DECEMBER 2019

Privacy policy and competition

Alex Marthews
National Chair, Restore The Fourth

Catherine Tucker
Sloan Distinguished Professor of Management, MIT

This report is available online at: https://www.brookings.edu

B | Economic Studies at BROOKINGS

The Brookings Economic Studies program analyzes current and emerging economic issues facing the United States and the world, focusing on ideas to achieve broad-based economic growth, a strong labor market, sound fiscal and monetary policy, and economic opportunity and social mobility. The research aims to increase understanding of how the economy works and what can be done to make it work better.

Contents

ABOUT THE AUTHOR

Alex Marthews is the National Chair for Restore The Fourth.
Catherine Tucker is the Sloan Distinguished Professor of Management and a Professor of Marketing at the MIT Sloan School of Management.

STATEMENT OF INDEPENDENCE

The authors did not receive financial support from any firm or person for this article or from any firm or person with a financial or political interest in this article. In her broader work, Tucker has received grants from multiple companies, institutes and associations including DARPA, Google, the MIT CryptoEconomics Lab, the National Science Foundation, National Institutes for Health, the Sloan Foundation, Tilburg Law and Economics Center, Time Warner, the Net Institute, WPP, the Marketing Science Institute, and the Computer & Communications Industry Association. Unrelated to this paper, she has also served as a consultant for ADT, CBS, Bausch and Lomb, Facebook, Microsoft, Lyft, Ripple, RTIC, Samsung, Verizon, ContextLogic, and Yahoo. Marthews is the elected National Chair of Restore The Fourth, a 501(c)(4) nonprofit that advocates against mass government surveillance; founder and coordinator of the Campaign for Digital Fourth Amendment Rights, a 501(c)(4) organization; and the founder and coordinator of Digital Fourth Amendment Research & Education, Inc., a 501(c)(3) nonprofit that seeks to foster research into the application of Fourth Amendment law to digital technologies.

Introduction

Commercial privacy protection and competition law have long been jointly regulated by a single authority - the FTC - in the US. Though managed separately, both types of law have at their heart a desire to protect consumers. Antitrust law tries to ensure that consumers' ability to choose between products is not restricted by anticompetitive acts. In the US, commercial privacy is protected through consumer protection law, which tries to ensure that consumers' ability to choose between products is not restricted by misleading information.

In general, these consumer protection and antitrust missions have developed in parallel with little overlap in the actions of regulatory agencies. However, this article will argue that data privacy will be the first domain where this overlap is unavoidable and will cause tension going forward as authorities have to face hard tradeoffs between ensuring effective competition and ensuring privacy.

Usually when privacy concerns are discussed with reference to competition policy, the general argument is whether the effect of, for example, market share on consumer privacy is something that should be considered as part of competition policy. This has been a main thrust of the argument used by German regulators in their actions towards Facebook. German regulators challenged Facebook by claiming that because of Facebook's large market share, consumers have no choice but to relinquish their privacy to be part of the social network. Therefore, their privacy effectively became a price they had to pay in the absence of a monetary price. As a result, the German competition authority has reduced the extent to which Facebook is able to share data across its various web properties.[1]

Some commentators in the US have advanced similar arguments, by saying that privacy is a sufficiently different part of product quality that the effect of the transaction on privacy should be considered independently in potential antitrust and merger review. This was an argument made by the non-profit Electronic Privacy Information Center (EPIC), for example, in its comments to the FTC on the acquisition of WhatsApp by Facebook.[2] However, then-FTC commissioner Maureen Ohlhausen rejected this viewpoint, arguing that the traditional focus of antitrust was sufficient.

By contrast, this paper will consider how privacy regulation itself, designed to help consumers, can raise competition concerns, and will also examine how government efforts to shape the extent of privacy consumers have from itself may also shape the competitive environment. Ultimately, we argue that there can be tradeoffs between promoting competition and protecting consumer privacy.

. . .

1. https://www.nytimes.com/2019/02/07/technology/germany-facebook-data.html

2. https://epic.org/privacy/ftc/whatsapp/WhatsApp-Complaint.pdf

Is there a positive role for privacy protection and competition?

Before turning to potential tensions, it makes sense to start off with some more hopeful analysis of whether privacy protections and competition policy can complement one another.

In theory, given the treatment and discussion by antitrust authorities, especially in Europe, consumer-facing pricing policy can be seen as another input which determines the product or service's quality. If consumers value privacy, then in theory increased competition should increase the provision of desired privacy protection in the market, much in the same manner that we would expect increased competition in the microprocessor market to increase the provision of speed in processors as well as lowering prices.

However, as of yet there are few examples where creating markets for privacy appear to have been effective.

DuckDuckGo is an alternative search engine which does not collect data.[3] As such, it represents a privacy-preserving alternative to users of search engines such as Google or Bing whose privacy policies are less strict and collect more user data. For example, the Google privacy policy (which includes other data it collects in addition to search) is 7,156 words long.[4] It has also seen 31 major changes in Google's history.[5]

However, an examination of the usage of DuckDuckGo, as shown in Figure 1, does not suggest that privacy competition between it and major search engines are driving usage.

For example, in early 2012 Google shifted its privacy policy so that it would combine all information about an individual over YouTube, Maps, Calendar Gmail and Search. For users worried about the risks of an expansion of scope of data collection, and data combination this shift in policy should have provided a clear impetus for them to shift their usage towards DuckDuckGo. However, the time trend suggests no change in rate of adoption of DuckDuckGo at the time of the shift in Google privacy policy.

Another example that might propel consumers towards using a more privacy-protective search engine would have been the revelations of government surveillance of search engine data in mid-2013. These revelations were the result of the whistleblower Edward Snowden revealing details of a PRISM surveillance program that used data from both the Google and Bing Search engines, but not DuckDuckGo, for National Security Agency surveillance.[6]

. . .

3. DuckDuckGo does not collect or share personal information. That is our privacy policy in a nutshell.' https://duckduckgo.com/privacy'

4. https://policies.google.com/privacy?hl=en

5. https://policies.google.com/privacy/archive?hl=en

6. https://www.theguardian.com/world/interactive/2013/nov/01/snowden-nsa-files-surveillance-revelations-decoded#section/2

However, again, Figure 1 does not show a large uptick in DuckDuckGo usage as a result of these revelations.

