

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: The Economics of Privacy

Volume Authors/Editors: Avi Goldfarb and Catherine E. Tucker, editors

Volume Publisher: University of Chicago Press

Volume ISBNs: 9780226834078 (cloth), 9780226834085 (electronic)

Volume URL:

<https://www.nber.org/books-and-chapters/economics-privacy>

Conference Date: October 14, 2022

Publication Date: August 2024

Chapter Title: The Economics of Privacy: An Agenda

Chapter Author(s): Catherine E. Tucker

Chapter URL:

<https://www.nber.org/books-and-chapters/economics-privacy/economics-privacy-agenda>

Chapter pages in book: p. 5–20

The Economics of Privacy

An Agenda

Catherine E. Tucker

1.1 The Challenge of Privacy for Economics

The economics of privacy is a challenging field in which to be an economist. There are two reasons for this.

The first challenge stems from the definition of privacy. What is privacy? My favorite definition is “freedom from unwarranted intrusion.” This definition stems from Warren and Brandeis (1890), who defined privacy as the “right to be left alone.” They were famously inspired to write their influential essay by the rise of gossip columnists capturing photos with the new technology of portable cameras. This is important, as I will argue in this essay that our conceptions of privacy, and therefore what is important to study as economists, are inextricably tied up with progress in technology. Concepts of privacy are constantly challenged by new technologies that parse personal information in new and unexpected ways. Therefore, unlike a field such as labor economics where the definition of what labor supply is is relatively unchanging, our ideas of what privacy is and should be are in constant flux.

The second challenge stems from our need as economists to at essence conceptualize any question in economics through the lens of a utility function. Farrell (2012) describes the issue very usefully. Typically, in the theoretical literature in economics we tend to think about intermediate preferences for privacy—these reflect the anticipation that if we share our data with a firm, it can be used potentially for things that harm us, like price discrimina-

Catherine Tucker is the Sloan Distinguished Professor of Management Science and Professor of Marketing at MIT Sloan, and a research associate of the National Bureau of Economic Research.

For acknowledgments, sources of research support, and disclosure of the author’s material financial relationships, if any, please see <https://www.nber.org/books-and-chapters/economics-privacy/economics-privacy-agenda>.

tion. By contrast, the vast majority of the literature outside of economics thinks about privacy as a right or something where people should just have a fixed intrinsic taste for keeping certain types of information private. Indeed, often the descriptions of tastes for privacy outside of economics suggest a distaste for creepiness (Richards and Hartzog 2015), or a taste for data being only used in the same context (Nissenbaum 2004). While of course a taste for anything can be included in a utility function, it is unsatisfactory for a discipline that has tried to always model utility functions based on first principles.

1.2 The Focus of the Economics of Privacy Literature So Far

When trying to write an essay on the economics of privacy, it is important to highlight that this has already been done stupendously well by Professor Alessandro Acquisti of Carnegie Mellon University and co-authors, culminating in an essay published in the *Journal of Economic Literature* (Acquisti, Taylor, and Wagman 2016). What is attractive about this conception of the history of privacy is that they define a variety of decades of schools of thought and how this has progressed over time.

The first wave identified by Acquisti et al. (2016) is that of the Chicago School in the 1970s, led by theorists such as Stigler and Posner. In this literature, privacy was defined as a propensity toward secrecy—and in a world where information is generally beneficial to welfare, these models evaluated how tastes for privacy itself could lead to harm to welfare (Posner 1978, 1981; Stigler 1980). Of course, the wave of information economics that characterized theory in the 1980s in economics, itself questioned the idea that more information is always beneficial initiated by some of the idea in theories of signaling and information cascades (Spence 1973; Hirshleifer 1971).

The second wave identified by Acquisti et al. (2016) is also a theoretical literature but this time led by information economists who were interested in questions of technology. Varian (2002) shifted the question of privacy from being one of simply secrecy in what information is shared with other people, to being one firmly about data. This led to new questions such as what should be secondary use rights associated with data. As such it gave rise to what Acquisti et al. (2016) describe as the third wave of theoretical literature, which is interested in questions such as price discrimination (through the use of cookies) (Acquisti and Varian 2005) and targeting effects in online advertising (Johnson 2013; Bergemann and Bonatti 2011, 2015).

