NBER WORKING PAPER SERIES

# ECONOMIC RESEARCH ON PRIVACY REGULATION: LESSONS FROM THE GDPR AND BEYOND

Garrett Johnson

Working Paper 30705
http://www.nber.org/papers/w30705

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2022

I thank Samuel Goldberg and Scott Shriver as well as participants at both the NBER's 2022 Privacy Tutorial and the European Commission's Joint Research Centre (Digital Economy Unit) for providing helpful comments. I thank the NBER for providing funding for this work. I dedicate this work to Luke. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Economic Research on Privacy Regulation: Lessons from the GDPR and Beyond
Garrett Johnson
NBER Working Paper No. 30705
December 2022
JEL No. K2,L51


## ABSTRACT

This paper reviews the economic literature on the European Union's General Data Protection Regulation (GDPR). I highlight key challenges for studying the regulation including the difficulty of finding a suitable control group, variable firm compliance and regulatory enforcement, as well as the regulation's impact on data observability. The economic literature on the GDPR to date has largely—though not universally—documented harms to firms. These harms include firm performance, innovation, competition, the web, and marketing. On the elusive consumer welfare side, the literature documents some objective privacy improvements as well as helpful survey evidence. The literature also examines the consequences of the GDPR's design decisions and illuminates how the GDPR works in practice. Finally, I suggest opportunities for future research on the GDPR as well as privacy regulation and privacy-related innovation more broadly.



Garrett Johnson
Questrom School of Business
595 Commonwealth Avenue
Boston, MA 02215
United States
garjoh@bu.edu

# 1 Introduction

Privacy is a conundrum. Privacy and the data economy are two sides of the same coin. Viewed from each side, progress on the respective dimension can seem obvious. Nevertheless, the two are often at cross-purposes.[1] Economic researchers can illuminate our understanding of privacy, the data economy, and the trade-offs involved. Policy-makers and regulators worldwide wrestle with crafting and enforcing privacy regulation. Economic research can inform their difficult task.

The European Union's General Data Protection Regulation (GDPR) is a landmark privacy regulation that elevated the tension between privacy and the data economy. The European Union (EU) passed the GDPR in April 2016, but delayed enforcement until May 25, 2018. In many ways, the GDPR set the privacy regulation agenda globally. More than a dozen countries have since passed privacy regulation including Brazil, China, India, and New Zealand. At the heart of the regulation, the GDPR defines personal data expansively to include all data relating to an individual. The regulation provides EU residents with multiple data rights, like the right to access and delete their data. The GDPR imposes responsibilities on firms like data auditing and data-breach notification. The regulation also lays out multiple legal bases—including consent—for processing personal data. The GDPR's maximum fines of 4% of a firm's annual revenue ensured it caught the attention of firms and the wider public. As a landmark and influential regulation, the GDPR is of great interest to economists.

However, the GDPR poses three key challenges for empirical research. First, economists often examine the GDPR as an event study, but may lack a suitable control group in certain settings. In particular, the GDPR covers most of Europe and also has substantial global spillovers that contaminate candidate control group members. Second, GDPR compliance and enforcement vary by industry, compliance requirement, firm size, country, and over time. This creates gaps between the regulation as written and the regulation in practice which, in

[1] The World Bank devoted its 2021 World Development Report to exploring this tension.

turn, complicates the conclusions we can draw from GDPR research. Third, the GDPR may directly restrict the availability—and selection into—individual-level data that economists can use to understand the impact of the regulation. As we will see, economists have proposed various solutions to and workarounds for these challenges.

Almost five years after the GDPR's enforcement deadline, economic research on the GDPR is maturing. To date, much economic research examines the GDPR's impact on firms. The GDPR hurt firm performance by imposing costs, decreasing revenue, and thereby hurting profitability. Venture funding for technology firms fell—particularly for more data-related ventures. In some cases, the GDPR both accelerated market exit and slowed entry. At the same time, the GDPR created an opportunity to test hypotheses about the consequences of privacy regulation for firm competition and innovation (see e.g., Goldfarb & Tucker, 2012). Research shows that the GDPR hurt competition by creating greater harms for smaller firms and by increasing market concentration in the data vendor market. The evidence for innovation is more mixed, though several studies suggest that the GDPR constrained data-related innovation. Research shows that the GDPR reduces the share of individual-level data available to firms. When firms rely on consent to process data, consumer data becomes self-selected though consenting consumers tend to be favorably selected. On the web, studies show a decrease in EU traffic to websites after the GDPR as well as a short-lived reduction in sites' use of third-party vendors. However, the GDPR had no apparent effect at the Internet's connectivity layer or on website content provision. Finally, the GDPR seemed to constrain firms marketing activities for personalized channels like email and online display advertising.

Fewer studies examine the GDPR's consequences for consumers, though this gap largely reflects the inherent measurement challenge. Survey evidence quantifies consumer valuations for their data rights as well consumer's awareness of privacy and perceived control over their personal data. Structural modeling suggests consumer harm from the GDPR's adverse impact on innovative product development. Theory evidence suggests varying consequences

of certain elements of the GDPR for both firms and consumers. Finally, empirical research shows post-GDPR reductions in data collection that suggest objective improvements in consumer privacy.

The economics literature also illuminates the consequences of the GDPR's design decisions. The literature documents important spillovers of the GDPR outside of the EU. In particular, research shows that foreign firms that serve EU consumers can exhibit higher compliance levels. These firms fall under the GDPR's extraterritoriality component and may be leery of GDPR fines that are based on global revenue rather than EU revenue alone. Research also shows indirect spillovers like global firms implementing their compliance efforts worldwide, so that non-EU consumers benefit. Though the GDPR intended to harmonize regulation within the EU, several scholars document differences in regulatory impact by the perceived strictness of EU country-level regulators.

This review is far from the last word on the GDPR, as the literature and the practical application of the regulation are both still evolving. I focus on the economic literature and the empirical economic literature in particular. However, the study of the GDPR is inherently interdisciplinary, so I occasionally draw on research from law and computer science. This review was commissioned for the NBER Privacy Tutorial in October 2022. As such, my emphasis on research challenges and future research opportunities in part stems from that tutorial's doctoral student audience. Nonetheless, I think that this emphasis is helpful for understanding the literature and the shape it has taken so far. For future research, I indicate more privacy-related changes—whether through regulation or platform policies—that provide possible event studies for empiricists. I also suggest that economists should study privacy-enhancing technologies that are beginning to be commercialized, as these technologies improve the tradeoff between privacy and economic uses of data.