Figure 1: DuckDuckGo traffic

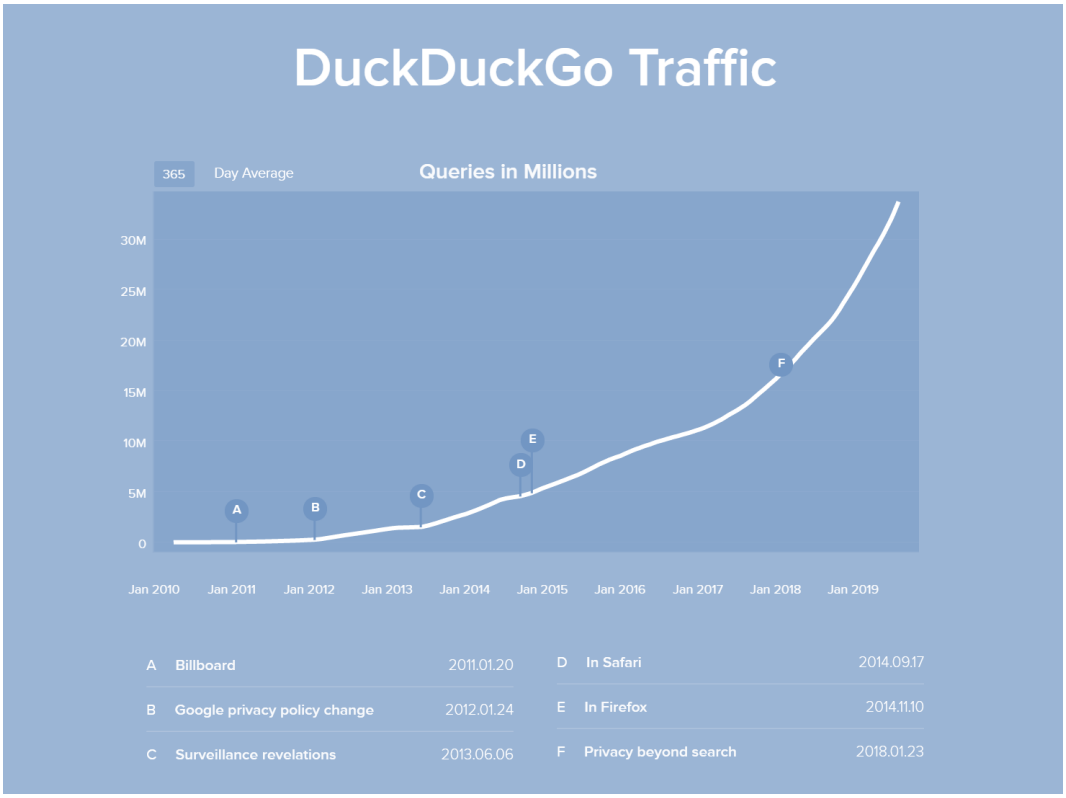

Another example of a social media network that differentiated itself in privacy terms (or at least increased perceived control over data sharing) was Google Plus. When it launched in 2011, in order to get a large initial installed base, all existing Gmail users were automatically signed up to Google Plus. This created pushback from users,[7], and Google's response was to highlight its enhanced privacy controls. It named these enhanced privacy controls as '+Circles' with the tagline 'share what matters, with the people who matter most.'[8] In spite of these efforts, Google Plus sputtered; its emphasis on increased control over data sharing didn't seem to help. Some commentators suggested that the idea of categorizing friends and acquaintances by what information an individual wanted to reveal to them seemed

. . .

7. http://techland.time.com/2012/01/20/want-a-google-account-now-youre-automatically-signed-up-for-google/

8. https://googleblog.blogspot.com/2011/06/introducing-google-project-real-life.html

'strange' and part of the reason that Google Plus never gained traction.[9] The service shut down in 2018.[10]

A natural question of course is why there are few examples showing why firms that compete using better privacy policies or more privacy controls have not succeeded in thriving. One potential explanation is a phenomenon known as the 'privacy paradox.' The privacy paradox reflects the insight that though consumers often express concerns about their privacy, they rarely act in ways which are consistent with their stated preferences. An illustration of the privacy paradox is, who show that though MIT students in general acted in a way which accorded with their stated privacy preferences in terms of sharing information, when these students were offered pizza they started sharing information even if previously they had stated a greater preference for keeping their information private.

There are two ways of interpreting this result. One is that, despite consumers saying they care a lot for privacy, in reality they don't, and their actions reflect that. Another way of interpreting this result is that, though consumers do cherish their privacy, in real-life situations involving material incentives, they often act in a way which is not in accordance with their actual desires. The former interpretation suggests that greater privacy protection is not needed, and the latter interpretation suggests that greater privacy protection is needed. However, the implication of both interpretations for whether competition will enhance privacy is similar. If consumers do not make choices which accord with their stated privacy preferences and instead choose small convenience benefits or monetary benefits over privacy, then a firm that offers superior privacy protections is unlikely to attract many consumers by virtue of its superior privacy protections.

An additional element here is that tech offerings that provide high levels of privacy control, like Mastodon, a decentralized social network, or Signal, an encrypted messaging service, are often also open-source and non-commercial. Even commercial offerings that consciously refrain from certain types of exploitation of consumer data, like DuckDuckGo, are signaling that they are not simply focused on maximizing shareholder value. Commercial, profit-oriented offerings which are, or which aim at being publicly traded, intentionally monetize more of people's data, and can thereby raise more money for advertising, expanding and adding features to their networks using resources from venture capital firms and the stock market. It seems likely that some of the lack of adoption of highly privacy-conscious software applications stems as well from this inbuilt advantage for commercial and proprietary offerings.

There seem to be few efforts by firms to compete via superior privacy protections even in sectors where one might expect privacy concerns to foster a large role in consumer choice. For example, the cloud computing sector is one which is very likely to be most affected (out of many technology sectors) by the recent EU General Data Privacy Directive (GDPR), given the restrictions on the sharing of data it implies. In a 2017 survey of 500 IT decision-

. . .

9. https://blumint.co/stop-thinking-google-plus-dead

10. One reason given for the shutdown was a bug that potentially allowed outside developers to access user data. https://www.theverge.com/2019/4/2/18290637/google-plus-shutdown-consumer-personal-account-delete

makers,[11] 92% expressed themselves worried about issues of GDPR compliance and the cloud, but when it came to choosing cloud-computing software their paramount concern was scalability (41%), followed by reputation for innovation (34%), and for only 26% was compliance the major concern. This may have reflected an assessment that there would not in fact be serious consequences to firm revenues for non-compliance, or an assessment that their customers would not care much about GDPR compliance; perhaps, for many of the firms involved, a relatively small proportion of their customers were EU citizens or were located in the EU.

Theoretically, privacy protection could enhance competition among providers on the privacy dimension. This is because, if firms were forced to disclose details of their privacy policies, consumers would have an effective means of comparing privacy protections across firms. The privacy-oriented nonprofit the Electronic Frontier Foundation (EFF) does part of this work for consumers by publishing an annual consumer-facing "Who Has Your Back" report, comparing major tech firms' privacy policies. However, there is little empirical evidence that this approach shapes usage at scale. Indeed, existing empirical work such as suggests instances where consumers receive some information about privacy and security can actually lead consumers to make less privacy-secure choices going forward.

Evidence on effects of privacy regulation on U.S. domestic competition

There is little evidence that competition itself appears to enhance privacy, so we now turn to a different question, which is whether privacy protections themselves can theoretically hurt competition.

Privacy and domestic competition in theory

The fixed cost argument

From an economics perspective, when modeling the effects of privacy regulation on the ability of firms to compete, one starting point is the observation that in theory, any regulation that imposes any fixed costs on firms will have an anti-competitive effect. This will be the same whether it is environmental protection, protection for workers' rights, or privacy protections. The concern is that if compliance has a fixed cost, then that fixed cost will be more heavily felt by a smaller firm with smaller revenues, putting smaller firms at a cost disadvantage relative to larger firms, or at least only weakly increasing in firm size.

Many privacy regulations (as do other regulations) implicitly acknowledge the potential for regulation to impose relatively higher costs on smaller firms than larger firms, by generally reserving the strictest provisions of privacy protection for larger firms. For example, HIPAA (the Health Information Privacy and Accountability Act) was not judged to apply

. . .

11. "The Clock is Ticking The truth about GDPR compliance", Calgigo, 2019

to health plans that had less than 50 participants that were administered solely by the employer.[12]

The GDPR is somewhat an exception to that practice, in that it allows very few exemptions for smaller businesses. There are a few small record-keeping exemptions for small businesses with fewer than 250 employees, but in general GDPR's policies apply to all firms, and this appears intentional. For example, the first draft of GDPR limited the requirement of appointing a Data Protection Officer to companies that had more than 250 employees or processed more than 5000 personal records. However, the final record included no such exemption.[13] This means that, unlike many prior privacy regulations, GDPR did not adjust the potential costs of compliance by firm size.

Early evidence suggests that actual total GDPR compliance costs appear to be similar no matter the size of the firm. For example, a 2017 survey suggested that on average 10 people were working on GDPR compliance for the average firm which had more than 100 employees in the UK.[14] The average cost of compliance was £1.67 million. For firms between 100 and 249 employees, the average investment in GDPR compliance was £947,000, and the average compliance investment of firms with more than 1,000 employees was £2.3 million. There was more variation in costs from sector to sector: There were three sectors that spent less than £500 thousand on average (Healthcare, Education and Arts and Culture) and three sectors that spent more than £2 million on average (Finance, Travel, and Legal.) In general, this evidence suggests that the cost of GDPR per employee will be proportionately greater for smaller firms.