The other large shift in the last decade of research has been a proliferation of empirical work in privacy. As described by Goldfarb and Tucker (2012a), much of this work has tried to quantify the effects of privacy regulation on the economy, with much of the literature asking questions about advertising markets (Goldfarb and Tucker 2011; Chiou and Tucker 2012; Johnson, Shriver, and Du 2020; Jia, Jin, and Wagman 2018; Peukert et al. 2020; Johnson, Shriver, and Goldberg 2022; Godinho de Matos and Adjerid

2022), though some of the literature also asks about online behavior (Zhao et al. 2021), financial markets (Kim and Wagman 2015), and health markets (Miller and Tucker 2009, 2011; Adjerid et al. 2016; Miller and Tucker 2017).

These few brief paragraphs do not of course do justice to the literature on the economics of privacy. However, it is fair to say that as yet the number of researchers and number of publications are relatively small given its potential importance in the digital economy. Recently recognizing this, the NBER, and with support from the Sloan Foundation, has instituted a one-off conference on the economics of privacy and also a PhD tutorial to try and inspire more work in this area. This chapter of the handbook that reflects this work tries to offer some suggestions about how economists might be able to deepen and broaden this current literature.

1.3 Outstanding Questions

This handbook is aimed at young researchers who are starting off their careers. Therefore, it makes sense to focus on some of the big questions that researchers in economics have not yet tackled (or have only tackled in part).

1.3.1 The Value of Privacy

1.3.1.1 *Measuring Positive Consequences of Privacy Regulation*

Much of the empirical wave of research on privacy has focused on the question of how privacy regulation hurts economic outcomes—by restricting advertising effectiveness (Goldfarb and Tucker 2011; Johnson, Shriver, and Du 2020), leading to market concentration (Peukert et al. 2020; Johnson, Shriver, and Goldberg 2022), exacerbating inequality (Kim and Wagman 2015), or hurting health outcomes (Miller and Tucker 2009; Adjerid et al. 2016; Miller and Tucker 2017). However, given the large literature on how privacy regulation has large negative economic consequences, the paucity of literature on the benefits of privacy regulation is surprising.

Therefore, it may make sense for researchers to also think about situations or contexts where privacy rights and regulations might have clear positive consequences for individuals. Some I have thought of include:

- Data concerning reproductive health
- Data concerning mental illness
- Data concerning disability that might be used to disqualify potential employees from jobs they could do well
- Data concerning past crimes that are orthogonal to a current question that requires judgment

One thing which all these things have in common is that they concern questions where a stigma exists that is unrelated to potential economic output or the economic quality of a match. In such cases, if privacy regulation

temper data diffusion about something that has an irrational stigma, then it must be the case that privacy regulation benefits individuals. If this is the case, privacy regulation or privacy protections should have positive effects on consumer welfare.

Other occasions where it should be straightforward to document benefits from privacy regulation include instances where data itself might be used for coercion: This might include:

- Targeting those who suffer from addictive behavior to pursue their addiction
- Targeting those who have struggled managing their credit in the past, with further unwise credit offerings
- Targeting those who suffer eating disorders, with weight loss products

These examples share the theme that if prompted, an individual might pursue a course that is not ultimately utility-maximizing for them. As such a restriction of data that means they are not likely to be targeted with prompts may benefit them.

1.3.1.2 Measuring Tastes for Privacy

If we are to truly understand, though, whether privacy regulation has benefits to consumers, we have to return to measurement of key parameters in the utility function. If we assume that a taste for privacy is built into a consumer's utility function, then by definition any regulation that caters to this taste improves consumer welfare. However, this implies we have to actually measure relative intrinsic tastes for privacy. One of the first empirical papers that has attempted to do so is Lin (2022). This paper finds that in general there is a lot of heterogeneity in intrinsic tastes for privacy and perhaps the magnitudes are smaller than might be expected given the privacy literature.

It is clear that the more that can be inferred about underlying tastes for privacy, given observed consumer choices over privacy decisions regarding their data, the more informed this debate can be.

One issue that has thwarted attempts at measuring preferences for privacy is something known as the privacy paradox (Athey et al. 2017). This reflects the observed phenomenon that often while consumers express a desire for privacy when asked about it, it appears they are willing to share their data very readily in a way which seems to contradict this. Of course, in economics tension between stated preferences and revealed preferences are not new, and economists by disposition tend to trust more revealed preferences. But privacy is a domain where trying to unpack this tension appears worthwhile. How is the privacy paradox moderated by the knowledge of consumers? Does the privacy paradox ever reflect consumers engaging in some type of behavioral distortion, which means their stated preferences are closer to the truth? In what domains is the privacy paradox most important, and how does that affect our attempts to evaluate privacy regulation?

