PRIVACY ECONOMICS: SURVEY OF THE LITERATURE AND IMPLICATIONS FOR PUBLIC POLICY

Daniel P. O'Brien and Rodrigo Montes[1]

June 30, 2019

Abstract

Consumer demand for privacy has led to the development of a range of market solutions governing the exchange of consumer information. Examples include detailed disclosures by firms of their general policies regarding the protection and use of customer information and contracts between e-commerce platforms and their customers specifying how their data can be used. Nonetheless, governments around the world face increasing pressure to regulate the collection, protection, and use of personal information to protect consumer privacy.

The justification for increased regulation requires showing that markets for privacy fail, that existing regulation is insufficient to address the failure, and that additional regulation can correct any remaining market failure and make society better off. In this paper, we describe the economic aspects of privacy relevant for making this assessment, and we review the theoretical and empirical economic literature on privacy with this assessment in mind. We reach the following conclusions. First, the theoretical literature does not support the view that current U.S. market institutions for dealing with privacy concerns—truthful disclosure and customer choice with market-driven provisions of privacy—are likely to perform poorly. In particular, there is no basis for concluding that markets are likely to systematically under-supply privacy protections or give consumers too few options for choosing how their data will be protected. In fact, the literature indicates that constraints on firms' use of customer data would often harm consumers. Second, the empirical literature on privacy does not establish empirically that markets for privacy systematically fail or that broad restrictions on the use of customer data would make society better off. We conclude that there currently is no economic basis for macro-level increases in the regulation of the collection, protection, and use of personal information in the U.S.

[1] Funding was provided by the Computer and Communications Industry Association ("CCIA").

CONTENTS

1. INTRODUCTION

Long before the internet, businesses, data brokers, and other organizations collected information about consumers, often in large amounts, and adopted security policies relating to that information. Companies used some of the information to target to specific customers with advertisements via traditional "snail" mail and/or telemarketing or to determine what prices to charge specific customers or customer groups. These activities raised privacy concerns: consumers wanted certain data kept secure;[2] some did not like certain advertising directed toward them;[3] some did not enjoy paying prices based on information firms collected about them;[4] and some simply did not like having certain third parties hold information about them, for whatever reason.[5] Of course, consumers benefited then, as now, from sellers' use of consumer information to help match consumers with their most preferred products.

E-commerce has not changed the qualitative nature of these issues, but it has undoubtedly raised the prominence of the issue. Firms, governments, and other organizations collect, store, use, and distribute far more information today than they did before the internet, which has elevated privacy concerns. At the same time, the benefits to consumers from firms' use of their information have also risen[6]—and the potential cost of mis-guided regulation has risen in tandem.

The scope of the privacy problem in the internet age and the appropriate policy response are the subjects of intense debate. One view is that consumers have a basic right to privacy, and that Congress should either pass legislation that controls how consumer information is used or give consumers rights to determine how it is used.[7] Another view is that market forces combined with the "notice and choice" regulatory environment go a long way toward providing

---

2 See, e.g., the 1984 breach of TRW, a credit reporting company (www.nytimes.com/1984/06/22/business/credit-file-password-is-stolen.html, retrieved on 05/23/2019).

3 These concerns gave rise to the National Do Not Call Registry (www.consumer.ftc.gov/articles/0108-national-do-not-call-registry, retrieved on 05/23/2019).

4 See, e.g., the effect of credit scores on mortgage rates (www.nytimes.com/2013/05/04/your-money/credit-scores/vantagescore-ignores-paid-collections-in-setting-a-credit-score.html, retrieved on 06/28/2019)

5 "[P]rivacy has elements of both a final good (one valued for its own sake), and an intermediate good…" (See Acquisti, et al., 2016, p. 447).

6 The business models of many of the largest firms in the economy rely on the collection and use of consumer information, and the contribution of these firms to social welfare is obviously high. However, discussions in the popular press on privacy seem to focus much more on the potential harms from non-privacy than potential benefits of using consumer information.

7 This position is supported by many privacy advocates. See, for example, Letter from U.S. NGO's to U.S. Government Leaders Letter from U.S. NGOs to U.S. Government Leaders ("On the Need to Modernize and Update EU and US Privacy Law"), February 4, 2013, available at https://epic.org/privacy/intl/NGOs-to-US-Gov-re-EU-US-Privacy.pdf.

the right amount of privacy in the marketplace, and that the cost of additional regulation limiting the use of consumer information or dictating the nature of consumer control would likely outweigh the benefits.[8]

The purpose of this paper is to review the economic literature on privacy and discuss the implications of this literature for the debate over privacy regulation. Section 2 provides an economic definition of privacy that embodies the economic concerns expressed in policy discussions about privacy. Section 3 discusses the theoretical economic literature on privacy and identifies key factors that determine how well markets for privacy perform on their own, the role of enforceable disclosure of privacy policies, the role of contracts, and the limited scope for improving market performance through additional regulation. Section 4 reviews empirical literature and discusses its implications for privacy regulation in the context of the theoretical literature. Section 5 discusses policy issues in the context of both economic literatures.

Based on our review of the economic literature, we conclude that the literature does not provide a basis to conclude that current U.S. market institutions for dealing with privacy concerns—notice and choice with market-driven provisions of privacy choices—are likely to perform poorly overall, or that broad restrictions on the use of customer data are likely to increase social welfare. In an important sense, privacy is a demand attribute. Economic theory predicts:

- Companies have incentives to choose privacy policies that account for their impact on consumers so that consumers are willing to use their services. Under linear demand, constant marginal cost, and a constant value of information per customer, the privacy preferences of sellers/data collectors and consumers are aligned. This cuts against the notion that markets for privacy systematically fail in ways that call for more stringent privacy protections.

- Simulations of a seller's privacy choice in the demand attribute model suggest that over-supply of privacy is more likely than under-supply, but that both occur infrequently on average across a broad range of market scenarios.

- Simulations also suggest that globally imposing privacy in lieu of allowing the firm to make its profit-maximizing choice could reduce total and consumer welfare by a large amount.

- Even when companies cannot commit to privacy policies in advance, recent economic theory shows that economic policies that tax information sales, improve the transparency of privacy policies, and give consumers more control over their personal information have either mixed or neutral effects on consumer welfare.

Economic theory predicts that markets may under-supply privacy in limited circumstances, and this may justify targeted regulation for specific industries or to address

[8] See, e.g., Lenard and Rubin (2009).

specific issues. However, neither the theoretical nor empirical literature justifies a broad increase in privacy regulation.

2. DEFINITION AND CORE ECONOMIC FRAMING

A. Economic Definition of Privacy

Privacy can be difficult to define, but a careful definition is helpful to facilitate a discussion of economic issues. In this paper, privacy is a situation in which information conveyed in an interaction between two entities (persons or firms) is not used outside of that interaction in a way that has economic consequences. By "economic consequences" we mean that using the information outside of the interaction affects profits or consumer satisfaction (consumer "utility") in some way. When information conveyed in an interaction is used outside of the interaction and has economic consequences, we say that some degree of non-privacy exists.[9]

This definition encompasses the economic concerns commonly expressed in policy discussions about privacy. For example, web browsing conveys information about the browsing consumer to the internet service provider, the search engine, and websites the consumer visits. One or more collectors of data created in a browsing interaction may use it to provide the consumer with future information about products, make price offers to the consumer for new products, or make it easier for the consumer to find other information that might be valuable to the consumer based on their browsing history. In each case, information is used in ways that has economic consequences that presumably are good for the information collector (or they would not collect and use it) and may be good or bad for the consumer depending on the consumer's preferences. Thus, use of the information involves a degree of non-privacy according to our definition. Information that has not been conveyed through an interaction does not create a situation of privacy or non-privacy under our definition.

B. Core Economic Framing: Privacy is a Relationship Attribute

Privacy is not a typical good or service available for individual purchase. Instead, it is an attribute of an interaction that has economic consequences due to potential additional interactions with either the same party or third parties that acquire the information. A useful way to frame the economic aspects of privacy for welfare analysis is to think of privacy as a *relationship attribute.* This perspective provides a framework for understanding different aspects of the economic literature on privacy.[10]

---

[9] Stigler (1980) discussed non-privacy as "the possession and acquisition of knowledge about people" (p. 624). Posner (1981) observed the term privacy is used in three senses: "the concealment of information," "peace and quiet," and "freedom and autonomy." (p. 405). Acquisti et al. (2016) suggest that a common thread in definitions of privacy is that they "pertain to the boundaries between the self and the others" (p. 443).

[10] There is no new economics embodied in this terminology. It is simply a framing device for discussing different contributions to the economic literature on privacy.

Figure 1 provides a schematic illustrating this framing perspective. An individual has an initial interaction, in this case an online shopping experience. This interaction creates a relationship. If the relationship is private, illustrated by the lower set of blue arrows, it terminates after the initial interaction. If the relationship is non-private to some degree, it may involve subsequent interactions that have economic consequences. In the example here, a specific degree of non-privacy may lead to greater security,[11] a targeted advertisement, or a targeted price offer, perhaps in a subsequent browsing experience. The economic consequences of the use the information obtained in the initial interaction may include a guided purchase, improved security protection for subsequent browsing, or another external effect on the customer. If the information is private, on the other hand, there is no improved security, no guided purchase, and no other externality. The picture is simplified, as the degree of non-privacy may vary, but it illustrates the main aspect of our relationship attribute framing of privacy

Figure1. Privacy as a Relationship Attribute

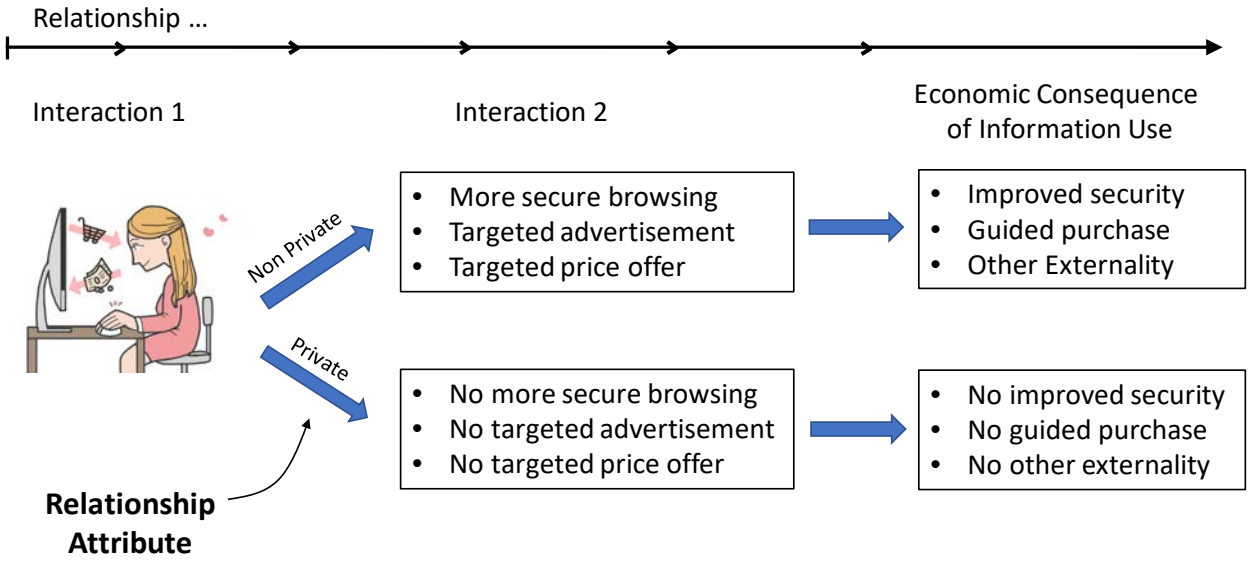

The interaction that creates the relationship need not involve a product purchase or an online shopping experience. A decision to use a particular search engine to surf the web at an

[11] For example, customers that use the Chrome browser with the Google search engine can opt in to reporting "additional data relevant to security to help improve Safe Browsing and security on the Internet." If the customer opts in, Chrome sends an incident report to Google every time the customer visits a suspicious page. See https://www.google.com/chrome/privacy/whitepaper.html.

out-of-pocket price of zero creates an interaction that may convey information that has economic consequences outside the interaction or the interaction may be with a social network.

