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 thinnest of understanding about familiar and ordinary objects or policies in their world. Amusingly, most people cannot explain how plumbing 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). This suggests limits to reliance on responses to broad and decontextualized survey questions about privacy to guide specific provisions of privacy policies.

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, some 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 this case, the privacy concern regards 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. 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.

A good example of a policy promoting boundary regulation rather than just concealment of personal data is the Trusted Exchange Framework and Common Agreement (TEFCA). Epic had faced criticism for limited data exchange with non-Epic systems, and TEFCA required interoperability with competitors' systems and hospitals served by those systems, extending the benefits of "Share Everywhere" (Fox 2024).

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. McDonald and Cranor (2008) estimated the national opportunity cost of lost time to read privacy policies at $781 billion annually, roughly 40 minutes per internet user per day. 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? [2]https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3305331

[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.

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.

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 (lune), 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.

Fox A (2024) Epic Nexus connects 625 hospitals to TEFCA. *Health Care IT News*, December 16 2024, available at: https://www.healthcareitnews.com/news/epic-nexus-connects-625-hospitals-tefca

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.

McDonald AM, Cranor LE (2008) The cost of reading privacy policies. *I/S: J. of Law and Policy for Information Society*.4 *4*, p.543-568.

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).