

The Intended and Unintended Consequences of Privacy Regulation for Consumer Marketing

Jean-Pierre Dubé, University of Chicago
John G. Lynch, University of Colorado-Boulder
Dirk Bergemann, Yale University
Mert Demirer, Massachusetts Institute of Technology
Avi Goldfarb, University of Toronto
Garrett Johnson, Boston University
Anja Lambrecht, London Business School
Tesary Lin, Boston University
Anna Tuchman, Northwestern University
Catherine Tucker, Massachusetts Institute of Technology

Abstract

As businesses increasingly rely on granular consumer data, the public has increasingly pushed for enhanced regulation to protect consumers' privacy. We provide a perspective based on the academic marketing literature that evaluates the various benefits and costs of existing and pending government regulations and corporate privacy policies. We make four key points. First, data-based personalized marketing is not automatically harmful. Second, consumers have heterogeneous privacy preferences, and privacy policies may unintentionally favor the preferences of the rich. Third, privacy regulations may stifle innovation by entrepreneurs who are more likely to cater to underserved, niche consumer segments. Fourth, privacy measures may favor large companies who have less need for third-party data and can afford compliance costs. We also discuss technology platforms' recent proposals for privacy solutions that mitigate some of these harms, but, again, in a way that might disadvantage small firms and entrepreneurs.

Keywords

Privacy, consumer protection, discrimination, digital exclusion, competition, regulation

1. Introduction

Several recent initiatives have re-invigorated the debate about consumer digital privacy in the U.S. As we write, 19 states in the U.S. have enacted comprehensive privacy laws that emulate the European Union’s (EU) General Data Protection Regulation (GDPR): 11 of these laws come into effect in 2025 or 2026. Privacy regulation has broad and bipartisan support, as with the recent draft of the American Privacy Rights Act in 2024. The June 2024 draft would preempt state laws. Congress is debating the American Innovation and Choice Online Act (AICOA) which includes restrictions to enhance consumer privacy for both private and publicly listed companies. President Biden’s Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence highlighted the Federal Government’s commitment to enforce consumer protection laws and enact appropriate safeguards “against fraud, unintended bias, discrimination, infringements on privacy, and other harms from AI” (Biden 2023).

The Federal Trade Commission (2024, p. vii) recently recommended federal privacy legislation like GDPR for social media and streaming platforms, recommending “*minimizing data collection to only that data which is necessary for their services,*” “*data retention and data deletion policies,*” “*limiting data sharing with affiliates, other company-branded entities, and third parties*” and “*clear, transparent, and consumer-friendly privacy policies.*” They cited no academic literature on consumer privacy in marketing or economics, and they discussed neither potential harms from the proposed measures nor empirical evidence of benefits of past regulation. Marketing scholars are well-positioned to fill this gap by offering data-driven insights.

This *discussion paper*¹ synthesizes emerging empirical findings from the academic literatures in

¹ Prior *Marketing Science* discussion papers include Bass (1995), Farley et al (1995), Rossi and Allenby (2003), Chintagunta et al (2006), and Mela (2011). The Marketing Science Institute convened this group of experts on the economics of privacy at a Brookings Institution workshop to distill key themes from the academic literature pertinent to deliberations of technology platforms, digital marketing firms, and regulators.

marketing, economics, and behavioral science about the intended versus unintended consequences of existing and pending privacy regulations for consumer markets (Bleier et al. 2020).² Regulation invariably involves tradeoffs, including the unintended costs that are often omitted from proposals for data privacy regulation. Our intended audience includes: a) policymakers; b) managers who set corporate policies and who might collaborate in creating public policy, and c) privacy scholars outside of marketing science. Given the likely re-evaluation of pending initiatives by the Trump administration, such a synthesis is timely.

Two high-level conclusions emerge from our discussion. First, policy analysis has focused heavily on restricting data flows. But humans seek **boundary regulation** -- sharing information when they wish and restricting it when they do not (Acquisti 2023; Goldfarb and Que 2023). Second, marketing privacy regulations often favor the powerful in ways not acknowledged by legal experts driving the privacy debate.

After outlining key intended benefits of privacy regulation in Section 2, we elaborate four key points leading to this conclusion.

- Section 3. Some privacy advocates assume, incorrectly, that personalized marketing based on granular consumer data is automatically harmful -- a zero-sum game in which value is transferred from consumers to firms. Personalization can be win-win. Moreover, in domains like pricing, *lack* of personalization may favor those most able to pay.
- Section 4. Heterogeneous privacy preferences create the central problem in privacy regulations. Some consumers lack established privacy preferences, so surveys like those cited by FTC (2024) may be unreliable for policy. Among consumers with established preferences, interests conflict. Current regulations tend to favor high-income consumers with stronger

² We do not consider legal arguments for consumer privacy as a fundamental right or concerns about access to personal data by malign actors or governments.

privacy preferences. Low-income consumers are already digitally invisible to the point that firms implicitly discriminate by excluding them from outreach. Blanket restrictions on data-sharing exacerbate that inequality.

- Section 5. Privacy measures stifle the wave of innovation and entry of direct-to-consumer online business, especially by small entrepreneurs with offerings targeted to niche and underserved segments.
- Section 6. Privacy measures tilt competition in favor of large, incumbent firms that have less need for third-party data and can afford the large compliance costs. A growing literature shows differential harm to small businesses from government policies and tech platform policies like Apple's App Tracking Transparency (ATT).

Section 7 outlines the promise and problems of new Privacy Enhancing Technologies. Section 8 concludes and recommends future research.

2. Intended Benefits of Digital Marketing Privacy Regulation and Pertinent Regulations

There are several important reasons why consumers may benefit from oversight of the use of their personal data by marketers and why the current consent-based regime fails to protect consumers (Acquisti 2024).

- a. Consumers might be harmed if firms possess and act on incorrect information about them and if such data use lacks transparency and the ability to correct (CFPB 2022a). Consumer scores can suffer from biases due to social inequities and intrinsic bias in the data themselves, and from inaccuracies and out-of-date information. Their usage can lead to seemingly unfair differential treatment.
- b. Firms might use personal data to discriminate against disadvantaged consumers or protected classes (CFPB 2022b). Consider the U.S. Justice Department suit "alleging that Meta's

housing advertising system discriminates against Facebook users based on their race, color, religion, sex, disability, familial status and national origin” (Civil Rights Litigation Clearinghouse 2022). Even without an intent to discriminate, market forces may cause online advertising algorithms to underserve certain groups of consumers and deny them full access to the digital economy (Lambrecht and Tucker 2019).

- c. Firms might price discriminate against consumers with higher valuations of a product or service. The Council of Economic Advisors (2014) explains: *“Consumers have a legitimate expectation of knowing whether the prices they are offered for goods and services are systematically different than the prices offered to others.”*
- d. “Notice and consent” regimes may offer insufficient protection to consumers. Sellers and buyers may have asymmetric information about the consequences of data-sharing unforeseen by the buyer (Acquisti et al. 2016; Clark 2020). Moreover, consumers suffer from “consent fatigue” due to the large number of consent requests, often to get access to information on a website (Acquisti 2024; Miller and Tucker 2018). Arguably, it is unreasonable to expect any consent to be “informed” and meaningful (Utz et al. 2019).

Motivated by concerns like those above, the EU proactively launched the GDPR in 2018, with far-reaching implications for over 20 million companies spanning dozens of countries. GDPR puts a high bar on a firm’s ability to collect and process personal individual data and to guarantee transparency. For example, personal data such as sex and gender should only be collected and processed when it is necessary for the task and not used beyond the original purpose. Our Web Appendix provides further details.

The U.S. has adopted a more decentralized patchwork of federal and state laws, along with industry-specific regulations that are enacted at the federal level but apply to specific sectors, such as HIPAA which governs health data. Most U.S. privacy measures have been implemented in a

fragmented way across a variety of state laws like California’s Privacy Rights Act and Colorado’s Privacy Act.

Perhaps in anticipation of heightened regulations, many American firms proactively strengthened consumer protections. Apple’s “Ask-Not-To-Track” option in its ATT framework blocks apps from tracking an individual’s behavior on other companies’ apps and websites without opt-in consent (Kesler 2023). In 2024, Google ended its five-year effort to phase out third-party cookies in its Chrome Browser, aiming instead to introduce tools for consumers to make informed choices and adjust privacy settings easily.

While issues (a)-(d) are indeed concerns, we discuss below how existing privacy regulations and policies may have unintended consequences on both the supply and demand sides that offset anticipated consumer welfare gains. For instance, Apple’s ATT appears to have reduced the number of fraud complaints (Bian et al. 2023) but increased product prices and market concentration and reduced ad spending (Deisenroth et al. 2025).