This review builds on previous review articles on the economics of privacy and complements other work by great scholars in this volume. For instance, Acquisti et al. (2016) provide a general introduction to the economics of privacy. In this volume, the chapter

by Miller (2022) on health information privacy describes important antecedents to GDPR research that often exploit changes in health privacy regulation. Carrière-Swallow & Haksar (2019) and the World Bank (2021) examine data policy from an economic perspective. Goldfarb & Tucker (2012) discuss the economics of privacy and innovation. Notably, Prasad & Perez (2020) provide an early review of the economic literature on the GDPR.

The rest of this guide is organized as follows. Section 2 begins by providing a background on the GDPR. Section 3 discusses key challenges that the GDPR poses for empirical research. Section 4 reviews the economic literature to date on the GDPR. Section 5 highlights some avenues for future research on the economics of privacy regulation. Section 6 concludes.

# 2 GDPR background

The GDPR is a lengthy and multifaceted regulation, which opens many avenues for economic research. In this section, I share background on the regulation *as written* for economists. The GDPR contains 99 articles and is supported by an additional 173 recitals. Jones & Kaminski (2020) provide a helpful background for those who are more familiar with the American legal context. Jones & Kaminski (2020) point out the the GDPR is situated within a broader legal context that includes the EU Charter, complementary EU and national privacy regulations, EU privacy regulator guidance, EU judicial rulings, and the EU's 1995 Data Protection Directive that preceded the GDPR.

I begin by laying out the regulation's essential features. The GDPR takes a broad approach to data protection regulation by defining personal data as all data relating to a person (Article 4(1)). This extends beyond personally-identifiable information like a name or address to include pseudonymous identifiers and online identifiers. For brevity, I refer to personal data as simply data below. The GDPR refers to the "processing" of data which includes data collection, storage, use, analysis, sharing, and more (Article 4(2)). The GDPR further distinguishes what it refers to as "special category data" as being particularly privacy-sensitive.

This includes data on health, genetics, sexual orientation, political opinions, religious beliefs, and more (Article 9(1)). Though this review focuses on firms, the GDPR covers all individuals and institutions (e.g., governments and non-profit organizations) that process personal data.[2]

The GDPR establishes six data rights for EU residents (Articles 12-23). Under this regulation, residents gain the right to access data that a firm has about them. Furthermore, residents gain the right to correct or delete their data—the latter is often referred to as the "right to be forgotten." Residents even have the right to port their data to another firm. Residents have the right to object to data processing and even the right to object to decisions made on the basis of automated processing.

The GDPR imposes a number of responsibilities on firms (Articles 24-43). Firms have to fulfill the above rights-related responsibilities in a timely manner. Firms need to audit their data processing activities—also known as a Data Protection Impact Assessment. Firms need to minimize their data processing activities—i.e., data protection by default—which is also a key principle of the GDPR (Article 5(1c)). Firms must encrypt and pseudonymize the data they process—i.e., data protection by design. In the event of a data breach, firms must notify the regulator and affected consumers within 72 hours. Finally, firms should designate a data protection officer—either an employee or an external consultant—to oversee their data protection-related activities.

Though consent sometimes plays an outsize role in discussions about the GDPR (Jones & Kaminski, 2020), consent is only one of the GDPR's six legal bases for processing data (Article 6(1)). These legal bases are consent, contractual obligation, legitimate interest, legal obligation, vital interest of an individual, and public interest. For most firms, the first three bases are most relevant. As an example, an ecommerce website could use contractual obligation as a legal basis for processing a consumer's name and address information for

[2]The GDPR distinguishes between data controllers and data processors (Article 4(7-8)). This distinction refers to cases where, e.g., firm X delegates data processing to firm Y, but firm X retains decision rights regarding the data processing. In this example, firm X is the data controller and firm Y is the data processor.

the purpose of shipping products to the consumer. Legitimate interest is the most flexible of the legal bases, but it is not a carte blanche as it should not override an individual's right to privacy (ICO, 2021). Legitimate interest carries additional duties like carrying out and documenting a "legitimate interest assessment" that weighs the firm's specific interest against consumers' privacy interest (ICO, 2021). Regardless of the legal basis, the firm should provide information to the consumer including the purpose(s) of data processing, the relevant legal basis, the contact information of the data protection officer (where applicable), and the identities of all third-party data recipients (Article 13). Note that special category data has additional restrictions (Article 11) as does child's consent (Article 7).

The GDPR sets a high standard for consent (Article 7). Consent should be an unambiguous, affirmative act like ticking a box on a website: pre-ticked boxes or inactivity do not indicate valid consent (Recital 32). Consumers must be able to withdraw consent at any time and just as easily as they provided consent. In obtaining consent, firms must inform consumers using plain language. Consent should be granular to the purpose(s) of processing (Recital 32). As mentioned above, this includes listing all third-party data recipients. Consent should be freely given in that the firm should not condition its consumer offerings on consent when these do not require data processing. Finally, firms must be able to show a record of the consumer's consent.

The GDPR also covers data transfer outside of the European Economic Area (i.e., EU plus Iceland, Liechtenstein, and Norway). Article 45 permits the transfer to countries that have adequate data protection. As of now, the European Commission deems 14 countries as adequate including: Argentina, Canada, Israel, Japan, New Zealand, South Korea, Switzerland, the United Kingdom, and Uruguay. Articles 45 to 50 lay out alternative data transfer arrangements including the foreign firms' adherence to standard contractual clauses adopted by the European Commission. Data transfers to the United States remain a thorny issue, however. The 2016 "EU-US Privacy Shield" permitted data transfers to certified firms, but was invalidated by the Court of Justice of the European Union in 2020. In 2022, the Euro-

pean Commission and US agreed in principle to a new data-transfer arrangement, but this is still being finalized.

