NBER WORKING PAPER SERIES

THE ECONOMICS OF DIGITAL PRIVACY

Avi Goldfarb
Verina F. Que

Working Paper 30943
http://www.nber.org/papers/w30943

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2023

This paper is forthcoming in the Annual Review of Economics. It will be available at https://doi.org/10.1146/annurev-economics-082322-014346. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w30943

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

The Economics of Digital Privacy
Avi Goldfarb and Verina F. Que
NBER Working Paper No. 30943
February 2023
JEL No. L51,L86

ABSTRACT

There has been increasing attention to privacy in the media and in regulatory discussions. This is a consequence of the increased usefulness of digital data. The literature has emphasized the benefits and costs of digital data flows to consumers and firms. The benefits arise in the form of data-driven innovation, higher quality products and services that match consumer needs, and increased profits. The costs relate to intrinsic and instrumental values of privacy. Under standard economic assumptions, this framing of a cost-benefit tradeoff might suggest little role for regulation beyond ensuring consumers are appropriately informed in a robust competitive environment. The empirical literature thus far has focused on this direct cost-benefit assessment, examining how privacy regulations have affected various market outcomes. However, an increasing body of theory work emphasizes externalities related to data flows. These externalities, both positive and negative, suggest benefits to the targeted regulation of digital privacy.

Avi Goldfarb
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
and NBER
agoldfarb@rotman.utoronto.ca

Verina F. Que
Rotman School of Management
University of Toronto
105 St. George Street
Toronto, ON M5S 3E6
f.que@rotman.utoronto.ca

1 Introduction

Privacy is increasingly in the media and the subject of regulatory discussions. This rise of attention to privacy stems from the increased usefulness of data. Digitization has reduced the cost of collection, storage, transmission, and analysis of data (Goldfarb and Tucker (2019)). This, in turn, has led to the expanded use of digital data in decision-making (Brynjolfsson and McElheran (2016)). For consumers, data enables personalized services and products at a much lower cost, which significantly enhances consumer welfare. The use of this data can help firms improve profits and it can help individuals get higher quality products and services that better match their needs.

The use of such data, however, has some negative consequences. Some of these negative consequences are direct. Individuals have an intrinsic distaste for the collection and use of information about them. Firms face direct costs of obtaining and protecting consumer data. Both individuals and firms may also find this information used against them.

Some of these negative consequences arise from externalities. Information about one individual can be informative about another. For instance, Erlich et al. (2018) show that a genetic database needs to cover only 2% of the target population to identify nearly everyone. This negative externality is similar to the data spillovers described in Tucker et al. (2018). When people take photos of their car with geocodes to take note of their parking spot, they can record other people and cars. This information could be used in ways that harm people who did not take photos. These negative externalities generate many of the concerns about digital privacy.

In short, privacy affects economic outcomes.

Privacy is a difficult term to define. In the 19th century, privacy denoted "the right to be let alone", which was recognized as fundamental to human existence and inherent in human nature (Warren and Brandeis (1890)). After World War II, technological developments led to debates on the tradeoffs between privacy and surveillance. Westin (1967) describes privacy as the control over

and safeguard of personal information. Altman (1975) refers to the boundaries between self and others, between private and shared or public features of one's life. More recently, Solove (2008) emphasizes that there are many different facets of privacy including information collection, information processing, information dissemination, and invasion. Nissenbaum (2009), in contrast, depicts privacy to be "the right to appropriate the flow of personal information". The appropriateness of the information flow depends on the context-relative information norms, which determine how information should flow within particular social contexts (Bleier et al. (2020)).

From an economic perspective, the digital privacy literature has focused on the benefits and costs of restricting information flows (see Acquisti et al. (2016) for a comprehensive summary of the history of economic analysis on the trade-offs associated with privacy). Data is fundamentally information (Farboodi et al. (2019)), which is a tool to reduce uncertainty about unknown outcomes. In the following sections, we will use data and information interchangeably, and use the term "digital privacy" to denote a restriction on digital data flows. Because digital information can be copied for near zero marginal cost without degrading quality, it is non-rival in the absence of effort to exclude (Goldfarb and Tucker (2019)). This non-rivalry can generate both positive and negative externalities from data flows. Our economics-focused discussion of digital privacy will move away from the concepts such as "control", "autonomy", "secrecy", and "right to be let alone" in the classical privacy literature. Instead, we concentrate on the trades-off for consumers and firms, both directly and in terms of externalities.

In the next section, we discuss the direct benefits and costs of data flows to consumers. We follow this with a similar discussion for firms. We then turn to externalities, both negative and positive. Next, we turn our attention to a discussion of current regulatory and engineering approaches to privacy, discussing consequences in light of the benefits, costs, and externalities. We conclude with a brief summary and a discussion of open questions.

2 Consumers' Decision Making Under Privacy

2.1 Benefits of Privacy

2.1.1 Valuation of Privacy

Privacy preferences can be divided into two types. Those where privacy itself is treated as an intrinsic right (Warren and Brandeis (1890)), and those where privacy is an instrument for protecting agents from revealing their type in a way that could impact the payoff of their economic activities (Stigler (1980); Posner (1981)). The formal microeconomic models of privacy that started appearing in the early 2000s focused on the latter type, where consumers care about privacy in order to avoid price discrimination in a repeat-purchase scenario (Taylor (2004); Acquisti and Varian (2005); Hann et al. (2008)).

Specifically, online retailers have rich data on purchase history, address, and browsing history. This information could be used for price discrimination. There is an extensive stream of literature that investigates how the extraction and storage of consumer information could be utilized to design personalized pricing and targeted advertising strategies (Villas-Boas (1999, 2004); Taylor (2004); Zhang and Krishnamurthi (2004)). For example, Taylor (2004) provides an early and influential perspective on analyzing consumer privacy and the market for customer information. He defines the value of customer information to be the firm's ability to identify individuals for personalized prices. This can harm or benefit consumers, depending on whether the firm has market power.

Privacy also interacts with advertising technology. Consumer preferences for privacy depend on the sophistication of the firm's advertisement targeting technology. Johnson (2013) finds that, without any intrinsic preference for privacy, consumer preferences for increased targeting are not monotone. Instead, consumer utility has a U shape in the accuracy of targeting. This shape highlights the consumer attitude change towards advertising as technology advances. When targeting is not accurate, incremental improvement in targeting accuracy only leads to further frustration and prompts consumers to block ads. When targeting has improved sufficiently, consumers may

eventually welcome it. The stage of the targeting technology therefore greatly influences consumer attitudes and preferences about privacy.

Building on Becker (1980)'s framework, Lin (2022) models the intrinsic and instrumental components of consumer privacy preferences, and empirically estimates them through a lab experiment. The intrinsic component of "taste" includes consumers' characteristics or behaviors to be kept secret. The instrumental part comes from consumer's anticipated surplus or economic loss from disclosing private information to the firm. In Lin's model, consumer i has a vector of personal variables $D_i = [d_{i1}, d_{i2}, ..., d_{ik}]$ with k types of data, and a sharing decision with equal length S_i: each sharing decision $(s_{i1}, ..., s_{ik})$ indicates whether the individual shares the associated variable. The decision S_i brings an intrinsic privacy cost $C_i = [c_1, c_2, ..., c_k]$, a type-induced payoff from sharing, baseline compensation, and a random utility shock ϵ_{ik} to the consumer's utility specification:

$$
\begin{aligned}
U(\mathcal{S}_i; C_i, D_i) = \sum_k - \underbrace{c_k(X) \cdot s_{ik}}_{\text{intrinsic preference}} \\
+ 1_{inst} \cdot 1_{k\in\{1,2\}} \cdot \underbrace{\beta \cdot p_i \cdot w_k \cdot \hat{E}[d_{ik}|\mathcal{S}_i, D_i]}_{\text{type-induced payoff}} \\
+ \underbrace{\beta \cdot p_i \cdot s_{ik}}_{\text{util from compensation}} + \epsilon_{ik}
\end{aligned}
$$

In the above utility specification, each intrinsic privacy cost c_k could be expanded to a function of observables X in line 1. In line 2, 1_{inst} is an indicator for instrumental privacy concerns, and $1_{k\in\{1,2\}}$ indicates the information sharing decisions influenced by the instrumental incentives. While data type k could include many types, Lin emphasizes and measures two–income and purchase intent. In the type-induced payoff, consumer i has two types of beliefs– *first-order belief* and *higher-order belief*. The consumer's *first-order belief* is w_k –their expected increase in the percentage winning probability for an adjacent, higher type; while $\hat{E}[\cdot]$ is their *higher-order belief*– consumer's expectation of the firm's expectation about his type. In addition, β stands for the marginal utility of monetary rewards and p_i represents the compensation offered for the data. In line 3, the utility

from baseline compensation is $\beta \cdot p_i$, which in Lin's specification is proportional to the number of shared variables s_{ik}.

