemma A.1).

The connection between information structures and segmentations also runs backward. Every Bayes-plausible segmentation supported on $X(\epsilon)$ can be generated by a private information structure. To see the construction, use a message for each segment x and assign a type-v_k consumer to it with probability $\sigma(x)x_k/x^*_k$. The numerator is the mass of these consumers needed in that segment, and the denominator is their total mass. Bayes plausibility ensures that these assignment probabilities sum to one. Appendix A.1 verifies the privacy constraint.

Thus the designer can equivalently choose an information structure or a segmentation from

$$\Sigma(\epsilon) := \left\{\sigma \in \Sigma : \operatorname{supp}(\sigma) \subseteq X(\epsilon)\right\}.$$

We work with segmentations below because they make the firm's pricing incentives and the resulting surplus transparent. The information structure construction ensures that each segmentation has an implementable disclosure rule behind it.

2.3 Pricing and Welfare

The timing is as follows. The designer commits to an ϵ-DP information structure; a message is generated from each consumer's valuation and disclosed to the firm; the firm chooses a message-contingent price; the consumer observes her price and decides whether to buy. Consumers purchase when indifferent, so a consumer with valuation v_j buys at price p if and only if $v_j \geq p$.

Although the firm may charge any nonnegative price, restricting attention to prices in V is without loss.[9]

For each v_k and market x,

$$R_k(x) := v_k \sum_{j=k}^{K} x_j$$

is the firm's revenue from charging v_k in market x. At that price, consumers with valuations at least v_k buy, and their surplus is

$$U_k(x) := \sum_{j=k}^{K} (v_j - v_k) x_j.$$

Raising the price increases the margin on consumers who continue buying, but excludes those whose valuations fall below it. The composition of the market determines which price best balances these effects. Price v_k is *optimal* in market x if $R_k(x) \geq R_i(x)$ for every $i = 1, \ldots, K$, and

$$X_k := \{x \in X : R_k(x) \geq R_i(x), \text{ for every } i = 1, \ldots, K\}$$

denotes the set of markets in which v_k is optimal. Given the aggregate market x^*, select

$$i^* \in \operatorname{argmax}_{k \in \{1,\ldots,K\}} R_k(x^*),$$

write $v^* := v_{i^*}$ for the uniform monopoly price, and let $X^* = X_{i^*}$ be its optimality region. The following regularity condition keeps the aggregate market away from a price tie.

[9] For $k = 2, \ldots, K$, a price in $(v_{k-1}, v_k]$ sells to the same consumers as v_k at weakly lower revenue. A price below v_1 is dominated by v_1, and a price above v_K sells nothing.

Assumption 1. *The aggregate market x^* lies in the interior of X^* relative to X, i.e., $x^* \in \mathrm{ri}_X(X^*)$.*

Assumption 1 says the aggregate market lies in the interior of the region in which v^* is optimal. It rules out a tie at the no-disclosure benchmark and makes v^* unique (Lemma A.2).

Under $\epsilon-$DP, the markets in which v_k is optimal are

$$X_k(\epsilon) := X_k \cap X(\epsilon).$$

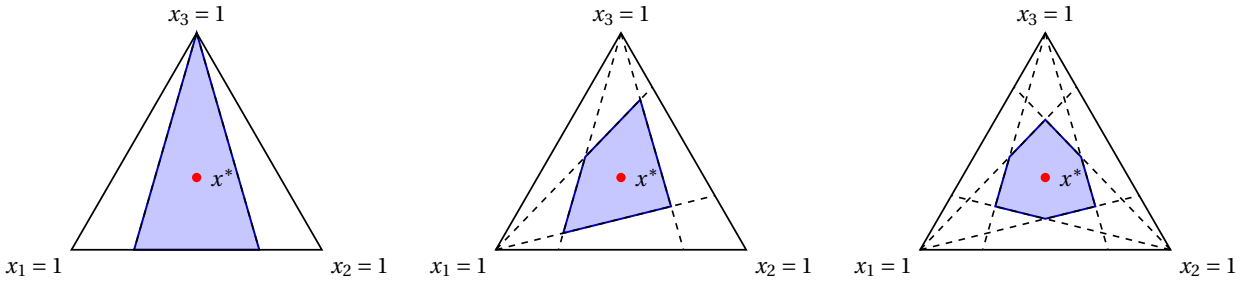
An empty $X_k(\epsilon)$ means that no admissible report can make the firm willing to charge v_k. If the intersection is nonempty, the price is available in at least one private market. Each nonempty intersection is a polytope (Lemma A.3); its boundaries reflect either a binding privacy constraint or the firm's indifference between prices.

The following example, adapted from BBM, shows how privacy reshapes the simplex of markets.

Example 1 (Three valuations with a uniform aggregate market). The consumers have three types of valuations, $V = \{1,2,3\}$, with equal proportion in the aggregate market, $x^* = (1/3, 1/3, 1/3)$. The budget is $\epsilon = \log 3$: no segment may more than triple the odds between any two valuations. As in BBM, the set of all markets is represented by a two-dimensional simplex. Figure 1 shows how $\epsilon-$DP progressively reduce the simplex. The left panel bounds the relative shares of valuations 1 and 2. The middle panel bounds the shares of 2 and 3, and the right panel imposes the remaining bound between 1 and 3. Together these restrictions leave a hexagon around the aggregate market. Its distance from the simplex's edges reflects the impossibility of ruling out any valuation. Figure 2 illustrates the monopoly-price regions: a feasible market's position inside the hexagon determines which price the firm is willing to charge. This example will also illustrate the welfare consequences of privacy in Section 5.

For a segmentation σ, a *pricing rule*

$$\phi : \mathrm{supp}(\sigma) \to \Delta(V)$$



(a) $\left|\log \frac{x_1}{x_2} - \log \frac{x_1^*}{x_2^*}\right| \leq \log 3$ (b) Adding $\left|\log \frac{x_2}{x_3} - \log \frac{x_2^*}{x_3^*}\right| \leq \log 3$ (c) Adding $\left|\log \frac{x_1}{x_3} - \log \frac{x_1^*}{x_3^*}\right| \leq \log 3$

Figure 1: ϵ-Differentially Private Markets

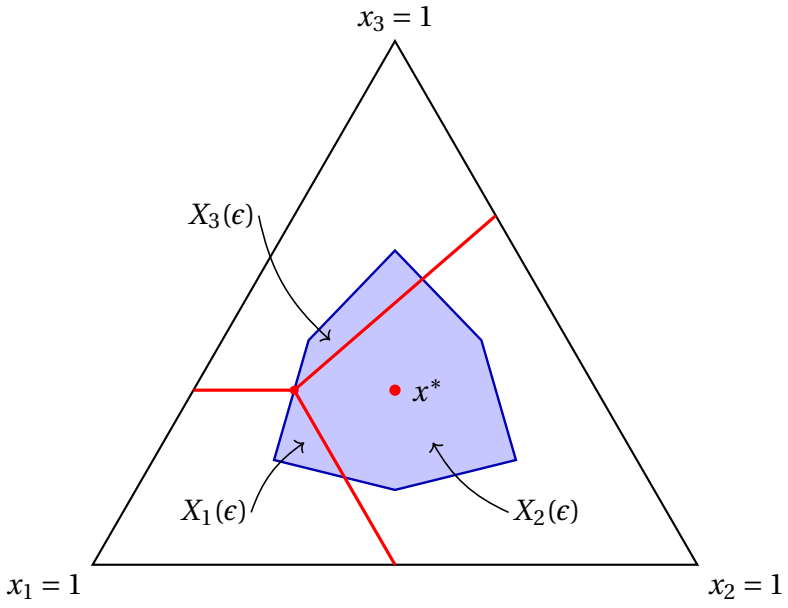


Figure 2: The Simplex of Markets

specifies a distribution over prices in every market in the support of σ, with $\phi_k(x)$ the probability of charging v_k in market x. A pricing rule is *optimal* if

$$v_k \in \text{supp}(\phi(x)) \implies x \in X_k \qquad \text{for every } x.$$

A pair (u, π) of consumer and producer surplus is a *welfare outcome*. Under the uniform monopoly price v^*, the welfare outcome is (u^*, π^*), where

$$u^* := \sum_{j=i^*}^{K} (v_j - v^*) x_j^*$$

is consumer surplus and

$$\pi^* := R_{i^*}(x^*) = v^* \sum_{j=i^*}^{K} x_j^*$$

is the uniform monopoly profit.

Given a segmentation σ and an associated pricing rule ϕ, consumer surplus is

$$u(\sigma, \phi) := \sum_{x \in \text{supp}(\sigma)} \sigma(x) \sum_{k=1}^{K} \phi_k(x) \sum_{j=k}^{K} (v_j - v_k) x_j,$$

producer surplus is

$$\pi(\sigma, \phi) := \sum_{x \in \text{supp}(\sigma)} \sigma(x) \sum_{k=1}^{K} \phi_k(x) R_k(x),$$

and total surplus is

$$w(\sigma, \phi) := \sum_{x \in \text{supp}(\sigma)} \sigma(x) \sum_{k=1}^{K} \phi_k(x) \sum_{j=k}^{K} v_j x_j.$$

In particular, $w(\sigma, \phi) = u(\sigma, \phi) + \pi(\sigma, \phi)$.

The set of attainable welfare outcomes under $\epsilon-$DP is

$$\mathcal{W}(\epsilon) := \left\{ (u(\sigma, \phi), \pi(\sigma, \phi)) : \sigma \in \Sigma(\epsilon),\ \phi \text{ is optimal} \right\}.$$

Without a privacy restriction, BBM show that three bounds completely describe attainable wel-

fare. Consumers can always decline to buy, and the firm can ignore disclosure and charge v^*. Consumer surplus must therefore be nonnegative, and producer surplus cannot fall below the uniform monopoly profit. Finally, total surplus is bounded above by the value of serving every consumer.

Theorem 0 (BBM)**.** *There exists a segmentation and associated optimal pricing rule ϕ that implement (u,π) if and only if*

$$u \geq 0, \qquad \pi \geq \pi^*, \qquad u+\pi \leq \sum_{j=1}^{K} v_j x_j^*.$$

In this paper, we do not impose a particular objective on the information designer. Our goal is the characterization of the set $\mathcal{W}(\epsilon)$, and the segmentation and pricing rules that attain its boundary. The question is how much of BBM's welfare triangle remains available once every disclosure must protect consumers' privacy.

3. Preliminary Results

We first establish two general properties of the set of welfare outcomes under $\epsilon-$DP. The first is inherited directly from the unrestricted BBM benchmark.

Observation 1. *For any feasible segmentation σ satisfying* $\text{supp}(\sigma) \subseteq X(\epsilon)$ *and any associated optimal pricing rule ϕ, the resulting producer surplus is weakly greater than π^*.*

This follows from Theorem 0, since every private segmentation is also feasible without privacy. Because $X(\epsilon)$ is a subset of the unrestricted set of markets, every segmentation satisfying $\epsilon-$DP is also feasible in the BBM environment. Since π^* is a lower bound on producer surplus over the larger, unrestricted set of segmentations, the same lower bound necessarily applies to the smaller set of $\epsilon-$DP segmentations.

Although the lower bound of producer surplus is preserved under $\epsilon-$DP, the privacy budget determines how much the firm's optimal price can vary across segments. In particular, sufficiently stringent privacy regulation makes all disclosed information irrelevant for pricing.

Observation 2. *There exists $\underline{\epsilon} > 0$ such that, for every $\epsilon \in [0, \underline{\epsilon})$, the set of welfare outcomes is the singleton:*

$$\mathcal{W}(\epsilon) = \left\{(u^*, \pi^*)\right\}.$$

The intuition follows from the strict optimality of the uniform monopoly price. Because x^* lies in the relative interior of the set of markets in which v^* is optimal, the price v^* has a strict profit advantage over every competing price in the aggregate market. This advantage continues to hold for markets sufficiently close to x^*.

When the privacy budget is small, the ϵ-DP constraint restricts a market to be sufficiently similar to the aggregate market. Thus, all feasible markets remain within the region in which v^* is the unique optimal price. Because the firm charges v^* in every segment, price discrimination disappears. Therefore, the consumer surplus and producer surplus coincide with their no-disclosure levels, u^* and π^*.

More generally, tightening the privacy constraint reduces the set of markets, and therefore the set of attainable welfare outcomes. It limits the producer-surplus gains that may arise from segmentation, but Observation 1 shows that producer surplus cannot fall below π^*. Its effect on consumer surplus is more subtle because restricting price discrimination may affect both trade and the division of surplus. To examine these effects more closely, we first consider the case of binary valuations.

4. Binary Benchmark

With two valuations, the role of privacy regulation depends on whom the firm serves before receiving information. When $v^* = v_2$, only high-valuation consumers buy, and they keep no surplus. Information can benefit consumers by inducing the firm to offer the low price to some of them. When $v^* = v_1$, everyone already buys at the low price; information instead creates opportunities to charge some high-valuation consumers more. Privacy limits the first change but protects consumers against the second. Once privacy is relaxed enough for the firm to strictly prefer the alternative price in some segment, the set of welfare outcomes form a triangle.

Its three vertices correspond to maximum producer surplus, maximum consumer surplus, and minimum consumer surplus.