The GDPR charges Data Protection Authorities (DPAs) in all EU countries with enforcing the regulation (Articles 51-59). DPAs are charged with regulating data processing by firms that are located in their country, that substantially affects their country's residents, or for which they have received a complaint by a resident or organization in their country (Articles 4(22), 57). Though the GDPR was intended to harmonize EU-wide regulation, regulators vary in resources by country (EDPB, 2020). For multi-national firms, the GDPR's "one-stop shop mechanism" allows firms to select a country as their lead regulator by locating their headquarters in that country (Article 56). The lead regulator mechanism simplifies the firm's dealings with EU regulators, though firms may therefore prefer to locate their headquarters in countries which they believe have weaker DPAs. Nevertheless, other EU DPAs retain considerable rights in multi-national cases (Article 60). The GDPR also establishes an EU-wide European Data Protection Board consisting of the European Data Protection Supervisor and the head of each country's DPA (Article 68). The board issues guidelines, promotes cooperation between DPAs, issues opinions on draft DPA decisions, and resolves disputes between DPAs (Articles 65, 70).

The GDPR stipulates that firms can be fined up to the greater of €20 million or 4% of their global annual revenue (Article 83(5)). For lesser infractions, the maximum fines are halved (Article 83(4)). Presthus & Sønslien (2021) provide an analysis of the first two years of GDPR fines. They use Enforcementtracker.com as one data source, as this site maintains a list of GDPR fines that are made public. As of September 2022, this site lists 1,279 fines totaling €2 billion and averaging €1.6 million per fine. The majority of these fines are €10,000 or less. The largest seven fines have all been issued to big technology firms: Amazon (the single largest fine), Meta (3 fines), and Alphabet/Google (3 fines). The countries that have issued the most fines are Spain (496), Italy (181), and Germany (115). The total value of fines are highest for Luxembourg (€746 million), Ireland (€649 million),

then France (€272 million). However, DPAs can instead handle cases by warning firms and these instances are usually not documented publicly. Beyond the administrative fines, the GDPR also includes a private right of action, whereby consumers can seek compensation for privacy-related damage suffered through their country's courts (Article 82).

In sum, the GDPR is a multi-faceted regulation that increases the legal risk and cost associated with data processing. In later sections, we will discuss still more features of the GDPR for researchers to consider. As we will see in Section 3.2, the GDPR is further complicated by the sometimes substantial gap between the regulation as written and the reality on the ground.

# 3 Research challenges

The GDPR represents a tremendous opportunity for economists to study privacy regulation and its impact. Nevertheless, the GDPR poses several challenges for research. Below, I focus on three key challenges and describe solutions devised from the literature.

Most economists study the GDPR as an event study. I begin by recalling a leading approach for analyzing event studies: difference-in-differences (see e.g., Cameron & Trivedi, 2005). Difference-in-differences combines two comparisons. First, we compare a treatment group that is subject to the policy with a control group that is not. These groups should satisfy the stable unit treatment value assumption (SUTVA), meaning that the GDPR does not affect the control group. Second, we compare outcomes before and after the policy. As the name suggests, the difference-in-differences approach estimates the policy's impact by subtracting the before-after means comparison in the control group from that of the treatment group. The identifying assumption is that the treatment and control groups' outcome variable would follow parallel trends after the policy, but for the policy's impact.

The GDPR poses several problems for this analysis framework. Section 3.1 discusses the potential challenge of finding a suitable control group that satisfies both SUTVA and

the parallel trends assumption. Section 3.2 notes that both firm compliance and regulatory enforcement were variable under the GDPR. This poses a problem for generalizing from the real-world estimated impact of the GDPR—or lack of thereof—to the regulation as written. Finally, Section 3.3 notes the GDPR's confounding impact on data observability. By construction, the GDPR creates a missing-data problem whereby observed individual-level data are selected and the corresponding aggregate statistics are incomplete.

## 3.1 Lack of a suitable control group

Most economists study the GDPR as an event study. Event studies should include 1) a suitable control group, and 2) a clear start date. These criteria are often to challenging to address satisfactorily. In the case of the GDPR, both criteria pose problems for research, though the first is unusually challenging.

The GDPR's scale and global scope can make a suitable control group difficult to find in many cases. First, the GDPR's large scale makes it appealing to study, but limits the set of suitable control countries. The GDPR covers 28 EU countries and another 3 European Economic Area countries. To put the problem starkly, a substantial idiosyncratic economic shock to the EU after May 2018 would bias many economic studies. Second, the GDPR has substantial spillovers outside of Europe because the regulation's scope includes not only EU firms, but also non-EU firms that target EU residents. For instance, a Canadian ecommerce site that offers shipping to customers in the EU is also subject to the regulation. Third, the GDPR may have indirect spillovers outside of the EU as well. International firms may choose to roll out their GDPR compliance efforts globally due to cost efficiencies in treating their customers and data uniformly. Furthermore, the GDPR raised the attention paid to privacy worldwide and—to some extent—raised global commercial compliance standards to the EU's high standard. Bradford (2020) refers to such phenomena that, in effect, export EU policy globally as the "Brussels Effect."

GDPR researchers need to also reflect on the appropriate timing to use. The GDPR has

two main start dates to consider: its passage in April 2016 and its enforcement deadline in May 2018. The GDPR affects all EU countries simultaneously, unlike past research that benefited from variation in the timing of privacy regulation (e.g., Miller & Tucker, 2009). Most studies focus on the latter enforcement date, but some consider both. For instance, firms may have incurred compliance costs before and after the enforcement deadline. If consumer-facing compliance efforts come online after the deadline, the GDPR's effect on revenue may manifest after the deadline. In some cases, anticipatory compliance may attenuate GDPR impact estimates. In other cases, firms may have delayed compliance until the enforcement deadline or even later. In sum, researchers should evaluate the relevant timing in their setting as a function of its underlying economics and its institutional realities.