3. Access to Consumer Data Can Increase Consumer Value via Personalization

One of the most contentious aspects of the use of personal data for marketing purposes is the personalization of the marketing mix. Consider personalized pricing. Some have questioned its legality (Ramasastry 2005). The popular press has been rife with headlines like “To Fight Surveillance Pricing, We Need Privacy First” (Noble 2024). Documented examples of such personalized pricing are scarce, yet public officials have expressed concerns: “[differential pricing] transfers value from consumers to shareholders, which generally leads to an increase in inequality and can therefore be inefficient from a utilitarian standpoint” (Council of Economic Advisors 2015, p. 6).

Theoretically, personalized marketing could harm consumers by transferring consumer value to firms or excluding segments of the population from valuable communications or regressive

pricing. However, economic theory also shows that price discrimination can increase consumer value when the total quantity of consumers served increases (e.g., Stole 2007). Moreover, oligopoly price discrimination can trigger price wars and a prisoner's dilemma whereby firms' prices and profits decline to the benefit of consumers (e.g., Stole 2007). Whether price discrimination increases both consumer value and firm value depends on the nature of the consumer segments defined by the data available (Bergemann et al. 2015; 2024).

Several empirical case studies demonstrate how personalized pricing can benefit consumers less able to pay. Dubé and Misra (2023) find that personalized pricing lowered prices for over 60% of the customers for a large digital human resources platform, primarily for the smallest enterprise customers. DellaVigna and Gentzkow (2019) find that supermarket prices in poor neighborhoods are 8% higher than they would be if large chains implemented more granular geographic price differences across stores in each city. Allcott et al. (2019) find that willingness-to-pay for healthy nutrients increases with a household's income, so personalized pricing could help reduce nutritional inequality. Glenn et al. (2022) discuss personalized pricing to help low-income households afford municipal fines and fees to avoid defaulting and accumulating municipal debt, while potentially increasing municipal revenues. Arslan et al. (2023) find that switching from uniform to variable pricing of National Football League tickets benefitted hometown teams with lower income and higher income diversity.

Looking beyond pricing, the application of large language models to personalize the California SNAP program's email campaign more than doubled enrollments for food stamps (Misra 2020). Disabling personalization for Alibaba customers led to less efficient search and lower purchase incidence, particularly for consumers with unusual tastes and niche merchants serving them (Sun et al. 2023). Personalized recommendations increased the diversity of digital music consumed (Datta et al. 2018). Regulators have acknowledged some of these non-price

personalization benefits (e.g., Council of Economic Advisors 2014, pp. 7-8). In sum, data-based personalized marketing can improve matching of customers with less common needs with appropriate sellers.

4. Which Consumers Care Most About Privacy, and Do Privacy Policies Unintentionally Favor the Privileged?

Consumers differ in privacy preferences and consequences of privacy policies. Policy analyses like FTC (2024) rely on consumer surveys that appear to show broad public support for more restrictive privacy regulation. In a Pew survey, "...81% of the public say that the potential risks they face because of data collection by companies outweigh the benefits" (Auxier et al. 2019). Pew finds bipartisan support for more government regulation by 68% of Republicans and 78% of Democrats (Faverio 2023).

How probative are these surveys for policy? A well-documented "privacy paradox" shows that stated attitudes may diverge from actual privacy behaviors (e.g., Goldfarb and Que 2023) and are contaminated by "socially desirable" responses (Larsen 2023). Is this low attitude-behavior correspondence a symptom of "nonattitudes" and low importance and knowledge for some consumers (cf. Howe and Krosnick 2017, Shuman and Presser 1980)? As we discuss in more detail in the Web Appendix, when survey respondents do not have a prior opinion on an issue, their responses are "constructed" on the spot. The hallmark of constructed preferences is their sensitivity to normatively irrelevant context, such as minor changes in wording or question order (Feldman and Lynch 1988).

4.1 Instrumental, Contextually Determined Privacy Preferences

Privacy preferences are likely heterogeneous and include both **instrumental** and **intrinsic** components, complicating consumer welfare analysis (Lin 2022). Consumers may have an intrinsic psychological distaste for sharing data, independent of how their data will be used. Consumers may

also have instrumental preferences driven by strategic or economic motives, such as a high-income consumer deciding not to share data for fear of receiving higher prices.

Data transmission can sometimes benefit consumers such that regulators should focus on **boundary regulation** instead of banning data tracking. While industry solutions like browser-level Global Privacy Controls seem well suited for consumers with strong “intrinsic” privacy preferences to say “no” or “yes” to all data sharing, consumers with instrumental privacy preferences will instead need a way to provide meaningful “batched consent” to accept some requests and decline others without experiencing “consent fatigue.” For instance, iOS users’ opt-in rates to share data are much higher for gaming apps than travel apps (Perik 2024). Successful examples of data-sharing between entities to improve service quality while remaining compliant with GDPR exist in child services and medical records (ICO 2024). The Web Appendix discusses other concrete policies that treat privacy as a matter of boundary regulation rather than concealment.

4.2 Weaker Privacy Preferences for Low-Income, Less Educated, and Younger Consumers

Richer, more educated, and older consumers tend to be willing to pay more for privacy on Facebook than poorer, less educated, and younger consumers (Lin and Strulov-Schlain 2024). These findings conceptually replicate Savage and Waldman (2015) and Varian et al. (2005).

Privacy preferences may also be constructed and sensitive to “choice architecture” or “nudges” (Thaler and Sunstein 2021) to encourage disclosure, such as using “opt out” rather than “opt in” as the default for data tracking (Johnson et al. 2002). Lin and Strulov-Schlain (2024) find that poorer, less educated, and younger consumers who value their privacy the least are also the most sensitive to choice architecture, again suggesting that their preferences are largely constructed on the spot. See also Mrkva et al. (2021).

These findings raise thorny consumer protection issues. Policymakers should consider whether the gains from data-sharing outweigh possible risks from privacy loss for lower-income, younger, and less educated consumers. If so, current privacy proposals like FTC (2024) may tilt toward the instrumental concerns of the privileged if personalization has redistributive effects in domains like pricing. Of course, one who believes that more disadvantaged consumers are wrong not to worry more about privacy –rather than simply having different priorities -- might argue for paternalistic solutions that make data-sharing even harder than in the current consent-based frameworks (cf. Acquisti et al. 2022).

4.3 Privacy and Marketing Inclusiveness

The regulatory push to limit or ban access to third-party data suggests a prevailing sentiment that firms know too much about consumers. We now review research suggesting that, counterintuitively, firms may not know enough about disadvantaged consumers to market them offers or to monitor for unintended discrimination.

4.3.1 Algorithmic Discrimination against Disadvantaged Consumers

The U.S. prohibits disparate treatment based on someone’s protected class status in markets for housing, credit, employment, public accommodations, and voting. Even when targeted marketing that discriminates on sensitive consumer attributes is permissible, firms may avoid it if such practices may be perceived as unfair or unethical by stakeholders.

Indeed, sometimes, algorithms can learn to discriminate absent any deliberate intent by the marketer (Netzer et al. 2019). A cold start problem in algorithmic learning can cause uneven outcomes for minority relative to majority groups (Lambrecht and Tucker 2024). Facebook’s bidding algorithm served more ads for a STEM career campaign to men than women even though the algorithm was blind to gender (Lambrecht and Tucker 2019).

Paradoxically, the policy objective of privacy protection and the policy objective of

nondiscrimination can conflict (Ali et al. 2019). Many privacy policies discourage collecting and storing personal attributes like race and gender. However, without knowing a customer's race and gender, it would be difficult to monitor digital marketing for unfair discrimination (King et al. 2023). Nor can firms readily correct for potentially unfair discrimination when algorithms use variables correlated with otherwise unobserved protected attributes (cf. Ascarza and Israeli 2022, who propose an adaptive trees algorithm to partial out unintended disparate algorithmic treatment of specific consumer segments statistically). The Web Appendix discusses limits and imperfections of methods for monitoring using inferred sensitive variables.

4.3.2 Is Privacy a Problem of the Privileged? Too Little Data for Disadvantaged Consumers

Marketers often pre-screen consumers for offers based on predicted profitability. The necessary data are often more limited and fragmented for poorer than rich consumers. Poorer consumers live in “data deserts” (Tucker 2023), causing *algorithmic exclusion* from algorithmic processing due to missing or fragmented data. This exclusion thwarts marketing outreach and may deprive them of offers, exacerbating data deserts and marginalization.

Privacy restrictions may inadvertently exacerbate the algorithmic exclusion of “invisible” poor consumers. For instance, Boston's “Street Bump” app disproportionately increased pothole repair rates in wealthy neighborhoods more likely to use the app (Tucker et al. 2023). John Hancock's Fitbit-based insurance discounts disproportionately benefited the wealthy who more commonly use fitness trackers (Goldfarb and Tucker 2017). In peer-to-peer lending, privileged consumers with more social connections were more likely to benefit (Freedman and Jin 2017).