In general, what is clear is that from an economics perspective the more we can examine and model actual consumer behavior regarding privacy data using individual-level decisions, the more we will be able to model and parse individual privacy preferences. In particular, the more we could have data on an individual's decision-making regarding the privacy of their data across different domains, the more informative this may be.

1.3.1.3 Privacy Preferences and Context

An appealing theory for understanding some apparent disconnects in privacy preferences is that of the idea of contextual integrity (Nissenbaum 2010). This states that privacy preferences or intrinsic tastes for privacy can be understood as reflecting five contextual parameters that help shape the view of privacy of the individual. These span who the sender of the data is; who the subject of the data is; who the recipient of the data is; the type of data that is sent; and what is referred to as the “transmission” principle, which reflects whether the data was obtained via consent, coercion, or by sale, or by law. Dr Nissenbaum is a philosopher, meaning much of the work is conceptual. As such it reflects a potentially rich testing ground for different theories of tastes for privacy (Bleier, Goldfarb, and Tucker 2020): For example:

1. How much does the same person have different tastes for privacy depending on the recipient? Or the type of data?
2. How does the original context in which data was given affect privacy principles?
3. Do we have similar privacy preferences over our associates' data as we do our own?

Mapping out all these parameters that affect privacy preferences conceptually seems a very useful exercise for empirical analysis.

1.3.1.4 Time-Inconsistency and Privacy Preferences

One important question that appears to have been neglected in the literature is the question of how privacy preferences evolve over time. This is particularly striking because the drop in costs of storing and parsing digital data means that it is virtually costless to store an individual's history of actions over time, rather than periodically deleting it. In other words, there is no reason to think that costs of storage will necessitate the deletion of data. However, there is evidence (Goldfarb and Tucker 2012b) that people's privacy preferences evolve as they grow older; as people get older, they get more privacy conscious. This means that data that young people create today may not reflect their privacy preferences when they get older, which may have negative consequences. In addition, it is of course possible that there are technology shocks, which means that there are unanticipated consequences of sharing data. For example, I might have made decisions

about sharing video footage of myself without predicting that advances in machine learning could lead such data to be decomposed in a manner that allows seamless prediction.

1.3.1.5 Inferential Privacy

Much of the privacy debate has focused on issues of data. And indeed this reflects lawmaking—most laws regarding privacy do not mention the use of algorithms or how that might affect privacy considerations. However, this could be an artifact of laws generally being backwards-looking rather than a prediction of the future.

If I was to speculate, I would argue that in the future we may see a realization that many potential privacy concerns are not a result of the data itself being transferred but instead a result of predictions that are made using this data. For example, though I might be happy to share my photos publicly, and the photos themselves do not cause me any privacy concerns, if algorithms were able to make predictions from these photos about my health, my financial status, my fertility, or other domains of data that I considered private, then I might want to object to using the data in that manner.

If this prediction comes to pass then, this opens up multiple different avenues for research. Indeed, there are already theory papers that are exploring these topics (Acemoglu et al. 2022; Bergemann, Bonatti, and Gan 2020; Goldfarb et al. 2020).

1.3.2 Markets and Privacy

1.3.2.1 Property Rights and Privacy

Whenever economists who have not studied the economics of privacy give interviews about privacy they tend to immediately and instinctively talk about property rights. After all, one of the central tenets of economics is the Coase theorem (Coase 1960), which suggests that many instances of inefficiencies in information markets can be solved by simply clarifying property rights (Farrell 1987). And the idea that all the tensions involved by trying to optimize privacy protections can be best solved by property rights is superficially an attractive one. And indeed my own research has shown that giving controls to people over their privacy—perhaps akin to property rights—helps address privacy concerns (Tucker 2014; Miller and Tucker 2017).

However, there are ultimately obvious flaws in thinking that property rights alone can address privacy concerns, which are themselves worth exploring as potential research topics:

- The idea that data is neatly binary does not fit current data markets. Instead, it makes sense in a world of spreadsheets where each person's data is neatly encased in a single row of data. Take, for example, a photo

I take of myself in a shopping center. This photo might—through facial recognition technologies—also place other individuals at that shopping center. However, even though I might have taken the data and therefore own the data, it is not clear that I have property rights over anyone else's image that might appear in the photo.

