A Theory of the Effects of Privacy*

Alessandro Bonatti† Yunhao Huang‡ J. Miguel Villas-Boas§

August, 2023

Abstract

The development of information technologies has heightened the debate over the potential loss of individual privacy and raised the importance of consumer rights to privacy. However, some of these technologies allow for substantial customization of communications and offerings to consumers. We present a theory of privacy based on the concavity of the indirect utility function with respect to market beliefs: concavity of the indirect utility function makes privacy valuable to individuals. We provide conditions on market settings that generate concavity (or convexity) of the indirect utility function. We apply the framework to product choice, price discrimination, hacking, and health insurance.

*Bonatti acknowledges financial support through NSF Grant SES-1948692. This paper benefited from conversations with and comments by Drew Fudenberg, Michael Katz, and Preston McAfee. All errors are our own responsibility.

†Sloan School of Management, MIT, bonatti@mit.edu.

‡University of California, Berkeley, huangyh@haas.berkeley.edu.

§University of California, Berkeley, villas@haas.berkeley.edu.

1 Introduction

The development of information technologies has introduced a vibrant debate over the loss of privacy by individual citizens. Currently available technologies allow the tracking of citizens' search for information, communication, movement, activities, and products or services bought. The potential for tracking means a loss of privacy for citizens, and it is important to understand conditions under which this loss of privacy can be harmful or beneficial to individuals and limit tracking when harmful.

The main idea of this paper is that a loss of privacy means more precise posterior beliefs for observers of the information about an individual's preferences or behavior and that if the derived payoff of the individual is concave (convex) on the posterior beliefs by the observers of the information, then a loss of privacy harms (benefits) the individual. A loss of privacy leads to a mean-preserving spread of the posterior beliefs about the individual's type or behavior, and therefore we simply need to apply Jensen's inequality to establish a relation between the concavity or convexity of the derived payoff function and the relative costs or benefits of the loss of privacy. The framework presented here is a unifying theory of privacy, which has thus far proven elusive. Indeed, several existing results on the effect of privacy in different situations can be viewed from this perspective.

The paper also derives conditions on the primitive payoff functions of individuals and how observers of the individual's type of behavior react to the individual's information such that the derived payoff function of the individual is concave or convex. These conditions depend on three effects that are ultimately quite intuitive. 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 primitive payoff function of the individuals with respect to the action of the observers of the information may be convex or concave, which has implications for the derived payoff function in terms of the posterior beliefs. Third, the action of the observers of information may react to their posterior beliefs in a concave or convex way.[1]

We consider the application of this framework to different settings including various forms of targeting effects, hacking, health insurance, and the case of privacy being a value in itself.

One interesting benchmark that we consider to formalize the general effects above is the case of a single decision-maker under uncertainty. In that case, more information about the

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

realization of uncertainty, which can be seen as a loss of privacy, is beneficial to the decision-maker, and thus, as is well-known, the derived payoff function of the individual is convex in the posterior beliefs. In terms of the problem of privacy, if the observer of information acts in the best interest of the individual, a loss of privacy is then beneficial to the individual.

As discussed in Acquisti, Taylor, and Wagman (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, Eliaz, and Mu (2021) measure privacy as the Kullback-Liebler 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, while most individuals state a preference for privacy, they require relatively small incentives to forgo some dimensions of their privacy, e.g., Gross and Acquisti (2005), Barnes (2006), Adjerid, Acquisti, Brandimarte, and Loewenstein (2013), and Athey, Catalini, and Tucker (2017).

There has been significant work on different dimensions of privacy with a focus on specific market applications. See Hui and Png (2006) and Acquisti, Taylor, and Wagman (2016) for surveys on the economics of privacy. Work on privacy has examined several issues including the following; contracting, e.g., Taylor (2004), Hermalin and Katz (2006), and Calzolari and Pavan (2006); markets with repeat purchases, e.g., Villas-Boas (1999), Fudenberg and Tirole (2000), Taylor (2003), Villas-Boas (2004), Acquisti and Varian (2005), Zhang (2011), and Conitzer, Taylor, and Wagman (2012); behavior-based and targeted advertising, e.g., Evans

(2009), Goldfarb and Tucker (2011), Goh, Hui, and Png (2015), de Cornière and de Nijs (2016), Shen and Villas-Boas (2018), and Villas-Boas and Yao (2021); and personalized pricing with product recommendations, e.g., Ichihashi (2020).

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, e.g., Calzolari and Pavan (2006), Montes, Sand-Zantman, and Valletti (2019), Choi, Jerath, and Sarvary (2020), Jullien, Lefouli, and Riordan (2020), and Argenziano and Bonatti (2023).[2]

Another strand of literature addresses how the presence of more or less information in a market affects the market interaction between different economic agents, e.g., Vives (1984, 1988), Gal-Or (1985), Shapiro (1986), and Raith (1996). 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 payoff functions and the market actions as a function of the information in the market. In this vein, the market segmentation analysis in Bergemann, Brooks, and Morris (2015), Yang (2022), and Elliott, Galeotti, Koh, and Li (2022) assesses the impact of privacy in monopolistic and oligopolistic settings. Finally, Rhodes and Zhou (2022) explore fully personalized vs. 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. 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 presents a motivating example and Section 3 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] 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, Liu, and Shuai (2022), and Gritckevich, Katona, and Sarvary (2022).

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. Consider the following motivating example, 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 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 receives an informative signal s of the type of each individual. Let $\mu(s) \triangleq \mathbb{E}[\theta \mid s]$ denote the mean of the firms' posterior distribution induced by signal s.