Data fragmentation is another obstacle for poorer consumers (Tucker 2023; 2024). Contact information like name, phone number, or email address is less consistently tracked for consumers facing economic instability. Large data brokers' databases are more likely to have missing or biased records for individuals with lower wealth, education, and home ownership rates (Neumann et al.

2024). Experian was 50% less likely to contain information about Hispanic and Asian individuals than White individuals (Kaplan et al. 2017). Credit scores are statistically noisier indicators of default risk for historically underserved groups who lack credit histories (Blattner and Nelson 2021).

While some consumers may benefit from reduced corporate and government surveillance, others may be harmed by digital exclusion in markets like credit, employment, or public services. Financial inclusion increased in India with the launch of the Aadhar national ID scheme (IMF 2023), and in Peru where credit access was expanded using retail data (Lee et al. 2024). Policymakers should consider how to level the playing field by incentivizing more equal understanding of poorer and richer consumers.

5 Privacy Measures May Stifle Entry and Innovation by Entrepreneurs and Small Businesses Who Are More Likely to Serve Niche Consumer Segments

Digital advertising now constitutes most ad spending (Cramer-Flood 2023). Data-driven targeting improves its efficiency over traditional media. Data-driven targeting often uses cross-site or cross-app user data. Major advertising platforms track such data with user identifiers, like third-party cookies, to optimize ad delivery. In addition to ‘onsite data’ – user data collected on the platform – some advertising platforms facilitate ‘offsite data’ -- collected off the advertising platform -- including browsing history, purchase events, and other online user actions. For example, online retailers can use a pixel or conversion API to transmit purchase data to Meta to optimize ad campaigns.

Greater marketing data generated a surge in the launch of valuable, disruptive new products for consumers, particularly by small businesses and entrepreneurs. For instance, the craft beer segment surged from 4% of U.S. sales to over 20% since 2005 (Bronnenberg et al. 2022). In 2018, 16,000 smaller CPG companies generated 19% of total U.S. sales, a \$2 billion (2 percentage point) increase over 2017 (eMarketer Editors 2019).

By reducing the costs that used to be required to build a new brand on television and other mass media, digital advertising eroded massive barriers to entry for new consumer brands (cf. Sutton 1991; Bronnenberg et al. 2009). Digital advertising saves small U.S. entrepreneurs \$163 billion annually, and over two-thirds of them lack a cost-effective alternative ad medium (Kerrigan and Keating 2019).

More broadly, without cross-site/app identity, consumers enjoy less free content (e.g., Johnson et al. 2024; Kircher and Foerderer 2023). GDPR likely hurt the European advertising-supported software industry and stifled innovation with an associated decline in new firms, venture capital investment, and new apps (Jia et al. 2021; Janssen et al. 2022). Apple's ATT substantially degraded digital advertising: firm revenue fell 37% more for more Meta-dependent firms (Aridor et al 2024) and the number of U.S. establishments fell 1.1% (roughly 91,000 U.S. establishments) in more affected industries. Privacy protection in healthcare could discourage IT adoption and lead to worse health outcomes (Derksen et al. 2021; Miller and Tucker 2009, 2011; cf. Adjerid et al. 2016). Security and privacy measures for financial services can discourage take-up of innovations such as online banking (Lambrecht et al. 2011). Goldfarb and Tucker (2012) discuss the tradeoff between privacy and innovation, though firms may adapt to privacy restrictions by innovating in ways that do not involve data (cf. Agrawal et al. 2019).

6 Privacy Policy May Harm Small Companies with Greater Need for Third-Party Data and Less Ability to Afford Compliance Costs

Privacy regulations also impact competition among businesses that rely on digital marketing. Dozens of papers that consider the economic impact of GDPR largely document its harms to firm performance, competition, innovation, the web, and marketing (Johnson 2024).

GDPR increased the cost of collecting and storing data by requiring firms to enhance data protection, imposing penalties for data breaches, and requiring more transparency to consumers about tracking and data usage. Demirer et al. (2024) provide a case study of one of the largest global

cloud-computing providers between 2015 and 2021, estimating that EU firms store, on average, 26% less data than comparable US firms two years after the GDPR. Interestingly, EU firms decrease their computation relative to comparable US firms by 15%---implying that firms became less data-intensive after GDPR. Demirer et al. estimated that GDPR increased average total data storage costs by 20% -- again, disproportionately more for smaller firms.

Restrictions to limit the effectiveness of digital advertising would likely disproportionately disadvantage small businesses, since nine out of ten predominantly use digital advertising, especially on Meta (Kerrigan and Keating 2019). The GDPR and company-initiated policies, such as the deprecation of third-party cookies and Apple's ATT, limit the collection of 'offsite data'. For instance, banning offsite data for advertising on Meta disproportionately harms small advertisers, with the median small advertiser losing 4.6 times as many customers per \$1,000 spent on advertising than the median large advertiser (Wernerfelt et al. 2024). Apple's ATT policy also disproportionately harmed small businesses (Aridor et al. 2024; Deisenroth et al. 2024). Disabling access to targetable data in Safari and Chrome "*...disproportionately hurt price responsive consumers and small/ niche product sellers*" on an e-commerce retail platform (Korganbekova and Zuber 2023). It also disproportionately harmed merchants serving niche consumers on Alibaba (Sun et al. 2023).

Data restrictions have also eroded digital advertising effectiveness. The EU's e-Privacy Directive EC/2002/58 was associated with a 65% decrease in online advertising effectiveness (Goldfarb and Tucker 2011). Without cookies, the value created by advertising *falls* 52%, a loss shared roughly proportionately between market advertisers, publishers, and ad tech intermediaries (Johnson et al. 2020).

A broader concern is that privacy restrictions that limit or ban the use of data used for advertising could inadvertently increase concentration in the advertising market. GDPR increased concentration in EU digital advertising markets shortly after its implementation (Johnson et al. 2023;

Peukert et al. 2022), with Google and Facebook both gaining market share. Apple's ATT used a different opt-in prompt for Apple apps than for third-party apps; this may have led to higher opt-in rates for Apple apps, potentially giving Apple more access to targetable user data (Baviskar et al. 2024). Regulators should consider this tradeoff between the protection of consumer privacy and the potential harm from increased market concentration and the potential stifling of the recent wave of innovation among small, disruptive consumer brands.

Regulators seem to have noted the disproportionate impact on small businesses. For instance, the most recent EU Data Protection Act recognizes the differential compliance burden for small businesses (Beveridge 2024). Its requirements are lessened for small and medium businesses.

7 A Path Toward Acceptable Data Processing

Privacy regulation steers both the data economy and firm compliance by defining acceptable data processing. For instance, HIPAA specifies health data storage and transfer requirements between covered parties, which can include encryption, de-identification, written agreements, and breach notification. COPPA restricts processing children's data but establishes a safe harbor program for firms to coordinate self-regulation. The GDPR prioritizes privacy while imposing substantial compliance costs on firms because the GDPR defines personal data broadly, imposes multiple data-related responsibilities on firms, and prescribes a high consent standard for many marketing purposes. Forward-looking privacy regulation should consider the role of privacy-enhancing technologies in defining acceptable data processing.

7.1 Privacy-Enhancing Technologies (PETs)

As discussed in the Web Appendix, marketing research about PETs is growing. PETs include diverse technologies such as adding noise to data (i.e., differential privacy), cohorting consumers (e.g., K -anonymity), decentralizing data processing (e.g., federated learning, on-device

computation), limiting data flows (e.g., zero-knowledge proof), and privacy-safe data combination (e.g., secure multi-party computation). The U.S. Census uses differential privacy to anonymize its public statistics. Google uses federated learning to implement keyboard next-word predictions (Hard et al. 2018).

The online advertising industry is on a path to replace cross-site/app identifiers with PETs (Geng et al. 2023; Johnson et al. 2022). For example, the Google (2022) “Privacy Sandbox” consists of multiple technologies for ad targeting (Topics API, Protected Audience API), ad measurement (Attribution Reporting API), and fraud detection (Privacy State Tokens). Similarly, Microsoft launched its Ad Selection API for ad targeting, Apple offers AdAttributionKit for ad measurement, and Facebook and Mozilla jointly proposed their Interoperable Private Attribution approach for ad measurement. These PET applications appear performant (Kobayashi et al. 2024) and see moderate adoption among websites and AdTech vendors (Johnson 2024).

PETs have limitations. First, computer science has treated “re-identifiability” as the focal problem (cf. Ponte et al. 2024). However, consumers do not perceive dramatically lessened privacy violations from PET solutions like Google’s “Topics” and “Protected Audience APIs” (Jerath and Miller 2024).

Second, PETs may have competitive consequences because firms with fewer customers have fewer data points. Small firms may require greater transformation of those data to protect individual privacy, exacerbating challenges for inference (Komarova and Nekipelov 2020). Consequently, many real-world applications choose permissive privacy parameters that effectively sacrifice privacy for utility (Blanco-Justicia et al. 2022; Williams and Bowen 2023). Many small firms also lack the technical expertise to implement PETs.