The initial interaction that creates a relationship is typically voluntary as are the additional interactions that were consented to during the initial interaction, but some additional interactions that form part of the relationship may be involuntary. For example, a customer that makes an online purchase may receive marketing material via email, telephone, or even snail mail about certain products if the seller's privacy policy permits this use of the data. The marketing material might come from the seller involved in the initial interaction, or it might come from a third party that purchases data from the seller involved in the initial interaction if the seller's privacy policy allows transferring the data. On the other hand, the customer may voluntarily return to the store involved in the initial interaction hoping to receive information about related products. In either case, the firm's privacy choice typically is not directly mediated by price, which means that the choice creates externalities. For example, ads received by the customer based on information transferred in the first interaction may benefit or annoy the customer, i.e., they may create positive or negative externalities. One theme in the theoretical literature on privacy (discussed in Section 3.A) is the extent to which the seller in the initial interaction takes this into account in developing privacy policies and how the seller's choices compare with socially optimal choices.

The initial interaction may develop information about the customer's willingness to pay for certain products that leads to targeted price offers from seller in the first interaction or from a third party. This could benefit or harm the customer depending on whether the targeted price offer is lower or higher than the price offer would be if the seller did not use information to target. A second theme in the theoretical literature on privacy (Section 3.B) is whether sellers have incentives to use information in this way and if so the economic effects.

A customer's willingness to participate in an interaction depends on value the customer expects from the relationship created by the interaction, which in turn depends on the nature of relationship attributes, including privacy. For example, if a customer is concerned that an online bookseller will collect information during a purchase that might cause the customer to receive ads it does not want, the customer might choose to visit a local bookstore instead. On the other hand, a customer that that views a seller's use of customer information as beneficial might be more inclined to use the online seller. The presence of other competitors and the degree of competition, therefore, matters when assessing the incentives firms have to offer more privacy choices.

Sellers recognize that the willingness of consumers to participate and pay for interactions that raise privacy issues depend on how the seller handles these issues. A third theme in the theoretical literature (Section 3.C) is the extent to which effective contracts arise between sellers and consumers to address privacy concerns and how best to facilitate the development of such solutions in the presence of transaction costs and incomplete information.

In the era of big data, it may be possible to use information from consumers who share their data to draw inferences about consumers that do not share their data. If these inferences lead to actions that affect the consumers who do not share their data, externalities across consumers may occur that lead firms to capture too much or too little data depending in part on

whether the externalities are positive or negative. Such externalities take the analysis out of a pure relationship attribute framework. A fourth theme in the theoretical literature (Section 3.D) is the implication of across-consumer externalities for privacy regulation.

3. ECONOMIC LITERATURE ON PRIVACY – THEORY

Earlier economic literature made informal arguments regarding the efficiencies and inefficiencies from privacy.[12] Much of the more recent formal theoretical literature that examines the issues can be classified into three broad themes within the relationship attribute framework: (A) privacy as demand attribute; (B) privacy as it relates to price discrimination; and (C) and contractual solutions to non-privacy externalities. All three themes were foreshadowed in a seminal thought piece on privacy by Varian (1996), and subsequent literature has built on these themes. Sub-themes within these categories include the role of consumer rationality in determining outcomes and the interaction between privacy and competition. In the remainder of this section we discuss implications of the economic literature relating to each theme.

A. Privacy as a Demand Attribute

Consumer information conveyed in a consumer/seller[13] interaction conveys information whose use by the seller or a third party (if the information is subsequently shared) may benefit or harm the consumer.[14] Possible benefits include the consumer receiving more relevant information about future products or services. Possible harms include using the data in ways the consumer does not like, costs associated with identity theft, or a dis-utility associated with non-privacy.[15] The potential for a consumer's information to be used this way, which depends in part

---

12 For example, Posner (1978, 1981) argued that privacy may conceal information that, if revealed, would lead to more efficient trade. Stigler (1980) observed that regulation aimed at keeping information private faces the hurdle that consumers with information favorable to them have incentives to reveal it, and this may effectively reveal information about other consumers that choose not to reveal their information (because their information was favorable, they would reveal it).

13 We use the term "seller" to mean any commercial entity that offers a product or service to consumers at any out-of-pocket price, including zero, and is exposed to consumer information that may be used in ways that have economic consequences outside of the initial interaction. This includes online retailers, social networking services, search engine providers, internet service providers, credit card networks, banks, software providers, and cellphone apps. It also includes offline, brick-and-mortar retailers who store consumer information through loyalty programs, security cameras, etc.

14 As Varian (1996) observed, consumers may rationally want certain information about them to be known to other parties. For example, a consumer interested in purchasing a product benefits if a third party with knowledge of the consumer's interest provides information that helps the consumer make a better decision. On the other hand, consumers may find certain ads annoying.

15 Carrascosa et al. (2015) document that online behavioral advertising is a frequent practice, with categories valued more by advertisers are more intensely targeted.

on the seller's privacy and data security policies, raises economic issues relating to externalities and demand attributes.

If the use of the information is not independently mediated by price—i.e., if the seller and each customer in the interaction do not agree to terms for the use of the information—then the use of the information creates an externality that may benefit or harm the customer.[16] The presence of externalities brings in classic arguments of Coase (1960), who showed that complete contracting can provide a solution to the externality problem. In the privacy context, if there were no asymmetric information and no transaction costs involved in negotiating the use of customer information, firms would contract with consumers over all aspects of the use of their information and come to a mutually beneficial agreement. If such contracting were feasible and costless, it would solve this externality problem associated with non-privacy.[17] (Discussed further in Section 3.C below.)

In practice, asymmetric information exists, transaction costs are not zero, and contracting over privacy is incomplete. Incomplete contracting over privacy means that a firm's privacy policy, which it offers to all consumers with whom it interacts, is a demand attribute (Dorfman and Steiner, 1954; Spence, 1975; Johnson and Myatt, 2006; O'Brien and Smith, 2014) that affects the value to consumers of the interactions. Rational consumers understand the potential for non-privacy externalities from the use of their information, and their expectations about the effects of these externalities depend on the firm's privacy policy. These expectations affect consumers' willingness to participate in and pay for interactions that convey personal information. Different consumers may experience different effects from non-privacy externalities—some may benefit, some may be harmed. The aggregate demand for an interaction across all consumers depends on the price of the product, the firm's privacy policy, and the distribution of consumer valuations for the product and privacy.[18]

Sellers have incentives to adopt and disclose privacy policies that affect the demand for their products and services because they understand that rational consumers avoid interactions they view as harmful. Reputation and legal constraints (e.g., the illegality of unfair or deceptive claims) help make such policies credible. From an economic perspective, privacy in this context

---

16 Varian identified externalities as a privacy concern in cases where a mailing list is sold to a third party whose interests are not as well aligned with consumers as the original seller (Varian, 1996, p. 4). We use the term "externality" more broadly to mean any situation where one party makes a decision that affects another party and the decision is not mediated by price. Our use of the term recognizes as externalities situations in which demand attribute decisions that could be determined through contract between the seller and customer (though perhaps at high cost) are not determined that way, and the entity that makes the attribute decision does not internalize the full effects of the decision on the profits or utility of the other.

17 Contracting between firms and consumers would not solve problems associated with "across-consumer" externalities, discussed in Section 3.D below.

18 Models that treat privacy as a demand attribute include O'Brien and Smith (2014), Casadesus-Masanell and Hervas-Drane (2015), Jullien et al. (2018), and Shen and Villas-Boas (2018).

is a non-price demand attribute that shifts and rotates consumer demand (Johnson and Myatt, 2006) for the interaction.

A threshold question for policy analysis is whether a seller that can commit to an enforceable privacy policy will choose too much or too little privacy relative to the socially optimal level.

3.A.1 Firms incentives to adopt privacy policies consumers want

Spence (1975) provided the seminal analysis of the welfare effects of a monopoly seller's quality decision when that decision is an observable commitment to consumers. Quality, like privacy, is a demand attribute that shifts and rotates the demand for a product. However, Spence's analysis is incomplete for addressing the welfare effects of privacy for at least three reasons. First, while quality and privacy are both demand attributes, all consumers benefit from higher quality whereas some consumers may benefit from and some may be harmed by privacy. Second, Spence asked a different welfare question than the one relevant for policy toward privacy. He addressed whether a seller provides too much or too little quality relative to the socially optimal level *conditional on a fixed quantity of the product sold in the interaction*. However, unless quantity is regulated, the quantity sold of a product or service that comes bundled with an attribute (quality or privacy) generally varies with the attribute choice. Indeed, the primary motivation for a seller to disclose and enforce a privacy policy is to increase the demand for its services in ways that allow it to sell more to consumers or advertisers at prices that reflect the value created by the privacy choice. Third, the welfare analysis relevant when the seller's price and quantity are not regulated depends critically on how the choice of the attribute (privacy) shifts and rotates the demand for the product. This was not a focus of Spence. Johnson and Myatt (2006) extended Spence's analysis of demand attributes to understand the implications of demand rotations on firms' attribute choices, but they did not provide a welfare analysis. O'Brien and Smith (2014) extend the model of Johnson and Myatt to study privacy and provide the missing welfare analysis.

In the analysis of O'Brien and Smith, an initial interaction between several consumers that each demand one unit of a good and a monopoly seller of that good conveys information about the consumers.[19] The information may have value for many reasons, but to be concrete, suppose the information is valuable to the seller because it can be used to reach a member of a specific consumer group through targeted advertising or sold to a third party for that purpose. The value of the information to the firm is s per customer if it uses all of the information (no privacy). If the firm chooses not to pursue t percent of the opportunities for targeted advertising—i.e., if the privacy level associated with the information is t percent—then the value of the information to the firm is $(1-t)s$ per customer. This reflects the idea that greater privacy is associated with less use of the information for potentially profitable advertising. For the purpose of this illustration, we imagine that the firm can offer two levels of privacy, t_0 and t_1, with $t_1 > t_0$.

[19] The "product" in their model could be a service like web browsing purchased at zero price, but we focus on the case where the product is sold at a positive price here for ease of exposition.