Let $V = \{v_1, v_2\}$ and write $\alpha = x_1^*/(1-x_1^*)$, so $x_1^* = \alpha/(1+\alpha)$. A market is described by its share x_1 of low-valuation consumers. Condition (1) implies

$$X(\epsilon) = \left\{(x_1, 1-x_1) : \frac{\alpha}{\alpha + e^\epsilon} \leq x_1 \leq \frac{\alpha e^\epsilon}{1+\alpha e^\epsilon}\right\}.$$

The lower endpoint is the market most concentrated on high valuations that privacy permits; the upper endpoint is the market most concentrated on low valuations. A larger budget moves the endpoints farther apart, but both types remain present at every finite budget.

The firm faces a trade-off between the revenue from serving more consumers and charging higher price to those willing to pay more. It is indifferent between charging v_1 to everyone and charging v_2 to high-valuation consumers at $x_1 = (v_2 - v_1)/v_2$. Each price is uniquely optimal in some feasible market exactly when this indifference market lies strictly inside the privacy interval:

$$\frac{\alpha}{\alpha + e^\epsilon} < \frac{v_2 - v_1}{v_2} < \frac{\alpha e^\epsilon}{1+\alpha e^\epsilon} \quad \Longleftrightarrow \quad \epsilon > \epsilon_B := \left|\log \frac{v_2 - v_1}{\alpha v_1}\right|.$$

Below ϵ_B, every feasible segment has the same optimal price. At ϵ_B, the alternative price becomes optimal at one endpoint of the interval. The firm is then indifferent, but consumers are not: at the indifference market, they receive zero if the firm charges v_2 and $v_1(v_2 - v_1)/v_2$ if it charges v_1. The low price gives high-valuation consumers a discount and brings low-valuation consumers into the market. The additional sales exactly compensate the firm for the lower markup. Thus welfare can change as soon as the second price becomes optimal, even though the firm need not gain any profit.

Figure 3 depicts producer surplus $\pi(x_1) = \max\{v_1, v_2(1-x_1)\}$. The blue portion is its restriction to $X(\epsilon)$; the red chord is its concave closure. At the aggregate share of low valuations, the height of this chord gives the largest average profit that a feasible segmentation can deliver.

The three boundary outcomes below are defined for $\epsilon \geq \epsilon_B$. Above the threshold, the firm strictly prefers the high price at the lower endpoint of the privacy interval and the low price at the upper

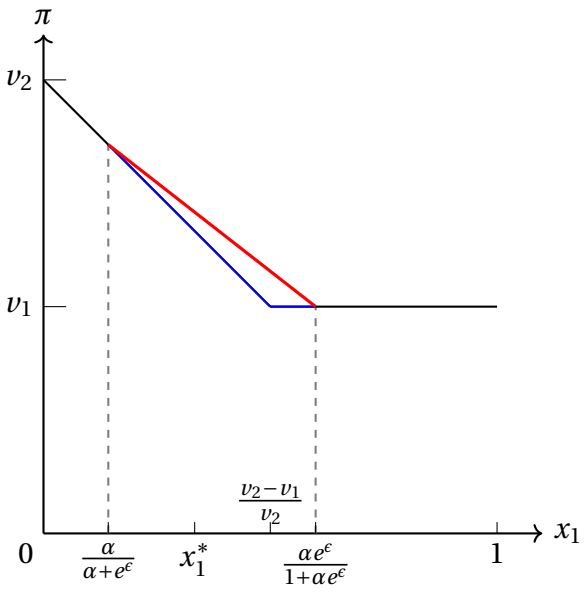


Figure 3: Producer surplus under ϵ-DP

endpoint. Choosing the two segments and their weights then determines how many consumers receive each price.

Maximum producer surplus. The firm gains most from separating the two types as much as privacy permits. The segment concentrated on high valuations receives the high price, while the other segment receives the low price. Notice that producer surplus is convex in x_1, so splitting the aggregate market between the endpoints of the privacy interval reaches the red chord in Figure 3. The segment weights preserve the aggregate share of each valuation. This segmentation yields

$$A_\epsilon = (u_A^\epsilon, \pi_A^\epsilon) = \left(\frac{v_2 - v_1}{(1+\alpha)(1+e^\epsilon)}, \frac{(1+\alpha e^\epsilon) v_1 + e^\epsilon v_2}{(1+\alpha)(1+e^\epsilon)} \right). \tag{2}$$

Privacy prevents this segmentation from perfectly separating the two types. High-valuation consumers who remain in the low-price segment keep a positive surplus. Low-valuation consumers who remain in the segment priced at v_2 do not buy. The first effect limits how much surplus the firm captures; the second destroys gains from trade. This exclusion creates a total-surplus loss of

$$\overline{w} - w(A_\epsilon) = \frac{\alpha v_1}{(1+\alpha)(1+e^\epsilon)} > 0, \tag{3}$$

where $\overline{w} = (\alpha v_1 + v_2)/(1+\alpha)$ is first-best surplus.

Consumer-surplus bounds. Minimum and maximum consumer surplus are attained on the profit floor $\pi = \pi^*$. The relevant segmentations keep the original uniform price optimal in both segments. One segment is an indifference market, so changing its price alters consumer surplus without changing the firm's revenue.

When $v^* = v_2$, low-valuation consumers are too scarce in the aggregate market to make the low price attractive. To induce a discount, the designer groups enough of them into an indifference market and assigns the remaining consumers to a segment concentrated on high valuations. Specifically, split the aggregate market between the lower endpoint $\alpha/(\alpha + e^{\epsilon})$ and the indifference market $(v_2 - v_1)/v_2$. Charging the high or low price at the latter gives, respectively,

$$B_\epsilon = \left(0, \frac{v_2}{1+\alpha}\right) \quad \text{and} \quad C_\epsilon = \left(\frac{\alpha v_1 (v_2 - v_1)(e^{\epsilon} - 1)}{(1+\alpha)((v_2 - v_1)e^{\epsilon} - \alpha v_1)}, \frac{v_2}{1+\alpha}\right). \tag{4}$$

No disclosure already gives consumers zero surplus, so privacy cannot raise the minimum. To attain the maximum, the designer wants as many consumers as possible to receive the low price at the indifference market. But the remaining segment must also contain some low-valuation consumers to satisfy privacy. Those consumers cannot be reassigned to support a discount elsewhere. This restriction reduces the maximum below its unrestricted value $\alpha v_1/(1+\alpha)$.

When $v^* = v_1$, the aggregate market contains enough low-valuation consumers for the firm to serve everyone. To reduce consumer surplus, the designer must now create a segment in which the high price becomes attractive. Combine the indifference market with the upper endpoint $\alpha e^{\epsilon}/(1+\alpha e^{\epsilon})$ instead. The two pricing choices give

$$B_\epsilon = \left(\frac{(v_2 - v_1)(\alpha v_1 - (v_2 - v_1))}{(1+\alpha)(\alpha v_1 e^{\epsilon} - (v_2 - v_1))}, v_1\right) \quad \text{and} \quad C_\epsilon = \left(\frac{v_2 - v_1}{1+\alpha}, v_1\right). \tag{5}$$

Privacy prevents complete separation of the two types: some high-valuation consumers must remain in the segment charging v_1. These consumers retain surplus $v_2 - v_1$, so minimum consumer surplus remains strictly positive. Maximum consumer surplus is unaffected. Charging

v_1 to everyone gives C_ϵ, where high-valuation consumers receive the largest possible surplus and everyone buys. This achieves the same consumer maximum as without privacy, along with first-best total surplus. When identifying these endpoints, we allow the firm to choose either price whenever it is indifferent.

Figures 4 and 5 plot the upper and lower envelopes of consumer surplus in the two cases.

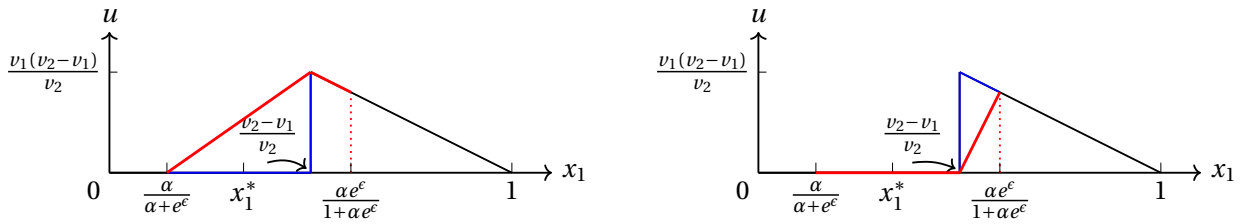


Figure 4: Maximum and minimum consumer surplus when $x_1^* < \frac{v_2-v_1}{v_2}$

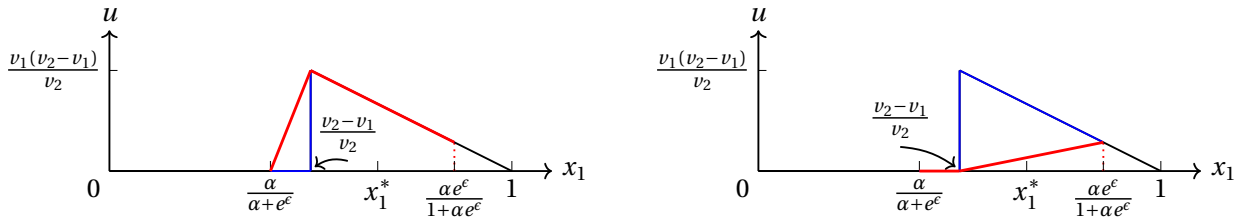


Figure 5: Maximum and minimum consumer surplus when $x_1^* > \frac{v_2-v_1}{v_2}$

The three outcomes also describe every trade-off between consumer and producer surplus. Combining feasible segmentations fills their convex hull; the linear surplus formulas within each price region imply that no feasible outcome lies outside it. At ϵ_B, the firm can vary consumer surplus by choosing between tied prices, but cannot yet earn more than the uniform monopoly profit.

Proposition 1 (Binary set of welfare outcomes)**.** *If $K = 2$, then*

$$\mathcal{W}(\epsilon) = \begin{cases} \{(u^*, \pi^*)\}, & \epsilon < \epsilon_B, \\ \operatorname{co}\{B_\epsilon, C_\epsilon\}, & \epsilon = \epsilon_B, \\ \operatorname{co}\{A_\epsilon, B_\epsilon, C_\epsilon\}, & \epsilon > \epsilon_B, \end{cases}$$

where the vertices are given in equations (2), (4), and (5). At ϵ_B, $A_\epsilon = C_\epsilon$ if $v^* = v_2$, and $A_\epsilon = B_\epsilon$ if $v^* = v_1$.

Above the threshold, A_ϵ maximizes producer surplus, B_ϵ minimizes consumer surplus, and C_ϵ maximizes both consumer and total surplus. The last coincidence has a direct economic explanation. High-valuation consumers buy at either price, so gains in total surplus come entirely from allowing more low-valuation consumers to buy. The segmentation attaining the consumer maximum also serves the largest feasible mass of low-valuation consumers. It excludes fewer of them than the producer-optimal segmentation, and excludes none when $v^* = v_1$. Thus C_ϵ is the unique total-surplus-maximizing welfare outcome, with strictly less total-surplus loss than A_ϵ.

Figure 6 places the two private welfare triangles inside the unrestricted BBM triangle. When the uniform price is high, privacy cuts off part of the consumer gains available under unrestricted disclosure. When the uniform price is low, it prevents the most severe consumer losses. In both cases, the firm cannot extract all gains from trade: full support rules out the unrestricted producer-optimal outcome $(0, \overline{w})$. The private set of welfare outcomes is therefore a strict subset at every finite budget.

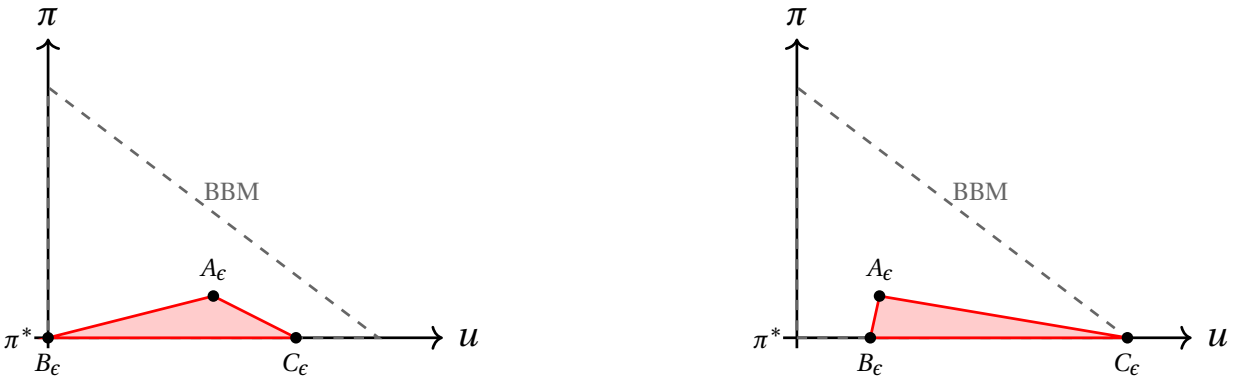


Figure 6: The binary ϵ-DP set of welfare outcomes when the uniform monopoly price is v_2 (left) and v_1 (right)

Relaxing privacy. A larger budget permits sharper separation of valuations. The firm can exploit this separation to increase profit, but its implications for consumers depend on which price prevails without disclosure.