Third, differential privacy may indirectly leak sensitive information, in particular for disadvantaged consumers (Chang and Shokri 2021). For these reasons, scholars criticize the use of differential privacy in the U.S. census (Hotz et al. 2022).

Furthermore, PETs do not alleviate some of the other unintended consequences of privacy regulation. Korganbekova and Zuber (2023) show that a privacy-preserving algorithm mitigates, but does not eliminate, the tendency for privacy restrictions to cause greater harm to small sellers and price-sensitive consumers. Moreover, extant PETs do not address the problem that firms know too little about disadvantaged consumers, effectively excluding them from the marketplace.

7.2 Forward-looking Regulation

Privacy regulation had the intended consequence of favoring innovation in PETs; but U.S. regulatory proposals to date (to our knowledge) omit PETs. For instance, the FTC's request for public comment on "Commercial Surveillance and Data Security" only mentions PETs in passing.

Since PETs are costly for firms to implement, forward-looking regulation should consider how to incentivize PET adoption and innovation further. For instance, regulation could include appropriate PET use as a sufficient legal basis for data processing. Regulation could stipulate that firms can forgo costly consent collection if they employ PETs. In settings where consent plays an important role, regulation could incentivize PET adoption by permitting consent defaults that advantage data collection (e.g., opt-out rather than opt-in consent). In contrast, the French regulator CNIL has stated that the consent standard should be the same for Privacy Sandbox-enabled online advertising as third-party cookies.

8 Conclusion

Herein, we have summarized key themes in the relevant academic marketing literature regarding potential intended and unintended consequences from current privacy measures. Many policies reduce the usefulness of consumer data to both consumers and firms by eliminating the

ability to track sources of heterogeneity. On the demand side, these policies weaken personalized marketing; this can reduce value creation for consumers with niche tastes and potentially exclude marginalized consumer segments. On the supply side, these regulations can stifle innovation and reduce the competitiveness of markets, especially hampering small businesses and entrepreneurs. While PETs offer potential to reduce some of these documented costs to consumers and firms, these technologies are likely to advantage larger firms.

Public policy needs to weigh the tradeoffs between the costs and benefits to consumer data privacy restrictions. The unintended costs are frequently omitted from data privacy regulation proposals. A similar balanced approach has been recommended in the past in the discussion of the tradeoffs between innovation and privacy (e.g., Goldfarb and Tucker 2012). Additionally, we offer several suggestions to better understand when surveys do or do not elicit reliable privacy preferences across contexts.

Three topics deserve further research by marketing scholars. First, few papers have quantified the tangible economic benefits of privacy regulations to consumers. Second, more work is needed on the redistributive effects of privacy regulation. Third, research and policy attention are needed on boundary regulation. Heterogeneous consumers need help to efficiently and accurately assess when data-sharing is in their interests and to ease the tasks of data sharing when it is helpful. Technical solutions are needed that allow consumers to “own” their own data locally and to share with selected providers only for a limited purpose and time.

References

- Acquisti A (2024) The economics of privacy at a crossroads. Goldfarb A, Tucker C, eds. *The Economics of Privacy* (University of Chicago Press, Chicago).
- Acquisti A, Brandimarte L, Hancock J (2022) How privacy's past may shape its future. *Science*. 375(6578):270–272.
- Acquisti A, Taylor C, Wagman L (2016) The economics of privacy. *J. Econom. Literature*. 54(2):442–492.
- Adjerid I, Acquisti A, Telang R, Padman R, Adler-Milstein J (2016) The impact of privacy regulation and technology incentives: the case of health information exchanges. *Management Sci.* 62(4):1042–1063.
- Agrawal A, Gans J, Goldfarb A (2019) Economic policy for artificial intelligence. Lerner J, Stern S, eds. *Innovation Policy and the Economy* (University of Chicago Press, Chicago).
- Ali M, Sapiezynski P, Bogen M, Korolova A, Mislove A, Rieke A (2019) Discrimination through optimization: How Facebook's ad delivery can lead to biased outcomes. *Proc. ACM Hum.-Comput. Interact.* 3(CSCW):1-30.
- Allcott H, Diamond R, Dubé JP, Handbury J, Rahkovsky I, Schnell M (2019) Food deserts and the causes of nutritional inequality. *Quart. J. Econom.* 134(4):1793–1844.
- Aridor G, Che YK, Hollenbeck, B, McCarthy D, Kaiser M (2024) Evaluating the impact of privacy regulation on e-commerce firms: evidence from Apple's App Tracking Transparency * (June 13, 2024). Available at SSRN: <https://ssrn.com/abstract=4698374>
- Arslan HA, Tereyağoğlu N, Yılmaz Ö (2023) Scoring a touchdown with variable pricing: Evidence from a quasi-experiment in the NFL ticket markets. *Management Sci.* 69(8):4435-56.
- Ascarza E, Israeli A (2022) Eliminating unintended bias in personalized policies using bias-eliminating adapted trees (BEAT). *Proc. Natl. Acad. Sci.* 119(11).
- Auxier B, Rainie L, Anderson M, Perrin A, Kumar M, Turner E (2019) Americans and privacy: Concerned, confused and feeling lack of control over their personal information. *Pew Research Center*. Accessed March 25, 2024, <https://www.pewresearch.org/internet/2019/11/15/americans-and-privacy-concerned-confused-and-feeling-lack-of-control-over-their-personal-information/>
- Bass FM (1995) Empirical generalizations and marketing science: A personal view. *Marketing Sci.* 14(3):G6-G19.
- Baviskar S, Chowdhury I, Deisenroth D, Li B, Sokol D (2024) ATT vs. personalized ads: user's data sharing choices under Apple's divergent consent strategies (June 28, 2024). Available at SSRN: <https://ssrn.com/abstract=4887872>.
- Bergemann D, Brooks B, Morris S (2015) The limits of price discrimination. *Amer. Econom. Rev.* 105(3):921–957.

Bergemann D, Brooks B, Morris S (2024) On the alignment of consumer surplus and total surplus under competitive price discrimination. Working Paper, Yale University, New Haven.

Beveridge, C (2024) European data act – key provisions and their implications. Accessed May 15, 2024, <https://www.bdo.co.uk/en-gb/insights/advisory/risk-and-advisory-services/european-data-act-key-provisions-and-their-implications#>.

Bian B, Pagel M, Tang H, Raval D (2023) Consumer surveillance and financial fraud. National Bureau of Economic Research, https://www.nber.org/system/files/working_papers/w31692/w31692.pdf

Biden, Joseph R. (2023) Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. Accessed May 15, 2024, <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.

Blanco-Justicia A, Sánchez D, Domingo-Ferrer J, Muralidhar K (2022) A critical review on the use (and misuse) of differential privacy in machine learning. *ACM Comput. Surv.* 55(8):1—16.

Blattner L, Nelson S (2021) How costly is noise? Data and disparities in consumer credit. arXiv:2105.07554.

Bleier A, Goldfarb A, Tucker C (2020) Consumer privacy and the future of data-based innovation and marketing. *Int. J. Res. Mark.* 37(3):466-80.

Bronnenberg BJ, Dhar SK, Dubé JP (2009) Brand history, geography, and the persistence of brand shares. *J. Political Econom.* 117(1):87-115.

Bronnenberg BJ, Dubé JP, Joo J (2022) Millennials and the takeoff of craft brands: Preference formation in the U.S. beer industry. *Marketing Sci.* 41(4):710-732.

Chang, H., & Shokri, R. (2021). On the privacy risks of algorithmic fairness. *IEEE European Symposium on Security and Privacy (EuroS&P)*. 292-303.

Chintagunta KP, Erdem T, Rossi PE, Wedel M (2006) Structural modeling in marketing: review and assessment. *Marketing Sci.* 25(6):604-616.

Civil Rights Litigation Clearinghouse (2022) Resource: Justice Department Secures Groundbreaking Settlement Agreement with Meta Platforms, Formerly Known as Facebook, to Resolve Allegations of Discriminatory Advertising. Accessed [May 15, 2024], <https://clearinghouse.net/resource/3807/>

Clark, C (2020) How “notice and consent” fails to protect our privacy. Accessed May 15, 2024, <https://www.newamerica.org/oti/blog/how-notice-and-consent-fails-to-protect-our-privacy/>.