The seller's marginal cost (here, the cost of selling to each customer) is c, and marginal cost net of the opportunity cost of privacy level t is $c - (1 - t)s$. Thus, a policy of no privacy ($t = 0$) yields a marginal cost of $c - s$, and a cost of complete privacy ($t = 1$) yields a marginal cost of c. Thus, the opportunity cost to the firm of complete privacy relative to no privacy is s per customer.

The (inverse) demand facing the seller is $P = \alpha(t) + \beta(t)(1 - q)$ where P is price, q is quantity, and α and β are demand parameters that depend on the privacy choice. The example here assumes linear demand and constant marginal cost for ease of exposition. Observe that the privacy choice t affects the demand intercept $\alpha(t) + \beta(t)$ and slope $\beta(t)$. Thus, the seller's choice of privacy can shift and rotate demand in arbitrary ways depending on how it affects the demand intercept and slope. How the privacy choice does this depends on the distribution of preferences of different consumers for the product and privacy.[20]

Figure 2 presents a standard textbook diagram for evaluating a monopoly seller's choice of price and privacy in this context. In this example, a higher level of privacy yields greater consumer demand because all consumers are assumed to value privacy. Additional privacy benefits high valuation consumers by more than it benefits low valuation consumers, as reflected by the outward shift and steepening of demand when privacy is increased from t_0 to t_1.

Figure 2. Privacy as a Demand Attribute

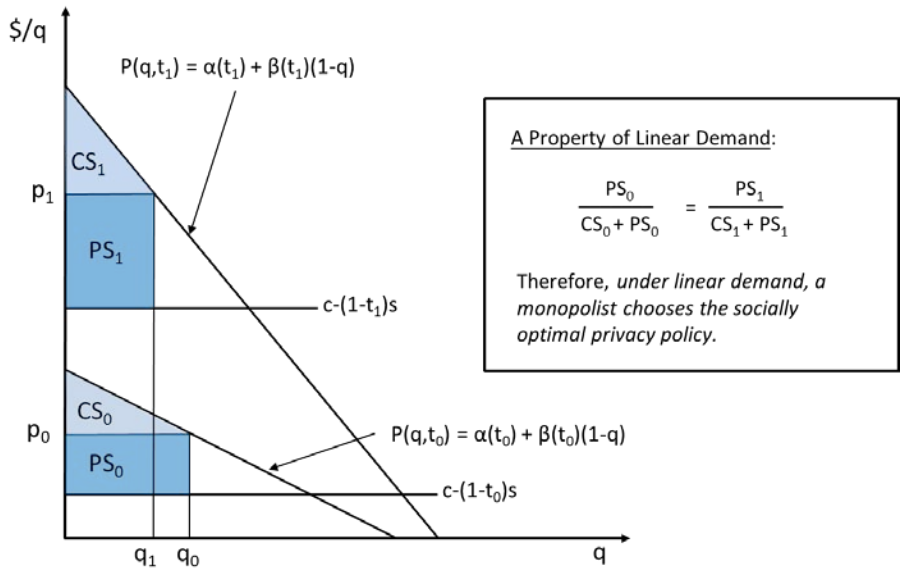

A critical policy question is whether the seller will choose too much or too little privacy relative to some measure of the socially desirable amount. This example yields a surprising

[20] Formally, this is the linear demand that arises when consumer preferences are distributed in the location-scale class. What is relevant for the example here is that we allow the privacy choice to shift and rotate demand in all possible ways. See O'Brien and Smith (2014) for more details.

answer: *in the absence of any regulation except a requirement for truthful disclosure, the seller has an incentive to choose the level of privacy that maximizes consumer surplus and total welfare (the sum of consumer and producer surplus).* That is, conditional on the restrictions of the model—that the seller offers the same privacy policy to every consumer and the seller's price and quantity are not regulated—the seller chooses the socially optimal level of privacy.[21]

To see why, recognize that the seller chooses price and the level of privacy to maximize its profits, which in this example is the same as the level of producer surplus, labelled PS in Figure 2. At privacy level t_0, the profit-maximizing price is p_0 and the seller's profit is PS_0.[22] At privacy level t_1, the profit-maximizing price is p_1 and the seller's profit is PS_1. The consumer surplus values associated with these privacy levels, which measure the value consumers receive from purchasing the product, are CS_0 and CS_1 respectively.[23] A property of the linear demand and constant marginal cost case examined in this example is that the ratio of consumer and producer surplus is constant – in particular, consumer surplus is exactly half the value of producer surplus. This means that the ratio of consumer surplus to total welfare (measured here as the sum of consumer and producer surplus) is also constant and equal to 1/3. These facts mean that any increase in producer surplus due to the shift and rotation in demand and the shift in marginal cost caused by a change in privacy policy is met by proportionate increases in consumer surplus and total welfare. Therefore, the privacy choice that maximizes profit is the same as the privacy choice that maximizes consumer surplus and total welfare.[24] [25]

---

21 Technically, the seller's privacy choice is a "second best" optimum conditional on the inability to specify different privacy levels for individual consumers and the absence of price regulation. In this analysis, society would be better off if the seller contracted bilaterally with each consumer over the degree of privacy afforded to that consumer. However, the focus of the demand attribute formulation of the privacy problem is the effect of the seller's choice of privacy policy when transaction costs prevent complete bilateral contracting over privacy.

22 Under linear demand and constant marginal cost, the profit-maximizing price lies halfway between marginal cost and the demand intercept, which is the price above which demand is zero.

23 Consumer surplus is the difference between what consumers would be willing to pay for a given quantity and the amount they actually pay. It is a common measure of the value to consumers of purchasing a product.

24 Readers familiar with Spence (1975) may wonder whether he anticipated this result in his analysis of a monopolist's optimal choice of quality. He did not. Spence examined the question of whether a monopolist produced too much or too little quality relative to the social optimum holding product quantity fixed. O'Brien and Smith (2014) examined whether a monopolist produces too much or too little privacy at the profit-maximizing quantities associated with each privacy choice. Spence's analysis is relevant in a regulated environment that constrains a firm's quantity, whereas O'Brien and Smith's analysis is relevant for unregulated environments where the seller's price and quantity are unconstrained.

25 An issue related to privacy but with somewhat different incentives is data security. Like privacy, data security is reasonably viewed as demand attribute that affects consumers' desires to participate in interactions that transfer information. A difference is that data security likely entails fixed costs of developing secure systems, while the cost of privacy is primarily a marginal

The conclusion that the preferences of the seller and consumers are *perfectly* aligned is an artifact of the linearity of the model, but the conclusion that these preferences are often very well aligned is not, as discussed further below. Nuances aside, this finding cuts against the notion that markets for privacy systematically fail in ways that call for more stringent privacy protections when privacy is a demand attribute. In the linear model discussed above—often the first model economists write down to generate economic intuition—regulating a monopolist's choice of privacy beyond the enforcement of truthful disclosure would reduce consumer and total welfare.

The interaction illustrated in Figure 2 involves the purchase of a product at a positive price, but many online interactions involve services in which the out-of-pocket price paid by consumers is zero and the seller monetizes the service through advertising revenue. The model of O'Brien and Smith captures these cases by incorporating a zero lower price bound that constrains the seller's price to be no less than zero. The idea is that the seller might like to offer a negative price and subsidize the service, but it is constrained from doing so because consumers would exploit a negative price, e.g., by making multiple purchases under different identities.

When the seller in the demand attribute framework is constrained by a zero lower price bound, the clean result that the preferences of the seller and consumers are perfectly aligned under linear demand no longer holds. The reason is that an increase or decrease in privacy that would benefit consumers may not be profitable for the firm when its price is constrained. In this case, the firm may under-supply privacy in some cases.[26] Although this suggests that the scope for welfare-improving privacy regulation may be higher in cases where the price of the service is zero, simulations conducted by O'Brien and Smith (discussed further in the next section) suggest that the frequency of under-supply is likely to be low and that a broad regulation requiring additional privacy is likely to reduce consumer and total welfare on average.

O'Brien and Smith also consider the predictions of the demand attribute model of privacy when firms compete by offering both a price and privacy level. They find that in this case, the market yields either the efficient level of privacy or *too much* privacy. In the absence of a zero lower price bound, competition generates the efficient level of privacy. The reason is that competition leads to outcomes where different firms offer different privacy levels and prices appropriate for consumers with different preferences. In the example here, a competitive market leads to privacy level t_0 offered to consumers at price $c - s(1 - t_0)$, and privacy level t_1 offered to consumers at price $c - (1 - t_1)s$. Consumers would then choose the price-privacy combination that is best for themselves. That is, competition in price and privacy leads to

---

opportunity cost from the inability to use information. In the demand attribute framework, fixed costs make it more likely that a seller will under-supply the attribute (See O'Brien and Smith (2014), Section III.F). Consistent with this, the FTC's enforcement efforts in the privacy and data security space have focused more on data security.

[26] For example, if additional privacy increases demand for the service but would reduce the profit-maximizing subsidy (in the absence of consumer exploitation of such a subsidy), then the constraint on the firm's ability to adjust the subsidy (because its price is zero) prevents the seller from benefitting from the reduced subsidy. In this case, the seller could have too little incentive to offer additional privacy.

multiple price-privacy combinations that optimally sort customers into two groups, those that should purchase with privacy and those that should not, based on their preferences.[27] The equilibrium product prices in a competitive market therefore indirectly mediate the exchange of information just as would occur if the seller could contract with each customer at zero transaction cost.

If competing firms are constrained by a zero lower price bound, competition leads to too much privacy relative to the social optimum. The intuitive reason is that the zero lower bound prevents subsidies to individuals who would experience greater benefits from having their information collected and used more intensely.

3.A.2 Robustness of alignment between firms and consumers' incentives

The example in the preceding section assumes linear demand and that the seller can commit to and disclose an enforceable privacy policy. Neither modification significantly enhances the scope for privacy regulation.

3.A.2.1 Relaxing the linear demand assumption

O'Brien and Smith (2014) consider demand functions generated by preference distributions in the location-scale class. In simple terms, this allows the demand for the product to have arbitrary curvature and assumes that a firm's privacy choice shifts and rotates that demand. Under this assumption, a monopoly seller may over- or under-supply privacy depending on the nature of the demand curvature and how privacy rotates demand. However, it turns out that with common assumptions economists make about demand, the scope for the under-supply of privacy is limited.

The conditions that determine whether the seller over- or under-supplies privacy turn out to depend on two rather arcane factors: (1) the nature of demand curvature (the extent to which the slope of demand in price increases or decreases with price), and (2) the effects of privacy on the slope of demand, which reflects the sensitivity of demand to price. If demand is linear, as in the example above, the preferences of the seller and consumers are perfectly aligned, and there is no scope for regulation in the demand attribute framework with a monopoly seller. If privacy shifts demand without changing its slope, the preferences of the seller and consumers are also aligned even if demand has curvature. This would be true if consumers all placed the same value on privacy, a result also recognized by Farrell (2012). For the privacy preferences of the seller and consumers to be mis-aligned, demand must have curvature (i.e., it must be non-linear) and privacy must alter the slope of demand. And preferences are not necessarily mis-aligned even if these conditions hold.

