Discussion Paper Series – CRC TR 224

Discussion Paper No. 798
Project B 04

Privacy in Search Markets

Laurenz Marstaller[1]

November 2026

[1]University of Bonn, Email: laurenz.marstaller@uni-bonn.de

Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Privacy in Search Markets*

Laurenz Marstaller†

University of Bonn

January 28, 2026

Abstract

Heterogeneous search costs enable price discrimination, which I study in the canonical Wolinsky (1986) sequential search setting. Firms observe a public signal of a consumer's search cost before posting a personalized price. The welfare effects of search-cost-based price discrimination depend on the distribution of search costs. For sufficiently small search costs, all consumers participate, and price discrimination reduces consumer surplus. When search costs are sufficiently dispersed, price discrimination reduces participation; its effect on consumer surplus is ambiguous and decomposed into three forces. This decomposition guides optimal information design: the consumer-surplus-maximizing policy is a binary signal that separates low- and high-search-cost consumers.

Privacy, Price Discrimination, Search Costs, Consumer Search Market

D83, L13, D18, D43

*I would like to thank my advisors, Stephan Lauermann, Andreas Kleiner, and Jidong Zhou, for their guidance. I am also grateful to Marina Halac for hosting me at Yale University, where part of this work was written. I have benefited from extended discussions with Carl-Christian Groh.

I also thank the participants of various conferences for their thoughtful comments and questions. This research was supported in part by the Bonn Graduate School of Economics (BGSE) and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under both Germany's Excellence Strategy – EXC 2126/1–390838866 and CRC TR 224 (Project B04).

†laurenz.marstaller@uni-bonn.de

1 Introduction

Virtually every online interaction leaves a digital footprint of clicks, searches, and device traits. Firms mine this vast pool of data to tailor recommendations and, more controversially, prices. Personalized pricing is documented in many industries.[1] As data accumulate and analytics become cheaper, the prospect of increased price discrimination has renewed regulatory and academic interest in its welfare and competitive effects.[2]

Many markets in which personalized prices and targeted offers arise are naturally described as consumer search markets. Consumers do not observe all available options and prices at once; they must invest time and attention inspecting offers. In e-commerce, consumers click on product pages to learn product details. Price discrimination has been documented in such settings (Hannak et al., 2014). Evidence also indicates substantial heterogeneity in search frictions: some consumers can continue searching at low cost while others face higher costs, particularly in online markets (Jolivet and Turon, 2018; Hong and Shum, 2006).[3] At the same time, firms observe rich digital traces that can be informative about (i) preferences (e.g., willingness to pay) and (ii) the search environment—outside options and the ease or inclination to solicit additional offers ("willingness to walk"), which is captured by the search cost c. These traces can reveal a consumer's propensity to keep searching, allowing firms to tailor prices to inferred search costs. Although much of the literature studies discrimination along the valuation dimension, this paper focuses on discrimination along the willingness-to-walk dimension. If search costs are partially inferable from behavior, such as search duration or click intensity, it is natural to study how outcomes change when firms observe this information, relative to an uninformative benchmark in which they do not observe it.

To study search-cost-based price discrimination, this article considers a sequential-

[1] Personalization is widespread across sectors (European Commission, 2018) and has been documented in e-commerce (Hannak et al., 2014) and tutoring (Larson et al., 2015); see also Kroger and Lidl Plus for personalized price discounts.

[2] See, e.g., CMA (2018); Ofcom (2020); European Parliament (2022); FTC (2024); OECD (2018).

[3] For example, Jolivet and Turon (2018) show that many purchases in French e-commerce are consistent with low search costs, while a nontrivial share requires high costs; heterogeneity is also needed to rationalize prices and consumer behavior online for books and offline for cars (Hong and Shum, 2006; Moraga-González et al., 2023; De los Santos, 2018); see also Blake et al. (2016); Ursu et al. (2023).

search model with horizontal differentiation, following Wolinsky (1986) and Moraga-González et al. (2017) (Section 2). A unit mass of consumers engages in costly sequential search to discover how much they value a firm's product. Consumers are heterogeneous in their search costs, which may be high enough to deter participation. A unit mass of firms offers horizontally differentiated products, and consumers learn prices and match values only after inspection. Before firms set prices for a given consumer, all firms observe a public signal of the consumer's search cost c (e.g., shared by a platform or purchased from a common data broker). Match values v are never observed; this focuses the analysis on search-cost-based price discrimination. The framework combines endogenous participation with search-cost-based price discrimination in a differentiated-goods sequential-search environment and permits analysis of public-signal design about the search cost, c. The analysis proceeds in two parts.

The first part of the analysis examines how firms' information about search costs affects consumer surplus and welfare when participation is endogenous. I begin by contrasting two benchmark information regimes: one in which firms observe search costs and fully personalize prices, and one in which firms observe no information and charge a uniform price. When all consumers have sufficiently low search costs, all consumers participate under either regime. In this case, withholding information about search costs benefits consumers while leaving total welfare unchanged. Intuitively, heterogeneous search costs generate heterogeneous price sensitivities, and equilibrium pricing admits an "average-sensitivity" characterization: the price for a pooled group is pinned down by the group's average price sensitivity. It will be shown that average prices are higher when firms price off individual sensitivities than when they must price off the average sensitivity (Proposition 1; Section 3). When search costs are sufficiently dispersed, some consumers do not participate under personalized pricing. Moving to an uninformative regime expands participation and raises welfare by serving additional consumers. The effect on aggregate consumer surplus, however, is ambiguous because it reflects three forces (Proposition 3; Section 3). The first is the same averaging force from the full-participation case, which tends to lower prices when information is removed. The second is the surplus created

for new participants who enter under uniform pricing. The third is a composition effect: the marginal entrants are relatively price-insensitive, so their entry changes the average sensitivity and raises the equilibrium price paid by all consumers, potentially enough to reduce aggregate consumer surplus. After characterizing equilibrium under the benchmarks of full observability and no observability of search costs, intermediate information structures are studied in which firms observe noisy signals of consumers' search costs. This permits comparative-static and welfare analysis across Blackwell-ordered information structures. Consistent with the benchmark comparisons, the effect of additional information depends on how participation responds: when participation is independent of the information structure, Blackwell-more informative signals reduce consumer surplus (Corollary 1; Section 4.1); when participation depends on the signal structure, consumer surplus can increase or decrease with informativeness, and can be non-monotone along Blackwell refinements (Proposition 5; Section 4.1).

The second part of the analysis characterizes the information structure about search costs that maximizes consumer surplus. The optimal information structure has a threshold form: firms observe one signal for consumers with search costs below a cutoff and another signal for consumers above it (Proposition 6, Section 4.2). Consumers above the cutoff do not participate, so further distinctions among high-search-cost types are irrelevant. Thus, a binary signal implements the consumer-optimal information structure.

This analysis complements the seminal result that, under competition, personalized pricing can intensify competition among firms and benefit consumers (as in Thisse and Vives (1988)), as well as the subsequent literature examining the limits of that conclusion (e.g., Rhodes and Zhou (2024)). The findings show that the source of price discrimination and the market structure matter for the regulation of personalized pricing, since welfare and consumer surplus effects can go in different directions depending on the setting. Moreover, the information-design results might offer one potential contributing factor for why intermediaries with rich behavioral data may nevertheless disclose only coarse information to sellers: in the model, relatively coarse public signals can maximize consumer surplus and expand participation.

Literature. This paper lies at the intersection of competitive price discrimination and consumer search. It builds on differentiated-goods models of consumer search (e.g., Wolinsky (1986); Anderson and Renault (1999)) and their extension to heterogeneous search costs (Moraga-González et al., 2017), and relates to oligopoly models of personalized pricing and price discrimination (e.g., Thisse and Vives (1988); Armstrong (2006)).[4]

Seminal work on the role of information in search environments is Bergemann et al. (2021), who characterize profit-maximizing private signals about consumers' search activity in a homogeneous-goods market with simultaneous search. Relatedly, Fabra and Reguant (2020) study discrimination based on quantity demanded in a homogeneous-goods market with simultaneous search when consumers differ in both search costs and quantities, and Mauring and Williams (2023) analyze discrimination based on search length in a homogeneous-goods market with heterogeneously impatient consumers, emphasizing the resulting Coasian price dynamics as the composition of remaining searchers changes over time. Mauring (2025) is particularly close in spirit, as it explicitly studies search-cost-based price discrimination. It examines a homogeneous-good Stahl (1989) shopper/non-shopper model in which firms receive private signals about consumers' search costs. The signal structure is such that, with some independent probability across searches, a consumer's type is perfectly revealed; otherwise, no information is revealed. Since the first search is free, all consumers participate. The article focuses on how the private nature and precision of these signals shape equilibrium pricing, as the signal reveals the consumer's type more frequently; equilibrium pricing reflects both the signal realization and the consumer's continuation value of waiting for a positive signal at a later firm.

In contrast to these articles, the present paper studies public signals about consumers' search costs in a differentiated-goods sequential-search market with endogenous participation. A central additional force is a selective-entry channel: information affects which consumers enter and search, and the induced composition of active searchers feeds back

[4] A large literature studies information and price discrimination under monopoly, going back to Pigou (1920) and including seminal modern contributions such as Bergemann et al. (2015). This paper instead focuses on oligopoly search markets, where competition interacts with partial observability of search costs.

into firms' pricing incentives and welfare, which in turn shapes the welfare consequences of alternative public information structures about search costs. Allowing for match-value differentiation also changes the nature of search; consumers trade off prices and match values, yielding a tractable pure-strategy pricing equilibrium with a transparent characterization of equilibrium prices that is useful for the subsequent welfare and information-design analysis.

A different strand studies price discrimination in search markets based on willingness to pay (WTP) rather than search costs. A central issue is the Diamond (1971) logic: when firms perfectly observe a consumer's WTP, they can extract the full surplus, so consumers have no incentive to search or participate. With imperfect signals, this logic can be mitigated (e.g., Groh (2026); Groh and von Wangenheim (2024)). Related work studies incentives to learn valuations and the role of privacy in search and decentralized trade (Lauermann, 2012; Choi and Rocheteau, 2024), as well as discrimination based on observables such as browsing and search histories that proxy for WTP (Groh and Preuss, 2025; Armstrong and Zhou, 2016; Armstrong, 2017).[5] The present paper is complementary: it studies price discrimination in search markets when discrimination is based on heterogeneity in search costs rather than WTP.

In frictionless oligopoly, a seminal contribution is Thisse and Vives (1988), which shows that WTP-based discrimination can benefit consumers when all consumers purchase. Related and subsequent work extends this insight in several directions (Chen and Iyer, 2002; Shaffer and Zhang, 2002; Montes et al., 2019; Esteves, 2014; Ali et al., 2023). Recent work highlights that participation and targeting frictions are crucial for the welfare conclusions (Rhodes and Zhou, 2024; Anderson et al., 2023). The present paper complements this literature by showing that when discrimination is based on search costs rather than WTP, consumers need not benefit from price discrimination even under full participation.

Finally, the paper relates to the privacy literature on data collection and targeting

[5]Preuss (2023) considers consumer search in a monopoly setting. The monopolist uses information about whether a consumer knew the product's valuation before arrival to personalize prices, which is indicative of a sufficiently high WTP.

in online markets (Braghieri, 2019; Esteves and Resende, 2016; Ichihashi, 2020; Conitzer et al., 2012; De Cornière and De Nijs, 2016). Whereas those works primarily consider personalization based on preferences and hence WTP, the present article focuses on heterogeneity in search costs and its implications for pricing and information policy.

Outline. Section 2 presents the model. Section 3 characterizes equilibrium under full versus no observability of c and decomposes the effect of privacy. Section 4 studies noisy public signals of c and characterizes the consumer-surplus-maximizing signal, enabling comparative statics and welfare comparisons as informativeness varies. Section 5 studies voluntary disclosure and establishes unraveling. Section 6 extends the analysis to joint signals about valuations and search costs and to finitely many firms. Section 7 discusses policy implications and relates the results to the literature. Section 8 concludes. The appendix contains all proofs except those for Section 6, which appear in the online appendix; the online appendix also includes extensions with voting, search-cost comparative statics, inference of search costs from data, and additional simulations.

2 Model and Equilibrium

The model builds on Moraga-González et al. (2017), which extends Wolinsky (1986) to heterogeneous search costs under uniform pricing (equivalently, an uninformative signal). I extend this framework by introducing partially informative signals about search costs, so that firms may condition prices on s. There is a unit mass of firms $j \in [0,1]$ and a unit mass of consumers $i \in [0,1]$. Firms produce a horizontally differentiated good at zero marginal cost. Consumer i values the product offered by firm j according to match value v_{ij}. These match values are i.i.d. draws from a continuous CDF F and density f on $[0,\overline{v}]$. Consumers have unit demand and an outside option 0; if consumer i buys from firm j at price p_{ij}, their utility is $v_{ij} - p_{ij}$.

Consumers search sequentially over firms to learn (v_{ij}, p_{ij}). The realized distribution of search costs c is described by the CDF G on $C = [\underline{c}, \overline{c}]$, $0 \leq \underline{c} < \overline{c}$. Search is characterized by uniform random matching and free recall. Each search step incurs a

cost c_i. Upon consumer i's visit to firm j, the firm observes a (public) signal s_i.

The information structure is a pair (S, σ), where $\sigma(c, \cdot)$ is a Markov kernel that, for each type c, returns a probability distribution over S which is the signal space. Elements of S are denoted by $s \in S$. Let the unconditional distribution over signals be M_σ and let the posterior distribution over c given $s \in S$ be $G(c \mid s)$. Because comparisons across information structures require integrating over signals, standard regularity conditions are imposed ensuring that the induced joint distribution of search costs and signals admits a well-defined posterior $G(c \mid s)$.[6]

Assumption 1 (Match-value regularity). $R(x) = x(1 - F(x))$ is concave on $[0, \overline{v}]$; $1 - F$ is strictly log-concave; F is twice differentiable on $(0, \overline{v})$.

Assumption 2 (No atom at zero). $G(0) = 0$.

Assumption 3 (Signal regularity). Except on an M_σ-measure-zero set of signals, $G(\cdot \mid s)$ has no atoms except possibly at the bottom of its support.

When G is atomless, many standard signal structures satisfy Assumption 3, including uninformative signals, perfectly informative signals, and monotone partitions. When $\underline{c} > 0$, Assumption 2 is trivially satisfied. Allowing atoms at the lower bound is feasible but complicates the exposition.

The timing is as follows. First, for each consumer $i \in [0, 1]$, nature draws a search cost $c_i \sim G$ (with support $C = [\underline{c}, \overline{c}]$) and then draws a public signal s_i according to the kernel $\sigma(c_i, \cdot)$. After observing (c_i, s_i), consumer i decides whether to enter the market by initiating a search. An entering consumer visits a random firm j. Upon the visit, the firm observes s_i and posts a personalized price p_{ij}, while a match value v_{ij} is realized. The consumer then observes (v_{ij}, p_{ij}) and chooses whether to purchase, to exit the market, or

[6] Assume that S is a compact, metrizable space and S and C are both endowed with their Borel σ-algebras, denoted by $\mathcal{B}$. A Markov kernel $\sigma : C \times \mathcal{B}(S) \to [0, 1]$ satisfies i) $c \mapsto \sigma(c, B)$ is measurable for every $B \in \mathcal{B}(S)$, ii) $B \mapsto \sigma(c, B)$ is a Borel probability measure on S for every $c \in C$. Given a prior $G \in \Delta(C)$, consider the probability space $(C \times S, \mathcal{B}(C \times S), \mathbb{P})$ with $\mathbb{P}$ satisfying

$$\mathbb{P}(B_C \times B_S) = \int_{B_C} \sigma(c, B_S) dG(c)$$

whose $S-$marginal is $M_\sigma(B_S) = \int_C \sigma(c, B_S) dG(c)$. Because the space is compact and metrizable, there exists a regular conditional distribution $G(\cdot \mid s)$ of C given $S = s$ which is unique for M_σ-almost every s (Klenke, 2013, Thms. 8.37 and 8.38).

to pay the search cost again and continue searching at a previously unvisited firm drawn at random.

The solution concept is essentially a symmetric pure-strategy Perfect Bayesian Equilibrium. A consumer's strategy is a search/purchase rule as a function of (c, s) and observed prices, and a firm j's strategy is a signal-contingent pricing function $p_j : S \to \mathbb{R}_+$. Beliefs are Bayes-consistent wherever possible given the prior and the information structure σ; off the equilibrium path, beliefs satisfy a no-signaling-what-you-don't-know restriction, so observing an off-equilibrium price at one firm does not affect beliefs about prices at unvisited firms. An equilibrium is a profile of strategies and beliefs such that strategies are sequentially optimal given beliefs. I restrict attention to symmetric equilibria with $p_j = p$ for all j and denote the resulting pricing rule by $p^*_\sigma(\cdot)$.

Two Benchmarks. The two extremes serve as benchmarks.

Perfectly informative: firms observe $s_i = c_i$, allowing for "perfect" price discrimination (PD). Formally, $S = C$ and $\sigma^{PD}(c, \{s\}) = \mathbb{1}\{s = c\} \quad \forall c$.

Uninformative: the signal is constant, so prices cannot depend on s, inducing a uniform price (U). Formally, $S = \{s_U\}$ and $\sigma^U(c, \{s\}) = \mathbb{1}\{s = s_U\} \quad \forall\, c \in [\underline{c}, \overline{c}]$.

2.1 Search and Prices

Fix an information structure σ. Consider a consumer with search cost c and signal s. In a symmetric pure-strategy equilibrium, all firms post the same price to consumers with signal s; denote it by $p^*_\sigma(s)$.

Because prices are the same at all visited firms, the value of an additional search comes only from improving the match value. Let the reservation value $\hat{x}(c)$ solve

$$\int_{\hat{x}(c)}^{\overline{v}} (v - \hat{x}(c))\, f(v)\, dv \;=\; c, \tag{1}$$

with the convention $\hat{x}(c) = 0$ if no solution exists. A consumer who currently holds a draw v benefits from one additional search if and only if $v < \hat{x}(c)$; this myopic cutoff rule is optimal in sequential search (Kohn and Shavell, 1974; Weitzman, 1979). The ex-ante

individual consumer surplus (CS) from starting to search can be shown to be $\hat{x}(c) - p^*_\sigma(s)$ (see, e.g., Moraga-González et al., 2017), so participation occurs if and only if $\hat{x}(c) \geq p^*_\sigma(s)$ and individual consumer surplus is given by $\mathbb{1}\{\hat{x}(c) \geq p^*_\sigma(s)\}\,(\hat{x}(c) - p^*_\sigma(s))$.

Define the participation cutoff

$$c_0(p) \;\equiv\; \max\big\{\, c \in [\underline{c}, \overline{c}] :\; \hat{x}(c) \geq p \,\big\}, \tag{2}$$

interpreting the set as empty (no participation) when $\hat{x}(\underline{c}) < p$.

Consider a firm j that posts p while all other firms post $p^*_\sigma(s)$. Upon visiting j, a consumer with $c \leq c_0(p^*_\sigma(s))$ buys if and only if

$$v - p \;\geq\; \hat{x}(c) - p^*_\sigma(s) \quad\Longleftrightarrow\quad v \;\geq\; \hat{x}(c) + p - p^*_\sigma(s), \tag{3}$$

which occurs with probability $1 - F\big(\hat{x}(c) + p - p^*_\sigma(s)\big)$. With a unit mass of firms and uniform random matching, the expected number of searches before stopping is $1/\big(1 - F(\hat{x}(c))\big)$; this is the arrival weight at any given firm.

Aggregating over active consumers $c \leq c_0(p^*_\sigma(s))$, the demand faced by firm j at price p is

$$D_\sigma\big(p;\, p^*_\sigma(s) \mid s\big) \;=\; \int_{\underline{c}}^{c_0(p^*_\sigma(s))} \frac{1 - F\big(\hat{x}(c) + p - p^*_\sigma(s)\big)}{1 - F(\hat{x}(c))}\, dG(c \mid s), \tag{4}$$

and profit is

$$\pi_j\big(p;\, p^*_\sigma(s) \mid s\big) \;=\; p \cdot D_\sigma\big(p;\, p^*_\sigma(s) \mid s\big). \tag{5}$$

For more details, see Appendix OA.1.

Lemma 1 (Equilibrium price and uniqueness)

Fix an information structure σ. For almost every signal s, any equilibrium in which at least some consumers are active satisfies:

$$G\big(c_0(p^*_\sigma(s)) \mid s\big) > 0,$$

$$p^*_\sigma(s) \;=\; \frac{1}{\mathbb{E}_{c \sim G(\cdot|s)}\Big[\dfrac{f(\hat{x}(c))}{1 - F(\hat{x}(c))} \mid \hat{x}(c) \geq p^*_\sigma(s)\Big]}. \tag{6}$$

The system of equations has a unique solution under Assumptions 1, 2, and 3.

Lemma 1 characterizes the equilibrium price and shows that the price is unique whenever there are consumers who participate. The remaining analysis will focus on these types of equilibria.[7]

The proof follows Moraga-González et al. (2017), but the assumptions on F and G differ. The proof involves two challenges. First, the equilibrium price is determined by a fixed point, where firms' optimal prices must be consistent with consumers' participation decision. Second, for the profit to be uniquely maximized at the first-order condition, the profit must be a concave function. Moraga-González et al. (2017) use the log-concavity of demand from each type of consumer; however, for a given signal, there are different types of consumers, so the demand that a firm faces integrates over different consumer types. In general, log-concavity is not preserved under integration, so conditions on G ensure that the aggregate demand is log-concave, which implies that the profit function is concave. Following Luther (2025), Lemma 1 employs a restriction on F, namely that $x(1-F(x))$ is concave to ensure that the profit from each type of consumer is concave. Concavity is preserved under expectations; the firm's profit function is also concave and therefore admits a unique maximizer.

The analysis tracks three aggregates: consumer surplus (CS), industry profits (Π), and welfare (W). Aggregate consumer surplus is the consumer surplus of all individual consumers, $CS_\sigma := \mathbb{E}\big[\mathbf{1}\{\hat{x}(C) \geq p^*_\sigma(S)\}\big(\hat{x}(C) - p^*_\sigma(S)\big)\big]$. Because all participating consumers eventually purchase a product, expected profits are the average traded prices scaled by the measure of participating consumers, $\Pi_\sigma := \mathbb{E}[\mathbf{1}\{\hat{x}(C) \geq p^*_\sigma(S)\}\, p^*_\sigma(S)]$. Welfare is the sum of consumer surplus and firm profits, $W_\sigma := \mathbb{E}[\mathbf{1}\{\hat{x}(C) \geq p^*_\sigma(S)\}\, \hat{x}(C)]$.

[7] They are the only equilibria that survive the D1 refinement. Under a D1 refinement, an off-path participation would be associated with the lowest c feasible under the signal s and thus the lowest price and largest incentive to enter. If there is no participation under the D1 refinement for a given signal s, there is no equilibrium with participation.