Consumer Financial Protection Bureau (2022a) Hold credit reporting companies accountable for incorrect reports and shoddy service. Accessed May 15, 2024, <https://www.consumerfinance.gov/about-us/blog/hold-credit-reporting-companies-accountable-incorrect-reports-shoddy-service>

- Consumer Financial Protection Bureau (2022b) CFPB targets unfair discrimination in consumer finance. Press release, Consumer Financial Protect Bureau, Washington, DC.
- Council of Economic Advisors (2014) Big data: Seizing opportunities, preserving values. Report, Council Econ, Advisors, Washington, DC.
- Council of Economic Advisors (2015) Big data and differential pricing.” Report, Council Econ. Advisors, Washington, DC.
- Cramer-Flood E (2023) Meta’s ad revenue share vastly exceeds its share of consumer time. eMarketer. Accessed May 15, 2024, <https://www.emarketer.com/content/meta-s-ad-revenue-share-vastly-exceeds-its-share-of-consumer-time>.
- Datta H, Knox G, Bronnenberg B (2018) Changing their tune: How consumers’ adoption of online streaming affects music consumption and discovery. *Marketing Sci.* 37(1):5-21.
- DellaVigna S, Gentzkow M (2019) Uniform pricing in us retail chains. *Quart. J. Econom.* 134(4):2011-2084.
- Demirer M, Hernández DJ, Li D, Peng S. (2024) Data, privacy laws and firm production: Evidence from the GDPR. Working paper, National Bureau of Economic Research, Cambridge.
- Derksen L, McGahan A, Pongeluppe L (2022) Privacy at what cost? Using electronic medical records to recover lapsed patients into HIV care. Goldfarb A, Tucker C, eds. *Proceedings NBER Workshop on the Economics of Privacy* (National Bureau of Economic Research).
- Deisenroth D, Manjeer U, Sohail Z, Tadelis S, Wernerfelt N (2024) Digital advertising and market structure: implications for privacy regulation. National Bureau of Economic Research. <https://www.nber.org/papers/w32726>
- Dubé JP, Misra S (2023) Personalized pricing and consumer welfare. *J. Political Econ.* 131(1):131-89.
- eMarketer Editors (2019) CPG industry struggles to find growth—eMarketer trends, forecasts & statistics.
- Farley, JU, Lehmann DR, Sawyer A (1995) Empirical marketing generalization using meta-analysis. *Marketing Sci.* 14(3):G36-G46.
- Faverio M (2023) Key findings about Americans and data privacy. Accessed March 25, 2024, <https://www.pewresearch.org/short-reads/2023/10/18/key-findings-about-americans-and-data-privacy/>.
- Feldman JM, Lynch JG (1988), Self-generated validity and other effects of measurement on belief, attitude, intention, and behavior. *J Applied Psych*, 73(3): 421-435.
- Freedman S, Jin G (2017) The information value of online social networks: Lessons from peer-to-peer lending. *Int. J. Ind. Organ.* 2017, 51:185-222.

FTC (2024) A Look Behind the Screens: Examining the Data Practices of Social Media and Video Streaming Services. *Federal Trade Commission*, 2024.

GDPR (2020) General Data Protection Regulation. Accessed May 15, 2024, <https://gdpr.eu/tag/gdpr/>

Geng, T., Dawson, M., & Nair, H. (2023) Effectively Combining the Event and Aggregate Summary Reports from the Privacy Sandbox Attribution Reporting API for Improving Ad-Measurement Fidelity. Technical report, Google Ads.

Glenn B, Dubé JP, Kavanagh SC (2022) Segmented pricing for fines and fees. Accessed May 15, 2024, <https://www.gfoa.org/materials/segmented-pricing>

Goldfarb A, Que VF (2023) The economics of digital privacy. *Annu. Rev. Econom.* 15:267-86.

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

Goldfarb A, Tucker C (2012) Privacy and innovation. Vol 12. Lerner J, Stern S, eds. *Innovation Policy and the Economy* (University of Chicago Press, Chicago), 65-90.

Goldfarb A, Tucker C. (2017) Inequality, privacy and digital market design. Kominers S, Teytelboym A eds. *Fair by Design* (Oxford University Press, Oxford).

Google (2022). The privacy sandbox: technology for a more private web. Accessed May 15, 2024, <https://privacysandbox.com>.

Hard A, Rao K, Mathews R, Ramaswamy S, Beaufays F, Augenstein S, Eichner H, Kiddon C, Ramage D (2018) Federated learning for mobile keyboard prediction. arXiv preprint arXiv:1811.03604.

Hotz VJ, Bollinger CR, Komarova T, Manski CF, Moffitt RA, Nekipelov D, Sojourner A, Spencer BD (2022) Balancing data privacy and usability in the federal statistical system. *Proc. Natl. Acad. Sci.* 119(31):e2104906119.

Howe LC, Krosnick JA (2017) Attitude strength. *Ann. Rev. Psych.* 68:327-51.

Information Commissioners Office (2024). Case studies and examples. Accessed at <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/case-studies-and-examples> on January 15, 2025.

Janssen R, Kesler R, Kummer ME, Waldfogel J (2022) GDPR and the lost generation of innovative apps. NBER Working Paper No. 30028, National Bureau of Economic Research, Cambridge, MA.

Jerath K, Miller KM (2024) Consumers' perceived privacy violations in online advertising. Marketing Science Institute Working Paper.

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

- Johnson, EJ, Bellman S, Lohse GL (2002) Defaults, framing and privacy: Why opting in-opting out, *Mark. Lett.* 13 (1), 5–15.
- Johnson GA (2024) Economic research on privacy regulation: Lessons from the GDPR and beyond. Goldfarb A, Tucker C eds. *The Economics of Privacy* (University of Chicago Press, Chicago).
- Johnson GA, Lin T, Cooper J, Zhong L (2024) COPPAcalypse? The YouTube settlement’s impact on kid’s content. Working paper.
- Johnson GA, Neumann N (2024) The advent of privacy-centric digital advertising: Tracing privacy-enhancing technology adoption. Working paper.
- Johnson GA, Runge J, Seufert E (2022) Privacy-centric digital advertising: Implications for research. *Cust. Needs Solut.* 9(1):49—54.
- Johnson GA, Shriver SK, Du S (2020) Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Sci.* 39(1):33—51.
- Johnson GA, Shriver SK, Goldberg SG (2023) Privacy and market concentration: intended and unintended consequences of the GDPR. *Management Sci.* 69 (10): 5695-5721.
- Kaplan L, Mislove A, Sapiezynski P (2017) Measuring Biases in a Data Broker’s Coverage. *Proceedings of FTC Conference* (Washington DC).
- Kesler R (2023) The Impact of Apple's App Tracking Transparency on App Monetization. Working paper, University of Zurich, Zurich, Switzerland.
- Kerrigan, K., and R. Keating (2019) Online Advertising Delivers BIG Benefits for Small Businesses. Accessed March 25, 2024, <https://sbecouncil.org/2019/09/10/online-advertising-delivers-big-benefits-for-small-businesses/>.
- Khera, Purva (2023). India’s financial system: building the Foundation for Strong and Sustainable Growth. *International Monetary Fund*, Accessed January 15, 2025 at <https://www.elibrary.imf.org/display/book/9798400223525/CH007.xml>
- King J, Ho D, Gupta A, Wu V, Webley-Brown H (2023) The Privacy-Bias Tradeoff: Data Minimization and Racial Disparity Assessments in US Government. *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Association for Computing Machinery, New York), 492-505.
- Kircher T, Foerderer J (2023) Does privacy undermine content provision and consumption? evidence from educational YouTube channels. Working paper, Technical University of Munich, Munich, Germany.
- Komarova T, Nekipelov D (2020) Identification and formal privacy guarantees. arXiv preprint arXiv:2006.14732.
- Korganbekova M, Zuber C (2023) Balancing user privacy and personalization. Working paper, Northwestern University Kellogg School of Management, Chicago.

- Lambrecht A, Seim K, Tucker C (2011) Stuck in the adoption funnel: The effect of interruptions in the adoption process on usage. *Marketing Sci.* 30(2):355-367.
- Lambrecht A, Tucker C (2019) Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. *Management Sci.* 65(7):2966-2981.
- Lambrecht A, Tucker C (2024) Apparent algorithmic discrimination and real-time algorithmic learning in digital search advertising. *Quant. Marketing & Econom.* (2024): 1-31.
- Larson RB (2023). Privacy concerns and social desirability bias. *Int. J. Mark. Res.* 0(0).
- Lee JY, Yang J, Anderson E (2024) Who benefits from alternative data for credit scoring? Evidence from Peru. Available at SSRN: <https://ssrn.com/abstract=4852032>
- Lin T (2022) Valuing intrinsic and instrumental preferences for privacy. *Marketing Sci.* 41(4):663–681.
- Lin T, Strulov-Shlain A (2024) Choice architecture, privacy valuations, and selection bias in consumer data. *Marketing Sci.* in press.
- Mela, Carl F. (20011) Data Selection and Procurement. *Marketing Sci.* 25(6):965-976.
- Miller AR, Tucker C (2009) Privacy protection and technology diffusion: The case of electronic medical records. *Management Sci.* 55(7):1077-1093.
- Miller AR, Tucker CE (2011) Can health care information technology save babies? *J. Political Econ.* 119(2):289–324.
- Miller AR, Tucker CE (2018) Privacy protection, personalized medicine, and genetic testing. *Management Sci.* 64(10):4648-4668.
- Misra S (2020) Algorithmic nudges. Working paper, University of Chicago.
- Mrkva K, Posner NA, Reek C, Johnson EJ (2021) Do nudges reduce disparities? Choice architecture compensates for low consumer knowledge. *J. Marketing.* 85(4):67-84.
- Neumann N, Tucker CE, Kaplan L, Mislove A, Sapiezynski P (2024) Data deserts and black boxes: The impact of socio-economic status on consumer profiling. *Man. Sci.* 0(0).
- Netzer O, Lemaire A, Herzenstein M (2019) When words sweat: Identifying signals for loan default in the text of loan applications. *J. Marketing Res.* (6):960-980.
- Noble, T (2024) To Fight Surveillance Pricing, We Need Privacy First. *Electronic Frontier Foundation.* August 5, <https://www.eff.org/deeplinks/2024/08/fight-surveillance-pricing-we-need-privacy-first>
- Perik N (2024) ATT Opt-in Rates in 2024 & Best Practices to Increase Them. MAF. Accessed January 23, 2025 at <https://maf.ad/en/blog/att-opt-in-rates-boost/>.
- Peukert C, Bechtold S, Batikas M, Kretschmer T (2022) Regulatory spillovers and data governance: Evidence from the GDPR. *Marketing Sci.* 41(4):746-68.