The most common preference distributions used in economics (e.g., Normal or Weibull) generate demand curves with a specific demand curvature such that a privacy policy that rotates demand clockwise, increasing its slope, tends to benefit consumers that place high value on the product relatively more than it benefits marginal consumers that are close to indifferent between consuming the product or not. If privacy rotates demand clockwise—it does so if consumers

[27] The considerable variety observed in firms' stated privacy policies supports this view. See Ramadorai et al. (2019).

with higher valuations for the product also have higher valuations for privacy—then the seller may under-supply privacy. The reason is that the seller cares about the impact of additional privacy on the purchase decisions of "marginal" consumers that are enticed to buy due to increased privacy, but it does not care about the impact of privacy on "inframarginal" consumers who would buy even without additional privacy protection. Because the seller under values the benefits of privacy to high valuation customers that also happen to place high valuation on privacy, it may offer too little privacy protection.

Unfortunately, the circumstances in which a seller will under-supply privacy are difficult to determine, as they depend on both demand curvature and how privacy rotates demand, both of which are difficult to measure. To provide some window into the likelihood that the seller may under-supply privacy, O'Brien and Smith simulate a range of market scenarios in which consumers have normally distributed preferences for the product and privacy. They start with essentially flat priors over the relevant parameters: the mean and variance of the values of the product and privacy, and the correlation between consumer values for the product and privacy. The parameter ranges considered are very broad and intended to be agnostic. Across 2 million simulated markets, they find that the seller fails to supply privacy when the planner wants it in less than one half of one percent of the cases, and it over-supplies privacy in just less than 1 percent of the cases. Thus, over-supply is more frequent than under-supply, but both are small on average across a broad range of market scenarios. Further, the social benefit of identifying and correcting under-supply averages only 3 percent across all cases in which under-supply occurs. An apparent reason for the relative infrequency and small magnitude of the distortions in the simulation is that the demand curve implied by the normal distribution is quite close to linear over much of its range. As shown in our example above, when demand is linear, (and marginal cost is constant), the privacy preferences of the seller and consumers are aligned.

O'Brien and Smith also examine the welfare cost of globally imposing privacy in lieu of allowing the firm to make its profit-maximizing choice. This regulation reduces welfare by a large amount—about 51 percent on average across all 2 million simulated markets.

Although a global imposition of privacy has strongly negative effects on average in this framework, there are specific cases where a seller may offer too little protection relative to the social optimum. For the case of normally distributed preferences for the product and privacy, this occurs when consumers' willingness to pay for the product and privacy are highly positively correlated. O'Brien and Smith observe that healthcare markets could be one case where sellers (hospitals, physicians, drugstores) might under-supply privacy protection due to such a correlation in preferences:

> "[O]ne can also imagine scenarios where consumers' product and privacy valuations are positively correlated, possibly strongly so. One example may be the privacy of health care information. Consumers that are sick likely place high value on health care, and they may also have strong desires to keep information about their health confidential to avoid possible discrimination by employers or other third parties."[28]

[28] O'Brien and Smith (2014), p. 30.

Consistent with this observation, healthcare is one area in the U.S. where regulation goes beyond the notice-and-choice approach by imposing specific rules on the use of consumers' protected health information.[29]

3.A.2.2 Relaxing the assumption that the seller can commit to and disclose an enforceable privacy policy

The strong results that tend to counsel against broad privacy regulation in the preceding section assume that the seller can credibly commit to and disclose a privacy policy. That assumption focuses the regulation question on whether additional regulation is called for in a regime where either reputation or truthful disclosure laws allow sellers to make such commitments. However, the assumption of perfect enforcement of disclosed policies is a strong one. If enforcement costs are high, or if penalties for violation are too low, the commitment assumption could be too strong.

Jullien et al. (2018) go to the other extreme by relaxing the commitment assumption completely. That is, they assume (in most of their analysis) that the seller's privacy choice is an independent choice of the seller that is not observed by consumers before they choose whether to participate in a given interaction.

In their model, consumers visit website at no charge to engage in some activity they value without knowing the online site's privacy policy or their own valuations for the privacy of the information conveyed in the visit. The online site receives a fixed dollar amount for each visit (e.g., from providing traffic for advertisers). Simultaneous with consumer visit decisions, the online site chooses its privacy policy. Consumers then receive additional advertisements with some probability that depends on the firm's privacy policy. Based on the advertisements they receive and whether they like them, consumers draw inferences about their own preferences for privacy over information transferred when they interact with the site. If consumers dislike the advertisements they receive enough, they choose not to visit the site a second time. Thus, the online site has an incentive to protect consumer privacy to some degree because more aggressive advertising increases the risk that consumers will have bad experiences and choose not to return to the site.

Jullien et al.'s findings indicate that the commitment assumption is perhaps not as important as one might have thought. Based on their analysis, they conclude:

> Greater privacy protection is a mixed blessing for consumers, who, on the one hand, are better protected from intrusions, but, on the other hand, may be deprived of positive matches with third parties and are less informed about their vulnerability to third-party intrusions. Consequently, it is difficult for authorities to regulate privacy protection in a way that reliably improves consumer welfare. For example, policies that tax information

[29] Title II of The Health Insurance Portability and Accountability Act (HIPAA) has a Privacy Rule that regulates the use and disclosure of Protected Health Information (PHI). Healthcare market participants can disclose PHI to certain parties to facilitate treatment, payment, or health care operations without a patient's express written authorization, but other disclosures of PHI require written authorization from the patient.

sales, improve the transparency of privacy policies, and give consumers more control over their personal information, all have either mixed or neutral effects on consumer welfare.[30]

The common market force present in both the commitment and no-commitment variants of the demand attribute model that tends to protect consumers is the seller's desire to maintain the consumer demand for interactions. In the model that assumes commitment is possible, the seller discloses an enforceable privacy policy that takes into account the impact of non-privacy on consumer demand for the interaction. Although the assumption in that model is that the firm's privacy policy is observed, a long run interpretation is that consumers come to understand what the policy is and that it will be enforced, either through laws that prevent unfair or deceptive practices or the firm's desire to maintain its reputation. In Jullien et al.'s model focusing on the case of no commitment, the seller's privacy policy affects consumer decisions about whether to make a return visit. Either way, the theoretical literature based on the demand attribute approach does not provide a basis for concluding that markets for privacy systematically fail in ways that regulation will improve.[31]

B. Price Discrimination

A second theme in the theoretical literature under the relationship attribute umbrella is that sellers may use consumer information from interactions to price discriminate. Taylor (2004), Acquisti and Varian (2005), and Hermalin and Katz (2006) rigorously examine this issue. They identify circumstances in which a seller would benefit from using consumer information to price discriminate and examine the welfare consequences of such discrimination.[32]

A simplification of the model presented in Acquisti and Varian conveys many of the main ideas. Imagine that a consumer is willing to pay either $10 or $6 for a seller's product in each of two periods, but the seller does not know which amount. Suppose the two possible consumer valuations are equally likely, and that the seller's costs are zero (for simplicity).

In this scenario, the seller would charge $10 (or a penny less) in both periods if it knew the consumer's valuation was $10, but it does not have this information. If the seller charges $10, it risks losing the sale if the consumer's valuation turns out to be only $6. One strategy the seller might attempt is to charge $10 in the first period hoping that the consumer is willing to pay it, and then charge $10 in the second period if the consumer purchases in the first period

---

30 Jullien et al., (2018), p. 32.

31 There may be specific exceptions, such as health care, the online activities of children, and financial information, as discussed in Section 5.

32 For evidence of price discrimination online see Mikians et al (2012) and Mikians et al. (2013). However, online price discrimination may not be as prevalent as thought. See Narayanan, Arvind (2013), "Online Price Discrimination: Conspicuous By Its Absence," https://33bits.wordpress.com/2013/01/08/online-price-discrimination-conspicuous-by-its-absence/, retrieved on 06/25/2019.

(knowing that the consumer is willing to pay that amount) and charge $6 in the second period if the consumer does not purchase in the first period. Two other possible strategies involve no price discrimination based on purchase history: charge $10 in both periods, or charge $6 in both periods. Which strategy is best for the seller?

Among the strategies that do not involve price discrimination, the best one is to charge $6 in both periods. This leads to certain sales in both periods at $6 and an expected profit of $12. If the seller were to charge $10 in both periods instead, its expected profit would be only $10 [=(1/2) x 10 + (1/2) x 10] because there is 50-50 chance the consumer will pay $10 in both periods and a 50-50 chance that will not make a purchase in either period. Thus, the seller's best non-discriminatory strategy is to charge $6 and earn an expected profit of $12.

What about the price discrimination strategy: charge $10 in the first period and then charge $10 in the second period if the consumer buys in the first period and charge $6 in the second period otherwise. Does this strategy increase the firm's profit? The answer turns out to depend on whether the consumer is myopic, meaning that it ignores the implications of its behavior for future pricing, or rational, meaning that it recognizes how its first period behavior affects price.

If the consumer is myopic, the price discrimination strategy yields an expected profit of $13 [=(1/2) x 20 + (1/2) x 6], as there is a 50-50 chance consumer will buy at a price of $10 in both periods and a 50-50 chance that it will not buy in the first period but then will buy at a price of $6 in the second period. This profit exceeds the $12 the seller would earn from its best non-discriminatory strategy. Thus, if the consumer is myopic, the seller benefits from conditioning the second period price on information collected in the first period. Acquisti and Varian show quite generally that if the consumer is myopic, the seller benefits from using information about the consumer's purchase behavior to price discriminate.

If the consumer is rational, on the other hand, the price discrimination strategy is not optimal for the seller.[33] Why? If the consumer's valuation is $10, a rational consumer knows that if it buys at that price in the first period, it will pay the same price in the second period and will earn zero surplus [=10-10], while if it chooses not to buy in the first period, it will purchase in the second period at $6 for a surplus of $4 [=10-6]. If the consumer's valuation is $6, the consumer also purchases at $6 in the second period. Thus, the rational response of the consumer to the seller's strategy, whether the consumer's valuation is high or low, is to wait until the second period and purchase at $6, which means the seller's expected profit from its price discrimination strategy when it faces rational consumers is only $6. Thus, *if the consumer is rational, the seller does not benefit from using the customer's information to price discriminate*. Moreover, total welfare in this example is higher when there is no price discrimination (the seller's optimal choice), so the seller will make the socially optimal choice if it can credibly commit to a policy of not conditioning price on purchase behavior.

---

[33] Under competition, some firms may refrain from using consumer data to avoid intense price competition. See, e.g., Chen and Iyer (2002).

A key lesson from this example is that rational, market-driven behavior by consumers can lead sellers to try to avoid using their purchase histories to make targeted price offers.[34] Many papers that model the price discrimination aspects of privacy find a similar result.[35] The market force behind this result is related to the market force discussed in Section 3.A above that prevents a seller from excessive targeted advertising based on consumer information: a rational seller recognizes that the demand for its product or service from rational consumers depends on its privacy policy. In both cases—targeted advertising and purchase history-based price discrimination—privacy that reduces such advertising and price discrimination is a relationship attribute, and sellers have incentives to choose this attribute in recognition of how rational consumers will respond. This market force can prevent sellers from under-supplying privacy.