The paper finds heterogeneous and right-skewed intrinsic preferences of consumers with a mean valuation of 10 dollars for sharing a demographic profile, and a 97.5% quantile of 30 dollars. Whether and how consumers opt to share data depends on the heterogeneity and correlation of the two main components in their preferences. This framework, which explicitly recognizes intrinsic and instrumental aspects of consumer privacy preferences, underlies our interpretation of much of the rest of the literature.

Tang (2019) estimates the value of privacy for online borrowers using large-scale field experiments. Her structural model, by linking individuals' disclosure, borrowing, and repayment decisions, is able to quantify the monetary value of personal data. She shows that individuals value privacy, and measures the intrinsic part of the privacy: the social network ID and employer contact information are valued at 32 dollars, accounting for 8% of the value of a foregone loan.

2.1.2 The Digital Privacy Paradox

Lin (2022), Tang (2019), and others provide evidence that consumers do care about digital privacy. Nevertheless, privacy preferences are context-dependent and have changed over time. For example, Goldfarb and Tucker (2012b) use survey data with 3 million responses from 2001 to 2008 to document that older consumers are more privacy-sensitive than younger consumers and that overall privacy concerns are rising over time. Privacy sensitivity is measured by refusal to provide personal information about income, age, or zip code in a survey. This change in privacy concerns over time appears to be driven by consumers expanding the types of data that are considered privacy. Specifically, consistent with Nissenbaum (2004)'s concept of contextual integrity, privacy concerns in non-personal contexts (e.g. entertainment, consumer-packaged goods) grew more rapidly as cross-context data exchange became more common. Acquisti et al. (2015) also measure

privacy using refusal to provide information in a survey, and similarly demonstrate an increasing concern about privacy over time. Both studies document that privacy concerns grow as consumers are immersed in more sophisticated data-sharing practices when using digital products and services.

Despite these measured benefits of privacy to individuals, and despite evidence of increasing concern for privacy, consumers continue to give out large quantities of personal information. There is often a gap between consumer's stated and revealed privacy preferences. This phenomenon is labeled the "privacy paradox" (Norberg et al. (2007)). Individuals' valuation of privacy is affected by contextual and nonnormative factors. Acquisti et al. (2013) establish a notable gap between individuals' willingness to accept (WTA) and willingness to pay (WTP) in a field experiment. They show that the cash-for-privacy exchange is larger when individuals consider getting money from trading out their data, while small when people pay for privacy. Athey et al. (2017) convincingly establish a digital privacy paradox through three main empirical findings– namely "small money", "small costs", and "small talk". "Small money" recognizes that people are willing to relinquish private data quite easily when they face small incentives though they claim they care about privacy. In the study, people gave up the email addresses of their friends in exchange for a slice of pizza. "Small costs" speaks to the fact that small frictions in navigation costs can efficiently reduce technology adoption, even with consequences transparently presented. People typically did not select a privacy-protecting option from a list of four when the privacy-protecting option was at the bottom of the list. "Small talk" shows that an irrelevant aspect of privacy in the particular context studied, encryption, could provide an illusion of protection and reduced privacy-enhancing behavior.

The privacy literature has suggested several explanations to understand the privacy paradox. Burtch et al. (2015) emphasize consumer ignorance or lack of attention about how data might be used, demonstrating that delayed presentation of privacy policies increases revenue in an online fundraising context. Adjerid et al. (2016) also find that reminders about privacy policies tend to lead individuals to opt out. They point out that privacy reminders reduce the usage of health information exchanges, unless combined with subsidies for adoption.

Related to the role of consumer ignorance in explaining the privacy paradox is the impact of giving consumers the perception of control. Tucker (2014) examines the relationship between users' perception of control over their personal data and the likelihood of them clicking on Facebook's advertisements. As shown in **Figure 1**, reproduced from Tucker (2014), personalized advertising was relatively ineffective before the introduction of policy that increased consumers' perceived control over personal data flows. After the policy was introduced, personalized advertising was nearly twice as effective at attracting users, even though these controls were not directly related to the way the data was used.

[Figure 1 about here.]

Chen et al. (2021) emphasize correlated preferences between benefiting from data flows and the desire to protect privacy. Combining survey and behavioral data on the Alipay platform, one of the largest Chinese online payment and lifestyle platforms with more than 900 million active users as of 2022, they find that users with stronger privacy concerns in the survey tend to give out authorization and use the app services more frequently and extensively. They argue this is driven by an instrumental value of privacy. It is the intense use of digital services by the most active users that creates a stronger preference for privacy. It is an open question whether this correlation is exogenous or whether digital preferences cause increased privacy concerns.

Solove (2021) highlights a methodological explanation. The observed behavior is measured in very specific contexts while self-reported privacy concerns tend to come from general surveys. Thus, the latter may not correlate closely with the former.

The literature overall suggests a privacy paradox, in the sense that individuals claim to care a great deal about how their data is used, but appear to act as if they do not care. A number of explanations have been put forward related to consumer ignorance, correlated preferences between privacy and the benefits of data, and methodological issues.

2.2 Benefits of Data flows to Consumers

Data flows in the digital world bring substantial economic benefits directly to consumers. Consumers enjoy new services such as search engines and recommendation systems, personalized advertising and offers, and targeted products and services. Moreover, customized communication can reduce information overload and assist customers in making informed decisions (Ansari and Mela (2003)). When firms have access to data, prices can fall. For example, Kummer and Schulte (2019) use data from 300,000 apps on the Google Play Store and document that paid apps ask for less consumer data than free apps. Apart from the economic benefits, consumers may experience direct psychological benefits from sharing data. Tamir and Mitchell (2012) discover that human self-disclosure activities, by sharing information with others, engage neural and cognitive mechanisms associated with rewards.

2.2.1 Better Service & Personalization

Open data flows grant firms more information to tailor their products, services, and communications to an individual customer, which we refer to as personalization. Personalization can reduce information overload, which aids consumers in making efficient decisions. By allowing firms to learn their preferences, consumers benefit from reduced customer search costs (Goldfarb and Tucker (2019)), so that the right products and messages could be delivered to the right person, at the right time.

To understand the value of personalization to consumers, Sun et al. (2021) conduct a large-scale field experiment on the Alibaba E-commerce platform, involving a random sample of 555,800 customers. By banning the use of personal data in the homepage recommendation algorithm, they observe a sharp decrease in both customer engagement (clickthrough rate and product browsing) and market transactions (sales volume and amount). Specifically, the customers' clickthrough rate on the recommended products drops immediately by 75%, and the customers' browsing behavior on the homepage was subsequently reduced by 33%. As a result of the two combined effects, purchases fall 81%. The analysis indicates that the value of personalization in e-commerce is large for the whole consumer group. Moreover, it disproportionately benefits newer customers, with less

purchase power, females, and those from developing regions.

