



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

A Theory of the Effects of Privacy

Alessandro Bonatti, Yunhao Huang, J. Miguel Villas-Boas

To cite this article:

Alessandro Bonatti, Yunhao Huang, J. Miguel Villas-Boas (2026) A Theory of the Effects of Privacy. *Management Science* 72(5):3925–3947. <https://doi.org/10.1287/mnsc.2024.05354>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Management Science*. Copyright © 2025 The Author(s). <https://doi.org/10.1287/mnsc.2024.05354>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2025 The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

A Theory of the Effects of Privacy

Alessandro Bonatti,^a Yunhao Huang,^b J. Miguel Villas-Boas^{c,*}

^aMIT Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142; ^bMarshall School of Business, University of Southern California, Los Angeles, California 90089; ^cHaas School of Business, University of California, Berkeley, California 94720

*Corresponding author

Contact: bonatti@mit.edu,  <https://orcid.org/0000-0002-9150-2334> (AB); yunhao.huang@marshall.usc.edu,

 <https://orcid.org/0000-0003-2197-750X> (YH); villas@berkeley.edu,  <https://orcid.org/0000-0003-1299-4324> (JMV-B)

Received: May 7, 2025

Revised: June 30, 2025

Accepted: July 1, 2025

Published Online in Articles in Advance:
September 2, 2025

<https://doi.org/10.1287/mnsc.2024.05354>

Copyright: © 2025 The Author(s)

Abstract. The growth of information technologies has intensified concerns about individual privacy and highlighted the importance of consumer privacy rights. At the same time, these technologies enable significant personalization of communications and offerings. We develop a theory of the effects of privacy based on the concavity of an individual's derived payoff function with respect to market beliefs: When this function is concave, privacy is valuable. We identify market conditions that lead to concavity or convexity in the derived payoff function. The framework is applied to contexts such as product choice, price discrimination, data breaches, and health insurance.

History: Accepted by Raphael Thomadsen, marketing.



Open Access Statement: This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as "Management Science. Copyright © 2025 The Author(s). <https://doi.org/10.1287/mnsc.2024.05354>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>."

Funding: A. Bonatti acknowledges financial support through the National Science Foundation [Grant SES-1948692].

Supplemental Material: The data files are available at <https://doi.org/10.1287/mnsc.2024.05354>.

Keywords: privacy • information • price discrimination • targeted advertising

1. Introduction

The development of information technologies makes it possible to track citizens' searches, communications, movements, and purchases in unprecedented detail. The potential for tracking means a loss of privacy, and understanding when tracking harms or benefits individuals has become a central policy and managerial issue.

However, the existing literature shows contrasting effects of privacy. Within a fixed market environment, the welfare effects of privacy can vary depending on market structure and which dimensions of information are revealed. For example, in simple price discrimination settings, a loss of privacy generally harms consumer welfare when firms learn about consumer preferences (Robinson 1933, Schmalensee 1981), but a loss of privacy can improve consumer welfare when firms are in competition or learn about cost differences across market segments (Thisse and Vives 1988, Chen and Schwartz 2015). Across environments, privacy can have opposite effects despite similar information being revealed. For example, information about consumer preference tends to harm consumers in price discrimination settings but can help in targeted advertising settings through improved match quality between products and preferences (Athey and Gans 2010, de Cornière and de Nijs 2016).

The divergent results within and across market environments call for a unifying framework that accommodates general information structures and market actions. This paper's main idea is that a loss of privacy increases the precision of observers' posterior beliefs about an individual's preferences or behavior. In particular, a loss of privacy leads to a mean-preserving spread in the distribution of their beliefs. An individual's welfare response therefore hinges on the curvature of their derived payoff in the observers' beliefs: Concavity implies that a loss of privacy is harmful, and convexity that it is beneficial. This single mechanism offers a unified explanation for existing results and clarifies when new information helps or hurts consumers.

A familiar benchmark illustrates the logic. When the observer is the individual herself, the case of a single decision maker under uncertainty, more information is always valuable; the decision maker's payoff is convex in their beliefs. In terms of privacy losses, if the observer is perfectly aligned with the individual (i.e., they act in their best interest), a loss of privacy is always beneficial to the individual.

This paper derives conditions on the indirect utility function of an individual and on observers' responses to information about the individual's type of behavior

such that the individual's derived payoff function is concave or convex. These conditions depend on three effects that collectively explain when and why privacy is beneficial or harmful within and across different market environments. First, the observers of the information may react in a way that is beneficial or harmful to the individual. Second, as more information leads the observers to react more steeply to any signals, the indirect utility function of the individual with respect to their actions may be convex or concave, which has implications for the derived payoff function in terms of the posterior beliefs. Third, the actions of the observers of the information may react to their posterior beliefs in a concave or convex way.¹

Applying the framework clarifies several longstanding puzzles and, more broadly, offers a unifying theory of the effects of privacy, which has thus far proven elusive. Indeed, several existing results on privacy in diverse settings can be reinterpreted through this lens. For example, it explains some conditions for price discrimination to benefit or harm consumers. If price discrimination is based on consumer preference, the first effect typically dominates. That is, higher types suffer more from higher prices, and markets charge higher prices for segments with higher types. In contrast, if price discrimination is based on marginal costs, the third effect may dominate. When the optimal prices are concave in marginal costs,² more information leads to lower average prices, which benefits consumers. We further consider the application of this framework to other settings, including various forms of targeting effects, health insurance, hacking, and the case of privacy being a value in itself. We show how the insights apply to these settings by identifying the shape of the individual's derived payoff function and the dominating effect in each application.

This unified framework offers implications for both managers and policymakers. For managers, it provides a systematic approach to predict the environments when individuals will resist sharing information and when they might welcome personalization, guiding data-collection strategies. For policymakers, the framework helps differentiate environments where privacy regulations should be stringent (when information harms consumers) from environments where flexibility might be beneficial (when information benefits consumers).

As discussed in Acquisti et al. (2016), privacy has been defined in different ways, from the protection of an individual's personal space (Warren and Brandeis 1890), to the control over one's personal information (Westin 1967), to a dimension of dignity and academic freedom (Schoeman 1992). Related to the change in beliefs, Eilat et al. (2021) measure privacy as the Kullback–Leibler divergence between the prior distribution of beliefs over the types of an agent and the

posterior beliefs induced by the equilibrium actions. In terms of this paper, we concentrate only on the aspects of observation by third parties of certain characteristics or behaviors of individuals (the effect on posterior beliefs). If privacy is seen as control over that information, then this paper can be understood as studying whether individuals should give access to that information, at what price, and to whom. If privacy is regarded as a dimension of dignity and academic freedom, then it can be considered a value in itself, and the paper also considers that case as an application, interpreting this possibility in terms of the framework presented here.

Regarding privacy as a value in itself, Farrell (2012) notes that privacy can be seen as a final good (value in itself) or an intermediate good, where privacy allows an individual to receive better offers from other agents. Regarding the value in itself, Farrell (2012) notes the “icky” feeling of being watched and tracked. In fact, at the extreme, most individuals may not find it pleasing if a technology is developed such that the individual is on camera at all hours of the day. However, some literature suggests that this value of privacy per se may not be very large, which has been called “the privacy paradox.” That is, although most individuals state a preference for privacy, they require relatively small incentives to forgo some dimensions of their privacy (Gross and Acquisti 2005, Barnes 2006, Adjerid et al. 2013, Athey et al. 2017, Johnson et al. 2020).

There has been significant work on different dimensions of privacy with a focus on specific market applications. Hui and Png (2006) and Acquisti et al. (2016) provide surveys on the economics of privacy. This literature reveals contrasting findings across different environments: In contracting environments, a loss of privacy may induce inefficient switching (Fudenberg and Tirole 2000, Taylor 2003), increase rent extraction (Calzolari and Pavan 2006), and facilitate adverse selection (Hermalin and Katz 2006), but may also prevent inefficient market exclusion (Hermalin and Katz 2006). In markets with repeat purchases, a loss of privacy may benefit consumers by intensifying competition (Villas-Boas 1999, Zhang 2011) or inducing consumers to strategically delay purchases (Villas-Boas 2004); the effect of privacy can depend on consumer sophistication (Taylor 2004) and the cost of anonymization technologies (Acquisti and Varian 2005, Conitzer et al. 2012). With behavior-based, targeted advertising, a loss of privacy may benefit consumers through improved ad relevance and more informed inferences (Athey and Gans 2010, Goh et al. 2015, Rafieian and Yoganarasimhan 2021, Shin and Yu 2021, Villas-Boas and Yao 2021) but may lead to misuse of data (Evans 2009), negative externalities (Goh et al. 2015), or higher prices (de Cornière and de Nijs 2016, Oery 2016, Ning 2021); the effect may vary with consumer annoyance of receiving advertising (Shen and Villas-Boas 2018, Lin 2020) and consumer

flexibility in product preferences (Zhu et al. 2025). In personalized pricing settings, the effect of privacy may depend on market competition (Shaffer and Zhang 1995), fairness concerns (Li and Jain 2016), quality-adjusted costs (Rhee and Thomadsen 2017), consumer valuations (Amaldoss and He 2019), firms' commitment power (Ichihashi 2020), network effects (Hajihashemi et al. 2022), and approaches to balance personalization with privacy (Chen et al. 2022, Lei et al. 2024). These contrasting results highlight the need for a unifying theoretical framework.

There is also literature on how agents (for example, consumers) react to privacy regulations, incurring costs to preserve privacy, or ceasing to interact with the market (Calzolari and Pavan 2006, Hann et al. 2008, Kummer and Schulte 2019, Montes et al. 2019, Chen et al. 2020, Jullien et al. 2020, Argenziano and Bonatti 2023, Choi et al. 2023, Fainmesser et al. 2023, Miklós-Thal et al. 2024).³

Another strand of literature addresses how the presence of more or less information in a market affects the market interaction between different economic agents (Vives 1984, 1988; Gal-Or 1985; Shapiro 1986; Raith 1996; Casadesus-Masanell and Hervas-Drane 2015; Johnson et al. 2023). In relation to that work, the presentation here does not focus on the interaction between economic agents but instead concentrates on the shape of the derived payoff functions and the market actions as a function of the information in the market. In this vein, the market segmentation analyses in Bergemann et al. (2015), Yang (2022), and Elliott et al. (2022) assess the impact of information design in settings with market power. In particular, Farboodi et al. (2025) use the concavity of the aggregate consumer surplus to sign the welfare effects of information facilitating third-degree price discrimination. Finally, Rhodes and Zhou (2024) explore fully personalized versus fully anonymous pricing in a competitive context.

One potential issue regarding privacy is that an individual may want some agents to have the individual's information, but for that information not to go to other agents. The analysis of that more complex situation is not considered here, with the focus being solely on the situation of an individual facing a certain type of agent.

Another issue is that if individuals can credibly disclose information about their type (for example, disclosing verifiable information), these problems of privacy can unravel because the "best" types have an incentive to reveal their type to receive better treatment, and the market may be able to infer that the types that do not reveal their information are "less good" types (Posner 1978, Stigler 1980). This yields a situation where individuals do not have privacy. We consider settings where the individual's type cannot be credibly communicated or where there are laws against such communication.

The remainder of the paper is organized as follows. Section 2 discusses a motivating example and Section 3 presents a general model tying an individual's *indirect* utility function to that individual's preferences for privacy. Section 4 derives conditions on the *primitives* of the problem, such that the derived payoffs of the individuals satisfy the conditions for privacy to be preferred or not to be preferred. Section 5 presents applications of the framework to several settings, and Section 6 concludes the paper.

2. Motivating Example: Health Insurance

Most privacy regulation involves sensitive personal data. One prime example of such sensitive data is health data. Part of the regulators' concerns is that such data may be used to price discriminate in health insurance markets. This motivating example illustrates clearly the tradeoffs involved. In Section 5, we present several general market applications, including price discrimination, product and price targeting, and quality pricing and screening.

Consider the following setting, based on the insurance model of Hirshleifer (1971). An individual consumer may or may not have a certain gene. The presence of the gene causes health complications that cost L to address. The individual starts with wealth w and has a strictly concave direct utility function $u(\cdot)$ over total wealth. Perfectly competitive, risk-neutral insurers provide insurance to this market.

Each individual consumer has the gene with probability θ . The risk parameter $\theta \in [\underline{\theta}, \bar{\theta}]$ is distributed in the consumer population with mean μ_0 . The competitive insurance industry has access to individual-level data, which we model as observing an informative signal s of the type of each individual. Let $\mu(s) \equiv \mathbb{E}[\theta|s]$ denote the mean of the firms' posterior distribution induced by signal s . By the law of total probability, we also have $\mathbb{E}[\mu(s)] = \mu_0$, that is, the posterior averages out to the prior.

Assume that the parameters are such that all consumers participate in the competitive insurance market, so there are no selection effects. This means that the following condition holds:

$$\underline{\theta} u(w - L) + (1 - \underline{\theta})u(w) < u(w - \bar{\theta} L).$$

Every individual who generates signal s is fully reimbursed in case of a loss L and pays a competitive insurance premium of

$$p(s) = \mu(s)L.$$

This results in a fully insured wealth level $w - \mu(s)L$.

Therefore, the aggregate consumer surplus generated under the signals s is given by

$$W = \mathbb{E}_s[u(w - \mu(s)L)]. \quad (1)$$

Under the prior distribution only (i.e., without additional signals), we instead obtain

$$W_0 = u(w - \mu_0 L),$$

and because $\mathbb{E}_s[\mu(s)] = \mu_0$ and $u(\cdot)$ is strictly concave, we immediately conclude that $W < W_0$. Therefore, consumer welfare (i.e., the aggregate consumer surplus) with a competitive insurance industry is larger under a privacy regime. This is intuitive—a lack of privacy exposes the consumer to classification risk. With more information, insurance companies can more precisely categorize consumers based on their risk levels, leading to varying premiums across different consumer segments. Because consumer utility is strictly concave in wealth (and therefore in premiums), a fully insured consumer prefers a premium with certainty over an uncertain premium with the same average. This risk aversion makes this classification risk on average welfare-reducing for consumers.

In the rest of this paper, we derive a more general model that applies beyond the insurance setting. In the context of this example, the richer model will enable us to explore the robustness of the above conclusion to assumptions such as a binary distribution of health expenditures for each consumer, competitive insurance markets, full market coverage, and no moral hazard.

3. Model

Consider a representative consumer who interacts with a firm or market. The consumer has a preference type $\theta \in \Theta$. These types can be discrete or continuous, depending on the application. They can also be seen as private actions taken previously (and nonstrategically) by the individuals. Let $U(\theta, a)$ be the consumer's indirect utility function if the consumer has type θ and the market takes action a .⁴ This indirect utility function will vary with the application considered and is discussed in the next section.⁵

The commonly known prior distribution of the consumer's type is given by $F_0(\theta)$. The market observes some data about the consumer, which we model as an informative signal $s \in S$ of the consumer's type θ . The signals s are drawn from a known information structure and allow the market to update its prior beliefs over θ . The market's action a is a function of its beliefs F_s about consumer type; that is, the market can tailor its actions to each consumer segment.⁶ Note that there is asymmetric information between the consumer and the market. The consumer knows θ , whereas the market only observes the signal s .