Assume that all consumers participate in the competitive insurance market, i.e., there are no selection effects ($\underline{\theta}u(W - L) + (1 - \underline{\theta})u(W) < u(W - \bar{\theta}L)$). Every individual who generates signal s pays a competitive insurance premium of

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

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

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

$$V = \mathbb{E}_s \left[u(w - \mu(s)L) \right]. \tag{1}$$

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

$$V_0 = u(w - \mu_0 L).$$

Finally, because $\mathbb{E}_s \left[\mu(s) \right] = \mu_0$ and $u(\cdot)$ is strictly concave, it follows that $V < V_0$. Therefore, 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*.

How do these conclusions depend on the assumptions of the model, such as the binary distribution of health expenditures for each consumer, competitive insurance markets, full market coverage, and no moral hazard? In the rest of the paper, we will derive these conclusions in a more general model, and we will explore the robustness of the above conclusion.

3 The 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 utility if she has type θ and the market takes action a. This utility will vary with the application considered and is discussed in the next section.[3]

The commonly known prior distribution of the consumer's type is given by

$$F_0(\theta).$$

The market receives an informative signal $s \in S$ of the consumer's type. 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 assumed to be measurable with respect to its beliefs about consumer type, i.e., the market can tailor its actions to each consumer segment.[4]

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

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

A segmentation is a mixture with weights π_s over distributions F_s that jointly satisfy

$$\int_s F_s(\theta)\pi_s \mathrm{d}s = F_0(\theta), \quad \forall \theta \in \Theta.$$

In a learning interpretation, see 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. 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 $\mathcal{S}$, each market segment s obtains an expected surplus of

$$V(F_s) \triangleq \int_\theta U(\theta, a(F_s)) \mathrm{d}F_s(\theta).$$

[3]We describe the function U as the 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.

[4]We treat a as one-dimensional in the general presentation, but a having more than one dimension can be important in several applications, such as those in Section 5.2.

The expected consumer surplus is then given by

$$W(\mathcal{S}) \triangleq \mathbb{E}_s[V(F_s)] = \int_s V(F_s)\pi_s \mathrm{d}s.$$

In the special case of *full privacy*, $\mathcal{S} = \varnothing$, the market takes a single action $a(F_0)$, and the consumer's total surplus is given by

$$W(\varnothing) \triangleq V(F_0) = \int_\theta U(\theta, a(F_0))\mathrm{d}F_0(\theta).$$

We are interested in whether a loss of privacy—a more informative segmentation $\mathcal{S}$ in the Blackwell order—is beneficial or harmful to consumer surplus. To abstract from unraveling issues, 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. In the population interpretation of the model, we assess the value of privacy from the perspective of the entire consumer population. Thus, 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 provides 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.

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, Lefouli, and Riordan, 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 her 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, Sand-Zantman, and Valletti, 2019). In that case, $a(F_s)$ could represent the equilibrium action of the market if the individual did not incur costs to preserve her 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.[5]

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.

Definition 1 (Privacy and the Blackwell Order)
Segmentation $\mathcal{S} = \{(\pi_s, F_s)\}_{s\in S}$ is more private than $\mathcal{S}' = \{(\pi'_s, F'_s)\}_{s\in S}$ if and only if $\mathcal{S}' \succ \mathcal{S}$ in the Blackwell order, i.e., if and only $\mathcal{S}'$ is a mean-preserving spread of $\mathcal{S}$.

Consider the case of full privacy vs. an informative segmentation. If the consumer retains her privacy, she obtains $W(\varnothing)$. If the individual consumer loses her privacy, her expected surplus is given by $W(\mathcal{S})$. Given that the distribution of the posteriors $\{F_s(\cdot)\}_{s\in S}$ is a mean-preserving spread of $F_0(\cdot)$, a simple application of Jensen's inequality establishes that $\mathbb{E}_s V(F_s(\cdot)) \geq (\leq)\ V(F_0)$ if $V(F)$ is a convex (concave) function of the belief distribution F.

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

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. However, when is V concave or convex? In the remainder of this section, let us discuss properties of the functions $U(\theta, a)$ and of the market's action $a(F)$ that determine the shape of V.

[5] Note that 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.

4.1 Binary Types

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 $f \triangleq \Pr[\theta_H]$, and compute the total 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

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

Consumer–Market Alignment As a warmup 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:

$$a(f) = \arg\max_a \left[(1-f)U(\theta_L, a) + fU(\theta_H, a)\right] \text{ for all } f \in [0,1]. \tag{2}$$

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 (2) is linear in f, and therefore, the maximum over a for each f 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 (2) 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

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

This is just the intuitive result that if the high type gains by increasing a more than the low type, then a market that acts in the consumer's interest will increase the action a when it assigns a higher probability to type θ_H.

General Market Actions Without knowledge of the market's objective function, our approach 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

$$
\begin{aligned}
V''(f) &= 2a'(f)[U_a(\theta_H, a(f)) - U_a(\theta_L, a(f))] && (3)\\
&\quad + a'(f)^2[(1-f)U_{aa}(\theta_L, a(f)) + fU_{aa}(\theta_H, a(f))] && (4)\\
&\quad + a''(f)[(1-f)U_a(\theta_L, a(f)) + fU_a(\theta_H, a(f))]. && (5)
\end{aligned}
$$

This expression highlights three effects of information: (i) assuming a is increasing, information helps the consumer if higher types also have a higher marginal utility of a;[6] (ii) information helps if U is convex in a on average; and (iii) information helps if a is convex in f and U is increasing in a on average.[7]

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

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

$$\int_\theta \lambda \left[U(\theta, a(F')) - U(\theta, a(\bar{F}))\right] \mathrm{d}F'(\theta) + \int_\theta (1-\lambda) \left[U(\theta, a(F'')) - U(\theta, a(\bar{F}))\right] \mathrm{d}F''(\theta) \geq 0.$$

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

[6] 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.