- Ponte GR, Wieringa JE, Boot T, Verhoef PC (2024) Where's Waldo? A framework for quantifying the privacy-utility tradeoff in marketing applications. *International Journal of Research in Marketing*.
- Ramasastri A (2005) Websites that charge different customers different prices: Is their 'price customization' illegal? Should it be? Accessed May 15, 2024, <https://supreme.findlaw.com/legal-commentary/websites-that-charge-different-customers-different-prices.html>.
- Rossi PE, Allenby GM (2003) Bayesian statistics and marketing. *Marketing Sci.* 22(3):304-328.
- Savage SJ, Waldman DM (2015) Privacy tradeoffs in smartphone applications. *Economics Letters*, 137 (Dec.): 171-175.
- Schwarz N (1999) Self-reports: How the questions shape the answers. *Amer. Psychologist*. 54(2):93.
- Schuman H, Presser S (1980) Public opinion and public ignorance: The fine line between attitudes and nonattitudes. *Am. J. Sociology.* ;85(5):1214-25
- Simmons CJ, Bickart B, Lynch JG (1993), "Capturing and creating public opinion in survey research. *J. Consumer Research*, 20(2): 316-329.
- Stole LA (2007) Price Discrimination and Imperfect Competition. , vol. 3, Armstrong M, Porter RH, eds. *Handbook of Industrial Organization* (North-Holland, Amsterdam), 2221–2299.
- Sun T, Yuan Z, Li C, Zhang K, Xu J (2023) The value of personal data in internet commerce: A high-stakes field experiment on data regulation policy. *Management Sci.* 70(4):2645-2660.
- Sutton J (1991) *Sunk costs and market structure: Price competition, advertising, and the evolution of concentration* (MIT Press, Cambridge).
- Thaler RH, Sunstein CR (2021) *Nudge: The Final Edition* (Yale University Press, New Haven).
- Tucker C (2023) Algorithmic exclusion: The fragility of algorithms to sparse and missing data. Report, Brookings Center on Regulation and Markets at Brookings, Washington, DC, 1-26.
- Tucker C (2024) "The economics of privacy: An agenda." *Economics of Privacy* (University of Chicago Press, Chicago)
- Tucker CE, Wang, Y, Yu, S (2023). Does IT lead to more equal treatment? An empirical study of the effect of smartphone use on customer complaint resolution. Accessed January 22, 2025, <https://ssrn.com/abstract=4499402>.
- Utz C, Degeling M, Fahl S, Schaub F, Holz T (2019) (un)informed consent: Studying gdpr consent notices in the field. *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security* (Association for Computing Machinery, New York), 973–990.
- Varian H, Wallenberg F, Woroch W (2005) The demographics of the do-not-call list [security of data]. *IEEE Security & Privacy*. 3(1): 34-39.

Wernerfelt N, Tuchmann A, Shapiro B, Moakler R A (2024) Estimating the value of offsite data to advertisers: evidence from Meta. *Mar. Sci.*, in press

Williams AR, Bowen CM (2023) The promise and limitations of formal privacy. *WIREs Comp. Stats.* 15(6):e1615.

Web Appendix – Background on Key Themes in “Intended and Unintended of Privacy Regulation for Consumer Marketing”

Section 1: Introduction

Privacy regulatory landscape

GDPR-like privacy laws have been implemented by Australia, Brazil, Canada, Chile, China, Egypt, India, Israel, Japan, Kenya, New Zealand, Nigeria, South Africa, South Korea, Switzerland, Thailand, Turkey, and the UK (Zafar 2023).

References

Zafar, F (2023) 18 countries with GDPR-like data privacy laws. Yahoo! Finance, September 14, 2023. Accessed March 25, 2024, <https://finance.yahoo.com/news/18-countries-gdpr-data-privacy-121428321.html>.

Section 2: Intended Benefits of Digital Marketing Privacy Regulation and Pertinent Regulations

Consumers may be harmed if they cannot correct inaccurate information about them

If firms act on inaccurate information about consumers, they may be harmed. For example, the U.S. Consumer Financial Protection Bureau recently fined Equifax for failing to investigate and process information disputed by consumers, causing inaccurate credit scores (CFPB 2025a). In another action, the CFPB fined Honda Finance for allowing customers to defer payments during COVID but then incorrectly telling credit reporting agencies that these customers were delinquent (CFPB 2025b).

Firms might discriminate against protected classes

Numerous additional examples exist beyond those in the main body of the paper showing firms discriminating against protected classes. Class actions have been filed against healthcare companies for allegedly disclosing protected health information to Meta's pixels, used to track individuals' browsing behavior (Asplund 2024). Online matching platforms rely on ratings of buyers and sellers. Human inputs to recommendation systems can also contribute to discrimination via those systems. A study of an online freelance worker platform found that female freelancers received lower rating scores than men (Bairathi et al. 2023). This gap may stem from discrimination and reflect the stereotypes of those individuals who submitted the ratings. The gap is wider in countries with lower gender equality and in markets with lower female labor force participation rates. Israeli and Ascarza (2020) list additional examples of bias organized around the 4 P's of product, price, place, promotion, often due to incomplete or unrepresentative training data for underrepresented groups.

Sapiezynski et al. (2022) show experimentally that merely removing demographic features from a real-world algorithmic system's inputs can fail to prevent biased outputs. As a result, organizations using algorithms to help mediate access to important life opportunities should consider other approaches to mitigating discriminatory effects. This paper provides justification for the Bias-Eliminating Adaptive Tree approach by Ascarza and Israeli (2022) cited in the main body of the paper.

Firms might price discriminate against consumers with higher valuations

Firms can infer consumer valuations from historic purchase data and use that information to choose personalized price or discount levels (Rossi, McCullough, and Allenby 1996). With more refined algorithmic personalized pricing, firms can unintentionally discriminate along socially controversial segment boundaries. For example, Princeton Review charged higher prices in ZIP codes with many Asians (Angwin et al. 2015). Behavioral-based price discrimination” (Fudenberg and Villas-Boas 2006) can lead to a “ratchet effect” even when consumers attempt to protect their privacy (Hart and Tirole 1988). For these reasons, the Council of Economic Advisors (2014) explains:

“Consumers have a legitimate expectation of knowing whether the prices they are offered for goods and services are systematically different than the prices offered to others.”

References

- Angwin C, Mattu S, Larson J (2015), Test prep is more expensive – for Asian students. *The Atlantic* (September 3, 2015). Accessed May 15, 2024.
<https://www.theatlantic.com/education/archive/2015/09/princeton-review-expensive-asian-students/403510/>
- Ascarza E, Israeli A (2022) Eliminating unintended bias in personalized policies using bias-eliminating adapted trees (BEAT). *Proc. Natl. Acad. Sci.* 119(11).