- When I take a genetic test, and create data, I am creating data that might affect my ancestors and my descendants. Though I might be able to sell my genetic data to an interested firm, what should be done about the spillovers this has and inferences that are created for my family members?
- Often my data is not particularly valuable, however, inferences from it may be (Miklós-Thal et al. 2023). Let us suppose I liked curly fries on Facebook and researchers were able to infer that this implied I was clever. Do I as the owner of the data also own rights to this inference—or do property rights to that inference belong to the researchers? As an aside, this correlation is based on real-life research (Kosinski, Stillwell, and Graepel 2013).

Therefore, perhaps a way for research in this area to succeed is to study the differences between data where there is a clear property right, and data where there is not. And understand the economic implications of both. This is an area where it seems to be that the talents of theorists would be particularly helpful.

1.3.2.2 *Individual Data Markets*

Though it is possible to think that the fuzziness of property rights when it comes to data could be one potential explanation for why current approaches to privacy have failed, it is also possible to think of more traditional sources of market failure, such as moral hazard and adverse selection. A useful place to study this is in current efforts to build up individual data markets. There are plenty of firms that have sought to set up businesses which would allow individuals to own their data and trade it for monetary value. For example, firms like CitizenMe, Streamlytics, and Clture have tried to establish individual data markets along these lines (see <https://www.citizenme.com/>, <https://www.streamlytics.co/>, and <https://www.clture.io/>). Firms like brave offer to pay people for their attention and data.¹ However, as of yet none of these efforts have thrived.

There is a fledgling literature that tries to understand some of the limitations from a privacy perspective of these markets (Spiekermann et al. 2015). There is also a theoretical literature that explores the consequences of these markets not existing (Jones and Tonetti 2020), being distorted by regulation (Fainmesser, Galeotti, and Momot 2022), or being plagued by externalities

1. See: <https://brave.com/compare/chrome/earning/>.

(Ichihashi 2021). But it seems clear that more papers are needed that try to study the diffusion of these attempts to create data markets and issues of adverse selection and moral hazard that might intuitively plague attempts to create such markets.

Another explanation that may be worth exploring is that the ubiquity of data and non-rivalry of data have also hampered the successful monetization of an individual's data.

1.3.2.3 Competitive Dynamics and Privacy

It is also useful to think about how privacy regulation or tastes for privacy might affect market dynamics and competition as a whole. Early theoretical work such as Campbell, Goldfarb, and Tucker (2015) sketched out theoretical reasons why privacy regulation might lead to concentration. Since then, a variety of work has appeared to confirm this (Miller and Tucker 2014; Peukert et al. 2020; Johnson, Shriver, and Goldberg 2022; Marthews and Tucker 2019a). However, this doesn't mean that the topic is closed to new research. Instead, it means it is time to broaden the number of contexts that such studies are conducted in—for example extending the insights to less studied industries where privacy matters—such as educational technology.

It is also possible to take this type of research and ask questions that illuminate competitive strategy. For example, it would be useful to study where a differentiation on the privacy dimension is a successful strategy, or whether, as appears to have been the case so far, it is ultimately a niche strategy. What types of privacy regulation might be most successful and curtailing the market power of firms, where their market power stems from data? For example, in the fledgling genetic and genomic health industry, can privacy regulations be designed in a way which will not cement market power for an incumbent?

1.3.2.4 The Market for Privacy-Enhancing Technologies

Just as technology has led to an increase in privacy concerns, there has also been an increase in the use of technologies to help individuals and firms institute privacy protections. In general, the work on economics that has considered the spread and importance of these technologies has focused on ad-blocking software (Shiller, Waldfogel, and Ryan 2018; Gritkevich, Katona, and Sarvary 2022). However, this vastly understates the breadth and depth of these technologies—especially the extent to which they are used by firms. The new suite or stack of technologies is often referred to by the label of “privacy-enhancing” technologies. As the Office of Science and Technology Policy recently said:²

2. See: <https://www.federalregister.gov/documents/2022/06/09/2022-12432/request-for-information-on-advancing-privacy-enhancing-technologies>.

Privacy-Enhancing Technologies (PETs) present a key opportunity to harness the power of data and data analysis techniques in a secure, privacy-protecting manner. This can enable more collaboration across entities, sectors, and borders to help tackle shared challenges, such as healthcare, climate change, financial crime, human trafficking, and pandemic response. PETs can also help promote continued innovation in emerging technologies in a manner that supports human rights and shared values of democratic nations, as highlighted during the Summit for Democracy in December 2021, which included an announcement that the US and the United Kingdom are collaborating to develop bilateral innovation prize challenges focused on advancing PETs.