The role of consent requirements beyond fixed compliance costs

When it comes to privacy protection, argues that regulators should take into consideration more than just the typical fixed cost argument. This is because many privacy regulations are based on gaining consent from a user. The paper argues that users are likely to base the likelihood of giving consent on the longevity of their relationship with the firm and the scope of the benefits they expect to receive from the firm. As such, consumers are more likely to give to consent for their data to be processed by firms that they have longstanding relationships with and where they can clearly understand the benefits of the service the firm provides.

A potential thought experiment to illustrate this point is to imagine the search engine market back in the late 1990s. Would a user be more likely to give their consent for their data to be processed by Altavista - a well-known search engine whose efficiency was a known quantity - or to Google, an upstart search engine whose efficiency was not a known quantity and whose website, as late as 1998, was marked as being in 'Beta'?[15]'

. . .

12. https://www.hhs.gov/sites/default/files/privacysummary.pdf
13. https://www.clarip.com/blog/gdpr-under-250-employees/
14. "The Clock is Ticking: The truth about GDPR compliance", Calgigo, 2019
15. https://www.salon.com/control/1998/12/21/straight_44/

This paper suggests that consent requirements, which are at the heart of both the EU Privacy Directive of 1995[16] and the original OECD privacy guidelines themselves,[17] can add beyond straightforward compliance costs a reason why privacy regulations may end up entrenching large firms and making entry more difficult for smaller firms.

Though not considered in the paper, one extension of this insight is that consent requirements themselves can exacerbate the effects of privacy protection in disadvantaging entrants, through their effect on the attractiveness of bundled services. For example, take a publisher of a small website that wants to monetize its content through advertising. To successfully execute on digital advertising, typically a firm needs a broad array of services - the publisher would need to use a SSP (supply-side platform), a DMP (data-management platform) and then link to a DSP (demand-side platform) which would facilitate the real-time purchase of ad space on its websites by advertisers. Now, in theory it is possible to purchase all of these services separately. Indeed, as shown by Figure 2, not only is it possible, but this description itself is very much a simplification of the variety of firms that might be involved with that single transaction.

Figure 2: LUMAscape

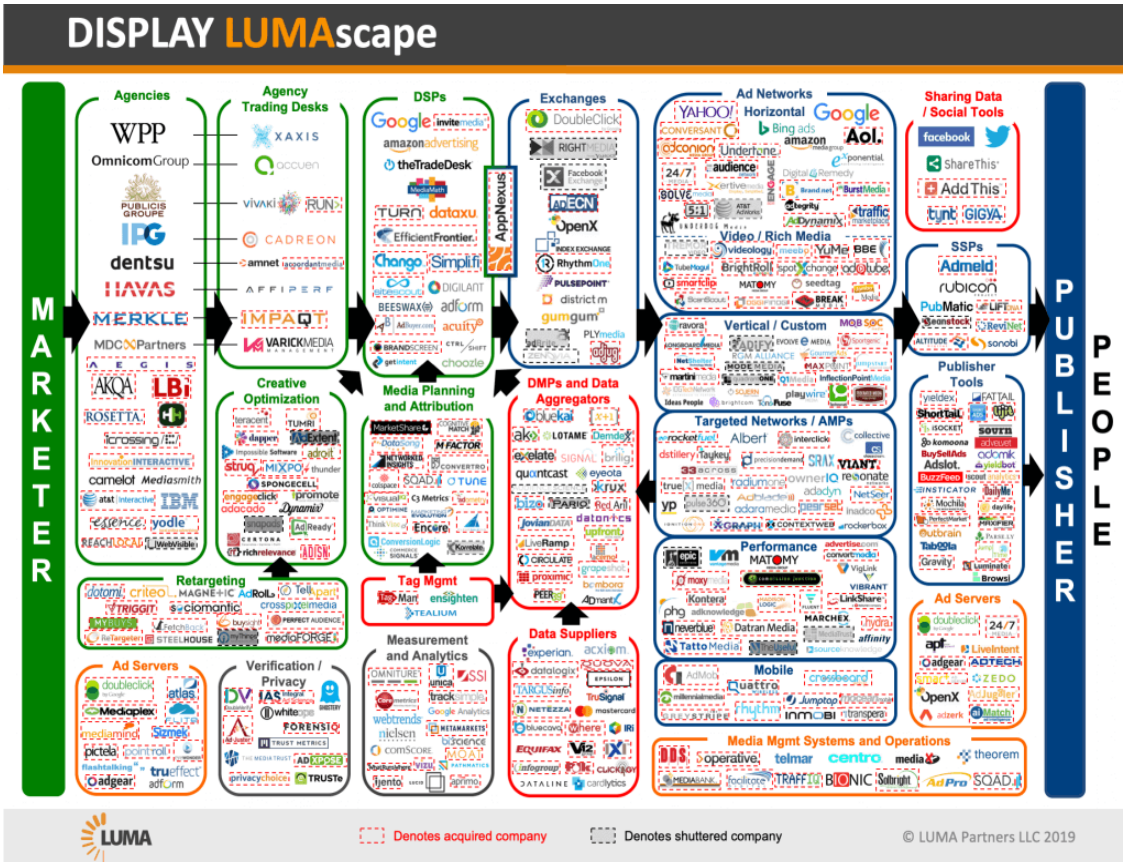

Source: https://lumapartners.com/content/lumascapes/display-ad-tech-lumascape/

. . .

16. https://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=CELEX:31995L0046:en:HTML

17. https://www.oecd.org/internet/ieconomy/oecdguidelinesontheprotectionofprivacyandtransborderflowsofpersonaldata.htm

Now, imagine the emergence of GDPR, which suggests that a firm needs to achieve privacy compliance with each partner who processes site visitor data that is used in targeted advertising. Compliance would be particularly an issue if the ad used online identifiers or IP addresses in its processing, as is typical in online advertising which requires detailed tracking of website visitors to ensure that they are likely to respond to an ad. Ensuring compliance on the part of a partner might involve making sure they assigned a data processing officer and that they are able to process a request from a data subject. Ensuring these compliance details is neither easy nor straightforward.

This small publisher has a choice between ensuring the compliance of multiple parties to a display ad, or instead using, for example, a vertically integrated service such as Google's AdSense, which would allow them to only have to verify the compliance of one ad partner.[18] From a time perspective, and a risk perspective, buying a bundle of ad services from one firm would be more efficient and cost-effective than taking on the burden of ensuring privacy compliance by multiple ad partners. The relative attractiveness of buying bundled services from a large vertically integrated platform was emphasized by a recent report issued by the UK Information Commissioner's Office. One issue that particularly concerned them was the existence of a 'Data Supply Chain' in digital advertising, noting:

> *In many cases there is a reliance on contractual agreements to protect how bid request data is shared, secured and deleted. This does not seem appropriate given the type of personal data sharing and the number of intermediaries involved.*

If they are right that the data-supply chain in the real-time bidding and placement of digital ads is not GDPR-compliant, then this has large implications for the competitive structure of the digital advertising industry. In particular, it suggests that the its current fragmentation and the existence of many small firms at each stage of the supply chain is not tenable. This implies that GDPR favors a lack of fragmentation and fewer firms within the data supply chain, and that GDPR will ultimately lead to pressures for vertical integration in the digital advertising industry.

Evidence on effects of different types of domestic privacy regulations on competition

Though due to the size and importance of GDPR, it is natural that this report should discuss its implications in depth, it is also useful to evaluate empirical studies which have attempted to measure the effects of US domestic privacy policies on competition. One feature of the US privacy regulation system is that it is sectoral in focus, so much of the empirical

. . .

18. Google has also offered complementary products that allow publishers to obtain GDPR-compliant user consent: https://searchengineland.com/googles-amp-project-announces-new-consent-component-ahead-of-gdpr-compliance-deadline-295633

work has been focused on the sectors such as health and finance which have seen the most stringent privacy regulation imposed.