Despite the prediction of many models that sellers often benefit from committing *not* to price based on purchase history, some online sellers have done so. One explanation could be that consumers have a degree of myopia. Acquisti and Varian (2005) show that in this case, the welfare effects of price discrimination are theoretically ambiguous. In the example above, conditioning on purchase history when the consumer is myopic raises expected profit and reduces expected consumer surplus and total welfare. Thus, in the case of myopic consumers with low dispersion between high value and low value customers, the seller will price discriminate because its expected profit is higher, but this is a suboptimal outcome because it decreases consumer surplus and total welfare. However, if the dispersion between high and low valuation consumers is larger, conditioning on purchase history raises welfare and the seller will make the right choice. For example, suppose the product is worth \$10 to the high valuation consumer but only \$4 to the low valuation consumer. In this case, the non-discriminatory price is \$10, which leads to an expected profit of \$10 [=(1/2) x 20], an expected consumer surplus of zero, and an expected total surplus of \$10. The optimal discriminatory prices (with myopic consumers) are \$10 in the first period and either \$10 or \$4 in the second period depending on whether the consumer purchases in the first period. Expected profits are \$12 [=(1/2) x 20 + (1/2) x 4], expected consumer surplus is zero, and expected total welfare is \$12. Thus, price

---

[34] As Acquisiti and Varian (2005) observe, the result in this example is a special case of a very general result in the intertemporal price discrimination literature showing that a seller does not benefit from intertemporal price discrimination to a fixed set of rational consumers. See Stokey (1979).

[35] See Villas Boas (2004) (A monopolist prefers to commit not to recognize its previous customers.); Conitzer et al. (2012) (When consumers that are engaged in repeat purchases can remain anonymous at zero cost and choose to do so, the seller may earn higher profit than when anonymity has high cost and consumers share their information.); De Corniere and Montes (2017) (When a firm can customize products for consumers using consumer data, the firm may be better off committing to a uniform price to induce data revelation rather than engaging in first-degree price discrimination, which would cause some consumers not to share their data.); Ichihashi (2018) (In a multi-product setting, disclosure benefits the consumer via more accurate product recommendations, but it also allows the seller to price discriminate. Such a seller is better off committing to a price in advance than it is pricing based on the consumer's disclosure, while the consumer is better off with price discrimination based on information disclosure. Total surplus can be lower if the seller is able to price discriminate, in which case a rule requiring privacy would increase total surplus but reduce consumer surplus.)

discrimination raises profits and welfare in this case, and the seller has an incentive to choose the socially optimal privacy policy. The fact that price discrimination does not increase consumer surplus in this case is an artifact of the simplicity of the example. If the low valuation consumers' willingness to pay was equally likely to be $4 and $4.50, the discriminatory price would still be $4 and consumer surplus would be positive.[36]

In some cases, sellers may find it difficult to commit to refrain from conditioning prices on purchase history. For example, the product offered in the second period might be an upgraded version of the product sold in the first period and it might be too costly to specify its price in advance. In this case, the seller may condition the price of the second product on the consumer's purchase history even if consumers are rational because it is too costly for the seller to commit not to do so. The welfare effects of requiring the seller to refrain from conditioning on the first period purchase decision in this case are ambiguous for the same reason as the case of myopic consumers. Discrimination can increase output by identifying and selling to lower valuation consumers at lower prices than would prevail without discrimination, or it can reduce output by causing the seller to charge higher first period prices in hopes of capturing profits from high valuation consumers.

The ambiguity of the welfare effects of conditioning price on purchase history mirrors the ambiguous effects of price discrimination that are well-known in the economics literature.[37] Just as there is no strong basis for laws forbidding price discrimination, the literature on the price discrimination effects of non-privacy does not provide a strong basis for preventing price discrimination based on purchase history.

The price discrimination-based models we have discussed so far abstract from competition between sellers. Several models show that competition can raise the benefits of information revelation for consumers and that sellers often have incentives to make commitments not to use customer information to price discriminate to soften competition.[38] As is

---

36 Hermalin and Katz (2006) examine a model in which consumers with private information (e.g., over their willingness to pay for a product) can choose whether or not to reveal an informative but not fully revealing signal of the information. "Privacy" means the consumer keeps the signal private, and "non-privacy" means the signal is revealed. They find that the effects of privacy that prevents conditioning on the signal has ambiguous welfare effects.

37 See, e.g., Varian (1989)

38 Corts (1998) identifies an important class of models in which price discrimination by imperfectly competitive firms leads to lower prices for all consumers. Villas-Boas (1999) studies competition between firms over time when firms can recognize their previous customers and use this recognition to price based on the whether the customer previously purchased from the firm or from a competitor. He shows that such recognition intensifies competition as firms attempt to attract their competitors' customers. Similarly, Fudenberg and Tirole (2000) study the incentive of competing firms to "poach" their competitors' customers with "behavior-based" price discrimination, which is the same thing as Acquisti and Varian's (2005) "conditioning on the purchase history" discussed above. They also find that such price discrimination can intensify competition. See also Cooper et al. (2005), Taylor and Wagman (2014), and Montes et al. (2019). The price discrimination strategies in these papers require knowledge of differences in

true for monopoly price discrimination, the literature on the price discrimination effects of non-privacy among competing firms does not provide a strong basis for preventing price discrimination based on purchase history. If anything, the literature suggests that preventing such discrimination could be more harmful than in the monopoly case by softening competition among competitors.

C. Contracting Over Privacy

3.C.1 The Coase Theorem

The third theme in theoretical literature on privacy involves contracting between the seller and consumers over specific aspects of the consumer's privacy. In a sense, all privacy concerns raised by information conveyed in commercial transactions involve externalities of some type. Coase (1960) established that if the parties involved in an externality do not have asymmetric information and can contract directly over the action that creates the externality, the outcome will be efficient—that is, the socially desirable amount the action that creates the externality will take place. In the context of non-privacy externalities, if it were possible for the seller and each consumer to negotiate complete contracts governing the privacy of the consumer's information, the non-privacy externalities we have been discussing would disappear, and the market would provide the socially optimal level of privacy.[39]

Of course, the real world has both incomplete information and transaction costs, and this places constraints on the Coase solution. The demand attribute and price discrimination literature discussed above takes these constraints into account and assumes that the seller and the consumer cannot contract over privacy considerations. The opposite case—complete contracting over every aspect of privacy—is not realistic, but a middle ground that involves an intermediate degree of contracting over privacy is possible and moves in the direction of a Coasian solution. To be sure, many online sites offer some consumer choices about how their information will be used.

Many of these solutions have emerged without specific regulation, which is not surprising given that both sellers and consumers stand to gain from agreeing to provisions that allow them to conduct mutually beneficial transactions. However, a question that arises with contractual solutions to externality problems is who holds the property rights: the seller who inflicts the externality, or the consumer who is affected by it (positively or negatively). The current situation in the U.S. is that sellers hold the property rights in the sense that, subject to constraints on fraud and regulations specific to certain industries and circumstances, sellers are free to use or sell information gathered in interactions with consumers for use in targeted advertising or price offers. An important question is whether this is the best allocation of

---

preferences between consumers or consumer groups, and suppressing this information often softens competition and raises prices to consumers. For a survey of the literature on price discrimination in environments with competing firms, see Stole (2007).

[39] Across-consumer externalities (Section 3.D) could still be present.

property rights, or whether transaction costs would be lower if consumers had more rights over their information.

Varian (1996) points out that the socially optimal assignment of property rights is the one that minimizes transaction costs. For the issue of targeted advertising, he observed that the relevant comparison involves the transaction cost to the individual of having his or her name removed from a list potentially at some cost (which could be non-monetary) versus the cost to the list owners (the seller) of soliciting permission from individuals to add them to the list at some price.[40] We discuss this further in Section 5 below.

3.C.2 Consumer incentives to reveal their information -- unravelling

An intuitive argument suggests that granting consumers property rights over their data could be an effective and efficient way to protect consumer privacy. In theory, consumers could sell their personal information to companies with whom they interact. But for privacy concerns associated with price discrimination, this strategy seems unlikely to work.

When consumers have private information about the value they place on a product, the optimal strategy for the seller typically involves charging a price above the value of the lowest valuation customer. Suppose there are two types of customers: some with a low valuation that is below the current price, and some with a high valuation that is above the current price. If the low valuation customer values the product more than cost, it has an incentive to reveal this to the seller to convince the seller to offer the product at a mutually agreeable price between cost and the customer's valuation. Because a low valuation customer has an incentive to reveal this, the failure of a customer to reveal its willingness to pay effectively reveals that it a high valuation customer. But then the seller knows it can charge those customers higher price.

The idea that the incentive for some set of consumers to reveal their information leads to the revelation of information about others is referred to as "unravelling" in the economics literature.[41] In the privacy context, if the policy concern is that it is unfair or harmful to some consumers for firms to condition prices on information about them, then the unravelling result suggests that giving consumers the property rights over their information is unlikely to address the issue, as Stigler (1980) observed.[42] For example, if a life insurance company conditions its premiums on fitness information, fit individuals that choose to disclose their information might

---

40 The transaction costs associated with the allocation of property rights can have second order effects. For example consent-based that creates consumer costs of providing the consent may make it harder firm firms to enter. See Campbell et al. (2015).

41 See, e.g., Grossman (1981) and Milgrom (1981).

42 Of course, it is possible that sellers would avoid conditioning price on consumer information in any event for reasons given earlier. If the seller cannot commit to refrain from price discrimination, the impact of price discrimination and consumers' incentives to hide or reveal their information interact in complex ways. See Belleflamme and Vergote (2016) and Koh et al. (2017).

receive lower prices than less fit individuals even if the latter choose not to disclose their fitness information.[43]

D. Externalities Across Consumers

The literature on privacy discussed to this point assumes the absence of explicit externalities across consumers from the use of consumer information. However, the use of information provided by one set of consumers could have spillover effects on other consumers who choose not to share their data. This may be more likely with the advent of artificial intelligence and big data analytics,[44] which improve the ability to make inferences about a set of consumers that do not share certain data by analyzing data from another set of similar consumers that do share their data.

Choi et al. (2019) examine a model where sharing of information by one group has significant external effects on another (note, they do not detail what these externalities are). Not surprisingly, they find that if there are significant negative externalities across consumers, sellers may collect and use too much data relative to the social optimum. On the other hand, if there are positive externalities, sellers may collect and use too little data. Standard, well-known economic arguments support these theoretical points.

However, the nature and scope of non-privacy externalities across consumers is not well-known, and Choi et al. provide no evidence of the magnitude or sign of these effects. It is not obvious that information gleaned from data shared by some consumers inflicts net negative externalities on other consumers, which is what is required in Choi et al.'s model to predict harm from non-privacy and motivate the need for privacy regulation. It is also seems clear that consumer information has scientific value, e.g., for use in scientific studies that create significant benefits for society.

E. Myopic Consumers

The theoretical literature summarized above shows that consumer responses to non-privacy give sellers incentives to design privacy policies to encourage consumers to participate in interactions with them—online purchases, web surfing, participation in social networks, etc. The alignment between seller and consumer interests is not perfect, but for some aspects of privacy, the alignment can be quite close, and the theoretical literature does not predict a systematic mis-alignment one way or another.