Many GDPR papers use difference-in-differences as their identification strategy and most use non-EU countries (or units therein) as a control. For instance, Aridor et al. (2020) examine data from travel websites and argue that these have "separate, country-specific, versions of their websites," so that the sites' requirement to comply with the GDPR is clear. Moreover, Aridor et al. (2020) use non-EU travel websites in northern hemisphere countries as a control group, so that these sites are both exempt from the GDPR and should have similar seasonal demand for travel. Similarly, Jia et al. (2021) examine the GDPR's effect on EU technology venture investment using the US as their primary control group, and a combination of remaining countries as a secondary control group for robustness. In this case, the free flow of capital between countries may create spillovers to the control group. Jia et al. (2021) therefore argue that they would overestimate (underestimate) the GDPR's impact if the GDPR decreases (increases) investment outside the EU. Johnson et al. (2022c) use a "panel differences" approach in their study of website traffic. This approach is essentially a difference-in-differences strategy that uses the same websites in the previous year as a control group. By construction, this approach rules out GDPR spillovers and accounts for firm-specific seasonal differences, but requires parallel trends across years.

Several GDPR papers instead apply identification strategies that do not depend on a

control group. Some authors argue that a sudden change in an outcome after the GDPR can be attributed to the regulation. For websites' use of technology vendors, Peukert et al. (2022) use essentially an interrupted time-series design whereas Johnson et al. (2022c) use before-after differences. An interrupted time-series design (see e.g., McDowall et al., 2019) assumes that the counterfactual outcome continues its baseline (e.g., linear) time trend, as established pre-GDPR. This approach attributes both post-GDPR changes the outcome's level and trend to the regulation. Lacking pre-trend data, Johnson et al. (2022c) instead compare outcome levels after the GDPR with a pre-GDPR baseline. The authors argue that unobserved time trends confound their estimates, so that short-run differences best reflect the causal impact of the GDPR. Other authors exploit variation in the degree of exposure to the GDPR. For example, Yuan & Li (2019) compare the financial performance of EU hospitals by whether the hospitals are more or less data-dependent. Chen et al. (2022) use variation in industry-level exposure to the EU using trade data to calculate the share of output sold to EU countries. Finally, Godinho de Matos & Adjerid (2022) use a GDPR-related marketing field experiment in order to avoid the event-study analysis entirely.

## 3.2 Variable firm compliance and regulatory enforcement

The European Commission (2019) status report on the GDPR acknowledges that the regulation fell short of its potential due to a lack of enforcement. The GDPR literature has shown variation in compliance efforts by industry, by country, by compliance requirement, by firm size, and over time. As a result, economists must critically examine the lessons that can be drawn from the GDPR in the context of variable compliance and enforcement.

In general, regulatory outcomes can be thought of as the product of a strategic interaction between firms and regulators. Compliance is costly to firms, and small and medium-sized firms in particular may lack the resources to comply. In surveys, a majority of firms reported that they were not compliant with the GDPR at the enforcement deadline and that their compliance efforts were a work in progress (TrustArc, 2018). At the same time, enforcement

is costly to the regulator and country-level DPAs vary in resources (EDPB 2020). GDPR fines to date also show that country DPAs vary in their strictness and tactics. We can therefore expect a gap between the regulation as written and the reality on the ground.

GDPR enforcement and compliance are especially challenging for a number of reasons. Unlike vehicle emissions standards, for instance, GDPR compliance is multidimensional and compliance outcomes can be difficult to observe.[3] Moreover, the GDPR is complex and enumerates many compliance options (e.g., bases for data processing), which make some compliance elements subjective. Relatedly, compliance norms may arise gradually and evolve over time (see e.g., Hils et al., 2020; Lefrere et al., 2022).[4] Since personal data is pervasive, the GDPR can be considered to be a "law of the whole economy." Regulators must therefore set enforcement priorities. Privacy regulators, unlike antitrust regulators for example, lack enforcement experience and established precedent to draw upon.

The compliance literature emphasizes that regulators can ensure compliance using a combination of fines and the probability of receiving a fine (see e.g., Polinsky & Shavell, 2000). The above points may reduce the probability of receiving a fine. Perhaps to offset this, the maximum fines under the GDPR are large.

Nevertheless, the cost of strict GDPR compliance may exceed even the maximum fines in some industries. Websites and the technology vendors that support them provide plausible examples. Many websites rely on advertising to generate revenue and some research shows that ad prices double when ad impressions contain a cross-site cookie identifier for users (Johnson et al., 2020; Ravichandran & Korula, 2019). Websites may therefore resist complying on dimensions that jeopardize their revenue model.[5] However, regulators are concerned about the privacy harm of this industry's use of online identifiers and have repeatedly criticized this industry's level of compliance (AP, 2019; CNIL, 2019; ICO, 2019; DPC, 2020).

[3]Of course, the observability of compliance outcomes also poses a problem for empirical research.

[4]This poses a challenge if we treat the GDPR as an event study.

[5]Beyond limiting technology vendors, website compliance strategies include notifying users of the presence of browser cookies, offering the user some consent choice, discontinuing the use of third-party cookies (at least prior to consent), and/or blocking EU users (Johnson et al., 2022c; Lefrere et al., 2022; Skiera et al., 2022).

Regulators complain that the industry loads vendor content and cookies prior to obtaining consent and that the industry's consent practices fall short of the GDPR's opt-in standard. Nevertheless, regulators did not fine this industry until the end of 2020. Several economic studies find that websites cut the number of vendors and/or third-party cookies in May 2018, but also find that these returned to pre-GDPR levels within a few months (Lukic et al., 2021; Johnson et al., 2022c; Lefrere et al., 2022; Peukert et al., 2022). These papers wrestle with what policy lessons can be drawn as a result, and most focus on the short-run changes. For instance, Johnson et al. (2022c) argue that the post-GDPR rebound can not be attributed to the GDPR alone because of some combination of low compliance, shifting compliance norms, lack of enforcement, and the industry's exogenous growth.

Nevertheless, the GDPR did meaningfully change the compliance and enforcement environment within the EU. The GDPR and its large fines in particular caught the attention of European firms (Martin et al., 2019). Even US firms increased their attention to data privacy—as evidenced by mentions in publicly listed firms' annual reports—particularly for those firms with a presence in the EU (Boroomand et al., 2022). Beforehand, EU enforcement of some privacy laws on the books was low, so non-compliance was a viable strategy for firms (Martin et al., 2019). I emphasize this, because it again shows that economists should not assume that firms comply with the letter of the law where privacy regulation is concerned. However, the GDPR increased political pressure on Data Protection Authorities to use their new powers to increase enforcement and thereby shifted firm beliefs about the probability of penalties (Martin et al., 2019).