Arguably, one of the US's first forays into explicit privacy regulation was the 'Fair Credit Reporting Act.' In its introspective into its experience with the FCRA, the FTC strikingly did not examine the competitive effects of the Act through its 40-year experience in its enforcement. However, it is notable that the four large credit reporting agencies (Equifax, Experian, Innovis and TransUnion) have seen little competitive entry over the past decades, despite huge shifts in the nature of data and data collection due to the digital revolution, and the incumbent firms experiencing significant scandals relating to consumer privacy. We speculate that a potential explanation is that it is hard for a smaller entrant in this space to achieve even the levels of consent and compliance practiced by the incumbent credit reporting agencies.

To illustrate this, consider the histories of the three (four) credit bureaus:[19]

1. Equifax was founded in 1899 as the Retail Credit Company.
2. Transunion was founded in 1968 as a holding company for railroad leasing, but in 1969 acquired the Credit Bureau of Cook County, leading it to be part of the credit reporting industry.
3. Experian was founded in 1968 as a sub-unit of a firm named TRW that focused on credit data. Of the three credit bureaus, it has shown the most evolution to its current position as a result of its acquisition by a UK company that initially focused on catalogs.
4. Innovis was founded in 1970 under the name the Associated Credit Bureaus (ACB). It is typically smaller than the other three and collects data from fewer sources.

All trace their historical roots to before the passing of the Fair Credit Reporting Act in 1970; there has been essentially no entry into the industry since then. This is despite the fact the digital revolution has theoretically vastly decreased the potential costs to any entrant of collecting, collating and parsing large swathes of consumer credit data and the size of the market has expanded dramatically. One interpretation of this is that not only are compliance costs with FCRA high, but also that in this industry firms that rely on credit reports are far more likely to share data with and rely on the compliance of incumbent agencies, rather than using an entrant agency.

In, the authors study the effects of privacy regulation in the financial sector. Though it is not the major focus of their study, which looks at variations in the adoption of local financial privacy ordinances in five California Bay Area counties and its effects on mortgage denial outcomes, there are some implications for the nature of competition. They show that in general, the effect of privacy laws on the issuance of loans was greater for loans that were not backed by the government agencies of Fannie Mae and Freddie Mac. One potential interpretation is that the larger effect of privacy regulation on the issuance of non-conforming loans is that these could be interpreted as the riskier product, where information could be more useful. As such, this work highlights that privacy regulation is likely to suppress

. . .

19. https://www.creditrepair.com/blog/credit-score/credit-bureau-history/

the supply of riskier and more marginal products. Though this is not straightforwardly applicable to the typical entrant-vs-incumbent dichotomy that we focus on when thinking about competition, it is suggestive about privacy regulations' potential at the margins to restrict the supply of products created for higher-risk individuals.

Outside of the financial sector, there are signals of this pattern occurring in other sectors governed by tighter privacy regulation such as health. studies the effect of privacy protections and concerns on the sharing of health information. Their major finding is that large hospital systems are far more likely to share data internally than externally. In contrast, smaller hospitals are more likely to share patient data externally. Combined with earlier work on privacy regulation and its effects on suppressing data sharing in general such as, this suggests that privacy protection in healthcare may inadvertently encourage larger hospital systems to create data silos, which exclude smaller providers.

Though there is no empirical study yet, another sector where privacy compliance may also be making it harder for entrants to compete is the educational technology sector. As shown by Figure 3, there are multiple complex laws with which educational technology firms need to comply.

Figure 3: Question compliance references

QUESTION COMPLIANCE REFERENCES

Statute	Name	Frequency
Children's Online Privacy Protection Act (COPPA)[11]	COPPA	79
Student Online Personal Information Protection Act (SOPIPA)[12]	SOPIPA	39
Family Educational Rights and Privacy Act (FERPA)[13]	FERPA	38
California Online Privacy Protection Act (CalOPPA)[14]	CalOPPA	29
California AB 1584 - Privacy of Pupil Records (AB 1584)[15]	AB 1584	17
California Privacy Rights for Minors in the Digital World (CalPRMDW)[16]	CalPRMDW	12
California Privacy of Pupil Records (CalPPR)[17]	CalPPR	8
California Data Breach Notification Requirements (DataBreach)[18]	DataBreach	6
General Data Protection Regulation (GDPR)[19]	GDPR	4
Children's Internet Protection Act (CIPA)[20]	CIPA	4
Early Learning Personal Information Protection Act (ELPIPA)[21]	ELPIPA	4
Protection of Pupil Rights Act (PPRA)[22]	PPRA	3
The Communications Decency Act of 1996 (CDA)[23]	CDA	3
California "Shine the Light" (ShineTheLight)[24]	ShineTheLight	2
California Electronic Communications Privacy Act (CalECPA)[25]	CalECPA	2
Digital Millennium Copyright Act (DMCA)[26]	DMCA	2
Copyright Act of 1976 (Copyright)[27]	Copyright	2
The National School Lunch Act (NSLA)[28]	NSLA	1
California Revised Uniform Fiduciary Access to Digital Assets Act (RUFADAA)[29]	RUFADAA	1
California Electronic Commerce Act (CalECA)[30]	CalECA	1
Controlling the Assault of Non-Solicited Pornography and Marketing Act of 2003[31]	CAN-SPAM	1

Figure 3 illustrates the frequency of statutory and regulatory laws that reference questions.

Source:https://www.commonsense.org/education/sites/default/files/tlr-blog/cs-state-of-edtech-privacy-report.pdf, p30

Figure 4 show the results of the degree of compliance with these laws as assessed by a non-profit established to ensure that digital tools protect children's privacy. It is notable that the top-scoring firms from a privacy-compliance perspective are Google and Apple. These are also by far the largest firms. Therefore, these results show the potential for schools who wish to ensure privacy compliance of their educational software to be be driven to buy from large incumbents rather than smaller entrants. That is not to say that children's privacy is not paramount, but that this case illustrates the stark tradeoffs between protecting the privacy of a vulnerable group of individuals and fostering competition.[20]

. . .

20. There are of course other large fixed costs in providing educational software which may also lead to a tendency towards fewer smaller firms entering the space.

Figure 4: Educational products provided by larger firms are rated the most compliant

Product	Tier	Score	Effective Date	Data Sold	Third-Party Marketing	Behavioral Ads	Third-Party Tracking	Track Users	Ad Profile
Apple School Manager	✓	72							
Google Classroom	✓	71							
Kiddom	✓	59							
Blackboard Learn	!	58							
MasteryConnect	!	58							
Edmodo	!	52							
D2L Brightspace	!	50							
Engrade	!	50							
Schoology	!	44							
Neo	!	40							

Figure 4 illustrates better, worse, and unclear practices for the selected transparency questions. The color blue means the policies disclose better practices, red means the policies disclose worse practices, and orange means the policies are unclear as to whether or not the vendor engages in the respective practice.

Source:https://www.commonsense.org/education/articles/privacy-evaluation-of-top-10-district-wide-edtech-products

Evidence on effects of domestic privacy regulation on international competition

In a digital economy, firms of course do not have to be locally based to sell services and products to consumers. This means that there are potential effects on the relative competitiveness of domestic firms with international firms if there are privacy regulations or safeguards which are specific to the domestic economy.

In theory, given the nature of most privacy regulation this should not be a challenge. This is because most privacy regulation regulates the privacy rights of the citizens of its own country, rather than the privacy obligations of firms within its borders.[21] However, if the degree of international enforcement is uncertain, privacy regulation may have a greater

. . .

21. An exception to this is upcoming privacy regulation in California which will affect states outside of California - see https://www.americanbar.org/groups/business_law/publications/committee_newsletters/bcl/2019/201902/fa_9/

effect on the privacy compliance of firms within a country's borders than on the privacy protections given to consumers by firms outside of the country's borders.