3 Perfectly informative and uninformative signals

Consider the perfectly informative (PD) and perfectly uninformative (U) signal benchmarks. When signals are perfectly informative, $s = c$, meaning that firms know exactly a consumer's search cost. $G(\cdot|s)$ is degenerate at c and Lemma 1 implies that for all active consumers the equilibrium price is $p^*_{PD}(c) = \frac{1-F(\hat{x}(c))}{f(\hat{x}(c))}$. Under Assumption 1, the function $c \mapsto p^*_{PD}(c)$ is strictly increasing on the set of participating (active) consumers. Define $h(c) := \frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}$, the hazard rate at the reservation value $\hat{x}(c)$. It summarizes consumer c's demand responsiveness to marginal price changes (see Section OA.1). With perfectly uninformative signals, the equilibrium uniform price satisfies $p^*_U = \frac{1}{\mathbb{E}_{c\sim G}[\frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}|\hat{x}(c)\geq p^*_U]}$.

Under U, the equilibrium price is the harmonic mean of the discriminatory prices $p^*_{PD}(c)$ for consumers who participate under U ($\hat{x}(c) \geq p^*_U$).

Lemma 2 (Relationship between p^*_U and p^*_{PD})

p^*_U and p^*_{PD} satisfy:

i) p^*_U is the harmonic mean of $p^*_{PD}(c)$ for the consumers who satisfy $\hat{x}(c) \geq p^*_U$.

ii) The uniform price p^*_U lies in the range of the discriminatory prices, p^*_{PD}:

$$\min_{c:\hat{x}(c)\geq p^*_U} p^*_{PD}(c) \leq p^*_U \leq \max_{c:\hat{x}(c)\geq p^*_U} p^*_{PD}(c) \tag{7}$$

The inequalities are strict whenever the set of active participants is non-degenerate, $c_0(p^*_U) > \underline{c}$.

iii) If $\hat{x}(\underline{c}) > p^*_{PD}(\underline{c})$, then $c_0(p^*_U) > \underline{c}$.

Lemma 2 shows that under U the price equals the harmonic mean of the prices each participant would face under PD. This means that the uniform price lies between the discriminatory prices that the lowest- and highest-search-cost consumers who participate would be charged.

Moving from PD to U alters prices and, through those prices, the set of active participants.

The next lemma shows that privacy, the absence of price discrimination, weakly expands participation relative to perfect price discrimination. Let $c_0^U := c_0(p_U^*)$ denote the highest-cost participant under privacy. Under PD, the highest-cost participant satisfies $c_0^{PD} = c_0(p_{PD}^*(c_0^{PD}))$. Both cutoffs exist and are unique whenever $\hat{x}(\underline{c}) \geq p_{PD}^*(\underline{c})$.

The set of active consumers under privacy is larger:

Lemma 3 (Privacy weakly expands participation)

The participation cutoffs under perfect information c_0^{PD} and no information c_0^U satisfy

$$c_0^{PD} \leq c_0^U, \tag{8}$$

with strict inequality whenever $\underline{c} < c_0^{PD} < \overline{c}$.

Intuitively, under price discrimination, firms adjust prices based on consumers' search costs, charging higher prices to those with higher search costs. This implies that low-search-cost consumers would face lower prices under price discrimination and higher prices under privacy, while the reverse holds for high-search-cost consumers. The marginal participating consumer under price discrimination has a search cost of c_0^{PD} and is indifferent between participating and not participating. Under privacy, pooling yields a uniform price—a weighted average of discriminatory prices—that is below the PD price faced by the marginal participant, so participation expands and $c_0^U \geq c_0^{PD}$.

If search costs are sufficiently low, everyone participates, even with discrimination.[8] Otherwise, only some consumers participate.

Two regimes, therefore, arise.

A market is *fully covered* when $c_0^{PD} = \overline{c}$. Similarly, a market is *partially covered* if $c_0^{PD} < \overline{c}$.

[8]For example, when $v \sim U(0,1)$, and $\overline{c} \leq \frac{1}{8}$, then the market is fully covered. In other words, the market is fully covered if the maximal search costs are at most a quarter of the expected value of a random draw.

3.1 Full Coverage

All consumers participate when the market is fully covered, regardless of the privacy regime. Privacy affects prices but not participation. Recall that welfare is the sum of industry profits and consumer surplus: $W = CS + \Pi$.

Proposition 1 (Full Coverage)
Under full coverage, it holds that

i) $\mathbb{E}_{c\sim G}[p^*_{PD}(c)] > p^*_U$

ii) $CS_U > CS_{PD}$

iii) $\Pi_U < \Pi_{PD}$

iv) $W_U = W_{PD}$

Proof. Because of full coverage, all consumers participate. This means that

$$\mathbb{E}_{c\sim G}[p^*_{PD}(c)] = \mathbb{E}_{c\sim G}\left[\frac{1}{\frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}}\right] > \frac{1}{\mathbb{E}_{c\sim G}\left[\frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}\right]} = p^*_U$$

The inequality holds because $\frac{1}{x}$ is strictly convex for $x > 0$, which holds a.s., $f/(1-F)$ is not constant, and G is non-degenerate. Thus, Jensen's inequality is strict.

Welfare is defined as:

$$\begin{aligned} W &= CS + \Pi \\ &= \mathbb{E}[\mathbb{1}\{\hat{x}(C) \geq p^*_\sigma(S)\}\,(\hat{x}(C) - p^*_\sigma(S) + p^*_\sigma(S))] \\ &= \mathbb{E}[\mathbb{1}\{\hat{x}(C) \geq p^*_\sigma(S)\}\,\hat{x}(C)]\,. \end{aligned}$$

Therefore, welfare only depends on the set of active consumers, which is the same under full coverage.

From the average traded price, it follows that $CS_U > CS_{PD}$ and $\Pi_U < \Pi_{PD}$. □

There are two explanations why privacy reduces the average price in the market, an effect that may be called the *direct privacy effect*. Technically, under privacy the market

price is a harmonic mean of discriminatory prices, which is below their arithmetic mean. Economically, prices are inversely related to price sensitivity, $\frac{f}{1-F}$, and are convex in that sensitivity. Pricing on the average sensitivity therefore lowers the average price by Jensen's inequality: In equilibrium, firms compete fiercely for consumers with high price sensitivity (low c), resulting in low prices, while consumers with low price sensitivity (high c) face higher prices. If a group's sensitivity increases marginally, an increase in price sensitivity for already highly sensitive consumers has only a minor impact on equilibrium prices. This is because prices for these consumers are already close to zero, leaving firms with little incentive to lower prices further to attract demand. However, a similar change in the price sensitivity of a low-sensitivity group leads to a significant price reduction, as firms are eager to capture more demand at high prices. Combining two groups, one with low search costs and another with high search costs, results in a uniform price based on their average price sensitivity. This raises prices for low-search-cost consumers but lowers them significantly for high-search-cost consumers. Jensen's inequality captures this effect and provides insight into why consumers, on average, are harmed by price discrimination.

The insight that pooling groups with different price elasticities benefits consumers (in fully covered markets) applies to settings beyond search markets.[9] However, with partial coverage the interaction between price discrimination and privacy is more complex.

With full coverage, welfare is unchanged; changes in average prices merely reallocate surplus between consumers and firms.

3.2 Partial Coverage

When coverage is partial ($c_0^{PD} < \overline{c}$), the direct privacy effect persists, but endogenous participation adds additional forces.

[9]The effect of pooling consumers with different search costs is reminiscent of pricing consumers with heterogeneous tastes in Hotelling-Duopoly (Armstrong and Vickers, 2001; Armstrong, 2006). It is shown that price discrimination based on consumer tastes yields discriminatory prices that also take the form of a harmonic mean, and that privacy (pooling) can increase consumer surplus. Although the contexts and sources of elasticity differ, search-cost-based rather than preference-based, pooling benefits consumers in both models. Because the models are highly related, later results on partial coverage and optimal information design should also apply to the preference-based price discrimination setting on the Hotelling line.

Proposition 2 (Partial coverage)

Under partial coverage, it holds that $W_U > W_{PD}$.

Proof. Recall that welfare is

$$W_U = \int_{\underline{c}}^{c_0^U} \hat{x}(c) dG(c) > \int_{\underline{c}}^{c_0^{PD}} \hat{x}(c) dG(c) = W_{PD} \tag{9}$$

The inequality follows from $\hat{x}(c) > 0$ for all $c \leq c_0^U$ and $G(c_0^U) - G(c_0^{PD}) > 0$. □

Thus, privacy increases participation and, in turn, welfare.

The effects on consumer surplus and profits are ambiguous: CS_U can exceed or fall below CS_{PD}, and likewise for profits. For example, when $\overline{c} - \epsilon < c_0^{PD} < \overline{c}$ for small ϵ, the full-coverage conclusions obtain in the limit.

To understand how consumer surplus is impacted by privacy, it is useful to decompose the change in consumer surplus, $\Delta CS = CS_U - CS_{PD}$, into three effects: the previously discussed *Direct Privacy Effect* and two new effects: *Direct Participation Effect* and *Participation Externality*.

Proposition 3 (Decomposition)

$$\Delta CS = \underbrace{\int_{\underline{c}}^{c_0^{PD}} \left(p_{PD}^*(c) - p_U^{before} \right) dG(c)}_{\textit{Direct Privacy Effect}} - \underbrace{\int_{\underline{c}}^{c_0^{PD}} \left(p_U^{after} - p_U^{before} \right) dG(c)}_{\textit{Participation Externality}}$$
$$+ \underbrace{\int_{c_0^{PD}}^{c_0^U} \left(\hat{x}(c) - p_U^{after} \right) dG(c)}_{\textit{Direct Participation Effect}}$$

Whenever $c_0^{PD} < c_0^U$*:*

$$p_U^{before} = \frac{1}{\mathbb{E}_{c \sim G}\left[\frac{f(\hat{x}(c))}{1 - F(\hat{x}(c))} \,\middle|\, c \leq c_0^{PD} \right]} < p_U^{after} = \frac{1}{\mathbb{E}_{c \sim G}\left[\frac{f(\hat{x}(c))}{1 - F(\hat{x}(c))} \,\middle|\, c \leq c_0^{U} \right]}.$$

The direct privacy effect and the direct participation effect are positive; the participation externality enters negatively.

This decomposition allows us to construct extreme cases.

Consumer surplus may decrease due to privacy. This is the case whenever the direct privacy effect and the direct participation effect are insufficiently large relative to the participation externality.

Example 1

Let G_B be a binary distribution with mass m at $c_l < c_0^{PD}$ and mass $1-m$ at $\hat{x}(c_h) = p_U^*$. Then, the direct participation effect is zero. The set of always active consumers, $c \leq c_0^{PD}$, is a singleton, and hence there is no direct privacy effect. Hence, the only non-zero effect is due to the participation externality, which reduces consumer surplus.[10]

This example also establishes that firms' profits can increase due to privacy. The uniform price p_U^* exceeds the PD price for the marginal participant, and participation increases.

The direct privacy effect is large whenever there is a large set of active consumers and their price sensitivity is very dispersed. Consumers with search costs above c_0^{PD} selectively participate depending on the information structure. Consumers with search costs close to c_0^{PD} but far from c_0^U have sizable consumer surplus, thereby increasing aggregate consumer surplus. At the same time, the higher c, the higher $p_{PD}^*(c)$, which influences the uniform price p_U^*. Higher search costs exert a stronger externality on consumers.

In order for the negative price externality to dominate, a distribution G should not exhibit much dispersion in the region $c \leq c_0^{PD}$ and have a lot of mass close to c_0^U, where the price externality is the strongest and the direct participation effect is lowest.

Only if these properties are satisfied can privacy decrease consumer surplus. One distribution that does not exhibit these properties is the uniform distribution: the density is constant, so there is dispersion at $c \leq c_0^{PD}$ and no concentration of mass near c_0^U, limiting the scope for the negative price externality.

[10] To ensure equilibrium existence, one can either show directly that an equilibrium exists. Alternatively, to ensure that the example fits the model's regularity conditions (in particular, atomlessness), approximate G_B by a sequence of atomless distributions $(G_n)_{n\geq 1}$ obtained by "spreading" each point mass over an interval of width $\varepsilon_n \downarrow 0$ (and, if desired, adding a vanishing absolutely continuous component to maintain full support on $[\underline{c}, \bar{c}]$). Then $G_n \Rightarrow G_B$ weakly as $n \to \infty$, and for each n an equilibrium exists under the maintained assumptions. Moreover, $\Delta CS(G_n) \to \Delta CS(G_B)$, the direct privacy and direct participation terms converge to 0, while the price effect (the participation externality) remains bounded away from 0. Hence, for all sufficiently large n, $\Delta CS(G_n) < 0$.

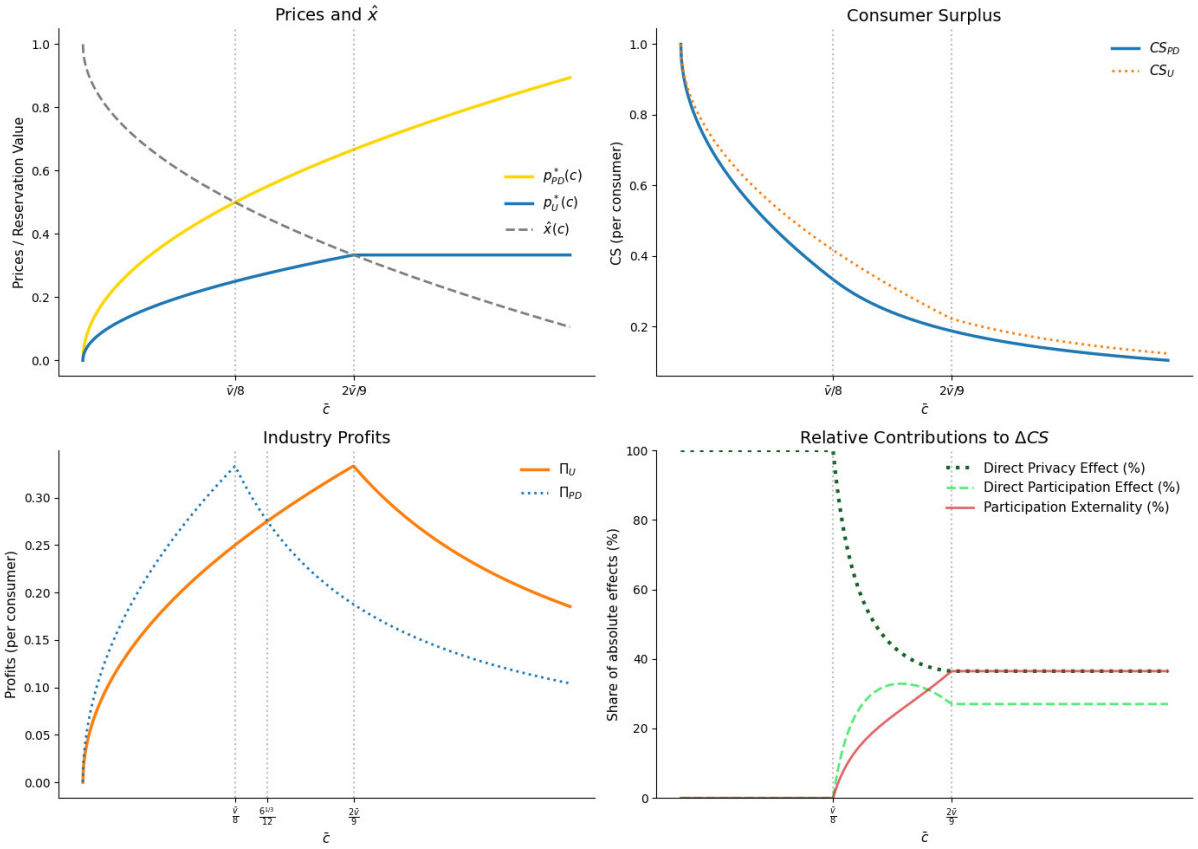


Figure 1: Prices, profits, and the decomposition of CS for $c \sim U[0, \bar{c}]$, $v \sim U[0, \overline{v}]$

Proposition 4 (Uniform Distributions)

When match values are uniformly distributed, $v_{ij} \sim U(0, \overline{v})$, and search costs are uniformly distributed, $c_i \sim U(0, \bar{c})$, then:

i) $\Delta CS > 0 \quad \forall \overline{v}, \bar{c} > 0$

ii) $\Delta \Pi \geq 0 \iff \frac{\bar{c}}{\overline{v}} \geq \frac{6\frac{1}{3}}{12}$

Figure 1 visualizes Proposition 4 and decomposes the changes in consumer surplus into the different effects. The upper left graph plots the equilibrium price under privacy p^*_U as a function of $\bar{c}$ as well as $p^*_{PD}(c)$ and $\hat{x}(c)$. For $c = \frac{1}{8}\overline{v}$, $p^*_{PD}(c)$ and $\hat{x}(c)$ intersect, and consumers with higher search costs therefore would not participate under price discrimination. The price with privacy is lower, and consumers stop participating if their search costs exceed $\frac{2}{9}\overline{v}$. Because the active set does not expand further, p^*_U becomes flat beyond that point: the distribution conditional on participation is unchanged as $\bar{c}$ increases.

Regardless of coverage, $CS_U > CS_{PD}$ (upper right panel). A decomposition of the effect into the three components characterized in Proposition 3 can be found in the lower right panel. Initially, under full coverage, only the privacy effect is present. Once $\overline{c} > \frac{1}{8}\,\overline{v}$, the price and CS effects occur. They have similar magnitudes, but the privacy effect dominates both other effects. The value of privacy for consumers $c \in [0, \frac{1}{8}\overline{v}]$ is the driving force. Lastly, full coverage implies $\Pi_U \leq \Pi_{PD}$. Once the market is only partially covered and privacy can induce sufficiently many new consumers to participate, $\Pi_{PD} < \Pi_U$. Afterward, profits fall, as an increase in $\overline{c}$ implies that more consumers do not participate.

4 Noisy and Optimal Signals

4.1 Noisy Signals

The privacy and price discrimination settings represent extreme points on the spectrum of information accuracy. Instead of firms having perfect or no information, they may only observe an imprecise signal of consumers' search costs.

The benchmark analysis delivered two main results. Under full coverage, more information reduced consumer surplus; under partial coverage, its effect was ambiguous. This ambiguity was organized around a three-part decomposition of the surplus change into a privacy effect benefiting consumers who always participated, a participation effect from the entry of consumers with positive individual surplus, and a negative participation externality in which entry raised prices for all active consumers.

The same decomposition remains informative, and the comparative statics mirror the benchmarks.

The set of all potential information structures is very large and lacks structure. Therefore, to sharpen some results, the class of monotone partitions will be examined. Monotone partitions divide the cost space into intervals, and the signal indicates which interval the search cost falls in. For example, consumers might be grouped by past search intensity. More formally, a monotone partition is characterized by thresholds

$\underline{c} = t_0 < t_1 < \cdots < t_K = \overline{c}$ and intervals $I_1 := [t_0, t_1]$ and $I_k := (t_{k-1}, t_k]$ for $k = 2, \ldots, K$. A monotone-partition information structure assigns signal s_k to consumers with $c \in I_k$, i.e., $\sigma(c, \{s_k\}) = \mathbb{1}\{c \in I_k\}$.

These monotone partitions have several useful properties.

Lemma 4 (Properties of monotone partitions)

For any monotone partition,

i) Price monotonicity: If $k < \ell$ and both $\{S = s_k,\ \hat{x}(c) \geq p^*_\sigma(s_k)\}$ and $\{S = s_\ell,\ \hat{x}(c) \geq p^*_\sigma(s_\ell)\}$ have positive probability, then $p^*_\sigma(s_k) < p^*_\sigma(s_\ell)$.

ii) Participation threshold: There exists a $c_0^m \in [\underline{c}, \overline{c}]$ such that consumers $c \leq c_0^m$ are active and consumers $c > c_0^m$ are inactive.

iii) Blackwell order of monotone partitions: Let σ and σ' be monotone partitions of C. Then σ' is Blackwell more informative than σ if and only if each interval of σ' is a subinterval of an interval of σ (i.e., σ' is obtained by subdividing the intervals of σ).

These properties mean that for the intervals where there is active participation, prices increase as the interval index increases. Because signals are deterministic and prices equal the (harmonic) mean of an increasing function of c, higher signals entail higher prices. Prices weakly increase in search cost (and the signal); thus, participation follows a cutoff strategy with consumers $c \leq c_0^m$ participating and all others refraining. A monotone partition is more informative if and only if it takes the existing intervals and divides them into smaller intervals.

In the fully covered market setting, it was observed that increasing the informativeness of the information structure decreases consumer surplus. With two different information structures, a given consumer is associated with different signals under the two information structures. A consumer with cost c will be associated with signal s and s' under the two different information structures. Let the set of active consumers under σ be $A_\sigma = \{(c, s, s') : \hat{x}(c) \geq p^*_\sigma(s)\}$ and assume that a positive measure of consumers is active.[11]

[11] If $\sigma' : C \to \Delta(S')$ is Blackwell more informative than $\sigma : C \to \Delta(S)$, then there exists a garbling

G_A and M_σ^A are respectively the C and S marginals conditional on the consumer being active.

Lemma 5 (Blackwell refinement weakly increases the average price)

If σ' is Blackwell more informative than σ, $\sigma' \succ_B \sigma$, and the set of active participants has positive measure and is the same for σ and σ', $A_\sigma = A_{\sigma'} = A$, then

i) $\mathbb{E}_{s\sim M_\sigma^A}[p_\sigma^*(s)] \leq \mathbb{E}_{s'\sim M_{\sigma'}^A}[p_{\sigma'}^*(s')]$

ii) The inequality is strict iff there exists a measurable set $B \subseteq S$ with $M_\sigma^A(B) > 0$ such that, for M_σ^A-a.e. $s \in B$,

$$\mathrm{Var}_{s'\sim q(\cdot|s)}\big(p_{\sigma'}^*(s')\big) > 0.$$

Equivalently, $p_{\sigma'}^*(S')$ is not a.s. constant conditional on $S = s$.

Whenever the signals of consumers are refined such that the prices associated with the refined signals differ and participation remains the same, consumer surplus strictly decreases with more informative signals. In the fully covered case ($c_0^{PD} = \bar{c}$), the active set is unchanged: the worst-off consumer ($c = \bar{c}$) cannot face a higher price than under PD and therefore still participates.

Corollary 1 (Full coverage)

When there are two information structures $\sigma' \succ_B \sigma$ and there is full coverage ($c_0^{PD} = \bar{c}$), then

$$CS_\sigma \geq CS_{\sigma'}$$

In general, a strict Blackwell refinement need not change equilibrium prices, so consumer surplus may be unchanged. To obtain sharper implications, one may consider monotone partitions.