- Asplund, John (2024) VillageMD facing privacy lawsuit over use of Meta's 'Pixels', *Crain's Chicago*, April 11, 2024. Accessed at https://www.chicagobusiness.com/health-pulse/villagemd-facing-privacy-lawsuit-over-use-metas-pixels?utm_source=Sailthru&utm_medium=email&utm_campaign=Newsletter-Health-BreakingNews-20240411 on April 11, 2024.
- Bairathi M, Lambrecht A, Zhang X (2023). Gender Disparity in Online Reputation: Evidence from an Online Freelance Platform. Working paper.
- Council of Economic Advisors (2014) Big data: Seizing opportunities, preserving values. Report, Council Econ, Advisors, Washington, DC.
- CFPB (2025a) CFPB orders Equifax to pay \$15 million for improper investigations of credit reporting errors. US Consumer Financial Protection Bureau, Jan 17, 2025. <https://www.consumerfinance.gov/about-us/newsroom/cfpb-orders-equifax-to-pay-15-million-for-improper-investigations-of-credit-reporting-errors>
- CFPB (2025b) CFPB orders Honda's auto financing arm to pay \$12.8 million for COVID-19 and other credit reporting failures. US Consumer Financial Protection Bureau, Jan 17, 2025. <https://www.consumerfinance.gov/about-us/newsroom/cfpb-orders-hondas-auto-financing-arm-to-pay-128-million-for-covid-19-and-other-credit-reporting-failures>
- Fudenberg D, Villas-Boas JM (2006) Behavior-based price discrimination and customer recognition. *Handbook on Economics and Information Systems* (Elsevier Science, Oxford), 377-436.
- Hart OD, Tirole J (1988) Contract renegotiation and Coasian dynamics. *The Review of Economic Studies*. 55(4):509-40
- Israeli A, Ascarza E. (2020) "Algorithmic bias in marketing. Harvard Business School Technical Note 521-020, September 2020. (Revised July 2022.)
- Rossi PE, McCulloch RE, Allenby GM (1996) *The value of purchase history data in target marketing. Marketing Science*. 15(4):321-40.
- Sapiezynski P, Ghosh A, Kaplan L, Rieke A, Mislove A. Algorithms that " Don't See Color" Measuring Biases in Lookalike and Special Ad Audiences. In Proceedings of the 2022 AAAI/ACM Conference on AI, Ethics, and Society 2022 Jul 26 (pp. 609-616).

Section 4: Which Consumers Care Most About Privacy, and Do Privacy Policies Unintentionally Favor the Privileged?

Policy reliance on survey data

It is commonplace to use survey results to justify privacy policy recommendations (see e.g. Wheeler 2018, 2023). These surveys appear to show broad support for further privacy regulation (e.g. Arbanas et al. 2023; Consumer Reports 2017).¹

We cited the FTC (2024) report as an example of a policy analysis relying on consumer survey data rather than marketing science and empirical causal evidence quantifying benefits and costs of privacy restrictions for consumers. For example, FTC (2024, p. 40) cited a “Responsible Tech” survey (Greenberg et al 2021), concluding:

“In recent survey data, consumers expressed strong concerns about the use of certain categories of information for ad targeting. For example, according to a 2021 study, 73% of consumers were opposed to companies tracking online behavior and collecting personal data in order to serve targeted ads. This same study revealed that 56% of consumers surveyed were opposed to companies displaying ads based on age, gender, and general location.”

In our paper, we questioned the wisdom of relying on these surveys as a foundation for policies further discouraging data sharing. We pointed to many studies documenting a “privacy paradox” – survey respondents profess caring deeply about privacy but willingly share data (e.g. Spiekerman et al. 2001). We see the privacy paradox as an instance of the general phenomenon of

¹ Below we offer a general criticism of how survey research has been used in setting privacy policy. The Greenberg et al (2021) survey cited by FTC has the further problem of being a textbook example of leading questions and biased “advocacy research.” Questions about privacy and personalization in advertising came later in the survey after asking questions about a list of controversial figures and companies, each likely to provoke distrust in some political subset of Americans.

attitude-behavior inconsistency (Ajzen et al. 2018). It is well known that attitudes / values and behavior divergent most when the attitude is not activated when making decisions and when knowledge is low (Davidson et al. 1985). Dalmia and Diehl (2024) documented that consumers faced with decisions where they could choose whether to share data did not even think about privacy; its importance paled next to other priorities in online search and shopping.

Moreover, surveys typically focus mostly on information security concerns (e.g., fraud, sharing with malign actors, etc.) that are not tied to marketing and issues of privacy. Surveys cited rarely ask about specific marketing tactics associated with the broad “privacy” label. Most surveys do not provide clear guidance to regulators about which consumers believe they would benefit by reducing online retailer personalization, by removing access to data from prior months on the same website, or by blocking an online merchant’s ability to track the consumer’s path from clicking on an ad to a purchase of the advertised product.

Surveys are unreliable foundation for policy if respondents construct responses on the spot

It is naïve, we claim, to assume that consumers have attitudes and opinions on any and every given topic. One must grapple with the fact that many survey respondents have almost no relevant knowledge and have not thought about the issues queried before being confronted by survey researchers. In the political and policy arena, Converse (1960) introduced the concept “non-attitudes” -- responses given by individuals who answer survey questions without understanding the issue or having a genuine opinion. This paper was the inspiration for Schumann and Presser (1980), cited in the body of our manuscript, who claimed that political poll respondents have well-articulated views only on a small number of highly personally relevant issues.

Many lines of evidence support that conclusion. Sloman and Fernbach (2017) summarize research showing that people operate effectively in the world despite the thinness of understanding about familiar and ordinary objects or policies in their world. Amusingly, most people cannot explain how a toilet works nor draw a bicycle in a way that plausibly shows how it creates locomotion. More materially, across a range of policy issues, most people cannot explain why their preferred policy would cause some claimed outcome (Fernbach et al. 2013, 2019).

When the recommendation of a privacy measure uses consumer opinion surveys to justify policy details, the authors often overlook the following important and well-established point: when survey respondents lack pre-determined attitudes or cannot retrieve them in the moment, they “construct” preferences on the spot. The hallmark of constructed preferences is that behaviors or survey responses are highly sensitive to seemingly minor changes in context or question wording (Bettman, Luce, and Payne 1998; Feldman and Lynch 1988; Schwarz 1999, Simmons, Bickart, and Lynch 1993). Expressed values are “labile.” Acquisti et al. (2013) found evidence that the level of instability for valuations of privacy was greater than that for normal consumer goods.

The privacy paradox is a disconnect between stated preferences and “revealed preference” via behavior. When preferences are constructed, both stated and revealed preferences are constructed, but based on different inputs. Neither is a gold standard showing some underlying true and stable preference, because stable preferences do not exist. In those circumstances, consumer choices are particularly influenced by “nudging” and “choice architecture” (Thaler and Sunstein 2021). We noted in our paper that an opt-in policy that uses “no” as the default option for data tracking leads to much lower consent than an opt-out policy that uses “yes” as the default option for data tracking (Johnson, Bellman, and Lohse 2002).

Thaler and Sunstein characterize nudges as very light touch “libertarian paternalism.” However, McKenzie et al. (2006) show that consumers interpret “opt in” vs “opt out” as advice from policymakers that prudent consumers should choose the default. In section 4 we cited data showing heterogeneity of consumer privacy preferences. The normative message conveyed by opt in requirements in GDPR and similar policies raises concerns that opt in is not neutral and favors those for whom the costs of data sharing outweigh the benefits. We develop this argument further in the next section.

Boundary regulation vs. concealment as the key in privacy

Altman (1977) argues that privacy regulation mechanisms are boundary control process where people sometimes make themselves open and accessible to others and sometimes close themselves off from others. Though his analysis applies to interpersonal relationships, we believe it also may be extended to relationships of individuals with nonhuman technological systems. People may experience motives to be open or closed with digital marketing powered by algorithms blazing away at billions of transactions per second in remote data centers when no human is plausibly looking at their personal data.

Similarly, “contextual integrity” refers to the phenomenon of wanting disclosed information to be used only in approved contexts (Nissenbaum 2004). For instance, people who say “I have nothing to hide” may nonetheless tilt their laptop screen away from strangers seated next to them on a plane. The same people might gladly share their laptop screens with colleagues on collaborative Zoom calls or upload confidential Zoom call transcripts to ChatGPT for summarization. People who prefer that their work colleagues not know about a health condition might wish that HIPAA medical privacy regulations made it easier to share their unified health histories with new specialists or with family members and their doctors. In these examples, privacy is about whether one’s

personal data is used in contexts and by other entities where the consumer reasonably expects it to be shared; but not otherwise. Therefore, we regard sweeping statements that consumers do or do not value privacy as unhelpful because they overlook the role of context.

Privacy policies support “boundary regulation” by reducing the cost of sharing when sharing is desired, not just minimizing data sharing when one would prefer not to share. Consider health information privacy. For decades, technical barriers to interoperability in disparate Electronic Health Record systems and a lack of financial incentives deterred providers from addressing the problem. As a result, patients would find it laborious to allow a current health provider to access the records from a former provider. A patient would also find it cumbersome to share health information with family members whose own care might be informed by family health history. Patients visiting their primary care physicians would be repeatedly asked the same questions on each visit, forcing them to reenter health history manually and based on memory. All of these examples are failures of boundary regulation. In the last decade, the Epic electronic health records system introduced two new features to make safe data sharing easier. “Invite Friends & Family” reduces the frictions to authorizing a close friend or family member to have access to one’s medical records. “Share Everywhere” gives one-time access to any clinician in the world. These changes are very much in the spirit of boundary regulation.