One extreme is the case in which the signal is fully informative of θ . Another extreme case is that of an uninformative signal, that is, full consumer privacy. More generally, the signals induce market *segmentation*

$$S = \{(\pi_s, F_s)\}_{s \in S}.$$

A segmentation is a mixture with weights π_s over distributions F_s that jointly satisfy the law of total probability

$$\int_s F_s(\theta) \pi_s ds = F_0(\theta), \quad \forall \theta \in \Theta. \quad (2)$$

In a learning interpretation (Yang 2022), a segmentation is the probability distribution over the market's *posterior beliefs* over the consumer's type. Equivalently, each signal realization s corresponds to a segment of the consumer population.⁷ The size of each segment s is given by π_s , and its composition in terms of the consumer types is given by $F_s(\theta)$.

Under segmentation S , each market segment s obtains an expected surplus of⁸

$$V(F_s) \equiv \int_{\theta} U(\theta, a(F_s)) dF_s(\theta).$$

We define *consumer welfare* as the aggregate consumer surplus across all market segments. It is then given by

$$W(S) \equiv \mathbb{E}_s[V(F_s)] = \int_s V(F_s) \pi_s ds.$$

Note that $W(S)$ represents the *ex ante* expected surplus before consumers learn their types. We use consumer welfare to refer to this aggregate (or *ex ante*) measure W , whereas consumer surplus refers to $V(F_s)$, the expected surplus for a specific segment s with distribution F_s . In the special case of *full privacy*, $S = \emptyset$,⁹ the market takes a single action $a(F_0)$, and consumer welfare is given by

$$W(\emptyset) \equiv V(F_0) = \int_{\theta} U(\theta, a(F_0)) dF_0(\theta).$$

We are interested in whether a loss of privacy—a more informative segmentation S in the Blackwell order, which we define below—is beneficial or harmful to consumer welfare. We evaluate the question of privacy before individuals learn their types or take their actions. That is, privacy policies are set when the market and individuals have symmetric information about θ , after which the market observes a signal s and takes action a , and finally, the consumer observes their type θ and may respond accordingly.

This setup aligns with many real-world environments, such as platform-mediated interactions. In such settings, consumers engage with multiple firms through a common platform, such as an online marketplace or a content ecosystem. When consumers make decisions about their privacy settings, they typically cannot anticipate which specific firms they will eventually interact with or what particular preferences they might have during future transactions. Moreover, the signals that firms receive about consumer types are derived from behaviors across multiple products and interactions, making strategic withholding or distortion of information difficult for any specific anticipated transaction.

Given this structure, we assess the value of privacy from the perspective of the entire consumer population, evaluating expected outcomes across the distribution of consumer types θ . We recognize that privacy regulation likely has heterogeneous effects on the population, and we take an aggregate welfare perspective.

Because the type space and distribution are unrestricted, the model presented here can be seen as providing a unifying framework to understand the effects of privacy in a multidimensional setting. This unifying framework can be seen as relatively straightforward, but making progress in this direction has been relatively slow.

This framework also offers an interpretation of the monetary value of privacy. When consumer utility is quasi-linear in money, the difference $W(\emptyset) - W(S)$ represents the consumer's willingness to pay to maintain privacy to prevent information disclosure corresponding to segmentation S . From the market or an observer's perspective, this same quantity can be understood as the cost of acquiring the corresponding information.

The model does not formally consider the strategic interaction between the consumer and the market and summarizes that interaction by the market action $a(F_s)$. That is, the point of the paper is just to consider the effects of the market action $a(F_s)$ and the consumer utility properties on whether privacy is beneficial or harmful to the consumer. Considering how the strategic interaction between the consumer and the market affects the market action $a(F_s)$ is beyond the scope of this paper and will depend on the specificities of the market interaction. In the applications considered below, this strategic interaction between the consumer and the market has to be formally considered to understand the properties of the market action $a(F_s)$ in that particular application.

Finally, the relationship between signals and posterior beliefs characterized by mean-preserving spreads of the prior is a core result in the Bayesian persuasion literature (Kamenica and Gentzkow 2011), where signals chosen by the sender induce distributions over the receiver's posteriors. In that framework, the sender's benefit from persuasion at every prior depends on the curvature (concavity or convexity) of their utility with respect to the receiver's posterior beliefs.

3.1. Discussion of Limitations

The setting presented here does not account directly for the possibility of individuals taking actions to protect their privacy. For example, it could be that individuals decide not to participate in the market to protect their privacy (Jullien et al. 2020). In such a case, $a(F_s)$ could represent the equilibrium action of the market if the individual did participate in the market and lost their

privacy, and $U(\theta, a(F_0))$ could account for the decision of the individual not to participate in the market.

Similarly, individuals could incur costs to preserve their privacy (Montes et al. 2019). In that case, $a(F_s)$ could represent the equilibrium action of the market if the individual did not incur costs to preserve their privacy, and $U(\theta, a(F_0))$ could account for the decision of incurring costs to protect one's privacy. One possibility, not accounted for in the model above, is that the individuals could distort their type in response to market behavior (Calzolari and Pavan 2006). This is an interesting avenue for future research.

Another issue to consider is that this framework allows comparisons between privacy regimes in which one captures beliefs that are a mean-preserving spread of the beliefs generated by the other regime, which can be a relevant comparison on the degree of privacy. However, the framework is not informative when comparing privacy regimes that cannot be ordered in terms of the mean-preserving spread.¹⁰

The setup considered here assumes that individuals are fully informed about the privacy regime in which they are making decisions. In the real world, individuals may have incomplete information about the privacy regime and only learn through experience about how they are exposed. This is an interesting general issue to consider when modeling the effects of privacy, but it is beyond the scope of this paper.

4. Effects of Privacy

As mentioned, we are interested in conditions under which a loss of privacy is harmful (or beneficial) to consumers as a whole. To formalize the comparison between different degrees of privacy, we use the Blackwell ordering (Blackwell 1951), which provides a precise way to rank information structures according to their informativeness.

Definition 1 (Privacy and the Blackwell Order). Segmentation $S = \{(\pi_s, F_s)\}_{s \in S}$ is more private than $S' = \{(\pi_{s'}, F_{s'})\}_{s' \in S'}$ if and only if $S' > S$ in the Blackwell order, that is, if and only if S' is a mean-preserving spread of S . Formally, S' is a mean-preserving spread of S if and only if there exists a transition kernel K ,¹¹ such that for all $F_s \in S$

$$\int_{S'} F_{s'}(\theta) K(s|s') \pi_{s'} ds' = F_s(\theta) \pi_s, \quad \forall \theta \in \Theta. \quad (3)$$

To build intuition for this definition, let us first examine a simple special case. Consider a scenario where segmentation S represents the prior belief, that is, the signal s is completely uninformative and $S = \{(1, F_0)\}$, and segmentation S' is induced by an informative signal s' . In this case, (3) holds with transition kernel $K(s|s') = 1$, which simplifies to the law of total probability (2). This

indicates that when market participants update their beliefs in a Bayesian manner with consistent priors, the resulting distribution of posteriors is always a mean-preserving spread of the prior distribution.¹² When comparing segmentations induced by different signals, the definition of mean-preserving spread through (3) can be seen as a generalization of the law of total probability.

To get further intuition, consider an example with binary types, in which the consumer's type is one-dimensional and takes one of two values, $\theta \in \{\theta_L, \theta_H\}$, with $\theta_H > \theta_L$. In this case, the distribution over types F_s for signal s is fully characterized by a scalar $f_s \equiv \Pr[\theta_H]$. Denote the prior by f_0 . Suppose that under the prior, both types are equally likely, $f_0 = 1/2$. The market may receive a high signal s_H or a low signal s_L .¹³ For a segmentation that is completely uninformative,¹⁴ we would have $f_{s_H} = f_{s_L} = f_0 = 1/2$. In other words, the posterior distribution over low and high types after receiving a high signal would be $(1/2, 1/2)$ and the distribution over types after receiving a low signal would also be $(1/2, 1/2)$.

Compare this now with a segmentation that is completely informative. In that case, the posterior beliefs after receiving a low signal would be $(1, 0)$, and the posterior beliefs after receiving a high signal would be $(0, 1)$. That is, the distribution of posteriors in the perfectly informative segmentation case (i.e., perfectly informative signal) is a mean-preserving spread over the prior distribution, which is the same as the distribution over posteriors of the completely uninformative segmentation (i.e., completely uninformative signal).

In the case of binary types, as the realization of posteriors is uni-dimensional, we can also order the signals s , such that a higher s corresponds to a higher posterior belief f_s on the type θ_H . We can then have that the mean-preserving spread condition (3) is equivalent to the familiar one-dimensional condition $\int_0^f (\Pi'_f - \Pi'_f) \geq 0$, $\forall f$, where Π'_f represents the cumulative distribution function of the posterior beliefs $f_{s'}$ under segmentation S' , and Π_f represents the cumulative distribution function of the posterior beliefs f_s under segmentation S .¹⁵

Note that the definition of privacy considered in this paper is based on the Blackwell order. Thus, it is equivalent to the statement that the more private segmentation S is a garbling of S' , that is, it can be seen as adding noise to the information constructed with segmentation S' through the transition kernel $K(s|s')$. Using Rothschild and Stiglitz (1970), this is also equivalent to the result that if the expected consumer surplus V is concave (convex) in the posteriors, then consumers prefer more (less) privacy, that is, $W(S) \geq (\leq) W(S')$.¹⁶

Now we turn to the central question of this paper: the relationship between information revelation (privacy loss) and consumer welfare. Consider first the case of full privacy versus an informative segmentation. If the consumer retains their privacy, they obtain $W(\emptyset)$. If the

individual consumer loses their privacy, their welfare is given by $W(S)$. Given that the distribution of the posteriors $\{F_s\}_{s \in S}$ is a mean-preserving spread of the prior F_0 , a simple application of Jensen's inequality establishes that $\mathbb{E}_s V(F_s) \geq (\leq) V(F_0)$ if $V(F)$ is a convex (concave) function of the belief distribution F .

The first result uses the definition of a mean-preserving spread to generalize this observation to segmentations induced by different signals. Consider any two segmentations S and S' . We say that consumers prefer more (less) private segmentations if $W(S) \geq (\leq) W(S')$ whenever S is more private than S' .

Proposition 1 (Value of Privacy). *Consumers prefer more (less) private segmentations if $V(\cdot)$ is concave (convex).*

Intuitively, the curvature of V determines the consumer's preferences for privacy. When V is concave, geometrically, this means the consumer surplus function "curves downward" as the market forms more varied beliefs about consumer types (represented by the spread of posteriors). For example, in the earlier binary type example, this means moving from a single point $(1/2, 1/2)$ representing the prior belief to posterior beliefs $(1, 0)$ and $(0, 1)$ reduces consumer welfare. However, what determines whether V is concave or convex? In the remainder of this section, we discuss properties of the functions $U(\theta, a)$ and of the market's action $a(F)$ that determine the shape of V . Using these properties, we provide a decomposition of three sources that explain why consumers prefer more (less) private segmentations when V is concave (convex).

4.1. Binary Types

We first consider the case of binary types to build intuition. Suppose that the consumer's type is one-dimensional and takes one of two values, $\theta \in \{\theta_L, \theta_H\}$, with $\theta_H > \theta_L$. Then, let $f \equiv \Pr[\theta_H]$, and compute the expected surplus of a consumer segment with distribution $(1 - f, f)$. With a slight change of notation (i.e., defining V to be a function of f rather than the distribution), we obtain

$$V(f) = (1 - f) U(\theta_L, a(f)) + f U(\theta_H, a(f)).$$

4.1.1. Consumer–Market Alignment. As a warm-up exercise, suppose that the market acts in the interest of the individual. This is the traditional case of a single decision maker under uncertainty. The market therefore chooses the following action:

$$a(f) = \arg \max_a [(1 - f) U(\theta_L, a) + f U(\theta_H, a)] \quad \text{for all } f \in [0, 1]. \quad (4)$$

In this case, it is a standard result that V is convex in f , and hence convex in the distribution of types. The typical proof of this result is that the right-hand side of (4) is linear in f , and therefore, the maximum V over a has to be convex in f .

It may be nonetheless useful to consider a derivation of this result with differentiation techniques. We will assume that $U(\theta, a)$ is twice continuously differentiable, that a is chosen from a continuous open set, and that the market's problem in (4) is concave.

By the envelope theorem, we have

$$V'(f) = U(\theta_H, a(f)) - U(\theta_L, a(f)),$$

Differentiating the first-order condition for a , we have

$$a'(f) = -\frac{U_a(\theta_H, a(f)) - U_a(\theta_L, a(f))}{(1-f)U_{aa}(\theta_L, a(f)) + fU_{aa}(\theta_H, a(f))},$$

and we can obtain

$$V''(f) = a'(f)[U_a(\theta_H, a(f)) - U_a(\theta_L, a(f))] \geq 0.$$

Thus, if the high-type benefits from a higher action more than the low type, then a market that acts in the consumer's interest increases a when it assigns a higher probability to θ_H .

4.1.2. General Market Actions. Without knowledge of the market's objective function, the model yields a more general expression that highlights additional economic forces. In particular, we differentiate $V(f)$ twice and let subscripts denote the partial derivatives of U . We then have

$$V'''(f) = 2a'(f)[U_a(\theta_H, a(f)) - U_a(\theta_L, a(f))] \quad (5)$$

$$+ a'(f)^2[(1-f)U_{aa}(\theta_L, a(f)) + fU_{aa}(\theta_H, a(f))] \quad (6)$$

$$+ a''(f)[(1-f)U_a(\theta_L, a(f)) + fU_a(\theta_H, a(f))]. \quad (7)$$

This expression highlights three effects of information: (5) captures that assuming a is increasing, information helps the consumer if higher types also have a higher marginal utility of a ¹⁷; (6) captures that information helps if U is convex in a on average; and (7) captures that information helps if a is convex in f and U is increasing in a on average.¹⁸

4.2. General Types

For more general type spaces, we must use a different approach to highlight these economic forces. By definition, V is convex if and only if, for any $\lambda \in [0, 1]$ and any two distributions of types F' and F'' , it holds that

$$\lambda V(F') + (1-\lambda)V(F'') \geq V(\bar{F}),$$

where $\bar{F} \equiv \lambda F' + (1-\lambda)F''$.