Such statements make it clear that policy makers believe that these technologies may help unravel the traditional trade-off between privacy regulation and economic efficiency documented by economists. Therefore, it makes sense for economists to both explore the extent to which such privacy-enhancing technologies are successful at achieving these aims and also any barriers that exist to their adoption. Indeed, there are a whole set of technologies designed to help firms manage their privacy internally, such as “consent managers,” “privacy assessment managers,” and “de-identification tools.”³ Economists are well placed to assess the extent to which these tools actually enhance privacy of customer, grounding such research on the insights of both organizational economics and enterprise-level diffusion of technology. It is also worth exploring the extent to which such tools reflect the deadweight welfare loss of document privacy regulation compliance relative to actual enhancements of privacy protection for consumers.

1.3.2.5 Decentralized Markets

The twin popular waves of cryptoeconomics and web 3 emphasize the emergence of decentralized markets. As such any discussion of markets and data should at least consider the potential consequences of decentralization of markets on privacy.

In general, I have expressed some skepticism about the extent to which many of the underpinning technologies or principles of blockchain technologies are naturally privacy enhancing (Marthews and Tucker 2019b, 2023). In particular the qualities of verifiability of data and immutability of data that are inherent in the blockchain appear to both restrict users’ ability to control their data or privacy principles such as the right to be forgotten.

However, it is certainly the case that firms and individuals within this community are hopeful that there are potential ways of using these technologies to enhance privacy. For example, firms like Meeco and Solid are both hoping to establish privacy-compliant data markets and data owner-

3. See: <https://www.trustradius.com/data-privacy-management>.

ship structures (see <https://www.meeco.me/platform> and <https://github.com/solid/solid>). These new technologies and these new settings themselves present opportunities for researchers—given the promise that they themselves will generate data which allow us to study privacy-related behaviors and underlying preferences.

1.3.3 The Broader Economy and Privacy

1.3.3.1 *Privacy and Inequality*

There appears to be a positive correlation between privacy concerns, the enactment of privacy regulation, and GDP. We also know that when we use proxy measures for privacy concerns such as sign ups to the do-not-call list—an antitelemarketing innovation—this proxy for privacy preferences correlates with household income and is negatively correlated with demographic group indicators that have been historically disadvantaged in the US (Varian, Wallenberg, and Woroch 2005).

However, despite these striking correlations there is little work that tries to understand why there is this relationship between economic prosperity and privacy concerns. Correspondingly, there is little work which investigates whether there are distributional consequences of privacy protections or privacy regulations. For example, one provocative way of thinking about the advertising-supported Internet is that it is hugely redistributive. Rich people's data is valued by advertisers, and it is these high valuations which allow advertising-supported platforms to supply their services for free to many low-income households both domestically and internationally. Privacy regulation might restrict this redistribution. Whether or not the reader agrees with this rather provocative characterization, it does suggest that the question of whether privacy regulation or protection has distributional consequences is an important one to answer. After all, in economics we are interested in studying phenomena that affect both efficiency and equity.

Some initial research in this area has tried to at least establish some facts about the scale and accuracy of data collection (Neumann et al. 2022). This suggests that low-income households, less-educated households, and renters are far less likely to have demographic information accurately filed and also actually have data available to be collected about them.

1.3.3.2 *Privacy and Discrimination*

In general, as discussed in this agenda, it has been difficult to measure and calibrate privacy harms. This may explain the shift in the policy debate toward questions of algorithmic bias or discrimination. Just by terminology alone, the potential for such phenomena is alarming and also may reflect the untrammeled use of individual data by organizations and corporations. The growing algorithmic fairness literature is beyond the scope of this chapter;

see Cowgill and Tucker (2019) for an overview of the topic from an economics perspective. However, it is useful to think about how and whether privacy regulation reduces, doesn't affect, or augments the potential for algorithmic discrimination. It would seem from a theoretical perspective that any of these outcomes are possible. Privacy regulation might demand that firms reduce the amount of personally identifiable information that is available—this might hinder firms and governments' ability to audit their algorithms and identify instances of bias. Privacy regulation could also restrict the use of data by algorithms which give rise to algorithmic discrimination. Since the direction of the interaction between privacy regulation and algorithmic bias is unclear, this makes it an important area for empirical research.

1.3.3.3 *Political Economy: Surveillance and Privacy*

A clear gap in the focus of the current economics literature is that it has virtually all been focused on the question of commercial surveillance rather than government surveillance. However, the consequences of these two types of surveillance are not equal—unlike firms, governments can put you in jail and confiscate your property. As such understanding how the digital revolution has affected our conclusions about government and privacy is important.