Variable compliance and enforcement can obfuscate the lessons that can be drawn from empirical GDPR research. What is clear is that scholars should not assume uncritically that the GDPR *as written* actually happens *in practice*. Instead, scholars should investigate the reality of the GDPR on the ground. In particular, scholars must grapple with how firms comply with the GDPR in their setting. Cost-benefit analysis can illuminate the economics of a firm's compliance decisions. Scholars should also examine regulator's public statements

and regulatory actions to understand regulator priorities in their setting. On the consumer side, scholars should not assume, for instance, that consumers make use of their new data rights under the GDPR in economically meaningful quantities (DataGrail, 2020; Presthus & Sørum, 2021).

The literature grapples with these issues in several ways. Researchers look for domains where compliance activities are stronger or at least quantifiable. Finally, scholars acknowledge the variable nature of both compliance and regulation, and the difficulties this presents for generalizing from the short- and long-run impact of the GDPR.

## 3.3 GDPR's impact on data observability

The GDPR limits personal data processing, which creates problems for empirical researchers. The GDPR may increase the cost of accessing data for researchers or prevent data access altogether (Greene et al., 2019).[6] When consent is the legal basis for collecting data, this introduces self-selection into the data. Consent-based selection is more challenging than data missingness alone, because an unknown quantity of individual data will be altogether absent from the database. These data issues pose a challenge for many applied microeconometricians who use individual-level data to deliver economic insight.

Researchers have navigated this problem with a variety of approaches. To begin, economists can still use non-personal data—like accounting or macroeconomic data—which the GDPR should not affect (Jia et al., 2021; Chen et al., 2022). Alternately, Zhao et al. (2021) use individual data from a panel of consenting consumers to study the GDPR's impact on online search behavior. Though such panels are themselves selected—e.g, presumably panelists have a lower preference for privacy—panels at least are complete.

Other researchers embrace the GDPR's impact on consent-based missingness as interesting in its own right. For instance, Aridor et al. (2020) investigate the impact of the GDPR on online user data. Aridor et al. (2020) obtain data from a marketing intermediary that

[6]Relatedly, Yom-Tov & Ofran (2022) document a shift in clinical trials out of the EU and towards countries with weaker data protections after the implementation of the GDPR.

sends offers to users on a large collection of online-travel-agency websites around the world. These travel websites share user-level, travel-related search data with the intermediary that then makes targeted offers to users based on the user's predicted purchase probability. After the GDPR, the intermediary receives less data, which Aridor et al. (2020) attribute to a segment of users that refuse consent for data sharing. Aridor et al. (2020) show that the remaining consenting users are favorably selected in that they have longer search histories. Aridor et al. (2020) attribute this to two explanations. First, privacy-sensitive users obfuscate their browsing histories (e.g., by clearing cookies), so that they appear as multiple user identifiers with short browsing histories prior to the GDPR. Second, user willingness to consent may be correlated with user's travel-website activity, for instance, because users that like the site may be more willing to both browse the site and provide consent. After the GDPR, the intermediary can no longer see or sell to non-consenting users, which hurts its revenue. Aridor et al. (2020) point out an interesting silver lining: as the consenting user data is longer and higher quality, the intermediary may have an easier time predicting user behavior and making successful offers to consenting users.

Goldberg et al. (2022) work with similar data from a large number of websites globally from Adobe Analytics. Websites use Adobe Analytics to measure outcomes like site visits, pageviews, and ecommerce revenue. Goldberg et al. (2022) show that these outcomes—as recorded by Adobe—fell by about 12% after the GDPR. As in Aridor et al. (2020), Adobe may see less data because of non-consenting users after the GDPR. However, Adobe would also record less site data if the GDPR actually hurt the real outcomes for these sites. Goldberg et al. (2022) grapple with this identification problem by constructing bounds on the relative contributions of the consent and real effects of the GDPR to the drop in recorded site outcomes.

# 4 Literature review

In this section, I review the economics literature on the GDPR. Section 4.1 examines the GDPR from the perspective of consumers. Section 4.2 turns to the GDPR's impact on firms. This literature is larger, so we first consider the GDPR's impact on firms' economic performance measures before diving deeper into the GDPR's impact on competition, innovation, the web, and marketing. Section 4.3 discusses the lessons learned about the GDPR's constituent parts and how they work in practice.

At the outset, I point out that the GDPR literature is still maturing. Five years after the enforcement deadline, a small number of economics papers have appeared in print. As such, many of the papers I discuss below are working papers, and will therefore continue to evolve in the future.

## 4.1 Impact on consumers

The economics literature has explored the GDPR's consumer impact, though privacy economists generally find that consumer privacy preferences are difficult to ascertain (e.g., Athey et al., 2017). One approach is to survey consumers and ideally to do so prior to the GDPR for a baseline comparison. For instance, Presthus & Sørum (2021) surveyed a cross-section of Norwegian university students annually from 2018 to 2020. However, this evidence failed to show the GDPR's expected improvements: the surveys show no increase in general awareness of privacy or perceived control over personal data.

Sobolewski & Paliński (2017) implement a stated preference discrete choice experiment prior to the GDPR. By surveying Polish university students, Sobolewski & Paliński (2017) obtain willingness-to-pay estimates for four individual data rights under the GDPR.[7] This study reveals a similar average willingness to pay for the right to be forgotten, the right to object to profiling, and the GDPR's extended information obligations. However, the willingness to pay for data portability was negative and statistically insignificant. The authors

[7]See Presthus & Sørum (2019) for related survey evidence.

provide an estimate of the welfare benefit of the GDPR by summing consumer willingness to pay for these four rights. Sobolewski & Paliński (2017) thus estimate that the GDPR provides a value of €6.50 per person per month.

Other economic papers speak to the consumer welfare impact of the GDPR or show objective improvements in consumer privacy. Janssen et al. (2022) argue that the GDPR hurts consumer surplus by reducing innovation in consumer products. To show this, they use a structural demand model for apps to examine the consumer consequences of the GDPR to the app market. In theoretical work, Ke & Sudhir (2022) and Wang et al. (2022) investigate the welfare consequences of the GDPR for both firms and consumers.