In the domain of personalized advertising and personalized retail recommendations, a chief obstacle to boundary regulation is an effort / accuracy tradeoff faced by the consumer. Consumers routinely trade off the costs and benefits of search, deciding when to have some loss in “accuracy” to achieve less effort to make some decision (Payne et al. 1993). For many decisions, it may often be optimal to engage in little or no search (Moorthy et al. 1997). Alba et al. (1997) argue:

“Retailers and retail formats compete in the types of information they convey effectively to customers. Just as in Erlich and Fisher’s (1982) analysis of ‘derived demand for advertising’, we analyze derived demand for retailer information about products. Erlich and Fisher argue that buyers demand information from sellers reduce ...the costs of obtaining information about products and of dissatisfaction from disappointing purchases. Consumers demand information that reduces this wedge. Such information alternatively can be derived from their own prior knowledge, advertising, or “other selling efforts”-notably information from retailers.” (pp. 40-41)

Alba et al. (1997 p. 42) address whether sellers can learn enough about a customer’s preferences to offer value via algorithmic personalization. These authors compare methods relying on passive tracking of a consumer’s revealed preferences versus asking consumers about their preferences directly – at a cost of more effort by the consumer. They note: *“The usefulness of these customized approaches will depend on the consumer effort necessary to calibrate the screening mechanism and the accuracy with which the mechanism correlates with the consumer’s full utility function.”*

The same framework can be used to analyze how elements of consumer privacy protection facilitate the consumer’s boundary regulation goals. Data deletion, data minimization policies, and constraints on the use of third-party data all reduce the accuracy of algorithmic predictions of what a consumer might like, potentially increasing transaction costs. For example, deleting old orders from a food delivery app increases search costs for a consumer who wants to reorder favorites. Deleting credit card information increases transaction costs. However, for many but not all consumers, this sacrifice might be worthwhile if either they attach little value to those benefits or they perceive a disproportionately high intrinsic or instrumental privacy benefits.

Responding to notice and consent request and opt in policies creates costs of thinking. The increasingly ubiquitous use of consent notification seems to imply that regulators assume benefits of

careful consideration typically outweigh the cognitive costs. However, if consumers do not perceive a material benefit to scrutinizing a given opt in disclosure, they will exert very low effort to respond to the required opt-in query.

Our concern is that privacy policies are often less focused on “privacy as boundary regulation” and more focused on “privacy as concealment.” The net effect is to decrease the likelihood of data sharing or transmission. If regulators were focused on “boundary” regulation, they would consider methods to reduce the costs of deciding about individual data sharing requests while maximizing the “accuracy” of those decisions in terms of the consumers’ experienced utility.

Once again, choice architecture can influence a consumer’s calculus about the costs vs. benefits of data sharing; albeit in potentially socially undesirable ways. Farronto et al. (2024) examine the effects of “dark patterns” -- interface designs that encourage data sharing – for a web browser extension. In their field experiment, participants installed a browser extension that randomized cookie consent interface designs as they browsed the internet. In the absence of dark patterns, consumers accept all cookies over half of the time. Further, over half of the consumers vary their choices across websites, demonstrating the need for boundary regulation. However, dark patterns that introduce frictions into the consent decision increase the cookie accept rate considerably. Using a choice model, the authors infer a cost from clicking ‘customize settings’ that outweighs consumers’ utility of choosing the preferred sharing option.

Algorithmic discrimination, and the effect of inadequate data on detecting and correcting algorithmic discrimination

Bogen et al (2022, p.492) discuss the challenges in resolving algorithmic exclusion: “Some companies are required by law to collect sensitive attribute data, while others are prohibited from doing so. Still others, in the absence of legal mandates, have determined that collection and

imputation of these data are appropriate to address disparities.” Rieke et al. (2022) compare existing methods to infer the distribution of sensitive variables for monitoring purposes, finding mixed results on effectiveness. For instance, inferring a consumer’s race is much easier when the data include photos. When the firm observes a consumer’s surname and geographic location, Bayesian Improved Surname Geocoding (BISG) methods can be used to infer race. Even with these unique data, Rieke et al (2022) find these methods work better for some demographic groups than others. Finally, Cecere et al (2025) show that algorithmic monitoring needs to be on-going as one-time audits might generate errors simply because of a lack of predictability of the algorithms themselves. PETs may also facilitate the monitoring of algorithmic discrimination without observing a consumer’ race, gender, or other sensitive data fields (Juarez and Korolova (2023). However, as we discuss in section 7.1 of the paper, PETs also face several limitations.

References

- Acquisti A, John LK, Loewenstein G (2013b) What is privacy worth?
²https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3305331
- Alba J, Lynch J, Weitz B, Janiszewski C, Lutz R, Sawyer A, Wood S (1997) Interactive home shopping: consumer, retailer, and manufacturer incentives to participate in electronic marketplaces. *J. Marketing*. 61(3):38-53.
- Altman, I. (1977) Privacy regulation: Culturally universal or culturally specific? *J. Soc. Issues*. 33(3): 66-84.
- Arbanas J, Silverglate PH, Hupfner S, Loucks J, Raman P, Steinhart M (2023), Data privacy and security worries are on the rise, while trust is down. *Deloitte’s Connected Consumer Survey 2023*. Accessed March 25, 2024, <https://www2.deloitte.com/us/en/insights/industry/telecommunications/connectivity-mobile-trends-survey/2023/data-privacy-and-security.html>.

² The same authors published a paper with the same title in *J. Legal Studies* (2013). The SSRN working paper version has additional studies and analysis beyond what appeared in the journal article, including a finding of bimodal distributions of privacy valuations, with some respondents placing very high values and others very low values.

- Bettman JR, Luce MF, Payne JW (1998) Constructive consumer choice processes. *J. Consumer Res.* 25(3):187-217.
- Bogen M, Rieke A, Ahmed S (2022) Awareness in Practice: Tensions in Access to Sensitive Attribute Data for Antidiscrimination. *FAT* '20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency.* 492-500.
- Cecere G, Jean C, Manant M, Tucker C (2025) The Need for Repeated Testing in Algorithmic Auditing: The Example of Algorithms' Preference for Headless Women. MIT Working Paper.
- Consumer Reports (2017) Americans Want More Say in the Privacy of Personal Data. CR's second Consumer Voices Survey reveals deep concerns about how info is collected and used. May 17, 2017, available at: <https://www.consumerreports.org/electronics-computers/privacy/americans-want-more-say-in-privacy-of-personal-data-a5880786028/>
- Converse PE (1964). The nature of belief systems in mass publics. In D. E. Apter (Ed.), *Ideology and discontent* (pp. 75-169). New York: Free Press.
- Ehrlich E, Fisher L (1982) The derived demand for advertising: A theoretical and empirical investigation, *American Economic Review*, 72 (June), 366-88.
- Farronato C, Fradkin A, Lin T (2024) Data sharing and website competition: the role of “dark patterns.” Dec.16, 2024. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4920040
- Feldman JM, Lynch JG (1988), Self-generated validity and other effects of measurement on belief, attitude, intention, and behavior. *J Applied Psych*, 73(3): 421-435.
- Fernbach PM, Light N, Scott SE, Inbar Y, Rozin P (2019) Extreme opponents of genetically modified foods know the least but think they know the most. *Nature Human Behaviour*, 3(3):251-256.
- Fernbach PM, Rogers T, Fox CR, Sloman SA (2013) Political extremism is supported by an illusion of understanding. *Psychological Sci.*,24(6):939-46.
- Juarez M, Korolova A. (2023) You can't fix what you can't measure: Privately measuring demographic performance disparities in federated learning. *Proceedings of the Workshop on Algorithmic Fairness through the Lens of Causality and Privacy 2023 (PMLR)*, 67-85.
- McKenzie, CRM, Liersch, MJ, Finkelstein SR (2006). Recommendations implicit in policy defaults. *Psychological Science*, 17(5), 414-420.
- Nissenbaum H (2004) Privacy as contextual integrity. *Wash. Law Rev.* 79(1):119–57.
- Payne JW, Bettman JR, Johnson EJ (1993) *The adaptive decision maker: Effort and accuracy in choice.* Cambridge University Press: Cambridge.
- Rieke A, Southerland V, Svirsky D, Hsu M (2022). Imperfect inferences: a practical assessment. *FAT* '20: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency.*767-777.

Schuman H, Presser S (1980) Public opinion and public ignorance: The fine line between attitudes and nonattitudes. *Amer. J. Sociology.* ;85(5):1214-25

Schwarz N (1999) Self-reports: How the questions shape the answers. *Amer. Psychologist.* 54(2):93.

Simmons CJ, Bickart B, Lynch JG (1993), "Capturing and creating public opinion in survey research. *J. Consumer Research*, 20(2): 316-329.

Sloman S, Fernbach PM (2017), *The Knowledge Illusion: Why We Never Think Alone*. New York: Riverhead Books.

Thaler RH, Sunstein CR (2021) *Nudge: The Final Edition* (Yale University Press, New Haven).