There are two exceptions to this gap which are instructive. The first is a paper that explores the extent to which government surveillance of commercial searches associated with the PRISM scandal had chilling effects on customer behavior (Marthews and Tucker 2017). This area of work is important as it suggests that knowledge and fear of government surveillance actions can shape the commercial landscape, putting this topic squarely in the realm of microeconomics and industrial organization. Recent work by Beraja, Yang, and Yuchtman (2020) examined the relationship between government surveillance and economic success but taking more of a macroeconomics perspective. This is useful as it suggests that understanding government surveillance can help shape our understanding of important macroeconomic questions such as growth and trade.

Indeed, there appear to be many questions to uncover when it comes to privacy and trade. In the US there have been multiple attempts to try and bridge trade barriers with the EU caused by government surveillance, among other things. Indeed, the Safe Harbor Framework ended up having to be replaced by the Privacy Shield Framework due to the inadequacy of the current regime in complying with EU privacy regulations.⁴ This uncertainty over trade and compliance has almost certainly had consequences that are important to study but have not yet been evaluated by economists.

4. See: <https://www.ftc.gov/business-guidance/privacy-security/us-eu-safe-harbor-frame-work>.

1.4 The Future

This chapter has been an attempt to set an agenda on privacy. However, it is written by a researcher who has been working on these topics for two decades. She anticipates both that she has missed things that are important and also made many wrong predictions about what is important. As a result, this chapter concludes by expressing excitement about what the research that young researchers who read this paper will do in the future.

Appendix

The FTC recently shared an Advance Notice of Proposed Rulemaking. Here are some illustrative questions. Answering these questions would indeed move the field forward:

1. Which practices do companies use to surveil consumers?
2. Which measures do companies use to protect consumer data?
3. Which of these measures or practices are prevalent?
4. Are some practices more prevalent in some sectors than in others?
5. How, if at all, do these commercial surveillance practices harm consumers or increase the risk of harm to consumers?
6. Are there some harms that consumers may not easily discern or identify? Which are they?
7. Are there some harms that consumers may not easily quantify or measure? Which are they?
8. How should the Commission identify and evaluate these commercial surveillance harms or potential harms? On which evidence or measures should the Commission rely to substantiate its claims of harm or risk of harm?
9. Which areas or kinds of harm, if any, has the Commission failed to address through its enforcement actions?
10. Has the Commission adequately addressed indirect pecuniary harms, including potential physical harms, psychological harms, reputational injuries, and unwanted intrusions?
11. Which kinds of data should be subject to a potential trade regulation rule? Should it be limited to, for example, personally identifiable data, sensitive data, data about protected categories and their proxies, data that is linkable to a device, or non-aggregated data? Or should a potential rule be agnostic about kinds of data?
12. Which, if any, commercial incentives and business models lead to lax data security measures or harmful commercial surveillance practices? Are some commercial incentives and business models more likely to protect

consumers than others? On which checks, if any, do companies rely to ensure that they do not cause harm to consumers?

13. Lax data security measures and harmful commercial surveillance injure different kinds of consumers (e.g., young people, workers, franchisees, small businesses, women, victims of stalking or domestic violence, racial minorities, the elderly) in different sectors (e.g., health, finance, employment) or in different segments or “stacks” of the Internet economy. For example, harms arising from data security breaches in finance or healthcare may be different from those concerning discriminatory advertising on social media, which may be different from those involving education technology. How, if at all, should potential new trade regulation rules address harms to different consumers across different sectors? Which commercial surveillance practices, if any, are unlawful such that new trade regulation rules should set out clear limitations or prohibitions on them? To what extent, if any, is a comprehensive regulatory approach better than a sectoral one for any given harm?

References

Acemoglu, D., A. Makhdoomi, A. Malekian, and A. Ozdaglar. 2022. “Too Much Data: Prices and Inefficiencies in Data Markets.” *American Economic Journal: Microeconomics* 14 (4): 218–56.

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman. 2016. “The Economics of Privacy.” *Journal of Economic Literature* 54 (2): 442–92.

Acquisti, A., and H. R. Varian. 2005. “Conditioning Prices on Purchase History.” *Marketing Science* 24 (3): 367–381.

Adjerid, I., A. Acquisti, R. Telang, R. Padman, and J. Adler-Milstein. 2016. “The Impact of Privacy Regulation and Technology Incentives: The Case of Health Information Exchanges.” *Management Science* 62 (4): 1042–1063.