The GDPR should improve consumer privacy by improving data security and reducing data processing. These objective improvements in privacy may be difficult to quantify across firms and at large scale. Nevertheless, Demirer et al. (2022) show that EU firms reduce both their data storage and computation activity on Microsoft's cloud service after the GDPR. As we will see in Section 4.2.3, several researchers find that websites reduced data sharing after the GDPR, though these privacy improvements were short-lived. A small segment of consumers appears to be exercising their consent privilege by opting out of data collection online (Aridor et al., 2020; Goldberg et al., 2022).

By drawing attention to data protection, the GDPR may have influenced how firms measure and report their data-protection activities. For example, the GDPR's data-breach notification requirement should have reduced the number of data breaches. This would be challenging to evaluate empirically, however, as the notification requirement should also increase the number of breaches that firms both notice and report. Similarly, the GDPR's encryption requirement should reduce the privacy risk from data breaches; nonetheless, Miller & Tucker (2011) show that (public) data-breach incidents actually increased after the American medical sector adopted data encryption.

## 4.2 Impact on firms

Several scholars document that the GDPR harmed a variety of firms' outcomes including: profits, revenue, investment, market exit, and entry. I first discuss this evidence for firm performance before turning to the GDPR's impact on competition (Section 4.2.1), innovation (Section 4.2.2), the web (Section 4.2.3), and marketing (Section 4.2.4).

Multiple studies examine accounting data and attribute a reduction in firms' profit and/or revenue to the GDPR. For instance, Koski & Valmari (2020) examine nearly 267,000 EU and US firms from 2014 to 2018. The authors use difference-in-differences with US firms as a control and 2018 as the treatment year. Koski & Valmari (2020) find a statistically insignificant effect on profit margins in their full sample, but a statistically significant -1.9% reduction in profit margins among data-intensive sectors in the EU (i.e., information and communications sector, banks, and other financial services). Chen et al. (2022) examine almost 700,000 firms across 61 countries and 34 industries. By comparing firms by their sector's revenue exposure to the EU, they attribute a decline in profits and a reduction in sales by the firm's degree of GDPR exposure. Yuan & Li (2019) use difference-in-differences to compare the financial performance of hospitals in the EU by the importance of information, communication, and telecommunication to their business. They find lower operating revenue (scaled by total assets) for more data-intensive hospitals during the GDPR's transition period from passage to enforcement (2016-2018).

Survey evidence finds that firms incurred significant costs in order to comply with the GDPR. The International Association of Privacy Professionals (IAPP 2017) estimated that Fortune 500 global firms would spend $7.8 billion on compliance.[8] DataGrail (2020) find that 74% of small- and mid-sized organizations spent more than $100,000 on compliance.

Jia et al. (2021) show that the GDPR reduced investment for EU technology ventures.[9] Using the difference-in-differences strategy described in Section 3.1, they find that the num-

[8]The IAPP figure extrapolates from survey evidence in the IAPP and Ernst & Young 2017 report.

[9]Note that Lambrecht (2017) also finds a reduction in venture investment in certain sectors after the EU's e-Privacy Directive.

ber of EU venture deals fell by 26% after the GDPR enforcement deadline. Jia et al. (2021) also document that the most affected firms are: early-stage ventures, data-related ventures, business-to-consumer (versus business-to-business) ventures, and ventures in the healthcare and finance industries. These patterns are consistent with a GDPR effect as we may expect the GDPR to have greater effects for ventures that use data, especially consumer data, health data (i.e., special category data), and in heavily-regulated industries. Jia et al. (2020) build on this research by examining differences between EU and foreign investors. Jia et al. (2020) find an increase in investor home bias post-enforcement: that is, foreign investment in EU technology ventures falls by more than local investment. Jia et al. (2020) argue that this is consistent with foreign investors having greater uncertainty about the financial consequences of the GDPR.

Finally, Janssen et al. (2022) show both an increase in market exit and and a decrease in entry for mobile apps on the Android platform after the enforcement deadline. Janssen et al. (2022) examine app data from the Google Play Store using a before-after comparison and supplement their findings by surveying German app developers. Kircher & Foerderer (2021) document a small increase in closures of US app start-ups post-GDPR as well as a small reduction in venture capital transactions for US app startups (relative to US enterprise software startups).

### 4.2.1 Impact on competition

Several observers warned of a potential tradeoff between privacy regulation and competition (e.g., Brill, 2011; Goldfarb & Tucker, 2012; Phillips, 2019). Indeed, the GDPR literature repeatedly confirms this hypothesis. In general, regulation can have effects on competition if firms experience returns to scale in compliance. For privacy regulation, consent requirements may also favor large established firms if consumers are more likely to provide consent to such firms (Campbell et al., 2015) or to consent to smaller lists of third-party data recipients. Gal & Aviv (2020) and Geradin et al. (2020) discuss several potential channels through which

the GDPR may affect competition.

Several researchers find that the GDPR disproportionately hurts smaller firms (e.g. Bessen et al., 2020; Jia et al., 2020; Koski & Valmari, 2020; Zhao et al., 2021; Chen et al., 2022). Goldberg et al. (2022) provide indirect evidence that smaller websites obtain lower consent rates. Both Johnson et al. (2022c) and Peukert et al. (2022) find that the market for technology vendors that serve websites became more concentrated right after the GDPR's enforcement deadline. This provides evidence for a new anticompetitive mechanism: when privacy regulation restricts business-to-business data transfers, firms may prefer to retain their larger vendors.[10]

### 4.2.2 Impact on innovation

Goldfarb & Tucker (2012) argue that a trade-off exists between privacy and innovation. They support their argument with numerous studies focusing on the online-advertising and health-care sectors. Supported by interviews of startups and lawyers in 2018, Martin et al. (2019) point out that the GDPR can both support and suppress innovation. For instance, the interviews suggested that the GDPR spurred privacy-related innovation as well as increased demand for "regulation-exploiting innovation:" e.g., diffusing compliance management software and encryption capabilities. However, Martin et al. (2019) also reveal claims that the GDPR led startups to abandon products, discouraged entrepreneurs, and limited innovator's access to input data (e.g., for artificial-intelligence applications).