In addition, Chan et al. (2022) show the great benefit of expanded credit access from digitally verified data. They document that the better verified data increase average loan origination rate by 35.5%, without substantially raising the interest rates charged on these loans. The effect is especially significant for deep subprime and subprime consumers, with a 146% and a 44% raise in the loan rate respectively. On the lender side, they also enjoy an estimated 19.6% increase in profit from the expanded credit access.

Similar results have been found in healthcare. When data flows are easier, electronic medical record (EMR) adoption is higher and patients benefit. Miller and Tucker (2011) show that the improved monitoring capacity resulting from adopting EMR can reduce neonatal mortality rates. Furthermore, Derksen et al. (2021) find that the introduction of an EMR system to track down HIV patients in Malawi immediately enhances the number of patients actively in care and reduces patient mortality. The use of this system is limited by privacy permissions. Therefore, patients' privacy preferences can inhibit the effectiveness of the EMR system significantly. In this sense, data flows improve patient health.

2.2.2 Price Discrimination and Data Flows

Price discrimination is central to the instrumental value that consumers receive from privacy. However, in equilibrium, under certain circumstances data flows can increase consumer surplus.

Conitzer et al. (2012) recognized that one way for data flows to increase consumer surplus is related to the Coase conjecture (Coase (1972)). Consider a monopolist who can track individual past purchasing patterns in order to price discriminate, and consumers can in turn conceal their personal data at a cost. When maintaining anonymity is costly, the seller has better capability to identify old customers and to price discriminate. Knowing that, consumers will hesitate to make the initial purchase. Anticipating this, the seller is constrained to provide a lower initial price,

which dominates the profit increases arising from future price discrimination. As a result, the seller would prefer to commit to a no-price-discrimination case. When the cost of maintaining anonymity is low, say consumers can freely anonymize themselves, all individuals will choose to do so, resulting in the highest profit for firms. In this "hide and seek" game, Conitzer et al. (2012) therefore provide a distinct perspective to the debate. When privacy is costly for consumers, they can be better off. As such, providing privacy protection can reduce consumer surplus and social surplus when the cost of maintaining anonymity is low.

Absent these dynamic considerations, competition can serve to prevent price discrimination via personalized pricing. While a monopoly firm with access to consumer data can make consumers worse off through the improved match values and more aggressive pricing, Loertscher and Marx (2020) demonstrate that in this setting, if the price is regulated, a reduction in privacy will always benefit the consumer because of the improved matches. In this model, maintaining competition is more important than privacy protection to advance consumer surplus. Furthermore, Miklós-Thal and Tucker (2019) demonstrate that better consumer information can decrease collusion and foster competition in the market. With better demand forecasting, colluding firms face a higher temptation to deviate to a lower price. The overall effect suggests that better forecasting from data flows leads to lower prices and higher consumer surplus.

Data flows may also incentivize pro-social behavior in a way that reduces prices and maximizes consumer surplus. Usage-based car insurance (UBI) can provide individualized price discounts based on driving behavior. Safe drivers self-select into the UBI program to pay a lower premium. The program with its economic incentives can motivate UBI participants to adopt better driving habits (Jin and Vasserman (2021)). UBI has a measurable effect on reducing fatal auto accidents (Reimers and Shiller (2019)) and promoting good driving habits in the long run. The average daily hard-brake frequency dropped by 21% for UBI customers after six months. Young drivers and female drivers show more improvements and benefit more from the program (Soleymanian et al. (2019)). When consumers choose to drive safer to get a lower price, they also increase the total

social welfare to the extent that safe driving generates positive externalities.

Thus, while price discrimination can give rise to instrumental privacy concerns, there are situations open data flows may increase consumer surplus even when those data flows facilitate price discrimination. Ultimately, the impact of data flows on consumer surplus (in an instrumental privacy sense) depends on the particular context.

3 Firms

3.1 Benefits of Data to Firms

Data flows from consumers have created new opportunities for firms. Firms can set personalized pricing, send targeted advertisements, and improve customer relation management. Data has created new markets and, in some circumstances, increased market power.

3.1.1 Personalized Pricing

Consumer data allows personalization. This suggests a potential for first-degree price discrimination (Shapiro and Varian (1998); Smith et al. (2001)). A broad theoretical literature has arisen around suggesting opportunities for firms (and welfare losses for consumers) from digital price discrimination (Acquisti and Varian (2005); Chen and Iyer (2002); Taylor (2004); Hermalin and Katz (2006)). However, as noted in Goldfarb and Tucker (2019), the theoretical literature on the use of data for enabling digital price discrimination appears to be more developed than both the empirical studies and the industry practices. Well-documented examples of first-degree price discrimination are limited.

3.1.2 Target Advertising

Data flows enable targeted advertising, which benefits firms, particularly small firms (Goldfarb (2014)). Unlike personalized pricing, there is a great deal of evidence that firms use data flows to target online advertising to consumers. Targeting allows the firm to endogenously increase differentiation in the market and avoid "wasted" advertising. In other words, targeting improves

advertising effectiveness (Iyer et al. (2005)). For example, Rafieian and Yoganarasimhan (2021) explore the targeted advertising in the mobile in-app advertising context. Their proposed machine learning framework of an efficient targeting policy is estimated to improve the average click-through rate by 66.80% over the current system.

Targeted advertising can affect market power in ways that benefit advertisers (and perhaps consumers) at the expense of firms. Athey and Gans (2010) use a model with local and general outlets to analyze the impact of targeting on the supply and the price of advertising. The authors specify that the improved efficiency of information allocation from targeting leads to a demand increase. However, it may reduce the market power of each individual publisher if advertising space or advertiser capacity are not constrained. Bergemann and Bonatti (2011) model competition between online and offline media, whose main difference lies in the targeting ability based on consumer data. Better targeting improves consumer-product matches, and thus, the social value of advertising. While at the same time, greater targeting amplifies the concentration of firms advertising in each market, which eventually leads to lower advertising prices received by the advertising market due to lack of competition between the advertisers.

3.1.3 Customer Relationship Management

Data enables firms to understand customer needs. Consumers' granular activity data can aid firms in implementing proactive retention strategies. For example, in subscription services, consumer data can reveal with customers are at risk of stopping their subscription. This data can also reveal the marginal impact of different interventions aimed at retaining the customer (Ascarza (2018)).

3.1.4 New Types of Firms

Digital data flows have enabled a market for data in which a new type of firm, the data intermediary, plays an important role. Aside from the direct data flows from consumers, firms also benefit from the third-party data collected, aggregated, and organized by data intermediaries. Bergemann and Bonatti (2019) survey the growing literature on data markets, and emphasize the role of data

intermediaries that sell user information ranging from direct sales of lists of consumers with certain characteristics to indirect sales of data through sponsored search and retargeting. A growing theory literature models these data intermediaries, and their optimal mechanisms in interacting with consumers and advertisers (Bergemann and Bonatti (2015); Bergemann et al. (2018); Yang (2022)).

In addition to being third-party data brokers, the data seller can be a platform with advertisers and consumers on both sides. De Corniere and De Nijs (2016) consider a setting where a platform makes a decision on disclosure or privacy—that is, whether to sell the consumer information gathered from one side of the platform to the advertisers on the other side of it. The model shows that disclosure improves the match between advertisers and consumers but raises prices, even without price discrimination. Disclosing information, in certain conditions, could increase the total profits of the platform and the advertisers, while leaving an information rent to the winning bidder. The results are in line with Bergemann and Bonatti (2015) that it is not optimal for data intermediaries to disclose the finest consumer information to firms since the informational rent is passed on to firms. There are certain conditions in which the intermediary optimizes its profits with an intermediate level of privacy.