$\gamma : S' \to \Delta(S)$ such that $\sigma = \gamma \circ \sigma'$ (Blackwell, 1953). This defines a joint law $\mathbb{P}$ on $C \times S \times S'$. Assume that $\mathbb{P}(A_\sigma) > 0$. This induces a joint probability distribution $\mathbb{P}^A(\cdot) = \frac{\mathbb{P}(\cdot \cap A_\sigma)}{\mathbb{P}(A_\sigma)}$. By the previous reasoning, regular conditional distributions under $\mathbb{P}^A$ exist; in particular, the conditional distribution of C given $S = s$, $G_A(\cdot \mid S = s)$, and the conditional distribution (kernel) of S' given $S = s$,

$$q(B' \mid s) := \mathbb{P}^A(S' \in B' \mid S = s) \qquad \text{for measurable } B' \subseteq S'.$$

.

Corollary 2 (Monotone partitions: refinement, surplus, participation)

Let σ and σ' be monotone partitions and let $\sigma' \succ_B \sigma$ be a strict Blackwell refinement (so at least one interval of σ is strictly subdivided on a set of positive G_A-measure).

i) Under full coverage ($c_0^{PD} = \bar{c}$) consumer surplus is strictly decreasing in a signal's informativeness: $CS_{\sigma'} < CS_{\sigma}$.

ii) Under partial coverage, participation weakly decreases: if c_0 and c_0' are the cutoffs under σ and σ', then $c_0' \leq c_0$.

Intuitively, under full coverage, the set of active consumers is fixed, so informativeness affects consumer surplus only through prices. For monotone partitions, any strict refinement is an interval subdivision (Lemma 4(iii)) that induces distinct prices across participating signals (Lemma 4(i)), so the strict case of Lemma 5 applies. Under partial coverage, monotone partitions imply a cutoff $c_0 = c_0^m$ such that consumers $c \leq c_0$ participate and $c > c_0$ do not (Lemma 4(ii)). If a strict refinement subdivides the interval I_k with $c_0 \in I_k$, then the posterior faced by types near the margin becomes truncated from below, which weakly raises the equilibrium price in the subinterval containing c_0. Hence, participation weakly decreases, i.e. $c_0' \leq c_0$.

With partial coverage, the effect of more informative signals on aggregate consumer surplus is ambiguous. The change in consumer surplus can be decomposed into three effects: the direct effect of privacy, the direct participation effect, and the negative participation externality. This decomposition remains useful even with noisy signals.

Suppose the initial signal structure is sufficiently uninformative (posteriors have positive density for all types and not all types participate). Then the marginal participant under that signal has no direct participation effect but imposes a negative participation externality. Constructing a signal that perfectly reveals the marginal participant's type decreases the average price paid by other consumers without reducing the direct participation effect. Therefore, there exists a more informative signal structure that increases total consumer surplus.

At the same time, making signals more informative for low-search-cost consumers erodes the value of pooling (direct privacy effect) and can induce some intermediate

search-cost types, who would have participated under a coarser signal and had a positive consumer surplus there, to exit. Therefore, increasing informativeness need not increase consumer surplus.

Proposition 5 (Local refinement raises CS and weakly lowers participation)
Assume there is a signal or collection of signals with positive probability such that for each signal s:

(i) the posterior $G(\cdot \mid s)$ has full support on an interval $[\underline{c}_s, \overline{c}_s] \subseteq [\underline{c}, \overline{c}]$; and

(ii) the market is not fully covered given s, i.e. $\hat{x}(\overline{c}_s) - p_\sigma(s) < 0$.

Then:

1. *There exists a strict Blackwell refinement $\sigma' \succ_B \sigma$ such that $CS(\sigma') > CS(\sigma)$ and participation weakly decreases (equivalently, the participation cutoff weakly falls).*

2. *There exists a sequence $(\sigma_n)_{n\geq 1}$ with $\sigma_{n+1} \succ_B \sigma_n$ for all n such that $(CS(\sigma_n))_{n\geq 1}$ is non-monotone.*

The following example illustrates the two refinements using monotone partitions.

Example 2

Let there be three types of consumers with different search costs:[10]

$$c_L, c_M \leq c_0^{PD} < c_H \quad \text{and} \quad \hat{x}(c_H) = p_U^*$$

A completely uninformative signal structure pools all three types together. Type c_H generates no positive CS from direct participation or privacy and only a negative price effect. Thus, c_H reduces total consumer surplus, and the negative participation externality dominates.

One strict Blackwell refinement is a monotone partition that separates c_H while pooling c_L and c_M: choose $t_1 \in (c_M, c_H)$ and $t_2 = \overline{c}$. This increases consumer surplus by excluding c_H from the market (since $c_H > c_0^{PD}$). Any further strict Blackwell refinement must split the remaining pooled block and thus separate c_L from c_M; equivalently, it yields a

perfectly revealing partition with thresholds $t_1' \in (c_L, c_M)$, $t_2' = t_1$, and $t_3 = \overline{c}$. The lower interval of the intermediate (two-interval) partition $[\underline{c}, t_1]$ is fully covered, so refining it erodes the value of pooling without any participation gain. Hence, this further increase in informativeness reduces consumer surplus.

4.2 Optimal Information Design

In fully covered markets, these insights imply that uninformative signals maximize consumer surplus. However, the results from the previous section also show that neither perfectly uninformative signals nor perfectly informative signals are optimal. Because consumer surplus can be non-monotone in informativeness, an interior level of information is optimal.

Proposition 6 (Consumer Surplus Maximization)

The consumer-surplus-maximizing information structure σ^* is characterized by a threshold m^*. It assigns a common pooling signal s_L to all consumers with $c \in [\underline{c}, m^*]$. Consumers with $c > m^*$ receive any signal $s \neq s_L$.

The threshold m^* is the unique solution to:

$$\hat{x}(m^*) - 2p(m^*) + p(m^*)^2 \frac{f(\hat{x}(m^*))}{1 - F(\hat{x}(m^*))} = 0, \tag{10}$$

where $p(m^*) = (\mathbb{E}[1/p^*_{PD}(C) \mid C \leq m^*])^{-1}$. The cutoff satisfies $c_0^{PD} \leq m^* \leq c_0^U$, with strict inequalities whenever the market is partially covered.

Proposition 6 implies that all active consumers are pooled into a single signal that induces a uniform price. Consumers above the threshold are assigned other signals and remain inactive. Therefore, signal design selects the active set. In particular, a binary information structure with one pooling signal for $c \leq m^*$ and one signal for $c > m^*$ implements σ^*.

More precise information generally harms consumers, so pooling is desirable. Yet high-search-cost consumers impose a negative price externality on others, so excluding them by assigning them a distinct signal is beneficial.

The design problem reduces to choosing which consumers participate. Because high search-cost consumers are least attractive (they generate low reservation value and consumer surplus but a large negative externality), the optimal policy is a cutoff rule. The cutoff satisfies the first-order condition in (10), balancing the marginal participant's direct surplus, $\hat{x}(c) - p$, against the price externality imposed on others, $p'(G(m^*))$, where p' denotes the price increase from including type m.

Example 3

If $c \sim U[0, \overline{c}]$ and $v \sim U[0, \overline{v}]$ with $\overline{c}, \overline{v} > 0$, then the consumer-surplus-maximizing information structure is to pool

$$c \in [0, m^*], \qquad m^* = \min\left\{\overline{c},\ \tfrac{8}{49}\,\overline{v}\right\}.$$

The example shows that only consumers with $c \leq m^*$ are active under the information structure. Consistent with Proposition 6, the cutoff satisfies $c_0^{PD} \leq m^* \leq c_0^U$. This is intuitive because the marginal participant under privacy makes no contribution to consumer surplus but exerts a negative externality on prices, which cannot be optimal.

Profit or Welfare Maximizing Signals. Profit- or welfare-maximizing signal structures need not be as simple. Within the closure of monotone partitions, however, the optimal structures are tractable. Monotone partitions with finitely many intervals achieve perfect information only in the limit as intervals shrink; taking the closure therefore admits perfectly informative signals.

Proposition 7 (Welfare and Profit Maximization)

Within the closure of the class of monotone partitions:

i) Welfare is maximized by the uninformative signal σ^U (complete pooling).

ii) Industry Profit is maximized by an upper-censorship policy characterized by a threshold $m^{**} \in [\underline{c}, \overline{c}]$. The optimal structure perfectly reveals all search costs $c \leq m^{**}$ and pools all consumers with $c > m^{**}$ into a single signal.

In a fully covered market ($c_0^{PD} = \overline{c}$), the profit-maximizing threshold is $m^{**} = \overline{c}$, implying

the perfectly informative signal σ^{PD} is optimal.

In fully covered markets, firms prefer fully informative signals. In general, pooling an interval of positive measure can induce participation by consumers who would not participate under perfect information. In monotone partitions, participation is unaffected when the signal structure changes for intervals that do not contain the marginal participant, so increasing the informativeness of the lower-cost intervals increases firm profits.

Welfare is maximized by maximizing participation. In monotone partitions, a single large uninformative signal minimizes the marginal participant's price and thus maximizes participation.

5 Voluntary Disclosure by Consumers

The optimal information structure involves partial or no disclosure. So far, it has been assumed that consumers cannot reveal any additional information beyond the signal chosen by the planner. In practice, however, privacy regulation is based on voluntary disclosure and consent: consumers can decide which information firms are allowed to have and use. In the present framework, this can be interpreted as a consumer deciding whether to enable firms to use rich browsing data (e.g. their entire browser history), from which firms can infer the consumer's search cost c and hence implement the price-discrimination schedule $p^*_{PD}(c)$ or whether they do not share their data and do not consent to their data being used to personalize the price.

The opt-in and opt-out decision for data sharing is modeled as a simple disclosure choice. Before searching, consumers observe their search cost c. Each consumer then chooses whether to consent to data processing by all firms. If the consumer consents, every firm effectively observes c. If the consumer does not consent, no firm receives additional information beyond the uninformative public signal. Firms set prices that are optimal given the posterior over c induced by this consent decision and the resulting signal.

Proposition 8 (Voluntary disclosure and unraveling)

When consumers can decide whether to share their data or not, then all consumers share their data, and the equilibrium prices are the same as under perfect information, $p^*_{PD}(c)$.

The argument is the standard unraveling logic of Milgrom (1981). Whenever consumers with different search costs do not consent to data sharing, they are pooled and face a common price determined by the posterior distribution of search costs in that pool. Because the perfect-information price $p^*_{PD}(c)$ is strictly increasing in c, the consumer with the most favorable type (the lowest search cost) in any non-degenerate pool faces a price that is strictly above the personalized price and thus has a strict incentive to disclose the type. The pool shrinks, and the same reasoning applies to the new most favorable type. Iterating this argument implies that all consumers disclose, and the outcome coincides with the price-discrimination benchmark.

Hence, an opt-in/opt-out regime for data sharing effectively provides firms with perfect information and enables price discrimination, despite the presence of formal privacy protections.

Robustness. The unraveling logic is general and does not rely on firms starting from an uninformative signal. Fix any information structure and a realized signal s. The induced posterior pools multiple types, and the equilibrium price without further disclosure lies between the lowest and highest personalized prices within that pool. The lowest-search-cost type in the pool, therefore, strictly benefits from disclosing additional information, which initiates unraveling within each s and restores the price-discrimination outcome $p^*_{PD}(c)$ whenever disclosure is feasible. In particular, even if a planner commits to a coarse consumer-surplus-maximizing structure, allowing consumers to consent to data use leads to unraveling within each signal realization.

The same incentives apply if consumers can disclose only a noisy signal about costs rather than c itself, and they do not require a once-and-for-all disclosure decision. If disclosure is made firm-by-firm, low-cost consumers still reveal their information to obtain lower prices. Frictions, such as the probabilistic availability of verifiable evidence or disclosure costs, limit but do not eliminate unraveling. Consumers with sufficiently

favorable evidence and sufficiently low disclosure costs disclose, which reduces pooling, weakens the privacy benefit, and can reduce participation by raising prices for higher-cost types.

6 Extensions

This section studies two extensions. First, it examines how the findings change when signals contain information about both search costs and willingness to pay. If the correlation is not too strong, the main effects persist; however, the composition of participants under no information versus full information can differ from the baseline in which signals reveal only search costs.

The second extension analyzes a fully covered market with finitely many firms. When search costs are low enough that all consumers search (even if some do not purchase), the results in Section 3.1 obtain provided either the number of firms is sufficiently large or search costs are sufficiently small.

The online appendix examines voting by consumers (see Section OA.3) and the effects of a more dispersed distribution of search costs (see Section OA.5.1).

6.1 Correlation between search costs and valuations

To relax independence between match values and search costs, let the match-value distribution depend on the consumer's search cost c. In particular, let $v_{ij} \overset{\text{i.i.d.}}{\sim} F_c \equiv U(0, a - bc)$ with $a > 0$, $b \in \mathbb{R}$. Throughout, the analysis is restricted to $\{c : a - bc > 0\}$; for the linear case with $b > 0$ this requires $c < a/b$.

Lemma 6 (PD prices and participation with correlated valuations)
Suppose $v_{ij} \sim U(0, a - bc)$ with $a > 0$ and $b \in \mathbb{R}$. Then:

i) Participation under PD follows a cutoff rule: there exists $c_0^{PD} \in [\underline{c}, \overline{c}]$ such that consumers with $c \leq c_0^{PD}$ search and those with $c > c_0^{PD}$ do not.

ii) If $b \leq 8$, then the discriminatory price schedule $p_{PD}^*(c)$ is strictly increasing in c on the active set $[\underline{c}, c_0^{PD}]$.

iii) If $b > 8$, there exists a turning point $c_{\text{tp}} \in (\underline{c}, c_0^{PD})$ such that $p^*_{PD}(c)$ is strictly increasing on $[\underline{c}, c_{\text{tp}}]$ and strictly decreasing on $[c_{\text{tp}}, c_0^{PD}]$; that is, $p^*_{PD}(c)$ is single-peaked on the set of participating types.

If the signal is uninformative, firms observe only whether a consumer participates, not her search cost (or the signal for her WTP). As under the baseline model, the price under the uninformative information structure takes on the familiar form of a harmonic mean.

Lemma 7 (Uniform price under an uninformative signal)

Suppose $v_{ij} \sim U(0, a - bc)$ with $a > 0$ and $b \in \mathbb{R}$ and that the assumptions of Lemma 6 are satisfied. Then

$$p^*_U = \left(\mathbb{E}\left[\frac{1}{p^*_{PD}(c)} \;\middle|\; \hat{x}(c) \geq p^*_U \right] \right)^{-1}. \tag{11}$$

Combining the shape of the PD price schedule in Lemma 6 with the harmonic-mean characterization in Lemma 7 yields the following implications for participation when moving from price discrimination to the uninformative signal. The term partial coverage will be used in a more general sense, referring to the fact that participation decisions differ across information structures, not necessarily to the idea that participation is lower under the discriminatory price.

Corollary 3 (Participation under privacy)

(i) Full coverage. *If all consumers participate under both price discrimination and the uninformative signal, then moving from price discrimination to the uninformative signal strictly increases aggregate consumer surplus.*

(ii) Partial coverage. *If participation is incomplete under at least one of the two information structures, then pooling can change the set of participating types. The participation effect depends on the correlation parameter b:*

*a) **Baseline analogue** ($-8 \leq b \leq 8$). If $p^*_{PD}(c)$ is increasing on the PD-active set, then pooling weakly expands participation at the high–search-cost margin.*

*b) **Selection reversal** ($b < -8$). Pooling can reduce participation at the low–search-cost margin.*

c) ***Non-monotone PD prices*** ***($b > 8$).*** *Pooling may increase or decrease participation at the high–search-cost margin, depending on the distribution of c on the PD-active set.*

Equation (11) implies that p_U^* is a harmonic mean across participating types and therefore lies strictly below the highest discriminatory price among them whenever the active set has positive measure. In fully covered markets, pooling thus reduces the average price paid and strictly increases aggregate consumer surplus.

When participation is endogenous, correlation affects both the redistribution of prices within the active set and the composition of that set. In particular, participation need not be nested across regimes: moving to the uninformative signal can expand participation at some margins while inducing exit at others. For $-8 \leq b \leq 8$, discriminatory prices are increasing in c on the active set, so pooling raises prices for low-c consumers and lowers them for high-c consumers; the latter effect dominates at the extensive margin and tends to expand participation among high search-cost consumers. For $b < -8$, the correlation is reversed—low search costs are associated with low match values—so pooling can raise the uniform price above the willingness to pay of the lowest-c types, inducing exit at the low-c margin. Finally, for $b > 8$ the discriminatory price schedule can be non-monotone; then the pooled price p_U^* may lie above or below the price faced by the marginal discriminatory participant, so pooling can either expand or contract participation at the high-c margin.

Finally, if match values v_{ij} are perfectly predicted by search costs (perfect correlation), the Diamond (1971) Paradox arises. Due to frictions in the search, firms end up charging the monopoly price and extracting the full surplus. Because each firm does so, no consumer searches. In anticipation of full rent extraction, consumers incur no search costs and refrain from participation. A uniform price restores trade and increases consumer surplus in that case as well.

6.2 Finite Firms

In the baseline with a continuum of firms, all consumers eventually purchase and there is no recall. With finitely many firms, some consumers do not purchase because their

realized matches are low. The probability of purchase depends on the prices.

For simplicity, consider a standard uniform distribution of match values, $v_i \sim U(0,1)$, in the benchmark comparison of full information and no information. The market is assumed to be fully covered in the sense that all consumers who are active in the market choose to search (though they may not purchase).

If the equilibrium price is p^*, then the probability that a consumer purchases a product conditional on searching is $1-(F(p^*))^n$, the probability that at least one match-value realization is above p^*.

Then, the privacy effect holds:

Proposition 9

For the benchmarks of full information and no information and full coverage, it holds that if $v_{ij} \sim U(0,1)$, then with no information,

i) The average price is lower: $\mathbb{E}[p_U^*] < \mathbb{E}[p_{PD}^*(c)]$

ii) There will be more trade

iii) If $n \geq 5$ or $\bar{c}$ is sufficiently small:

a) The average traded price is lower under no information:

$$\bar{p}_U^{trade} = p_U^* < \bar{p}_{PD}^{trade} := \frac{\mathbb{E}\left[p_{PD}^*(C)\left(1-F(p_{PD}^*(C))^n\right)\right]}{\mathbb{E}[1-F(p_{PD}^*(C))^n]}.$$

b) Aggregate consumer surplus satisfies $CS_U > CS_{PD}$

With finitely many firms, the intuition from the continuum case, that pooling types with different price sensitivities lowers the average price, remains. However, different prices also induce different purchase probabilities: low-search-cost consumers facing low prices are more likely to buy. This heterogeneity in purchase probabilities makes the comparison of average traded prices less immediate. When there are sufficiently many firms, the gap in purchase probabilities and prices becomes small enough that the direct privacy effect from the benchmark setting dominates, so the average traded price falls and aggregate consumer surplus rises under no information.

7 Discussion

This section interprets the main results and places them in the broader literature on price discrimination and privacy. In the sequential search environment studied above, observability of heterogeneous search costs determines prices via an inverse price-sensitivity rule. Consequently, pooling types through privacy lowers the average price but also distorts participation incentives, creating a tension between the benefits of pooling and the externalities of participation. The effect on consumer surplus can therefore be decomposed into a direct privacy effect for always-participating consumers, a direct participation effect for selectively active consumers, and a participation externality operating through the equilibrium price response to additional entry. These mechanisms differ in important ways from the more familiar case in which firms discriminate based on willingness to pay (WTP) rather than search costs, and they have distinct implications for how regulators should evaluate the welfare consequences of more detailed consumer information.

Relationship to WTP Discrimination. Search-cost discrimination differs from WTP-based discrimination on both the intensive and extensive margins. Holding participation fixed, more information about search costs lowers prices for low-search-cost (more price-sensitive) consumers and raises prices for high-search-cost (less price-sensitive) consumers, shifting surplus from the latter to the former. By contrast, richer WTP information primarily facilitates rent extraction (Choi and Rocheteau, 2024) and, in the limit, delivers the Diamond outcome with no search on the equilibrium path and full rent extraction (Diamond, 1971). On the Hotelling line, WTP discrimination can intensify competition because firms react to rivals' prices: Thisse and Vives (1988) show that, under full coverage, it raises consumer surplus by enabling firms to match competitors' offers. By contrast, in the present model with full coverage, privacy and coarse signals benefit consumers because pooling types lowers the average price.

The extensive-margin logic is also distinct in partially covered markets. With WTP discrimination, firms can profitably offer low prices to marginal consumers who would otherwise not participate, thereby expanding participation; consumer gains then depend

on coverage (Rhodes and Zhou, 2024). With search-cost discrimination, firms face a commitment problem: conditional on entry, marginal high-search-cost consumers are charged high prices even though firms would like to commit ex ante to lower prices to induce participation. Coarsening information about search costs partially relaxes this problem and raises participation and overall welfare. Its effect on consumer surplus, however, depends on the balance between the direct privacy effect, the direct participation effect, and the participation externality.

These conclusions parallel monopoly results on WTP discrimination. In Bergemann et al. (2015), when all gains from trade are realized, more information shifts surplus from consumers to the monopolist, consistent with Proposition 1. When trade is inefficient, the effect on consumer surplus is ambiguous, consistent with Proposition 3.

In practice, consumer information often contains both WTP and search costs, with potentially positive or negative correlation. For example, a higher opportunity cost of time may imply higher search costs and be associated with higher WTP, but it may also be associated with sophistication and lower effective search costs. Section 6.1 shows that this correlation affects participation. When WTP is observed nearly perfectly, Diamond logic applies, and search-cost discrimination becomes irrelevant. When WTP information is weak, pooling increases participation as in the baseline model. If search costs are strongly positively correlated with WTP, the incentives generated by WTP-based discrimination can dominate, leading to reduced participation when coarse signals are used. Even then, pooling price sensitivities lowers the average price under full coverage.

Generality. The model's key ingredients are unit demand and an inverse price-sensitivity rule: in symmetric equilibrium, the price charged to a consumer group is inversely related to that group's average price sensitivity. Price sensitivity is pinned down by exogenous type through the search cost. Hence, the direct privacy effect does not depend on the search protocol: pooling types replaces a schedule of discriminatory prices with a single price equal to the (harmonic) mean.

This mechanism extends to other search environments. In Mauring (2025), consumers have either zero or positive search costs and, with some probability, firms observe a signal

of the consumer's cost type. The two types differ in price sensitivity, so pooling remains beneficial: never revealing the signal is preferred to always revealing it. The article focuses on private signals and shows that partial revelation can raise consumer surplus. When the signal is absent with positive probability, a seller may offer a low price to avoid losing a low-search-cost consumer who would otherwise keep searching.