The empirical evidence for the GDPR's impact on innovation is somewhat mixed. As we have seen, the GDPR reduced technology venture funding (Jia et al., 2021) and hurt the mobile app market (Kircher & Foerderer, 2021; Janssen et al., 2022). Bessen et al. (2020) survey artificial intelligence startups. Bessen et al. (2020) find that GDPR imposes costs on these firms in terms of adding new position(s), reallocating resources, and deleting

[10]Contrary to Campbell et al. (2015), Johnson et al. (2022c) find no evidence that consent drives increased concentration. However, the simple explanation is that sites rarely make the list of third-party data firms prominent when requesting consent.

data. Despite the GDPR's requirements on firms, Bessen et al. (2020) find that the use of various data protection means does not differ by whether the firm has customers in Europe. Venkatesan et al. (2022) provide evidence that the GDPR increased the return on assets from acquisitions of AI technology companies—particularly for acquisitions related to customer experience and cyber security. Perhaps counter to expectations, Chen et al. (2022) find that patenting among IT service firms increased 30%, though this figure is imprecisely estimated.

Blind et al. (2022) examine innovation using an annual survey of German firms from 2011 to 2020. Examining the 2018 survey, Blind et al. (2022) note that 35.0% of firms report that data protection regulation hampers their innovation activities whereas only 4.7% report the opposite. Perhaps in contrast with other GDPR research, the share of firms that report either an innovation-facilitating or innovation-complicating role seems to increase with firm size. Blind et al. (2022) also find evidence that the GDPR shifts innovation to become more incremental and less radical in nature.

#### 4.2.3 Impact on the web

The web uses personal data to personalize websites, content, and advertising. At a basic level, the Internet requires IP addresses—which the GDPR considers to be personal data—to function. For researchers, the internet and websites therefore provide an opportunity to study an industry that is both targeted by the regulation and provides data for empirical study.

Researchers have examined the GDPR's impact on site traffic, site vendor use, site content creation, internet infrastructure, and online search. Several researchers find that the GDPR reduced sites' use of vendors and/or data sharing using third-party cookies (Lukic et al., 2021; Johnson et al., 2022c; Lefrere et al., 2022; Peukert et al., 2022). Several computer science researchers concur with these findings (e.g. Libert et al., 2018; Urban et al., 2020). Despite this, or perhaps due to the rapid post-GDPR bounce-back, Lefrere et al. (2022) find no impact on news and media websites' production of new content or social sharing of that

content. Using data from Adobe Analytics, Goldberg et al. (2022) argue that real website pageviews and ecommerce revenue from EU users falls by at least about 0.5% post-GDPR. Using third-party site-traffic data, Schmitt et al. (2021) find a larger (5-10%) reduction in site visits, Congiu et al. (2022) find an even larger (15%) reduction in 2019, but Lefrere et al. (2022) find that EU site traffic measures are relatively stable except for a small decline in pageviews per user.[11] Finally, Zhao et al. (2021) examine the browsing behavior of a panel of online users. Zhao et al. (2021) find that EU users increase their online search intensity after the GDPR relative to their non-EU counterparts.

The GDPR limits international data transfers—particularly to the majority of countries that do not meet the EU's adequacy requirements. As such, we might expect that the GDPR affected data flows between the EU and the rest of the world. Zhuo et al. (2021) investigate this possibility by obtaining data at the Internet's infrastructure level to monitor physical investments in international data flows. However, Zhuo et al. (2021) find no GDPR effect in the EU on the Internet's interconnectivity layer. This finding is further notable because it arises despite the reductions—albeit modest—in site traffic and vendor use documented above. Though the authors lack more granular data on the type of data flows, the authors suggest that growth in, for instance, data-heavy video traffic may mask the observed reduction in other web-related data flows.

#### 4.2.4 Impact on marketing

The GDPR was expected to increase the difficulty of matching consumers and firms using marketing. In particular, the GDPR's data processing restrictions were expected to hurt personalized marketing channels like email and online display advertising. Consistent with this, Goldberg et al. (2022) find larger reductions in recorded EU site traffic originating from

[11] These authors assume that their data fully captures real site outcomes (i.e., the ground truth). Nevertheless, it is unclear how their data sources—SimilarWeb (Schmitt et al., 2021; Congiu et al., 2022) and Alexa web information services (Lefrere et al., 2022)—address traffic from non-consenting users (see Section 3.3). In particular, SimilarWeb explains that it somehow models traffic using a variety of data sources, which include site analytics data (which must exclude non-consenting users) shared by websites as well as a panel of browser extension users.

email or display ad clicks relative to visits that directly navigate to the website.[12] Aridor et al. (2020) highlight that the GDPR can limit personalized marketing opportunities, but favorably-selected data from consenting users can improve the firm's individual marketing response predictions.

Godinho de Matos & Adjerid (2022) and D'Assergio et al. (2022) examine email permissioning campaigns. Godinho de Matos & Adjerid (2022) run a marketing field experiment with a large European telecommunications firm. Many marketers sought to bring their marketing consent up to the GDPR standard by running a permissioning campaign to (re-)obtain consent. In this case, the firm ran an experiment that sent out a permissioning email in the treatment group, and sent that email after a delay in the control group. Godinho de Matos & Adjerid (2022) show that the permissioning campaign succeeded at increasing the share of consumers to which the firm can market. Moreover, Godinho de Matos & Adjerid (2022) show that the firm was able to subsequently both increase the marketing messages it sent to treated consumers and increase revenue from these consumers.

D'Assergio et al. (2022) collect and categorize 1,506 different permissioning emails. They find that 29% of these tried to persuade users (e.g., with discount offers or discussing benefits of data sharing), 35% only used an informative approach, and 20% combined both approaches. D'Assergio et al. (2022) too partner with an European firm to run an email field experiment. The authors find evidence that persuasive tactics can improve opt-in rates and that combining this with informative tactics can further improve opt-in rates. However, the authors find no significant differences in the amount of personal data shared across conditions.