[7] Information also helps if a is concave in f and U is decreasing in a on average.

Proposition 2 (Convexity of V)
Consumer surplus $V(F)$ is convex (concave) if

$$\lambda(1-\lambda)\int_\theta [U(\theta,a(F'))-U(\theta,a(F''))]\,[\mathrm{d}F'(\theta)-\mathrm{d}F''(\theta)] \tag{6}$$

$$+\int_\theta [\lambda U(\theta,a(F'))+(1-\lambda)U(\theta,a(F''))-U(\theta,\bar{a})]\,\mathrm{d}\bar{F}(\theta) \tag{7}$$

$$+\int_\theta \left[U(\theta,\bar{a})-U(\theta,a(\bar{F}))\right]\mathrm{d}\bar{F}(\theta)\geq(\leq)\ 0, \tag{8}$$

for all F', F'', λ, where

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

To gain some intuition for Proposition 2, we consider particular cases. For instance, suppose that the market action is linear in F (e.g., because $a(F)=\mathbb{E}_F[\theta]$) and that U is linear in a (e.g., because the consumer has unit demand and a is a price variable). We then obtain the following simpler conditions.

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

$$\int_\theta [U(\theta,a(F'))-U(\theta,a(F''))]\,[\mathrm{d}F'(\theta)-\mathrm{d}F''(\theta)]\geq 0, \tag{9}$$

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

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 the Riesz representation of a and $U_a(\theta,a)$, with respect to θ.[8] If the high types of consumer dislike increases in a more than low types and the market chooses a greater a when it has a higher belief distribution, the consumer prefers more privacy. Conversely, V is convex if the types that prefer action $a(F')$ are also more likely under distribution F', i.e., if the interests of the market and the agent are sufficiently aligned.

For more general functions U and a, expressions (6)-(8) uncover the further effects of a nonlinear U and nonlinear a, respectively. In particular, the term in (7) captures the effect that with more information, market actions are more varied, and that can be harmful or beneficial to consumers depending on whether the function $U(\theta,a)$ is concave or convex in

[8]The Riesz representation of a can be seen as the derivative of a with respect to F, and we assume proper continuity conditions in a to ensure it exists. See Appendix A.1 for derivations and more discussion.

a. Finally, the term in (8) captures the effect of whether the market action is convex or concave in the market's beliefs.

Therefore, if U is convex in the market action, the consumer dislikes privacy *ceteris paribus*; if the consumer on average likes a higher market action, then convexity of $a(\cdot)$ also induces the consumer to dislike privacy.

For example, suppose that $U(\theta, a)$ is linear in a, that $a(F)$ is not linear in F, and that the market chooses a lower a than the consumer would like. Then, the convexity of $a(F)$ would be a force for the consumer to prefer less privacy. Consider the following particular case: suppose that $a(F)$ is linear in F and that $U(\theta, a)$ is convex in a. Then, the convexity of U is a force for the consumer to prefer less privacy.

4.3 Differential Approach

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 $\Delta F \triangleq F' - F$, then the Gateaux derivative $\partial_{\Delta F} V(F)$ of V at F in the direction ΔF is defined as

$$\partial_{\Delta F} V(F) \triangleq \lim_{t \to 0} \frac{V(F + t\Delta F) - V(F)}{t} = \frac{\mathrm{d}}{\mathrm{d}t} V(F + t\Delta F)\bigg|_{t=0}.$$

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

$$\partial^2_{\Delta F} V(F) \triangleq \partial_{\Delta F} \left(\partial_{\Delta F} V(F)\right) = \frac{\mathrm{d}^2}{\mathrm{d}t^2} V(F + t\Delta F)\bigg|_{t=0}.$$

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

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

$$\partial^2_{\Delta F} V(F) \geq (\leq)\ 0. \tag{10}$$

Applying the Gateaux derivative operator to V twice leads to the following differential characterization of convex and concave surplus functions V.[9] 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 3 (Differential Version of V's Convexity)
Consumer surplus $V(F)$ is convex (concave) if

$$2\partial_{\Delta F}\, a(F) \int_\theta U_a(\theta, a(F))\mathrm{d}\Delta F(\theta) \tag{11}$$

$$+\,(\partial_{\Delta F}\, a(F))^2 \int_\theta U_{aa}(\theta, a(F))\mathrm{d}F(\theta) \tag{12}$$

$$+\,\partial^2_{\Delta F}\, a(F) \int_\theta U_a(\theta, a(F))\mathrm{d}F(\theta) \geq (\leq)\ 0, \tag{13}$$

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

To see that (3), (4), and (5) 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 (3), (4), and (5).

These expressions are related to (6), (7), and (8), respectively. The first effect in (6) and (11) captures the correlation between a's derivative and $U_a(\theta, a)$. The second effect in (7) and (12) captures the benefits or harms due to the convexity or concavity of U in a. The third effect in (8) and (13) captures the convexity or concavity of the market reaction in F.

5 Applications

We now consider several applications of the general framework we presented, beginning with a richer version of our health insurance example.

5.1 Health Insurance Revisited

An important assumption in the motivating 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, as a reduced form for a model with partial contractibility

[9] To simply 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.

(e.g., moral hazard), that 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.