Corollary 1. *For $\epsilon > \epsilon_B$, maximum producer surplus strictly increases with ϵ. If $v^* = v_2$, maximum consumer surplus strictly increases while the minimum remains zero. If $v^* = v_1$, minimum*

consumer surplus strictly decreases while the maximum remains u^*.

At the producer-optimal outcome, sharper separation sends more high-valuation consumers to the high-price segment and more low-valuation consumers to the low-price segment. The firm earns more, consumers keep less surplus, and fewer low-valuation consumers are excluded. Accordingly, A_ϵ moves up and to the left in the figures even as total surplus rises.

Along the profit floor, a larger budget permits more consumers to be assigned to the indifference market. If $v^* = v_2$, more consumers can receive the low price, so C_ϵ moves right. If $v^* = v_1$, more high-valuation consumers can be charged the high price, so B_ϵ moves left. The other endpoint stays fixed because no disclosure continues to attain it. When $v^* = v_1$, maximum total surplus also remains at its first-best level. As $\epsilon \to \infty$, all three vertices converge to their unrestricted counterparts.

5. General Finite Valuations

With more than two valuations, disclosure can induce several discounts and markups relative to the uniform price. The set of welfare outcomes need not remain a triangle, but the main result for consumers survives: every segmentation that maximizes their surplus leaves the firm with exactly π^*. Giving the firm additional profit cannot help consumers attain a higher surplus. At the other extreme, a profit-preserving segmentation attains minimum consumer surplus under a bound on the revenue that higher prices could generate.

The analysis separates the prices that privacy makes available from the welfare outcomes that the designer can achieve with them. This distinction matters because a newly available price may be used in only a small part of the market. Throughout this section, the privacy budget is finite.

5.1 A finite description of welfare outcomes

The designer may use many messages, but welfare depends on the prices those messages induce. Segments that receive the same price can be pooled: the price remains optimal, privacy is preserved, and neither surplus changes. It therefore suffices to record how many consumers of each valuation face each price.

Let q_{kj} denote the mass of valuation-v_j consumers assigned to price v_k. A feasible assignment $q = (q_{kj})$ has nonnegative entries and satisfies the following three conditions:

$$\sum_{k=1}^{K} q_{kj} = x_j^* \qquad \text{for every } j, \tag{6}$$

$$q_{kj} x_\ell^* \le e^\epsilon q_{k\ell} x_j^* \qquad \text{for every } k, j, \ell, \tag{7}$$

$$v_k \sum_{j=k}^{K} q_{kj} \ge v_\ell \sum_{j=\ell}^{K} q_{kj} \qquad \text{for every } k, \ell. \tag{8}$$

Condition (6) is Bayes plausibility, condition (7) imposes $\epsilon-$DP within every price-indexed segment, and condition (8) makes the price indexing that segment optimal.

Write $Q(\epsilon)$ for the polytope defined by these conditions. Once consumers are assigned to prices, welfare is linear in their masses. Buyers retain the difference between valuation and price, and the firm receives the price from everyone who buys:

$$u(q) = \sum_{k=1}^{K} \sum_{j=k}^{K} (v_j - v_k) q_{kj},$$
$$\pi(q) = \sum_{k=1}^{K} v_k \sum_{j=k}^{K} q_{kj}.$$

Proposition 2 (The welfare polygon)**.** *The set of attainable welfare outcomes is*

$$\mathcal{W}(\epsilon) = \left\{ (u(q), \pi(q)) : q \in Q(\epsilon) \right\}.$$

It is a compact convex polygon, possibly a point or a line segment. Every outcome can be implemented using at most K market–price pairs, at most one per price.

Each feasible assignment gives both a welfare outcome and the segmentation that implements it.[10] Convexity has a direct interpretation: randomizing between two feasible disclosures gives the corresponding average of their welfare outcomes. To find a boundary point, solve

$$\max_{q \in Q(\epsilon)} \{au(q) + b\pi(q)\}.$$

The weights (a, b) describe how the designer values the two surpluses; varying them traces the polygon. With several valuations, a change in these weights can alter which groups receive discounts, which face markups, and which privacy constraints bind. These additional choices explain why the boundary can have more corners than in the binary case.

Relaxing privacy enlarges $Q(\epsilon)$ and therefore $\mathcal{W}(\epsilon)$. Maximum profit, maximum consumer surplus, and maximum total surplus weakly rise, while minimum consumer surplus weakly falls. These are changes in what disclosure can achieve; they do not imply that every consumer benefits from weaker privacy. For every finite budget, the welfare polygon is a strict subset of the BBM triangle. It approaches that triangle as the privacy restriction vanishes. Appendix A.3 gives the limit argument and the geometric representation used below.

5.2 Consumer surplus and the profit floor

The profit floor provides a useful starting point for understanding how disclosure helps or hurts consumers. If the firm remains willing to charge v^* in every segment, its expected profit stays at the uniform-monopoly-price benchmark. It can still choose a different price when that price ties v^*, changing both who buys and what buyers retain.

Let $u_{\min}(\epsilon)$ and $u_{\max}(\epsilon)$ be the smallest and largest consumer surplus in $\mathcal{W}(\epsilon)$, and let $\overline{w} = \sum_j v_j x_j^*$ be first-best total surplus. Call a price *active* if it is optimal in some private market, and write

$$\mathcal{A}(\epsilon) = \{k : X_k(\epsilon) \neq \emptyset\}.$$

An active price need not be uniquely optimal; a tie is enough for the firm to charge it.

[10] A nonzero row q_k describes a segment of weight $s_k = \sum_j q_{kj}$, posterior q_k / s_k, and price v_k.

A segmentation is *profit-preserving* if v^* remains optimal in every segment:

$$\Sigma^*(\epsilon) = \left\{\sigma \in \Sigma(\epsilon) : \text{supp}(\sigma) \subseteq X^* \cap X(\epsilon)\right\}.$$

Proposition 3 (The profit floor)**.** *Fix a finite privacy budget ϵ. Under any private segmentation $\sigma \in \Sigma(\epsilon)$ and any associated optimal pricing rule, profit is at least π^*, with equality if and only if v^* is optimal in every segment of σ. The set of consumer surpluses attainable with profit π^* is a closed interval $[u^-(\epsilon), u^+(\epsilon)]$ that contains u^*. If a segmentation and optimal pricing rule implement $(u^-(\epsilon), \pi^*)$, the rule charges the highest optimal price in every segment; if they implement $(u^+(\epsilon), \pi^*)$, it charges the lowest.*

The firm's option to keep charging v^* explains the equivalence. Any strict improvement over this fallback in one segment raises expected profit, because no other segment can offset it with a loss. Profit therefore equals π^* precisely when the firm is indifferent between its chosen price and v^* in every segment.

This is the same indifference that supports the bottom edge of the BBM triangle. A higher tied price reduces consumer surplus through higher payments and exclusion; a lower tied price increases it through discounts and additional purchases. Privacy limits the compositions and weights of the segments that sustain these ties. The uninformative disclosure preserves profit, so

$$u_{\min}(\epsilon) \leq u^-(\epsilon) \leq u^* \leq u^+(\epsilon) \leq u_{\max}(\epsilon).$$

The interval on the profit floor could, in principle, cover only part of the attainable consumer-surplus range. At its upper endpoint, the next proposition rules out that possibility.

Proposition 4 (Consumer-optimal disclosure)**.** *Every consumer-surplus-maximizing segmentation preserves uniform monopoly profit. Hence*

$$u_{\max}(\epsilon) = u^+(\epsilon).$$

Proof sketch. By Lemma A.7, every maximizer uses only prices at most v^*. Consumers valued

at least v^* therefore buy in every segment. Collapsing them into a single type valued at v^* preserves revenues at these prices and subtracts only the constant u^* from aggregate consumer surplus. Restoring the tail in its aggregate proportions preserves privacy and keeps higher prices strictly less profitable. Lemma A.9 thus reduces the problem to a market with $v^* = v_K$, without changing producer surplus.

In this reduced market, fix an arbitrary maximizer $\bar{q}$. Among maximizers with producer surplus $\pi(\bar{q})$, choose q to minimize the auxiliary objective $\mathcal{P}(q) = \sum_{r,j} q_{rj}^2 / x_j^*$. If a row earns more at its assigned price than at v_K, some lower valuation must have a larger likelihood multiplier than the highest type; otherwise, the aggregate market's strict revenue ranking would make v_K more profitable. Bayes plausibility supplies another row with the opposite multiplier inequality.

Working down the valuation tail, the proof exchanges small masses of these two types across the two rows. The changed multipliers move toward one another within each row. All-pairs privacy bounds the largest multiplier by e^ϵ times the smallest, so keeping every multiplier within its previous range preserves every privacy inequality. Pricing incentives are checked separately: comparisons that can bind remain unchanged or improve, while adverse changes affect only comparisons with strict slack. A sufficiently small exchange therefore preserves both privacy and optimal pricing.

The exchange either raises consumer surplus or preserves both surpluses while strictly reducing $\mathcal{P}$. The first possibility contradicts consumer optimality; the second contradicts the auxiliary minimization. These contradictions rule out every positive row profit gap, so $\pi(q) = \pi^*$. Because $\pi(q) = \pi(\bar{q})$ and $\bar{q}$ was arbitrary, every maximizer preserves profit. Appendix A.4 gives the full argument. □

When consumer-optimal disclosure induces a discount, revenue from additional buyers exactly offsets the lower payments of consumers who would have bought at v^*. Consumers benefit from lower prices and expanded trade while the firm retains its uniform monopoly profit.

At the opposite extreme, a segmentation that minimizes consumer surplus never uses a price below v^* (Lemma A.7). If no price above v^* is active, the uninformative segmentation is there-

fore a minimizer, and $u_{\min}(\epsilon) = u^-(\epsilon) = u^*$. Otherwise, the question is whether the firm can charge higher prices while remaining indifferent between each of them and v^*. In a segment priced above v^*, this indifference relies on excluded buyers: consumers who buy at v^* but at no higher active price. They add revenue at v^* but not at any higher active price.

Proposition 5 (Minimum consumer surplus under an upper-tail bound)**.** *Let v_a and v_b be the lowest and second-highest active prices above v^*, with $v_b = v_a$ if only one is active. Suppose the firm would earn no more than π^* even if every consumer valued at least v_a paid v_b:*

$$v_b \sum_{j=a}^{K} x_j^* \le \pi^*. \tag{9}$$

Then there exists a profit-preserving segmentation that minimizes consumer surplus, so $u_{\min}(\epsilon) = u^-(\epsilon)$. Condition (9) holds whenever at most two prices above v^* are active.

Since $\pi^* = v^* \sum_{j \ge i^*} x_j^*$, condition (9) is equivalent to

$$(v_b - v^*) \sum_{j=a}^{K} x_j^* \le v^* \sum_{j=i^*}^{a-1} x_j^*.$$

The left side bounds what the firm could gain by charging v_b instead of v^* to consumers who can afford a higher active price. The right side is the revenue that excluded buyers generate at v^*. The condition therefore requires enough excluded buyers relative to the upper tail. The highest active price does not enter; the proof handles it separately. With one or two active prices above v^*, $v_b = v_a$ and the left side of (9) equals $R_a(x^*) < \pi^*$, so the condition holds automatically. With at most three valuations, $u_{\min}(\epsilon) = u^-(\epsilon)$ at every finite budget.[11]

The proof selects, among minimizers, one with the lowest profit and, among those, the least informative one. It then shows that no segment earns more than it would at v^*, in three steps. First, such a segment contains at least its proportional share of excluded buyers. Otherwise, moving some of them into it from another segment would preserve privacy and every price,

[11] For $K \ge 2$, a sufficient condition that is independent of the privacy budget is $v_{K-1} \sum_{j>i^*} x_j^* \le \pi^*$. When at least two higher prices are active, $v_b \le v_{K-1}$ and $a > i^*$; with fewer, the equality already holds. Conditional on buying at the uniform price, this condition requires the share of consumers valued exactly at v^* to be at least $1 - v^*/v_{K-1}$.

and would lower consumer surplus or profit or make the segmentation less informative. Second, the segment's share of consumers valued at least v_a is then no larger than in the aggregate market, so condition (9) applies within the segment: no active price above v^* other than the highest earns more there than v^*. Third, the highest active price cannot be uniquely optimal at a minimum, since adding a small share of a lower-priced segment would keep that price optimal and lower consumer surplus. Appendix A.4 gives the details.

Under the upper-tail bound, the profit floor therefore spans the entire attainable consumer-surplus range, $[u_{\min}(\epsilon), u_{\max}(\epsilon)]$. The minimum result guarantees the existence of a profit-preserving minimizer; the maximum result applies to every maximizer. Since the former uses no price below v^*, its consumer-surplus loss relative to uniform pricing equals the loss of total surplus from forgone trades. Consumers can therefore be harmed even when the firm earns no additional profit.

The equality $u_{\min}(\epsilon) = u^-(\epsilon)$ can fail if the upper-tail bound is violated.