The pooling effect also appears outside search. In Armstrong (2006)'s fully covered Hotelling model, consumers differ in transportation costs (or preference intensity), which determine price sensitivity. Coarser information that pools transportation-cost types lowers equilibrium prices and increases consumer surplus. The argument relies on unit demand and on prices, mainly reallocating surplus rather than changing quantities. Extending Armstrong (2006) to allow selective participation would preserve the applicability of the decomposition (Proposition 3) and of the associated information design exercise, although details may differ.

In richer environments with selective participation, the same trade-off continues to shape optimal information structures. Under the condition that the least price-sensitive consumers also generate the lowest welfare and consumer surplus (for example, because they face high search or transportation costs), consumer-surplus-maximizing policies should take simple cutoff forms that pool types on one side of a threshold and separate those that impose the largest participation externalities.

Observability of the Signal s_i. Consumer i is assumed to observe s_i (and c_i) when deciding whether to initiate search.[12] If s_i is deterministic in c_i, this is without loss: consumers can infer the implied equilibrium price $p(s_i)$ from their type, so participation depends only on that price.

When s_i is stochastic, making it observable weakly increases consumer surplus relative to keeping it unobserved. To see this, consider an alternative timing in which the firm observes s_i after the consumer pays c_i and posts $p = p(s_i)$, while the consumer never observes s_i and only sees the posted price. Suppose signals are stochastic: $p(s_i) \in$

[12]This is also consistent with data-access rights (e.g., GDPR/CCPA), which may allow consumers to infer the data used for personalization.

$\{p_L, p_H\}$ with $p_L < p_H$ and each outcome occurring with strictly positive probability. Consider a consumer with costs c_i who initiates search only because $\Pr(p_L) > 0$ but, due to high search costs, would not search again once they expect future prices to be p_H. In a symmetric equilibrium, observing p_H at the first firm reveals the high-signal state, so the consumer expects subsequent firms also to charge p_H and does not continue searching. Thus, the first firm faces less elastic demand when posting p_H. Anticipating this, a firm that privately observes the low signal may have an incentive to deviate and raise the price from p_L to p_H: upon observing p_H, the consumer infers the high state and stops searching. Hence, firms have an incentive to charge higher prices, so equilibrium prices are weakly higher under non-observability, which reduces consumer surplus. Consequently, making s_i observable weakly increases aggregate consumer surplus.

Policy and redistribution. The analysis establishes an instrumental value of privacy. In fully covered markets, hiding search-cost information maximizes consumer surplus because pooling types lowers the average traded price. In partially covered markets, a coarse signal maximizes consumer surplus. In contrast, privacy (or sufficiently coarse signals) increases welfare by expanding participation relative to the fully informative setting; within the class of monotone partitions, the uninformative signal is welfare-maximizing.

Because current privacy regimes such as GDPR and CCPA give consumers control over data sharing (via opt-in or opt-out mechanisms), standard unraveling in this work's model implies that voluntary disclosure yields near-full disclosure, preventing the price and participation benefits of privacy from being realized.[13] This highlights the importance of data-use limitations and intermediaries in managing user data. If browsers or platforms place sufficient weight on aggregate consumer surplus or on maximizing participation, they have an incentive to restrict data sharing and implement coarse signals that approximate the uninformative or consumer-optimal information structures. This pattern aligns with current industry practice: Chrome's Privacy Sandbox API limits targeting to broad interest categories, and Safari and Firefox have long restricted cross-site

[13]This is consistent with descriptive evidence that cookie acceptance rates are high and appear to have increased over time, although there are many other possible explanations and there is no clear consensus on trends in tracking acceptance.

tracking cookies. Similarly, Amazon, despite having ample information about consumers' search behavior, refrains from sharing that data with third-party sellers.[14] The model thus offers a rationale for why intermediaries with access to rich data often share relatively little with third parties and why, despite the abundance of consumer data, personalized price discrimination appears less prevalent than one might expect.

These policy considerations have a distributional dimension. Privacy redistributes surplus not only between firms and consumers in aggregate, but also across consumers. In fully covered markets, consumers with low search costs face lower prices under search-cost-based discrimination but higher prices under privacy, while high-search-cost consumers experience the opposite pattern. With partial coverage, marginal high-search-cost participants gain from privacy, but their participation raises prices for always-participating consumers through the participation externality. Whether stronger privacy is socially desirable, therefore, depends on how search costs correlate with welfare weights. If search costs reflect the opportunity cost of time, then consumers with low search costs are likely "time-rich" but "income-poor." Under price discrimination, these consumers pay low prices. Privacy pools them with high-search-cost ("time-poor," "income-rich") consumers, thereby raising prices for lower-income consumers. Thus, while privacy raises total welfare, it may have adverse distributional consequences. This contrasts with much of the discussion around WTP-based discrimination, where allowing price discrimination can enable participation by low-WTP (and often less well-off) consumers who would otherwise be priced out of the market, as emphasized by Council of Economic Advisers (2015).

8 Conclusion

This article examines price discrimination based on heterogeneous consumer search costs in a sequential search environment, à la Wolinsky (1986). When consumers differ in their search costs, those costs induce them to differ in their price sensitivity. In equilibrium,

[14] Gutierrez Gallardo (2022) estimates that Amazon places substantial weight on consumer surplus relative to industry profits.

prices follow an inverse price-sensitivity rule. Hence, if firms cannot observe search costs perfectly (privacy), the impact on consumer surplus is determined by three forces. The direct privacy effect increases consumer surplus because average prices are lower when based on average price sensitivity rather than individual price sensitivity. The novel feature of selective participation by higher-search-cost consumers generates two additional forces. Privacy can induce high-search-cost consumers to participate by lowering the price they face relative to discriminatory pricing; their participation yields a direct effect on aggregate consumer surplus and welfare. At the same time, because prices depend on the price sensitivity of all participating consumers, additional high-cost participants exert a negative participation externality on the price.

In markets where consumers' search costs are sufficiently low, all consumers participate regardless of the information structure (full coverage), so privacy unambiguously increases consumer surplus. When some consumers have high search costs, participation depends on the information firms observe (partial coverage). Privacy increases total welfare by raising participation. However, prices may rise enough that the participation externality outweighs the gains from increased participation and the direct privacy effect, reducing aggregate consumer surplus.

The optimal information policy balances these forces. In partially covered markets, the consumer–surplus-maximizing policy is neither complete privacy nor perfect information, but a coarse signal. One signal pools all consumers whose search costs are low enough to be active, while the other signal deters other consumers from participating. This prevents high-search–cost types from imposing negative price externalities on active consumers while preserving the price benefits from pooling among active types.

When consumers can opt into data sharing and therefore can voluntarily share their data with firms, standard unraveling implies that neither outright privacy nor the optimal coarse policy is sustainable in equilibrium. The benefits of privacy, greater participation, and, in fully covered markets, lower average prices need not materialize under privacy laws that permit voluntary sharing (such as GDPR or CCPA). The model suggests a rationale for stricter privacy laws or limits on the use of search-cost information. Fur-

thermore, it highlights the importance of intermediaries, such as platforms and browsers, in maximizing consumer surplus and welfare. If platforms place substantial weight on participation or aggregate consumer surplus, they are incentivized to implement coarse information structures that improve welfare and benefit consumers.

These platform design issues raise further research questions: under what conditions will platforms share data with third-party sellers, and how does platform competition influence data-sharing incentives?

The analysis shows that discrimination based on search costs differs qualitatively from oligopolistic WTP discrimination. Contrary to Thisse and Vives (1988), consumers in aggregate benefit from privacy in fully covered markets. While WTP discrimination can stimulate participation, search-cost-based price discrimination reduces participation. With search-cost-based price discrimination, privacy serves as a commitment device, ensuring that consumers with a low value of participating in the market are not charged high prices. This contrasts with WTP discrimination, where privacy prevents charging lower prices to consumers with low valuations, thereby hindering their participation.

These differences have empirical and policy relevance. It is important for regulators to determine the extent to which observed personalization is based on WTP versus search costs, given their different welfare implications. In addition, the incentives under WTP-based price discrimination differ from those under search-cost-based discrimination in ways that extend beyond the mechanisms studied here, warranting further examination. For example, the incentives to invest in data acquisition under search-cost-based price discrimination appear in simulations to be substitutes rather than complements, in contrast to WTP discrimination.

References

Ali, S. N., G. Lewis, and S. Vasserman (2023). Voluntary disclosure and personalized pricing. *Review of Economic Studies 90*(2), 538–571.

Anderson, S., A. Balik, and N. Larson (2023). Price discrimination in the information

age: Prices, poaching, and privacy with personalized targeted discounts. *Review of Economic Studies 90*(5), 2085–2115.

Anderson, S. P. and R. Renault (1999). Pricing, product diversity, and search costs: A Bertrand-Chamberlin-Diamond model. *RAND Journal of Economics 30*(4), 719–735.

Armstrong, M. (2006). Recent developments in the economics of price discrimination. In R. Blundell, W. K. Newey, and T. Persson (Eds.), *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, Volume 2, pp. 97–141. Cambridge University Press.

Armstrong, M. (2017). Ordered consumer search. *Journal of the European Economic Association 15*(5), 989–1024.

Armstrong, M. and J. Vickers (2001). Competitive price discrimination. *RAND Journal of Economics 32*(4), 579–605.

Armstrong, M. and J. Zhou (2016). Search deterrence. *Review of Economic Studies 83*(1), 26–57.

Bergemann, D., B. Brooks, and S. Morris (2015). The limits of price discrimination. *American Economic Review 105*(3), 921–957.

Bergemann, D., B. Brooks, and S. Morris (2021). Search, information, and prices. *Journal of Political Economy 129*(8), 2275–2319.

Blackwell, D. (1953). Equivalent comparisons of experiments. *The Annals of Mathematical Statistics 24*(2), 265–272.

Blake, T., C. Nosko, and S. Tadelis (2016). Returns to consumer search: Evidence from eBay. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 531–545.

Braghieri, L. (2019). Targeted advertising and price discrimination in intermediated online markets. SSRN Working Paper 3072692, SSRN.

Chen, Y. and G. Iyer (2002). Research note: Consumer addressability and customized pricing. *Marketing Science 21*(2), 197–208.

Choi, M. and G. Rocheteau (2024). Information acquisition and price discrimination in dynamic, decentralized markets. *Review of Economic Dynamics 53*, 1–46.

CMA (2018). Pricing algorithms: Economic working paper on the use of algorithms to facilitate collusion and personalised pricing. Economic working paper, Competition and Markets Authority, London.

Conitzer, V., C. R. Taylor, and L. Wagman (2012). Hide and seek: Costly consumer privacy in a market with repeat purchases. *Marketing Science 31*(2), 277–292.

Council of Economic Advisers (2015). Big data and differential pricing. Report, Council of Economic Advisers, Executive Office of the President, Washington, DC.

De Cornière, A. and R. De Nijs (2016). Online advertising and privacy. *The RAND Journal of Economics 47*(1), 48–72.

De los Santos, B. (2018). Consumer search on the Internet. *International Journal of Industrial Organization 58*, 66–105.

Diamond, P. A. (1971). A model of price adjustment. *Journal of Economic Theory 3*(2), 156–168.

Esteves, R.-B. (2014). Price discrimination with private and imperfect information. *The Scandinavian Journal of Economics 116*(3), 766–796.

Esteves, R.-B. and J. Resende (2016). Competitive targeted advertising with price discrimination. *Marketing Science 35*(4), 576–587.

European Commission (2018). Consumer market study on online market segmentation through personalised pricing/offers in the european union. Consumer Market Study EAHC/2013/CP/04, Request for Specific Services 2016 85 02, Publications Office of the European Union, Luxembourg.

European Parliament (2022). Personalised pricing. Study PE 734.008, European Parliament, Policy Department for Economic, Scientific and Quality of Life Policies, Luxembourg.

Fabra, N. and M. Reguant (2020). A model of search with price discrimination. *European Economic Review 129*, 103571.

FTC (2024). Issue spotlight: The rise of surveillance pricing. Issue spotlight, Federal Trade Commission, Washington, DC.

Groh, C.-C. (2026). Competitive price discrimination, imperfect information, and consumer search. *Journal of Economic Theory 231*, 106133.

Groh, C.-C. and M. Preuss (2025). Tracking, price discrimination, and disclosure in sequential search markets. SSRN Working Paper 4249099, SSRN.

Groh, C.-C. and J. von Wangenheim (2024). Pigou meets Wolinsky: Search, price discrimination, and consumer sophistication. Discussion Paper 527, CRC TR 224 Discussion Paper Series.

Gutierrez Gallardo, G. (2022). The welfare consequences of regulating Amazon. SSRN Working Paper 3965566, SSRN. Latest version July 21, 2022.

Hannak, A., G. Soeller, D. Lazer, A. Mislove, and C. Wilson (2014). Measuring price discrimination and steering on e-commerce web sites. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, pp. 305–318.

Hong, H. and M. Shum (2006). Using price distributions to estimate search costs. *The RAND Journal of Economics 37*(2), 257–275.

Ichihashi, S. (2020). Online privacy and information disclosure by consumers. *American Economic Review 110*(2), 569–595.

Jolivet, G. and H. Turon (2018). Consumer search costs and preferences on the Internet. *Review of Economic Studies 86*(3), 1258–1300.

Klenke, A. (2013). Probability measures on product spaces. In *Probability Theory: A Comprehensive Course*, pp. 273–293. Springer.

Kohn, M. G. and S. Shavell (1974). The theory of search. *Journal of Economic Theory 9*(2), 93–123.

Larson, J., S. Mattu, and J. Angwin (2015). Unintended consequences of geographic targeting. *Technology Science 2015090103*.

Lauermann, S. (2012). Asymmetric information in bilateral trade and in markets: An inversion result. *Journal of Economic Theory 147*(5), 1969–1997.

Luther, L. (2025). *Essays in Economic Theory and Industrial Organization*. Ph. D. thesis, Rheinische Friedrich-Wilhelms-Universität Bonn.

Matzkin, R. L. (2007). Nonparametric identification. In J. J. Heckman and E. E. Leamer (Eds.), *Handbook of Econometrics*, Volume 6, pp. 5307–5368. Elsevier.

Mauring, E. (2025). Search and price discrimination online. *Journal of Political Economy Microeconomics 3*(2), 138–176.

Mauring, E. and C. Williams (2023). Coasian dynamics in sequential search. CEPR Discussion Paper 17907, Centre for Economic Policy Research.

Milgrom, P. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics 12*(2), 380–391.

Montes, R., W. Sand-Zantman, and T. Valletti (2019). The value of personal information in online markets with endogenous privacy. *Management Science 65*(3), 1342–1362.

Moraga-González, J. L., Z. Sándor, and M. R. Wildenbeest (2017). Prices and heterogeneous search costs. *The RAND Journal of Economics 48*(1), 125–146.

Moraga-González, J. L., Z. Sándor, and M. R. Wildenbeest (2023). Consumer search and prices in the automobile market. *The Review of Economic Studies 90*(3), 1394–1440.

OECD (2018). Personalised pricing in the digital era. Background Note DAF/COMP(2018)13, Organisation for Economic Co-operation and Development.

Ofcom (2020). Personalised pricing for communications: Making data work for consumers. Discussion paper, Ofcom, London.

Pigou, A. C. (1920). *The Economics of Welfare*. London: Macmillan.

Preuss, M. (2023). Search, learning, and tracking. *The RAND Journal of Economics 54*(1), 54–82.

Rhodes, A. and J. Zhou (2024). Personalized pricing and competition. *American Economic Review 114*(7), 2141–2170.

Shaffer, G. and Z. J. Zhang (2002). Competitive one-to-one promotions. *Management Science 48*(9), 1143–1160.

Stahl, D. O. (1989). Oligopolistic pricing with sequential consumer search. *American Economic Review 79*(4), 700–712.

Thisse, J.-F. and X. Vives (1988). On the strategic choice of spatial price policy. *American Economic Review 78*(1), 122–137.

Ursu, R. M., Q. Zhang, and E. Honka (2023). Search gaps and consumer fatigue. *Marketing Science 42*(1), 110–136.

Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica 47*(3), 641–654.

Wolinsky, A. (1986). True monopolistic competition as a result of imperfect information. *The Quarterly Journal of Economics 101*(3), 493–511.

A Proofs of Sections 2 to 5

Proof of Lemma 1. **Outline.** This proof follows Luther (2025) and Moraga-González et al. (2017). Whenever G has no atom at 0 and $G|s$ has no atom at the bottom of the support, I will show that there exists a unique symmetric equilibrium whenever there is an active consumer. First, Claim 1 shows that any firm prefers to charge a price

weakly below the monopoly price. This implies that in equilibrium, the equilibrium price cannot be above the monopoly price. Then, Claim 2 shows that there exists a unique candidate equilibrium. This candidate equilibrium satisfies both that the price must be optimal, i.e. satisfies a first-order condition, and that participation is consistent with the candidate price. Lastly, Claim 3 shows that the demand function of each firm is concave, so that profits are maximized at the first-order condition, establishing that the candidate equilibrium is the unique symmetric equilibrium.

Claim 1
When $1-F$ is log-concave, then $p^* \in [0, p^m]$, where p^m is the solution to $p^m = \frac{1-F(p^m)}{f(p^m)}$.

Proof. Suppose that $p > p^m$. Profits are

$$\begin{aligned} pD(p, p^* \mid s) &= \int_{\underline{c}}^{c_0(p^*)} p \frac{1 - F(\hat{x}(c) + p - p^*)}{1 - F(\hat{x}(c))} \, dG(c \mid s) \\ &< \int_{\underline{c}}^{c_0(p^*)} p^m \frac{1 - F(\hat{x}(c) + p^m - p^*)}{1 - F(\hat{x}(c))} \, dG(c \mid s) = p^m D(p^m, p^* \mid s) \end{aligned}$$

The inequality holds because $p(1-F(p+z))$ is decreasing in p for $z \geq 0$ and $p \geq p^m$. To see this, observe that the derivative is

$$1 - F(p+z) - f(p+z)p = f(p+z)\left(\frac{1-F(p+z)}{f(p+z)} - p\right) \leq f(p+z)\left(\frac{1-F(p)}{f(p)} - p\right) < 0$$

The first inequality is due to the log-concavity of $1-F$, which implies a decreasing reverse hazard rate. The second inequality follows because $p > p^m$ and p^m renders the third expression exactly equal to 0.

This concludes the proof. For more details, see Lemma 1.3 and Lemma 1.5 in Luther (2025). □

Claim 2
When $1-F$ is log-concave and $G|s$ has no atom except possibly at the bottom of the support, then the equilibrium candidate price is

$$p^* = \frac{G(c_0(p^*)|s)}{\int_{\underline{c}}^{c_0(p^*)} \frac{f(\hat{x}(c))}{1-F(\hat{x}(c))} \, dG(c \mid s)}$$

and unique.

Proof. Consider a potential equilibrium price p^* and assume that the equilibrium exhibits participation.

The firm's profit is the price times the demand that the firm faces when charging, p while all others charge p^*. Following the demand derived in Appendix OA.1 for $p \leq p^*$, the profit is

$$p \int_{\underline{c}}^{c_0(p^*)} \frac{1 - F(\hat{x}(c) + p - p^*)}{1 - F(\hat{x}(c))} dG(c \mid s)$$

In equilibrium at $p = p^*$, the first-order condition for profit maximization can be expressed

as

$$p^* \;=\; \frac{G(c_0(p^*)|s)}{\int_{\underline{c}}^{c_0(p^*)} \dfrac{f(\hat{x}(c))}{1-F(\hat{x}(c))}\, dG(c \mid s)}. \tag{6}$$

because the derivatives of demand for $p > p^*$ and $p < p^*$ coincide at $p = p^*$ when there is no atom at 0: $G(\{0\}|s) = 0$ which holds for almost all signals due to Assumption 2.

It remains to show that the solution is unique. A challenge is that it involves a fixed point. The price depends on the participation cutoff $c_0(p^*)$, while the individual optimal price depends on $c_0(p^*)$.

Case 1: $\hat{x}(\underline{c}) < \frac{1-F(\hat{x}(\underline{c}))}{f(\hat{x}(\underline{c}))}$ where $\underline{c}$ is the infimum of the support of $G \mid s$. Then, there is no equilibrium with active participation.

Case 2: At the equilibrium price, all consumers participate $G(c_0(p^*)|s) = 1$. Then, the FOC gives an explicit (and unique) solution for p^*.

Case 3: There is only partial participation given signal s: $G(c_0(p^*)|s) < 1$. Only consumers with $\hat{x}(c) \geq p \iff c \leq c_0(p^*)$ participate. One must show that there exists a unique fixed point p^* that is consistent with the induced participation cutoff $c_0(p^*)$.

By Claim 1, one can restrict attention to $p \in [0, p^m]$ since all higher prices are strictly dominated.

$$L(p) \;:=\; \int_C \left[1 \;-\; p\,\frac{f\big(\hat{x}(c)\big)}{1-F\big(\hat{x}(c)\big)}\right] \mathbf{1}\{\hat{x}(c) \geq p\}\, \mathrm{d}G(c \mid s).$$

Then $L(0) = \int_C \mathbf{1}\{\hat{x}(c) \geq 0\}\, dG(c \mid s) > 0$ (if some participation is possible) and

$$L(p^m) = \int_C \left[1 - p^m\,\frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}\right] \mathbf{1}\{\hat{x}(c) \geq p^m\}\, \mathrm{d}G(c \mid s) \;\leq\; 0,$$

The inequality is because $x \mapsto \frac{f(x)}{1-F(x)}$ is nondecreasing (log-concavity of $1-F$) and $p^m = \frac{1-F(p^m)}{f(p^m)}$. By the no-atoms except at the bottom assumption, $p \mapsto L(p)$ is continuous; hence by the intermediate value theorem, there exists $p^* \in [0, p^m]$ such that $L(p^*) = 0$.

Fix $0 \leq p_1 < p_2 \leq p^m$. Using $\mathbf{1}\{\hat{x} \geq p_1\} = \mathbf{1}\{\hat{x} \geq p_2\} + \mathbf{1}\{p_1 \leq \hat{x} < p_2\}$,

$$\begin{aligned}
L(p_2) - L(p_1) &= \int_C \left[1 - p_2\,\frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}\right] \mathbf{1}\{\hat{x}(c) \geq p_2\}\, \mathrm{d}G(c \mid s) \\
&\quad - \int_C \left[1 - p_1\,\frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}\right] \mathbf{1}\{\hat{x}(c) \geq p_1\}\, \mathrm{d}G(c \mid s) \\
&= -(p_2 - p_1) \int_C \frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}\, \mathbf{1}\{\hat{x}(c) \geq p_2\}\, \mathrm{d}G(c \mid s) \\
&\quad - \int_C \left[1 - p_1\,\frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}\right] \mathbf{1}\{p_1 \leq \hat{x}(c) < p_2\}\, \mathrm{d}G(c \mid s) \;\leq\; 0,
\end{aligned}$$

because, on $\{p_1 \leq \hat{x} < p_2\}$,

$$p_1\,\frac{f(\hat{x}(c))}{1-F(\hat{x}(c))} \;\leq\; p_1\,\frac{f(p^m)}{1-F(p^m)} \;=\; \frac{p_1}{p^m} \;\leq\; 1.$$

Notice that whenever not everyone participates, if $p_1 < p^* < p_2$, the set $\{p_1 < \hat{x}(c) < p_2\}$ has positive measure (due to $\hat{x}$ being continuous in c and full support).