This is something that is explored in a new paper, which looks at the practical implications of The Children's Online Privacy Protection Act of 1998 in the development of apps targeted at very young children (such as toddlers and preschoolers). The paper documents that the intrusiveness of data collection by these apps is very much driven by the location of the developer's country rather than local law. Many apps targeted at very young children are developed in countries such as the Ukraine, which allows firms to gather very detailed data on their young users. One potential explanation of this result is that the stringency of privacy protection for children in the US. This has led the FTC to launch several cases aimed at protecting children's privacy such as its case against VTECH, a children's toy manufacturer that collected children's data as part of its attempt to distribute digital toys.[22] In turn this stringency may have led developers in the US to be more reluctant to develop apps targeted at the children's market, leading to an opening for international developers to secure market share. The paper also suggests that consumer protection may sometimes be more effectively advanced by regulators trying to influence Apple's and Google's app stores' global policies towards vulnerable populations, rather than by focusing on changing the regulatory regime within a single country.

Evidence on effects on competition of non-U.S. privacy regulation

In this section we discuss the effects of Non-US privacy regulation on domestic competitiveness and international trade. Given its significance, we mostly focus on the recent EU General Data Protection Regulation (GDPR).[23]

Safe harbor, privacy shield, and international trade

Earlier iterations of European privacy regulation did not in general affect the operations of US firms. Indeed, documents that in response to the 2002 Privacy Directive, internet sites responded very much on the basis of the website's geography. If the website was located in Europe, compliance with the Privacy Directive was observed, while if it was located in the US there was no compliance. This was shown empirically in by documenting that US visitors to EU websites experienced the same suppression of ad-effectiveness experienced by EU visitors to that website, but that when EU visitors went to a US website there was no such suppression of advertising effectiveness.

. . .

22. https://www.fenwick.com/publications/Pages/5-Takeaways-from-the-FTCs-First-COPPA-Settlement-Over-Internet-Connected-Toys.aspx

23. This is also important because elements of the California Consumer Privacy Act which is due to go into effect in 2020 echo GDPR. Internationally, China's proposed 'Data Security Administrative Measures' also echoes elements of GDPR.

There have been two successive policies which intended to add nuance to the role of EU privacy protection in the operation of US firms.

The first policy was known as Safe Harbor and had relatively low compliance costs because it relied on self-certification rather than multiple layers of consent. This policy was judged as legally adequate in 2000 for managing the differences in US and EU privacy policy, however, in 2015 this policy was declared inadequate due to the Snowden revelations that the US government was conducting surveillance on EU citizens' data via large digital US platforms such as Google and Facebook. This led to the rushed replacement of Safe Harbor by a new policy, named Privacy Shield, in 2016.[24] The most striking change between the programs is that any transfer of data to third parties must offer the same level of privacy protection as offered by the original company. However, the future of Privacy Shield is in question because of its potential lack of compliance with GDPR's protections for EU users.[25] This indicates the extent to which GDPR is replacing Privacy Shield as an effective global privacy standard.[26]

The reason this is the case is twofold. First, as indicated by IAB Senior Manager for Privacy and Public Policy Matthias Matthiesen at IAB's annual leadership conference, 'because of the intricacies of compliance, many publishers and organizations coming into compliance are taking a "GDPR everywhere" approach because parsing out consent for users "of unknown citizenship" is actually the more complicated path.'[27] In other words, trying to figure out which website visitor may be subject to GDPR is independently costly, so for some firms it is more cost-effective to apply GDPR to all visitors regardless of origin. Second, there is a lot of uncertainty surrounding to whom GDPR applies. There is also a lot of misinformation. For example, a search for information about how to apply GDPR to a business within the US returns a website that suggests that 'Even a pizza shop in Nebraska, if an EU data subject gives their personal information while on holiday, could invoke GDPR when they return to the EU and receive an email from the pizza shop.'[28] The EU clarified that this was not the case in its November 2018 guidelines on the territorial scope of GDPR.[29] However, the fact that this guidance was only given six months after enactment illustrates that, as for any large and complex regulation, there will be significant uncertainty as to optimal compliance by firms for some time after passage, leading firms to apply GDPR more expansively than intended in order to save costs or reduce regulatory risk.

Some data from the US supports this viewpoint that GDPR is ultimately the first privacy regulation that has directly affected the actions of US firms. As shown in Figure 5, which reports the results of a survey from August 2018 of 145 corporate directors of public

. . .

24. Privacy Shield currently has 4742 organizations who have self-certified as providing substantially similar protections for their users as EU law provides. https://www.privacyshield.gov/list

25. https://iapp.org/news/a/european-parliament-voted-to-suspend-privacy-shield-now-what/

26. ttps://www.siliconrepublic.com/enterprise/privacy-shield-analysis

27. https://searchengineland.com/googles-amp-project-announces-new-consent-component-ahead-of-gdpr-compliance-deadline-295633

28. https://www.christopherspenn.com/2018/04/you-ask-i-answer-gdpr-101-for-marketers/

29. https://edpb.europa.eu/sites/edpb/files/consultation/edpb_guidelines_3_2018_territorial_scope_en.pdf

companies, a large proportion of larger US companies appear to have taken steps to comply with GDPR.

Figure 5: Compliance steps for GDPR

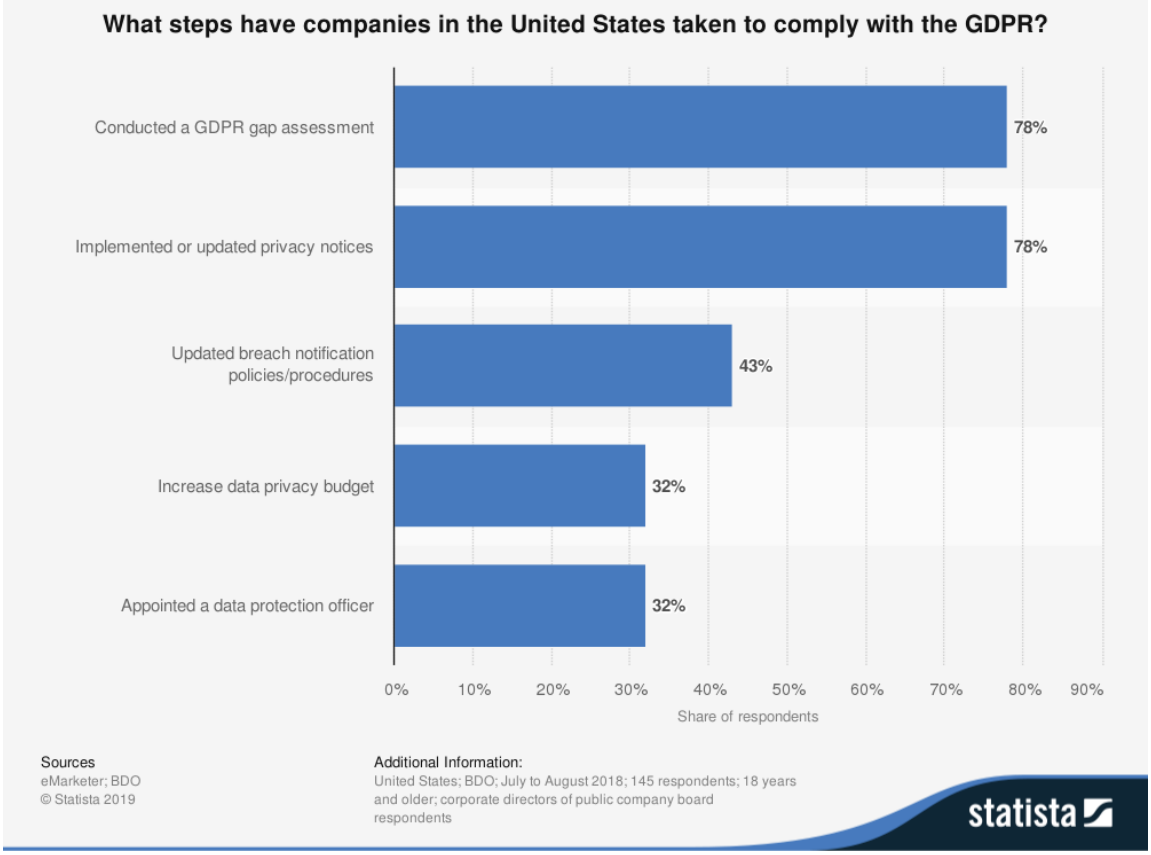

Source: eMarketer. (October 19, 2018). What steps have companies in the United States taken to comply with the GDPR? [Graph]. Retrieved August 02, 2019, from https://www.statista.com/statistics/946800/steps-taken-companies-taken-comply-gdpr-usa/

In addition, GDPR is unusual in its international scope in that its penalties are for 4% of the firm's global revenues, not just the revenues focused on Europe. This means that a US-based firm with a relatively small proportion of EU customers or visitors to its website, could have a far larger exposure than might be suggested by the amount of EU citizen data it processed.