Example 2. Let $V = \{57, 58, 59, 60\}$, $x^* = (4, 7, 9, 180)/200$, and $\epsilon = \log 3$. The uniform price is $v_1 = 57$, so $\pi^* = 57$. All three higher prices are active, and the upper-tail bound fails because $v_3(1 - x_1^*) = 57.82 > 57$. Minimum consumer surplus is $u_{\min}(\epsilon) = 15093/31400 \approx 0.4807$, strictly below the profit-floor minimum $u^-(\epsilon) = 660339/1372600 \approx 0.4811$, so every minimizer gives the firm strictly more than π^*.

Privacy also limits how far these extrema can reach. Every disclosed segment must retain every valuation, so consumers receiving different prices cannot be completely separated by valuation.

Proposition 6 (Extreme uniform prices)**.** *Fix a finite privacy budget ϵ. Then $u_{\min}(\epsilon) = 0$ if and only if $v^* = v_K$, and $u_{\max}(\epsilon) = \overline{w} - \pi^*$ if and only if $v^* = v_1$.*

At any price below v_K, the highest-valuation consumers retain some surplus, and privacy places some of them in every segment. Zero consumer surplus therefore requires charging v_K everywhere. At the other extreme, any price above v_1 excludes some buyers, so first-best total surplus requires charging v_1 everywhere. Bayes plausibility permits either outcome only when the same

price is optimal in the aggregate market.

It follows that maximum total surplus equals $\overline{w}$ exactly when $v^* = v_1$, and maximum profit is always strictly below $\overline{w}$. For an interior monopoly price, privacy prevents both complete surplus extraction and an efficient allocation that gives consumers all surplus above π^*.

5.3 When privacy changes consumer surplus

When privacy is sufficiently strict, the firm charging v^* in every segment and segmentation then has no welfare effect. Consumer surplus can change only once another price becomes optimal in some segment. To locate that point, the activation threshold of a price records the smallest privacy budget that makes a price available.

For a full-support market x, define

$$\rho(x) = \max_j \log \frac{x_j}{x_j^*} - \min_j \log \frac{x_j}{x_j^*},$$

and set $\rho(x) = +\infty$ on the boundary of the simplex. The quantity $\rho(x)$ measures the largest difference between the log likelihoods with which two valuations generate this posterior. It is the privacy budget needed to admit that market. The privacy constraint is exactly

$$X(\epsilon) = \{x : \rho(x) \le \epsilon\}. \tag{10}$$

For price v_k, let

$$\epsilon_k^{\text{act}} = \min_{x \in X_k} \rho(x). \tag{11}$$

These minima are attained. The uniform price has threshold zero and every other price has a positive, finite threshold, with

$$k \in \mathcal{A}(\epsilon) \quad \Longleftrightarrow \quad \epsilon \ge \epsilon_k^{\text{act}}. \tag{12}$$

Appendix A.5 computes each threshold with a linear program.

The direction of the price change determines whether segmentation can benefit or harm con-

sumers. A lower active price makes a discount possible; a higher active price makes a markup or exclusion possible. Define the first lower- and higher-price thresholds by

$$\epsilon^- = \min_{k<i^*} \epsilon_k^{\text{act}}, \qquad \epsilon^+ = \min_{k>i^*} \epsilon_k^{\text{act}}, \tag{13}$$

taking the minimum of an empty set to be $+\infty$.

Proposition 7 (Activation and consumer surplus)**.** *Fix a finite privacy budget ϵ. Then $u_{\min}(\epsilon) \le u^-(\epsilon) < u^*$ if and only if $\epsilon \ge \epsilon^+$. Moreover, $u_{\max}(\epsilon) > u^*$ if and only if $\epsilon \ge \epsilon^-$. Below the corresponding threshold, each endpoint equals u^*.*

An active price can be used in a small segment while leaving the rest of the market close enough to x^* that the firm still charges v^*. A lower price then raises the surplus of some buyers without changing the price faced by everyone else. A higher price reduces surplus in that segment through higher payments or exclusion. A tie suffices for either construction, which is why the inequalities hold at the thresholds themselves. The first gains and losses can also be achieved while preserving uniform profit.

The proposition describes the onset and direction of welfare changes. The size of those changes depends on how many consumers can be assigned to the new prices while respecting privacy and preserving the aggregate market. A larger budget can permit greater gains or losses even when the set of active prices is unchanged. Once both thresholds are reached, $u_{\min}(\epsilon) < u^* < u_{\max}(\epsilon)$: disclosure can then help consumers or hurt them, depending on its design.

5.4 Producer surplus and the rest of the polygon

To characterize maximum producer surplus, we begin with the firm's response to more detailed information. Splitting a segment into smaller segments that satisfy privacy cannot reduce expected profit: the firm can continue charging its original price in each segment or choose prices better suited to their demand. This allows us to construct a producer-optimal disclosure from the vertices of the set of private markets $X(\epsilon)$.

These vertices have a simple economic interpretation. Each places as much weight as privacy permits on a particular group of valuation types. Let S be a nonempty proper subset of $\{1,\dots,K\}$, and write $x^*(S) = \sum_{j\in S} x_j^*$ for this group's share in the aggregate market. We denote the corresponding market by $z^S(\epsilon)$. Its share of valuation-v_j consumers is

$$z_j^S(\epsilon) = \begin{cases} \dfrac{e^\epsilon x_j^*}{1+(e^\epsilon-1)x^*(S)}, & j \in S, \\ \dfrac{x_j^*}{1+(e^\epsilon-1)x^*(S)}, & j \notin S. \end{cases} \tag{14}$$

Starting from the aggregate market, this construction multiplies the shares of types in S by e^ϵ and then normalizes all shares to sum to one. It preserves relative shares within each group while increasing the odds of membership in S by the largest factor privacy allows.

For $\epsilon > 0$, these markets are exactly the vertices of $X(\epsilon)$, as established in Lemma A.5. They generalize the two endpoints of the privacy interval in the binary model. Every private market is a mixture of these vertices, so refining segments into them weakly raises profit. Maximizing producer surplus therefore reduces to choosing their weights.

Proposition 8 (Producer-optimal segmentation)**.** *Maximum producer surplus is attained on at most K markets $z^S(\epsilon)$ and equals*

$$\max_{\lambda\geq 0}\left\{\sum_S \lambda_S \max_k R_k(z^S(\epsilon)) : \sum_S \lambda_S z^S(\epsilon) = x^*\right\}.$$

The firm benefits weakly from replacing a pooled market with more detailed reports, because it can always ignore the distinction and keep its old price. Formally, profit is convex in the posterior: the firm chooses the best of several linear revenue functions. Dividing a market into the privacy vertices that average to it therefore weakly raises profit. The designer can repeat this refinement until only privacy vertices remain, then choose their weights to match the aggregate market. These weights are the remaining choice in the proposition's linear program.

Let

$$\overline{\pi}_\epsilon(u) = \max\{\pi : (u,\pi) \in \mathcal{W}(\epsilon)\}$$

be the largest profit compatible with consumer surplus u. This upper boundary records the trade-offs faced by a designer who wants to raise profit while delivering a given surplus to consumers. It is concave and piecewise linear, since $\mathcal{W}(\epsilon)$ is a polygon. Each corner marks a change in the segmentation that best balances the two objectives.

If segmentation can raise profit, the upper boundary lies strictly above π^* in the interior of the consumer-surplus range. At $u_{\max}(\epsilon)$ it returns to π^* by Proposition 4: reaching the greatest consumer surplus leaves no room for an additional profit gain. Away from the producer optimum, boundary outcomes can require price-indifference markets as well as privacy vertices.[12]

When the upper-tail bound in Proposition 5 holds, both consumer-surplus endpoints are attainable at π^*. Convexity then gives

$$\mathcal{W}(\epsilon) = \left\{(u,\pi) : u_{\min}(\epsilon) \le u \le u_{\max}(\epsilon),\ \pi^* \le \pi \le \overline{\pi}_\epsilon(u)\right\}. \tag{15}$$

In particular, this description applies whenever at most two prices above v^* are active, including every market with at most three valuations.

Example 1 continued. For $V = \{1,2,3\}$ and $x^* = (1/3, 1/3, 1/3)$, the uniform price is 2, with $\pi^* = 4/3$ and $u^* = 1/3$. Both other prices activate at $\epsilon = \log 2$. Below this threshold, every consumer faces price 2, so the set of welfare outcomes is a point. At the threshold, the firm can offer a discount to 1 or a markup to 3 through price ties, and the set becomes a line segment on the profit floor. Above it, disclosure can also increase profit, and the set is a pentagon.

Only one price above v^* can be active, so Proposition 5 requires no additional restriction here. Whenever price 3 is active, $v_a = v_b = 3$ and the bound holds automatically: $v_b x_3^* = 1 < 4/3 = \pi^*$. Thus both consumer-surplus endpoints lie on the profit floor, as in equation (15). At $\epsilon = \log 3$, the vertices are

$$B = \left(\frac{1}{15}, \frac{4}{3}\right), \quad C = \left(\frac{3}{5}, \frac{4}{3}\right), \quad D = \left(\frac{1}{12}, \frac{17}{12}\right), \quad D' = \left(\frac{1}{2}, \frac{17}{12}\right), \quad A = \left(\frac{4}{15}, \frac{22}{15}\right).$$

[12] Appendix A.3 describes these relative extreme points and shows that at most K market–price pairs generate any welfare vertex.

Figure 7 displays the resulting trade-offs. The endpoints B and C lie on the profit floor. At B, a share $8/15$ of consumers faces price 3, so only the highest valuation buys in that segment, retaining no surplus. At C, the same share faces price 1, so every valuation buys and buyers above 1 receive a discount. The remaining consumers face the uniform price 2. Both nonuniform prices tie 2 in their respective segments, keeping expected profit at $4/3$.

The producer-optimal point A uses privacy vertices throughout. The intermediate vertices D and D' show how the additional pricing option changes the boundary: D is implemented with prices 2 and 3, whereas D' uses prices 1 and 2. They give the firm the same profit but leave consumers with different surplus. The pentagon thus reflects the possibility of designing disclosure around either markups or discounts, as well as using both. Appendix A.7 gives the vertex formulas and implementing segmentations, and proves completeness through a facet description.

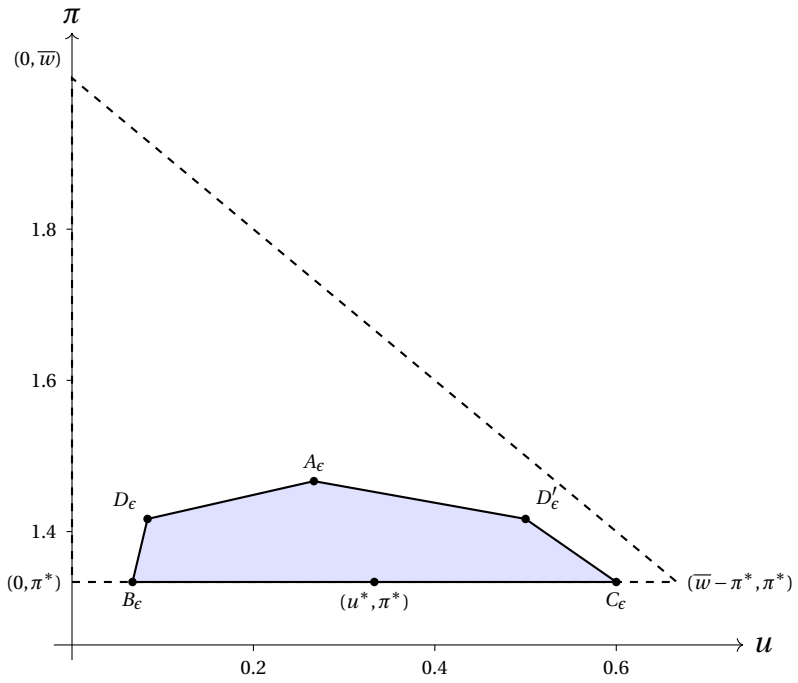


Figure 7: The welfare pentagon for $V = \{1,2,3\}$, $x^* = (1/3,1/3,1/3)$, and $\epsilon = \log 3$, inside the unrestricted BBM triangle

6. Conclusion

In this paper, we study the welfare effects of monopoly price discrimination when market segmentation must satisfy differential privacy. We characterize the attainable combinations of consumer and producer surplus and identify the privacy thresholds at which discounts and price increases become feasible.

As in BBM, every consumer-optimal segmentation keeps the firm at its uniform monopoly profit, although privacy regulation generally rules out efficient segmentation. Under the upper-tail bound, a segmentation can also minimize consumer surplus while keeping the firm's profit at this level. The resulting consumer loss relative to uniform pricing equals the value of forgone trades. Segmentation can therefore harm consumers even when the firm earns nothing more.

Discounts and price increases need not become feasible at the same privacy budget. If discounts become possible first, some budgets allow segmentation to benefit consumers while ruling out any consumer loss. If price increases come first, the initial departures from uniform pricing can only hurt consumers. Whether privacy regulation can preserve opportunities for discounts without allowing higher prices thus depends on the market.

Two features of the model suggest further questions. We abstract from consumers' intrinsic preferences for privacy. Evaluating a guarantee requires weighing its product-market consequences against the value of privacy itself, as in the accuracy–privacy tradeoff studied by Abowd and Schmutte (2019). We also apply the same likelihood-ratio bound to every pair of valuations. An extension could be protecting nearby valuations more strongly than distant ones, letting the firm distinguish consumers whose valuations are far apart.