Hence, L is nonincreasing on $[0, p^m]$ and strictly decreasing around the root. Since $L(0) > 0$ and $L(p^m) \leq 0$ and L is continuous, this implies a unique root. □

Claim 3
The profit function $\Pi(p, p^* \mid s) = pD(p, p^* \mid s)$ is strictly concave in p over $p \in [0, p^m]$ for all p^*

Proof. $x(1 - F(x))$ is concave by assumption, which implies that the second derivative is strictly negative: $-2f(x) - xf'(x) < 0 \quad \forall x \in [0, \overline{v}]$.

The firm's revenue from a consumer with search cost c who has received signal s is $p\frac{1-F(\hat{x}(c)+p-p^*)}{1-F(\hat{x}(c))} = \frac{1}{1-F(\hat{x}(c))}\, p\big(1 - F(\hat{x}(c) + p - p^*)\big)$.

To show concavity, notice that the first derivative of demand is continuous. This means that for the points where the second-order derivative of demand does not exist, the first derivative does not jump. The same holds for the profit function. It remains to show that the second-order derivative is negative almost everywhere. By Appendix OA.1, the demand is twice differentiable almost everywhere.

Using the concavity of $x(1 - F(x))$, which implies

$$-2f(x) - xf'(x) \leq 0 \quad \forall x \in [0, \overline{v}],$$

the profit is concave in p. For $p < p^*$, the second derivative of the profit function satisfies

$$\begin{aligned} &2\frac{\partial D(p, p^*)}{\partial p} + p\frac{\partial^2 D(p, p^*)}{\partial p^2} \\ &= \int_C \frac{-2f\big(\hat{x}(c)+p-p^*\big) - pf'\big(\hat{x}(c)+p-p^*\big)}{1-F(\hat{x}(c))} \mathbb{1}\{\hat{x}(c) \geq p^*\}\mathbb{1}\{\hat{x}(c)+p-p^* \leq \overline{v}\}\, dG(c|s) \leq 0, \end{aligned} \tag{12}$$

since for $y = \hat{x}(c) + p - p^*$ we have $y \geq p$ and

$$-2f(y) - pf'(y) = \big(-2f(y) - yf'(y)\big) + (y - p)f'(y) \leq 0, \tag{13}$$

where the first term is nonpositive by concavity of $x(1 - F(x))$, and the second is nonpositive when $f'(y) \leq 0$, while if $f'(y) \geq 0$ then already $-2f(y) - pf'(y) \leq 0$.

When $p > p^*$

$$\begin{aligned} &2\frac{\partial D(p, p^*)}{\partial p} + p\frac{\partial^2 D(p, p^*)}{\partial^2 p} \\ &= \int_C \frac{-2f(\hat{x}(c) + p - p^*) - pf'(\hat{x}(c) + p - p^*)}{1 - F(\hat{x}(c))} \mathbb{1}\{\hat{x}(c) \geq p^*\}\mathbb{1}\{\hat{x}(c) + p - p^* \leq \overline{v}\}\, dG(c \mid s) \\ &\quad + p\, \frac{f(\overline{v})}{1 - F\big(\overline{v} - p + p^*\big)}\, \frac{\partial}{\partial p}\, G\big(\{c : \hat{x}(c) \leq \overline{v} - p + p^*\} \mid s\big) \leq 0 \end{aligned} \tag{14}$$

The first part is pointwise nonpositive and thus nonpositive overall, as argued previously. The second term, resulting from the change in the boundary, is negative. A fortiori, the entire term is nonpositive.

Hence, all terms of the second derivative of the profit function are negative. Therefore, the second derivative of the profit function is negative almost everywhere (except at $p = p^*$).

The profit function is therefore continuously differentiable everywhere, twice differentiable almost everywhere, and the second derivative is strictly negative almost everywhere (except possibly at $p = p^*$). Hence, the profit function is strictly concave. □

This concludes the proof. □

Proof of Lemma 2. **Part 1:**Recall that $p^*_{PD}(c) = \left(\frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}\right)^{-1}$ and that $p^*_U = \left(\mathbb{E}_{c\sim G}\left[\frac{f(\hat{x}(c))}{1-F(\hat{x}(c))} \;\middle|\; \hat{x}(c) \geq p^*_U\right]\right)^{-1}$. This means it is the harmonic mean over the set of active consumers, those that satisfy $\hat{x}(c) \geq p^*_U$.
Part 2: Suppose that the set $\{c : \hat{x}(c) \geq p^*_U\}$ is non-degenerate. Because consumers follow a threshold strategy as discussed in section 2.1, the set of active consumers is $[\underline{c}, c_0(p^*_U)]$.

Because p^*_{PD} is strictly increasing and c is nondegenerate on $[\underline{c}, c_0(p^*_U)]$,

$$p^*_{PD}(\underline{c}) \leq p^*_{PD}(c) \leq p^*_{PD}(c_0(p^*_U)). \tag{15}$$

Taking the expectation, and observing that for $\mathbb{P}\big(p^*_{PD}(\underline{c}) < p^*_{PD}(c) < p^*_{PD}(c_0(p^*_U))\big) > 0$ implies

$$\frac{1}{p^*_{PD}(c_0(p^*_U))} \;<\; \mathbb{E}_{c\sim G}\left[\frac{1}{p^*_{PD}(c)} \;\middle|\; c \leq c_0(p^*_U)\right] \;<\; \frac{1}{p^*_{PD}(\underline{c})}. \tag{16}$$

Taking reciprocals (all terms are > 0) reverses the inequalities and yields the claim.$p^*_{PD}(\underline{c}) < p^*_U < p^*_{PD}\big(c_0(p^*_U)\big)$.
Part 3: Assume $c_0(p^*_{PD}(\underline{c})) > \underline{c}$. If, toward a contradiction, $c_0(p^*_U) = \underline{c}$, then by Part 1 applied to the (singleton) participant set,

$$\frac{1}{p^*_U} = \frac{f(\hat{x}(\underline{c}))}{1-F(\hat{x}(\underline{c}))} = \frac{1}{p^*_{PD}(\underline{c})} \quad \Rightarrow \quad p^*_U = p^*_{PD}(\underline{c}), \tag{17}$$

which implies $c_0(p^*_U) = c_0(p^*_{PD}(\underline{c})) > \underline{c}$, a contradiction. Hence $c_0(p^*_U) > \underline{c}$.

□

Proof of Lemma 3 . By Lemma 2 part (2), if $c^U_0 > \underline{c}$ then $p^*_U < p^*_{PD}(c^U_0)$. Because $\hat{x}(c^U_0) = p^*_U$, it implies $\hat{x}(c^U_0) < p^*_{PD}(c^U_0)$. $\hat{x}(c) - p^*_{PD}(c)$ is strictly decreasing in c, because $\hat{x}(c)$ is strictly decreasing in c and p^*_{PD} is strictly increasing in c. $\hat{x}(c) - p^*_{PD}(c)$ equals zero at c^{PD}_0 which implies $c^{PD}_0 < c^U_0$.

In the corner cases $c^{PD}_0 = \underline{c}$ (no PD participation) or $c^{PD}_0 = \overline{c}$ (full participation), only the weak inequality $c^{PD}_0 \leq c^U_0$ holds. (i) If $c^{PD}_0 = \underline{c}$, then $\hat{x}(\underline{c}) \leq p^*_{PD}(\underline{c}) \leq p^*_U$, by the weak version of Lemma 2. Hence, no type above $\underline{c}$ can participate because $\hat{x}(c)$ is strictly decreasing which implies $\hat{x}(c) < p^*_U \quad \forall c > \underline{c}$ (ii) If $c^{PD}_0 = \overline{c}$ (full PD participation), then for all c we have $\hat{x}(c) \geq p^*_{PD}(c)$. As p^*_U is the harmonic mean of $\{p^*_{PD}(c)\}$ over $[\underline{c}, \overline{c}]$, it satisfies $p^*_{PD}(\underline{c}) \leq p^*_U \leq p^*_{PD}(\overline{c})$. Hence $\hat{x}(\overline{c}) \geq p^*_{PD}(\overline{c}) \geq p^*_U$, so $c^U_0 = \overline{c}$ and equality holds: $c^{PD}_0 = c^U_0 = \overline{c}$. $\hat{x}(c)$ is decreasing in c, so all consumers with $c \leq \overline{c}$ also participate ($\hat{x}(c) \geq p^*_U$). □

Proof of Proposition 3. Partition consumers into the always-active $[\underline{c}, c^{PD}_0]$ and the participants $[c^{PD}_0, c^U_0]$, and add–subtract p^{before}_U on $[\underline{c}, c^{PD}_0]$. This yields exactly the displayed three-term decomposition. For the direct privacy effect, apply Proposition 1 to the fixed set $\{c \leq c^{PD}_0\}$: holding participation fixed, the uniform price p^{before}_U is weakly below the average of $p^*_{PD}(c)$ on this set. Hence $\int_{\underline{c}}^{c^{PD}_0} (p^{PD}(c) - p^{before}_U) g(c)\, dc \geq 0$. Define, for any

cutoff $m \in [\underline{c}, \bar{c}]$, $p_U(m) := \left(\mathbb{E}\left[\frac{1}{p^*_{PD}(C)} \,\middle|\, C \le m\right]\right)^{-1}$, the price for coverage $\{C \le m\}$ with an uninformative signal. Because $c \mapsto \frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}$ is strictly decreasing, $m \mapsto p_U(m)$ is strictly increasing. Therefore, $p_U^{after} = p_U(c_0^U) > p_U^{before} = p_U(c_0^{PD})$ whenever $c_0^{PD} < c_0^U$. This establishes the participation externality.

For every $c \in [c_0^{PD}, c_0^U]$, participation under p_U^{after} requires $\hat{x}(c) \ge p_U^{after}$, so $\int_{c_0^{PD}}^{c_0^U} (\hat{x}(c) - p_U^{after}) g(c)\, dc \ge 0$. □

Proof of Proposition 4. With $v_{ij} \sim U[0, \bar{v}]$, $\hat{x}(c) = \bar{v} - \sqrt{2\bar{v}\, c}$. Under PD, $p^{PD}(c) = \sqrt{2\bar{v}\, c}$ and the marginal condition $\hat{x}(c_0^{PD}) = p^{PD}(c_0^{PD})$ gives $c_0^{PD} = \min\{\bar{c}, \bar{v}/8\}$. Under privacy, $p_U^* = \sqrt{\bar{v}\, c_0^U/2}$ and $\hat{x}(c_0^U) = p_U^*$ gives $c_0^U = \min\{\bar{c}, 2\bar{v}/9\}$.

For any coverage $t \le \bar{c}$, integrating demand minus price over $c \in [0, t]$ yields $CS_{PD}(t) = \frac{\bar{v}\, t - \frac{4}{3}\sqrt{2\bar{v}}\, t^{3/2}}{\bar{c}}$ and $CS_U(t) = \frac{\bar{v}\, t - \frac{7}{3\sqrt{2}}\sqrt{\bar{v}}\, t^{3/2}}{\bar{c}}$. (1) If $\bar{c} \le \bar{v}/8$, then $c_0^{PD} = c_0^U = \bar{c}$ and $CS_U - CS_{PD} = \frac{1}{3\sqrt{2}}\sqrt{\bar{v}\, \bar{c}} > 0$. (2) If $\bar{c} \ge 2\bar{v}/9$, then $c_0^U = 2\bar{v}/9$ and $c_0^{PD} = \bar{v}/8$, so $CS_U - CS_{PD} = (\frac{4}{81} - \frac{1}{24})\bar{v}^2/\bar{c} = \frac{5}{648}\,\bar{v}^2/\bar{c} > 0$. (3) If $\bar{v}/8 \le \bar{c} \le 2\bar{v}/9$, then $c_0^{PD} = \bar{v}/8$, $c_0^U = \bar{c}$ and $\Delta CS(\bar{c}) = \bar{v} - \frac{7}{3\sqrt{2}}\sqrt{\bar{v}\, \bar{c}} - \frac{\bar{v}^2}{24\bar{c}}$ is concave on this interval because $\Delta CS''(\bar{c}) = \frac{\bar{v}^2}{24\bar{c}^3}\big(-2 + 7\sqrt{2}(\bar{c}/\bar{v})^{3/2}\big) < 0$ (because $(\bar{c}/\bar{v})^{3/2} \le (2/9)^{3/2} = 2\sqrt{2}/27$ makes the bracket $-2 + 28/27 < 0$). Hence, the minimum is at the endpoints, where (1)–(2) gives positivity. Therefore $CS_U > CS_{PD}$ for all $\bar{v}, \bar{c} > 0$.

Profits follow by integrating prices over coverage: $\Pi_{PD} = \frac{2}{3}\sqrt{2\bar{v}}\,(c_0^{PD})^{3/2}/\bar{c}$ and $\Pi_U = \sqrt{\bar{v}/2}\,(c_0^U)^{3/2}/\bar{c}$. Thus $\Pi_{PD} \ge \Pi_U$ iff $c_0^U/c_0^{PD} \le (4/3)^{2/3}$. As $c_0^{PD} = \min\{\bar{c}, \bar{v}/8\}$ and $c_0^U = \min\{\bar{c}, 2\bar{v}/9\}$, the ratio $\min\{\bar{c}, 2\bar{v}/9\}/\min\{\bar{c}, \bar{v}/8\}$ is nondecreasing in $\bar{c}$ (equal to 1 for $\bar{c} \le \bar{v}/8$, linear $= 8\bar{c}/\bar{v}$ on $[\bar{v}/8, 2\bar{v}/9]$, and $= 16/9$ thereafter), so there is a unique switch at $\bar{c}/\bar{v} = (1/8)(4/3)^{2/3} = 6^{1/3}/12$. Hence $\Pi_{PD} \ge \Pi_U$ iff $\bar{c}/\bar{v} \le 6^{1/3}/12$. □

Proof of Lemma 4. **Prices:** Under a monotone partition, $S = s_k$ implies $c \in I_k$. If $k < \ell$, then I_k lies (weakly) to the left of I_ℓ, so every $c \in I_k$ is strictly smaller than every $c' \in I_\ell$. Since $p^*_{PD}(c)$ is strictly increasing in c, it follows that

$$\frac{1}{p^*_{PD}(c)} \;>\; \frac{1}{p^*_{PD}(c')} \qquad \text{for all } c \in I_k,\ c' \in I_\ell.$$

Conditioning on participation only restricts attention to subsets of these intervals, namely $I_k \cap \{\hat{x}(c) \ge p_\sigma(s_k)\} \subseteq I_k$ and $I_\ell \cap \{\hat{x}(c) \ge p_\sigma(s_\ell)\} \subseteq I_\ell$, so the ordering is preserved. Therefore,

$$\mathbb{E}\left[\frac{1}{p^*_{PD}(c)} \,\middle|\, S = s_k,\ \hat{x}(c) \ge p_\sigma(s_k)\right] > \mathbb{E}\left[\frac{1}{p^*_{PD}(c)} \,\middle|\, S = s_\ell,\ \hat{x}(c) \ge p_\sigma(s_\ell)\right] \tag{18}$$

whenever both conditioning events have positive probability, and hence $p_\sigma(s_k) < p_\sigma(s_\ell)$. If an interval has no participation, the conditional expectations are undefined and prices are irrelevant, as no trade takes place.

Participation: Signals are a deterministic function of c, $\sigma(c)$, and prices are a piecewise constant function of s. Higher costs are associated with a weakly higher index, and a higher index is associated with a weakly higher price by part (i). Since $\hat{x}$ is decreasing in c, the net surplus

$$c \;\longmapsto\; \hat{x}(c) - p_\sigma(\sigma(c))$$

is weakly decreasing. Therefore, participation follows a cutoff $c_0^m \in [\underline{c}, \bar{c}]$: consumers with $c \le c_0^m$ are active and those with $c > c_0^m$ are inactive.

Blackwell order: For two information structures σ' and σ with finite signal spaces S', S, we say $\sigma' \succeq_B \sigma$ if σ is a garbling of σ', i.e. there exists a Markov kernel $\gamma : S \times S' \to [0,1]$ such that for all $c \in C$ and $s \in S$,

$$\sigma(c, \{s\}) = \sum_{s' \in S'} \gamma(s, s')\, \sigma'(c, \{s'\}). \tag{19}$$

Sufficiency. Assume that σ' is obtained by subdividing the intervals of σ: for every $s' \in S'$ there exists an (a.s. unique) $\phi(s') \in S$ with $I'_{s'} \subseteq I_{\phi(s')}$. For $s \in S$, $s' \in S'$, define

$$\gamma(s, s') := \mathbb{1}\{s = \phi(s')\}. \tag{20}$$

Because σ' is a partition, for each c there is a unique $s'_c \in S'$ with $c \in I'_{s'_c}$, so $\sigma'(c, \{s'\}) = \mathbb{1}\{s' = s'_c\}$ and

$$\sum_{s' \in S'} \gamma(s, s')\, \sigma'(c, \{s'\}) = \gamma(s, s'_c) = \mathbb{1}\{s = \phi(s'_c)\}. \tag{21}$$

By construction $I'_{s'_c} \subseteq I_{\phi(s'_c)}$, so $c \in I_{\phi(s'_c)}$. Since $\{I_s\}_{s \in S}$ is a partition of C, for this c there is a unique $s(c) \in S$ with $c \in I_{s(c)}$. Hence $s = \phi(s'_c)$ if and only if $c \in I_s$. Therefore,

$$\mathbb{1}\{s = \phi(s'_c)\} = \mathbb{1}\{c \in I_s\} = \sigma(c, \{s\}).$$

Thus, for each c, the probability measures $\sigma(c, \cdot)$ and $\sum_{s' \in S'} \gamma(\cdot, s')\, \sigma'(c, \{s'\})$ on the finite set S coincide on all singletons, and therefore on all subsets. Hence σ is a garbling of σ', so $\sigma' \succeq_B \sigma$.

Necessity. Conversely, suppose $\sigma' \succeq_B \sigma$, so there exists $\gamma : S \times S' \to [0,1]$ satisfying equation 19. Since σ' is a monotone partition, for each $s' \in S'$ there is an interval $I'_{s'}$ with $\sigma'(c, \{s'\}) = \mathbb{1}\{c \in I'_{s'}\}$. Fix $s' \in S'$ and any $c \in I'_{s'}$. Then $\sigma'(c, \{s''\}) = \mathbb{1}\{s'' = s'\}$, and equation 19 reduces to

$$\sigma(c, \{s\}) = \gamma(s, s') \qquad \forall\, s \in S. \tag{22}$$

Thus $\sigma(c, \cdot)$ is the same for all $c \in I'_{s'}$. Because σ is also a monotone partition, there is for each c a unique $s(c) \in S$ with $\sigma(c, \{s(c)\}) = 1$, so all $c \in I'_{s'}$ share the same signal; denote it by $\phi(s') \in S$. Hence $\sigma(c, \{\phi(s')\}) = 1$ for all $c \in I'_{s'}$, which means $I'_{s'} \subseteq I_{\phi(s')}$ (up to G-null boundaries). Therefore each σ'–interval is contained in a σ–interval, i.e. σ' is obtained by subdividing the intervals of σ. □

Proof of Lemma 5. For $\tau \in \{\sigma, \sigma'\}$ the price follows the standard expression:

$$p^*_\tau(s) \;:=\; \Big(\mathbb{E}_{c \sim G_A} \big[\, 1/p^*_{PD}(c) \,\big|\, S_\tau = s \,\big] \Big)^{-1}, \tag{23}$$

which is finite and positive as $1/p^*_{PD}(c) > 0$ a.s.[15]. Because $\sigma = \gamma \circ \sigma'$ and given the conditioning on the same active event, the active posterior under σ is a mixture of active

[15] Because $G(0) = 0$, there can be only a set of signals with measure 0 that has an atom at $c = 0$ and thus $\frac{1}{p^*_{PD}(0)} = 0$. Because all prices must be finite, the measure 0 set of signals that does not satisfy $\frac{1}{p^*_{PD}(c)}$ being finite a.s. can be excluded from the analysis without impacting the results.

posteriors under σ':

$$G_A(\cdot \mid S = s) \;=\; \int_{S'} G_A(\cdot \mid S' = s')\, q(\mathrm{d}s' \mid s) \qquad \text{for } M_\sigma^A\text{-a.e. .} \tag{24}$$

Writing $k(s) := \mathbb{E}_{c\sim G_A}\big[\,1/p^*_{PD}(c)\,\big|\,S = s\big]$, $k(s') := \mathbb{E}_{c\sim G_A}\big[\,1/p^*_{PD}(c)\,\big|\,S' = s'\big]$, it implies $k(s) = \int k(s')\, q(\mathrm{d}s' \mid s)$. The map $x \mapsto 1/x$ is convex on $(0,\infty)$, hence by Jensen's inequality,

$$p^*_\sigma(s) \;=\; \frac{1}{k(s)} \;\le\; \int \frac{1}{k(s')}\, q(\mathrm{d}s' \mid s) \;=\; \int p_{\sigma'}(s')\, q(\mathrm{d}s' \mid s), \quad \text{for } M_\sigma^A\text{-a.e. } s, \tag{25}$$

with strict inequality precisely when $\mathrm{Var}_{s'\sim q(\cdot|s)}(p^*_{\sigma'}(s')) > 0$.