Some research suggests that in some cases consumer myopia or behavioral considerations (bounded rationality) may play an important role in privacy. For example, Athey et al. (2017) report results from a field experiment documenting distortions in consumer behavior that suggests that consumers' ability to safeguard their privacy through their behavior in a notice and choice regime like that in force in the U.S. is limited. As we discussed above, intertemporal price discrimination is not profitable in Acquisti and Varian's baseline model when consumers

---

[43] See https://www.washingtonpost.com/business/2018/09/25/an-insurance-company-wants-you-hand-over-your-fitbit-data-so-they-can-make-more-money-should-you/?utm_term=.6904ebf99034, retrieved on 05/30/2019.

[44] For a review of the impacts of artificial intelligence on the privacy debate, see Tucker (2019).

are rational, but it is when they are myopic, and this can lead to price discrimination that harms consumers in some circumstances. And the demand attribute model discussed above assumes rational, forward looking consumers.

It might be tempting to argue based on this evidence that behavioral considerations such as myopic consumers increase the need for privacy regulation. However, just as the use of consumer information that occurs due to consumer myopia might harm consumers in some cases, it might also benefit them in others. The welfare effects of a relationship attribute like privacy when consumers are myopic depend on the distribution of benefits and harms from the use of consumer information when consumer myopia is present in the market. As far as we know, this question has not been addressed in the literature.

The simple analytics of this question can be examined in the demand attribute model introduced above. A set of heterogenous consumers each demand one unit of a product, and the value they place on the interaction in which the purchase is made depends on privacy level. Figure 3 illustrates this situation, where D_0 is the true demand when consumers accurately compute their willingness to pay at privacy level t_0, and D_1 is the true demand when consumers accurately compute that the privacy level is t_1. The height of the demand curve is the willingness to pay for the product of the marginal consumer at the privacy level corresponding to that demand curve. For example, consumer q' is willing to pay p' for the product at both privacy levels, while consumers to the left of q' are willing to pay more at privacy level t_1 and consumers to the right of q' are willing to pay more at privacy level t_0. The assumption that privacy rotates demand about a given point reflects the idea that some consumers benefit from greater privacy while others are harmed.

Figure 3: Privacy Incentives with Consumer Myopia

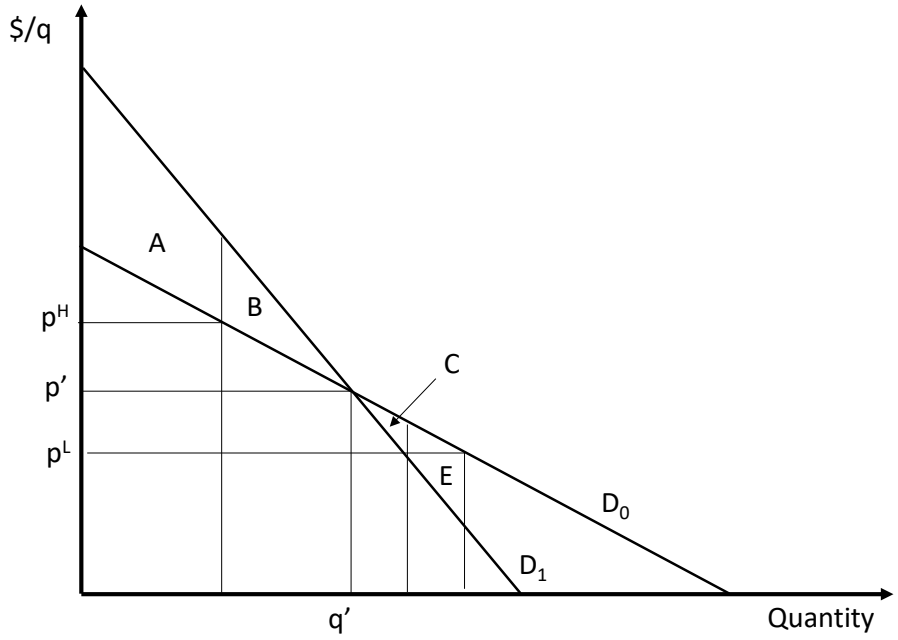

To understand the impact of privacy on naïve or myopic consumers, suppose consumers behave as though the privacy level offered will be t_0 and are oblivious to the seller's desire to choose t_1 when consumers behave that way. The effect of the seller's choice of t_1 on consumers depends on whether price is higher or lower than the price p' at which the demand curves cross. If price exceeds the p' (e.g., p^H), then the seller's choice of t_1 benefits consumers relative to the choice t_0 because the true demand of consumers that purchase the product shifts out. The increase in consumer surplus is area A in Figure 3. Observe that fewer consumers purchase than would be the case if their behavior reflected the true privacy choice. And the consumer surplus of those that purchase rises when the seller chooses t_1 instead of t_0. However, the firm will have too little incentive to offer t_1 rather than t_0 because it will not elicit the demand response if consumers do not understand the effects of changing the privacy policy on their surplus.

On the other hand, if price is below p', (e.g., p^L), then the seller's choice of t_1 rather than t_0 benefits some consumers and harms others. Those with valuations above p' benefit by the area A+B, while those with valuations below p' are harmed by area C+E. The net benefit to consumers from the choice of t_1 rather than t_0 is area A+B-C-E, which is positive as drawn but could be positive or negative. If it is positive, the firm will have too little incentive to choose t_1 rather than t_0, and if it negative, the firm will have too much incentive to do so.

In developing this example, we did not indicate whether t_0 or t_1 involves greater privacy. We did not do so for a reason. As we observed in our early discussion of demand attributes, privacy may rotate demand clockwise or counter-clockwise depending on the distribution of consumer valuations for the product and the benefits and costs of privacy. In this example as drawn, if t_1 is a higher level or privacy than t_0, then myopia causes the firm to under-supply privacy. On the other hand, if t_1 is a lower level of privacy than t_0, then myopia causes the firm to over-supply privacy. The effects of privacy in the naïve consumer case generally depend on whether demand rotates clockwise or counter-clockwise, whether price is above or below the valuation of the consumer who is indifferent between higher and lower privacy, and other factors.

A complete analysis of the effects of privacy in the demand attribute model with myopic consumers is beyond the scope of this paper. However, the example above in combination with research on the price discrimination aspects of privacy when consumers are myopic indicate that consumer myopia does not by itself provide a basis for concluding that markets systematically under-supply privacy.

4. ECONOMIC LITERATURE ON PRIVACY -- EMPIRICAL RESEARCH

In the preceding section, we framed privacy as a relationship attribute that comes bundled with interactions that have other aspects, such as the sale of a product, web browsing, social networking, etc. This framing largely encompasses the theoretical economic literature (except that which is focused on externalities *across* consumers). Because the value to consumers of interactions depend on privacy, firms have incentives to develop and disclose privacy policies as well as make commitments to have them enforced. The analyses of the different types of relationship attributes surveyed above help to identify specific conditions under which interaction participants may over-supply or under-supply privacy relative to a social ideal. These conditions generally depend in complex ways on the distribution of consumer preferences for

products and privacy and the opportunity costs to sellers of privacy. The conditions can also depend in complex ways on the extent of consumer myopia.

The current federal regulatory regime in the U.S. is one of notice and choice (with the FTC regulating unfair and deceptive practices within the notice and choice framework) plus specific privacy regulations for the healthcare and financial industries and children's online activities. We are not aware of empirical work attempting to assess the broad effects of additional privacy regulation beyond notice and choice on social welfare. Instead, most of the empirical literature focuses on specific markets or issues, and much of it focuses on sub-issues such as the benefit side or the cost side of what is ultimately a cost-benefit analysis and the extent of consumer myopia regarding privacy choices. That is, the literature tends to identify specific harms and benefits from privacy, sometimes evaluating welfare effects in specific contexts, and the departure of consumers from rational behavior, but without fully assessing the tradeoffs relevant for assessing broad privacy regulation. This type of research is clearly useful for understanding specific issues, but as we will see, the empirical research has not developed to the point where it provides a good basis for determining the efficacy of broad regulation.

We divide our survey of the existing empirical literature into two parts: (A) research focusing on the impact of privacy or security decisions or policies on one or more aspects of market performance, and (B) research on behavioral economics relating to privacy.

A. Market Impact of Privacy, Security, and Regulation

The EU Privacy and Electronic Communications Directive (2002/58/EC—sometimes referred to as the "e-Privacy Directive") on online advertising limited the ability of companies to track users and use their data to target advertising. Goldfarb and Tucker (2011) found that "banner ads have experienced, on average, a reduction in effectiveness of 65% in terms of changing [consumers] stated purchase intent."[45] Tucker (2012), reporting on the same study, concludes that "on the basis of this evidence, it is reasonable to say that privacy regulation could have sizable effects for the advertising-supported internet."[46] This evidence identifies one effect a specific privacy regulation—the impact of limiting targeted advertising—relevant for a welfare analysis of such a policy, but it does not evaluate the overall impact on consumers: some consumers may have been harmed by the regulation (e.g., less successful matching) while some may have benefitted (e.g., fewer unwanted ads). A welfare analysis of the effects of the e-Privacy Directive would have to develop a more complete model of decision-making, determine the impact on consumers, and weigh the two in the model.

Tucker (2014) conducted a field experiment that measured how consumers responded to personalized ads on the internet when they were given more control over how firms will handle their privacy matters. She found that giving consumers more control over how their data is used increased click rates in response to personalized ads. This is consistent with the demand attribute framework—changes in privacy policy impact consumer participation in interactions, and profit-maximizing firms would take this into account in developing their privacy policies but does not

45 Goldfarb and Tucker (2011), p. 57.

46 Tucker (2012), p. 266.

address the question of whether firms supply too much or too little privacy relative to some welfare benchmark.

Chiou and Tucker (2017) studied whether the length of time that search engines retain their server logs affected the accuracy of subsequent searches. They find little evidence that reducing the length of storage of past search engine searches affected the accuracy of the search. They interpret this finding as suggesting that limits on data retention such as that imposed by the "right to be forgotten" is less costly than is sometimes supposed. This is one cost that would factor into a welfare analysis of privacy regulations.

The Do Not Call ("DNC") registry in the U.S. was created by the FTC in 2003 in response to consumer complaints about unwanted telemarketing calls. This registry allowed consumers to register their phone numbers with the FTC to indicating that they did not want to receive calls from telemarketers. Based on the economic theory of unravelling discussed earlier, one might expect telemarketers to infer that consumers who did not register were less averse on average to receiving telemarketing calls and more likely to make a purchase based such a call. In turn, one might expect that this would cause telemarketers to increase their marketing efforts to those not on the list. Goh et al. (2015) provide evidence that this in fact occurred—after the DNC registry was implemented, members signed up, and marketing calls increased. This evidence does not provide clear welfare implication, however. Arguably, those who registered benefitted, at least in expectation, but those who did not register and did not want an increase in the number of calls may have been harmed.