3.2 Benefits of Privacy to Firms

To the extent that consumers value privacy and purchase from firms that have strong privacy policies, firms benefit from privacy directly. In addition, firms can benefit through reduced costs associated with data and through market power.

3.2.1 Direct Cost of Data-flows

Collecting and securely storing data is costly. Companies face challenging legal obligations and compliance requirements. They incur costly investments in protecting stored consumer data from malicious access from third parties, such as cyber-attacks. The data protection comes in the forms of API updating, improved firewalls, and vulnerability checking from company-hired hackers. The benefits of data may be small relative to this cost Shy and Stenbacka (2016). For example, Chiou

and Tucker (2017) investigate whether larger quantities of historical data affect the accuracy of subsequent searches, and thus, the firm's ability to maintain market share. Historical data is costly to store, and it creates security risks. They find no empirical evidence that reducing the length of data retention would harm the accuracy of search results. Similarly, Yoganarasimhan (2020) demonstrates that the returns to search personalization are concavely increasing with the length of user history data. In a field study, Neumann et al. (2019) find third-party consumer profiling often to be economically unattractive due to the high additional costs of targeting solutions and their limited accuracy. Bajari et al. (2019) provide both theoretical guidance and empirical support of the countervailing force of diminishing return to data.

3.2.2 Market Power

Although data flows help firms acquire consumers, too much data can reduce a firm's market power. Therefore, it can be beneficial for firms to maintain an intermediate level of privacy for consumers. Choe et al. (2018) consider a two-period model in which two firms compete dynamically through acquiring consumer information in their first-period purchases and offering personalized pricing in the second period. No matter whether product differentiation is exogenously or endogenously chosen, both firms end up worse off compared to when they use simpler pricing strategies or commit to substantial product differentiation. When the use of customer information is solely for pricing, more customer information is typically bad for competing firms because of the intensified competition in the first period of information gathering.

Data-enabled price targeting could intensify price competition, which may hurt the competing seller with better quality. A higher-quality firm can be worse off with personalized pricing (Choudhary et al. (2005)). Casadesus-Masanell and Hervas-Drane (2015) demonstrate that the existence of a data market can also lead a low-quality firm to translate its competitive pressure to consumer data disclosure. When firms have two revenue sources– the sales revenue from products and the disclosure revenue from trading consumer data, the presence of this additional revenue stream from data sales harms the quality-improvement incentives. Consumer data soften the intensity of

competition when consumers are heterogeneous, and firms focus on differentiating their privacy policies.

4 Externalities

Thus far, the focus of this review has been on the direct impacts of data flows and privacy protection on consumers and firms. However, data has externalities. Much of the more recent research on digital privacy has focused on these externalities (Bergemann et al. (2022); Acemoglu et al. (2019); Choi et al. (2019)). Data externalities occur when an individual shares data and the data also reveal information about others. The externalities can be negative or positive.

4.1 Negative Externalities of Data

"Data is the pollution problem of the information age, and protecting privacy is the environment challenge." –Schneier (2015)

The negative externalities of data provide insights into answering the questions: Why do consumers tend to allow some forms of data collection even if they are fully aware of the potential for the data to cause harm? Under what circumstances will firms collect too much personal data, both for their customers and for themselves? What is the role of policy in regulating data flows?

One person's data provides information about others in three possible ways. The first is direct. A person's list of contacts includes information about that person (who their friends are) but also information about their contacts. Similarly, a social media feed contains information about the account owners' posts and likes, but also the posts and likes of their connections, and sometimes the connections of their connections. Thus, an individual's decision to share data may directly affect others. The second and third are indirect. The second way is that individual preferences and behaviors are correlated. Therefore, information about one individual provides probabilistic information about others. The third way lies not in the data itself but in the data generating process. By choosing to withhold information, consumers may reveal their types in the market activities as

well. In this regard, individuals may provide information about the instrumental value of their data.

Choi et al. (2019) emphasize the second type of externality. They model a monopolist platform's data collection. The market equilibrium is characterized by the excessive collection of personal information, in particular, the excessive collection of sensitive information with negative externalities. The primary mechanism is that individuals do not take into account the spillover effects of their data sharing. Even fully-informed agents make individually optimal decisions, the outcome may not be socially efficient. Bergemann et al. (2022) demonstrate that this data externality means that each individual has little incentive to keep their data private, and that selling the data to a data intermediary will yield near zero compensation, even if the individual knows the information can be sold to a firm that seeks to extract her surplus. Acemoglu et al. (2019) model a data market where a monopoly digital platform can trade users' data. They show that there exists an equilibrium where too much data are shared, and the price of data is depressed. This equilibrium is directly due to data's negative externalities, which leads to excessive use of data by platforms and firms. Furthermore, the data prices will no longer reflect consumers' value of data and privacy. In addition to the overuse of data, the externalities also shift surplus from users to the platform and firms.

Ichihashi (2020) examines how a commitment to avoid price discrimination could make the seller better off and the consumer worse off. The mechanism is that under the commitment regime, certain consumers choose greater disclosure that leads to higher prices, which lowers the welfare of other consumers. As consumers fail to internalize this negative effect, they opt for the highest level of disclosure, even though they could benefit from collectively withholding information. In related work, Ichihashi (2021) provides a model of how the firm and consumers divide the surplus created by data externalities. The impact of the data collection depends on what data externalities consumers impose on each other. It could be beneficial or harmful, depending on whether the allocation of data is substitutable or complementary. However, firms tend to collect too much data because of the spillover effect from one consumer to another.

Overall, the literature on emphasizing externalities from correlated behavior and preferences has emphasized that firms usually collect too much data relative to a welfare maximizing benchmark.

There is also information in the choice not to provide data. For example, by not providing data, consumers may reveal their types in terms of advertising responsiveness and willingness to pay (Bergemann and Bonatti (2015)). This effect depends on the proportion of consumers who withhold data because of their intrinsic value for privacy. If enough consumers have an intrinsic value of privacy, the firm cannot infer the instrumental value of an individual's data based on the decision to withhold that data.

Despite the extensive theory work, there is still a lack of empirical evidence on the data externalities that could guide the policy debate.[1] Empirically documenting the negative externalities from data highlighted in the large and growing theory literature is a promising area for future research.

4.2 Positive Externalities of Data

Data also has positive externalities. This can benefit consumers and firms, increasing overall welfare. For example, the benefits of Google search results come from the data flows from all users' search activities. Likewise, the recommendation functions on Spotify, and the "frequently bought together" section on retail websites are made available because of other users' data on consumption patterns.

4.2.1 Productivity & Data Economy

A growing macroeconomic literature emphasizes the data economy (see Veldkamp and Chung (2019) for a survey). Jones and Tonetti (2020) examine how property rights for data determine its use in the economy. Data serves as an input into the development of high-quality ideas and the non-rival nature of data means that there are social gains to the wide use of data. If firms own property

[1] There is empirical work on externalities of a different type. Goh et al. (2015)) document a negative externality from the decision to withhold data on others in the context of the U.S. Do Not Call registry. They show that when more consumers register to block marketing calls, the remaining consumers receive more calls.

rights over data, then data might be hoarded, and the social gains will go unrealized. In contrast, if consumers control the property rights, the data may be used more broadly. While this paper does not emphasize the negative externalities to consumers highlighted above, a key implication is that the direct benefits that consumers get from privacy, and the externalities that firms get from data, can be addressed through giving property rights to consumers.

Farboodi and Veldkamp (2022) emphasize the role of data as a byproduct of economic activity on productivity and economic growth. Data serves to improve predictions and thereby optimize business processes. In the short run, data may have increasing returns as firms with many customers gather data which in turn improves productivity. This allows the firm to attract even more customers. In the long run, however, the data economy does not generate sustained growth. Data has diminishing returns to improving predictions. With respect to externalities from data, the core takeaway from Farboodi and Veldkamp (2022) is a pair of opposing forces—increasing and decreasing returns to data—that despite non-rivalry, increasing returns, and the production of data as a byproduct of economic activity in the short run, data production is efficient in equilibrium. With a natural bound on the prediction error, data has diminishing returns in the long run. Therefore, the positive externalities do not create incentives to subsidize data. Negative externalities, as expected, generate too much data in equilibrium.