References

ABOWD, J. M., R. ASHMEAD, R. CUMINGS-MENON, S. GARFINKEL, M. HEINECK, C. HEISS, R. JOHNS, D. KIFER, P. LECLERC, A. MACHANAVAJJHALA, B. MORAN, W. SEXTON, M. SPENCE, AND P. ZHURAVLEV (2022): "The 2020 Census Disclosure Avoidance System TopDown Algorithm," *Harvard Data Science Review*, Special Issue 2.

ABOWD, J. M., AND I. M. SCHMUTTE (2019): "An Economic Analysis of Privacy Protection and Statistical Accuracy as Social Choices," *American Economic Review*, 109(1), 171–202.

ACQUISTI, A., C. TAYLOR, AND L. WAGMAN (2016): "The Economics of Privacy," *Journal of Economic Literature*, 54(2), 442–492.

ALI, S. N., G. LEWIS, AND S. VASSERMAN (2023): "Voluntary disclosure and personalized pricing," *The Review of Economic Studies*, 90(2), 538–571.

APPLE (2017): "Learning with Privacy at Scale," Discussion paper, Apple Differential Privacy Team.

BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): "The Limits of Price Discrimination," *American Economic Review*, 105(3), 921–957.

BERGEMANN, D., AND S. MORRIS (2019): "Information Design: A Unified Perspective," *Journal of Economic Literature*, 57(1), 44–95.

BONATTI, A., AND G. CISTERNAS (2020): "Consumer Scores and Price Discrimination," *Review of Economic Studies*, 87(2), 750–791.

CHEN, X., S. MIAO, AND Y. WANG (2023): "Differential privacy in personalized pricing with nonparametric demand models," *Operations Research*, 71(2), 581–602.

DING, B., J. KULKARNI, AND S. YEKHANIN (2017): "Collecting Telemetry Data Privately," in *Advances in Neural Information Processing Systems 30*.

DUBÉ, J.-P., AND S. MISRA (2023): "Personalized Pricing and Consumer Welfare," *Journal of Political Economy*, 131(1), 131–189.

DWORK, C., F. MCSHERRY, K. NISSIM, AND A. SMITH (2006): "Calibrating Noise to Sensitivity in Private Data Analysis," in *Theory of Cryptography (TCC 2006)*, vol. 3876 of *Lecture Notes in Computer Science*, pp. 265–284. Springer.

EILAT, R., K. ELIAZ, AND X. MU (2021): "Bayesian Privacy," *Theoretical Economics*, 16(4), 1557–1603.

ELLIOTT, M., A. GALEOTTI, A. KOH, AND W. LI (2026): "Market segmentation through information," *Review of Economic Studies*, p. rdag067.

ERLINGSSON, Ú., V. PIHUR, AND A. KOROLOVA (2014): "RAPPOR: Randomized Aggregatable Privacy-Preserving Ordinal Response," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1054–1067.

FALLAH, A., M. I. JORDAN, A. MAKHDOUMI, AND A. MALEKIAN (2024): "The Limits of Price Discrimination Under Privacy Constraints," arxiv:2402.08223.

HAGHPANAH, N., AND R. SIEGEL (2022): "The Limits of Multiproduct Price Discrimination," *American Economic Review: Insights*, 4(4), 443–58.

HIDIR, S., AND N. VELLODI (2021): "Privacy, Personalization, and Price Discrimination," *Journal of the European Economic Association*, 19(2), 1342–1363.

ICHIHASHI, S. (2020): "Online Privacy and Information Disclosure by Consumers," *American Economic Review*, 110(2), 569–595.

KAIROUZ, P., S. OH, AND P. VISWANATH (2016): "Extremal mechanisms for local differential privacy," *Journal of Machine Learning Research*, 17(17), 1–51.

KAMENICA, E., AND M. GENTZKOW (2011): "Bayesian Persuasion," *American Economic Review*, 101(6), 2590–2615.

KASIVISWANATHAN, S. P., H. K. LEE, K. NISSIM, S. RASKHODNIKOVA, AND A. SMITH (2011): "What Can We Learn Privately?," *SIAM Journal on Computing*, 40(3), 793–826.

PAN, Y., Z. S. WU, H. XU, AND S. ZHENG (2025): "Differentially Private Bayesian Persuasion," in *Proceedings of the ACM on Web Conference 2025*, WWW '25, p. 1425–1440.

RHODES, A., AND J. ZHOU (2024): "Personalized Pricing and Competition," *American Economic Review*, 114(7), 2141–2170.

SCHMUTTE, I. M., AND N. YODER (2024): "Information Design for Differential Privacy," arXiv:2202.05452.

SHILLER, B. R. (2020): "Approximating Purchase Propensities and Reservation Prices from Broad Consumer Tracking," *International Economic Review*, 61(2), 847–870.

STRACK, P., AND K. H. YANG (2024): "Privacy-Preserving Signals," *Econometrica*, 92(6), 1907–1938.

TANEVA, I. (2019): "Information Design," *American Economic Journal: Microeconomics*, 11(4), 151–85.

TAYLOR, C. R. (2004): "Consumer Privacy and the Market for Customer Information," *RAND Journal of Economics*, 35(4), 631–650.

Appendix

A. Proofs and supporting results

A.1 Supporting results for the model

Implementation of private segmentations. Given $\sigma \in \Sigma(\epsilon)$, take $M = \text{supp}(\sigma)$ and define

$$\psi(x \mid v_k) = \frac{\sigma(x)x_k}{x_k^*}.$$

Bayes plausibility gives $\sum_x \psi(x \mid v_k) = 1$. Message x occurs with probability $\sigma(x)$ and induces posterior x. Since

$$\log \frac{\psi(x \mid v_{k'})}{\psi(x \mid v_k)} = \log \frac{x_{k'}}{x_k} - \log \frac{x_{k'}^*}{x_k^*},$$

condition (1) gives ϵ-DP. Conversely, Bayes' rule implies that every ϵ-DP information structure induces a segmentation in $\Sigma(\epsilon)$.

Lemma A.1 (Privacy budgets)**.** *$X(0) = \{x^*\}$, and $X(\epsilon_1) \subseteq X(\epsilon_2)$ whenever $\epsilon_1 \leq \epsilon_2$.*

Proof. At zero budget, (1) makes x proportional to x^*; normalization gives $x = x^*$. Increasing ϵ weakens every constraint. □

Lemma A.2 (Unique uniform price)**.** *Under Assumption 1, $R_{i^*}(x^*) > R_k(x^*)$ for every $k \neq i^*$.*

Proof. For $k \neq i^*$, $R_{i^*} \geq R_k$ holds on X^* and is strict at the market concentrated on v^*. Its equality set is therefore a proper face of X^* and cannot contain an interior point. □

Lemma A.3 (Feasible price regions)**.** *For finite ϵ, $X(\epsilon)$ and each nonempty $X_k(\epsilon)$ are polytopes consisting of full-support markets.*

Proof. The inequalities $x_i x_j^* \leq e^\epsilon x_j x_i^*$ define a bounded intersection of finitely many closed half-spaces within the simplex. A zero coordinate forces all coordinates to vanish, so every

feasible market has full support. Intersecting with the linear inequalities $R_k \geq R_\ell$ gives $X_k(\epsilon)$. □

Proof of Observation 2. Let $\widehat{\pi} = \max_{k \neq i^*} R_k(x^*) < \pi^*$ and $\underline{\epsilon} = \log(\pi^*/\widehat{\pi})$. For $x \in X(\epsilon)$, note that $\frac{\max_j x_j/x_j^*}{\min_j x_j/x_j^*} \leq e^\epsilon$, thus

$$\frac{R_k(x)}{R_{i^*}(x)} \leq \frac{\max_j x_j/x_j^*}{\min_j x_j/x_j^*} \frac{R_k(x^*)}{\pi^*} \leq e^\epsilon \frac{\widehat{\pi}}{\pi^*} < 1, \qquad k \neq i^*,\ \epsilon < \underline{\epsilon}.$$

Every segment therefore has the unique optimal price v^*. Linearity of R_{i^*} and U_{i^*} and Bayes plausibility give $(u, \pi) = (u^*, \pi^*)$. □

A.2 The binary benchmark

Proof of Proposition 1. Let $\ell = \alpha/(\alpha + e^\epsilon)$, $h = \alpha e^\epsilon/(1 + \alpha e^\epsilon)$, and $c = (v_2 - v_1)/v_2$. At $x_1 < c$, the unique optimal price is v_2, giving $(u, \pi) = (0, v_2(1 - x_1))$. At $x_1 > c$, it is v_1, giving $((v_2 - v_1)(1 - x_1), v_1)$. Both prices are optimal at c, and randomization fills the interval between their surplus pairs.

If $\epsilon < \epsilon_B$, $[\ell, h]$ lies strictly on one side of c. Both surpluses are affine there, so Bayes plausibility fixes their expectations at (u^*, π^*).

For $\epsilon > \epsilon_B$, $\ell < c < h$. The segment-level surplus pairs are affine on each side of c, so their lifted convex hull is generated by

$$P_\ell = (\ell, 0, v_2(1 - \ell)), \qquad P_c^H = (c, 0, v_1),$$
$$P_c^L = (c, (v_2 - v_1)(1 - c), v_1), \qquad P_h = (h, (v_2 - v_1)(1 - h), v_1),$$

with coordinates (x_1, u, π). Feasible outcomes are the slice at $x_1 = x_1^*$. For any two of these points with first coordinates $s < x_1^* < t$, the unique Bayes-plausible weights are

$$\lambda_s = \frac{t - x_1^*}{t - s}, \qquad \lambda_t = \frac{x_1^* - s}{t - s}.$$

If $x_1^* < c$, the pairs (P_ℓ, P_h), (P_ℓ, P_c^H), and (P_ℓ, P_c^L) give A_ϵ, B_ϵ, and C_ϵ, respectively. If $x_1^* > c$, the corresponding pairs are (P_ℓ, P_h), (P_c^H, P_h), and (P_c^L, P_h). Substituting these weights gives (2)–(5).

These pairs exhaust the slice vertices. Indeed, the mixing weights satisfy two equations: they sum to one and have mean x_1^*. An extreme feasible weight vector has at most two positive entries; with three or more, a nonzero perturbation preserves both equations and remains feasible in both directions. Since none of the four first coordinates equals x_1^*, every such pair must straddle x_1^*. The pairs listed above are the only ones that do. Their welfare images generate the entire slice. Above the threshold, $\pi_A^\epsilon > \pi^*$ and $u_B^\epsilon < u_C^\epsilon$, so all three points are vertices. At $\epsilon = \epsilon_B$, either $h = c$ and $P_h = P_c^L$, or $\ell = c$ and $P_\ell = P_c^H$. The same weights then give the stated segment and the coincident vertices.

For the total-surplus comparison, suppose $\epsilon > \epsilon_B$. High-valuation consumers always buy. If $v^* = v_1$, C_ϵ serves everyone, whereas A_ϵ excludes a positive mass of low-valuation consumers. If $v^* = v_2$, the masses of low-valuation consumers served at C_ϵ and A_ϵ are, respectively,

$$\frac{c(x_1^* - \ell)}{c - \ell} \quad \text{and} \quad \frac{h(x_1^* - \ell)}{h - \ell}.$$

The first is larger because $0 < \ell < c < h$. Thus $w(C_\epsilon) > w(A_\epsilon)$ in either case. Also, B_ϵ and C_ϵ have the same profit and $u_B^\epsilon < u_C^\epsilon$, so C_ϵ uniquely maximizes total surplus. At any finite budget, full support rules out $(0, \overline{w})$: zero consumer surplus requires price v_2 in every segment, excluding all low-valuation consumers. Hence $\mathcal{W}(\epsilon)$ is a strict subset of the unrestricted triangle. □

Proof of Corollary 1. Differentiating (2) gives

$$\frac{du_A^\epsilon}{d\epsilon} = -\frac{(v_2 - v_1)e^\epsilon}{(1+\alpha)(1+e^\epsilon)^2} < 0, \qquad \frac{d\pi_A^\epsilon}{d\epsilon} = \frac{(v_2 - v_1 + \alpha v_1)e^\epsilon}{(1+\alpha)(1+e^\epsilon)^2} > 0.$$

If $v^* = v_2$, $v_2 - v_1 > \alpha v_1$, and (4) fixes B_ϵ and gives

$$\frac{du_C^\epsilon}{d\epsilon} = \frac{\alpha v_1 (v_2 - v_1)(v_2 - v_1 - \alpha v_1)e^\epsilon}{(1+\alpha)((v_2 - v_1)e^\epsilon - \alpha v_1)^2} > 0.$$

If $v^* = v_1$, the inequality reverses, and (5) fixes C_ϵ and gives

$$\frac{du_B^\epsilon}{d\epsilon} = -\frac{\alpha v_1(v_2 - v_1)(\alpha v_1 - (v_2 - v_1))e^\epsilon}{(1+\alpha)(\alpha v_1 e^\epsilon - (v_2 - v_1))^2} < 0.$$

Equation (3) gives $dw(A_\epsilon)/d\epsilon > 0$. Taking limits in the vertex formulas yields the unrestricted triangle. □

A.3 Feasible markets and welfare outcomes

Proof of Proposition 2. Given a feasible segmentation and optimal pricing rule, collect consumers by the price they face:

$$q_{kj} = \sum_{x \in \mathrm{supp}(\sigma)} \sigma(x)\phi_k(x)x_j.$$

Bayes plausibility gives (6). Whenever $\phi_k(x) > 0$, $x \in X_k(\epsilon)$; summing its privacy and pricing inequalities with weights $\sigma(x)\phi_k(x)$ gives (7) and (8). The surplus formulas follow from the same summation.