Taking expectations with respect to M_σ^A and using the tower property under $\mathbb{P}^A$ gives

$$\mathbb{E}_{s\sim M_\sigma^A}\big[p_\sigma(s)\big] \;\le\; \mathbb{E}_{s\sim M_\sigma^A}\Big[\mathbb{E}_{s'\sim q(\cdot|s)}\big[p^*_{\sigma'}(s')\big]\Big] \;=\; \mathbb{E}_{s'\sim M_{\sigma'}^A}\big[p^*_{\sigma'}(s')\big], \tag{26}$$

which proves the weak inequality in the statement. The strictness criterion follows from the strict Jensen condition holding on a set of s of positive M_σ^A-measure. □

Proof of Corollary 1. The set of active consumers does not change regardless of the information structure. The price under any information structure σ is weakly below $p^*_{PD}(\bar{c})$ and $\bar{c}$ participates. This means:

$$0 \le \hat{x}(\bar{c}) - p^*_{PD}(\bar{c}) \le \hat{x}(\bar{c}) - p^*_\sigma(s) \le \hat{x}(c) - p^*_\sigma(s) \forall c, s \tag{27}$$

Hence, all consumers participate regardless of the information regime. □

Proof of Corollary 2. **Part i):** With full coverage, $A_\sigma = A_{\sigma'} = A = C$ and everyone is active under any information structure because

$$0 \le \hat{x}(\bar{c}) - p^*_{PD}(\bar{c}) \le \hat{x}(\bar{c}) - p^*_\sigma(s) \le \hat{x}(c) - p^*_\sigma(s) \qquad \forall c, s. \tag{28}$$

Hence, since $\sigma' \succ_B \sigma$ and the refinement is strict, there exists an interval I_k of σ with $G(I_k) > 0$ that is strictly subdivided by σ' into subintervals of positive measure, inducing (at least) two refined signals with distinct prices. By Lemma 4, prices are strictly increasing across (refined) intervals, and since the active set is the same under σ and σ', Lemma 5 (strict case) implies $\mathbb{E}_{M_\sigma^A}[p_\sigma] < \mathbb{E}_{M_{\sigma'}^A}[p_{\sigma'}]$. Under full coverage, consumer surplus equals welfare minus expected price, so $CS_{\sigma'} < CS_\sigma$.

Part ii): Let c_0 and c_0' be the participation cutoffs under σ and σ' (cutoff existence under monotone partitions follows from Lemma 4). Suppose for contradiction that $c_0' > c_0$ and set $t := c_0'$. Let $J = [\underline{J}, \overline{J}]$ be the σ'-cell containing t. Since σ' refines σ, there is a σ-cell $I = [\underline{I}, \overline{I}]$ with $J \subseteq I$, hence $\underline{J} \ge \underline{I}$.

In cell J, active types form the initial segment $[\underline{J}, t]$, so by the equilibrium price formula (standard expression, conditional on participation),

$$p^*_{\sigma'}(\sigma'(t)) = \left(\mathbb{E}\left[\frac{1}{p^*_{PD}(C)} \;\middle|\; C \in [\underline{J}, t]\right]\right)^{-1}.$$

Let $u(c) := 1/p^*_{PD}(c)$, which is decreasing. Then shifting the left endpoint right and

expanding the right endpoint can only weakly lower the conditional mean, so

$$\mathbb{E}[u(C) \mid C \in [\underline{J}, t]] \leq \mathbb{E}[u(C) \mid C \in [\underline{I}, t]] \leq \mathbb{E}[u(C) \mid C \in [\underline{I}, c_0]].$$

Taking inverses gives

$$p^*_{\sigma'}(\sigma'(t)) \geq (\mathbb{E}[u(C) \mid C \in [\underline{I}, c_0]])^{-1} = p^*_{\sigma}(\sigma(c_0)) = \hat{x}(c_0),$$

where the last equality uses that c_0 is marginal under σ. Since $\hat{x}$ is decreasing and $t > c_0$, it must be that $\hat{x}(t) < \hat{x}(c_0) \leq p^*_{\sigma'}(\sigma'(t))$, contradicting that t participates under σ'. Hence $c'_0 \leq c_0$.

□

Proof of Proposition 5. **Construct a refinement that raises CS:** Fix a signal s with $\Pr(S = s) > 0$ such that $G(\cdot \mid s)$ has full support on $[\underline{c}_s, \overline{c}_s]$ and the market is not fully covered under s, i.e., $\hat{x}(\overline{c}_s) - p_\sigma(s) < 0$ or fix a collection of signals with positive measure. Let $h(c) := \frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}$, which is strictly decreasing in c, and

$$p_\sigma(s) = \Big(\mathbb{E}_{c\sim G_A}[h(c) \mid S = s]\Big)^{-1}. \tag{29}$$

Write $c^s_0 := c_0(p^*_\sigma(s))$ for the highest active type under σ and signal s.

To prove the first result, that consumer surplus strictly increases, the proof will proceed by first introducing a signal that separates low and high search cost consumers from the set of active consumers, and will then compare the consumer surplus when the low signal incorporates ever more consumers.

For $\alpha \in (0, 1)$, let $q_s(1-\alpha)$ denote the $(1-\alpha)$–quantile of $C \mid S = s$ restricted to the originally active set $\{C \leq c^s_0\}$, and define

$$H_\alpha(s) := \{q_s(1-\alpha) \leq C \leq c^s_0\}, \qquad R_\alpha(s) := \{C < q_s(1-\alpha)\} \cap \{C \leq c^s_0\}. \tag{30}$$

Refine s into (s, H) and (s, R) by revealing the indicator $\mathbf{1}_{\{C\in H_\alpha(s)\}}$; assign all originally inactive types $C > c^s_0$ the signal (s, H) (so they remain inactive). Dropping the indicator garbles back to s, hence the refinement is strict. Let $m := \mathbb{E}_{c\sim G_A}[h(c) \mid s]$, $\quad m_H(\alpha) := \mathbb{E}_{c\sim G_A}[h(c) \mid s,\, c \in H_\alpha(s)]$, $\quad m_R(\alpha) := \mathbb{E}_{c\sim G_A}[h(c) \mid s,\, c \in R_\alpha(s)]$.

As h is strictly decreasing and $H_\alpha(s)$ is the top slice, $m_R(\alpha) > m > m_H(\alpha)$, so

$$p_{\sigma'}(s, R) = \frac{1}{m_R(\alpha)} < \frac{1}{m} = p_\sigma(s) < \frac{1}{m_H(\alpha)} = p_{\sigma'}(s, H). \tag{31}$$

Moreover, using $m = (1-\alpha)m_R(\alpha) + \alpha m_H(\alpha)$ and $h > 0$,

$$p_\sigma(s) - p_{\sigma'}(s, R) = \frac{1}{m} - \frac{1}{m_R(\alpha)} = \frac{\alpha\,[\,m - m_H(\alpha)\,]}{(1-\alpha)\,m\,m_R(\alpha)} \;\geq\; \frac{\alpha\,[\,m - m_H(\alpha)\,]}{m^2}. \tag{$*$}$$

Only consumers who previously saw s are affected.

Gains on $R_\alpha(s)$. Each $c \in R_\alpha(s)$ remains active because $p_{\sigma'}(s, R) < p_\sigma(s)$, and gains at least the common price drop $(*)$. The aggregate gain on $R_\alpha(s)$ is thus $\geq (p_\sigma(s) - p_{\sigma'}(s, R))\,(1-\alpha)$.

Losses on $H_\alpha(s)$. Before the refinement, the surplus of $c \in H_\alpha(s)$ was $\Delta(c, s) :=$

$\hat{x}(c) - p_\sigma(s) \le \eta(\alpha)$ where

$$\eta(\alpha) := \sup_{u \in [q_s(1-\alpha),\, c_0^s]} \big(\hat{x}(u) - p_\sigma(s)\big) \xrightarrow[\alpha \downarrow 0]{} 0 \tag{32}$$

(by continuity of $\hat{x}$ and $q_s(1-\alpha) \uparrow c_0^s$). After the refinement, these types face the higher price $p_{\sigma'}(s, H)$, and any who exit have ex–post surplus 0, so the aggregate loss on $H_\alpha(s)$ is at most $\eta(\alpha)\,\alpha$.

Because $R_\alpha(s)$ has positive mass and h is strictly decreasing, $m > \inf_{u \le c_0^s} h(u) = h(c_0^s)$. As $\alpha \downarrow 0$, $m_H(\alpha) \to h(c_0^s)$, so there exists $\bar{\alpha} \in (0,1)$ with

$$m - m_H(\alpha) \;\ge\; \tfrac{1}{2}\big(m - h(c_0^s)\big) \qquad \forall\, \alpha \in (0, \bar{\alpha}]. \tag{33}$$

Take $\alpha \le \min\{\bar{\alpha},\, 1/2\}$. Combining with $(*)$ gives

$$\underbrace{(p_\sigma(s) - p_{\sigma'}(s,R))(1-\alpha)}_{\text{gain on } R_\alpha(s)} \;\ge\; \frac{\alpha}{2} \cdot \frac{m - h(c_0^s)}{m^2}. \tag{34}$$

Finally, choose α even smaller so that $\eta(\alpha) \le \frac{1}{4}\,\frac{m-h(c_0^s)}{m^2}$. Then

$$CS_{\sigma'} - CS_\sigma \;\ge\; \frac{\alpha}{2} \cdot \frac{m - h(c_0^s)}{m^2} \;-\; \eta(\alpha)\,\alpha \;\ge\; \frac{\alpha}{4} \cdot \frac{m - h(c_0^s)}{m^2} \;>\; 0. \tag{35}$$

No consumer in $R_\alpha(s)$ exits (their price falls), and some consumers in $H_\alpha(s)$ exit. This is because the price rises and the distribution has full support, so there will be a positive measure of consumers refraining from participating under the new signal structure.

Non-monotone CS: Set $\sigma_1 = \sigma$ and let $\sigma_2 \succ_B \sigma_1$ be the local refinement constructed above, so $CS(\sigma_2) > CS(\sigma_1)$. Refine further to σ_3 by fully revealing c for all $c \ge c_0^{PD}$, and finally to σ_4 by fully revealing all types. Each step is a strict Blackwell refinement. Either $CS(\sigma_3) < CS(\sigma_2)$, or else $CS(\sigma_4) < CS(\sigma_3)$ by Lemma 5; in either case, the sequence $(\sigma_n)_{n=1}^4$ has consumer surplus that first increases and then decreases, so $(CS(\sigma_n))_{n \ge 1}$ is non-monotone. □

Proof of Proposition 6. **Outline:** The proof proceeds in three steps: First, I show that it is optimal to pool the set of active consumers together. Then, the problem can be expressed as a maximization problem over the set of participating consumers. This yields the optimal information structure. Lastly, it is checked that the information structure can be implemented with the desired active set of consumers.

Step 1 (Pooling improves consumer surplus). Fix σ and let $A_\sigma := \{(c,s) : \hat{x}(c) \ge p_\sigma^*(s)\}$ be the set of participating consumers under σ. Define the average traded price $\bar{p}_\sigma := \mathbb{E}[p_\sigma^*(S) \,|\, (C,S) \in A_\sigma]$. Construct σ' by assigning a common signal $\ell \notin S$ to all $(c,s) \in A_\sigma$, while leaving signals unchanged for $(c,s) \notin A_\sigma$.

Then,

$$\begin{aligned} CS_\sigma &= \mathbb{E}\big[\mathbf{1}\{\hat{x}(C) \geq p^*_\sigma(S)\}\big(\hat{x}(C) - p^*_\sigma(S)\big)\big] \\ &= \mathbb{E}\big[\mathbf{1}_{A_\sigma}\big(\hat{x}(C) - \bar{p}_\sigma\big)\big] \\ &\leq \mathbb{E}\Big[\mathbf{1}_{A_\sigma}\big(\hat{x}(C) - \bar{p}_\sigma\big)_+\Big] \\ &\leq \mathbb{E}\Big[\mathbf{1}_{A_\sigma}\big(\hat{x}(C) - p^*_{\sigma'}(\ell)\big)_+\Big] \\ &\leq \mathbb{E}\big[\mathbf{1}\{\hat{x}(C) \geq p^*_{\sigma'}(S')\}\big(\hat{x}(C) - p^*_{\sigma'}(S')\big)\big] = CS_{\sigma'}. \end{aligned} \tag{36}$$

The first equality is the definition of consumer surplus. The second uses that replacing $p^*_\sigma(S)$ by its conditional mean on A_σ leaves total payments on A_σ unchanged, i.e. $\mathbb{E}[\mathbf{1}_{A_\sigma} p^*_\sigma(S)] = \mathbb{P}(A_\sigma)\bar{p}_\sigma$. The first inequality uses $(y)_+ \geq y$. The last inequality holds because under σ' consumers outside A_σ can always opt out, so their contribution to $CS_{\sigma'}$ is weakly nonnegative.

The key step is the middle inequality in equation (36), which follows from the claim $p^*_{\sigma'}(\ell) \leq \bar{p}_\sigma$. The claim is true because of the following: Let $h(c) := \frac{f(\hat{x}(c))}{1-F(\hat{x}(c))} = \frac{1}{p^*_{PD}(c)}$. By the equilibrium pricing formula, for any signal s that induces participation, $p^*_\sigma(s) = \left(\mathbb{E}[h(C) \,|\, S = s,\ (C,S) \in A_\sigma]\right)^{-1}$. Hence,

$$\begin{aligned} \bar{p}_\sigma &= \mathbb{E}\big[\left(\mathbb{E}[h(C) \,|\, S,\ (C,S) \in A_\sigma]\right)^{-1} \,\big|\, (C,S) \in A_\sigma\big] \\ &\geq \left(\mathbb{E}[h(C) \,|\, (C,S) \in A_\sigma]\right)^{-1} \\ &\geq \left(\mathbb{E}[h(C) \,|\, (C,S) \in A_\sigma,\ \hat{x}(C) \geq p^*_{\sigma'}(\ell)]\right)^{-1} = p^*_{\sigma'}(\ell), \end{aligned} \tag{37}$$

where the first inequality is Jensen's (since $x \mapsto 1/x$ is convex), and the second uses that h is decreasing in c (equivalently $p^*_{PD}(c)$ increasing), so conditioning on the participation event $\{\hat{x}(C) \geq p^*_{\sigma'}(\ell)\}$ truncates C to lower costs and weakly increases the conditional mean of h. This proves $p^*_{\sigma'}(\ell) \leq \bar{p}_\sigma$ and therefore completes Step 1.

Step 2: Optimal participation set Given that all active consumers receive the same signal under the optimal information structure, the problem is to find the set of active consumers. The maximization problem is

$$\max_{\alpha(c)} \quad CS(\alpha) = \int_{\underline{c}}^{\overline{c}} \big(\hat{x}(c) - p(\alpha)\big)\, \alpha(c)\, dG(c), \tag{38}$$

$$\text{s.t.} \quad 0 \leq \alpha(c) \leq 1, \quad \forall c \in [\underline{c}, \overline{c}], \tag{39}$$

where $\alpha(c) \in [0,1]$ represents the participation rate of a consumer with search cost c, and where

$$p(\alpha) = \frac{\int \alpha(c)\, dG(c)}{\int \frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}\, \alpha(c)\, dG(c)}. \tag{40}$$

Consider a perturbation $\alpha_\epsilon = \alpha + \epsilon\eta$, with η integrable. Then

$$\frac{d}{d\epsilon} CS(\alpha_\epsilon)\Big|_{\epsilon=0} = \int [\hat{x}(c) - p(\alpha_\epsilon)]\, \eta(c)\, dG(c) - \int \alpha_\epsilon(c)\, dG(c) \frac{d}{d\epsilon}\Big(p(\alpha_\epsilon)\Big)\Big|_{\epsilon=0}, \tag{41}$$

$$\int \alpha_\epsilon(c)\, dG(c) \frac{d}{d\epsilon}\Big(p(\alpha_\epsilon)\Big)\Big|_{\epsilon=0} = \int \alpha(c)\, dG(c) \left(\frac{\int \eta(c)\, dG(c)}{\int \frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}\, \alpha(c)\, dG(c)} - \frac{\int \alpha(c)\, dG(c)}{\left(\int \frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}\, \alpha(c)\, dG(c)\right)^2} \int \frac{f(\hat{x}(c))}{1-F(\hat{x}(c))} \eta(c)\, dG(c) \right)$$

Collecting terms,

$$dCS(\alpha;\eta) = \int \left[\hat{x}(c) - 2p(\alpha) + p(\alpha)^2 \tfrac{f(\hat{x}(c))}{1-F(\hat{x}(c))}\right]\eta(c)\, dG(c). \tag{42}$$

Optimality implies

$$\hat{x}(c) - 2p(\alpha) + p(\alpha)^2 \tfrac{f(\hat{x}(c))}{1-F(\hat{x}(c))} \begin{cases} \geq 0, & \text{if } \alpha(c) = 1, \\ \leq 0, & \text{if } \alpha(c) = 0, \\ = 0, & \text{if } \alpha(c) \in (0,1) \end{cases} \tag{43}$$

For fixed α, the inner term is strictly decreasing in c as $\hat{x}(c)$ and $\frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}$ are decreasing. Hence, there exists a unique cutoff m such that

$$\alpha^*(c) = \begin{cases} 1, & c < m, \\ 0, & c > m, \end{cases} \tag{44}$$

with $\alpha^*(m) \in [0,1]$. The cutoff m^* is determined by

$$\hat{x}(m^*) - 2p_U(m^*) + p_U(m^*)^2 \frac{f(\hat{x}(m^*))}{1 - F(\hat{x}(m^*))} = 0 \tag{45}$$

with the previously defined $p_U(m) := \frac{1}{\mathbb{E}[\frac{f(\hat{x}(C))}{1-F(\hat{x}(C))} | C \leq m]}$. The cutoff m^* is unique because equation 45 has a unique solution. To observe that, consider the derivative with respect to m of the LHS.

$$\hat{x}'(m) + p_U(m)^2 \frac{d}{dm}\left(\frac{f(\hat{x}(m))}{1 - F(\hat{x}(m))}\right) + 2p_U'(m)\left(p_U(m)\frac{f(\hat{x}(m))}{1 - F(\hat{x}(m))} - 1\right) < 0$$

The inequality follows because $p_U(m^*)/p^*_{PD}(m^*) < 1$ since $p_U(m^*)$ is the harmonic mean of the p^*_{PD} and because $p_U'(m^*) > 0$, $\hat{x}'(m^*) < 0$, and $1/p^*_{PD}(c) = \left(\frac{f}{1-F} \circ \hat{x}\right)(c)$ is strictly decreasing in c.

The objective $CS(\alpha)$ is concave in α, so the global maximum is attained:

$$CS(\alpha) = C(\alpha) - \frac{A(\alpha)^2}{B(\alpha)}$$
$$C(\alpha) = \int \hat{x}(c)\alpha(c)dG(c), \quad A(\alpha) = \int \alpha(c)dG(c), \quad B(\alpha) = \int h(c)\alpha(c)dG(c)$$

A, B and C are linear in α, taking on the quadratic-over-linear form and therefore $CS(\alpha)$ is concave, so that the FOC yields the global optimum.

Notice that $c_0^{PD} \leq m^* \leq c_0^U$ and that the first inequality is strict whenever $c_0^{PD} < \overline{c}$. The second inequality is also strict when $c_0^U < \overline{c}$.

To see that $c_0^{PD} < m^*$ for $c_0^{PD} < \overline{c}$, suppose for contradiction that $m^* \leq c_0^{PD}$. Then:

$$\hat{x}(m^*) - 2p_U(m^*) + p_U(m^*)^2 \frac{f(\hat{x}(m^*))}{1 - F(\hat{x}(m^*))} \geq \hat{x}(c_0^{PD}) - 2p_U(m^*) + \frac{p_U(m^*)^2}{p^*_{PD}(c_0^{PD})} \tag{46}$$
$$> \hat{x}(c_0^{PD}) - p^*_{PD}(c_0^{PD}) = 0$$

This violates the optimality condition: the LHS is continuous in m^*, so there exists a $m' > m^*$ where $\hat{x}(m') - 2p_U(m') + p_U(m')^2 \frac{f(\hat{x}(m'))}{1-F(\hat{x}(m'))} > 0$, but $\alpha(m') = 0$.
The strict inequality holds because

$$-2p_U(m^*) + \frac{p_U(m^*)^2}{p^*_{PD}(c_0^{PD})} > -p^*_{PD}(c_0^{PD}) \tag{47}$$
$$\iff \left(p_U(m^*) - p^*_{PD}(c_0^{PD})\right)^2 > 0 \tag{48}$$

The inequality holds strictly whenever $p^*(m)$ is defined by pooling different types, i.e. $m^* > \underline{c}$.

Further, suppose that $m^* \geq c_0^U$. Then:

$$\hat{x}(m^*) - 2p_U(m^*) + \frac{p_U(m^*)^2}{p^*_{PD}(m^*)} < \hat{x}(m^*) - p_U(m^*) < \hat{x}(m^*) - p^*_{PD}(m^*) \leq 0 \tag{49}$$

This yields a contradiction with the optimality condition. The first inequality is strict because p^*_{PD} is strictly larger than the price when all lower types are pooled together, see Lemma 2.

Step 3: Feasibility Check

1. Active Types: The optimality condition ensures that $\hat{x}(c) - p_U(m^*) > 0$ for all $c \leq m^*$.

2. Non-Participation of Others: If others would participate were they to receive a separate signal, this would induce a minimum price of $p^*_{PD}(m^*) \geq p^*_{PD}(c_0^{PD})$ by Lemma 2. By construction, no consumer $c > c_0^{PD}$ would participate under such a price because c_0^{PD} consumers are already indifferent ($\hat{x}(c_0^{PD}) - p^*_{PD}(c_0^{PD}) = 0$).

□

Proof of Example 3. With $v \sim U[0, \bar{v}]$, the cutoff rule implies $\hat{x}(c) = \bar{v} - \sqrt{2c\bar{v}}$ and, at $\hat{x}(c)$, the inverse hazard is $\frac{f(\hat{x}(c))}{1-F(\hat{x}(c))} = 1/\sqrt{2c\bar{v}}$. Under a pooling cutoff m, the price that implements m is $p(m) = \frac{\int_0^m 1\,dc}{\int_0^m \frac{1}{\sqrt{2c\bar{v}}}\,dc} = \sqrt{\frac{m\bar{v}}{2}}$. The first-order condition at m is $\hat{x}(m) - 2p(m) + p(m)^2 \cdot \frac{1}{\sqrt{2m\bar{v}}} = 0$, i.e. $\bar{v} - \sqrt{2m\bar{v}} - 2\sqrt{\frac{m\bar{v}}{2}} + \frac{m\bar{v}}{2\sqrt{2m\bar{v}}} = \bar{v} - \frac{7}{2\sqrt{2}}\sqrt{m\bar{v}} = 0$. Hence $\sqrt{m\bar{v}} = \frac{2\sqrt{2}}{7}\bar{v}$ and so $m = \frac{8}{49}\bar{v}$. As c is supported on $[0, \bar{c}]$, the effective cutoff is $m^* = \min\{\bar{c}, \frac{8}{49}\bar{v}\}$. □

Proof of Proposition 7. **Welfare:** Recall that $W = \mathbb{E}_{c,s}[\mathbb{1}\{\hat{x}(c) \geq p^*_\sigma(s)\}\hat{x}(c)]$. Any informative monotone partition refines complete pooling, and by Corollary 2 a refinement weakly lowers participation. Welfare increases in the participation set as all consumers entering satisfy $\hat{x}(c) \geq p_\sigma(s) > 0$, hence complete pooling maximizes welfare.
Industry profits: *Full coverage.* If $c_0^{PD} = \bar{c}$, then all consumers participate under any monotone partition, so welfare, the sum of consumer and producer surplus, is unaffected.