4.2.2 Socially Beneficial Behaviors

Information (data) disclosure is not always harmful to individuals. It may in some cases decrease information distortion, incentivize prosocial behaviors, and improve social welfare. Related to Bénabou and Tirole (2006) that explore individuals' incentives in the pursuit of the prosocial activity, Daughety and Reinganum (2010) develop a model of the economics of privacy, where individuals' action generate externalities. Under the regime of privacy, agents choose their full information optimal actions; while under publicity, they distort to enhance others' perceptions of themselves. The trade-off arises between the expected disutility due to signaling and the increased contribution to the public good. The model considers three primary elements: an intrinsic value

for the activity, esteem (and, by contrast, social disapproval), and the consumption of public goods that arise from the aggregate activity of all other individuals. When the disutility of distortion is low relative to the marginal utility of the public good, a policy of publicity is optimal. Of course, the use of data to incentivize pro-social behavior can have negative consequences, an issue related to what Tirole (2021) labels "Digital Dystopia".

5 Implications

We organize the discussion of the implications from two major perspectives– regulatory implications and non-regulatory privacy protections. The externalities highlighted above suggest a role for targeted government intervention. In addition, consumer preferences create incentives for firms to invest in non-regulatory privacy protections.

5.1 Models of regulation

Privacy regulations restrict the flow of data. In the process, they can protect consumers from direct and indirect harms from data flows. They can also encourage firms to avoid competitive pressures that may degrade consumer privacy and security. Privacy regulations also can have negative consequences on market outcomes, particularly with respect to competition, innovation, and both producer and consumer surplus.

5.1.1 Consumer Behavioral Factors

Most current privacy protection regulations are designed in the spirit that consumers are uncertain and vulnerable to firms' actions. An extensive stream of literature has established that trust can reduce privacy concerns. Kummer and Schulte (2019) detail that privacy concerns of consumers in the app installation process are influenced by the reputation of those application developers. Similarly, Chen et al. (2021) show through an experiment that trust plays an important role in people's data sharing intent on the Alipay platform. The opt-in rate dropped due to the Alipay logo removal, which did not change any actual contracts but the perception of trust. The effectiveness of regulations varies due to different levels of consumer perception of control. Miller and Tucker (2018)

focus on how different features of privacy regulation lead to various effects. Among three alternative approaches to protecting patient privacy, they find notice and consent deters individuals from obtaining genetic tests, while an approach that grants users control over redisclosure encourages the spread of genetic testing. The positive effect of redisclosure on data flows stems from its ability to provide consumers with the perception of control over the use of consented data. Additionally, Baye and Sappington (2020) show that the impact of privacy policies can depend on consumer sophistication. If sophisticated customers know enough to make optimal decisions in the absence of regulatory protection, then regulations aimed at protecting unsophisticated customers may do so at the expense of sophisticated consumers.

5.1.2 Competition and Innovation

Privacy regulations could increase market concentration. Campbell et al. (2015) theoretically investigate the relationship between privacy protection and market structure. The results of the model suggest that the commonly used consent-based approach may disproportionately benefit generalist firms over specialist firms. Privacy regulation may be anti-competitive due to the nature of its transaction costs. This negative effect is strongest in industries with little price flexibility, such as the advertising-supported internet.

Data is an input into innovation (Goldfarb and Tucker (2012a)). In online advertising, healthcare, and a number of other fields, digital data generates better products and services and more efficient production. Therefore, restrictions on data flows will have an impact on the rate and direction of innovation. For instance, privacy protection of patients could discourage healthcare IT adoption efforts, and consequently lead to worse health outcomes (Adjerid et al. (2016); Derksen et al. (2021)).

5.1.3 Mitigation of Negative Externalities

As noted above, a growing literature examines how negative externalities may mean that even fully informed and rational consumers provide data to firms in excess of the welfare-maximizing amount

((Bergemann et al. (2022); Acemoglu et al. (2019); Choi et al. (2019)). Therefore, giving consumers control rights over their data, which is the spirit of many existing and proposed regulations including Europe's General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA), is insufficient. However, outlawing the selling of data entirely would also be harmful to consumers (Jones and Tonetti (2020)). The literature suggests a number of alternative tools. Speaking directly to the negative externality, Bergemann et al. (2022) call for consumer unions, at the segment level instead of the individual level, to internalize the data externality when bargaining with powerful digital platforms. Fainmesser et al. (2022) advocate a two-pronged regulatory policy, combining a minimal data protection requirement and a tax proportional to the data collected, to restore optimal efficiency. Jones and Tonetti (2020) emphasize the value of consumer property rights over data. Ali et al. (2022) highlight granting consumers granular control instead of an all-or-nothing form such as the opt-in option in GDPR. Consumer's selective disclosure of data can amplify competition or prompt a monopolist to lower price. Montes et al. (2019) advocate regulators' focus to be on how information is transacted (eg. the data contracts between data suppliers and users), not directly facilitating consumer privacy.

5.2 Empirical Impact of Regulation

A number of privacy regulations exist. The impact of these regulations has been examined across a variety of contexts. Goldfarb and Tucker (2011) examine an early digital privacy regulation, the EU ePrivacy Directive, which came into effect in 2004. The authors use the responses of 3 million survey takers who had been randomly exposed to nearly 10 thousand online display (banner) advertising campaigns to explore how privacy regulation influenced the effectiveness of advertising. They document that following the ePrivacy Directive, banner ads experienced a reduction in effectiveness of over 65 percent, in terms of changing the difference between treatment and control groups in stated purchase intent. Thus, the privacy regulation appears to have worked. Firms likely used less data, and therefore the ads became less effective. However, to the extent that advertising-supported software is an important industry, the regulation likely reduced the growth of that industry in Europe relative to the United States.

Many papers have focused on the impact of the GDPR. The empirical challenge in this work is that the GDPR was meant to have a global impact, and so there is no straightforward control group. As a consequence, the best papers in this stream use the heterogeneous impact of the regulation on different types of firms and different types of data flows to build a convincing argument. **Table 1** provides a summary of the various papers in the literature. Overall, the conclusion from these papers is that the GDPR led to an immediate reduction in web visits and revenue (Goldberg et al. (2022); Aridor et al. (2022)) and a reduction in the efficiency of online search (Zhao et al. (2021)). It also appears to have reduced the firm's ability to target advertising and track consumers (Godinho de Matos and Adjerid (2022); Peukert et al. (2022)). Competition appears to have decreased in the online advertising market (Johnson et al. (2022); Zhao et al. (2021)) and there was a decline in new firms, venture capital investment, and new apps (Jia et al. (2021); Janssen et al. (2022)). In summary, the early evidence in the aftermath of the GDPR is that it worked, in the sense that firms were using less data in the year following the law's passing. This had costs in terms of firm profits, the consumer online experience, innovation, and competition. There is some suggestive evidence that the impact has declined over time, with both less consumer protection and less impact on concentration (Johnson et al. (2022)).