Recent empirical work on aspects of privacy in the healthcare industry establishes the importance of privacy as a demand attribute in healthcare and provides evidence consistent with benefits of privacy regulation in that industry. Adjerid et al. (2016) study consumer participation in Health Information Exchanges ("HIEs"), which are designed to foster coordination of patient care across the U.S. healthcare systems to generate better health outcomes for patients. They found that among all states with laws creating incentives for the creation of HIE exchanges, only states that included consent requirements (requirements that the exchange give patients the right to consent to share their information) saw a net increase in operational HIEs, and those states reported decreased levels of privacy concern relative to HIEs in states with other legislative approaches. This suggests that privacy regulation increased the formation of HIEs relative to a case without regulation. If one assumes that HIE's increases the quality of patient care without raising prices, the regulation would benefit patients. This is consistent with (but does not establish) the prediction of demand attribute model that sellers may under-supply privacy when valuations for the interaction and privacy are highly positively correlated.

Miller and Tucker (2018) compare three approaches taken by states regarding their privacy laws related to genetic testing: (i) notification regarding privacy risks on the part of the individual and their consent to those risks; (ii) explicitly restricting the use of genetic data by health insurers, employers, or providers of long-term life care or insurance; and (iii) limiting re-disclosure without the consent of the individual or defining genetic data as the "property" of the individual. They find that that an approach that gives users control over redisclosure encourages the spread of genetic testing, whereas an approach of notice and consent deters individuals from obtaining genetic tests. If one assumes that more genetic testing is good for consumers (e.g., because it helps obtain better health outcomes) and that the increase in genetic testing was not

accompanied by large price increases, Miller and Tucker's findings would imply that stronger privacy regulation regarding genetic testing benefited consumers. This is consistent with (but does not establish) the prediction of demand attribute model that sellers may under-supply privacy when valuations for the interaction and privacy are highly positively correlated.

B. Behavioral Issues

Another strand of the empirical literature has focused on how users react to privacy-related incentives from a behavioral economics standpoint. Acquisti et al. (2013) conducted a field experiment to investigate how much consumers were willing to pay for privacy, with the question posed in two different ways: (i) they were asked whether they were willing to pay a given amount to protect their data; and (ii) they were asked whether they would accept a given payment in exchange for their data. The authors found that the default case matters, as subjects were "more likely to reject cash offers for their data if they believed that their privacy would be, by default, protected than if they did not have such a belief."[47] The authors suggest that their findings raise the question of whether notice-and-consent solutions (or similar self-regulatory approaches) are sufficient to guarantee consumers' privacy protection.

However, what Acquisti et al. (2013) measure is a framing effect that is common in experiments. The authors do not assess whether under either framing of the question the incentive of the seller to provide privacy is too high or too low and therefore cannot say based on their findings whether privacy would be over- or under-supplied in their example.

Athey et al. (2017) also found that the context matters when individuals make privacy decisions. They conducted an experiment where participants (students) were asked to select a digital wallet from a set had different privacy settings. They found that participants' choices depended on the order in which the choices were presented and that "small frictions in navigation costs surrounding privacy choices can have large effects in terms of technology selected, even in the presence of transparent information about the privacy consequences of those choices."[48] They state: "Even in an environment where students could maximize privacy in a way that was consistent with their stated preferences, the ordering - potentially combined with inattention - seemed to drive many of the participants' decisions."[49]

Although Athey et al. motivate the paper as presenting evidence about "a variety of distortions in the notice and choice process," they observe:

> "[O]ur empirical results can be used to support two highly contrasting stances towards privacy protection. The first policy stance is that our results could be taken as suggesting that consumers' behavior regarding privacy—as disclosed through their stated privacy preferences in our surveys—is slanted away from their actual normative preferences. This might suggest that consumers need to be protected from themselves, above and beyond the protection given by a notice and choice regime, to ensure that small incentives, search

47 Acquisti et al. (2013), p. 252.

48 Athey et al. (2017), p. 2.

49 *Id*, p. 14.

costs or misdirection are not able to slant their choices away from their actual normative preferences. Firms that want to implement a privacy strategy that is effective in the long-run, need to be aware of these distortions to avoid being penalized by consumers later.

The second policy stance our results document is that there is a disconnect between stated privacy preferences and revealed preference, but that revealed preference is actually closest to the normative preference. When expressing a preference for privacy is essentially costless as it is in surveys, consumers are eager to express such a preference, but when faced with small costs this taste for privacy quickly dissipates. This would suggest that basing privacy protection on stated preference data regarding privacy expressed in surveys is misguided, especially since such policies have been documented to impose costs on firms." (Athey et al., 2017, p. 4)

Two points are notable about Athey et al.'s description of implications of their results. First is the recognition that in the "long run," firms would have to be aware of the distortions in consumer decision-making to "avoid being penalized by consumers later." This is consistent with the firm's privacy policy being a demand attribute in the long run, in which case a welfare analysis of the firm's privacy choice depends on the factors discussed in Section 3.A above. Second, Athey et al. recognize that the revealed preference for the degree of privacy offered in actual markets may be a better measure of their preferences for privacy than in an experiment.

Jia et al. (2019) studies the impact of the European Union's General Data Protection Regulation ("GDPR"; see Section 5.B below) on new technology venture investment in Europe. Their empirical strategy is a difference-in-differences approach that compares the change in such investment in Europe with the change in the U.S. after the GDPR was rolled out in 2018. They find that the GDPR reduced the amount raised to fund European investment deals, the number of deals, and the amount raised per individual deal. This finding indicates that the GDPR has reduced the expected profits of new technology ventures in Europe relative to new ventures in the U.S. The findings do not address the welfare effects of the GDPR because it does not measure the impact of the policy on consumers.

5. POLICY ISSUES

A. Current Policy, U.S. – Notice and Choice Combined with Industry-Specific Regulation

The current approach to federal privacy policy in the U.S. is largely through "notice and choice," with a few specific regulations targeted to specific issues or industries. Notice and choice involves firms giving notice as to their privacy policies and consumers choosing to interact with firms under their stated policies or not. Notice and choice is assumed as the baseline case in most of the theoretical literature on privacy that we have discussed.

The theoretical literature implies that targeted privacy regulation may be warranted when the interaction that conveys information tends to be valued most by consumers that value privacy most and the correlation between these valuations is high, or when consumer myopia takes a form that is likely to lead to too little privacy protection. U.S. regulators have imposed targeted privacy regulations in some markets that meet these criteria.

5.A.1 The Health Insurance Portability and Accountability Act ("HIPAA")

The Health Insurance Portability and Accountability Act, passed by Congress in 1996, requires certain participants in the healthcare industry to standardize healthcare-related data, including patients' health data, and to adopt certain privacy requirements to protect patients' data. In particular, the law prevents healthcare providers, health plans, and clearinghouses from disclosing health information except as authorized by the patient or specifically permitted by the regulation.[50]

As discussed earlier, the demand attribute framework indicates that a firm is more likely to under-supply privacy protection the greater the correlation between the consumer's demand for the interaction and privacy. A possible justification for HIPAA that is consistent with the demand attribute framework is that very sick patients that place high value on healthcare may also place high value on keeping their healthcare information private.

5.A.2 Children's Online Privacy Protection Act ("COPPA")

The Children's Online Privacy Protection Act, passed in 1998, requires websites to obtain consent from parents before collecting the data of children under 13 years of age or sharing it with third parties.[51] COPPA is consistent with a notice and choice approach, but it assigns choices for children to parents under the assumption that children cannot make an informed decision about whether to use a website and share their data. A justification for this regulation is that children are more likely than their parents to be myopic in ways that are likely to lead to too little privacy in a notice and choice environment.

5.A.3 Gramm Leach Bliley Act

The Gramm-Leach-Bliley Act, passed in 1999, limits the ability of certain financial institutions to disclose certain consumer financial information to nonaffiliated third parties.[52] Financial institutions must notify their customers about their information-sharing practices and tell consumers of their right to opt-out if they don't want their information shared with certain nonaffiliated third parties. In addition, any entity that receives consumer financial information from a financial institution may be restricted in its reuse and redisclosure of that information.

---

50 See "The HIPAA Privacy Rule," U.S. Department of Health & Human Services, available at https://www.hhs.gov/hipaa/for-professionals/privacy/index.html?language=en, retrieved on 06/27/2019. The law permits disclosure of protected health information to law enforcement officials for law enforcement purposes as required by law and disclosure to certain parties to facilitate treatment, payment, or health care operations without a patient's express written authorization. Any other disclosures of PHI require the covered entity to obtain written authorization from the individual for the disclosure.

51 See FTC (2002), *Protecting Children's Privacy Under COPPA: A Survey on Compliance*, available at https://www.ftc.gov/sites/default/files/documents/rules/children%E2%80%99s-online-privacy-protection-rule-coppa/coppasurvey.pdf, retrieved on 06/27/2019.

52 See FTC (2002). How to Comply with the Privacy of Consumer Financial Information Rule of the Gramm-Leach-Bliley Act, p. 2.

Intuition suggests that financial information may be another area where consumers that place high value on interactions may also place high value on keeping the information private. For example, a person with high credit card debt may place high value on the use of credit but may also have strong preferences to keep their debt confidential.

B. Europe – General Data Protection Regulation (GDPR)

The Europeans have taken a different approach to the regulation of privacy than the U.S. The EU General Data Protection Regulation ("GDPR") imposes strict new rules on entities that collect and control data ("data controllers"), one of which is to obtain consent from individuals before using their data.[53] In addition, the GDPR gives consumers the "right to be forgotten" by having their data erased (Article 17) and the right to "data portability," i.e., to have their data transferred to other data controllers (Article 20).[54]

C. Should the U.S. Adopt The European Model?

The intuitive appeal of GDPR is understandable—it gives consumers significant control over their data. If the transaction costs of tailored privacy were low, this transfer might be interpreted as giving the consumer privacy property rights that they might use to collect a larger share of the benefits from efficient, Coasian, privacy policies with sellers. However, consumers and firms typically do not negotiate privacy policies in exchange for cash because the transaction costs are too high.[55] The correct interpretation of the GDPR is that it imposes through regulation aspects of the privacy policy menu that firms must offer to consumers, and it establishes a default in the event consumers fail to make choices that limit firms' use of their data.

The issue is whether such regulation is warranted. The underlying assumptions are (i) the notice and choice regime systematically under-supplies privacy, and (ii) the informed consent

---

53 For a brief summary see Choi, J., D. Jeon, and B. Kim (2019). "Privacy and Personal Data Collection with Information Externalities," Journal of Public Economics, 173: 113-124.

54 Article 17 applies when the data "are no longer necessary in relation to the purposes for which they were collected or otherwise processed." Article 20 states that "[t]he data subject shall have the right to receive the personal data concerning him or her, which he or she has provided to a controller, in a structured, commonly used and machine-readable format and have the right to transmit those data to another controller without hindrance from the controller to which the personal data have been provided[.]" See Chapter III in GDPR, Official Journal of the European Union, https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX:32016R0679&from=EN.

55 Companies are in the business of offering consumers a menu of options regarding the treatment of their information, but consumers are not in the business of charging the online sites they visit for the use of their information. This does not mean that consumers could not develop such procedures, but it seems likely that effectively requiring consumers to become entrepreneurs in the businesses selling information about themselves to control it would be costly endeavor for many individuals. In fact, it seems likely that the transaction costs would sufficiently high for some consumers that they would fail to set up a mechanism to do this, in which case, sellers would not have rights to use their information even in ways that benefits the consumer.

regime of the GDPR rectifies the problem. But neither economic theory nor empirical research indicates that these assumptions are likely to hold.