Athey, S., C. Catalini, A. Moehring, and C. Tucker. 2023. “The Digital Privacy Paradox: Small Money, Small Costs, Small Talk.” NBER Working Paper 23488. Cambridge, MA: National Bureau of Economic Research.

Beraja, M., D. Y. Yang, and N. Yuchtman. 2020. “Data-Intensive Innovation and the State: Evidence from AI Firms in China.” NBER Working Paper 27723. Cambridge, MA: National Bureau of Economic Research.

Bergemann, D., and A. Bonatti. 2011. “Targeting in Advertising Markets: Implications for Offline versus Online Media.” *RAND Journal of Economics* 42 (3): 417–443.

Bergemann, D., and A. Bonatti. 2015. “Selling Cookies.” *American Economic Journal: Microeconomics* 7 (3): 259–294.

Bergemann, D., A. Bonatti, and T. Gan. 2020. “The Economics of Social Data.” arXiv preprint arXiv:2004.03107.

Bleier, A., A. Goldfarb, and C. Tucker. 2020. “Consumer Privacy and the Future of Data-Based Innovation and Marketing.” *International Journal of Research in Marketing* 37 (3): 466–480.

Buckman, J. R., I. Adjerid, and C. Tucker. 2023. "Privacy Regulation and Barriers to Public Health." *Management Science* 69 (1): 342–350.

Campbell, J., A. Goldfarb, and C. Tucker. 2015. "Privacy Regulation and Market Structure." *Journal of Economics & Management Strategy* 24 (1): 47–73.

Chiou, L., and C. Tucker. 2012. "Data Storage, Data Privacy and Search Engines." Mimeo, MIT.

Coase, R. H. 1960. "The Problem of Social Cost." *The Journal of Law and Economics* 56 (3): 1–40.

Cowgill, B., and C. E. Tucker. 2019. "Economics, Fairness and Algorithmic Bias." Mimeo, MIT.

Fainmesser, I. P., A. Galeotti, and R. Momot. 2022. "Digital Privacy." *Management Science* 69 (6): 3157–3173.

Farrell, J. 1987. "Information and the Coase Theorem." *Journal of Economic Perspectives* 1 (2): 113–129.

Farrell, J. 2012. "Can Privacy Be Just Another Good?" *Journal of Telecommunications and High Technology Law* 10: 251.

Godinho de Matos, M., and I. Adjerid. 2022. "Consumer Consent and Firm Targeting after GDPR: The Case of a Large Telecom Provider." *Management Science* 68 (5): 3330–3378.

Goldfarb, A., A. Haviv, J. Miklos-Thal, and C. Tucker. 2020. "Digital Hermits." Mimeo, Rochester University.

Goldfarb, A., and C. Tucker. 2012a. "Privacy and Innovation." *Innovation Policy and the Economy* 12 (1): 65–90.

Goldfarb, A., and C. Tucker. 2012b. "Shifts in Privacy Concerns." *American Economic Review: Papers and Proceedings* 102 (3): 349–53.

Goldfarb, A., and C. E. Tucker. 2011. "Privacy Regulation and Online Advertising." *Management Science* 57 (1): 57–71.

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

Hirschleifer, Jack. 1971. "The Private and Social Value of Information and the Reward to Inventive Activity." *American Economic Review* 61 (4): 561–74.

Ichihashi, S. 2021. "The Economics of Data Externalities." *Journal of Economic Theory* 196: 105316.

Jia, J., G. Z. Jin, and L. Wagman. 2018. "The Short-Run Effects of GDPR on Technology Venture Investment." NBER Working Paper 25248. Cambridge, MA: National Bureau of Economic Research.

Johnson, G., S. Shriver, and S. Goldberg. 2022. "Privacy & Market Concentration: Intended & Unintended Consequences of the GDPR." Available at SSRN 3477686.

Johnson, G. A., S. K. Shriver, and S. Du. 2020. "Consumer Privacy Choice in Online Advertising: Who Opt Out and At What Cost to Industry?" *Marketing Science* 39 (1): 33–51.

Johnson, J. P. 2013. "Targeted Advertising and Advertising Avoidance." *RAND Journal of Economics* 44 (1): 128–144.

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

Kim, J.-H., and L. Wagman. 2015. "Screening Incentives and Privacy Protection in Financial Markets: A Theoretical and Empirical Analysis." *RAND Journal of Economics* 46 (1): 1–22.

Kosinski, M., D. Stillwell, and T. Graepel. 2013. "Private Traits and Attributes Are Predictable from Digital Records of Human Behavior." *Proceedings of the National Academy of Sciences* 110 (15): 5802–5805.