Binary Health Expenditures In our binary-expenditure example, the expected utility 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

$$V(\mu) = \mu\ u(w - (1-\alpha)\mu L - \alpha L) + (1-\mu)\ u(w - (1-\alpha)\mu L). \tag{14}$$

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 our differential approach of (3)-(5) and obtain

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

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 second term is negative because giving firms greater information will introduce variability in the insurance premium faced by each consumer. The first term is negative because, with informed insurance companies, the premium will on average increase more for consumers who are in fact more likely to incur the loss L. These consumers have a larger marginal utility of income.

This allows us to reach an intuitive but important conclusion, namely that consumers prefer full privacy in this model. In other words, consumer surplus is maximized in this example by letting the insurance market operate under the prior distribution. The above conclusion, however, relies on the binary nature of the uncertainty facing each consumer type, i.e., the health expenditure level $\{0, L\}$, as we show below.

General Health Expenditures Now, assume that the health expenditures L of each consumer type θ follow a more general distribution with *cdf* $G(L;\theta)$. 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, i.e.,

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

where

$$\mu_s(L) = \int_\theta G(L;\theta)\mathrm{d}F_s(\theta).$$

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

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

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 $\bar{\mu} = \lambda\mu' + (1-\lambda)\mu''$ denote the convex combination of the two distributions with weights λ and $1-\lambda$, respectively. Finally, we let $\bar{L}'$, $\bar{L}''$ and $\bar{L} = \lambda\bar{L}' + (1-\lambda)\bar{L}''$ 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

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

and simplify it as follows:

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

Note that the integrand of the second term is negative for all L because $u(\cdot)$ is strictly concave. Regarding the first term, the difference

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

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''.

Alternatively, we can also apply Proposition 3 to show the same result. For distribution μ and direction $\Delta\mu$,

$$\partial^2_{\Delta\mu} V(\mu) = 2\partial_{\Delta\mu} p(\mu) \int_L -u'(w - p(\mu) - \alpha L)\mathrm{d}\Delta\mu(L) + (\partial_{\Delta\mu} p(\mu))^2 \int_L u''(w - p(\mu) - \alpha L)\mathrm{d}\mu(L).$$

Analogous to the reasoning above, the second term is negative, and regarding the first term, $-u'(w - p(\mu) - \alpha L)$ is a decreasing function of L. We also have

$$\partial_{\Delta\mu} p(\mu) = (1-\alpha) \int_L L \mathrm{d}\Delta\mu(L).$$

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.

Proposition 4 (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. More generally, this need not hold, as we show 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*

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

5.2 Price Discrimination with Linear Demand

We now study price discrimination with linear consumer demand, and we apply our results in Proposition 3 to a setting with multidimensional consumer types and market actions.

For this application, consider a consumer with a two-dimensional type, and let the 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.[10] Given the utility function, the consumer purchases the following quantities:

$$q_i = \frac{\theta_i - p_i + \gamma(\theta_j - p_j)}{1 - \gamma^2}, \quad i = 1, 2. \tag{15}$$

The consumer's indirect utility function 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 (15) 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:

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

Because the market actions (p_1, p_2) are linear in distribution F, we can then apply Proposition 3 (which readily extends to multidimensional a) and obtain the following proposition.

[10]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.

Proposition 5 (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).

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 (15) are given by

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

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

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

Proposition 6 (Duopoly Price Discrimination)
In the duopoly model with linear demand, consumers unambiguously like privacy.

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

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

Consumer type θ's demand is then given by

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

A monopolist seller can produce quality y at a quadratic cost

$$c(y) = c\,y^2/2, \quad c > 1/2.$$

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

$$p^* = \frac{\mathbb{E}_F[\theta]\,c}{2c-1} \quad \text{and} \quad y^* = \frac{\mathbb{E}_F[\theta]}{2c-1}.$$

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

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

where

$$\beta \triangleq \frac{1-c}{2c-1}.$$

Note that the first term (the expectation of the random variable θ^2) is by definition linear in the probability distribution F. The second term (the square of the expectation of θ) is convex in F. Therefore, the sign of β determines the convexity/concavity of V.

Proposition 7 (Quality and Price Discrimination)
In the model with quality and price discrimination, consumers like privacy if and only if the cost of investing in quality satisfies $\beta < 0$, i.e., $c > 1$.

Intuitively, in this example, the first term (6) dominates: θ determines whether the consumer likes higher or lower actions. In any case, the market increases its action linearly in F (so the third term (8) is nil), and the utility function is convex. One could have imagined that convexity is a force that drives the value of privacy down, but in this quadratic model, this is insufficient to overcome the first term.

5.3 Targeting Captive Consumers

Several papers, beginning with the classic work of Thisse and Vives (1988), recently generalized by Rhodes and Zhou (2022) have analyzed the impact of demand information on competitive price targeting. Here, we develop a simpler version of the exogenous consideration sets model of Armstrong and Vickers (2019) and Narasimhan (1988), which follows the exposition in Shi and Zhang (2022).

There are three types of consumers: captives of firm 1, captives of firm 2, and shoppers. All consumers have unit demand and a willingness to pay normalized to one for either good.

Therefore, we summarize the distribution of consumer types by (q_1, q_2), which denotes the fractions of captive consumers for each firm. Shoppers are of measure $1 - q_1 - q_2 > 0$.