Conversely, for each nonzero row of $q \in Q(\epsilon)$, set $s_k = \sum_j q_{kj}$ and $x_j^k = q_{kj}/s_k$. The constraints give $x^k \in X_k(\epsilon)$ and $\sum_k s_k x^k = x^*$. Assign weight s_k and price v_k to this market. If several rows give the same market, combine their weights and randomize over their prices in those proportions. This implements $(u(q), \pi(q))$ using at most one market–price pair per price. Finally, $0 \le q_{kj} \le x_j^*$, so $Q(\epsilon)$ is a compact polytope. Its linear welfare image is a compact convex polygon, possibly degenerate. □

Lemma A.4 (The unrestricted limit)**.** *The sets $\mathcal{W}(\epsilon)$ expand with ϵ and are strict subsets of the BBM triangle at every finite budget. Their limit satisfies*

$$\overline{\bigcup_{\epsilon<\infty} \mathcal{W}(\epsilon)} = \left\{(u,\pi) : u \ge 0,\ \pi \ge \pi^*,\ u + \pi \le \overline{w}\right\}.$$

Proof of Lemma A.4. Increasing ϵ weakens (7), so $Q(\epsilon)$ and its welfare image are nested. At a finite budget, every nonzero row has full support. Zero consumer surplus then requires price v_K in every row, excluding all lower valuations. Thus $(0, \overline{w})$ is infeasible, and the inclusion in the

BBM triangle is strict.

Let $Q(\infty)$ retain Bayes plausibility and optimal pricing but omit privacy. For each $k \neq i^*$, choose a full-support market close enough to the market concentrated on v_k that v_k remains strictly optimal. Give these markets small positive weights. The residual market is then close to x^*, so it has full support and v^* remains strictly optimal. This constructs $q^0 \in Q(\infty)$ with every entry positive. For any $q \in Q(\infty)$,

$$q^\delta = (1-\delta)q + \delta q^0, \qquad 0 < \delta < 1,$$

remains in $Q(\infty)$ and satisfies privacy for some finite budget because all its entries are positive. As $\delta \downarrow 0$, its welfare converges to that of q. Every unrestricted outcome is therefore a limit of outcomes under finite budgets; Theorem 0 gives the reverse inclusion. □

Lemma A.5 (Privacy vertices)**.** *For $\epsilon > 0$, the vertices of $X(\epsilon)$ are exactly the $2^K - 2$ markets $z^S(\epsilon)$ in (14).*

Proof of Lemma A.5. Write $\eta_j = x_j / x_j^*$. Feasibility requires $\sum_j x_j^* \eta_j = 1$ and $\eta_j \le e^\epsilon \eta_\ell$ for all j, ℓ, with every $\eta_j > 0$. Let $m = \min_j \eta_j$ and $M = \max_j \eta_j$. If $M < e^\epsilon m$, all privacy inequalities have strict slack, so the point is not a vertex.

If $M = e^\epsilon m$, let H and L contain the coordinates attaining M and m, and let I contain the others. A privacy inequality binds exactly when its left coordinate belongs to H and its right coordinate belongs to L: equality in $\eta_j \le M = e^\epsilon m \le e^\epsilon \eta_\ell$ forces both conditions. The active equalities therefore impose one common low multiplier, a high multiplier e^ϵ times as large, and leave the $|I|$ intermediate multipliers free. Normalization removes one degree of freedom, since the coefficient on the common low multiplier is

$$e^\epsilon \sum_{j \in H} x_j^* + \sum_{j \in L} x_j^* > 0.$$

If $I \neq \emptyset$, a nonzero direction remains. Small perturbations in both signs preserve normalization, positivity, the active equalities, and all inactive inequalities by their strict slack. The point is not a vertex. If $I = \emptyset$, the active equalities and normalization uniquely determine the point, making

it a vertex. With $S = H$, normalization gives

$$m = \frac{1}{1 + (e^{\epsilon} - 1)x^*(S)},$$

hence the point is $z^S(\epsilon)$. Every nonempty proper S yields a distinct feasible vertex. □

Let $\mathcal{V}_k(\epsilon)$ denote the vertices of $X_k(\epsilon)$ and $\mathcal{V}(\epsilon) = \bigcup_k \mathcal{V}_k(\epsilon)$. These *relative extreme points* are either privacy vertices or intersections at which price-indifference constraints bind. Linearity of U_k and R_k on $X_k(\epsilon)$ gives

$$\mathcal{W}(\epsilon) = \left\{ (u, \pi) : (x^*, u, \pi) \in \operatorname{co} \bigcup_k \{(x, U_k(x), R_k(x)) : x \in \mathcal{V}_k(\epsilon)\} \right\}. \tag{16}$$

Lemma A.6 (Generating welfare vertices)**.** *Every vertex of $\mathcal{W}(\epsilon)$ is generated by at most K market–price pairs with markets in $\mathcal{V}(\epsilon)$. They can be chosen with linearly independent market vectors, i.e., their weights are unique.*

Proof of Lemma A.6. Choose a linear objective $au + b\pi$ uniquely maximized at the given welfare vertex. By (16), maximize it over nonnegative weights on the finite collection of pairs (x, v_k) with $x \in \mathcal{V}_k(\epsilon)$, subject to

$$\sum_k \sum_{x \in \mathcal{V}_k(\epsilon)} \lambda_{kx} x = x^*.$$

These K equations also imply that the weights sum to one. A basic optimal solution selects at most K linearly independent market columns and must give the chosen welfare vertex. Bayes plausibility uniquely determines the selected weights. □

Unlike the direct representation, this vertex representation may use several markets for one price: pooling them preserves welfare but may move the posterior away from a relative extreme point.

A.4 Consumer-surplus endpoints

For unnormalized market masses, let

$$D(\epsilon) = \left\{ y \geq 0 : y_j x_\ell^* \leq e^\epsilon y_\ell x_j^* \text{ for all } j, \ell \right\},$$

which is a closed convex cone. Every nonzero private row has full support. Privacy bounds its largest likelihood multiplier by e^ϵ times its smallest; keeping all multipliers within their original range therefore preserves privacy. The profit-preserving assignments form

$$Q^*(\epsilon) = \left\{ q \in Q(\epsilon) : R_k(q_k) = R_{i^*}(q_k) \text{ for every } k \right\}.$$

Proof of Proposition 3. Bayes plausibility gives

$$\pi(\sigma, \phi) - \pi^* = \sum_x \sigma(x) \left[\max_k R_k(x) - R_i(x) \right].$$

Each summand is nonnegative, so equality holds exactly when v_i is optimal in every positive-weight segment. In direct form, these assignments constitute $Q^*(\epsilon)$, a nonempty compact convex face of $Q(\epsilon)$. Its linear image under u is the interval $[u^-, u^+]$.

For any nonzero private row y and $k < m$,

$$U_k(y) - U_m(y) = \sum_{j=k}^{m-1} (v_j - v_k) y_j + (v_m - v_k) \sum_{j=m}^{K} y_j > 0.$$

Hence a floor endpoint must select the highest tied price for the minimum and the lowest for the maximum; otherwise, repricing the same segmentation improves consumer surplus while preserving profit. □

Lemma A.7 (Prices at consumer-surplus extrema)**.** *Every q attaining $u_{\min}(\epsilon)$ has $q_k = 0$ for $k < i^*$. Every q attaining $u_{\max}(\epsilon)$ has $q_k = 0$ for $k > i^*$.*

Proof. Let $H = \sum_{j \geq i^*} x_j^*$ and $H_r = \sum_{j \geq i^*} q_{rj}$. The two relaxations below retain Bayes plausibility and privacy. Both are compact and contain the uninformative assignment.

For the minimum, set rows below i^* to zero and require only $q_{i^*} \in D(\epsilon)$; higher rows remain obedient. Pooling all rows $k \le i^*$ into row i^* satisfies this relaxation and weakly lowers consumer surplus, strictly if a lower row is nonzero.

Take a relaxed minimizer. Keep its tail cells and set

$$\widehat{q}_{rj} = x^*_j H_r / H, \qquad j < i^*.$$

Column sums and consumer surplus are unchanged. Each new multiplier H_r/H averages the unchanged tail multipliers, so privacy is preserved. Every nonzero row satisfies

$$R_\ell(\widehat{q}_r) = \frac{H_r}{H} R_\ell(x^*) < \frac{H_r}{H} R_{i^*}(x^*) = R_{i^*}(\widehat{q}_r), \qquad \ell < i^*.$$

Higher rows remain obedient because their tail revenues are unchanged. If row i^* preferred a higher price, repricing and merging it into that price's row would be feasible and strictly lower consumer surplus, contradicting the relaxed minimum. Thus the relaxation is exact. The strict initial pooling inequality excludes lower rows from every original minimizer.

For the maximum, set rows above i^* to zero, retain obedience below i^*, and require only $q_{i^*} \in D(\epsilon)$. Pooling rows $k \ge i^*$ into row i^* weakly raises consumer surplus, strictly if a higher row is nonzero. In a relaxed maximizer, keep the cells below i^* and set

$$\widehat{q}_{rj} = x^*_j H_r / H, \qquad j \ge i^*.$$

This averaging preserves privacy, column sums, and consumer surplus, since

$$u(q) = \sum_{r \le i^*} \sum_{j=r}^{i^*-1} (v_j - v_r) q_{rj} + \sum_{j \ge i^*} v_j x^*_j - \sum_{r \le i^*} v_r H_r.$$

Revenues at prices at most v^*_i are unchanged, while every nonzero row satisfies

$$R_h(\widehat{q}_r) = \frac{H_r}{H} R_h(x^*) < \frac{H_r}{H} R_{i^*}(x^*) = R_{i^*}(\widehat{q}_r), \qquad h > i^*.$$

If row i^* preferred a lower price, repricing and merging it there would improve on the relaxed maximum. Hence this relaxation is also exact, and no original maximizer uses a higher row. □

Lemma A.8 (Pooling at an extremum)**.** *At a consumer-surplus minimum, the firm charges the highest optimal price in each segment. If positive segments x, x' receive prices $v_k < v_m$, some $v_h < v_m$ is also optimal in x' and satisfies $R_h(x) > R_m(x)$.*

At a consumer-surplus maximum, the firm charges the lowest optimal price. For segments priced at $v_k < v_m$, some $v_h > v_k$ is optimal in the lower-priced segment and satisfies $R_h(x') > R_k(x')$.

Proof. Strict monotonicity of U_k gives the tie-breaking claims. At a minimum, suppose every lower price tied with v_m in x' satisfied $R_h(x) \le R_m(x)$. Then v_m remains optimal after adding a sufficiently small amount of x to x': tied lower prices do not overtake it, and all other prices have strict slack, including higher prices by tie-breaking. Convexity preserves privacy, and transferring mass between these segments preserves Bayes plausibility. Charging the transferred consumers v_m strictly lowers their surplus, a contradiction. At a maximum, the same argument transfers mass from x' to x; tied higher prices cannot overtake v_k unless one satisfies the stated inequality, and all other prices have strict slack. □

Lemma A.9 (Directional reduction)**.** *Let $H = \sum_{j=i^*}^{K} x_j^*$. The tail market*

$$V_T = (v_i, \dots, v_K), \qquad x_T = (x_i^*/H, \dots, x_K^*/H)$$

has unique monopoly price v_i and satisfies

$$u_{\min}(\epsilon) = H u_{\min,T}(\epsilon), \qquad u^-(\epsilon) = H u_T^-(\epsilon).$$

The collapsed market

$$V_C = (v_1, \dots, v^*), \qquad x_C = (x_1^*, \dots, x_{i^*-1}^*, H)$$

also has unique monopoly price v^* and satisfies

$$u_{\max}(\epsilon) = u^* + u_{\max,C}(\epsilon), \qquad u^+(\epsilon) = u^* + u_C^+(\epsilon).$$

Both reductions keep the same privacy budget. Their active prices are, respectively, the original active prices at or above v^* and at or below v^*.

Proof of Lemma A.9. The tail market scales aggregate revenues at prices $v_r \geq v^*$ by $1/H$; the collapsed market leaves revenues at prices $v_r \leq v^*$ unchanged. Both retain the unique monopoly price v^*.

By Lemma A.7, a consumer-surplus minimizer uses only rows $r \geq i^*$. For any feasible assignment with this support, the restriction

$$\bar{q}_{rj} = q_{rj}/H, \qquad r, j \geq i^*,$$

preserves Bayes plausibility and obedience in the tail market. Likelihood multipliers are unchanged, so privacy holds, and $u(q) = Hu_T(\bar{q})$.