Using the definition of $V(\cdot)$, we can rewrite this condition as

$$\begin{aligned} & \int_{\theta} \lambda [U(\theta, a(F')) - U(\theta, a(\bar{F}))] dF'(\theta) \\ & + \int_{\theta} (1-\lambda) [U(\theta, a(F'')) - U(\theta, a(\bar{F}))] dF''(\theta) \geq 0. \end{aligned}$$

We can expand the expression above to obtain the following characterization of convex and concave surplus functions V .¹⁹

Proposition 2 (Convexity of V). *Consumer surplus $V(F)$ is convex (concave) in F if and only if*

$$\lambda(1-\lambda) \int_{\theta} [U(\theta, a(F')) - U(\theta, a(F''))][dF'(\theta) - dF''(\theta)] \quad (8)$$

$$+ \int_{\theta} [\lambda U(\theta, a(F')) + (1-\lambda)U(\theta, a(F'')) - U(\theta, a(\bar{a}))] d\bar{F}(\theta) \quad (9)$$

$$+ \int_{\theta} [U(\theta, \bar{a}) - U(\theta, a(\bar{F}))] d\bar{F}(\theta) \geq (\leq) 0, \quad (10)$$

for all F', F'' , and $\lambda \in [0, 1]$, where

$$\bar{a} \equiv \lambda a(F') + (1-\lambda)a(F'').$$

Note that the terms in (8) are nil when consumer indirect utility U is additively separable in type θ and market action a . Moreover, (9) is zero when U is linear in a , and (10) is zero when a is linear in the belief distribution F . All three conditions hold simultaneously in the knife-edge case where V is both convex and concave, making consumers indifferent to privacy. In the health insurance example, this occurs when consumers' direct utility function u is linear, leading to an indirect utility function U that is additively separable in θ and a , and with a linear in F . To gain some intuition for Proposition 2, we consider three particular cases, each isolating one of the effects captured by (8), (9), and (10).

4.2.1. Degree of Consumer-Market Interest Alignment.

The first example is regarding (8), which measures whether consumer's and market's interests are aligned. When high-type consumers benefit more from certain market actions than low-type consumers, and markets take those actions more frequently when they believe they face high-type consumers, privacy becomes less valuable.

Suppose that consumer indirect utility U is linear in market action a and that the market action a is linear in belief distribution F . In this case, the effects (9) and (10) are zero. We then obtain the following simpler conditions.

Corollary 1 (Linear Utility and Actions). *If U is linear in a and a is linear in F , the condition for V 's convexity is given by*

$$\int_{\theta} [U(\theta, a(F')) - U(\theta, a(F''))][dF'(\theta) - dF''(\theta)] \geq 0, \quad (11)$$

for all F', F'' .

When $U(\theta, a)$ is linear in a and $a(F)$ is linear in F , the shape of $V(F)$ is determined by the correlation between consumer marginal utility for market actions and the market's response to consumer types.²⁰ If the high

consumer types like increases in a more than the low types and the market chooses a greater a when it has a higher belief distribution, the consumer prefers less privacy. Conversely, V is concave and privacy is preferred if the types that prefer action $a(F')$ are also less likely under distribution F' , that is, if the interests of the market and the agent are misaligned.

As an illustrative example, consider a content platform that customizes the quality level a based on user types. The consumer's indirect utility function is given by²¹

$$U(\theta, a) = \theta a.$$

This means that the higher types gain more utility from quality improvements. The platform's revenue comes from advertising, which increases with user engagement time (proportional to θa). With quadratic costs of quality provision, the platform's quality choice becomes proportional to its posterior mean:

$$a(F) \propto \mathbb{E}_F[\theta].$$

In particular, the platform provides higher-quality content when it believes its audience has stronger preferences. Applying Condition (11) yields

$$(\mathbb{E}_F[\theta] - \mathbb{E}_{F'}[\theta])^2 \geq 0,$$

which is always satisfied. Therefore, V is convex, and consumers benefit from less privacy. Intuitively, when the platform knows more about consumer preferences, it can better tailor content quality levels to match those preferences, creating more value for both parties.²²

This example illustrates that privacy preferences can emerge even when consumers are risk-neutral in market actions (as indicated by the linear utility function). In this case, these preferences stem from the correlation between consumer types and market actions rather than from risk aversion. For more general functions U and a , Expressions (9) and (10) uncover the further effects, including those related to risk attitudes, which we explore in the following examples.

4.2.2. Consumer Preferences for Variability in Market Actions. The second example shows that the term in (9) captures the effect that with more information, market actions are more varied and that can be beneficial or harmful to consumers depending on whether the indirect utility function $U(\theta, a)$ is convex or concave in a . Specifically, the convexity of U in a is a force for the consumer to prefer less privacy.

To isolate this effect, suppose that consumer indirect utility U is additively separable in consumer type θ and market action a , meaning type does not affect how consumers value the action and that the market action a is linear in belief distribution F . This nullifies Effects (8) and (10), leading to the following.

Corollary 2 (Separable Utility and Linear Actions). *If U is additively separable in θ and a , and a is linear in F , and then V is convex if and only if $U(\theta, a)$ is convex in a , for a in the range of $a(F)$.*

When $U(\theta, a)$ is additively separable in θ and a and $a(F)$ is linear in F , the shape of $V(F)$ is determined by the convexity/concavity of $U(\theta, a)$ in a . If consumers prefer more variable market actions, they prefer less privacy.

As an illustrative example, consider a content platform similar to the previous example, but where consumer utility $U(\theta, a)$ is additively separable in θ and a . When the platform chooses quality a , the consumer's indirect utility function is given by²³

$$U(\theta, a) = a^2.$$

That is, the consumer's indirect utility is independent of types, and the consumer prefers content with higher variability. The platform's advertising revenue is proportional to the time the consumer spends on the platform (proportional to a) multiplied by the ads' click-through rate (proportional to θ). With quadratic quality costs, the platform's quality choice is proportional to its posterior mean ($a(F) \propto \mathbb{E}_F[\theta]$), that is, it provides higher-quality content when it has a higher belief distribution, as in the earlier example. Then, Corollary 2 implies that V is convex and the consumer benefits from less privacy. Intuitively, this is because a loss of privacy leads to higher variability in the platform's quality choice, and the consumer prefers this variation because they value quality improvements at an increasing rate.²⁴

4.2.3. Market Response Curvature to Information. Finally, the term in (10) captures the effect of whether the market action is convex or concave in the market's beliefs. Therefore, if the consumer on average likes a higher market action, then the convexity of a also induces the consumer to dislike privacy.

To isolate this effect, we consider a setting where consumer indirect utility U is additively separable in consumer type θ and market action a , and is linear in the market action a .

Corollary 3 (Separable and Linear Utility). *If U is additively separable in θ and a , and U is linear and increasing in a , V is convex if and only if a is convex.*

When $U(\theta, a)$ is additively separable in θ and a and linear in a , the shape of $V(F)$ is determined by the convexity/concavity of a in F . Suppose that $a(F)$ is convex in F and consumers prefer higher market actions, then, with more information, the market chooses a higher a on average. Thus, the convexity of $a(F)$ would be a force for the consumer to prefer less privacy.

As an illustrative example, we again consider the content platform example but now with a linear consumer utility in a . The consumer's indirect utility

function under quality a is given by²⁵

$$U(\theta, a) = a.$$

That is, consumer indirect utility is independent of types. Similar to the second setup, the platform's advertising revenue is proportional to the time a the consumer spends on the platform multiplied by the consumer's type θ .

Unlike the earlier examples, we now vary the platform's cost structure. The platform incurs costs proportional to $a^{1+\rho}$, where $\rho > 0$ is a constant. The parameter ρ captures different production technologies. A small ρ can capture the economies of scale or learning-by-doing effects, and a large ρ can capture resource constraints. This cost structure leads to the platform's quality choice

$$a(F) \propto (\mathbb{E}_F[\theta])^{\frac{1}{\rho}}.$$

Then, Corollary 3 implies that when $0 < \rho < 1$, V is convex, and the consumer benefits from less privacy. Intuitively, economies of scale or learning-by-doing effects enable the platform to disproportionately increase quality when it has more precise information about high-value users. In contrast, when $\rho > 1$, resource constraints make $a(F)$ concave, leading consumers to prefer privacy.

5. Applications

We now consider several applications of the general framework that we present.

5.1. Price Discrimination with Linear Demand

Consider the case of price discrimination with linear consumer demand. Let us apply the results in Proposition 2 to a setting with multidimensional consumer types and market actions.

For this application, consider a consumer with a two-dimensional type (θ_1, θ_2) . The market first observes a signal of the consumer's type and determines the prices p_1 and p_2 of two products. Then, the consumer observes their type and chooses the quantities purchased. Let the direct utility function be given by

$$u(\theta, p, q) = (\theta_1 - p_1)q_1 + (\theta_2 - p_2)q_2 + \gamma q_1 q_2 - q_1^2/2 - q_2^2/2,$$

where p is the vector of prices of the two products and q is the vector of quantities purchased of each product. The parameter $\gamma \in [-1, 1]$ denotes the degree of complementarity between the two products. We assume that θ_1 and θ_2 are independently distributed in the prior and on any posterior distribution after any signal. We also assume that the costs of production are zero. Furthermore, we assume that $\mathbb{P}_F(\theta_i \geq \mathbb{E}_F[\theta_i]/2) = 1$, that is, θ_1 and θ_2 are not too low. This ensures that the demands for both products from any consumer type are nonnegative. Given the direct utility function, the

consumer purchases the following quantities:

$$q_i = \frac{\theta_i - p_i + \gamma(\theta_j - p_j)}{1 - \gamma^2}, \quad i = 1, 2. \quad (12)$$

The consumer's indirect utility function, when taking into account their optimal quantities purchased, is then given by

$$U(\theta, p) = \frac{(\theta_1 - p_1)^2 + (\theta_2 - p_2)^2 + 2\gamma(\theta_1 - p_1)(\theta_2 - p_2)}{2(1 - \gamma^2)}.$$

We separately consider the case of monopoly (one firm sells both goods) and price competition (two firms, each selling one good). The monopoly prices for Demand System (12) under distribution F are given by, for all $\gamma \in (-1, 1)$,

$$p_i^*(F) = \mathbb{E}_F[\theta_i]/2.$$

This immediately yields an expression for the consumer's value function in a segment with type distribution F :

$$V(F) = \frac{\mathbb{E}_F[\theta_1]^2 + \mathbb{E}_F[\theta_2]^2 + 2\gamma \mathbb{E}_F[\theta_1] \mathbb{E}_F[\theta_2] + 4(\text{Var}_F[\theta_1] + \text{Var}_F[\theta_2])}{8(1 - \gamma^2)}.$$

We can then apply Proposition 2 and obtain the following proposition. Note that the market actions (p_1, p_2) are linear in distribution F , so the third term (10) is nil.

Proposition 3 (Monopoly Price Discrimination). *In the monopoly pricing model with linear demand, consumers unambiguously like privacy.*

This result mirrors the classic analysis of third-degree price discrimination by Robinson (1933) and Schmalensee (1981). This effect comes from the fact that the first term (8) is negative, the second term (9) is positive, and the first term dominates the second term. Intuitively, the first term, measuring the degree to which the consumer's and the market's interests are aligned, is negative because higher types suffer more from higher prices, but with an informed monopolist, the prices will increase more for segments with higher types on average. The second term, measuring consumer preferences for variability in market actions, is positive because of the convexity of the indirect utility function in prices. This convexity arises because consumers optimally choose higher quantities when prices are lower, creating an increasing benefit from price reductions. One could have imagined that this convexity is a force that drives the value of privacy down, but in this quadratic model, this is insufficient to overcome the first term.

For managers and policymakers, this result suggests that privacy regulation protecting consumer data from monopolistic sellers could benefit consumers overall. It may prevent firms from segmenting markets and

charging personalized prices that would extract more surplus from high-value customers. Managers of monopolistic firms seeking to implement personalized pricing strategies should recognize that such practices create consumer welfare losses and may trigger regulatory scrutiny or consumer backlash.

We now turn to the case of price competition and show that the result extends to this case. The Nash equilibrium prices for a commonly known distribution of types F and Demand System (12) are given by

$$p_i^{\text{NE}} = \frac{(2 - \gamma^2) \mathbb{E}_F[\theta_i] + \gamma \mathbb{E}_F[\theta_j]}{4 - \gamma^2}, \quad i, j = 1, 2.$$

Substituting into the utility function and taking expectations over types θ under distribution F , we obtain

$$V^{\text{NE}}(F) = \frac{(4 - 3\gamma^2)(\mathbb{E}_F[\theta_1]^2 + \mathbb{E}_F[\theta_2]^2) + 2\gamma^3 \mathbb{E}_F[\theta_1] \mathbb{E}_F[\theta_2]}{2(4 - \gamma^2)^2(1 - \gamma^2)} + \frac{\text{Var}_F[\theta_1] + \text{Var}_F[\theta_2]}{2(1 - \gamma^2)}.$$

Proposition 4 (Duopoly Price Discrimination). *In the duopoly model with linear demand, consumers unambiguously like privacy.*

The intuition of Proposition 4 is similar to that of Proposition 3: In the duopoly setting, the first term (8) also dominates the second term (9). Interestingly, competition between firms does not change the privacy preferences of consumers in this setting. This suggests that competitive forces alone are insufficient to address consumer privacy concerns regarding price discrimination. Managers in competitive markets should note that access to consumer data does not necessarily benefit consumers. Policymakers should consider that privacy regulations may benefit consumers in seemingly competitive markets.

Finally, we consider the case of a price- and quality-discriminating monopolist introduced by Argenziano and Bonatti (2023). Let the consumer's direct utility function from consuming q units of good of quality y be given by

$$u(\theta, p, q, y) = (\theta + by - p)q - q^2/2,$$

where b is a parameter that measures every consumer's type marginal valuation for quality. Consumer type θ 's demand is then given by

$$q(\theta, y, p) = \theta + by - p.$$

A monopolist seller can produce quality y at a quadratic cost

$$c(y) = cy^2/2, \quad c > b^2/2.$$

Argenziano and Bonatti (2023) show that the optimal price and quality levels for a monopolist with beliefs F

over the consumer's type are given by

$$p^* = \frac{\mathbb{E}_F[\theta]c}{2c - b^2} \quad \text{and} \quad y^* = \frac{\mathbb{E}_F[\theta]b}{2c - b^2}.$$

This allows us to write the expected consumer surplus in a segment F as

$$V(F) = \frac{1}{2}[\mathbb{E}_F[\theta^2] + \beta(2 + \beta)\mathbb{E}_F[\theta]^2],$$

where

$$\beta \equiv \frac{b^2 - c}{2c - b^2}.$$

Note that the term $\mathbb{E}_F[\theta^2]$ (the expectation of the random variable θ^2) is by definition linear in the probability distribution F . The term $\mathbb{E}_F[\theta]^2$ (the square of the expectation of θ) is convex in F . Also, the assumption $c > b^2/2$ implies $2 + \beta > 0$. Therefore, the sign of β determines the convexity/concavity of V .

Proposition 5 (Quality and Price Discrimination). *In the model with quality and price discrimination, consumers like privacy if and only if $c > b^2$, that is, the cost of investing in quality is high, or consumers' marginal valuation for quality is low.*

In this example, the third term (10) is also nil, and whether the first term (8) dominates the second (9) depends on the parameter values. Intuitively, higher types suffer more from higher prices and benefit more from higher quality levels. When the cost of investing in quality is high, or consumers' marginal valuation for quality is low, the prices are sensitive to consumer types (which harms consumers), whereas the quality levels are not sensitive enough to consumer types (which benefits consumers, but insufficiently). This would be a force for the first term to be more negative, resulting in an effect of preference for privacy.