In terms of our earlier exposition, the market action is multidimensional and generically stochastic. Indeed, for all nontrivial segmentations, the equilibrium is in mixed strategies: it consists of two price distributions with the cumulative distribution denoted by G_i for each seller. In particular, without loss of generality, letting $q_1 \geq q_2$, the equilibrium of the pricing game is given by

$$G(p_1) = \frac{1-q_1}{1-q_1-q_2}\left(1 - \frac{q_1}{1-q_2}\frac{1}{p_1}\right), \quad p_1 \in \left[\frac{q_1}{1-q_2}, 1\right]$$

and

$$G_2(p_2) = \frac{1-q_2}{1-q_1-q_2}\left(1 - \frac{q_1}{1-q_2}\frac{1}{p_2}\right), \quad p_2 \in \left[\frac{q_1}{1-q_2}, 1\right].$$

Producer surplus is then given by

$$\Pi(q_1, q_2) = q_1 \mathbb{E}_{G_1}[p_1] + q_2 \mathbb{E}_{G_2}[p_2] + (1-q_1-q_2)\mathbb{E}_{G_1,G_2}[\min\{p_1, p_2\}],$$

and because consumers are homogeneous (in terms of valuations) and they buy with probability 1, we also have consumer surplus as

$$V(q_1, q_2) = 1 - \Pi(q_1, q_2).$$

Finally, for this setting, Shi and Zhang (2022) show that consumer surplus is reduced in any symmetric segmentation. That is, while the welfare results for arbitrary segmentations (q_1, q_2) are ambiguous, they show that consumer surplus with symmetric shares $V(q, q)$ is a concave function of q.

5.4 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 0 or 1. The consumer's location is her 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 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 $s \in \{0, 1\}$ of the location of the consumer, summarized by the posterior probability $\beta = \Pr[\theta = 0 \mid s]$.

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 be the location that the firm chooses after receiving a signal s. The firm therefore solves the following problem:

$$\max_{P,x} P\left[\beta(2-x-P)+(1-\beta)(1+x-P)\right]. \tag{16}$$

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

$$2V(\beta) = (2-P^*(\beta))^2 \max\{\beta, 1-\beta\} + (1-P^*(\beta))^2 \min\{\beta, 1-\beta\}. \tag{17}$$

Figure 1 illustrates our result for this example.

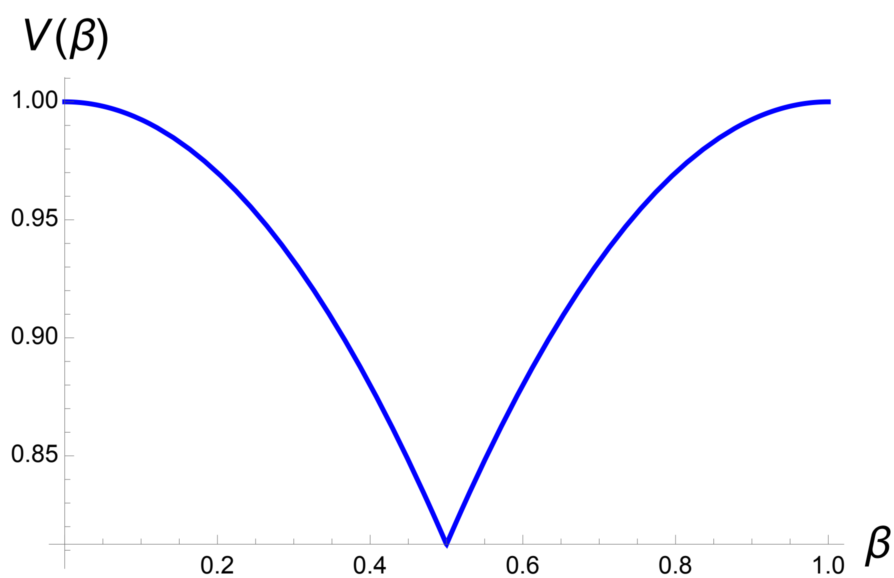

Figure 1: Consumer Surplus

Let us now consider using Proposition 2 for this application. From the discussion above, we can see that an increase in β for a certain type can induce a discrete change in the firm's location. The direction of this change benefits the consumer on average. This effect will tend to dominate in (6). To see that an increase in μ 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 location decision, an increase in β makes the firm increase its price, so the welfare effect of information through pricing is negative. As this effect is small compared to the location decision in this application, the first term (6) is positive.

Furthermore, the second term (7) is positive because $U(\theta, P)$ is convex in both location and price. Finally, if the revelation of information does not modify the firm's location choice, the third term (8) is zero at the firm's optimum, as the pricing decision is linear in β. We summarize our discussion in the following proposition.

Proposition 8 (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.

It is noteworthy that the model can be extended to show that when the type space is multidimensional, consumers may prefer privacy in one dimension while disliking privacy in another dimension. Consider a consumer with a two-dimensional type (θ, η), where θ denotes consumer location and η denotes consumer taste, and they are independent given any signals. 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.

In this extension, the consumer may dislike privacy in θ but prefer privacy in η. The intuition is consistent with the main model: the consumer may benefit from two-signal segmentations that induce the firm to choose different locations. However, as shown in the price discrimination model in Section 5.2, the revelation of information regarding taste η is always harmful to consumers under linear demand. See Appendix B for additional discussion.

Finally, the main model in this subsection has another interpretation 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 it is displayed (P). Denote the mass of the consumers visiting the platform as $D(P) = 2 - |\theta - x| - P$. It is decreasing in P because if the platform shows ads more frequently, consumers will visit the platform less often; it is also decreasing in

$|\theta - x|$ because an ad for a product the consumers are not interested in is more harmful to consumers. Each time the ad is displayed, the platform receives a fixed commission,[11] so the platform's profit is proportional to the mass of consumers who visit the platform, $D(P)$, times the frequency the ad is displayed (P). The rest of the derivation is the same, and the proposition under this setting can be interpreted as that consumers can only benefit from the revelation of information if it modifies the platform's ad targeting decision.