Since consumer surplus falls with informativeness under full coverage (Corollary 2), profits are maximized by the most informative structure, i.e. σ^{PD}; hence $m^{**} = \bar{c}$.

Partial coverage. Fix any monotone partition σ with cutoff c_0^σ and let I_k be the cell containing c_0^σ. Consider two changes:

Step 1 (pooling the top). Let $I_k = [t_{k-1}, t_k]$ be the cell that contains the participation cutoff c_0^σ. Consider the coarsening that merges all cells $I_k, I_{k+1}, \ldots$ into a single top cell $\tilde{I}_k = [t_{k-1}, \bar{c}]$ (leaving $I_1, \ldots, I_{k-1}$ unchanged).

If c_0^σ lies in the interior of I_k, then all types in $I_{k+1}, I_{k+2}, \ldots$ are inactive, so this pooling does not change the posterior over active consumers and therefore does not change the equilibrium prices or profits.

If instead $c_0^\sigma = t_k$ is at the upper endpoint of I_k, then pooling I_{k+1} into I_k strictly coarsens the information structure at the margin. By Corollary 2, a coarsening weakly increases participation (the cutoff weakly rises). Moreover, the price attached to the pooled top signal is weakly higher (it is the equilibrium price for a signal whose posterior shifts weight to higher-cost types). Hence, industry profit weakly increases: either nothing changes, or additional consumers enter and pay a larger price.

Step 2 (refine the sure participants). After Step 1, there is a single "top" signal for all $c > t_{k-1}$. All consumers with $c \leq t_{k-1}$ are strictly below the participation cutoff, hence active irrespective of the details of partitioning of the cost-space $\{c < t_{k-1}\}$. Refining the partition on $[\underline{c}, t_{k-1}]$ leaves (a) the set of arriving consumers and (b) the posterior and price associated with the top signal unchanged. By Lemma 5 a refinement weakly increases industry profits. Taking the refinement to the limit (working in the closure of monotone partitions) yields perfect revelation on $[\underline{c}, m]$ with $m = t_{k-1}$.

Thus, any monotone partition is weakly dominated by an *upper-censorship* policy: perfectly reveal all $c \leq m$ and pool all $c > m$ into one signal.

Existence of an optimum. Let $\Pi(m)$ denote the equilibrium industry profit under this upper-censorship policy. Since $m \in [\underline{c}, \bar{c}]$ and $\Pi(m)$ varies continuously with m, Weierstrass' theorem implies that $\Pi(m)$ attains a maximizer $m^{**} \in [\underline{c}, \bar{c}]$. □

Proof of Proposition 8. Let each consumer with type c choose a report $r \in \{\emptyset\} \cup [\underline{c}, \bar{c}]$, where $r = c$ denotes full disclosure and $r = \emptyset$ denotes non-disclosure. After reporting, the consumer decides whether to participate. Let $G(\cdot \mid r)$ be the Bayes posterior over c induced by the disclosure strategy, and let $G_A(\cdot \mid r)$ denote the posterior over active types, i.e. $G(\cdot \mid r)$ truncated to those c that participate under report r.

For any report r that induces positive participation, the equilibrium price $p(r)$ is characterized as in Lemma 1 using the posterior over arriving types:

$$p(r) = \left(\mathbb{E}[h(c) \mid r,\ \hat{x}(c) \geq p(r)]\right)^{-1} = \left(\mathbb{E}_{c \sim G_A(\cdot|r)}\left[\frac{1}{p^*_{PD}(c)}\right]\right)^{-1},$$

where the second equality uses $p^*_{PD}(c) = 1/h(c)$ under full disclosure.

Recall $p^*_{PD}(c)$ is strictly increasing on the set of participating types. Consider any report r that pools a non-degenerate set of active types. Let c_L be the lowest cost in the support of $G_A(\cdot \mid r)$. Since r pools at least two active types and p^*_{PD} is strictly increasing, $p^*_{PD}(c) > p^*_{PD}(c_L)$ with positive probability under $G_A(\cdot \mid r)$. Hence the pooled price satisfies $p(r) > p^*_{PD}(c_L)$. Hence, the consumer c_L payoff under disclosure is strictly higher. This contradicts the optimality of r for type c_L.

Therefore, no non-degenerate pool of active types can be sustained: disclosure unravels among participants, and on-path prices coincide with $p^*_{PD}(c)$. Types who do not

participate are payoff-irrelevant. Thus, the conclusion holds for all participants and for almost all active types. □

Online Appendix

OA.1 Demand and Profit Calculations

This section examines the derivation of the demand and equilibrium profits in greater detail, following Moraga-González et al. (2017).

Suppose there is a symmetric equilibrium, one in which firms charge the same price. For a consumer, it is only profitable to continue the search when the expected value of an additional search is not negative. The value of an additional search when the current best item is x is

$$\int_x^{\overline{v}} (v - x) dF(v)$$

while the search costs are c.

Upon visiting a firm charging price p while expecting all other firms to charge p^*, the consumer compares buying to continuing to search. The consumer stops searching and purchases iff $v_i - p \geq \hat{x}(c) - p^*$, equivalently iff $v_i \geq \hat{x}(c) + p - p^*$; otherwise the consumer continues searching.

Thus, for a consumer with search cost c, the probability that a consumer at firm j stops is $\Pr(v_{ij} \geq \hat{x}(c) + p - p^*) = 1 - F(\hat{x}(c) + p - p^*)$. The probability that the consumer arrives at firm j in the course of the search process is the probability that firm j is the nth firm in the search sequence and the consumer has not stopped before, summed over n: $\sum_{n=1}^{\infty} F(\hat{x}(c))^{n-1} = \frac{1}{1-F(\hat{x}(c))}$.

The demand from a given consumer who is active is therefore:

$$\frac{1 - F(\hat{x}(c) + p - p^*)}{1 - F(\hat{x}(c))}$$

Moraga-González et al. (2017) shows that the expected consumer surplus for a consumer with search costs c in a market with a price p is

$$CS(c) = \hat{x}(c) - p$$

Consequently, consumers with search costs $c \leq c_0(p^*)$ will enter. This quantity is endogenous, and for a more detailed discussion, see Moraga-González et al. (2017).

The resulting demand is therefore:

$$D(p, p^*) = \int_C \frac{1 - F(\hat{x}(c) + p - p^*)}{1 - F(\hat{x}(c))} \mathbb{1}\{\hat{x}(c) \geq p^*\} \mathbb{1}\{\hat{x}(c) + p - p^* \leq \bar{v}\} dG(c|s)$$

The demand is well defined when $G(\{c : \hat{x}(c) = \bar{v}\}) = 0$. Otherwise, the demand might not be well-defined (it can be infinite).

The first-order condition (FOC) is:

$$\frac{\partial D(p,p^*)}{\partial p} = \begin{cases} -\int_C \frac{f\left(\hat{x}(c)+p-p^*\right)}{1-F(\hat{x}(c))} \mathbb{1}\{\hat{x}(c) \geq p^*\} \mathbb{1}\{\hat{x}(c)+p-p^* \leq \bar{v}\} \, \mathrm{d}G(c|s), & p < p^*, \\ -\int_C \frac{f\left(\hat{x}(c)+p-p^*\right)}{1-F(\hat{x}(c))} \mathbb{1}\{\hat{x}(c) \geq p^*\} \mathbb{1}\{\hat{x}(c)+p-p^* \leq \bar{v}\} \, \mathrm{d}G(c|s) & p > p^*. \\ -\mathbb{1}\{\hat{x}(c) \geq p^*\} G(\{c : \hat{x}(c) = \bar{v}-p+p^*\} \mid s) \frac{1-F\left(\bar{v}^-\right)}{1-F(\bar{v}-p+p^*)} \end{cases}$$

When F is continuous, the second part drops because $1-F(\bar{v})=0$ for all $p > p^*$. In a symmetric equilibrium $p = p^*$, and it holds that $G(\{c : \hat{x}(c) = \bar{v}-p+p^*\} \mid s) = G(\{0\}|s) = 0$, i.e. there is no atom at $c=0$, following Assumption 2. Thus, in either case, the limits of the right and left derivative coincide for $p = p^*$.

Thus, the demand is continuous and differentiable everywhere. It is twice differentiable almost everywhere (except possibly at $p = p^*$).

Then for $p < p^*$,

$$\frac{\partial^2 D}{\partial p^2}(p,p^*) = -\int_C \frac{f'\left(\hat{x}(c)+p-p^*\right)}{1-F(\hat{x}(c))} \mathbb{1}\{\hat{x}(c) \geq p^*\} \, \mathrm{d}G(c \mid s). \tag{OA.50}$$

For $p > p^*$,

$$\begin{aligned} \frac{\partial^2 D}{\partial p^2}(p,p^*) = &-\int_C \frac{f'\left(\hat{x}(c)+p-p^*\right)}{1-F(\hat{x}(c))} \mathbb{1}\{\hat{x}(c) \geq p^*\} \mathbb{1}\{\hat{x}(c) \leq \bar{v}-p+p^*\} \, \mathrm{d}G(c \mid s) \\ &+ \frac{f(\bar{v})}{1-F\left(\bar{v}-p+p^*\right)} \frac{\partial}{\partial p} G\left(\{c : \hat{x}(c) \leq \bar{v}-p+p^*\} \mid s\right) \end{aligned} \tag{OA.51}$$

OA.2 Proofs of Section 6

OA.2.1 Correlation

Proof of Lemma 6. A consumer with search cost c searches if the expected surplus from searching exceeds the cost. The reservation value $\hat{x}(c)$ solves

$$\int_{\hat{x}(c)}^{\bar{v}(c)} (v-\hat{x}(c)) \, dF_c(v) = c. \tag{OA.52}$$

For $F_c = \mathrm{Uni}[0,\bar{v}(c)]$ with density $f_c(v) = 1/\bar{v}(c)$, the integral equals $(\bar{v}(c)-\hat{x}(c))^2/(2\bar{v}(c)) = c$, so $\hat{x}(c) = \bar{v}(c) - \sqrt{2c\,\bar{v}(c)}$.

Using $p^*_{PD}(c) = \frac{1-F_c(\hat{x}(c))}{f_c(\hat{x}(c))} = \bar{v}(c) - \hat{x}(c)$ gives $p^*_{PD}(c) = \sqrt{2c\,\bar{v}(c)}$. Consumer surplus is $CS^{PD}(c) = \hat{x}(c) - p^*_{PD}(c) = \bar{v}(c) - 2\sqrt{2c\,\bar{v}(c)}$. Participation under PD requires $CS^{PD}(c) \geq 0$, i.e., $\bar{v}(c) \geq 8c$; the participation cutoff c_0^{PD} solves $\bar{v}(c_0^{PD}) = 8c_0^{PD}$. In the linear case, $c_0^{PD} = \frac{a}{b+8}$; for $b > 0$, $c_0^{PD} < \frac{a}{b}$, so the cutoff lies in the domain where $\bar{v}(c) > 0$.

For monotonicity (on $\{c : \bar{v}(c) > 0\}$), differentiate:

$$\frac{dp^*_{PD}(c)}{dc} = \frac{\bar{v}(c) + c\,\bar{v}'(c)}{\sqrt{2c\,\bar{v}(c)}}. \tag{OA.53}$$

With $\bar{v}(c) = a - bc$ and $\bar{v}'(c) = -b$, the sign is $\operatorname{sgn}(a - 2bc)$. Hence $p^*_{PD}(c)$ is increasing for $c < c_{tp} := \frac{a}{2b}$ and decreasing for $c > c_{tp}$ when $b > 0$; if $b < 0$, then $p^*_{PD}(c)$ is increasing on its domain. □

Proof of Lemma 7. The harmonic-mean relation (11) follows from the aggregation result in the main model and applies verbatim to the correlated setting. □

Proof of Corollary 3. If $-8 \le b \le 8$ and $p^*_{PD}(c)$ is increasing on $[\underline{c}, c_0^{PD}]$, the marginal PD participant pays $\max_{c \in [\underline{c}, c_0^{PD}]} p^*_{PD}(c) = p^*_{PD}(c_0^{PD})$. By equation 11, p^*_U is a harmonic mean over active types and is strictly below this maximum whenever the active set is non-degenerate; the marginal type strictly gains surplus, and by continuity, a neighborhood of types participates. Hence, participation weakly expands.

If $b < 0$, then $p^*_{PD}(c)$ is increasing on the relevant domain and p^*_U strictly exceeds $\min p^*_{PD}$ on any non-degenerate active set, raising the price for low-c types. For $b \le -8$ individual consumer surplus, $\hat{x}(c) - p^*_{PD}(c)$ under full information, is decreasing in c. For sufficiently small a, low-c consumer surplus can turn negative under pooling, causing exit at the bottom.

If $b > 8$, then $c_{tp} < c_0^{PD}$ and there exists a region of active types with $c > c_{tp}$ where $p^*_{PD}(c)$ is decreasing. Replacing personalized prices by the pooled price p^*_U can therefore raise or lower the price faced by the marginal participant c_0^{PD}. In particular, the sign of $p^*_U - p^*_{PD}(c_0^{PD})$ depends on whether $\mathbb{E}\Big[\frac{1}{p^*_{PD}(C)} \mid C \le c_0^{PD}\Big]$ lies above or below $\frac{1}{p^*_{PD}(c_0^{PD})}$, which in turn is determined by the distribution of C on $[\underline{c}, c_0^{PD}]$. If $p^*_U < p^*_{PD}(c_0^{PD})$, participation expands at the high-c margin; if $p^*_U > p^*_{PD}(c_0^{PD})$, the marginal type exits and participation reduces. □

OA.2.2 Finite Firms

Proof of Proposition 9. Write

$$w(p) := \frac{p}{1 - F(p)^n}, \qquad \phi := w^{-1}, \qquad \psi(h) := \phi\Big(\frac{1}{h}\Big). \tag{OA.54}$$

Under full coverage and symmetry, the first-order conditions for a market with finitely many firms imply (Wolinsky, 1986)

$$w\big(p^*_{PD}(c)\big) = \frac{1}{h(c)} \qquad \text{and} \qquad w\big(p^*_U\big) = \frac{1}{\mathbb{E}_c[h(c)]}. \tag{OA.55}$$

Equivalently,

$$p^*_{PD}(c) = \psi\big(h(c)\big), \qquad p^*_U = \psi\big(\mathbb{E}_c[h(c)]\big). \tag{OA.56}$$

Average price and trade: By convexity of ψ (Claim 4 below) and Jensen's inequality,

$$\mathbb{E}_c\big[p^*_{PD}(c)\big] = \mathbb{E}_c\big[\psi(h(c))\big] \;\ge\; \psi\big(\mathbb{E}_c[h(c)]\big) = p^*_U, \tag{OA.57}$$

so the average price is (weakly) higher under discrimination.

With $F(p) = p$ on $[0, 1]$, the probability of trade at price p is $1 - p^n$. Under discrimination, it equals $1 - \mathbb{E}_c\big[(p^*_{PD}(c))^n\big]$, while under the uniform price it is $1 - (p^*_U)^n$. Since

$x \mapsto x^n$ is convex and $\mathbb{E}[p^*_{PD}(c)] \geq p^*_U$, $\mathbb{E}\big[(p^*_{PD}(c))^n\big] \;\geq\; \big(\mathbb{E}[p^*_{PD}(c)]\big)^n \;\geq\; (p^*_U)^n$, hence

$$1 - \mathbb{E}\big[(p^*_{PD}(c))^n\big] \;\leq\; 1 - \big(\mathbb{E}[p^*_{PD}(c)]\big)^n \;\leq\; 1 - (p^*_U)^n. \tag{OA.58}$$

Thus, trade is (weakly) higher under the uniform price, with strict inequality if the distribution of $p^*_{PD}(c)$ is non-degenerate.

Average traded prices: The average traded price at a given price p under discrimination and no information is

$$\bar{p}^{\text{trade}}_{PD} = \frac{\mathbb{E}_c\big[p^*_{PD}(c)\big(1 - F(p^*_{PD}(c))^n\big)\big]}{\mathbb{E}_c\big[1 - F(p^*_{PD}(c))^n\big]}, \qquad \bar{p}^{\text{trade}}_{U} = p^*_U. \tag{OA.59}$$

As shown above, the denominator $\mathbb{E}_c[1-F(p^*_{PD}(c))^n]$ is (weakly) smaller than $1-F(p^*_U)^n$, so the denominator is larger under the uniform price. A sufficient condition for the average traded price to be lower under no information is therefore that the numerator is smaller under no information.

Specialize to $F(p) = p$ on $[0,1]$, and define

$$r(p) := p(1-p^n), \qquad R(h) := r\big(\phi(1/h)\big). \tag{OA.60}$$

Applying Claim 5 (that $R(h)$ is convex whenever $n \geq 5$ or $\bar{c}$ is sufficiently small):

$$\mathbb{E}\big[r(p^*_{PD}(c))\big] = \mathbb{E}\big[R(h(c))\big] \;\geq\; R\big(\mathbb{E}[h(c)]\big) = r(p^*_U) = p^*_U\big(1-(p^*_U)^n\big), \tag{OA.61}$$

so the numerator of the average traded price is (weakly) larger under discrimination than under the uniform price. Combined with the denominator comparison, this implies that the average traded price is (weakly) lower under no information than under discrimination, with strict inequality whenever there is dispersion in the search costs.

Consumer surplus: Consumer surplus in a Wolinsky (1986) model for a type c facing price p can be written as

$$CS(c;p) = \hat{x}(c)\big(1 - F(\hat{x}(c))^n\big) + \int_p^{\hat{x}(c)} v\, dF_{1:n}(v) - p\big(1 - F(p)^n\big), \tag{OA.62}$$

and aggregate consumer surplus is its expectation over c.

Comparing the uniform price p^*_U with discriminatory prices $p^*_{PD}(c)$,

$$\begin{aligned} CS_U - CS_{PD} &= \mathbb{E}\big[CS(c;p^*_U) - CS(c;p^*_{PD}(c))\big] && \text{(OA.63)}\\ &= \mathbb{E}\big[p^*_{PD}(c)\big(1 - F(p^*_{PD}(c))^n\big)\big] - p^*_U\big(1 - F(p^*_U)^n\big) && \text{(OA.64)}\\ &\quad + \mathbb{E}\Big[\int_{p^*_U}^{p^*_{PD}(c)} v\, dF_{1:n}(v)\Big]. && \text{(OA.65)} \end{aligned}$$

For $F(v) = v$ on $[0,1]$, $F_{1:n}(v) = v^n$ and $dF_{1:n}(v) = nv^{n-1}dv$, so

$$\int_{p^*_U}^{p^*_{PD}(c)} v\, dF_{1:n}(v) = \frac{n}{n+1}\big((p^*_{PD}(c))^{n+1} - (p^*_U)^{n+1}\big), \tag{OA.66}$$

and therefore

$$CS_U - CS_{PD} = \Big(\mathbb{E}[p^*_{PD}(c)(1-(p^*_{PD}(c))^n)] - p^*_U(1-(p^*_U)^n)\Big) \tag{OA.67}$$

$$+ \frac{n}{n+1}\Big(\mathbb{E}[(p^*_{PD}(c))^{n+1}] - (p^*_U)^{n+1}\Big). \tag{OA.68}$$

The first bracket is nonnegative (and, under the convexity condition on $R(h)$, strictly positive), by the argument above for $\mathbb{E}[r(p^*_{PD}(c))] \geq r(p^*_U)$. The second bracket is nonnegative because $x \mapsto x^{n+1}$ is convex and $\mathbb{E}[p^*_{PD}(c)] \geq p^*_U$, so

$$\mathbb{E}[(p^*_{PD}(c))^{n+1}] \;\geq\; \big(\mathbb{E}[p^*_{PD}(c)]\big)^{n+1} \;\geq\; (p^*_U)^{n+1}. \tag{OA.69}$$

Thus $CS_U \geq CS_{PD}$, with strict inequality whenever there is dispersion in the search cost c which induces dispersion in p^*_{PD}.

Proof of Claims:

Claim 4

ψ is convex for the relevant price range $p \in [0, p^*_{PD}(\overline{c})]$.

Proof. Differentiate using inverse-function identities:

$$\psi'(h) = \phi'(t)\Big(-\frac{1}{h^2}\Big), \qquad \psi''(h) = \underbrace{\phi''(t)}_{-\frac{w''(p)}{[w'(p)]^3}}\frac{1}{h^4} + \underbrace{\phi'(t)}_{\frac{1}{w'(p)}}\frac{2}{h^3}. \tag{OA.70}$$

Thus, the sign of ψ'' is the sign of

$$S(p) \;:=\; \frac{2\,w'(p)^2}{w(p)} - w''(p)\,. \tag{OA.71}$$

It can be shown that $S(p)$ has the same sign as

$$2 - \big[(n-1)(n-2)\big]\cdot\big(F(p)\big)^n. \tag{OA.72}$$

In particular, ψ is convex at h (equivalently at the corresponding p) whenever

$$\big(F(p)\big)^n \;\leq\; \frac{2}{(n-1)(n-2)}. \tag{OA.73}$$

Given that equilibrium prices satisfy $p \leq p^m = \frac{1}{2}$, this condition holds, so ψ is convex over the relevant range. □

Claim 5

$R(h)$ is convex whenever $n \geq 5$ or $\overline{c}$ is sufficiently small.