Conversely, given a feasible tail assignment $\bar{q}$, let $s_r = \sum_{j=i^*}^{K} \bar{q}_{rj}$ and define

$$q_{rj} = H\bar{q}_{rj} \quad (j \geq i^*), \qquad q_{rj} = x^*_j s_r \quad (j < i^*), \qquad r \geq i^*,$$

with zero rows below i^*. Column sums equal x^*, and each added multiplier s_r averages the row's tail multipliers. Thus privacy holds. The added prices are strictly worse in every nonzero row:

$$R_\ell(q_r) = s_r R_\ell(x^*) < s_r R_{i^*}(x^*) = R_{i^*}(q_r) \leq R_r(q_r), \qquad \ell < i^*.$$

These maps prove $u_{\min} = Hu_{\min,T}$ and preserve profit gaps up to scale:

$$R_r(q_r) - R_{i^*}(q_r) = H\big[R_{r,T}(\bar{q}_r) - R_{i^*,T}(\bar{q}_r)\big].$$

The floor minimum also uses only prices at least v^* by Proposition 3, so the same maps give $u^- = Hu^-_T$.

For the maximum, consider a feasible assignment with zero rows above i^* and collapse each

tail into

$$\widetilde{q}_{ri^*} = \sum_{j=i^*}^{K} q_{rj},$$

retaining cells below i^*. The multiplier $\widetilde{q}_{ri^*}/H$ averages the tail multipliers, so privacy holds. Column sums and revenues at prices at most v^* are preserved, and $u(q) = B + u_C(\widetilde{q})$.

To lift a feasible collapsed assignment, retain its cells below i^* and set $q_{rj} = (x^*_j/H)\widetilde{q}_{ri^*}$ for $j \geq i^*$. Privacy, column sums, and lower-price revenues are preserved; every higher price is strictly worse in a nonzero row:

$$R_h(q_r) = \frac{\widetilde{q}_{ri^*}}{H} R_h(x^*) < \frac{\widetilde{q}_{ri^*}}{H} R_{i^*}(x^*) = R_{i^*}(q_r) \leq R_r(q_r), \qquad h > i^*.$$

Lemma A.7 now gives $u_{\max} = B + u_{\max,C}$. The maps preserve each row's profit gap, and the floor maximum uses only prices at most v_i by Proposition 3; hence $u^+ = B + u^+_C$.

The same restrictions and lifts apply to individual nonzero rows without imposing column sums. They preserve privacy and optimality of the relevant price, so normalization gives both active-price identities. □

Proof of Proposition 4. Lemmas A.7 and A.9 map every consumer-surplus maximizer to one in the collapsed market, preserving producer surplus. It suffices to consider $v^* = v_K$; the case $K = 1$ is immediate.

Fix any maximizer $\bar{q}$. Among maximizers with producer surplus $\pi(\bar{q})$, choose q to minimize

$$\mathcal{P}(q) = \sum_{r,j} \frac{q^2_{rj}}{x^*_j}.$$

The constraint set is nonempty and compact. In every nonzero row,

$$R_r(q_r) > R_h(q_r), \qquad h < r. \tag{17}$$

Otherwise, repricing at a tied lower price and merging with that price's row preserves feasibility and strictly raises consumer surplus.

Call $R_k(q_k) - R_K(q_k)$ the profit gap in row k. We prove by descending induction on $j = K-1, \ldots, 1$ that every row with a positive gap satisfies

$$\frac{q_{kj}}{x_j^*} \leq \frac{q_{kK}}{x_K^*}. \tag{18}$$

Assume the inequalities for indices above j (an empty hypothesis when $j = K-1$). For a positive-gap row k and $j < h < K$,

$$R_h(q_k) \leq \frac{q_{kK}}{x_K^*} R_h(x^*) < \frac{q_{kK}}{x_K^*} R_K(x^*) = R_K(q_k) < R_k(q_k).$$

The prior's unique monopoly price gives the first strict inequality. Full support gives $q_{kK}/x_K^* > 0$. Price v_K is also strictly worse by the positive gap. Hence $k \leq j$, and every price above v_j is strictly suboptimal in row k.

Suppose $q_{kj}/x_j^* > q_{kK}/x_K^*$. Since $\sum_r q_{rj}/x_j^* = \sum_r q_{rK}/x_K^* = 1$, another row r satisfies $q_{rj}/x_j^* < q_{rK}/x_K^*$. Change four cells:

$$\begin{aligned} q'_{kj} &= q_{kj} - \delta, & q'_{kK} &= q_{kK} + \delta, \\ q'_{rj} &= q_{rj} + \delta, & q'_{rK} &= q_{rK} - \delta. \end{aligned}$$

Column sums are unchanged. Choose

$$0 < \delta < \frac{\min\left\{q_{kj}/x_j^* - q_{kK}/x_K^*,\, q_{rK}/x_K^* - q_{rj}/x_j^*\right\}}{{x_j^*}^{-1} + {x_K^*}^{-1}}.$$

In each changed row, the two multipliers move toward one another without crossing. Let $m_t = \min_\ell q_{t\ell}/x_\ell^* > 0$ and $M_t = \max_\ell q_{t\ell}/x_\ell^*$ for $t \in \{k, r\}$. Every new multiplier lies in $[m_t, M_t]$, so all ordered pairs satisfy

$$q'_{t\ell}/x_\ell^* \leq M_t \leq e^\epsilon m_t \leq e^\epsilon q'_{th}/x_h^*.$$

Thus privacy and nonnegativity hold in both changed rows.

In row k, revenues at prices at most v_j are unchanged; higher-price revenues increase by $v_h\delta$.

Their strict slack, established above, preserves obedience for small δ.

In row r, lower-price revenues are unchanged and revenues above v_j fall by $v_h\delta$. If $r \le j$, obedience follows directly. If $r > j$, write $G_h = R_r(q_r) - R_h(q_r)$. The new gaps are

$$R_r(q'_r) - R_h(q'_r) = \begin{cases} G_h - v_r\delta, & h \le j, \\ G_h - (v_r - v_h)\delta, & j < h < r, \\ G_h + (v_h - v_r)\delta, & h > r. \end{cases}$$

The first two cases have $G_h > 0$ by (17); the last expression is nonnegative. There are finitely many comparisons, so one sufficiently small positive δ preserves obedience and privacy in both rows.

Since $k \le j$, consumer surplus changes by

$$u(q') - u(q) = \begin{cases} 0, & r \le j, \\ (v_r - v_j)\delta > 0, & r > j. \end{cases}$$

The second case contradicts consumer optimality. In the first, both assigned revenues are unchanged, so q' also has producer surplus $\pi(\bar{q})$. But

$$\left.\frac{d\mathcal{P}(q')}{d\delta}\right|_{\delta=0} = 2\left[(q_{kK}/x^*_K - q_{kj}/x^*_j) + (q_{rj}/x^*_j - q_{rK}/x^*_K)\right] < 0,$$

contradicting the choice of q for sufficiently small δ. This proves (18) and completes the induction.

A positive-gap row $k < K$ would therefore satisfy

$$R_k(q_k) \le \frac{q_{kK}}{x^*_K} R_k(x^*) < \frac{q_{kK}}{x^*_K} R_K(x^*) = R_K(q_k),$$

contrary to obedience. Row K has zero gap, so all rows preserve profit and $\pi(q) = \pi^*$. Because $\pi(q) = \pi(\bar{q})$ and $\bar{q}$ was arbitrary, every maximizer preserves profit. The reduction gives the result in the original market, including $u_{\max}(\epsilon) = u^+(\epsilon)$. □

Proof of Proposition 5. Lemma A.9 reduces the problem to the tail market with $v^* = v_1$. If no higher price is active, only v_1 can be charged and $u_{\min} = u^- = u^*$. Otherwise, relabel the tail. The reduction preserves active prices and scales the assumed bound to $v_b \sum_{j=a}^{K} x_j^* \le v_1$.

Among consumer-surplus minimizers, minimize producer surplus, then minimize

$$I(q) = \sum_{r:s_r>0} \sum_j q_{rj} \log \frac{q_{rj}}{x_j^* s_r}, \qquad s_r = \sum_j q_{rj}.$$

These minima exist by compactness: both surpluses are linear, and

$$I(q) = \sum_{r,j} q_{rj} \log(q_{rj}/x_j^*) - \sum_r s_r \log s_r$$

is continuous with $0 \log 0 = 0$.

Suppose row k has a positive profit gap, $R_k(q_k) > R_1(q_k)$. We claim

$$\frac{q_{kj}}{x_j^*} \ge s_k, \qquad j < a.$$

Otherwise, $\sum_r q_{rj}/x_j^* = \sum_r s_r = 1$ supplies a donor row r with $q_{rj}/x_j^* > s_r$. Transfer a small mass $\delta > 0$ of type j from r to k. The recipient's multiplier rises from below its row average and the donor's falls from above its average. Choose δ so both remain within their original row's multiplier range; this preserves every privacy inequality and nonnegativity. Column sums are unchanged.

Since $k \ge a > j$, recipient revenue is unchanged. Only competitors at or below v_j gain revenue: v_1 has strict slack by the positive gap, and every other such price is inactive, hence strictly suboptimal in any nonzero private row. Small δ preserves obedience.

If the donor charges a price at least v_a, its assigned revenue is also unchanged and competing revenues weakly fall. Both surpluses stay fixed, but

$$\left.\frac{dI}{d\delta}\right|_{\delta=0} = \log \frac{q_{kj}}{x_j^* s_k} - \log \frac{q_{rj}}{x_j^* s_r} < 0,$$

contradicting the choice of q. The only other possible donor price is v_1. It is uniquely optimal in that row by Lemma A.8, so a small transfer again preserves obedience. For $j > 1$, consumer surplus falls by $(v_j - v_1)\delta$; for $j = 1$, it is unchanged and profit falls by $v_1\delta$. Either contradicts the successive minimizations, proving the claim.

The claim implies

$$\sum_{j=a}^{K} q_{kj} \le s_k \sum_{j=a}^{K} x_j^*.$$

If at least two higher prices are active, every one except the highest has $a \le h \le b$, so

$$R_h(q_k) \le v_b \sum_{j=a}^{K} q_{kj} \le v_b s_k \sum_{j=a}^{K} x_j^* \le v_1 s_k = R_1(q_k) < R_k(q_k).$$

Price v_1 and every inactive price are also strictly worse. Thus v_k is the highest active price and is uniquely optimal in row k. With only one higher active price, the positive gap gives this conclusion directly.

Some other nonzero row must charge less than v_k; otherwise, summing obedience inequalities makes v_k optimal in the prior, contrary to $v^* = v_1$. Lemma A.8 then requires a lower price to tie v_k in row k, contradicting its unique optimality.

Hence every row of the selected minimizer preserves profit, giving $u_{\min} = u^-$. The tail-market lift gives the equality in the original market. □

Proof of Proposition 6. Every nonzero row contains every valuation. Thus zero consumer surplus requires charging v_K in every row. Aggregating obedience inequalities makes v_K optimal in the prior, so $v^* = v_K$. Conversely, the uninformative segmentation attains zero consumer surplus when $v^* = v_K$.

First-best total surplus requires charging v_1 in every row, since any higher price excludes some consumers. Aggregation then gives $v^* = v_1$. Also, $u \le \overline{w} - \pi^*$, with equality only if total surplus is first-best and profit is π^*. The uninformative segmentation attains both when $v^* = v_1$, proving the total-surplus and maximum-consumer-surplus claims.

Finally, $\pi = \overline{w}$ requires both zero consumer surplus and first-best total surplus, hence every row charging both v_K and v_1. This is impossible for $K \geq 2$; compactness makes the bound on maximum producer surplus strict. □

A.5 Price-activation thresholds

Lemma A.10 (Computing activation thresholds)**.** *For each k, the minimum in (11) is attained and*

$$e^{\epsilon_k^{\text{act}}} = \min_{\tau, r} \left\{ \tau : \begin{array}{ll} 1 \leq r_j \leq \tau & \text{for every } j, \\ v_k \sum_{j=k}^{K} x_j^* r_j \geq v_\ell \sum_{j=\ell}^{K} x_j^* r_j & \text{for every } \ell \end{array} \right\}. \tag{19}$$

The uniform price has threshold zero; all other thresholds are positive and finite. Price v_k is active exactly when $\epsilon \geq \epsilon_k^{\text{act}}$.

Proof of Lemma A.10. For any full-support $x \in X_k$, setting

$$r_j = \frac{x_j / x_j^*}{\min_k (x_k / x_k^*)}$$

makes $(e^{\rho(x)}, r)$ feasible in (19). Conversely, a feasible (τ, r) induces

$$x_j = \frac{x_j^* r_j}{\sum_k x_k^* r_k} \in X_k, \qquad \rho(x) \leq \log \tau.$$

Thus the two programs have the same infimum after exponentiation. A market sufficiently close to the point mass at v_k, mixed with a positive weight on x^*, has full support and keeps v_k strictly optimal. The linear program is therefore feasible at some finite τ_0. Its feasible set with $\tau \leq \tau_0$ is compact, so both minima are attained.

Since $\rho(x) = 0$ only at x^*, uniqueness of the uniform price gives $\epsilon_{i^*}^{\text{act}} = 0$ and $\epsilon_k^{\text{act}} > 0$ for $k \neq i^*$. Attainment and (10) give the activation equivalence, including equality at the threshold. □

Proof of Proposition 7. At a full-support market, $U_k(x)$ strictly decreases with k. If no higher price is active, every feasible outcome has $u \geq u^*$; if no lower price is active, every outcome has

$u \le u^*$. The uninformative segmentation attains u^*.