5.5 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 utility is given by θq, and the monopolist's cost function is $c(q) = q^2/2$.

A monopolist with beliefs $p \triangleq \Pr[\theta = \theta_H]$ offers the following menu of qualities

$$q_H = \theta_H, \tag{18}$$

$$q_L = \max\left\{0, \theta_L - \frac{p}{1-p}(\theta_H - \theta_L)\right\}. \tag{19}$$

Therefore, the expected utility of a consumer in a segment $s \in \Delta\Theta$ where $\Pr[\theta_H \mid s] = p$ is

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

Clearly, we have $CS(0) = CS(1) = 0$ since 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 9 *Zero privacy is never optimal for the consumer. Full privacy is optimal iff* $CS(p_0) = \text{cav}(CS(p_0))$*, i.e., iff*

$$p_0 \le p^* \triangleq \frac{\theta_L}{2\theta_H - \theta_L}.$$

Figure 2 shows CS and its concave closure $\text{cav}(CS)$ as a function of p.

[11] An alternative interpretation is that the products are owned by the platform, and there is a fixed probability that the consumer will make a purchase upon each impression.

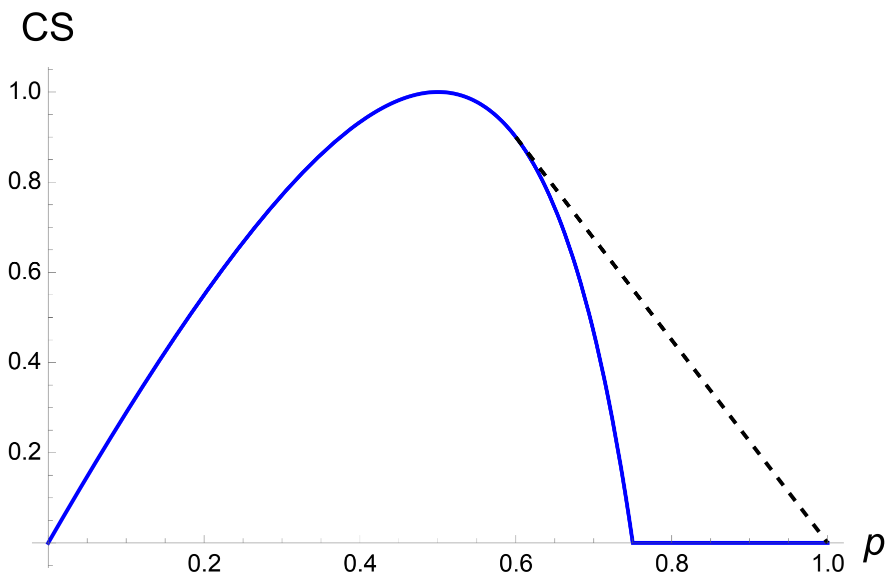

FIGURE 2: Consumer Surplus ($\theta_L = 3, \theta_H = 4$)

The intuition for Proposition 9 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 vs. the firm.

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 long literature on persuasion has shown) to multidimensional environments or environments where the posterior mean is not a sufficient statistic for the firm's problem.

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 $1/2$. 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 our framework, let the market action a represent the type θ to which the action 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, we have that

$$U(\theta, a) = v - \mathbf{1}[\theta = a]k(p).$$

The market action as a function of beliefs $a(p)$ is then the more likely type.

The surplus of a segment with a fraction p of type-A individuals is then 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.

5.7 Privacy as a Value in Itself

As discussed above, in some cases, individuals may dislike a lack of privacy per se, due to an "icky" effect. This effect can be seen as the same as the hacking example in the previous subsection, 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, while it is behavior B), the individual obtains utility $v > 0$.

In terms of the notation above, let the market action 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, we have that $U(\theta, a) = v - I[\theta = a]k$.We can then apply the analysis of the previous subsection to obtain that $V(p)$ is concave, and therefore the individual prefers more privacy.

6 Concluding Remarks

This paper presents a concavity/convexity theory of the effects of privacy. If the derived payoff function of the individual as a function of the market beliefs about the individual's type/behavior is concave (convex), individuals prefer more (less) privacy. In a variety of applications, we identify markets where consumer surplus is globally concave or convex as a function of the firm's posterior. In these cases, zero or full privacy is optimal for any value of the prior and any market segmentation.

In terms of the primitives of the problem, the paper identifies three effects regarding the shape of the derived payoff function: (1) how the market behaves in the interest, or against the interest, of the individual when it has more information about the individual's type; (2) the concavity/convexity of the primitive payoff function of the market actions; and (3) the concavity/convexity of the market actions as a function of the market beliefs.

One aspect not considered here, which would be interesting to investigate, is the effect of different allocations of information across different agents in the economy. Does an individual want a particular agent to have a precise signal about her type while having other agents with imprecise signals? The effects of privacy as a value in itself can be obtained relatively directly from the framework presented, but it would be interesting to have a better understanding of the mechanisms that generate that possibility in human interaction.

A Proofs

A.1 Proof of Corollary 1 and Discussion

Since U is linear in a,

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

so the term (7) is nil. Also, since a is linear in F,

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

so the term (8) is nil. The inequality is thus simplified to (9).

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. Since a is linear in F, assuming proper continuity conditions,[12] there exists a Riesz representation α of a such that

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

Applying this and the linearity of U in a to (9) leads to

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

The first expression $\int_\theta \alpha(\theta) \left[\mathrm{d}F'(\theta) - \mathrm{d}F''(\theta)\right]$ captures the average change in $\alpha(\theta)$ when the distribution of θ changes, and the second expression $\int_\theta U_a(\theta, a) \left[\mathrm{d}F'(\theta) - \mathrm{d}F''(\theta)\right]$ 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.