Table 1: Empirical Evidence of GDPR's Effects

Authors	**Journal**	**Main Findings**	**Implications**	**Data Setting**
Peukert et al. (2022)	*Marketing Science*	Websites reduce the number of third-party web technology providers they use, including websites not legally bound by the GDPR. The changes are disproportionally pronounced among less popular websites.	GDPR brought market concentration. All firms experience losses. However, the vendor leader, Google, incurs relatively smaller losses and greatly expands its market share in crucial markets like advertising and analytics.	Web technology industry

Godinho de Matos and Adjerid (2022)	*Management Science*	Consumer's consent for different data types improved when GDPR-compliant consent was obtained, leading to an increase in sales because of more effective targeted advertising.	GDPR may be effective to enhance consumer privacy protection while enabling companies to improve products that rely on consumers' personal data at the same time.	A large telecommunication provider with operations in Europe
Jia et al. (2021)	*Marketing Science*	The study finds negative short-term effects of the GDPR on investment in technology ventures. The effect is particularly pronounced in the period immediately after the GDPR's rollout, and for newer, data-related, and consumer-facing ventures.	GDPR had a disproportionally negative impact on venture capital investment into technology firms.	Venture capital investment
Zhuo et al. (2021)	*Telecomm. Policy*	All estimates show small or zero effects of GDPR: the number of observed agreements, the agreement types, the number of observed inter-connection points per agreement, the entry and the observed number of customers of networks.	GDPR had no visible short-term impact on internet interconnection layer.	Internet interconnection
Aridor et al. (2022)	*RAND Journal of Economics (forthcoming)*	The opt-in requirement of GDPR led to a 12.5% decrease in the consumer amount. However, the remaining consumers are trackable and predictable for a longer period of time. Their rising value to advertisers offsets part of the losses.	GDPR-enabled opt-out option increases the trackability of the opt-in consumers who choose to reveal their data, imposing an externality.	Online travel intermediary
Goldberg et al. (2022)	*American Economic Journal: Economic Policy (forthcoming)*	After GDPR's enforcement deadline, the platform recorded an approximately 12% reduction in both website pageviews and e-commerce revenue among EU users for 1,084 online firms.	GDPR not only reduced data recording but also harmed real economic outcomes.	Adobe's website analytics platform

Johnson et al. (2022)	*Management Science (forthcoming)*	After GDPR's enforcement deadline, the website use of web technology vendors decreased 15% among EU residents. At the same time, the concentration of vendor market increased by 17% since websites are more likely to drop smaller vendors.	GDPR increased market concentration among technology vendors in a B2B context.	Web technology vendors
Janssen et al. (2022)	NBER Working Paper	GDPR induced approximately one third of the available apps to exit and decreased the entry rate of new apps by half in the market.	GDPR reduced beneficial innovation.	Apps on Google play store
Zhao et al. (2021)	Working Paper	GDPR impacts consumers' online browsing and search behavior. The authors find a panelist exposed to GDPR has 21.6% more search terms for information and 16.3% more pages browsed for goods and services access, indicating higher friction.	GDPR increase friction in online search. The increased friction is heterogeneous across firms, where smaller e-commerce firms were hurt more.	Consumer online browsing, app usage and search activities

5.3 Non-regulatory Privacy Protection

Firms have incentives to protect consumer privacy, even in the absence of regulation. For example, Jullien et al. (2020) investigate the equilibrium privacy policy of websites that generate revenue from charging the third parties on user information. Therefore, customer retention incentivizes the website to be mindful in its monetization efforts or in investing resources to screen third parties. Furthermore, a firm's privacy protection choice can work as a competition-mitigation strategy (Lee et al. (2011)). While empirical work is still nascent, recent changes at Apple and Google are likely to lead to a richer understanding of firm incentives for privacy protection. Specifically, Apple and Google have restricted certain types of data flows from their devices and operating systems to third parties [2]. Apple's App Tracking Transparency (ATT) feature asks users' permission to be tracked from advertisers for every app they download on their iPhone. Similarly, Google introduces its privacy initiative—Privacy Sandbox against cookies. These restrictions protect their customers from

[2]https://www.bloomberg.com/news/articles/2021-04-26/how-apple-google-are-killing-the-advertising-cookie-quicktakexj4y7vzkg

data-related harms from third parties (whether intrinsic or instrumental) but may have negative consequences on data-driven innovation at other firms. Measuring these effects remains an open question.

Privacy technology solutions may complement the regulatory process. All-or-nothing forms of consumer control—such as track/do-not-track—need richer and more sophisticated technologies to benefit consumers (Ali et al. (2022)). A great deal of engineering effort has gone into enabling data-driven innovation while restricting the flow of personally identifiable data. One development is the use of "differential privacy" which preserves anonymity in data while trying to ensure the data can be used for statistical analysis (see Dwork and Roth (2014) for a review and Abowd and Schmutte (2019) for an example in economics). Another development includes decentralized data management through a distributed ledger (Zyskind et al. (2015)) or through anonymous transactions (Böhme et al. (2015)). Innovations in privacy-preserving machine learning solutions are growing, where consumer's privacy is protected while some valuable information can still be extracted in order to improve products and services (Sutanto et al. (2013); Zhou et al. (2020)).

6 Conclusion

The increasing concern for privacy is directly related to the increasing use of digital data. The digital privacy literature in economics has focused on the costs and benefits of restricting data flows. Data flows are useful. They allow firms to provide consumers with the products and services they want, at the time they want them. Data flows also have negative consequences. Many consumers have an intrinsic value for privacy, and so are intrinsically hurt by data flows. Data flows can also be used in ways that hurt consumers, and so there is an instrumental value to privacy.

The recent theory literature has emphasized positive and negative externalities from data flows. The empirical literature, however, has focused largely on the direct impact of regulation on consumers and firms so far. Looking forward, a key open question is the empirical relevance of the

various theories of data externalities in determining the nature and consequences of privacy regulations and the strategic benefits to firms of proactively restricting data flows absent regulation.

Both the theory and empirical literatures have hinted at digital competition as central to our understanding of digital privacy. The earlier theory work provided reasons to be optimistic that increased competition may generate a welfare-maximizing level of privacy. The more recent work on externalities suggests that competition may be insufficient. The empirical work on the consequences of GDPR broadly suggests a reduction of competition. It isn't clear, however, the degree to which this is a short-term phenomenon or driven by idiosyncratic aspects of the GDPR as a privacy regulation. There remain a variety of open questions for both theory and empirical work with respect to how privacy regulation will affect competition and how competition (and competition policy) will affect consumer privacy.

To conclude, the economics literature has emphasized that both data flows and privacy have benefits to consumers and firms. Privacy is not free but it is valuable. It affects economic outcomes. As governments consider new privacy regulations, and as firms develop privacy strategy, we hope that the perspective of economists—emphasizing costs, benefits, externalities, and competition—will be central to the discussion.

References

Abowd, J. M. and Schmutte, I. M. (2019). An economic analysis of privacy protection and statistical accuracy as social choices. *American Economic Review*, 109(1):171–202.

Acemoglu, D., Makhdoumi, A., Malekian, A., and Ozdaglar, A. (2019). Too much data: Prices and inefficiencies in data markets. Technical report, National Bureau of Economic Research.

Acquisti, A., Brandimarte, L., and Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221):509–514.

Acquisti, A., John, L. K., and Loewenstein, G. (2013). What is privacy worth? *The Journal of Legal Studies*, 42(2):249–274.

Acquisti, A., Taylor, C., and Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2):442–92.

Acquisti, A. and Varian, H. R. (2005). Conditioning prices on purchase history. *Marketing Science*, 24(3):367–381.

Adjerid, I., Acquisti, A., Telang, R., Padman, R., and Adler-Milstein, J. (2016). The impact of privacy regulation and technology incentives: The case of health information exchanges. *Management Science*, 62(4):1042–1063.

Ali, S. N., Lewis, G., and Vasserman, S. (2022). Voluntary disclosure and personalized pricing. *forthcoming Review of Economic Studies*.

Altman, I. (1975). The environment and social behavior: Privacy, personal space, territory, and crowding.

Ansari, A. and Mela, C. F. (2003). E-customization. *Journal of Marketing Research*, 40(2):131–145.