Lin, T. 2022. "Valuing Intrinsic and Instrumental Preferences for Privacy." *Marketing Science* 41 (4): 663–681.

Marthews, A., and C. Tucker. 2019a. "Privacy Policy and Competition." Brookings Paper.

Marthews, A., and C. Tucker. 2023. "What Blockchain Can and Can't Do: Applications to Marketing and Privacy." *International Journal of Research in Marketing* 40 (1): 49–53.

Marthews, A., and C. Tucker. 2017. "The Impact of Online Surveillance on Behavior." In *The Cambridge Handbook of Surveillance Law*, edited by D. Gray and S. Henderson. Cambridge: Cambridge University Press. doi:10.1017/9781316481127.019.

Marthews, A., and C. E. Tucker. 2019b. "Blockchain and Identity Persistence." In *Cryptoassets: Legal, Regulatory, and Monetary Perspectives*, edited by Chris Brummer. New York: Oxford Academic online edition Oct. 24, 2019. Accessed Oct. 29, 2023. <https://doi.org/10.1093/oso/9780190077310.001.0001>.

Miklós-Thal, J., A. Goldfarb, A. M. Haviv, and C. Tucker. 2023. "Digital Hermits." NBER Working Paper 30920. Cambridge, MA: National Bureau of Economic Research.

Miller, A., and C. Tucker. 2011. "Can Health Care Information Technology Save Babies?" *Journal of Political Economy* 119 (2): 289–324.

Miller, A., and C. Tucker. 2014. "Health Information Exchange, System Size and Information Silos." *Journal of Health Economics* 33 (2): 28–42.

Miller, A., and C. Tucker. 2017. "Privacy Protection, Personalized Medicine, and Genetic Testing." *Management Science* 64 (10): 4648–4668.

Miller, A. R., and C. Tucker. 2009. "Privacy Protection and Technology Diffusion: The Case of Electronic Medical Records." *Management Science* 55 (7): 1077–1093.

Neumann, N., C. E. Tucker, L. Kaplan, A. Mislove, and P. Sapiezynski. 2022. "Data Deserts and Black Boxes: The Impact of Socio-Economic Status on Consumer Profiling." Mimeo, MIT.

Nissenbaum, H. 2004. "Privacy as Contextual Integrity." Symposium. *Washington Law Review* 79: 119. Available at: <https://digitalcommons.law.uw.edu/wlr/vol79/iss1/10>.

Nissenbaum, H. 2010. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford Law Books. Stanford: Stanford University Press.

Peukert, C., S. Bechtold, M. Batikas, and T. Kretschmer. 2020. "European Privacy Law and Global Markets for Data." *Marketing Science* 41(4): 663–869.

Posner, R. A. 1979. "Privacy, Secrecy, and Reputation." *Buffalo Law Review* 1: 28.

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

Richards, N., and W. Hartzog. 2015. "Taking Trust Seriously in Privacy Law." *Stanford Technology Law Review* 19: 431.

Shiller, B., J. Waldfogel, and J. Ryan. 2018. "The Effect of Ad Blocking on Website Traffic and Quality." *RAND Journal of Economics* 49 (1): 43–63.

Spence, M. 1973. "Job Market Signaling." *Quarterly Journal of Economics* 87 (3): 355–374.

Spiekermann, S., A. Acquisti, R. Böhme, and K.-L. Hui. 2015. "The Challenges of Personal Data Markets and Privacy." *Electronic Markets* 25 (2): 161–167.

Stigler, G. J. 1980. "An Introduction to Privacy in Economics and Politics." *Journal of Legal Studies* 9 (4): 623–644.

Tucker, C. 2014. "Social Networks, Personalized Advertising, and Privacy Controls." *Journal of Marketing Research* 51 (5): 546–562.

Varian, H., F. Wallenberg, and G. Woroch. 2005. "The Demographics of the Do-Not-Call List [Security of Data]." *IEEE Security & Privacy* 3 (1): 34–39.

Varian, H. R. 2002. "Economic Aspects of Personal Privacy." In *Cyber Policy and Economics in an Internet Age*, 127–137. Boston, MA: Springer.

Warren, S. D., and L. D. Brandeis. 1890. "The Right to Privacy." *Harvard Law Review* 4 (5): 193–220.

Zhao, Y., P. Yildirim, and P. K. Chintagunta. 2021. "Privacy Regulations and Online Search Friction: Evidence from GDPR." Available at SSRN 3903599.