The result implies that managers should carefully assess their cost structure and consumer preferences when determining data collection strategies. In markets with high costs in quality improvements or where quality improvements deliver limited value to consumers, firms should expect resistance to personalization efforts. Policymakers may need to develop context-specific privacy regulations that consider industry cost structures rather than implementing one-size-fits-all approaches. This parameter-dependent result illustrates how the theoretical framework provides nuanced guidance for privacy regulation across different market conditions.

5.2. Targeting with Product and Price

A firm chooses both its location and its price. Consider a unit segment and a consumer who is located at either zero or one. The consumer's location is their type $\theta \in \{0, 1\}$. The type- θ consumer's downward-sloping

demand function is $D(p) = 2 - |\theta - x| - p$, if the product is located at x and the firm charges a price p . The consumer's indirect utility function is given by $U(\theta, x, p) = (2 - |\theta - x| - p)^2/2$.

The firm can only offer one product and has zero production costs. Initially, the firm has a prior belief over the consumer's location. In addition, the firm receives a signal s of the location of the consumer, which is fully summarized by the posterior probability

$$\beta \equiv \Pr[\theta = 0 | s].$$

The problem of the firm is to choose where to position itself and what price to charge. In this example, the market action has two dimensions: the location of the product, x , and the price charged, p . Let $x(\beta)$ and $p(\beta)$ be the location and price that the firm chooses after receiving a signal, respectively. The firm therefore solves the following problem:

$$\max_{x,p} p[\beta(2 - x - p) + (1 - \beta)(1 + x - p)]. \quad (13)$$

Given that the expected profit is linear in x , the firm will choose to locate at $x = 0$ if and only if $\beta > 1/2$. Therefore, the firm's optimal price choice is given by

$$p^*(\beta) = (1 + \max\{\beta, 1 - \beta\})/2.$$

Finally, the surplus of a consumer in a segment with proportion β of types $\theta = 0$ satisfies

$$2V(\beta) = (2 - p^*(\beta))^2 \max\{\beta, 1 - \beta\} + (1 - p^*(\beta))^2 \min\{\beta, 1 - \beta\}. \quad (14)$$

Figure 1 illustrates the result for this example.

Let us now consider using Proposition 2 for this application. From the discussion above, we can see that an increase in β , from below $1/2$ to above $1/2$, can induce a discrete change in the firm's location from $x = 1$ to $x = 0$. The direction of this change benefits the consumer on average. Intuitively, the revelation of information can make a firm choose a location that better matches consumer preferences. This effect will tend to dominate in (8). To see that an increase in β makes the

firm execute a substantial change in location, consider a variation of the model above where there are convex costs of the firm locating farther away from $1/2$. The lower the convexity of these costs is, the greater the effect an increase in β on the firm's change in location, and this application considers the case where this convexity is the lowest, so that the change in location is the highest.

Consider now the effect of the price decision. Holding fixed the firm's location at $x = 0$, an increase in β makes the firm increase its price, so the welfare effect of information through pricing is negative, similar to the intuition in the price discrimination examples. As this effect is small compared with the location decision in this application, the first term (8) is positive.


Furthermore, the second term (9) is positive because consumer indirect utility $U(\theta, x, p)$ is convex in both location and price. This convexity arises, similar to the price discrimination examples, because consumers can adjust their demand in response to prices and product locations. Finally, if the revelation of information does not modify the firm's location choice, the third term (10) is zero at the firm's optimum, because the pricing decision is linear in β . We summarize this discussion in the following proposition.

Proposition 6 (Location and Price Targeting). *Full privacy is optimal for the consumer if the revelation of information does not modify the firm's location choice. Conversely, the consumer may benefit from two-signal segmentations that induce the firm to choose different locations.*

This model can also be interpreted as targeted advertising. Consider two types of consumers that are interested in two products respectively, located at $\theta \in \{0, 1\}$. Each product has an ad, and a platform chooses which ad to display (x) and the frequency at which it is displayed (p). When consumers visit the platform, they observe which ad is displayed and spend time $D(p) = 2 - |\theta - x| - p$ on the platform. This engagement time $D(p)$ is decreasing in the mismatch between consumer interest and ad content $|\theta - x|$ and the ad frequency p . Moreover, the platform's revenue comes from advertising, which is proportional to user engagement time multiplied by ad frequency.²⁶ In this interpretation, the results show that consumers benefit from the revelation of information only when it causes the platform to significantly improve ad relevance. This illustrates the theoretical framework's first effect. When information revelation leads to better preference-content matching, consumers may dislike privacy. However, if information merely enables higher ad loads without improving relevance, consumers prefer privacy.

These examples illustrate that the revelation of the same piece of information can be welfare enhancing or welfare reducing, depending on how firms use it.²⁷ The results offer important implications for managers and

Figure 1. (Color online) Consumer Surplus



platforms. For managers, the findings suggest that consumers benefit from personalization when it significantly improves product positioning or ad relevance but are harmed when data are primarily used for price discrimination or personalizing ad load. To enhance consumer welfare, managers should prioritize using consumer data to better match products to consumer preferences or improve ad relevance. From a policy perspective, regulations should distinguish between different uses of consumer data rather than treating all data collection uniformly. Privacy policies that allow data use for improving product matching whereas limiting price discrimination could enhance consumer welfare. Policymakers might consider context-specific regulations that permit targeted content but restrict personalized pricing.

5.3. Price Competition with Targeting

Consider the case of Thisse and Vives (1988). The consumer's type θ is uniformly distributed on the segment $[0, 1]$ and two competing firms are located at the extremes of that segment. A consumer buys one unit and gets utility $v - x - p$ when buying from a firm located at a distance x from the consumer at a price p . Let p_0 and p_1 be the prices of the firms located at zero and one, respectively. The distance from the firm located at zero is $x = \theta$, and the distance from the firm located at one is $x = 1 - \theta$. We assume v is large and production costs are nil.

The firms receive a common signal about the location of a consumer. That signal is fully informative and tells the exact location of the consumer with probability α , and is completely uninformative with probability $1 - \alpha$. Part of the information of the signal is also whether the signal is fully informative or fully uninformative. After receiving a fully informative signal the posterior of the firms is that the location of the consumer is at θ with probability one, which we denote by $\mathbf{1}_\theta$. That is, after receiving a fully informative signal, the beliefs of the firms are one for type θ and zero for all other types. After receiving a fully uninformative signal, the posterior of the firms is the prior, which is the uniform distribution on the segment $[0, 1]$, which we denote by $\mathcal{U}$.

We interpret increasing α as decreasing the privacy in the market. When $\alpha = 0$, there is complete privacy; when $\alpha = 1$, there is no privacy. Firms compete on prices and can choose a price contingent on the signal received. For the consumers for whom the firms receive a fully uninformative signal, firms charge a price equal to one, and the consumer buys from the firm that is closer.²⁸

For consumers of type t for whom the firms receive a fully informative signal, the equilibrium has the firm further away from that consumer charging a price of zero, and the firm closer to the consumer charging a price of $p = 1 - 2\min\{\theta, 1 - \theta\}$. In either the case of fully

uninformative or fully informative case, consumers end up purchasing from the firm that is closer to them. Let m be the market price that is ultimately paid by the consumer, then the consumer's indirect utility is

$$U(\theta, m) = v - \min\{\theta, 1 - \theta\} - m.$$

Moreover, the above discussion shows that $m(\mathcal{U}) = 1$, and $m(\mathbf{1}_\theta) = 1 - 2\min\{\theta, 1 - \theta\}$.

From the calculations above, we can obtain the consumer welfare under full privacy as

$$W_0 = V(\mathcal{U}) = v - 5/4,$$

and with no privacy as

$$W = \mathbb{E}[V(\mathbf{1}_\theta)] = \int_0^1 V(\mathbf{1}_\theta) d\theta = v - 3/4.$$

As α represents the fraction of consumers for whom the firms receive a fully informative signal, we have that an increase in α , that is, a decrease in privacy, benefits the consumers.

The intuition behind this result is related to the consumer-market interest alignment. When firms receive a signal about consumer locations, they compete more intensely for consumers they can identify, leading to lower prices. In this case, better information leads to more targeted price competition rather than more effective surplus extraction. As information allows firms to compete more aggressively for each consumer segment, prices for identified consumers are lower than when all consumers are served under uniform pricing.

For managers, this finding suggests that in markets with differentiated products, more granular consumer data could be positioned as proconsumer because it intensifies price competition for identified consumers. For policymakers, the result indicates that privacy regulations might need to be more nuanced in competitive settings with differentiated products, as restricting information flows could potentially reduce consumer welfare by softening price competition.

5.4. Quality Pricing and Screening

We now introduce the effects of privacy in a canonical model of second-degree price discrimination. Consider a price-discriminating (Mussa and Rosen 1978) monopolist who faces a consumer with a binary type $\theta \in \Theta = \{\theta_L, \theta_H\}$. The consumer's direct utility for a product with quality q is given by θq , and the monopolist's cost function is $c(q) = q^2/2$.

The monopolist first observes a signal of the consumer's type and then offers a menu of qualities and prices. Afterward, the consumer observes their type and chooses which quality-price combination to

purchase. A monopolist with beliefs $p \equiv \Pr[\theta = \theta_H | s]$ for the fraction of high types in a segment $s \in \Delta\Theta$ offers the following menu of qualities:

$$q_H = \theta_H, \quad (15)$$

$$q_L = \max \left\{ 0, \theta_L - \frac{p}{1-p}(\theta_H - \theta_L) \right\}. \quad (16)$$

In this example, the market action consists of q_H and q_L . When q_H and q_L are set, the optimal prices for the two products are fully determined by incentive compatibility and participation constraints and do not depend directly on the belief p . The consumer indirect utility function, taking into account the monopolist's optimal prices and the consumer's optimal choice, is given by

$$U(\theta, q_H, q_L) = q_L \mathbf{1}[\theta = \theta_H](\theta_H - \theta_L),$$

that is, only the high type obtains an information rent $q_L(\theta_H - \theta_L)$. Then, the surplus of a consumer in a segment s where $\Pr[\theta_H | s] = p$ is

$$V(p) = q_L(p)p(\theta_H - \theta_L).$$

Clearly, we have $V(0) = V(1) = 0$; that is, the monopolist can extract all the consumer's surplus under complete information. Therefore, zero privacy is never optimal.

Now, let p_0 denote the prior probability of the high type θ_H . A straightforward application of concavification techniques yields the following result on the optimality of full privacy.

Proposition 7. *Zero privacy is never optimal for the consumer. Let $\tilde{V}$ be the concave closure of V . Full privacy is optimal if and only if $V(p_0) = \tilde{V}(p_0)$, that is, if and only if*

$$p_0 \leq p^* \equiv \frac{\theta_L}{2\theta_H - \theta_L}.$$

Figure 2 shows V and its concave closure $\tilde{V}$ as a function of p .

The intuition for Proposition 7 is straightforward: If the likelihood of a high type is sufficiently high, the

monopolist shuts down the low type. Therefore, if the high types are sufficiently likely in the population, consumers benefit on average from revealing some information (to increase their rents after one signal, as indicated by concavification). Otherwise, the strong intuition is that consumers benefit from holding asymmetric (superior) information versus the firm.

Applying Proposition 2 to this application, we can decompose the effects of information. When the market's belief for the high type increases, the first term (8), measuring the degree of consumer-market interest alignment, is negative because $q_L(p)$ is weakly decreasing in p . With a higher likelihood of a high type, the monopolist reduces the quality offered to low types, decreasing information rents for the high type. This reflects misaligned interests between the monopolist and the consumer, which is a force for consumers to prefer privacy. The second term (9) is nil because the indirect utility function is linear in qualities. The third term (10), measuring the curvature of market actions in beliefs, is negative when p is sufficiently low, because $q_L(p)$ is concave in p in this region. However, when p is sufficiently high, the shutdown of the low type creates a kink $q_L(p)$, making (10) potentially positive and dominating (8). This convexity of the market action explains why partial information revelation can benefit consumers when the prior probability of high types exceeds p^* .


This example allows for an immediate computation of the consumer-optimal information structure: it is either full privacy or a two-point distribution with support $\{p^*, 1\}$. This feature is difficult to generalize (as a large amount of literature on persuasion has shown) to multidimensional environments or environments where the posterior mean is not a sufficient statistic for the firm's problem.

Unlike earlier examples, this screening setting shows that neither zero privacy nor full privacy is universally optimal for consumers. When high-value consumers dominate the market, partial information disclosure can benefit consumers on average by inducing firms to offer better quality-price combinations for consumers in low-value segments. Product managers might consider offering consumers options to selectively reveal information about their preferences when the market composition would make such disclosure welfare enhancing. From a policy perspective, regulators should develop market-specific privacy frameworks that account for customer composition rather than applying uniform restrictions. Privacy regulations might be most valuable in markets with fewer high-value customers, where information revelation enables firms to extract more surplus.

5.5. Health Insurance Revisited

We now turn to a richer version of the health insurance example. An important assumption in the motivating

Figure 2. (Color online) Consumer Surplus ($\theta_L = 3, \theta_H = 4$)



example (Section 2) was that the competitive equilibrium in the insurance industry yields full coverage for every consumer type. We now relax that assumption by assuming that, as a reduced form for a model with partial contractibility (e.g., moral hazard), the insurance industry can cover only a fraction of the consumer's health expenditures. Let $\alpha \in (0, 1)$ denote the fraction of *uninsurable* expenditures.

5.5.1. Binary Health Expenditures. In the binary-expenditure example, the indirect utility function of a consumer of type θ who purchases partial coverage $(1 - \alpha)L$ at a premium p is given by

$$U(\theta, p) = \theta u(w - p - \alpha L) + (1 - \theta)u(w - p).$$

In this example, the insurance premium p is a function only of the average type in each segment. We therefore denote the market action ($a(F)$ in the previous section) as $p(\mu)$. In particular, the equilibrium premium for a market segment s with $\mathbb{E}_s[\theta] = \mu$ is given by

$$p(\mu) = (1 - \alpha)\mu L.$$

Therefore, the expected surplus of a consumer in market segment s is given by

$$V(\mu) = \mu u(w - (1 - \alpha)\mu L - \alpha L) + (1 - \mu)u(w - (1 - \alpha)\mu L). \quad (17)$$

We know from Proposition 1 that consumers as a whole prefer to maintain privacy vis-à-vis the insurance industry if $V(\mu)$ is concave and to reveal as much information as possible if $V(\mu)$ is convex. Although this example allows for an arbitrary distribution of types, the segment-level surplus is a function of one variable only ($\mu \in [0, 1]$). We can then apply the differential approach of (5)–(7) and obtain

$$\begin{aligned} V''(\mu) = & 2(1 - \alpha)L[u'(w - (1 - \alpha)\mu L) \\ & - u'(w - (1 - \alpha)\mu L - \alpha L)] \\ & + (1 - \alpha)^2 L^2 [\mu u''(w - (1 - \alpha)\mu L - \alpha L) \\ & + (1 - \mu)u''(w - (1 - \alpha)\mu L)]. \end{aligned}$$

The differential approach is useful in this case because both terms can be signed and interpreted. In particular, the strict concavity of the utility function u implies that both terms are negative. Intuitively, the first term corresponds to the first effect (5) in Proposition 1, measuring the alignment between consumer and market interests. This term is negative because, with informed insurance companies, the premium will on average increase more for consumers who are more

likely to incur the loss L . These higher-risk consumers have a larger marginal utility of income. The second term captures (6), measuring consumer preferences for variability in market actions (premiums). This term is negative because consumers' indirect utility function is strictly concave, making them prefer less variability in the premium. This is the same risk-pooling benefit observed in the motivating example. Moreover, the third effect (7) is nil in this case, as the premium is linear in the market belief.

