The Economics of Privacy: An Agenda

Catherine Tucker*

January 26, 2023

Abstract

This paper traces out some of the focus of the literature in the economics of privacy so far, and then boldly attempts to set forth an agenda of topics that seem important to study in the future.

*Catherine Tucker is the Sloan Distinguished Professor of Management Science at MIT Sloan School of Management, Cambridge, MA, and Research Associate at the NBER.

Contents

1 The Challenge of Privacy For Economics

The Economics of Privacy is definitely a challenging field in which to be an economist, for two main reasons.

The first challenge stems from defining 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'. Warren and Brandeis (1890) were famously inspired to write their influential essay by the rise of the use of cameras by gossip columnists to capture photos of notable people. Our conceptions of privacy, and therefore what is important to study as economists, are inextricably tied up with technological progress. Concepts of privacy are therefore constantly challenged by new technologies that parse personal information in new and unexpected ways. Therefore, unlike a field such as health economics, where the definition of 'health' 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 conceptualize any economic question 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 discrimination. By contrast, the vast majority of the literature outside of economics think about privacy as a right, or as something where people just have a fixed intrinsic taste for keeping certain types of information private. Often, descriptions of 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). A taste for anything can be included in a utility function, but this is unsatisfactory for a discipline that has tried to always model utility functions based on first principles.

2 The Focus of the Economics of Privacy Literature So Far

An excellent historical overview is (Acquisti et al., 2016), which defines 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 the Chicago School in the 1970s, led by theorists such as Stigler and Posner. In this literature, privacy was defined as a propensity towards secrecy - and in a world where information is generally beneficial to welfare, these models evaluated how tastes for privacy itself could harm welfare (Posner, 1978, 1981; Stigler, 1980). In the 1980s, the wave of information economics questioned the idea that more information is always beneficial, using theories of signaling and information cascades (Spence, 1978; Hirshleifer, 1978).

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 secondary use rights should be associated with data. This gave rise to a third wave of theoretical literature, interested in questions such as price discrimination (though 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, much of the literature asking questions about advertising markets (Goldfarb and Tucker, 2011; Chiou and Tucker, 2012; Johnson et al., 0; Jia et al., 2018; Peukert et al., 2020; Johnson et al., 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 of yet, the numbers of researchers and publications are relatively small given its potential importance in the digital economy. Recently recognizing this, the NBER, with support from the Sloan Foundation, 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.

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).

3.1   The Value of Privacy

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 et al., 0), leading to market concentration (Peukert et al., 2020; Johnson et al., 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 tempers data diffusion about something that has a 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.

3.1.2 Measuring Tastes for Privacy

If we are to truly understand whether privacy regulation has benefits to consumers, we have to return to measurement of principles. If we assume that a taste for privacy is built into a consumers' utility function, than 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 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 than can be inferred about underlying tastes from privacy, given observed consumer choices over privacy decisions regarding their data, the more informed this debate can be.

One issue which has thwarted attempts at measuring preferences for privacy is something know as the privacy paradox. This reflects the observed phenomenon that often while consumers (Athey et al., 2017) 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, 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 data we could have on individuals' decision-making regarding the privacy of their data across different domains, the more informed we may be.

3.1.3 Privacy Preferences and Context

An appealing theory for understanding some apparent disconnects in privacy preferences is '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, 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 et al., 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?

3.1.4 Time-Inconsistency and Privacy Preferences

One important, neglected question is how privacy preferences evolve over time. 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. However, there is evidence (Goldfarb and Tucker, 2012b) that people's privacy preferences evolve as they grow older, and that 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. People may also learn over time from experiences with how the collection and use of different types of data may lead to positive or negative consequences for them.

In addition, it is of course possible that there are technology shocks which mean 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 which allows seamless prediction or even its use to create a realistic deep fake version of myself.

3.1.5 Inferential Privacy

Much of the privacy debate has focused on issues of data. This reflects law-making - 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 were 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 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, I might want to object to the data's use 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., 2019, 2022; Bergemann et al., 2020; Goldfarb et al., 2020).

3.2 Markets and Privacy

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 clarifying property rights (Farrell, 1987). The idea that all the tensions involved by trying to optimize privacy protections can be best solved via property rights is superficially attractive, and indeed giving controls to people over their privacy akin to property rights can help address privacy concerns (Tucker, 2014; Miller and Tucker, 2017). However, there are 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. Though I might have taken the photo 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 tests, and create data, I am creating data that might affect my ancestors, relatives and descendants. I might be able to sell my genetic data to an interested firm, but what should be done about the spillovers this has and inferences that are made about my family members?

- Often my data is not particularly valuable,but inferences from it may be. Let us

10

supposed 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? This correlation is based on real-life research (Kosinski et al., 2013).

3.2.2 Individual Data Markets

Though it is possible to think of all the ways that property rights being fuzzy when it comes to data as being a potential explanation for why rights of property right 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 also being in play. A useful place to study this is in current efforts to build up individual data markets. There are plenty of firms who have sought to set up businesses which would allow individuals to own their data and trade it for monetary value. For example, firms like `https://www.citizenme.com/`, `https://www.streamlytics.co/` and `https://www.clture.io/` have tried to establish individual data markets along these lines. Firms like Brave offer to pay people for their attention and data.[1]. However, as of yet none of these efforts have gained mass traction.