Aridor, G., Che, Y.-K., and Salz, T. (2022). The effect of privacy regulation on the data industry: Empirical evidence from gdpr. *forthcoming RAND Journal of Economics*.

Ascarza, E. (2018). Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing Research*, 55(1):80–98.

Athey, S., Catalini, C., and Tucker, C. (2017). The digital privacy paradox: Small money, small costs, small talk. Technical report, National Bureau of Economic Research.

Athey, S. and Gans, J. S. (2010). The impact of targeting technology on advertising markets and media competition. *American Economic Review*, 100(2):608–13.

Bajari, P., Chernozhukov, V., Hortaçsu, A., and Suzuki, J. (2019). The impact of big data on firm performance: An empirical investigation. In *AEA Papers and Proceedings*, volume 109, pages 33–37.

Baye, M. R. and Sappington, D. E. (2020). Revealing transactions data to third parties: Implications of privacy regimes for welfare in online markets. *Journal of Economics & Management Strategy*, 29(2):260–275.

Becker, G. S. (1980). Privacy and malfeasance: A comment. *The Journal of Legal Studies*, 9(4):823–826.

Bénabou, R. and Tirole, J. (2006). Incentives and prosocial behavior. *American Economic Review*, 96(5):1652–1678.

Bergemann, D. and Bonatti, A. (2011). Targeting in advertising markets: implications for offline versus online media. *The RAND Journal of Economics*, 42(3):417–443.

Bergemann, D. and Bonatti, A. (2015). Selling cookies. *American Economic Journal: Microeconomics*, 7(3):259–94.

Bergemann, D. and Bonatti, A. (2019). Markets for information: An introduction. *Annual Review of Economics*, 11:85–107.

Bergemann, D., Bonatti, A., and Gan, T. (2022). The economics of social data. *The RAND Journal of Economics*, 53(2):263–296.

Bergemann, D., Bonatti, A., and Smolin, A. (2018). The design and price of information. *American economic review*, 108(1):1–48.

Bleier, A., Goldfarb, A., and Tucker, C. (2020). Consumer privacy and the future of data-based innovation and marketing. *International Journal of Research in Marketing*, 37(3):466–480.

Böhme, R., Christin, N., Edelman, B., and Moore, T. (2015). Bitcoin: Economics, technology, and governance. *Journal of economic Perspectives*, 29(2):213–38.

Brynjolfsson, E. and McElheran, K. (2016). The rapid adoption of data-driven decision-making. *American Economic Review*, 106(5):133–39.

Burtch, G., Ghose, A., and Wattal, S. (2015). The hidden cost of accommodating crowdfunder privacy preferences: A randomized field experiment. *Management Science*, 61(5):949–962.

Campbell, J., Goldfarb, A., and Tucker, C. (2015). Privacy regulation and market structure. *Journal of Economics & Management Strategy*, 24(1):47–73.

Casadesus-Masanell, R. and Hervas-Drane, A. (2015). Competing with privacy. *Management Science*, 61(1):229–246.

Chan, T., Hamdi, N., Hui, X., and Jiang, Z. (2022). The value of verified employment data for consumer lending: Evidence from equifax. *Marketing Science*.

Chen, L., Huang, Y., Ouyang, S., and Xiong, W. (2021). The data privacy paradox and digital demand. Working Paper 28854, National Bureau of Economic Research.

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

Chiou, L. and Tucker, C. (2017). Search engines and data retention: Implications for privacy and antitrust. Technical report, National Bureau of Economic Research.

Choe, C., King, S., and Matsushima, N. (2018). Pricing with cookies: Behavior-based price discrimination and spatial competition. *Management Science*, 64(12):5669–5687.

Choi, J. P., Jeon, D.-S., and Kim, B.-C. (2019). Privacy and personal data collection with information externalities. *Journal of Public Economics*, 173:113–124.

Choudhary, V., Ghose, A., Mukhopadhyay, T., and Rajan, U. (2005). Personalized pricing and quality differentiation. *Management Science*, 51(7):1120–1130.

Coase, R. H. (1972). Durability and monopoly. *The Journal of Law and Economics*, 15(1):143–149.

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

Daughety, A. F. and Reinganum, J. F. (2010). Public goods, social pressure, and the choice between privacy and publicity. *American Economic Journal: Microeconomics*, 2(2):191–221.

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

Derksen, L., McGahan, A., and Pongeluppe, L. (2021). Privacy at what cost? using electronic medical records to recover lapsed patients into hiv care. *Working Paper*.

Dwork, C. and Roth, A. (2014). The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.*, 9(3-4):211–407.

Erlich, Y., Shor, T., Pe'er, I., and Carmi, S. (2018). Identity inference of genomic data using long-range familial searches. *Science*, 362(6415):690–694.

Fainmesser, I. P., Galeotti, A., and Momot, R. (2022). Digital privacy. *HEC Paris Research Paper No. MOSI-2019-1351*.

Farboodi, M., Mihet, R., Philippon, T., and Veldkamp, L. (2019). Big data and firm dynamics. In *AEA papers and proceedings*, volume 109, pages 38–42.

Farboodi, M. and Veldkamp, L. (2022). A model of the data economy. Technical report, Working Paper, Columbia University.

Godinho de Matos, M. and Adjerid, I. (2022). Consumer consent and firm targeting after gdpr: The case of a large telecom provider. *Management Science*, 68(5):3330–3378.

Goh, K.-Y., Hui, K.-L., and Png, I. P. (2015). Privacy and marketing externalities: Evidence from do not call. *Management Science*, 61(12):2982–3000.

Goldberg, S., Johnson, G., and Shriver, S. (2022). Regulating privacy online: An economic evaluation of the gdpr. *forthcoming American Economic Journal: Economic Policy*.

Goldfarb, A. (2014). What is different about online advertising? *Review of Industrial Organization*, 44(2):115–129.

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

Goldfarb, A. and Tucker, C. (2012a). Privacy and innovation. *Innovation Policy and the Economy*, 12:65–90.

Goldfarb, A. and Tucker, C. (2012b). Shifts in privacy concerns. *American Economic Review*, 102(3):349–53.

Goldfarb, A. and Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, 57(1):3–43.

Hann, I.-H., Hui, K.-L., Lee, S.-Y. T., and Png, I. P. (2008). Consumer privacy and marketing avoidance: A static model. *Management science*, 54(6):1094–1103.

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

Ichihashi, S. (2020). Online privacy and information disclosure by consumers. *American Economic Review*, 110(2):569–95.

Ichihashi, S. (2021). The economics of data externalities. *Journal of Economic Theory*, 196:105316.

Iyer, G., Soberman, D., and Villas-Boas, J. M. (2005). The targeting of advertising. *Marketing Science*, 24(3):461–476.

Janssen, R., Kesler, R., Kummer, M. E., and Waldfogel, J. (2022). Gdpr and the lost generation of innovative apps.

Jia, J., Jin, G. Z., and Wagman, L. (2021). The short-run effects of the general data protection regulation on technology venture investment. *Marketing Science*, 40(4):661–684.

Jin, Y. and Vasserman, S. (2021). Buying data from consumers: The impact of monitoring programs in u.s. auto insurance. *Working Paper*.

Johnson, G., Shriver, S., and Goldberg, S. (2022). Privacy & market concentration: Intended & unintended consequences of the gdpr. *forthcoming Management Science*.

Johnson, J. P. (2013). Targeted advertising and advertising avoidance. *The RAND Journal of Economics*, 44(1):128–144.

Jones, C. I. and Tonetti, C. (2020). Nonrivalry and the economics of data. *American Economic Review*, 110(9):2819–58.

Jullien, B., Lefouili, Y., and Riordan, M. H. (2020). Privacy protection, security, and consumer retention. Technical report, Toulouse School of Economics.

Kummer, M. and Schulte, P. (2019). When private information settles the bill: Money and privacy in google's market for smartphone applications. *Management Science*, 65(8):3470–3494.