This allows us to reach an intuitive but important conclusion, namely that consumers prefer full privacy in this model. In other words, consumer welfare is maximized in this example by letting the insurance market operate under the prior distribution. This implies that privacy protections in health insurance markets can benefit consumers on average, particularly when insurers charge risk-based premiums. Although precision pricing might seem efficient from a market perspective, it undermines the welfare-enhancing function of insurance by transferring the classification risk to consumers.

The above conclusion, however, relies on the binary nature of the uncertainty facing each consumer type, that is, the health expenditure level $\{0, L\}$, as we show below.

5.5.2. General Health Expenditures. Now, assume that the health expenditures L of each consumer type θ follow a more general distribution with cumulative distribution function $G(L; \theta)$. Because the market is assumed to be perfectly competitive, the price of partial insurance in any segment s with composition $F_s(\theta)$ is proportional to the mean expenditure in that segment, that is,

$$p(\mu_s) = (1 - \alpha)\mathbb{E}_{\mu_s}(L) = (1 - \alpha) \int_L L d\mu_s(L),$$

where

$$\mu_s(L) \equiv \int_{\theta} G(L; \theta) dF_s(\theta).$$

Note that the premium is linear in market beliefs, so the third effect (10) is nil, as in the binary expenditure application.

The segment-level consumer surplus under an arbitrary distribution μ is then given by

$$V(\mu) = \int_L u(w - (1 - \alpha)\mathbb{E}_{\mu}(L) - \alpha L) d\mu(L).$$

We can apply Proposition 2 in this case to obtain conditions in which consumers are harmed or not. Consider two distributions, μ' and μ'' , and let $\bar{\mu} \equiv \lambda \mu' + (1 - \lambda) \mu''$ denote the convex combination of the two distributions with weights λ and $1 - \lambda$, respectively. Finally, we let $\bar{L}', \bar{L}''$ and $\bar{L} \equiv \lambda \bar{L}' + (1 - \lambda) \bar{L}''$ denote the means of these three distributions, respectively. We assume

without loss that $\bar{L}' > \bar{L} > \bar{L}''$. We can then write the difference

$$\Delta \equiv \lambda V(\mu') + (1 - \lambda) V(\mu'') - V(\bar{\mu}),$$

and simplify it as follows:

$$\begin{aligned} \Delta &= \lambda \int_L u(w - (1 - \alpha)\bar{L}' - \alpha L) d\mu' \\ &\quad + (1 - \lambda) \int_L u(w - (1 - \alpha)\bar{L}'' - \alpha L) d\mu'' \\ &\quad - \int_L u(w - (1 - \alpha)\bar{L} - \alpha L) d\bar{\mu} \\ &= \lambda(1 - \lambda) \int_L [u(w - \alpha L - (1 - \alpha)\bar{L}') \\ &\quad - u(w - \alpha L - (1 - \alpha)\bar{L}'')] (d\mu' - d\mu'') \\ &\quad + \int_L [\lambda u(w - \alpha L - (1 - \alpha)\bar{L}') \\ &\quad + (1 - \lambda) u(w - \alpha L - (1 - \alpha)\bar{L}'') \\ &\quad - u(w - \alpha L - (1 - \alpha)\bar{L})] d\bar{\mu}. \end{aligned}$$

Regarding the first term, the integrand

$$u(w - \alpha L - (1 - \alpha)\bar{L}') - u(w - \alpha L - (1 - \alpha)\bar{L}'')$$

is a decreasing function of L because u is concave and $\bar{L}' > \bar{L}''$. Therefore, a sufficient condition for the first term to be negative is that μ' first-order stochastically dominates μ'' . Intuitively, consumers with stochastically larger health expenditures are more sensitive to premium changes because of their higher marginal utility of income. When insurers gain more information, these consumers face disproportionately higher premiums. Moreover, the second term is negative because $u(\cdot)$ is strictly concave.²⁹ We summarize these effects in the following proposition.

Proposition 8 (Partial Health Insurance). *Consumers are harmed by market segmentations that rank the distributions of expenditure levels in each segment by first-order stochastic dominance. Conversely, there exist utility functions for which consumers are not unambiguously harmed by market segmentations that do not rank groups by first-order stochastic dominance.*

The last observation reconciles the results for the binary expenditure case and the more general case: with only two health expenditure levels per consumer $\{0, L\}$, every probability shift between the two outcomes is a first-order shift. However, in more general cases, the welfare effects of information can be ambiguous, as illustrated in the following example.

Example 1 (Partial Insurance). Let the consumer's utility function be given by $u(x) = u_0 x - x^2/2$, with $u_0 > w + L/6$. Consider two distributions μ' and μ'' . The distribution μ' places probability mass $1/3$ on L and

mass $2/3$ on $L/2$. The distribution μ'' has $1/2$ mass on L and $1/2$ on 0 , so neither distribution first-order stochastically dominates the other. Let $\lambda = 1/2$ so that $\bar{L}' = 2L/3$, $\bar{L}'' = L/2$ and $\bar{L} = 7L/12$. Then it holds that

$$\Delta = \frac{(1 - \alpha)L}{24} \left(u_0 - w + (2 - 3\alpha)\frac{L}{6} \right) > 0.$$

In this example, neither distribution first-order stochastically dominates the other, and the first effect (with informed insurance companies, the premium will on average increase more for consumers who are more likely to incur the loss) no longer appears here.

The results in the general case suggest that health insurance privacy regulations might be most important in markets where risks are clearly stratified, such as genetic conditions with high predictive power. In markets with more complex risk distributions, a more nuanced approach to privacy regulation may be warranted.

5.6. Hacking

We consider a minimal model of hacking to provide a microfoundation for privacy preferences under the threat of data leakages. Suppose that an individual can have two types $\theta \in \{A, B\}$. Hackers receive a signal of the type that leads to posterior beliefs $p = \Pr[\theta = A]$.

The hackers choose which type θ to target with their attack. Here, we assume that they target type $\theta = A$ if and only if $p \geq 1/2$. A successful attack (i.e., one directed at the true type) causes damage captured by $k(p)$, which is a positive, convex function that is symmetric about $1/2$. We normalize $k(1/2)$ to be zero. This formulation captures the idea that if the signal is not very accurate (i.e., if the hackers are not very confident in their information), then the signal only captures the less important components of the individual's identity. Therefore, the hackers inflict limited damage on the individual. As the signal becomes more accurate, the hackers obtain more valuable information about the individual and can inflict increasingly greater damage.

In terms of the framework, let the a represent the type θ to which the attack is directed and k the extent to which the type is harmed. Let v denote the utility of a type that suffers an unsuccessful (i.e., mismatched) attack. Then, the individual's indirect utility is given by $v - \mathbf{1}[\theta = a]k(p)$, where the attack action as a function of beliefs $a(p)$ is the more likely type. One can denote the market action as a two-dimensional vector

$$(h_A, h_B) \equiv (\mathbf{1}_{[p \geq 1/2]} k(p), \mathbf{1}_{[p < 1/2]} k(p)),$$

where the two components denote the harm to the two types. The individual's indirect utility is then

$$U(\theta, h_A, h_B) = v - \mathbf{1}_{[\theta = A]} h_A - \mathbf{1}_{[\theta = B]} h_B.$$

The surplus of a segment with a fraction p of type A individuals is given by

$$V(p) = v - p \mathbf{1}_{[p \geq 1/2]} k(p) - (1 - p) \mathbf{1}_{[p < 1/2]} k(p).$$

Using the symmetry of $k(\cdot)$, we can therefore write

$$V(p) = v - \max\{p, 1 - p\} k(\max\{p, 1 - p\}),$$

from which we can immediately see that $V(p)$ is concave: The individual likes more privacy.

We now apply the decomposition of effects from Proposition 2 to this application. The first term (8), measuring the degree of consumer-market interest alignment, is negative because $h_A(p) - h_B(p)$ is weakly increasing in p , meaning the harms increase as hackers become more confident about user types. This misalignment of interests is a force for individuals to prefer privacy, because more precise targeting of attacks disproportionately harms consumers. The second term (9) is nil as the indirect utility function is linear in the harms h_A and h_B . The third term (10), measuring the curvature of market actions in beliefs, is negative because both h_A and h_B are convex in the posterior beliefs p and individuals prefer lower h_A and h_B . This means that as hackers gain more precise information, the potential harm increases at an increasing rate, an additional force for information disclosure to be harmful to individuals.

5.7. Privacy as a Value in Itself

As discussed above, in some cases, individuals may dislike a lack of privacy per se, because of an “icky” effect. This effect can be seen as the same as the hacking example in the previous section, and we can recast that analysis in terms of the case of privacy as a value in itself.

Suppose that an individual can have two types, A and B , which can be better seen as “behaviors” by the individual that can potentially be observed by “Others.” Others receive a signal of the behavior of the individual that leads to a posterior belief p .

The individual loses utility $k(p)$ if Others’ signal is correct, for a convex symmetric function $k(p)$. This captures the idea that if the signal is not very accurate, it only captures the less important components of the individual’s behavior and is therefore less harmful to the individual. As the signal becomes more accurate, Others obtain increasingly valuable information about the individual’s behavior, which is increasingly harmful to the individual. If the signal is incorrect (for example, the signal states that it is behavior A , whereas it is behavior B), the individual obtains utility $v > 0$.

In terms of the notation above, let a represent the effect on the individual’s utility that the signal received by Others indicates behavior θ , and let k denote the extent to which the individual is harmed. Then, the individual’s indirect utility function is $v - \mathbf{1}[\theta = a] k(p)$. We can then apply the analysis of the previous section to obtain that the individual’s surplus $V(p)$ is concave

in the market beliefs, and therefore the individual prefers more privacy.

This formulation shows that seemingly subjective privacy preferences can be understood as rational responses to utility loss from accurate signal detection. For policymakers, these results imply that privacy regulations should account for both concrete economic harms and the disutility of accurate signal detection, particularly in contexts of widespread surveillance or tracking technologies that generate increasingly accurate signals about individual behaviors.

6. Concluding Remarks

This paper develops a theoretical framework to analyze the effects of privacy through the lens of the concavity or convexity of an individual’s derived payoff as a function of market beliefs about their type. We identify environments in which the individual’s derived payoff function with respect to the firm’s posterior belief is globally concave or convex. Under these conditions, either complete or zero privacy is optimal regardless of the initial prior or the market segmentation.

We further decompose the structure of the payoff function into three fundamental determinants: (i) whether market actions informed by better signals about individuals’ types align or conflict with individual welfare; (ii) the curvature—concave or convex—of individuals’ indirect utility from market actions; and (iii) the curvature of the market’s actions as functions of its beliefs about individuals.

The framework yields concrete implications for both managerial decision making and policy formulation. Managers can leverage this analysis to assess when the collection and use of consumer data generates or erodes economic value. Specifically, in contexts characterized by a concave derived payoff—typical in markets featuring price discrimination—firms should anticipate resistance from consumers and regulators to extensive data gathering. Conversely, when the derived payoff is convex—such as in highly targeted product recommendations—firms can enhance consumer value through increased personalization. For policymakers, the findings underscore the importance of tailoring privacy regulations to specific market contexts rather than adopting uniform privacy standards. For example, stringent privacy protections are desirable in markets like health insurance, yet less critical in markets where consumers benefit substantially from personalization.

Several extensions to the baseline model are promising. First, it would be valuable to study scenarios in which information is asymmetrically distributed across agents in the economy. Investigating whether and why individuals may prefer precise information disclosure to certain firms whereas favoring limited disclosure elsewhere could yield insights into preferences concerning not only personal privacy, but the privacy of others as well.

Second, incorporating selective information sharing into the model would enhance realism, as consumers often wish to control precisely which firms access their data. Allowing for heterogeneous segmentation structures, where firms differ systematically in the information they receive, would offer a richer characterization of privacy preferences.

Third, extending the model to incorporate dynamic considerations—such as the persistent consequences of current information revelation for future interactions or the potential for secondary information use by other parties—could explain the endogenous emergence of privacy norms and preferences.

Finally, the current analysis assumes Bayesian rationality and consistent market beliefs, yet real-world markets frequently exhibit deviations because of algorithmic bias, historical discrimination, or incomplete data collection, particularly across demographic groups. Relaxing this assumption by analyzing markets with systematically mis-specified beliefs represents another valuable direction for future research.

Acknowledgments

This paper benefited from conversations with and comments by Drew Fudenberg, Michael Katz, and Preston McAfee. All errors are our own responsibility.

Appendix A. Consumer Misreporting Types

Consider a model where consumers can misreport their types by incurring a cost and compare it with a benchmark that disallows consumer misreporting types. We show that in both models, in equilibrium, the distribution of a firm's posterior belief of consumer types is a mean-preserving spread of its prior belief.

Consider the case of monopoly price discrimination with linear consumer demand. Suppose that the consumer's type takes one of two values, $\theta \in \{\theta_L, \theta_H\}$, with $\theta_H > \theta_L$. The firm's prior belief is that $\Pr[\theta_H] = f$, so the prior distribution can be denoted as $(1-f, f)$. The firm observes a signal s , which leads to the segmentation $\{(\pi_s, F_s)\}_{s \in S}$. We will consider two generating processes of s : a benchmark case that disallows consumer misreporting and a case that allows consumer misreporting. After observing s , the firm chooses a price p_s for each segment s , and then consumers choose the quantity purchased.