There is a fledgling literature that tries to understand some of the limitations of these markets from a privacy perspectives (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 et al., 2022) or being plagued by externalities (Ichihashi, 2021). But it seems clear that more papers are needed that try to study the diffusion of these data markets and issues of transaction costs, adverse selection and moral hazard that might intuitively plague attempts to create them. It may be that the ubiquity of data and non-rivalry of data has also hampered the successful monetization of an individual's data.

---

[1]`https://brave.com/compare/chrome/earning/`

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 et al. (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 et al., 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, by 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 an differentiation on the privacy dimension is a successful strategy, or whether as appears to have been the case so far, that it 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?

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 et al., 2018; Gritckevich et al., 2022). However, this vastly understates the breath and depth of these technologies, and especially the extent to which they are used by firms. The new suite or stack of technologies are often referred to by the label of 'privacy-enhancing' technologies. As the Office of Science

and Technology Policy recently said:[2]

> 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 health care, 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
> United States 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 tradeoff 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 is 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 customers' privacy, 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 compliance relative to actual enhancements of privacy protection for consumers.

---

[2]https://www.federalregister.gov/documents/2022/06/09/2022-12432/
request-for-information-on-advancing-privacy-enhancing-technologies
[3]https://www.trustradius.com/data-privacy-management

3.2.5   Decentralized Markets

The twin popular waves of blockchain 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, 2022). In particular, the qualities of verifiability of data and immutability of data that are inherent in the blockchain, appear to restrict users' ability to control their data and hinder the implementation of 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 `https://www.meeco.me/platform` and `https://github.com/solid/solid` are both hoping to establish privacy-compliant data markets and data ownership structures. These new technologies and 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.

3.3   The Broader Economy and Privacy

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 signups to the U. S. 'do not call list' - an anti-telemarketing innovation - that this proxy for privacy preferences correlates with household income and is negatively correlated with demographic group indicators that have been historically disadvantaged in the USA (Varian et al., 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. Privacy regulation might restrict this redistribution. Whether or not the reader agrees with this, 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 equity as well as efficiency.

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.

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 towards 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 privacy literature is beyond the scope of this article - 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 then hinder firms and government's 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.

3.3.3 Political Economy: Surveillance and Privacy

A clear gap in the focus of the current economics literature it 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 explore 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 et al. (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 U. S. there have been multiple attempts to try and bridge trade barriers with the E. U. caused by, among other things, U. S. government surveillance. The European courts invalidated the Safe Harbor framework, and there has been an iterative process since then of new proposals that have each been invalidated in their turn, as the U. S. government tries to

grapple with providing adequate privacy protections to E. U. nationals.[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   The Future

This article 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 article concludes by expressing excitement about the research that young researchers who read this paper will do in the future.

---

[4]https://www.ftc.gov/business-guidance/privacy-security/us-eu-safe-harbor-framework

References

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

Acemoglu, D., A. Makhdoumi, A. Malekian, and A. Ozdaglar (2022). Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics 14*(4), 218–56.

Acquisti, A., C. R. Taylor, and L. Wagman (2016). The economics of privacy. *Forthcoming, Journal of Economic Literature*.

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, and C. Tucker (2017, June). The digital privacy paradox: Small money, small costs, small talk. Working Paper 23488, 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. Technical report, 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, aug). 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.

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. *preparation for: Journal of Economic Perspectives*.

Fainmesser, I. P., A. Galeotti, and R. Momot (2022). Digital privacy. *Management Science*.

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. *J. on Telecomm. & High Tech. L. 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, January). Privacy regulation and online advertising. *Management Science 57*(1), 57–71.

Gritckevich, A., Z. Katona, and M. Sarvary (2022). Ad blocking. *Management Science 68*(6), 4703–4724.

Hirshleifer, J. (1978). The private and social value of information and the reward to inventive activity. In *Uncertainty in economics*, pp. 541–556. Elsevier.

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. Technical report, 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 (0). Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Science 0*(0), null.

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. *The 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*.

Marthews, A. and C. Tucker (2019a). Privacy policy and competition. *Brookings Paper*.

Marthews, A. and C. Tucker (2022). What blockchain can and can't do: Applications to marketing and privacy. *International Journal of Research in Marketing*.

Marthews, A. and C. E. Tucker (2017). The impact of online surveillance on behavior. *Cambridge Handbook of Surveillance Law*.

Marthews, A. and C. E. Tucker (2019b). Blockchain and identity persistence. *Cryptoassets: Legal and Monetary Perspectives, Forthcoming*.

Miller, A. and C. Tucker (2011). Can healthcare information technology save babies? *Journal of Political Economy* (2), 289–324.

Miller, A. and C. Tucker (2014, January). 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*.

Miller, A. R. and C. Tucker (2009, July). 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. *Wash. L. Rev. 79*, 119.

Nissenbaum, H. (2010). *Privacy in context: technology, policy, and the integrity of social life*. Stanford Law Books.

Peukert, C., S. Bechtold, M. Batikas, and T. Kretschmer (2020). European privacy law and global markets for data. *Center for Law & Economics Working Paper Series 1*.

Posner, R. A. (1978). Privacy, Secrecy, and Reputation.

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. *Stan. Tech. L. Rev. 19*, 431.

Shiller, B., J. Waldfogel, and J. Ryan (2018). The effect of ad blocking on website traffic and quality. *The RAND Journal of Economics 49*(1), 43–63.

Spence, M. (1978). Job market signaling. In *Uncertainty in economics*, pp. 281–306. Elsevier.

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), pp. 623–644.

Tucker, C. (2014, October). 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*, pp. 127–137. Springer.

Warren, S. D. and L. D. Brandeis (1890, December). The right to privacy. *Harward 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*.

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