Lee, D.-J., Ahn, J.-H., and Bang, Y. (2011). Managing consumer privacy concerns in personalization: a strategic analysis of privacy protection. *Mis Quarterly*, pages 423–444.

Lin, T. (2022). Valuing intrinsic and instrumental preferences for privacy. *Marketing Science*.

Loertscher, S. and Marx, L. M. (2020). Digital monopolies: Privacy protection or price regulation? *International Journal of Industrial Organization*, 71:102623.

Miklós-Thal, J. and Tucker, C. E. (2019). Collusion by algorithm: Does better demand prediction facilitate coordination between sellers? *Management Science*, 65(4):1552–1561.

Miller, A. R. and Tucker, C. E. (2011). Can health care information technology save babies? *Journal of Political Economy*, 119(2):289–324.

Miller, A. R. and Tucker, C. E. (2018). Privacy protection, personalized medicine, and genetic testing. *Management Science*, 64(10):4648–4668.

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

Neumann, N., Tucker, C. E., and Whitfield, T. (2019). Frontiers: How effective is third-party consumer profiling? evidence from field studies. *Marketing Science*, 38(6):918–926.

Nissenbaum, H. (2004). Privacy as contextual integrity. *Wash. L. Rev.*, 79:119.

Nissenbaum, H. (2009). Privacy in context. In *Privacy in Context*. Stanford University Press.

Norberg, P. A., Horne, D. R., and Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1):100–126.

Peukert, C., Bechtold, S., Batikas, M., and Kretschmer, T. (2022). Regulatory spillovers and data governance: Evidence from the gdpr. *Marketing Science*.

Posner, R. A. (1981). The economics of privacy. *The American economic review*, 71(2):405–409.

Rafieian, O. and Yoganarasimhan, H. (2021). Targeting and privacy in mobile advertising. *Marketing Science*, 40(2):193–218.

Reimers, I. and Shiller, B. R. (2019). The impacts of telematics on competition and consumer behavior in insurance. *The Journal of Law and Economics*, 62(4):613–632.

Schneier, B. (2015). *Data and Goliath: The hidden battles to collect your data and control your world*. WW Norton & Company.

Shapiro, C. and Varian, H. R. (1998). *Information rules: A strategic guide to the network economy*. Cambridge: Harvard Business School Press.

Shy, O. and Stenbacka, R. (2016). Customer privacy and competition. *Journal of Economics & Management Strategy*, 25(3):539–562.

Smith, M. D., Bailey, J., and Brynjolfsson, E. (2001). Understanding digital markets: Review and assessment. *Understanding the Digital Economy*, page 99–136.

Soleymanian, M., Weinberg, C. B., and Zhu, T. (2019). Sensor data and behavioral tracking: Does usage-based auto insurance benefit drivers? *Marketing Science*, 38(1):21–43.

Solove, D. J. (2008). Understanding privacy.

Solove, D. J. (2021). The myth of the privacy paradox. *Geo. Wash. L. Rev.*, 89:1.

Stigler, G. J. (1980). An introduction to privacy in economics and politics. *The Journal of Legal Studies*, 9(4):623–644.

Sun, T., Yuan, Z., Li, C., Zhang, K., and Xu, J. (2021). The value of personal data in internet commerce: A high-stake field experiment on data regulation policy. *Forthcoming Management Science*.

Sutanto, J., Palme, E., Tan, C.-H., and Phang, C. W. (2013). Addressing the personalization-privacy paradox: An empirical assessment from a field experiment on smartphone users. *MIS quarterly*, pages 1141–1164.

Tamir, D. I. and Mitchell, J. P. (2012). Disclosing information about the self is intrinsically rewarding. *Proceedings of the National Academy of Sciences*, 109(21):8038–8043.

Tang, H. (2019). The value of privacy: Evidence from online borrowers. *Technical Report, HEC Paris*.

Taylor, C. R. (2004). Consumer privacy and the market for customer information. *RAND Journal of Economics*, pages 631–650.

Tirole, J. (2021). Digital dystopia. *American Economic Review*, 111(6):2007–48.

Tucker, C., Agrawal, A., Gans, J., and Goldfarb, A. (2018). Privacy, algorithms, and artificial intelligence. *The economics of artificial intelligence: An agenda*, pages 423–437.

Tucker, C. E. (2014). Social networks, personalized advertising, and privacy controls. *Journal of marketing research*, 51(5):546–562.

Veldkamp, L. and Chung, C. (2019). Data and the aggregate economy. *Journal of Economic Literature*.

Villas-Boas, J. M. (1999). Dynamic competition with customer recognition. *The Rand Journal of Economics*, pages 604–631.

Villas-Boas, J. M. (2004). Price cycles in markets with customer recognition. *RAND Journal of Economics*, pages 486–501.

Warren, S. D. and Brandeis, L. D. (1890). Right to privacy. *Harv. L. Rev.*, 4:193.

Westin, A. F. (1967). Privacy and freedom.

Yang, K. H. (2022). Selling consumer data for profit: Optimal market-segmentation design and its consequences. *American Economic Review*, 112(4):1364–93.

Yoganarasimhan, H. (2020). Search personalization using machine learning. *Management Science*, 66(3):1045–1070.

Zhang, J. and Krishnamurthi, L. (2004). Customizing promotions in online stores. *Marketing science*, 23(4):561–578.

Zhao, Y., Yildirim, P., and Chintagunta, P. K. (2021). Privacy regulations and online search friction: Evidence from gdpr. *Available at SSRN 3903599*.

Zhou, Y., Lu, S., and Ding, M. (2020). Contour-as-face framework: A method to preserve privacy and perception. *Journal of Marketing Research*, 57(4):617–639.

Zhuo, R., Huffaker, B., Greenstein, S., et al. (2021). The impact of the general data protection regulation on internet interconnection. *Telecommunications Policy*, 45(2):102083.

Zyskind, G., Nathan, O., et al. (2015). Decentralizing privacy: Using blockchain to protect personal data. In *2015 IEEE Security and Privacy Workshops*, pages 180–184. IEEE.

Figure 1: Comparison in Click-Through Rates Before and After

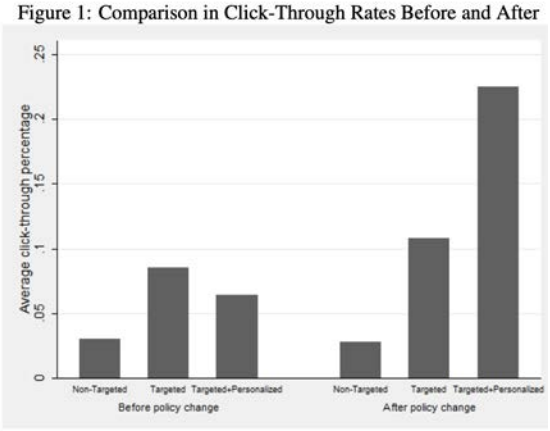

Figure 2: Comparison in Click-Through Rates Before and After

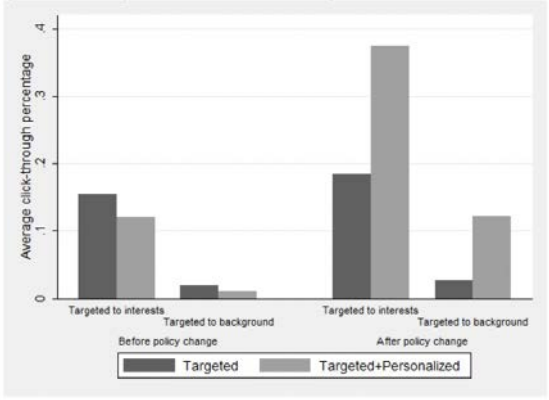

Figure 1: Consumer click-through rates before and after the introduction of policy