Could GDPR entrench existing digital firms?

There are several papers that attempt to quantify the early effects of GDPR on the digital economy.

Jia et al. (2018) investigates the trends of venture funding of EU relative to US firms. They find that in Europe there was a relative suppression in terms of dollar amounts, the number of details and the number of dollars raised per deal after the enactment of GDPR. The

results of this paper could be taken as evidence that the concern expressed in section 4.1 that GDPR would have an equal impact worldwide as it does in Europe has not come to pass. However, as the authors emphasize, their findings hold only for the data period they study, which spans July 2017 to September 2018. Another facet of their work which is of interest is the evidence on the effects of GDPR depending on the age of the firm. Their results suggest that, where there was a measurable effect on the amount raised per deal, it was similar across new and more established firms. Given that typically we would expect that a venture capitalist deal size would be larger for a larger firm, this suggests that there was a disproportionate drop in the relative size of venture capital funding for the smallest and youngest firms as a result of GDPR.

Goldberg et al. (2019) investigates the mechanics of GDPR by looking at how the regulation affected the ability of firms to measure the spending and visits of their visitors from the EU. They suggested that reported spending and visits declined by as much as 7% for EU visitors in 2018 over the period that GDPR was enacted relative to 2017. This paper is important because it shows the effects of privacy regulation on a little-thought-about input to effective advertising and website management, which is the ability to measure consumer behavior. Furthermore, they show that the effects were most profound for the smallest firms and that larger firms suffered less of an effect regarding their ability to record visits and revenues. Indeed, their estimates suggest that for e-commerce sites, the effects for small firms were roughly four times greater in terms of orders than recorded for larger firms.

Adjerid and de Matos (2019) provide some of the most detailed early evidence on how GDPR is reinforcing incumbency. They study a series of field experiments launched by a large telecommunications provider within Europe after GDPR. These field tests were designed to try and encourage users to consent within GDPR regulations. They show that these experiments were incredibly successful. Consequently, the telecoms provider ultimately was able to process more personal data after GDPR then it was prior to GDPR. This suggest that one artifact of GDPR and other consent-based privacy regulations may ultimately be that after gaining consent, firms are able to process a larger scope of user data than they previously processed as they optimize around ways of gaining consent. If consent is more likely to be given to incumbent firms (as is argued in Campbell et al. (2015)) then this is another mechanism by which incumbency can be reinforced by privacy regulation.

Does GDPR constrain international competition?

As of yet, there has been little work on whether indeed GDPR has acted to constrain the ability of firms outside of the EU to compete. Instead most empirical evidence is that of a termination in services of US firms as a result of GDPR. document that as a result of GDPR, several firms discontinued their use of the Adobe Cloud Websites analytics products. Though this was not central to their analysis - they ended up excluding the firms that discontinued using Adobe Cloud as a result of GDPR as they had no further data on them - this is direct evidence of the effects of a European privacy regulation on US trade since Adobe is a US firm.

Perhaps the most well-known and visible effects of the constraints posted by GDPR to US firms are the effects on the news media. Many European users have reported instances of US newspapers not displaying content within EU borders. As shown in Figures 6 and 7, this includes well-known newspapers with long histories and nationwide reach, such as the Chicago Tribune and the LA Times. As shown in Figure 8, USA Today developed a slimline version of their website with no ads, which showed only a few stories to visitors from the EU. Newspapers such as the Washington Post charged a higher price to EU residents and provided a non-tracking version of the newspaper.[30]

Figure 6: The Chicago Tribune unavailability notice

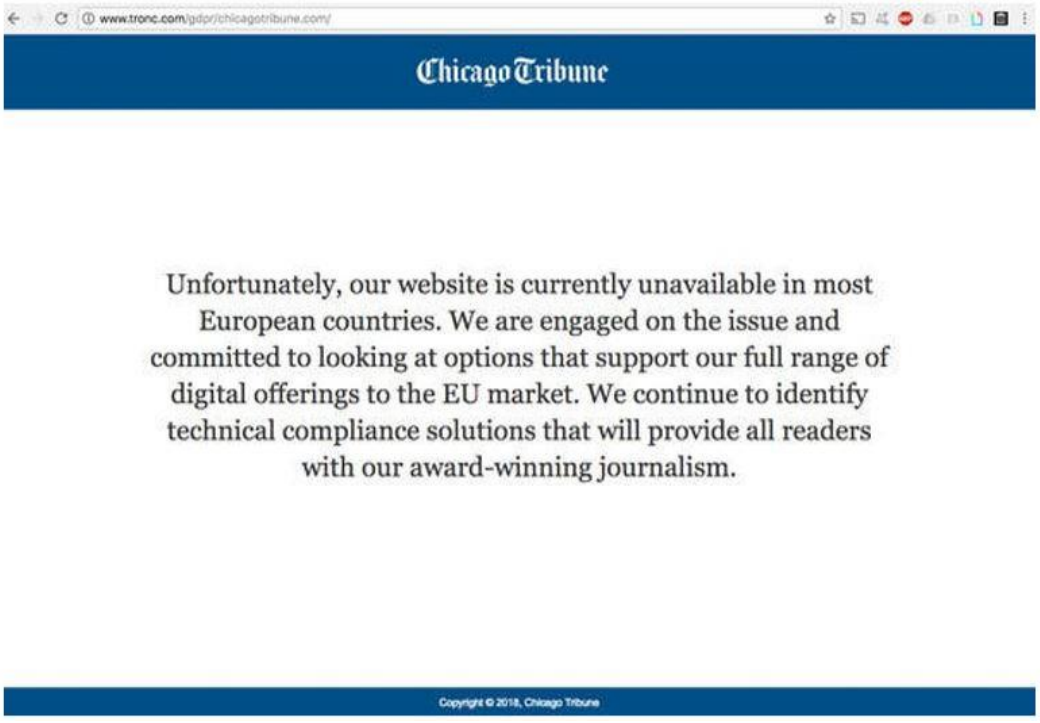

Though most of the discussion has centered on the ability of EU citizens to access US news, there has been less discussion about why the news media has responded to GDPR in this manner. The reasons can be found in the economics of the online news industry and advertising. As shown in, earlier rounds of EU privacy regulation that were constrained to EU firms had the most negative effects on the ability of the news media to support themselves through advertising. Relative to other types of content sites, such as websites devoted to babies and travel, it is difficult for newspaper sites to use news content as an effective ad targeting strategy. As a result, newspapers, more than any other type of website, need detailed user data to determine what ads to show to each pair of eyeballs that arrives on their website.

. . .