Proof. For $R(h) = r(\phi(1/h))$ we get

$$R'(h) = r'(p)\,\phi'(t)\,\frac{d}{dh}\left(\frac{1}{h}\right) = -\frac{r'(p)}{h^2 w'(p)}. \tag{OA.74}$$

Using $\dfrac{dp}{dh} = \phi'(t)\,\dfrac{dt}{dh} = \dfrac{1}{w'(p)}\left(-\dfrac{1}{h^2}\right)$,

$$R''(h) = \frac{2}{h^3}\frac{r'(p)}{w'(p)} + \frac{1}{h^4}\cdot\frac{r''(p)w'(p) - r'(p)w''(p)}{[w'(p)]^3} \tag{OA.75}$$

$$= \frac{1}{h^4[w'(p)]^3}\big\{\,2h\,r'(p)[w'(p)]^2 + r''(p)w'(p) - r'(p)w''(p)\big\}\,. \tag{OA.76}$$

As $h = 1/w(p)$ and $w'(p) > 0$ for $p \in (0,1)$, the sign of $R''(h)$ is the sign of

$$S(p) := r''(p)\,w'(p) \;+\; r'(p)\left(\frac{2[w'(p)]^2}{w(p)} - w''(p)\right). \tag{OA.77}$$

Equivalently (divide by $w'(p) > 0$),

$$R''(h) \;\gtrless\; 0 \quad\Longleftrightarrow\quad r''(p) \;+\; r'(p)\left(\frac{2w'(p)}{w(p)} - \frac{w''(p)}{w'(p)}\right) \;\gtrless\; 0. \tag{$\star$}$$

For $F(p) = p$,

$$w(p) = \frac{p}{1-p^n}, \quad w'(p) = \frac{1+(n-1)p^n}{(1-p^n)^2}, \quad w''(p) = \frac{n\,p^{n-1}\big((n+1)+(n-1)p^n\big)}{(1-p^n)^3}, \tag{OA.78}$$

and

$$r(p) = p(1-p^n), \quad r'(p) = 1-(n+1)p^n, \quad r''(p) = -n(n+1)p^{n-1}. \tag{OA.79}$$

Substituting these expressions into ($\star$), clearing the positive denominator, and simplifying yields

$$R''(h) \gtrless 0 \iff \Theta_n(p) \gtrless 0 \quad \text{with} \quad \Theta_n(p) := 1-(n^2+2)\,p^n + (1-n^2)\,p^{2n}. \tag{OA.80}$$

One checks that $\Theta_n'(p) < 0$ on $(0,1)$, so Θ_n is strictly decreasing. At $p = \frac{1}{2}$,

$$\Theta_n\Big(\tfrac{1}{2}\Big) > 0 \quad \text{for } n \geq 5. \tag{OA.81}$$

Because $p^*_{PD}(c) \leq p^m = \frac{1}{2}$ on the relevant range, R is convex whenever $n \geq 5$; if $\overline{c}$ is sufficiently small, convexity holds on $[0, p^*_{PD}(\overline{c})]$ as well. □

□

OA.3 Voting

Privacy policy is ultimately a political choice. With voluntary disclosure, unraveling can push the market to full revelation; in the benchmark, complete information lowers aggregate consumer surplus and, in partially covered markets, reduces welfare. This suggests that stringent, mandatory privacy (or limits on data use) may raise consumer surplus and welfare and may prevent the negative outcomes associated with voluntary disclosure and perfectly informative signals. A natural question is whether such a policy is implementable by majority vote.

To study the voting setting, adjust the timing of the model as follows. First, consumers observe their search costs c. Then, consumers will vote on whether to have privacy (uninformative signals) or no privacy (perfectly revealing signals). The voting rule is simple majority, and each consumer has one vote. Afterwards, the outcome is implemented, consumers decide to participate in the market, and firms set prices according to the information structure chosen and the signal received.

Moreover, I assume infinitesimally small voting costs or that consumers who have no stake in the outcome randomize their votes. This is equivalent to assuming that $c_0^U \leq \overline{c}$. Only those consumers who are active under at least one of the two regimes have an influence on the voting outcome.

The median voter theorem applies.

Lemma 8
There exists a cutoff c^* such that voters with $c \leq c^*$ vote against privacy (for no privacy), while voters with $c > c^*$ vote for privacy.

Thus, low-c consumers enjoy low personalized prices under no privacy and oppose privacy, while high-c consumers prefer the pooled price. Let c_m denote the median voter and $\hat{x}(c_m)$ the associated reservation value.

In principle, conditions on primitives for situations when the median voter (and the election) implements one or the other outcome are not very clear. However, one can find sufficient conditions for privacy or no privacy to be implemented.

Proposition 10
Let $h(c) := \frac{f(\hat{x}(c))}{1-F(\hat{x}(c))}$ and assume h is decreasing in c. The median voter prefers privacy iff $\mathbb{E}[h(c)] \geq h(c_m)$.

*A sufficient condition for **privacy** is:*

1. *h is (strictly) convex, and*
2. *$\mathbb{E}[c] \leq c_m$.*

*A sufficient condition for **no privacy** is:*

1. *h is (strictly) concave, and*
2. *$c_m \leq \mathbb{E}[c]$.*

For the majority to prefer privacy, the price under privacy must be lower than under disclosure for the median voter. When the distribution of price-sensitivities h is convex, then the average price-sensitivity in the population will be larger than the price-sensitivity for the consumer with the average search costs. If the median voter has even larger search costs and is thus less price-sensitive, the median voter will prefer privacy even more.

More generally, when the inverse hazard rate of F is convex in $\hat{x}$, then the pooled price will be higher than the disclosure price for the average $\hat{x}$. The median voter with c_m must therefore have very low costs relative to the population to ensure that $\hat{x}(c_m) \leq \mathbb{E}[\hat{x}]$. This is because $\hat{x}$ is convex, and it does not suffice if $c_m \leq \mathbb{E}[c]$.

This suggests that the conditions tend to be more conducive for the median voter to prefer privacy than price discrimination.

The previous proposition does not cover all possible distributions. For example, when h is convex, but c is very right-tailed, meaning that the median is very far left in the distribution, the median voter might prefer no privacy. Similarly, when h is concave, but the distribution is extremely left-tailed, the median voter might prefer privacy.

Comparative statics regarding the distribution G are briefly discussed in the appendix, and it just shows that the effects depend on the details.

The voting analysis assumes consumers know their own search cost when they vote and focuses on a single market. This is plausible if that market is representative and search costs are stable across markets (e.g., reflecting the opportunity cost of time), so a single vote sets the privacy policy for all market sets, or if votes are held market by market after costs are realized. If voters are ex-ante uninformed, each votes to maximize expected consumer surplus; with identical primitives across markets, this selects the consumer-surplus-maximizing policy. However, it need not coincide with welfare maximization, because gains from privacy are unevenly distributed across consumers and between consumers and firms. Moreover, when voters are aware of their own costs before voting, majority rule need not implement the aggregate consumer-surplus-maximizing policy: heterogeneity can cause the median voter's preferred policy to differ from the utilitarian optimum.

OA.3.1 Proofs of Section OA.3

Proof of Lemma 8. $p^*_{PD}(c) = \frac{1-F(\hat{x}(c))}{f(\hat{x}(c))}$ is increasing in c. Thus, consumers with a higher c benefit more from a common, fixed price p^*_U. Thus, there exists a threshold c_t such that all consumers with $c \geq c_t$ vote in favor of privacy (which yields p^*_U), and consumers with $c < c_t$ vote against privacy. □

Proof of Proposition 10. Under no privacy (perfect revelation), type c pays

$$p_{PD}(c) = \frac{1 - F(\hat{x}(c))}{f(\hat{x}(c))} = \frac{1}{h(c)}.$$

Under privacy (pooling), the price is

$$p_U = \frac{1}{\mathbb{E}[h(c)]}.$$

Thus the median voter prefers privacy iff

$$\hat{x}(c_m) - p_U \geq \hat{x}(c_m) - p_{PD}(c_m) \quad \Longleftrightarrow \quad p_U \leq p_{PD}(c_m) \quad \Longleftrightarrow \quad \mathbb{E}[h(c)] \geq h(c_m).$$

If h is convex, Jensen implies $\mathbb{E}[h(c)] \geq h(\mathbb{E}[c])$. If moreover $\mathbb{E}[c] \leq c_m$ and h is decreasing, then $h(\mathbb{E}[c]) \geq h(c_m)$. Hence $\mathbb{E}[h(c)] \geq h(c_m)$, so the median voter votes for privacy.

If h is concave, Jensen implies $\mathbb{E}[h(c)] \leq h(\mathbb{E}[c])$. If moreover $c_m \leq \mathbb{E}[c]$ and h is decreasing, then $h(\mathbb{E}[c]) \leq h(c_m)$. Hence $\mathbb{E}[h(c)] \leq h(c_m)$, so the median voter votes against privacy. □

OA.4 Microfoundation (Sketch)

In the main body of the work, firms have received some signal over search cost from consumers. One possible signal could be the number of searches or the percentile of searches a consumer ranked in other markets, which indicates whether a consumer tends to be an avid searcher with low search costs or a consumer with high search costs and

a low search intensity. This appendix sketches how, with enough panel data on a single consumer across many ex-ante symmetric markets, one can nonparametrically recover that consumer's search cost. The data requirements can certainly be relaxed if one employs data on many consumers, not just one.

Consider a sequence of ex-ante symmetric markets indexed by j, in each of which consumer i may conduct sequential search. For market j the total number of searches $N_{i,j} \in \{1, 2, \ldots\}$, as well as two observed shocks $(\gamma_{i,j}, \varepsilon_{i,j})$. γ shifts the per-search cost while ε shifts the price additively. The γ shock can be seen as an impediment to search, such as a slow internet connection or a new website layout. ε can be seen as special promotions or changes in shipping costs due to different locations of the buyer/seller. Let c_i denote the consumer's (time-invariant) search cost. Across markets j, we observe consumer i's total searches $N_{i,j} \in \{1, 2, \ldots\}$ and two shocks $(\gamma_{i,j}, \varepsilon_{i,j})$. The shock γ shifts the per-search cost; ε shifts the acceptance threshold additively (e.g., promotions or shipping/location surcharges). Let c_i be the consumer's time-invariant per-search cost, $\hat{x}(\cdot)$ the reservation value induced by the cost, and F the CDF of match values with density f. [16]

The consumer accepts in market j iff $v_{i,j} \geq \hat{x}(c_i + \gamma_{i,j}) + \varepsilon_{i,j}$. Thus, the per-draw stopping probability given the shocks, $Q(\gamma, \varepsilon)$, is

$$Q(\gamma, \varepsilon) \;=\; 1 - F\big(\hat{x}(c_i + \gamma) + \varepsilon\big),$$

and (with independent draws), the expected number of searches, $N_{i,j}$, is geometric with

$$\mathbb{E}[N_{i,j} \mid \gamma, \varepsilon] \;=\; \frac{1}{Q(\gamma, \varepsilon)} \;=:\; m(\gamma, \varepsilon).$$

Take partial derivatives holding the other shock fixed:

$$\partial_\varepsilon Q(\gamma, \varepsilon) \;=\; -f\big(\hat{x}(c_i + \gamma) + \varepsilon\big), \qquad \partial_\gamma Q(\gamma, \varepsilon) \;=\; -f\big(\hat{x}(c_i + \gamma) + \varepsilon\big) \cdot \hat{x}'(c_i + \gamma).$$

Hence, on the observed support,

$$\hat{x}'(c_i + \gamma) \;=\; \frac{\partial_\gamma Q(\gamma, \varepsilon)}{\partial_\varepsilon Q(\gamma, \varepsilon)}.$$

Integrating,

$$\hat{x}(z) = C_x + \int_{z_0}^{z} \hat{x}'(y)\, dy \quad \text{and} \quad F(x) = C_F + \int_{x_0}^{x} f(t)\, dt,$$

for arbitrary z_0 and x_0, with integration constants C_x, C_F.

The constants can be pinned down. Choose any reference pair $(\gamma_0, \varepsilon_0)$ with $0 < Q(\gamma_0, \varepsilon_0) < 1$. Define

$$x_0 := \hat{x}(c_i + \gamma_0) + \varepsilon_0, \qquad z_0 := c_i + \gamma_0.$$

[16] This is just a sketch, for more technical assumptions needed, such as that the ε and γ need to have differentiable densities, see Matzkin (2007)

Then from the definition of Q,

$$C_F = F(x_0) = 1 - Q(\gamma_0, \varepsilon_0).$$

For $\hat{x}$,

$$\hat{x}(z_0) = C_x \quad \text{and} \quad Q(\gamma_0, \varepsilon_0) = 1 - F\big(\hat{x}(z_0) + \varepsilon_0\big).$$

Invert F to solve for the level:

$$C_x = \hat{x}(z_0) = F^{-1}\big(1 - Q(\gamma_0, \varepsilon_0)\big) - \varepsilon_0 \ .$$

With C_F and C_x, $\hat{x}$ and F are pinned (on the empirically reached range) by the integrals above.

Recovering search cost. For any shocks (γ, ε),

$$c_i \ = \ \hat{x}^{-1}\Big(F^{-1}\big(1 - 1/m(\gamma, \varepsilon)\big) - \varepsilon\Big) \ - \ \gamma.$$

In particular, at $(\gamma, \varepsilon) = (0, 0)$ with $k := \mathbb{E}[N \mid 0, 0]$,

$$c_i \ = \ \hat{x}^{-1}\Big(F^{-1}\big(1 - 1/k\big)\Big).$$

To summarize, two independent and exogenous shifts, ε and γ, nonparametrically trace out F and $\hat{x}$. This is a standard shift-design argument common in the nonparametric literature. The argument requires only a single consumer tracked across many markets, and mild regularity (continuity of F and $\hat{x}$, independence of shocks). Rich shocks are used, and the identified range of $\hat{x}, F$ is limited when shocks are less rich. I am confident that with more consumers, all subject to the same $\hat{x}, F$, it is also possible to identify the entire distribution with fewer rich shocks.

OA.5 Comparative Statics with respect to search costs

In the following, changes in the distribution of search costs, G, will be examined. It will be studied how privacy's value and voting incentives depend on G.

OA.5.1 Value of Privacy

A natural starting point is to ask how equilibrium outcomes under privacy change when G changes (relative to perfect observability).

Proposition 11
Let $C \sim G$ and let C' be a mean-preserving spread (MPS) of C, i.e. $C' = C + \varepsilon$ with $\mathbb{E}[\varepsilon \mid C] = 0$. Assume the market is fully covered under privacy for both G and the post-spread distribution of C'. Then:

- *The uninformative price under privacy weakly decreases (increases), and profits weakly decrease (increase), when h is convex (concave).*
- *Welfare weakly increases.*
- *Consumer surplus weakly increases when h is convex.*

Moreover, if the spread is non-degenerate, welfare increases strictly. If in addition h is strictly convex (concave) on the relevant support, then the price and profit comparisons are strict; consequently, consumer surplus increases strictly when h is strictly convex.

The function h measures price sensitivity, depending on both the reservation value of search and the hazard rate of F, the distribution of match values. If h is convex in search costs, an MPS increases average price sensitivity, benefiting consumers (since higher price sensitivity reduces prices). Conversely, if h is concave, an MPS reduces price sensitivity, increasing prices to the detriment of consumers.

OA.5.2 Comparative Statics for Voting

In this subsection, the privacy policy is determined by a simple majority vote as in subsection OA.3. It was shown that the median voter determines the election outcome. A standard mean-preserving spread (MPS) of search costs can change the median type. To isolate the effect of a more dispersed G while holding the decisive type fixed, restrict attention to spreads that preserve both the mean and the median of G.

Let $C \sim G$ and $C' \sim G'$. G' is a median- and mean-preserving spread of G if $\text{Median}(C') = \text{Median}(C)$ and $C' = C + \varepsilon$ with $\mathbb{E}[\varepsilon \mid C] = 0$.

Following Proposition 11,

Corollary 4
Suppose G' is a median- and mean-preserving spread of G, and the market is fully covered under privacy for both G and G'. If h is convex (concave), then the median voter is weakly more (less) inclined to support privacy under G' than under G. The comparison is strict if the spread is non-degenerate and h is strictly convex (concave) on the relevant support.

The proof is straightforward. When h is convex, the uninformative price weakly decreases by Proposition 11. The median voter remains the same under G and G' by construction, and the discriminatory price faced by the median voter is unchanged. Thus, privacy (relative to perfect observability) becomes more attractive for the decisive voter.

If the market is not fully covered when voting occurs, then participation is endogenous and beyond the scope of these comparative statics.

OA.5.3 Proofs of Section OA.5

Proof of Proposition 11. **Prices:** In a fully covered market, prices determine how welfare is split between consumers and firms. Under a mean-preserving spread (MPS),

$$p_U^{MPS} = \frac{1}{\mathbb{E}[h(C')]} = \frac{1}{\mathbb{E}[h(C+\varepsilon)]} = \frac{1}{\mathbb{E}\big[\mathbb{E}[h(C+\varepsilon) \mid C]\big]}.$$

If h is convex, then by Jensen's inequality and $\mathbb{E}[\varepsilon \mid C] = 0$,

$$\mathbb{E}[h(C+\varepsilon) \mid C] \;\geq\; h(\mathbb{E}[C+\varepsilon \mid C]) = h(C),$$

hence $\mathbb{E}[h(C+\varepsilon)] \geq \mathbb{E}[h(C)]$ and therefore

$$p_U^{MPS} = \frac{1}{\mathbb{E}[h(C+\varepsilon)]} \;\leq\; \frac{1}{\mathbb{E}[h(C)]} = p_U^*.$$

For concave h, the inequalities reverse.

Welfare: With full coverage, welfare under privacy is $\mathbb{E}[\hat{x}(C)]$ (and analogously $\mathbb{E}[\hat{x}(C')]$ after the spread). Since $\hat{x}$ is (strictly) convex, an MPS weakly increases welfare; the increase is strict under a non-degenerate spread.

Profits and consumer surplus: In a fully covered privacy equilibrium, the uninformative price p_U is constant across consumers. With unit mass of consumers, industry profit equals p_U; hence profits move one-for-one with the uninformative price. Consumer surplus equals welfare minus industry profits. When h is convex, welfare weakly rises and profits weakly fall, so consumer surplus weakly increases. When h is concave, welfare rises but profits rise as well, so the net effect on consumer surplus is ambiguous. □

OA.6 Numerical Simulation

I conduct a numerical simulation assuming that F and G follow beta distributions. Since F must be log-concave, requiring $F_\alpha, F_\beta > 1$. The parameter values used are:

$$(F_\alpha, F_\beta) \in \{1.01, 1.1, 1.5, 2, 2.5, 3, 4, 5, 7, 10\}^2$$

For G, there are no restrictions beyond differentiability, which holds for the beta distribution. The corresponding simulation parameters are:

$$(G_\alpha, G_\beta) \in \{0.5, 0.9, 1, 1.1, 1.5, 2, 2.5, 3, 5, 10\}^2$$

Computing the consumer surplus under privacy (CS_P) and price discrimination (CS_D) for the Cartesian product of these parameter combinations results in $10^2 \times 10^2 = 10,000$ configurations.

Figure OA.2 shows that $CS_P > CS_D$ in almost all cases, except when F_α is large, F_β is small, and G_α is intermediate. This combination implies that the distribution of v_i is concentrated at a high value, making $\hat{x}(c)$ relatively insensitive to changes in c.

Additionally, when G_α is intermediate and G_β is large, mass concentrates near an intermediate value of c. If this concentration aligns closely with c_0^U, the entry effect on consumer surplus diminishes, while the price effect remains significant. When much of the probability mass is near c_0^U, the dispersion within $[0, c_0^D]$ is reduced.

Given that numerical simulations are subject to approximation errors, integrals are computed with a precision of 10^{-4}. To reduce classification errors, I only classify cases where $CS_P < CS_D - 0.0001$ as meaningful violations. However, since subsequent calculations depend on earlier approximations, cumulative numerical errors may still affect the results.

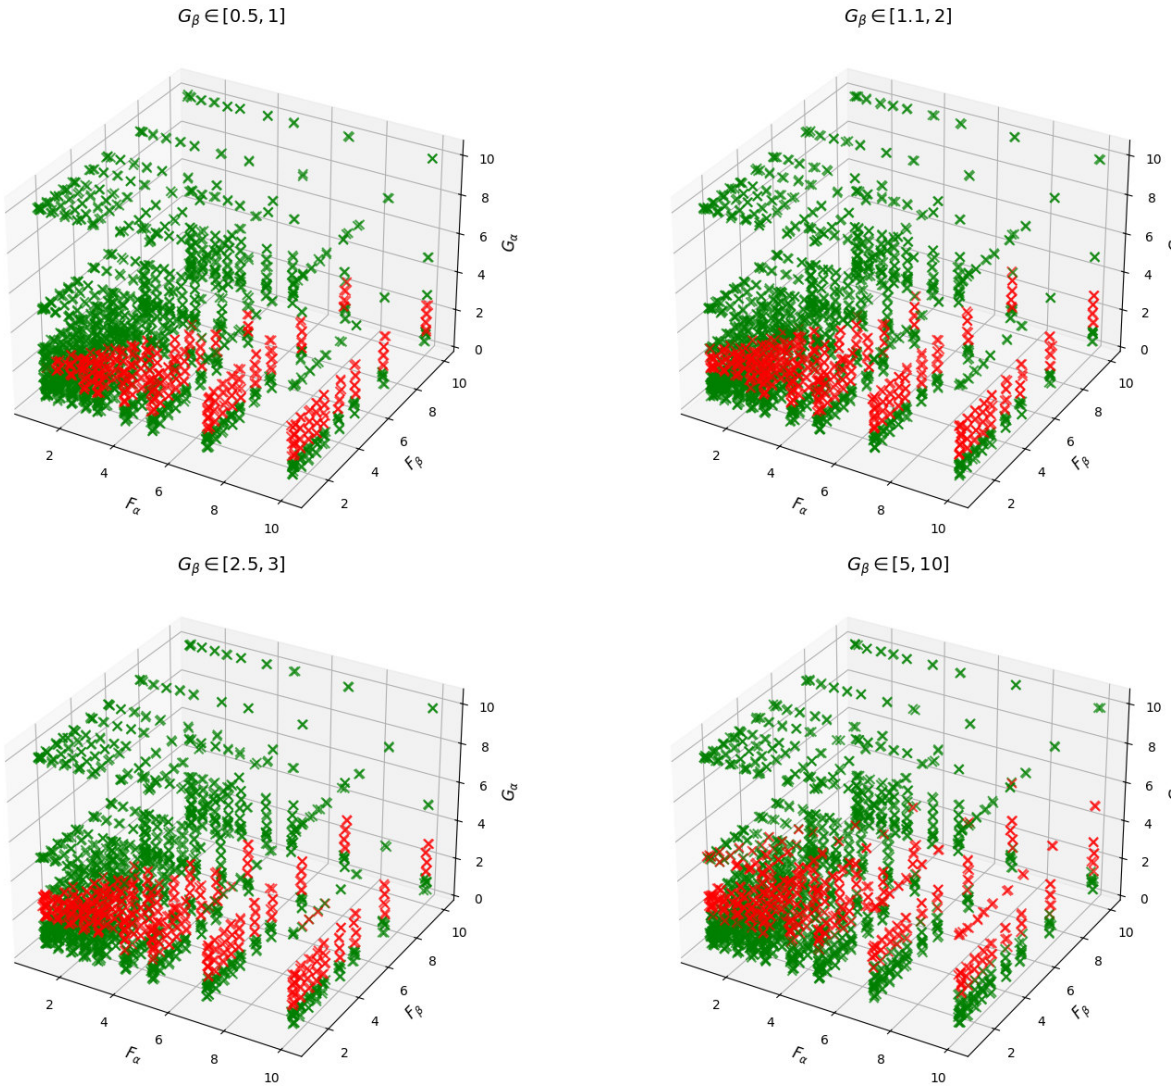


Figure OA.2: Figure OA.2: Consumer surplus under privacy vs. discrimination for G and F beta distributions

Note: Each point represents a parameter combination $(F_\alpha, F_\beta, G_\alpha, G_\beta)$. Points are colored red if $CS_U < CS_{PD} - 0.001$ and green otherwise. This threshold allows for numerical error from approximate integral calculations performed at precision 10^{-4}. Violations occur when F_α is large, F_β is small, and G_α is intermediate, reflecting reduced sensitivity in $\hat{x}(c)$ and concentrated mass in G.