[12] One example of such conditions is that a is continuous with respect to the weak-* topology in the following sense. Let $\langle F, g \rangle \triangleq \int_\theta g(\theta) \mathrm{d}F(\theta)$ 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 $\langle \cdot, g \rangle \triangleq$ is continuous for all g. Theorem IV.1.2 in Schaefer (1971) ensures the existence of the Riesz representation of a given that a is continuous with respect to the weak-* topology.

A.2 Proof of Lemma 1

We first show $V(F)$ is convex if and only if $\partial^2_{\Delta F} V(F) \geq 0$ 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

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

$\Leftarrow$: Suppose that $\partial^2_{\Delta F} V(F) \geq 0$ 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 $\partial^2_{\Delta F} V(F)$ is homogeneous of degree 2 in ΔF, thus, for all $t' \in (0, 1]$, we have

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

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

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

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

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

Since 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 $\partial^2_{\Delta F} V(F) \leq 0$ for all F, F' and $\Delta F = F' - F$.

A.3 Proof of Proposition 3

Since

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

applying the Gateaux derivative operator yields

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

then applying the Gateaux derivative operator again yields

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

A.4 Proof of Proposition 5

We can expand

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

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

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

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

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

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

A.5 Proof of Proposition 6

We can expand

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

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

$$\partial^2_{\Delta F} V^{\mathrm{NE}}(F) = \frac{-(\gamma^4-5\gamma^2+12)(\mathbb{E}_{\Delta F}[\theta_1]^2+\mathbb{E}_{\Delta F}[\theta_2]^2)+2\gamma^3\mathbb{E}_{\Delta F}[\theta_1]\mathbb{E}_{\Delta F}[\theta_2]}{(4-\gamma^2)^2(1-\gamma^2)}.$$

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

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

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

A.6 Proof of Proposition 8

Expanding (17) 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 $[0,\frac{1}{2}]$ and $[\frac{1}{2},1]$, respectively.

First, we show the first part of the proposition. The firm's optimal location choice when $\beta \in [0,\frac{1}{2}]$ is $x^*(\beta)=0$, while the firm's optimal location choice when $\beta \in [\frac{1}{2},1]$ 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 $[0,\frac{1}{2}]$ or $[\frac{1}{2},1]$. Since $V(\beta)$ 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

$$V(\pi\beta_1 + (1-\pi)\beta_2) \leq \pi V(\beta_1) + (1-\pi)V(\beta_2). \tag{20}$$

The sufficient condition we provide is that $\beta_1 < \frac{1}{2} < \beta_2$ and

$$V'(\beta_1) \leq \frac{V(\beta_2) - V(\beta_1)}{\beta_2 - \beta_1} \leq V'(\beta_2). \tag{21}$$

To show that these conditions imply (20), discuss two cases: (i) $\pi\beta_1 + (1-\pi)\beta_2 \leq \frac{1}{2}$; and (ii) $\pi\beta_1 + (1-\pi)\beta_2 > \frac{1}{2}$. In cases (i), since V is concave on $[0, \frac{1}{2}]$, 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 simplifed to (20). In cases (ii), since V is concave on $[\frac{1}{2}, 1]$, 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 simplifed to (20).

Moreover, expanding (21) 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 ensures the consumer benefits from two-signal segmentations.

B Extension: Product Targeting with Price Discrimination

In this section, we discuss an extension to the main model in Section 5.4.

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 $\beta = \Pr[\theta = 0 \mid s]$ denote the posterior probability of location, and G denote the distribution of η after receiving a signal s.[13]

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

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

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.

For a fixed G, $V(\beta, G)$ is different from (17) in a constant and a term $\mathbb{E}_G[\eta](1+\max\{\beta, 1-\beta\})/2$, which is linear on either $[0, \frac{1}{2}]$ and $[\frac{1}{2}, 1]$ 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 8.

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

References

ACQUISTI, A., C. R. TAYLOR, AND L. WAGMAN (2016): "The Economics of Privacy," *Journal of Economic Literature*, 54(2), 442–92.

ACQUISTI, A., AND H. R. VARIAN (2005): "Conditioning prices on purchase history," *Marketing Science*, 24(3), 367–381.

ADJERID, I., A. ACQUISTI, L. BRANDIMARTE, AND G. LOEWENSTEIN (2013): "Sleights of privacy: Framing, disclosures, and the limits of transparency," in *Proceedings of the ninth symposium on usable privacy and security*, pp. 1–11.

ARGENZIANO, R., AND A. BONATTI (2023): "Data Linkages and Privacy Regulation," Discussion paper, Essex and MIT.

ARMSTRONG, M., AND J. VICKERS (2019): "Discriminating against captive customers," *American Economic Review: Insights*, 1(3), 257–72.

ATHEY, S., C. CATALINI, AND C. TUCKER (2017): "The Digital Privacy Paradox: Small Money, Small Costs, Small Talk," Discussion Paper 23488, National Bureau of Economic Research.

BARNES, S. B. (2006): "A privacy paradox: Social networking in the United States," *First Monday*.

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

CALZOLARI, G., AND A. PAVAN (2006): "On the optimality of privacy in sequential contracting," *Journal of Economic Theory*, 130(1), 168–204.

CHEN, Y., AND Q. LIU (2022): "Signaling through advertising when an ad can be blocked," *Marketing Science*, 41(1), 166–187.

CHOI, W., K. JERATH, AND M. SARVARY (2020): "Consumer Purchase Journey, Ad Wearout, and Privacy Choices," Discussion paper, Columbia University.