Any active nonuniform price can change consumer surplus. Take $x \in X_k(\epsilon)$ with $k \neq i^*$, so $\epsilon > 0$. At x^*, all nontrivial privacy inequalities and all revenue comparisons favoring v^* are strict. Hence, for sufficiently small $\lambda > 0$,

$$y_\lambda = \frac{x^* - \lambda x}{1 - \lambda}$$

has full support, satisfies privacy, and has unique optimal price v^*: it converges to x^*, and there are only finitely many inequalities to preserve. Give x weight λ at price v_k and y_λ weight $1 - \lambda$ at price v^*. Bayes plausibility then gives

$$u - u^* = \lambda[U_k(x) - U_{i^*}(x)],$$

which is negative for $k > i^*$ and positive for $k < i^*$. This proves the claims for $u_{\min}$ and $u_{\max}$.

To obtain a loss on the profit floor, take an active higher-price market $x \in X_k(\epsilon)$, $k > i^*$. Recall that $H = \sum_{j \ge i^*} x_j^*$, let $T = \sum_{j \ge i^*} x_j$, and set

$$x_{T,j} = x_j^* / H, \qquad z_j = x_j / T, \qquad j \ge i^*.$$

The likelihood multipliers of z relative to x_T are $(H/T)(x_j / x_j^*)$, so z satisfies tail-market privacy. Since tail revenues equal the corresponding original revenues divided by T, v_k remains optimal among tail prices.

The path $z(t) = (1 - t)x_T + tz$ stays in the tail privacy polytope by convexity. Define

$$g(t) = \max_{h > i^*} \{R_{h,T}(z(t)) - v^*\}.$$

The lowest tail price earns v^* throughout. Uniqueness of the uniform price gives $g(0) < 0$, while optimality of v_k at z gives $g(1) \ge 0$. Continuity ensures that

$$t_0 = \min\{t \in [0,1] : g(t) \ge 0\} \in (0,1]\}, \qquad g(t_0) = 0.$$

At $z(t_0)$, every higher price earns at most v^* and some $v_h > v^*$ earns exactly v^*.

Lift this market by setting

$$\widehat{x}_j = \begin{cases} x^*_j, & j < i^*, \\ Hz_j(t_0), & j \geq i^*. \end{cases}$$

The tail multipliers $b_j = z_j(t_0)/x_{T,j}$ average to one. If $m = \min b_j$ and $M = \max b_j$, then $m \leq 1 \leq M \leq e^{\epsilon} m$. The added lower-type multipliers equal one, so every original multiplier lies in $[m, M]$ and every pair satisfies privacy. Higher-price revenues are H times their tail revenues; lower-price revenues satisfy

$$R_\ell(\widehat{x}) = R_\ell(x^*) < R_i(x^*) = R_{i^*}(\widehat{x}), \qquad \ell < i^*.$$

Thus $\widehat{x} \in X_{i^*}(\epsilon) \cap X_h(\epsilon)$: the tie is optimal against every original price.

Use the residual construction with $\widehat{x}$, charging v_h there and v^* in the residual market. Both segments keep v^* optimal, so profit is π^*, while

$$u - u^* = \lambda[U_h(\widehat{x}) - U_{i^*}(\widehat{x})] < 0.$$

Hence an active higher price implies $u^- < u^*$. The converse follows from $u^- \geq u_{\min}$. The threshold formulas now follow from (12) and (13); Proposition 4 also gives $u^+ = u_{\max}$. □

A.6 Producer-optimal segmentation

Proof of Proposition 8. At $\epsilon = 0$, all $z^S(0) = x^*$ and one market suffices. For $\epsilon > 0$, monopoly profit $f(x) = \max_k R_k(x)$ is convex. By Lemma A.5, write each private market as $x = \sum_S \mu_x(S) z^S(\epsilon)$. Replacing that segment with its vertex decomposition preserves Bayes plausibility and weakly raises profit:

$$f(x) \leq \sum_S \mu_x(S) f\big(z^S(\epsilon)\big).$$

The stated linear program therefore attains maximum producer surplus. Summing its Bayes-plausibility equations gives $\sum_S \lambda_S = 1$, so its feasible set is compact. A basic optimal solution uses at most K markets, the number of equations.

The upper boundary of a compact convex polygon is concave and piecewise linear, with finitely many kinks generated as in Lemma A.6. If $\overline{\pi}_\epsilon(u_A) > \pi^*$ somewhere, mix that point with an outcome at the opposite endpoint of the consumer-surplus range to reach any other interior u. Every feasible outcome has profit at least π^*, so this mixture has profit strictly above π^*. In the binary case, Proposition 1 gives the unique kink A_ϵ above the activation threshold. □

A.7 The three-valuation example

Return to $V = \{1,2,3\}$, $x^* = (1/3, 1/3, 1/3)$. Both discriminatory prices activate at $\epsilon = \log 2$. For $\epsilon > \log 2$, the welfare polygon has vertices

$$B_\epsilon = \left(\frac{1}{3(2e^\epsilon - 1)}, \frac{4}{3}\right), \qquad C_\epsilon = \left(\frac{4e^\epsilon - 3}{3(2e^\epsilon - 1)}, \frac{4}{3}\right),$$
$$D_\epsilon = \left(\frac{1}{3(e^\epsilon + 1)}, \frac{5e^\epsilon + 2}{3(e^\epsilon + 1)}\right), \qquad D'_\epsilon = \left(\frac{e^\epsilon + 3}{3(e^\epsilon + 1)}, \frac{5e^\epsilon + 2}{3(e^\epsilon + 1)}\right),$$
$$A_\epsilon = \left(\frac{4}{3(e^\epsilon + 2)}, \frac{6e^\epsilon + 4}{3(e^\epsilon + 2)}\right).$$

Write the rows of a feasible assignment as $a_j = q_{1j}$, $b_j = q_{2j}$, and $c_j = q_{3j}$. Thus $a_j + b_j + c_j = 1/3$, and

$$u = a_2 + 2a_3 + b_3, \qquad \pi = a_1 + a_2 + a_3 + 2b_2 + 2b_3 + 3c_3.$$

We show that the welfare set is exactly the intersection of the half-planes

$$\pi \geq \frac{4}{3}, \tag{20}$$

$$\pi - (2e^{\epsilon} - 1)u \leq 1, \tag{21}$$

$$(3e^{\epsilon} + 2)\pi - e^{\epsilon}(e^{\epsilon} - 2)u \leq \frac{14e^{\epsilon} + 4}{3}, \tag{22}$$

$$(e^{2\epsilon} + e^{\epsilon} + 2)\pi + e^{\epsilon}(e^{\epsilon} - 2)u \leq \frac{6e^{2\epsilon} + 2e^{\epsilon} + 4}{3}, \tag{23}$$

$$2e^{\epsilon}\pi + (2e^{\epsilon} - 1)u \leq 4e^{\epsilon} - 1. \tag{24}$$

Optimality against price 2 gives the profit floor:

$$\pi - \frac{4}{3} = (a_1 - a_2 - a_3) + (c_3 - 2c_2) \geq 0.$$

For the remaining inequalities, set

$$d = \frac{e^{\epsilon} + 2}{e^{\epsilon}(e^{\epsilon} - 1)}, \qquad \kappa = \frac{e^{2\epsilon} - 4}{e^{\epsilon} - 1}.$$

For $\epsilon \geq \log 2$, we have $0 < d \leq 2$ and $\kappa \geq 0$. Bayes plausibility and the surplus formulas give the following identities for the four upper-facet slacks:

$$
\begin{aligned}
&1 - \pi + (2e^{\epsilon} - 1)u \\
&\quad = (e^{\epsilon} a_3 - a_1) + (2e^{\epsilon} - 2)a_2 + e^{\epsilon} a_3 + 2(e^{\epsilon} b_3 - b_2), \\
&\frac{14e^{\epsilon} + 4}{3} - (3e^{\epsilon} + 2)\pi + e^{\epsilon}(e^{\epsilon} - 2)u \\
&\quad = (e^{\epsilon} + d)(e^{\epsilon} a_2 - a_1) + (2e^{\epsilon} + 2 - d)(e^{\epsilon} a_3 - a_1) \\
&\qquad + \kappa(e^{\epsilon} b_3 - b_2) + \left(5 + \frac{3}{e^{\epsilon} - 1}\right)(e^{\epsilon} c_2 - c_3), \\
&\frac{6e^{2\epsilon} + 2e^{\epsilon} + 4}{3} - (e^{2\epsilon} + e^{\epsilon} + 2)\pi - e^{\epsilon}(e^{\epsilon} - 2)u \\
&\quad = (2 + d)(e^{\epsilon} a_2 - a_1) + \left(3 + \frac{2}{e^{\epsilon}}\right)(e^{\epsilon} a_3 - a_1) \\
&\qquad + \kappa(e^{\epsilon} b_1 - b_2) + \kappa(e^{\epsilon} c_1 - c_3) + e^{\epsilon}(2 + d)(e^{\epsilon} c_2 - c_3), \\
&4e^{\epsilon} - 1 - 2e^{\epsilon}\pi - (2e^{\epsilon} - 1)u \\
&\quad = (e^{\epsilon} b_1 - b_2) + (e^{\epsilon} b_1 - b_3) + 2(e^{\epsilon} c_1 - c_3) + (4e^{\epsilon} - 1)c_2.
\end{aligned}
$$

Privacy and nonnegativity make every term on the right nonnegative, proving (20)–(24).

For $\epsilon > \log 2$, the four upper boundary lines have strictly decreasing slopes

$$
2e^{\epsilon} - 1 > \frac{e^{\epsilon}(e^{\epsilon} - 2)}{3e^{\epsilon} + 2} > 0 > -\frac{e^{\epsilon}(e^{\epsilon} - 2)}{e^{2\epsilon} + e^{\epsilon} + 2} > -\frac{2e^{\epsilon} - 1}{2e^{\epsilon}}.
$$

The consecutive intersections, including those with the profit floor, are $B_\epsilon, D_\epsilon, A_\epsilon, D'_\epsilon, C_\epsilon$, in strictly increasing order of consumer surplus. The five half-planes therefore intersect in exactly the proposed pentagon.

The following assignments attain its vertices; rows correspond to prices 1, 2, and 3:

$$q^B = \frac{1}{3(2e^\epsilon - 1)} \begin{pmatrix} 0 & 0 & 0 \\ e^\epsilon & e^\epsilon & 1 \\ e^\epsilon - 1 & e^\epsilon - 1 & 2(e^\epsilon - 1) \end{pmatrix}, \qquad q^C = \frac{1}{3(2e^\epsilon - 1)} \begin{pmatrix} 2(e^\epsilon - 1) & e^\epsilon - 1 & e^\epsilon - 1 \\ 1 & e^\epsilon & e^\epsilon \\ 0 & 0 & 0 \end{pmatrix},$$

$$q^D = \frac{1}{3(e^\epsilon + 1)} \begin{pmatrix} 0 & 0 & 0 \\ e^\epsilon & e^\epsilon & 1 \\ 1 & 1 & e^\epsilon \end{pmatrix}, \qquad q^{D'} = \frac{1}{3(e^\epsilon + 1)} \begin{pmatrix} e^\epsilon & 1 & 1 \\ 1 & e^\epsilon & e^\epsilon \\ 0 & 0 & 0 \end{pmatrix},$$

$$q^A = \frac{1}{3(e^\epsilon + 2)} \begin{pmatrix} e^\epsilon & 1 & 1 \\ 1 & e^\epsilon & 1 \\ 1 & 1 & e^\epsilon \end{pmatrix}.$$

Every column sums to $1/3$, and each nonzero row has maximum-to-minimum ratio e^ϵ or 2, so privacy holds for $\epsilon \geq \log 2$. Optimality at price 1 requires $a_1 \geq a_2 + a_3$ and $a_1 + a_2 \geq 2a_3$; at price 2, it requires $b_2 + b_3 \geq b_1$ and $2b_2 \geq b_3$; and at price 3, it requires $2c_3 \geq c_1 + c_2$ and $c_3 \geq 2c_2$. The displayed assignments satisfy these inequalities for $\epsilon \geq \log 2$ and yield the stated surplus pairs. Convexity makes the whole pentagon attainable.

In q^B, the price-3 row has posterior $(1/4, 1/4, 1/2)$ and weight $4(e^\epsilon - 1)/(3(2e^\epsilon - 1))$; the remaining row is $z^{\{1,2\}}(\epsilon)$ at price 2. In q^C, the same weight is placed on $(1/2, 1/4, 1/4)$ at price 1, with the remainder on $z^{\{2,3\}}(\epsilon)$ at price 2. At each indifference market the nonuniform price ties 2. The other assignments use only privacy vertices.

At $\epsilon = \log 2$, the same valid inequalities impose $\pi = 4/3$ and $1/9 \leq u \leq 5/9$; q^B and q^C attain the endpoints. For $\epsilon < \log 2$, a nonzero price-1 row would require $a_1 \geq a_2 + a_3 \geq 2\min\{a_2, a_3\}$, contrary to $a_1 \leq e^\epsilon \min\{a_2, a_3\}$. Likewise, a nonzero price-3 row would require $c_3 \geq 2c_2$, contrary to $c_3 \leq e^\epsilon c_2$. Hence only price 2 can be charged and the welfare set is $\{(1/3, 4/3)\}$. As $\epsilon \to \infty$, B_ϵ, C_ϵ, and A_ϵ approach the vertices of the BBM triangle, while D_ϵ and D'_ϵ approach $(0, 5/3)$ and $(1/3, 5/3)$ on its edges.