We first consider a benchmark that disallows misreporting. Let the signal s be a noisy version of consumer types,

$$s = \theta + \epsilon,$$

where ϵ is independent of θ and has a normal distribution with zero mean and standard deviation σ . The firm's posterior belief is then

$$\Pr[\theta_H | s] = \frac{f e^{\frac{\theta_H - \theta_L}{2\sigma^2}(2s - \theta_H - \theta_L)}}{f e^{\frac{\theta_H - \theta_L}{2\sigma^2}(2s - \theta_H - \theta_L)} + (1-f)}.$$

Note that the segmentation based on the posterior belief $\{(\pi_s, F_s)\}_{s \in S}$ is a mean-preserving spread of the prior. This

can be shown using the law of total probability³⁰

$$\mathbb{E}_s[\Pr[\theta_H | s]] = f.$$

Let the direct utility function of a consumer with type θ be given by

$$u(\theta, p, q) = (\theta - p)q - q^2/2,$$

where p is the price of the product and q is the quantity purchased. Given the utility function, a consumer with type θ purchases

$$q = \theta - p.$$

Thus, the consumer's indirect utility function, when taking into account their optimal quantities purchased, is given by

$$U(\theta, p) = \frac{(\theta - p)^2}{2},$$

and the monopoly prices under distribution F are given by

$$p^*(F) = \mathbb{E}_F[\theta]/2.$$

The expected surplus for a segment with distribution F is then

$$V(F) = \frac{\mathbb{E}_F[\theta]^2 + 4\text{Var}_F[\theta]}{8}.$$

As we have shown the segmentation based on the signal s is a mean-preserving spread of the segmentation based on the prior, we can apply Proposition 2 and obtain that consumers unambiguously like privacy in this model.³¹

We now turn to a case that allows consumers to distort their types. Specifically, after a consumer's type θ is realized, they can choose to distort their type to $\theta - \Delta$ or $\theta + \Delta$ with a cost c , which we assume to be additive. That is, distorting means that the signal for each realization of ϵ can decrease or increase by Δ . We assume that $0 < \Delta < \theta_H - \theta_L$. Note that in this setting both types want to report that they have a lower type, so, if distorting, any of the types will distort to $\theta - \Delta$. Additionally, note that the higher type θ_H has a higher benefit from distorting, as the consumer's value is convex in θ . Finally, for c close to zero, both types distort, and the firm can fully undo the distorting and that for c large neither type distorts, and we are as in the situation in which distorting is not possible. Therefore, for there to be an effect of distorting we need an intermediate value of c such that type θ_L decides not to distort and type θ_H distorts, and in what follows, we assume that c is in that intermediate region.³²

Let $\tilde{\theta}$ be the so-called reported type after potential distortion. That is, the signal observed by the firm is

$$s = \tilde{\theta} + \epsilon,$$

where ϵ is independent of $\tilde{\theta}$ and has a normal distribution with zero mean and standard deviation σ . Let $\tilde{\theta}_H$ and $\tilde{\theta}_L$ be the types reported by the two types of consumers in the proposed equilibrium, $\tilde{\theta}_H = \theta_H - \Delta$ and $\tilde{\theta}_L = \theta_L$. The firm's posterior belief is then

$$\tilde{\Pr}[\theta_H | s] = \frac{f e^{\frac{\tilde{\theta}_H - \tilde{\theta}_L}{2\sigma^2}(2s - \tilde{\theta}_H - \tilde{\theta}_L)}}{f e^{\frac{\tilde{\theta}_H - \tilde{\theta}_L}{2\sigma^2}(2s - \tilde{\theta}_H - \tilde{\theta}_L)} + (1-f)}.$$

Let $\tilde{F}_s$ denote the posterior distribution, and $\tilde{\pi}_s$ be the probability density of observing signal s , given by

$$\tilde{\pi}_s = f e^{-\frac{(s-\tilde{\theta}_H)^2}{2\sigma^2}} + (1-f) e^{-\frac{(s-\tilde{\theta}_L)^2}{2\sigma^2}},$$

then the segmentation based on the posterior belief $\{(\tilde{\pi}_s, \tilde{F}_s)\}_{s \in S}$ is a mean-preserving spread of the prior, as shown by

$$\mathbb{E}_s[\tilde{\text{Pr}}[\theta_H|s]] = \int_s \tilde{\text{Pr}}[\theta_H|s] \tilde{\pi}_s ds = f.$$

Intuitively, although misreporting distorts the signals, rational firms account for this strategic behavior in equilibrium, resulting in consistent beliefs about the population distribution of types.

As the cost of misreporting is assumed to be additive to the utility of consumption, the consumer demand and the monopoly prices take the same expressions as in the benchmark case. Now we turn to the consumer's type-reporting decision problem. If consumer θ reports $\tilde{\theta} \in \{\theta - \Delta, \theta, \theta + \Delta\}$, their expected utility would be

$$\mathbb{E}_s[U(\theta, p(\tilde{F}_s))] - c \mathbf{1}(\tilde{\theta} \neq \theta), \quad s \sim \mathcal{N}(\tilde{\theta}, \sigma^2).$$

In the expression above, we use the same notation $U(\theta, p)$ as in the benchmark model. As discussed earlier, reporting $\theta + \Delta$ is strictly dominated by reporting $\theta - \Delta$, so we only compare the other two cases.

Recall that c is assumed to have an intermediate value, such that consumer θ_H is better off misreporting their type as $\theta_H - \Delta$. Also, it is too costly for consumer θ_L to misreport their type. This argument then leads to the following proposition.

Proposition A.1. *For an intermediate value of c , there exists a unique equilibrium in the model that allows for consumer misreporting types. In the equilibrium, consumer θ_H misreports their type as $\theta_H - \Delta$, and consumer θ_L truthfully reports their type. The firm's posterior belief is given by*

$$\tilde{\text{Pr}}[\theta_H|s] = \frac{f e^{\frac{\theta_H - \theta_L - \Delta}{2\sigma^2}(2s - \theta_H - \theta_L + \Delta)}}{f e^{\frac{\theta_H - \theta_L - \Delta}{2\sigma^2}(2s - \theta_H - \theta_L + \Delta)} + (1-f)},$$

and the monopoly price is given by

$$p_s^* = [\theta_L + (\theta_H - \theta_L) \tilde{\text{Pr}}[\theta_H|s]]/2.$$

Using the same notation of the consumer surplus V as in the benchmark model, consumer welfare can be written as

$$W(S) = \mathbb{E}_s[V(\tilde{F}_s)] - f c.$$

Note that there is an additional loss $f c$ in consumer welfare, which captures that consumer θ_H misreports their type by incurring a cost c , and the probability of being type θ_H is f . Thus, we obtain that consumers unambiguously like privacy in this model.

It is noteworthy that the mean-preserving spread result requires the assumption that the firm's posterior belief is consistent with the consumers' equilibrium reporting actions. It could be interesting to study models where the firms' beliefs are mis-specified, but that is beyond the scope of this paper.

Appendix B. Proofs

B.1. Differential Approach to General Results

In some applications, a differential approach allows for an easier characterization of the convexity (concavity) of $V(F)$. This approach leverages Gateaux derivatives, which we introduce as follows. For all distributions F, F' , let $\Delta F \equiv F' - F$, and then the Gateaux derivative $\partial_{\Delta F} V(F)$ of V at F in the direction ΔF is defined as

$$\partial_{\Delta F} V(F) \equiv \lim_{t \rightarrow 0} \frac{V(F + t \Delta F) - V(F)}{t} = \frac{d}{dt} V(F + t \Delta F)|_{t=0}.$$

Note that $V(F + t \Delta F)$ here is well defined, as for $t \in [0, 1]$, $F + t \Delta F = (1 - t)F + tF'$ is also a distribution.

The second Gateaux derivative of V at F in the direction ΔF is defined as

$$\partial_{\Delta F}^2 V(F) \equiv \partial_{\Delta F} (\partial_{\Delta F} V(F)) = \frac{d^2}{dt^2} V(F + t \Delta F)|_{t=0}.$$

We can then link the convexity of V with its second Gateaux derivatives using the following lemma.

Lemma B.1 (Convexity and Second Gateaux Derivative). *Suppose that for all F, F' and direction $\Delta F = F' - F$, the second Gateaux derivative $\partial_{\Delta F}^2 V(F)$ exists. Then $V(F)$ is convex (concave) if and only if for all such F and ΔF ,*

$$\partial_{\Delta F}^2 V(F) \geq (\leq) 0. \quad (\text{B.1})$$

Proof. We first show $V(F)$ is convex if and only if $\partial_{\Delta F}^2 V(F) \geq 0$ for all F, F' and $\Delta F = F' - F$.

$\Rightarrow$: Suppose that $V(F)$ is convex. Fix any F, F' and let $\Delta F = F' - F$, then $V(F + t \Delta F)$ is convex in $t \in [0, 1]$, so

$$\partial_{\Delta F}^2 V(F) = \frac{d^2}{dt^2} V(F + t \Delta F)|_{t=0} \geq 0.$$

$\Leftarrow$: Suppose that $\partial_{\Delta F}^2 V(F) \geq 0$ for all F, F' and $\Delta F = F' - F$. Fix any F, F' and let $\Delta F = F' - F$, we now show $V(F + t \Delta F)$ is convex in $t \in [0, 1]$. Note that $\partial_{\Delta F}^2 V(F)$ is homogeneous of degree 2 in ΔF , thus, for all $t' \in (0, 1]$, we have

$$\frac{d^2}{dt^2} V(F + t \Delta F)|_{t=t'} = \frac{1}{t'^2} \partial_{-t' \Delta F}^2 V(F + t' \Delta F) \geq 0.$$

The inequality holds because $F + t' \Delta F$, and $\partial_{-t' \Delta F}^2 V(F + t' \Delta F) \geq 0$ in the direction $-t' \Delta F = F - (F + t' \Delta F)$. Also, for $t' = 0$, we have

$$\frac{d^2}{dt^2} V(F + t \Delta F)|_{t=0} = \partial_{\Delta F}^2 V(F) \geq 0.$$

Thus, $V(F + t \Delta F)$ is convex in $t \in [0, 1]$, which implies that for all $\lambda \in [0, 1]$

$$\lambda V(F) + (1 - \lambda) V(F') \geq V(\lambda F + (1 - \lambda) F').$$

Because F, F' , and λ are arbitrary, this shows that $V(F)$ is convex.

Finally, it is analogous to show $V(F)$ is concave if and only if $\partial_{\Delta F}^2 V(F) \leq 0$ for all F, F' and $\Delta F = F' - F$. $\square$

Assuming appropriate regularity conditions, applying the Gateaux derivative operator to V twice leads to the following differential characterization of convex and concave surplus functions V .³³ The proposition involves the

first and second Gateaux derivatives of a at F in the direction ΔF , which can be defined similarly as above.

Proposition B.1 (Differential Version of V 's Convexity). *Consumer surplus $V(F)$ is convex (concave) if and only if*

$$2\partial_{\Delta F} a(F) \int_{\theta} U_a(\theta, a(F)) d\Delta F(\theta) \quad (\text{B.2})$$

$$+ (\partial_{\Delta F} a(F))^2 \int_{\theta} U_{aa}(\theta, a(F)) dF(\theta) \quad (\text{B.3})$$

$$+ \partial_{\Delta F}^2 a(F) \int_{\theta} U_a(\theta, a(F)) dF(\theta) \geq (\leq) 0, \quad (\text{B.4})$$

for all F, F' and direction $\Delta F = F' - F$.

Proof. Because

$$V(F) = \int_{\theta} U(\theta, a(F)) dF(\theta),$$

applying the Gateaux derivative operator yields

$$\partial_{\Delta F} V(F) = \int_{\theta} U(\theta, a(F)) d\Delta F(\theta) + \partial_{\Delta F} a(F) \int_{\theta} U_a(\theta, a(F)) dF(\theta),$$

And then applying the Gateaux derivative operator again yields

$$\begin{aligned} \partial_{\Delta F}^2 V(F) &= 2\partial_{\Delta F} a(F) \int_{\theta} U_a(\theta, a(F)) d\Delta F(\theta) \\ &\quad + (\partial_{\Delta F} a(F))^2 \int_{\theta} U_{aa}(\theta, a(F)) dF(\theta) \\ &\quad + \partial_{\Delta F}^2 a(F) \int_{\theta} U_a(\theta, a(F)) dF(\theta), \end{aligned}$$

which completes the proof. $\square$

To see that (5), (6), and (7) are special cases under binary types, we can express $\Theta = \{\theta_L, \theta_H\}$. The probability mass distributions over Θ take the form $(1-f, f)$, and thus the only direction of the derivative, up to a scale, is $(-1, 1)$. Then we can evaluate the integrals and obtain (5), (6), and (7).

These expressions are related to (8), (9), and (10), respectively. The first effect in (8) and (19) captures the correlation between a 's derivative and $U_a(\theta, a)$. The second effect in (9) and (B.3) captures the benefits or harms due to the convexity or concavity of U in a . The third effect in (10) and (B.4) captures the convexity or concavity of the market action in F .

Moreover, for exposition purposes, we treat a as one-dimensional in the statement of Proposition B.1. However, it readily extends to multidimensional a , in which case U_a is a vector-valued function, U_{aa} is a matrix-valued function, and the multiplications in (B.2), (B.3), and (B.4) are understood as inner products.

B.2. Proof of Corollary 1 and Discussion

Because U is linear in a ,

$$\lambda U(\theta, a(F')) + (1 - \lambda) U(\theta, a(F'')) - U(\theta, \bar{a}) = 0,$$

so the term (9) is nil. Also, because a is linear in F ,

$$U(\theta, \bar{a}) - U(\theta, a(\bar{F})) = 0,$$

so the term (10) is nil. The inequality is thus simplified to (11).

Now we show that the sign is determined by the correlation between the Riesz representation of a and $U_a(\theta, a)$, with respect to θ . Because a is linear in F , assuming proper continuity conditions,³⁴ there exists a Riesz representation α of a such that

$$a(F) = \int_{\theta} \alpha(\theta) dF(\theta).$$

Applying this and the linearity of U in a to (11) leads to

$$\begin{aligned} &\int_{\theta} [U(\theta, a(F')) - U(\theta, a(F''))][dF'(\theta) - dF''(\theta)] \\ &= \int_{\theta} U_a(\theta, a)(a(F') - a(F''))[dF'(\theta) - dF''(\theta)] \\ &= (a(F') - a(F'')) \int_{\theta} U_a(\theta, a)[dF'(\theta) - dF''(\theta)] \\ &= \int_{\theta} \alpha(\theta)[dF'(\theta) - dF''(\theta)] \int_{\theta} U_a(\theta, a)[dF'(\theta) - dF''(\theta)]. \end{aligned}$$

The first expression $\int_{\theta} \alpha(\theta)[dF'(\theta) - dF''(\theta)]$ captures the average change in $\alpha(\theta)$ when the distribution of θ changes, and the second expression $\int_{\theta} U_a(\theta, a)[dF'(\theta) - dF''(\theta)]$ captures the average change in $U_a(\theta, a)$ when the distribution of θ changes. Thus, their product captures the correlation between $\alpha(\theta)$ and $U_a(\theta, a)$ with respect to θ .