CONITZER, V., C. R. TAYLOR, AND L. WAGMAN (2012): "Hide and seek: Costly consumer privacy in a market with repeat purchases," *Marketing Science*, 31(2), 277–292.

de Cornière, A., and R. de Nijs (2016): "Online Advertising and Privacy," *Rand Journal of Economics*, 47(1), 48–72.

Dukes, A., Q. Liu, and J. Shuai (2022): "Skippable ads: Interactive advertising on digital media platforms," *Marketing Science*, 41(3), 528–547.

Eilat, R., K. Eliaz, and X. Mu (2021): "Bayesian privacy," *Theoretical Economics*, 16, 1557–1603.

Elliott, M., A. Galeotti, A. Koh, and W. Li (2022): "Market Segmentation through Information," Discussion paper, Cambdrige University.

Evans, D. (2009): "The Online Advertising Industry: Economics, Evolution, and Privacy," *Journal of Economic Perspectives*, 23(3), 37–60.

Farrell, J. (2012): "Can Privacy Be Just Another Good?," *Journal on Telecommunications and High Technology Law*, 10(2), 251–264.

Fudenberg, D., and J. Tirole (2000): "Customer poaching and brand switching," *RAND Journal of Economics*, 31(4), 634–657.

Gal-Or, E. (1985): "Information Sharing in Oligopoly," *Econometrica*, 53(2), 329–343.

Goh, K.-Y., K.-L. Hui, and I. Png (2015): "Privacy and Marketing Externalities: Evidence from Do Not Call," *Management Science*, 61(12), 2982–3000.

Goldfarb, A., and C. Tucker (2011): "Privacy regulation and online advertising," *Management Science*, 57(1), 57–71.

Gritckevich, A., Z. Katona, and M. Sarvary (2022): "Ad Blocking," *Marketing Science*, 68(6), 4703–4724.

Gross, R., and A. Acquisti (2005): "Information revelation and privacy in online social networks," in *Proceedings of the 2005 ACM workshop on Privacy in the electronic society*, pp. 71–80.

Hermalin, B. E., and M. L. Katz (2006): "Privacy, property rights and efficiency: The economics of privacy as secrecy," *Quantitative Marketing and Economics*, 4(3), 209–239.

Hirshleifer, J. (1971): "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, 61(4), 561–574.

Hui, K.-L., and I. Png (2006): "The Economics of Privacy," in *Handbooks in Information Systems*, ed. by T. Hendershott, vol. 1. Emerald: Bingley, United Kingdom.

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

Jullien, B., Y. Lefouli, and M. Riordan (2020): "Privacy Protection, Security, and Consumer Retention," Discussion paper, Columbia University and TSE.

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

Mussa, M., and S. Rosen (1978): "Monopoly and Product Quality," *Journal of Economic Theory*, 18(2), 301–317.

Narasimhan, C. (1988): "Competitive promotional strategies," *Journal of Business*, pp. 427–449.

Raith, M. (1996): "A General Model of Information Sharing in Oligopoly," *Journal of Economic Theory*, 71, 260–288.

Rhodes, A., and J. Zhou (2022): "Personalized Pricing and Competition," Discussion paper, TSE and Yale.

Robinson, J. (1933): *The Economics of Imperfect Competition.* Macmillan, London.

Schaefer, H. (1971): "Topological Vector Spaces," *Graduate Texts in Mathematics.*

Schmalensee, R. (1981): "Output and welfare implications of monopolistic third-degree price discrimination," *American Economic Review*, 71(1), 242–247.

Schoeman, F. (1992): *Privacy and Social Freedom.* Cambridge University Press, Cambridge and New York.

Shapiro, C. (1986): "Exchange of Cost Information in Oligopoly," *Review of Economic Studies*, 53(3), 433–446.

Shen, Q., and J. M. Villas-Boas (2018): "Behavior-based advertising," *Management Science*, 64(5), 2047–2064.

Shi, X., and J. Zhang (2022): "Welfare of Price Discrimination and Market Segmentation in Duopoly," Discussion paper, University of Toronto.

Taylor, C. R. (2003): "Supplier Surfing: Competition and Consumer Behavior in Subscription Markets," *RAND Journal of Economics*, 34(2), 223–246.

——— (2004): "Consumer privacy and the market for customer information," *Rand Journal of Economics*, 35(4), 631–650.

Thisse, J., and X. Vives (1988): "On the Strategic Choice of Spatial Price Policy," *American Economic Review*, 78(1), 122–137.

Villas-Boas, J. (1999): "Dynamic competition with customer recognition," *RAND Journal of Economics*, 30(4), 604–631.

——— (2004): "Consumer Learning, Brand Loyalty, and Competition," *Marketing Science*, 23(1), 134–145.

Villas-Boas, J., and Y. Yao (2021): "A Dynamic Model of Optimal Retargeting," *Marketing Science*, 40(3), 428–458.

Vives, X. (1984): "Duopoly Information Equilibrium: Cournot and Bertrand," *Journal of Economic Theory*, 34, 71–94.

——— (1988): "Aggregation of Information in Large Cournot Markets," *Econometrica*, 56(4), 851–876.

Warren, S., and L. Brandeis (1890): "The Right to Privacy," *Harvard Law Review*, 4(5), 193–203.

Westin, A. (1967): *Privacy and Freedom*. New York: Atheneum Publishers.

Yang, K. H. (2022): "Selling Consumer Data for Profit: Optimal Market-Segmentation Design and Its Consequences," *American Economic Review*, 112(4), 1364–93.

Zhang, J. (2011): "The Perils of Behavior-Based Personalization," *Management Science*, 30(1), 170–186.