There is a risk that such regulation could create significant harm. Early evidence consistent with this potential is the Jia et al. (2019) study discussed above, which finds that the GDPR reduced investment in new technology ventures in Europe. Moreover, evidence discussed above indicates that privacy has aspects of a demand attribute. Economic theory indicates that the incentives of firms that choose privacy policies to create valuable interactions for consumers may be quite well aligned with the interests of consumers, at least to a first-order approximation, and privacy regulation would break this alignment. Simulations across a range of market scenarios involving privacy indicate that the imposition of requirements for stronger privacy protection in the demand attribute framework could reduce consumer and total welfare significantly. Price discrimination motives for non-privacy may exist, especially when consumers are myopic, but the welfare effects of privacy that prevent such discrimination are ambiguous. And although the literature is undeveloped, the implications of consumer myopia on firms' incentives to supply privacy appear to be ambiguous as a matter of economic theory.

Across-consumer externalities created by consumer information raise another issue that has not received careful attention in the literature. A firm's privacy policy in a notice and choice regime is more likely to internalize across-consumer externalities than consumers' own privacy choices in an informed consent regime as under the GDPR. The reason is that firms internalize some (or even all) of the impact on consumers of their privacy choices, whereas individual consumer choices ignore the impact on other consumers. To the extent that across-consumer externalities from non-privacy are positive—e.g., the use of consumer data for research can benefit all consumers—an informed consent regime may under-supply information, i.e., over-supply privacy.

Empirical research has not addressed the core question—does the notice and choice approach yield too little privacy protection generally? Until research exists providing answers to this key question, it is premature to consider a large change to a market institution that has evolved to meet consumer demand and which economic theory predicts is likely to perform well in most circumstances.

References

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman (2016), "The Economics of Privacy," *Journal of Economic Literature*, 54(2):442-492.

Acquisti, Alessandro and Hal R. Varian (2005), "Conditioning Prices on Purchase History," *Marketing Science*, 24(3):367-81.

Acquisti, Alessandro, Leslie K. John, and George Loewenstein (2013), "What Is Privacy Worth?" *The Journal of Legal Studies*, 42(2):249-74.

Adjerid, Idris, Alessandro Acquisti, Rahul Telang, Rema Padman, Julia Adler-Milstein (2016), "The Impact of Privacy Regulation and Technology Incentives: The Case of Health Information Exchanges," *Management Science*, 62(4):1042-1063.

Athey, Susan, Christian Catalini, and Catherine Tucker (2017), "The Digital Privacy Paradox: Small Money, Small Costs, and Small Talk," *NBER working paper* 23488.

Belleflamme, Paul and Wouter Vergote (2016), "Monopoly Price Discrimination and Privacy: The Hidden Cost of Hiding," *Economic Letters*, 149:141–144.

Campbell, James, Avi Goldfarb, and Catherine Tucker (2015), "Privacy Regulation and Market Structure," *Journal of Economics and Management Strategy*, 24(1):47-73.

Carrascosa, Juan M., Ruben Cuevas, Vijav Erramilli, Nikolaos Laoutaris, and Jakub Mikians (2015), "I Always Feel Like Somebody's Watching Me: Measuring Online Behavioural Advertising," *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies, CoNEXT'15* (ACM, New York), 1–13.

Casadesus-Masanell, Ramon and Andres Hervas-Drane (2015), "Competing with Privacy," *Management Science*, 61(1):229-246.

Chen, Yuxin and Ganesh Iyer (2002), "Research Note on Consumer Addressability and Customized Pricing," *Marketing Science*, 22(2):197-208.

Chiou, Lesley and Catherine Tucker (2017), "Search Engines and Data Retention: Implications for Privacy and Antitrust," *NBER working paper* 23815.

Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim (2019), "Privacy and Personal Data Collection with Information Externalities," *Journal of Public Economics*, 173:113-124.

Coase, R. (1960), "The Problem of Social Cost," *Journal of Law and Economics*, 3:1-44.

Conitzer, Vincent, Curtis Taylor, and Liad Wagman (2012), "Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases," *Marketing Science*, 31(2):277-292.

Cooper, James C., Luke Froeb, Daniel P. O'Brien, and Steven Tschantz (2005), "Does Price Discrimination Intensify Competition? Implications for Antitrust," *Antitrust Law Journal*, 72(2):327-373.

Corts, Kenneth S. (1998), "Third-degree Price Discrimination in Oligopoly: All-out Competition and Strategic Commitment," *RAND Journal of Economics*, 29(2):306-323.

De Corniere, Alexandre and Rodrigo Montes (2017), "Consumer Privacy and the Incentives to Price-Discriminate in Online Markets," *Review of Network Economics*, 16(3):291-305.

Dorfman, Robert, and Peter O. Steiner (1954), "Optimal Advertising and Optimal Quality," The *American Economic Review*, 44(5):826-836.

Farrell, Joseph (2012), "Can Privacy be Just Another Good?" *Journal on Telecommunications & High Technology Law*, 10:251-264.

Fudenberg, Drew and Jean Tirole (2000), "Customer Poaching and Brand Switching," *RAND Journal of Economics*, 31(4):634-657.

Goh, Khim-Yong, Kai-Lung Hui, and Ivan P. L. Png (2015), "Privacy and Marketing Externalities: Evidence from Do Not Call," *Marketing Science*, 61(12):2982-3000.

Goldfarb, Avi and Catherine Tucker (2011), "Privacy Regulation and Online Advertising," *Management Science*, 57(1):57-71.

Grossman, Sanford J. (1981), "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 24(3):461-483.

Hermalin, Benjamin and Michael Katz (2006), "Privacy, Property Rights and Efficiency: The Economics of Privacy as Secrecy," *Quantitative Marketing and Economics*, 4(3):209-239.

Ichihashi, Shota (2018), "Online Privacy and Information Disclosure by Consumers," mimeo, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3112905.

Jia, Jian, Ginger Zhe Jin, and Liad Wagman (2019). "The Short-Run Effects of GDPR on Technology Venture Investment," mimeo, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3278912.

Johnson, Justin P., and David P. Myatt (2006), "On the Simple Economics of Advertising, Marketing, and Product Design," *American Economic Review*, 96(3):756-784.

Jullien, Bruno, Yassine Lefouili, and Michael Riordan (2018), "Privacy Protection and Consumer Retention," Working Paper TSE-947, *Toulouse School of Economics*.

Koh, Byungwan, Srinivasan Raghunathan, and Barrie Nault (2017), "Is Voluntary Profiling Welfare Enhancing? *MIS Quarterly*, 41(1):23-41.

Lenard, Thomas M. and Paul H. Rubin (2009), "In Defense of Data: Information and the Costs of Privacy," mimeo, *Technology Policy Institute*.

Mikians, Jakub, Lazlo Gyarmati, Vijay Erramilli, and Nikolaos Laoutaris (2012), "Detecting Price and Search Discrimination on the Internet," *Proceedings of the 11th ACM Workshop on Hot Topics in Networks* (ACM, New York), 79–84.

Mikians, Jakub, Lazlo Gyarmati, Vijay Erramilli, and Nikolaos Laoutaris (2013), "Crowd-Assisted Search for Price Discrimination in E-commerce: First results," *Proceedings of the 9th ACM Conference on Emerging Networking Experiments Technology* (ACM, New York), 1–6.

Milgrom, Paul R. (1981), "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, 12(2):380-391.

Miller, Amalia R. and Catherine Tucker (2018), "Privacy Protection, Personalized Medicine and Genetic Testing." *Management Science*, 64(10): 4648-4668.

Montes, Rodrigo, Wilfried Sand-Zantman, and Tommaso Valletti (2019), "The Value of Personal Information in Online Markets with Endogenous Privacy," *Management Science*, 65(3):1342-1362.

O'Brien, Daniel. P. and Doug Smith (2014), "Privacy in Online Markets: A Welfare Analysis of Demand Rotations," *FTC Bureau of Economics Working Paper*, 323.

Posner, Richard A. (1978), "The Right of Privacy," *Georgia Law Review*, 12(3):393-422.

Posner, Richard A. (1981), "The Economics of Privacy," *American Economic Review*, 71(2):405-409.

Ramadorai, Tarum, Antoine Uettwiller, and Ansgar Walther (2019), "The Market for Data Privacy," mimeo, available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3352175.

Shapiro, Carl and Hal R. Varian (1997), "Government Information Policy," mimeo, presented at Highlands Forum, Department of Defense, June 8, 1997, Washington, DC. Sponsored by the Office of the Assistant Secretary of Defense (Command, Control, Communications and Intelligence).

Shen, Qiaowei and J. Miguel Villas-Boas (2018), "Behavior-Based Advertising," *Management Science*, 64(5):2047–2064.

Spence, A. Michael (1975), "Monopoly, Quality, and Regulation," *Bell Journal of Economics*, 6(2):417-429.

Stigler, George J. (1980), "An Introduction to Privacy in Economics and Politics," *Journal of Legal Studies*, 9(4):623-44.

Stokey, Nancy L. (1979), "Intertemporal Price Discrimination," *The Quarterly Journal of Economics*, 93(3):355-371.

Stole, Lars A. (2007), "Price Discrimination and Competition," *Handbook of Industrial Organization* 3:2221-2299.

Taylor Curtis and Liad Wagman (2014), "Customer Privacy in Oligopolistic Markets: Winners, Losers, and Welfare," *International Journal of Industrial Organization*, 34:80-84.

Taylor, Curtis R. (2004), "Consumer Privacy and the Market for Customer Information," *RAND Journal of Economics*, 35(4):631-50.

Tucker, Catherine (2012), "Empirical Research on the Economic Effects of Privacy Regulation," *Journal of Telecommunications & High Technology Law*, 10:265-271.

Tucker, Catherine (2014), "Social Networks, Personalized Advertising, and Privacy Controls," *Journal of Marketing Research*, 51:546-62.

Tucker, Catherine (2019), "Privacy, Algorithms and Artificial Intelligence," in *The Economics of Artificial Intelligence: An Agenda*, Ajay Agrawal, Joshua Gans, and Avi Goldfarb, eds., University of Chicago Press.

Varian, Hal R. (1989), "Price discrimination," in *Handbook of Industrial Organization* 1:597-654.

Varian, Hal R. (1996), "Economic Aspects of Personal Privacy," mimeo, *U.C. Berkeley*, available at https://pdfs.semanticscholar.org/0810/927d65f5ab358cd4c1bff34beff6198e78a5.pdf. Subsequently published in Privacy and Self-Regulation in the Information Age. Washington, DC: US Department of Commerce, National Telecommunications and Information Administration, 1997; Cyber Policy and Economics in an Internet Age, pp. 127-137. Springer, Boston, MA, 2002 and Internet Policy and Economics, pp. 101-109. Springer, Boston, MA, 2009.

Villas-Boas J. Miguel (1999), "Dynamic Competition with Customer Recognition," *RAND Journal of Economics*, 30(4):604-631.

Villas Boas, J. Miguel (2004), "Price cycles in markets with customer recognition," *RAND Journal of Economics*, 35(3):486-501.