B.3. Proof of Proposition 3

We can expand

$$\begin{aligned} V(F) &= \frac{\mathbb{E}_F[\theta_1]^2 + \mathbb{E}_F[\theta_2]^2 + 2\gamma \mathbb{E}_F[\theta_1]\mathbb{E}_F[\theta_2] + 4(\text{Var}_F[\theta_1] + \text{Var}_F[\theta_2])}{8(1 - \gamma^2)} \\ &= \frac{-3(\mathbb{E}_F[\theta_1]^2 + \mathbb{E}_F[\theta_2]^2) + 2\gamma \mathbb{E}_F[\theta_1]\mathbb{E}_F[\theta_2] + 4(\mathbb{E}_F[\theta_1^2] + \mathbb{E}_F[\theta_2^2])}{8(1 - \gamma^2)}, \end{aligned}$$

and because $\mathbb{E}_F[\theta_1^2] + \mathbb{E}_F[\theta_2^2]$ is linear in F , we obtain

$$\begin{aligned} \partial_{\Delta F}^2 V(F) &= \frac{-3(\mathbb{E}_{\Delta F}[\theta_1]^2 + \mathbb{E}_{\Delta F}[\theta_2]^2) + 2\gamma \mathbb{E}_{\Delta F}[\theta_1]\mathbb{E}_{\Delta F}[\theta_2]}{4(1 - \gamma^2)} \\ &= \frac{1}{4(1 - \gamma^2)} (\mathbb{E}_{\Delta F}[\theta_1] \quad \mathbb{E}_{\Delta F}[\theta_2]) \begin{pmatrix} -3 & \gamma \\ \gamma & -3 \end{pmatrix} \begin{pmatrix} \mathbb{E}_{\Delta F}[\theta_1] \\ \mathbb{E}_{\Delta F}[\theta_2] \end{pmatrix}. \end{aligned}$$

Here we use $\mathbb{E}_{\Delta F}[\theta_i]$ as a shortcut for the following integral:

$$\mathbb{E}_{\Delta F}[\theta_i] \equiv \int_{\theta} \theta_i d\Delta F(\theta).$$

Because $\frac{1}{4(1 - \gamma^2)} \geq 0$ and $\begin{pmatrix} -3 & \gamma \\ \gamma & -3 \end{pmatrix}$ is a negative definite matrix, we have $\partial_{\Delta F}^2 V(F) \leq 0$ and thus V is concave.

B.4. Proof of Proposition 4

We can expand

$$\begin{aligned} V^{\text{NE}}(F) &= \frac{(4-3\gamma^2)(\mathbb{E}_F[\theta_1]^2 + \mathbb{E}_F[\theta_2]^2) + 2\gamma^3 \mathbb{E}_F[\theta_1] \mathbb{E}_F[\theta_2]}{2(4-\gamma^2)^2(1-\gamma^2)} \\ &\quad + \frac{\text{Var}_F[\theta_1] + \text{Var}_F[\theta_2]}{2(1-\gamma^2)} \\ &= \frac{-(\gamma^4 - 5\gamma^2 + 12)(\mathbb{E}_F[\theta_1]^2 + \mathbb{E}_F[\theta_2]^2) + 2\gamma^3 \mathbb{E}_F[\theta_1] \mathbb{E}_F[\theta_2]}{2(4-\gamma^2)^2(1-\gamma^2)} \\ &\quad + \frac{\mathbb{E}_F[\theta_1^2] + \mathbb{E}_F[\theta_2^2]}{2(1-\gamma^2)}. \end{aligned}$$

Because $\mathbb{E}_F[\theta_1^2] + \mathbb{E}_F[\theta_2^2]$ is linear in F , we obtain

$$\begin{aligned} \partial_{\Delta F}^2 V^{\text{NE}}(F) &= \frac{-(\gamma^4 - 5\gamma^2 + 12)(\mathbb{E}_{\Delta F}[\theta_1]^2 + \mathbb{E}_{\Delta F}[\theta_2]^2) + 2\gamma^3 \mathbb{E}_{\Delta F}[\theta_1] \mathbb{E}_{\Delta F}[\theta_2]}{(4-\gamma^2)^2(1-\gamma^2)}. \end{aligned}$$

Because $\gamma^4 - 5\gamma^2 + 12 \geq 8$ for $\gamma \in (-1, 1)$, we have

$$\begin{aligned} \partial_{\Delta F}^2 V^{\text{NE}}(F) &\leq \frac{-8(\mathbb{E}_{\Delta F}[\theta_1]^2 + \mathbb{E}_{\Delta F}[\theta_2]^2) + 2\gamma^3 \mathbb{E}_{\Delta F}[\theta_1] \mathbb{E}_{\Delta F}[\theta_2]}{(4-\gamma^2)^2(1-\gamma^2)} \\ &= \frac{1}{(4-\gamma^2)^2(1-\gamma^2)} (\mathbb{E}_{\Delta F}[\theta_1] \quad \mathbb{E}_{\Delta F}[\theta_2]) \begin{pmatrix} -8 & \gamma^3 \\ \gamma^3 & -8 \end{pmatrix} \begin{pmatrix} \mathbb{E}_{\Delta F}[\theta_1] \\ \mathbb{E}_{\Delta F}[\theta_2] \end{pmatrix}. \end{aligned}$$

Because $\frac{1}{(4-\gamma^2)^2(1-\gamma^2)} \geq 0$ and $\begin{pmatrix} -8 & \gamma^3 \\ \gamma^3 & -8 \end{pmatrix}$ is a negative definite matrix, we have $\partial_{\Delta F}^2 V^{\text{NE}}(F) \leq 0$ and thus V^{NE} is concave.

B.5. Proof of Proposition 6

Expanding (14) leads to

$$V(\beta) = \begin{cases} \frac{1}{8}(1 + 6\beta - 3\beta^2) & \beta \geq \frac{1}{2} \\ \frac{1}{8}(4 - 3\beta^2) & \beta < \frac{1}{2} \end{cases},$$

so $V(\beta)$ is concave on $[0, 1/2]$ and $[1/2, 1]$, respectively.

First, we show the first part of the proposition. The firm's optimal location choice when $\beta \in [0, 1/2]$ is $x^*(\beta) = 0$, whereas the firm's optimal location choice when $\beta \in [1/2, 1]$ is $x^*(\beta) = 1$. Thus, the revelation of information does not modify the firm's location choice if and only if β stays on either $[0, 1/2]$ or $[1/2, 1]$. Because $V(\beta)$ is concave on either interval, in either case, full privacy is optimal for the consumer.

Then we show the second part of the proposition. We provide sufficient conditions under which the consumer benefits from two-signal segmentations. Denote the segmentation as $\{(\pi, \beta_1), (1-\pi, \beta_2)\}$; that is, the first signal indicates that the probability of $\theta = 0$ is β_1 , and the second signal indicates that the probability of $\theta = 0$ is β_2 , and the weights of the two signals are π and $1-\pi$, respectively. The consumer benefits from the two-signal segmentations if and only if

$$V(\pi\beta_1 + (1-\pi)\beta_2) \leq \pi V(\beta_1) + (1-\pi) V(\beta_2). \quad (\text{B.5})$$

The sufficient condition we provide is that $\beta_1 < 1/2 < \beta_2$ and

$$V'(\beta_1) \leq \frac{V(\beta_2) - V(\beta_1)}{\beta_2 - \beta_1} \leq V'(\beta_2). \quad (\text{B.6})$$

To show that these conditions imply (B.5), discuss two cases: (i) $\pi\beta_1 + (1-\pi)\beta_2 \leq 1/2$; and (ii) $\pi\beta_1 + (1-\pi)\beta_2 > 1/2$. In case (i), because V is concave on $[0, 1/2]$, we have

$$\frac{V(\pi\beta_1 + (1-\pi)\beta_2) - V(\beta_1)}{(\pi\beta_1 + (1-\pi)\beta_2) - \beta_1} \leq V'(\beta_1) \leq \frac{V(\beta_2) - V(\beta_1)}{\beta_2 - \beta_1},$$

which is simplified to (B.5). In case (ii), because V is concave on $[1/2, 1]$, we have

$$\frac{V(\beta_2) - V(\pi\beta_1 + (1-\pi)\beta_2)}{\beta_2 - (\pi\beta_1 + (1-\pi)\beta_2)} \geq V'(\beta_2) \geq \frac{V(\beta_2) - V(\beta_1)}{\beta_2 - \beta_1},$$

which is also simplified to (B.5).

Moreover, expanding (B.6) leads to

$$6\beta_1(\beta_1 - \beta_2) \leq 3\beta_1^2 - 3\beta_2^2 + 6\beta_2 - 3 \leq 6(1 - \beta_2)(\beta_2 - \beta_1),$$

that is, we show that (i) $\beta_1 < \frac{1}{2} < \beta_2$ and (ii) $6\beta_1(\beta_1 - \beta_2) \leq 3\beta_1^2 - 3\beta_2^2 + 6\beta_2 - 3 \leq 6(1 - \beta_2)(\beta_2 - \beta_1)$ jointly ensure the consumer benefits from two-signal segmentations.

Appendix C. Extension: Product Targeting with Price Discrimination

In this section, we discuss an extension to the main model in Section 5.2. This extension illustrates that the effect of privacy can be context dependent. That is, the exposure of information may potentially be detrimental or beneficial to consumers, depending on the specific piece of information.

A firm chooses both its location and its price. There is a unit segment and a consumer with a two-dimensional type (θ, η) , where $\theta \in \{0, 1\}$ denotes consumer location and η denotes consumer taste, and they are independent given any signal s .

The consumer's demand function is

$$D(p) = 2 + \eta - |\theta - x| - p,$$

if the product is located at x and the firm charges a price p .

Let $\beta = \Pr[\theta = 0|s]$ denote the posterior probability of location, and G denote the distribution of η after receiving a signal s .³⁵

Analogous to the derivations in the main model, the firm's optimal price is

$$p^*(\beta, G) = (1 + \max\{\beta, 1 - \beta\} + \mathbb{E}_G[\eta])/2.$$

Then, the surplus of a consumer in a segment with proportion β of types $\theta = 0$ and distribution G of η satisfies

$$\begin{aligned} 2V(\beta, G) &= (2 - P(\beta))^2 \max\{\beta, 1 - \beta\} + (1 - P(\beta))^2 \min\{\beta, 1 - \beta\} \\ &\quad + \mathbb{E}_G[\eta^2] - 3\mathbb{E}_G[\eta]^2/4 + \mathbb{E}_G[\eta] (1 + \max\{\beta, 1 - \beta\})/2 \end{aligned}$$

For a fixed β , the only term in $V(\beta, G)$ that is nonlinear in G is $-3\mathbb{E}_G[\eta]^2/4$, which is a concave function of G . This implies that consumers always benefit from privacy in η . This represents a case where the first effect (8) dominates, similar to the price discrimination examples in Section 5.1.

For a fixed G , $V(\beta, G)$ is different from (14) in a constant and a term $\mathbb{E}_G[\eta](1 + \max\{\beta, 1 - \beta\})/2$, which is linear on either $[0, 1/2]$ and $[1/2, 1]$ respectively. This implies that the shape of $V(\beta, G)$ is similar to the shape of $V(\beta)$ in Figure 1, leading to a similar conclusion regarding the privacy of β as in Proposition 6.

In this extension, the consumer may dislike privacy in θ but prefer privacy in η . The intuition is consistent with the main model: when information leads to better product positioning (location), the first effect (8) can be positive and dominate, benefiting consumers. However, when information only enables more effective price discrimination (through η), (8) becomes negative, and consumers prefer privacy.

This multidimensional extension has important implications. Firms should develop data strategies that distinguish between different types of consumer information rather than treating all data uniformly. Location data that enable better product positioning may be worth collecting even under strict privacy regulations, whereas taste data used primarily for price discrimination might face greater consumer and regulatory resistance. For policymakers, this suggests developing privacy frameworks that differentiate between types of data based on their usage rather than simply restricting data collection broadly.

Endnotes

¹ Of course, convexity or concavity does not determine a partition of the space of derived payoff functions, but the mechanisms considered here suggest effects on the evaluation of the benefits of privacy. Indeed, the objective of the paper is to find sufficient conditions under which the value of privacy is unambiguously positive or negative and to highlight the market forces behind these conditions.

² This can happen when the pass-through rates are nonincreasing in the marginal costs, for example, as discussed in Chen and Schwartz (2015).

³ A related issue is that consumers may use a technology, such as ad blocking, to limit the use of targeting. See Chen and Liu (2022), Dukes et al. (2022), and Gritkevich et al. (2022).

⁴ In some cases, after the consumer observes the market action a and her type θ , the consumer can take further actions x . For example, the consumer can choose a quantity to purchase after the firm sets the prices (see an example in Section 5.1). Denote the consumer's direct utility function as $u(\theta, a, x)$, and then the indirect utility function is defined as $U(\theta, a) \equiv \max_x u(\theta, a, x)$. In other words, in such cases, the indirect utility function is the consumer's best-response payoff if the consumer has type θ and the market takes action a .

⁵ We describe the function U as the indirect utility of the individual consumer, but this same analysis also applies if U is the social welfare payoff or the payoff of some other agent in the economy.

⁶ We treat a as one dimensional in the general presentation, but allowing for multidimensional a can be important in applications. For example, a can represent the equilibrium prices for two products in a monopoly setting or a duopoly setting. (See these examples in Section 5.1.)

⁷ Note that this framework applies when firms observe signals for only a subset of consumers. In such cases, consumers for whom no signal is observed can be considered as belonging to a specific segment. This segment can be represented using a specific signal s_0 denoting no observation, with π_{s_0} representing the proportion of consumers for whom no signal is observed. The law of total

probability (2) holds when including this segment, as it accounts for the entire population. In this setting, types for which firms do not receive signals are interpreted as types for which firms receive uninformative signals.

⁸ Note that $V(F_s)$ is defined over cumulative distribution functions rather than individual types. This means that the market does not need to observe signals for all individual types in the entire support.

⁹ In terms of Definition 1 below, this is denoted as $S = \{(1, F_0)\}$.

¹⁰ This is the same issue as when evaluating the expected utility of a risk-averse individual from two lotteries that cannot be ordered in terms of the mean-preserving spread.

¹¹ A function $K(s|s')$ is a transition kernel if and only if $K(s|\cdot)$ is measurable for all $s \in S$ and $K(\cdot|s')$ is a probability measure for all $s' \in S$.

¹² Note that the signal need not be unbiased. The signal can be biased as long as the market knows the information structure and can correctly compute posterior beliefs.

¹³ As another interpretation, the market may observe a signal for only a subset of consumers, which we denote as s_H . For the consumers for whom the market observes no signals, we denote this as s_L .

¹⁴ For example, this can happen when both the high types and the low types send the high and low signals with equal probability. In this case, the market observes the high signal and the low signal with equal probability. As another example, this can happen when the high types send the high signal with probability one, and the low type mimics the high type by also sending the high signal with probability one. In this case, the market observes the high signal and cannot distinguish between the high and low types. We include a more general case of consumer strategic misreporting in Appendix A, showing that the mean-preserving property holds assuming consistent beliefs.

¹⁵ See, for example, Rothschild and Stiglitz (1970).

¹⁶ Note that there may be other definitions of privacy that will not have the garbling property, and the result on the concavity/convexity effects. This Blackwell order is a partial order. However, complete privacy is always a garbling of any (partially) informative signal structure on the consumer's type.

¹⁷ A sufficient condition for this to hold is the Spence–Mirrlees condition, namely $U_{\theta a}(\theta, a)$ is nonnegative for all θ and a .

¹⁸ Information also helps if a is concave in f and U is decreasing in a on average.

¹⁹ Appendix B presents a differential approach for these results, which may be useful in some applications.

²⁰ Mathematically, (11) has the same sign as the correlation between $U_a(\theta, a)$ and the Riesz representation of a , with respect to θ . The Riesz representation α of a is the derivative of a with respect to F , and $\alpha(\theta)$ can be interpreted as the market's action when the probability of type θ is one. We assume proper continuity conditions in a to ensure it exists. See Appendix B.2 for derivations and more discussion.