30. https://digiday.com/media/washington-post-puts-price-data-privacy-gdpr-response-tests-requirements/

Figure 7: LA Times unavailability notice

Figure 8: US Today restricted content notice

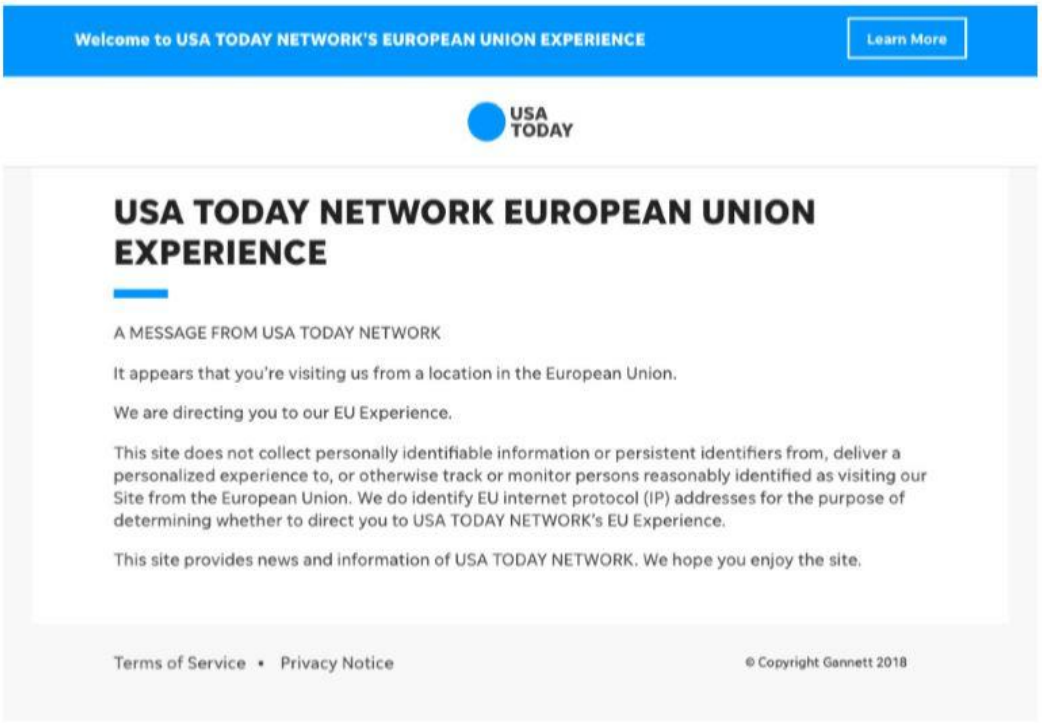

Though most of the coverage of the news media's response to GDPR has focused on well-known newspapers, a question of interest is the extent to which smaller newspapers have responded. In an analysis of 1361 news media websites that were initially blocked through GDPR, data suggests that only 252 (18 percent) have unblocked since the enactment of GDPR.[31]

Effects of government surveillance on competition

Thus far, we have focused on the effect on competition of consumer-oriented privacy regulation. However, the discussion of Privacy Shield above shows that there is a necessary additional dimension to consider, which is the effect on competition policy of governments' own interest in exploiting the personal data of their citizens, and citizens' interest in shielding that data from the government. In this section, we argue that government surveillance and its privacy implications can indeed shape competition.

With regard to surveillance, the consumer harm arising from surveillance tends to be couched in terms of "chilling effects," or the notion that consumer awareness of government surveillance will tend to "chill" their speech or actions. Before the Snowden revelations of 2013, the courts tended to treat plaintiffs' suggestions that they had been chilled by government surveillance as being "speculative." However, those revelations, by providing an exogenous shock to consumers' awareness of US government surveillance, enabled the kind of difference-in-difference analysis that had previously been impossible, and allowed researchers to begin to quantify the chilling effects from surveillance. found a chilling effect on consumers' Google searches from the Snowden revelations, particularly involving searches on health-related terms in countries allied to the US. extended our quantitative method to detect a chilling effect on Wikipedia and detected chilling effects from surveillance on Facebook in a lab setting. In further work, summarizes existing research on and conceptualizations of chilling effects and finds survey-based evidence of "networked chilling effects", where threats of surveillance received by a friend chill a user's behavior. It is now fairly well established that chilling effects from government surveillance and other privacy violations exist, and that significant proportions of consumers both describe themselves as being harmed by them and alter their actual behavior in response to that surveillance. What is not clearly established is whether governmental preference for deep surveillance-related relationships with incumbents possessing data of interest plausibly has had practical implications for the implementation of privacy protection or competition policy.

The central economic insight here is that, much like with consumer privacy and competition, governments incur fixed costs in establishing data-sharing relationships for surveillance purposes with potential data providers. As a result, it is theoretically more effective for government agencies in charge of surveillance if they develop long-term data-sharing relationships with large, entrenched, incumbent firms, rather than encouraging low

. . .

31. Based on data from https://data.verifiedjoseph.com/dataset/websites-not-available-eu-gdpr

barriers to entry in the telecommunications and technology sectors; in turn, data-sharing relationships with large partners impose costs, in terms of chilling effects, on hundreds of millions of users.

The PRISM program, begun in 2007 and revealed by Edward Snowden in June 2013, is perhaps the best way to explore the potential effect large-scale data-sharing relationships with government on the competitive landscape. PRISM involves NSA exploitation of tech firm consumer communications on a mass scale and is "the number one source of raw intelligence used for NSA analytic reports". Partner companies include Microsoft, Yahoo!, Google, Facebook, PalTalk, YouTube, Skype, AOL, Apple, and now also Amazon.[32]

This also illustrates another potential cost for larger firms of not complying with surveillance requests: Amazon and Microsoft are the finalists in bidding on a ten-year, $10 billion Department of Defense cloud computing and AI program; refusal to allow its servers to be transparent to PRISM might disadvantage Amazon or Microsoft in that process.[33]

A contrasting example of the kinds of decision faced by a small firm in cooperating or not with governmental requests for data is that of Lavabit. Lavabit was a small email provider that marketed itself as providing particularly private and heavily encrypted email services, with a premium offering that offered the highest level of encryption. This attracted Edward Snowden to sign up to their service. After he came forward, the FBI approached Lavabit's CEO, Ladar Levison, with a 'pen register' order for the metadata for Snowden's account. Levison explained that the account-level encryption Snowden had paid for made it impossible for Lavabit to read the metadata on his email. The FBI then ordered him to disclose his "developer-level keys," decrypting all Lavabit accounts so that they could reach Snowden's. Eventually, Levison provided the keys in a form the FBI could easily read but chose to shut down his service the next day, rendering those keys useless, rather than to "become complicit in crimes against the American people."

The distinction between Amazon, Microsoft and Lavabit here is instructive. Large companies easily become partners of the surveillance agencies, and their CEOs have an interest in demonstrating that they are good (corporate) citizens. Smaller firms seize on the data practices of larger firms to market themselves to privacy-conscious consumers even if, as suggested by the DuckDuckGo example we discussed earlier, this market may be limited. Lavabit was small enough for Levison to shut down and walk away from without causing collateral harm to tens of millions of users. Amazon or Microsoft, after declaring that it had not participated in PRISM, may have changed its mind in order to access contracting opportunities not available to smaller or more privacy-conscious firms. identified that smaller firms may find it harder than incumbents to respond to general privacy regulations, especially those that require the implementation of a consumer consent-based privacy framework. Here, this history indicates that larger firms may find it harder to resist surveillance mandates from government agencies, but that resisting those mandates may in some cases force smaller firms offering privacy-sensitive end-to-end encrypted solutions to shut down

. . .

32. https://techcrunch.com/2019/08/01/amazon-prism-transparency-data/

33. https://finance.yahoo.com/news/project-jedi-amazon-microsoft-pentagon-175143574.html

entirely. The end result may be a sharply limited ability of new firm entrants to compete on privacy, and an entrenchment of large firm partnerships with government.

Firms' reactions to governments' surveillance requests may also be a function of firm size, depending on the nature of requests. To take one example, at Google front-end law enforcement requests for user data have risen from around 9,000 in 2010 to over 22,000 in 2019.[34] Of course, it is larger firms who will be better placed to employ complex law enforcement data request response teams, and smaller firms who may bear larger costs by a poorly timed or overbroad data request.

Calibrating the tradeoffs between privacy regulation and competition

This paper has argued that there is a potentially irreconcilable conflict between the aims of privacy regulation and the aims of competition law. Though this may initially appear a pessimistic conclusion, it is probably no more pessimistic than some of the writing in economics about the tradeoffs between innovation and competition. The tradeoff of giving dynamic incentives for innovation, in the form of inventor monopoly rights over patented inventions, against the considerations of ensuring static competition at any one time, is widely acknowledged, and both elements of this tradeoff are worthwhile goals for regulation. Similarly, there are also well documented tradeoffs between regulation designed to ensure environmental protection and competition.

However, where this comparison falters is that the field of the economics of innovation has developed useful tools for calibrating that tradeoff that allow measurement of the benefits of intellectual property protection for innovation which can be compared to the negative effects of suppressing short-term competition. There are also well-established ways of measuring the benefits of environmental protection. Further, there are similarly well-developed tools for measuring the benefits of competition policy. By contrast, there are few tools or accepted methods of measuring the potential benefits of privacy regulation. Here, we discuss potential measures applicable to the world of commercial privacy, and their limitations.

There are a few outcomes that can be directly measured, such as increases or decreases in cases of identity theft, and this offers the benefit that only identity thieves want there to be more cases. Even here, however, the empirical evidence as to the efficacy of even straightforward regulatory protections such as breach notification requirements is weak, and there is a difference between how many cases are prosecuted and how many incidents of identity theft occur.

Another way of measuring outcomes is the number of data breaches reported. GDPR has increased the number of data breaches reported - since it passed, there have been over

. . .

34. https://transparencyreport.google.com/user-data/overview?hl=en&user_requests_report_period=series:requests,accounts;authority:US;time:&lu=user_requests_report_period

64,684 data breaches reported-[35], but it also increased incentives dramatically for reporting data breaches, meaning that the actual outcome of interest - whether a data breach occurs - is hard to measure. Normatively, there is the added complication that a minority of data breaches - such as leaks of information pertaining to government or corporate lawbreaking or abuse - may in fact be desirable.

So far there have been 94,622 complaints submitted to EU authorities under GDPR, with 55 million in fines recorded.[36] Data from the Irish Data Protection Commission suggests this is an increase in the number of complaints relative to complaints submitted under the older Data Protection Act. In 2017, it had received 2642 complaints in total for 12 months, and in 2018 for the 7-month period GDPR was in force it received 2864 complaints.

Further, though information security risks are related to the aims of privacy regulation, many of the aims of privacy regulation, such as dignity and consumer control, are not captured by measures of data security effectiveness. As of yet, there is some survey evidence regarding broader consumer responses to GDPR. There is evidence that GDPR has increased the number of cookie consent notices that EU subjects are exposed to by 16%. This represents an increase in control, but whether it increases consumer welfare depends on whether consumers value more highly the increase in a sense of control or a fluid customer experience. A French survey of 1,000 adults suggested that 50% of people accept the conditions for data collection notices without reading them, even though 76% are worried about the collection of their data.[37], suggesting that there is some difference between their stated preference for increased privacy and their revealed preference for a fluid customer experience.

A 2019 survey of 1008 people from France gives some impression of the sense of control they feel they have over their privacy, and what they feel about GDPR's role in protecting it. 9% of respondents felt that there was no action they could take in the wake of recent scandals related to the use of personal data. 26% of respondents chose a personal action that would improve the situation (a "massive boycott of the websites concerned"), whereas 66% of respondents cited policy actions (the two on offer in the survey were GDPR (14%) and "really punishing fines" (52%)).[38] This suggests, without knowing how much the respondents know about GDPR, that a majority in France view the problem with data breaches as not having been solved by GDPR, and that privacy concerns require further, more punitive state action, with coordinated consumer action playing a secondary role.

. . .

35. http://www.europarl.europa.eu.libproxy.mit.edu/meetdocs/2014_2019/plmrep/COMMITTEES/LIBE/DV/2019/02-25/9_EDPB_report_EN.pdf

36. http://www.europarl.europa.eu.libproxy.mit.edu/meetdocs/2014_2019/plmrep/COMMITTEES/LIBE/DV/2019/02-25/9_EDPB_report_EN.pdf

37. http://www.odoxa.fr/sondage/donnees-personnelles-francais-se-disent-preoccupes-utilisation-ne-se-premunissent-toujours/

38. Another response, "better technologies to protect personal data", attracted 34% of responses, but was poorly phrased, because we cannot tell whether the respondents took this as a cue that *they should adopt currently available technologies* or that *the government should develop and then respondents, or French people in general, should adopt more protective technologies*.

Our ability to understand the true tradeoffs between competition and privacy protection is hampered for the time being. In general, the privacy paradox suggests that simple measures such as stated 'willingness to pay' are unlikely to reveal true privacy preferences. Until we have effective measures of the benefits of privacy regulation, it is hard to compare them against a less competitive landscape.

References

(2013). NSA slides explain the PRISM data-collection program.

Adjerid, I. and M. G. de Matos (2019). Consumer Behavior and Firm Targeting after GDPR:

The Case of a Telecom Provider in Europe.

Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005). Competition and innovation: An inverted-U relationship. The Quarterly Journal of Economics 120(2), 701–728.

Athey, S., C. Catalini, and C. Tucker (2017). The Digital Privacy Paradox: Small Money, Small Costs, Small Talk. National Bureau of Economic Research.

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

Cecere, G., F. L. Guel, V.Lefrere, C.Tucker, and P. Yin (2019). Privacy and children: What drives digital data protection for young children?

Degeling, M., C. Utz, C. Lentzsch, H. Hosseini, F. Schaub, and T. Holz (2018). We Value Your Privacy... Now Take Some Cookies: Measuring the GDPR's Impact on Web Privacy.arXiv preprint arXiv:1808.05096.

Dolmen (2019). Les franais et les scandales lis aux donnes personnelles.

Franceschi-Bicchierai, L. (2016, March). Lavabit's Forgotten Encryption Fight Looms Over the Apple Case.

FTC (2011). Forty Years of Experience with Fair Credit Reporting Act.

Fuller, C. S. (2017). The perils of privacy regulation. The Review of Austrian Economics 30(2), 193–214. 33

Goldberg, S., G. Johnson, and S. Shriver (2019). Regulating Privacy Online: The Early Impact of the GDPR on European Web Traffic & E-Commerce Outcomes. Mimeo, Northwestern.

Goldfarb, A. and C. Tucker (2011a, May). Online display advertising: Targeting and obtru siveness. Marketing Science 30, 389–404.

Goldfarb, A. and C. E. Tucker (2011b, January). Privacy regulation and online advertising. Management Science 57(1), 57–71.

ICO. Update report into adtech and real time bidding.

Jia, J., G. Z. Jin, and L. Wagman (2018). The short-run effects of GDPR on technology venture investment.

Kim, J.-H. and L. Wagman (2015). Screening incentives and privacy protection in finacial markets: a theoretical and empirical analysis. The RAND Journal of Economics 46(1), 1–22.

Marthews, A. and C. Tucker (2014). Government surveillance and internet search behav ior. Mimeo, MIT.

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

Ohlhausen, M. K. and A. P. Okuliar (2015). Competition, consumer protection, and the right [approach] to privacy. Antitrust Law Journal 80(1), 121–156. 34

Penney, J. W. (2016). Chilling effects: Online surveillance and wikipedia use. Berkeley Technology Law Journal 31, 117.

Penney, J. W. (2017). Internet surveillance, regulation, and chilling effects online: A com parative case study. Internet Policy Review.

Romanosky, S., R. Telang, and A. Acquisti (2011). Do data breach disclosure laws reduce identity theft? Journal of Policy Analysis and Management 30(2), 256–286.

Sampat, B. and H. L. Williams (2019). How do patents affect follow-on innovation? Evi dence from the human genome. American Economic Review 109(1), 203–36.

Stoycheff, E. (2016). Under surveillance: examining Facebook's spiral of silence effects in thewake of NSA internet monitoring. Journalism & Mass Communication Quar terly 93(2), 296–311.

Vedder, H. (2003). Competition law and environmental protection in Europe: towards sustainability? Volume 3. Europa Law Publishing.

Economic Studies
at BROOKINGS

The Brookings Economic Studies program analyzes current and emerging economic issues facing the United States and the world, focusing on ideas to achieve broad-based economic growth, a strong labor market, sound fiscal and monetary policy, and economic opportunity and social mobility. The research aims to increase understanding of how the economy works and what can be done to make it work better.