²¹ This can be micro-founded by assuming consumer's direct utility function $u(\theta, a, q) = \sqrt{2\theta a q - q^2}$, where q represents the consumer's time spent on the platform.

²² In contrast, when there is a negative correlation between consumer marginal utility for market actions and the market's response to consumer types, *ceteris paribus*, we have that consumers benefit from privacy. An example is when a denotes prices and the high type suffers more from higher prices (see Section 5.1).

²³ This can be micro-founded by assuming consumer's direct utility function $u(\theta, a, q) = 2aq - q^2$, where q represents the consumer's time spent on the platform.

²⁴ In contrast, *ceteris paribus*, when the consumer's indirect utility function is concave in the market action; that is, the consumer prefers less variation, we have that the consumer prefers privacy.

²⁵ This can be micro-founded by assuming consumer's direct utility function $u(\theta, a, q) = \sqrt{2aq - q^2}$, where q represents the consumer's time spent on the platform.

²⁶ One interpretation is that the platform receives a fixed commission each time the ad is displayed. An alternative interpretation is that the products are owned by the platform, and upon each impression, consumers will purchase with a fixed probability.

²⁷ See Appendix C for an extension of the model, in which the revelation of the different pieces of information can be welfare-enhancing or welfare-reducing.

²⁸ For a different modeling of the continuum from full privacy to no privacy see, for example, Chen and Iyer (2002).

²⁹ Alternatively, we can also apply Proposition B.1 (presented in Appendix B, the differential approach of Proposition 2) to show the same results. Specifically, for distribution μ and direction $\Delta\mu$, $\partial_{\Delta\mu}^2 V(\mu) = 2\partial_{\Delta\mu} p(\mu) \int_L -u'(w - p(\mu) - \alpha L) d\Delta\mu(L) + (\partial_{\Delta\mu} p(\mu))^2 \int_L u''(w - p(\mu) - \alpha L) d\mu(L)$. Analogous to the reasoning above, regarding the first term, $-u'(w - p(\mu) - \alpha L)$ is a decreasing function of L , and the second term is negative. We also have $\partial_{\Delta\mu} p(\mu) = (1 - \alpha) \int_L L d\Delta\mu(L)$. Note that the integrand L is an increasing function of L . Thus, a sufficient condition for the first term to be negative is that the distributions stochastically dominate one another.

³⁰ Consumer types are binary, so the posterior distribution F_s can be summarized by $\Pr[\theta_H | s]$.

³¹ Moreover, this result can be generalized to compare different signal structures, e.g., when the signals have noises with different standard deviations. Let the $S(\sigma)$ be the segmentation when the standard deviation of the noise is σ , then one can show that $S(\sigma)$ is more private than $S(\sigma')$ if $\sigma \geq \sigma'$. We can apply Proposition 2 and obtain that consumers unambiguously like privacy, that is, a segmentation with larger σ .

³² More details are available upon request. For example, when $\theta_H = 2$, $\theta_L = 1$, $f = 1/2$, $\sigma = 1$, and $\Delta = 1/2$, the intermediate region of c includes $[0.015, 0.035]$.

³³ To simplify notation, we present this proposition with unidimensional market actions a . However, the proposition is readily extended to multidimensional market actions, simply replacing scalar multiplications with dot products and matrix multiplications.

³⁴ One example of such conditions is that a is continuous with respect to the weak-* topology in the following sense. Let $\langle F, g \rangle \equiv \int_{\theta} g(\theta) dF(\theta)$ be a mapping that maps any finite signed measure F and bounded measurable function g to a real number. The weak-* topology is the weakest topology on finite signed measures such that $\langle \cdot, g \rangle \equiv$ is continuous for all g . Theorem IV.1.2 in Schaefer and Wolff (1971) ensures the existence of the Riesz representation of a given that a is continuous with respect to the weak-* topology.

³⁵ We assume that $\mathbb{P}_G(\eta \geq \mathbb{E}_G[\eta]/2) = 1$, that is, η is not too low.

References

- Acquisti A, Varian HR (2005) Conditioning prices on purchase history. *Marketing Sci.* 24(3):367–381.
- Acquisti A, Taylor CR, Wagman L (2016) The economics of privacy. *J. Econom. Literature* 54(2):442–492.
- Adjerid I, Acquisti A, Brandimarte L, Loewenstein G (2013) Sleights of privacy: Framing, disclosures, and the limits of transparency. *Proc. 9th Sympos. Usable Privacy Security* (Association for Computing Machinery, New York), 1–11.
- Amaldoss W, He C (2019) The charm of behavior-based pricing: When consumers' taste is diverse and the consideration set is limited. *J. Marketing Res.* 56(5):767–790.
- Argenziano R, Bonatti A (2023) Data markets with privacy-conscious consumers. *AEA Papers Proc.* 113:191–196.
- Athey S, Gans JS (2010) The impact of targeting technology on advertising markets and media competition. *Amer. Econom. Rev.* 100(2):608–613.
- Athey S, Catalini C, Tucker C (2017) The digital privacy paradox: Small money, small costs, small talk. NBER Working Paper No. 23488, National Bureau of Economic Research, Cambridge, MA.
- Barnes SB (2006) A privacy paradox: Social networking in the United States. *First Monday* 11(9).
- Bergemann D, Brooks B, Morris S (2015) The limits of price discrimination. *Amer. Econom. Rev.* 105(3):921–957.
- Blackwell D (1951) Comparison of experiments. Neyman J, ed. *Proc. 2nd Berkeley Sympos. Math. Statist. Probab.*, vol. 2 (University of California Press, Berkeley), 93–103.
- Calzolari G, Pavan A (2006) On the optimality of privacy in sequential contracting. *J. Econom. Theory* 130(1):168–204.
- Casadesus-Masanell R, Hervas-Drane A (2015) Competing with privacy. *Management Sci.* 61(1):229–246.
- Chen Y, Iyer G (2002) Consumer addressability and customized pricing. *Marketing Sci.* 21(2):197–208.
- Chen Y, Liu Q (2022) Signaling through advertising when an ad can be blocked. *Marketing Sci.* 41(1):166–187.
- Chen Y, Schwartz M (2015) Differential pricing when costs differ: A welfare analysis. *RAND J. Econom.* 46(2):442–460.
- Chen Z, Choe C, Matsushima N (2020) Competitive personalized pricing. *Management Sci.* 66(9):4003–4023.
- Chen X, Simchi-Levi D, Wang Y (2022) Privacy-preserving dynamic personalized pricing with demand learning. *Management Sci.* 68(7):4878–4898.
- Choi WJ, Jerath K, Sarvary M (2023) Consumer privacy choices and (un)targeted advertising along the purchase journey. *J. Marketing Res.* 60(5):889–907.
- Conitzer V, Taylor CR, Wagman L (2012) Hide and seek: Costly consumer privacy in a market with repeat purchases. *Marketing Sci.* 31(2):277–292.
- de Cornière A, de Nijs R (2016) Online advertising and privacy. *RAND J. Econom.* 47(1):48–72.
- Dukes A, Liu Q, Shuai J (2022) Skippable ads: Interactive advertising on digital media platforms. *Marketing Sci.* 41(3):528–547.
- Eilat R, Eliaz K, Mu X (2021) Bayesian privacy. *Theoret. Econom.* 16(4):1557–1603.
- Elliott M, Galeotti A, Koh A, Li W (2022) Market segmentation through information. Discussion paper, Cambridge University, Cambridge, UK.
- Evans D (2009) The online advertising industry: Economics, evolution, and privacy. *J. Econom. Perspectives* 23(3):37–60.
- Fainmesser IP, Galeotti A, Momot R (2023) Digital privacy. *Management Sci.* 69(6):3157–3173.
- Farboodi M, Haghighpanah N, Shourideh A (2025) Good data and bad data: The welfare effects of price discrimination. Discussion paper, Massachusetts Institute of Technology, Cambridge, MA.
- Farrell J (2012) Can privacy be just another good? *J. Telecomm. High Tech. Law* 10(2):251–264.
- Fudenberg D, Tirole J (2000) Customer poaching and brand switching. *RAND J. Econom.* 31(4):634–657.
- Gal-Or E (1985) Information sharing in oligopoly. *Econometrica* 53(2):329–343.
- Goh K-Y, Hui K-L, Png I (2015) Privacy and marketing externalities: Evidence from do not call. *Management Sci.* 61(12):2982–3000.
- Gritkevich A, Katona Z, Sarvary M (2022) Ad blocking. *Marketing Sci.* 68(6):4703–4724.
- Gross R, Acquisti A (2005) Information revelation and privacy in online social networks. Sabrina De Capitani DV, Dingledine R, eds. *Proc. ACM Workshop Privacy Electronic Soc.* (Association for Computing Machinery Location of publisher, New York), 71–80.

- Hajihashemi B, Sayedi A, Shulman JD (2022) The perils of personalized pricing with network effects. *Marketing Sci.* 41(3):477–500.
- Hann I-H, Hui K-L, Lee S-YT, Png IP (2008) Consumer privacy and marketing avoidance: A static model. *Management Sci.* 54(6):1094–1103.
- Hermalin BE, Katz ML (2006) Privacy, property rights and efficiency: The economics of privacy as secrecy. *Quant. Marketing Econom.* 4(3):209–239.
- Hirshleifer J (1971) The private and social value of information and the reward to inventive activity. *Amer. Econom. Rev.* 61(4):561–574.
- Hui K-L, Png I (2006) The economics of privacy. Hendershott T, ed. *Handbooks in Information Systems*, vol. 1 (Emerald, Bingley, UK), 471–497.
- Ichihashi S (2020) Online privacy and information disclosure by consumers. *Amer. Econom. Rev.* 110(2):569–595.
- Johnson GA, Shriver SK, Du S (2020) Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Sci.* 39(1):33–51.
- Johnson GA, Shriver SK, Goldberg SG (2023) Privacy and market concentration: Intended and unintended consequences of the GDPR. *Management Sci.* 69(10):5695–5721.
- Jullien B, Lefouli Y, Riordan M (2020) Privacy protection, security, and consumer retention. Discussion paper, Columbia University and TSE, New York.
- Kamenica E, Gentzkow M (2011) Bayesian persuasion. *Amer. Econom. Rev.* 101(6):2590–2615.
- Kummer M, Schulte P (2019) When private information settles the bill: Money and privacy in Google's market for smartphone applications. *Management Sci.* 65(8):3470–3494.
- Lei Y, Miao S, Momot R (2024) Privacy-preserving personalized revenue management. *Management Sci.* 70(7):4875–4892.
- Li KJ, Jain S (2016) Behavior-based pricing: An analysis of the impact of peer-induced fairness. *Management Sci.* 62(9):2705–2721.
- Lin S (2020) Two-sided price discrimination by media platforms. *Marketing Sci.* 39(2):317–338.
- Miklós-Thal J, Goldfarb A, Haviv AM, Tucker C (2024) Frontiers: Digital hermits. *Marketing Sci.* 43(4):697–708.
- Montes R, Sand-Zantman W, Valletti T (2019) The value of personal information in online markets with endogenous privacy. *Management Sci.* 65(3):1342–1362.
- Mussa M, Rosen S (1978) Monopoly and product quality. *J. Econom. Theory* 18(2):301–317.
- Ning ZE (2021) List price and discount in a stochastic selling process. *Marketing Sci.* 40(2):366–387.
- Oery A (2016) Consumers on a leash: Advertised sales and intertemporal price discrimination. Discussion paper, Cowles Foundation, Yale University, New Haven, CT.
- Posner RA (1978) Economic theory of privacy. *Regulation* 2(3):19–26.
- Rafieian O, Yoganarasimhan H (2021) Targeting and privacy in mobile advertising. *Marketing Sci.* 40(2):193–218.
- Raith M (1996) A general model of information sharing in oligopoly. *J. Econom. Theory* 71(1):260–288.
- Rhee K-E, Thomadsen R (2017) Behavior-based pricing in vertically differentiated industries. *Management Sci.* 63(8):2729–2740.
- Rhodes A, Zhou J (2024) Personalized pricing and competition. *Amer. Econom. Rev.* 114(7):2141–2170.
- Robinson J (1933) *The Economics of Imperfect Competition* (Macmillan, London).
- Rothschild M, Stiglitz JE (1970) Increasing risk: I. A definition. *J. Econom. Theory* 2(3):225–243.
- Schaefer HH, Wolff MP (1971) *Topological Vector Spaces* (Springer New York, New York).
- Schmalensee R (1981) Output and welfare implications of monopolistic third-degree price discrimination. *Amer. Econom. Rev.* 71(1):242–247.
- Schoeman F (1992) *Privacy and Social Freedom* (Cambridge University Press, Cambridge, UK).
- Shaffer G, Zhang ZJ (1995) Competitive coupon targeting. *Marketing Sci.* 14(4):395–416.
- Shapiro C (1986) Exchange of cost information in oligopoly. *Rev. Econom. Stud.* 53(3):433–446.
- Shen Q, Villas-Boas JJ (2018) Behavior-based advertising. *Management Sci.* 64(5):2047–2064.
- Shin J, Yu J (2021) Targeted advertising and consumer inference. *Marketing Sci.* 40(5):900–922.
- Stigler GJ (1980) An introduction to privacy in economics and politics. *J. Legal Stud.* 9(4):623–644.
- Taylor CR (2003) Supplier surfing: Competition and consumer behavior in subscription markets. *RAND J. Econom.* 34(2):223–246.
- Taylor CR (2004) Consumer privacy and the market for customer information. *RAND J. Econom.* 35(4):631–650.
- Thisse J, Vives X (1988) On the strategic choice of spatial price policy. *Amer. Econom. Rev.* 78(1):122–137.
- Villas-Boas J (1999) Dynamic competition with customer recognition. *RAND J. Econom.* 30(4):604–631.
- Villas-Boas J (2004) Price cycles in markets with customer recognition. *RAND J. Econom.* 35(3):486–501.
- Villas-Boas J, Yao Y (2021) A dynamic model of optimal retargeting. *Marketing Sci.* 40(3):428–458.
- Vives X (1984) Duopoly information equilibrium: Cournot and Bertrand. *J. Econom. Theory* 34(1):71–94.
- Vives X (1988) Aggregation of information in large Cournot markets. *Econometrica* 56(4):851–876.
- Warren S, Brandeis L (1890) The right to privacy. *Harvard Law Rev.* 4(5):193–203.
- Westin A (1967) *Privacy and Freedom* (Atheneum Publishers, New York).
- Yang KH (2022) Selling consumer data for profit: Optimal market-segmentation design and its consequences. *Amer. Econom. Rev.* 112(4):1364–1393.
- Zhang J (2011) The perils of behavior-based personalization. *Management Sci.* 30(1):170–186.
- Zhu R, Moorthy S, Shi X (2025) Advertising platforms and privacy. *Marketing Sci.*, ePub ahead of print May 9, <https://doi.org/10.1287/mksc.2024.0879>.