## 4.3 Elements of the GDPR in practice

One challenge in studying the GDPR is that the regulation contains so many elements. Since these elements were all applied at once, the event-study nature of most GDPR research limits

[12] On this basis, Goldberg et al. (2022) find the aforementioned lower bound of about a 0.5% reduction in real site traffic and ecommerce revenue, which the authors attribute to degraded marketing ability.

how much can be learned about the GDPR's constituent parts. Nevertheless, unpacking these elements is useful for evaluating the regulation and designing effective privacy regulation. Several researchers have shown patterns that appear to reveal some consequences of the GDPR's design decisions and features of the regulation in practice.

The GDPR intended to harmonize data regulation within the European Union, and this was thought to be a source of efficiencies for firms that serve multiple EU countries (European Commission, 2012). However, we have seen that regulators vary in their resources and enforcement strategies. Several authors have found that the size of the GDPR's impact is correlated with firms' beliefs about data-protection regulatory strictness at the country level (Jia et al., 2020, 2021; Goldberg et al., 2022; Johnson et al., 2022c). To establish this, these studies use a European Commission (2008) survey of data processors by EU country that asked whether their local data protection regulator was more or less strict than regulators in the rest of the EU. By this metric, the strictest data regulators are Germany and Sweden, and the laxest regulators are Bulgaria and Greece. Though this regulatory strictness measure is dated, it appears to predict the depth of the GDPR's impact.[13]

Other research examines international spillovers from the GDPR. Peukert et al. (2022) highlight the spillovers to non-EU residents using website data collected from the vantage point of a US user. Non-EU residents see the largest vendor reductions on websites located in the EU that serve primarily an EU audience. This suggests that EU-focused firms roll their compliance efforts to all their consumers, which benefits their (limited) foreign audience. Non-EU websites cut their vendors vis-à-vis US users, though very little for sites that primarily serve a non-EU audience. Johnson et al. (2022c) instead scan websites from the perspective of a French user, using a VPN service. They find that—from the perspective of an EU user—foreign sites with a small share of EU users make deeper cuts to their vendors than sites that primarily serve EU users. Johnson et al. (2022b) attribute this pattern of results to the design of the GDPR fines, which reach 4% of a firm's *global* revenue. In

[13]Though country-level strictness is correlated with per capita income, these papers show that the strictness result is robust to including income as a model covariate.

particular, the benefit of exploiting user data is relatively small for sites with a small share of EU users, but otherwise equivalent sites would face the same fine. Perhaps due to these differing incentives, Johnson et al. (2022b) remark that EU firms here do less to protect EU residents than non-EU firms. Note that Lefrere et al. (2022) complement the above studies by scanning 909 news and media publisher websites from the vantage point of both EU and US users. Lefrere et al. (2022) largely confirm the above results using third-party cookies as their dependent variable.

Sørum & Presthus (2020) examine the GDPR's data access and portability rights by initiating personal data access requests from 15 firms. They find that almost all these firms responded quickly and provided personal data, though the data provided fell short of the letter of the law (i.e., all eight items regarding data access under Article 15).

Finally, several researchers show that firms that rely more on consumer data and sensitive data exhibit greater harms from the GDPR (e.g. Jia et al., 2021; Li et al., 2019). This may oversimplify the picture for certain industries though, as established firms with experience handling sensitive data may instead have lower adjustment costs. Koski & Valmari (2020) discuss this lower adjustment cost as a potential explanation for their findings.

# 5 Future opportunities for research

The GDPR is an important and relatively recent regulation. We will undoubtedly see more related research in the future. In the conclusion, I suggest some directions for future research. Below, I suggest two key opportunities for privacy research. Section 5.1 enumerates recent and future privacy-related changes to regulation and technology platforms. Section 5.2 introduces privacy-enhancing technologies and discusses opportunities for economists to improve these technologies and study their adoption.

## 5.1 More privacy regulations & changes on the horizon

Though the GDPR received most of the literature's attention in recent years, several other regulations and interventions have since passed or are on the horizon. Nevertheless, compliance and enforcement issues (Section 3.2) loom large here: the realized privacy results will vary.

First of all, the GDPR remains a worthwhile subject of research. Future research may extend beyond the GDPR's enforcement deadline. Given the GDPR's compliance and enforcement issues, future crackdowns may present opportunities to study the impact of the GDPR. For instance, potential "mini" GDPR events include regulator enforcement deadlines, major regulatory actions, major court decisions, voluntary changes in compliance strategies (e.g. self-regulatory changes), and private actions (e.g., noyb, 2022). For example, Johnson et al. (2022c) examine the French regulator's enforcement deadline for websites (April 2021) as well as a self-regulatory update to the web vendor industry's consent mechanism (Fall 2020).[14] Also, the United Kingdom is considering whether to revisit the GDPR in light of that country's exit from the EU. This may provide opportunities to study the impact of undoing certain elements of the GDPR.

Second, proposed and enacted regulations worldwide provide additional opportunities for research. Many countries have passed, enforced, and/or updated privacy regulation since the GDPR was passed including: Bahrain, Brazil, Burkina Faso, China, India, Israel, Japan, Kenya, Mauritius, New Zealand, Nigeria, Qatar, Singapore, South Africa, South Korea, Switzerland, Thailand, Turkey, and Uganda. The EU passed the Digital Services Act and Digital Markets Act in 2022, which contain relevant provisions. For instance, the Digital Services Act largely bans targeted online ads to children under 18. The EU's proposed ePrivacy Regulation will build on the GDPR by establishing particular privacy regulations for electronic communication in the EU. The ePrivacy Regulation will build on its predecessor—

[14]These results (in an online appendix) show that these GDPR-like events replicated the key findings of Johnson et al. (2022c): the GDPR simultaneously reduces vendor use and increases vendor market concentration.

the ePrivacy Directive—which Goldfarb & Tucker (2011) and Lambrecht (2017) study. In the US, Congress has considered several privacy laws while five states have passed privacy laws: California, Colorado, Connecticut, Virginia, and Utah. For instance, Abis et al. (2022) study the California Consumer Privacy Act and its impact on voice-AI firms. Also, the Federal Trade Commission (FTC) has telegraphed its desire to more aggressively protect consumer privacy with its 2022 Advanced Notice of Proposed Rule-making on "Commercial Surveillance and Data Security."

Third, some large technology firms responded to increased privacy-related regulatory scrutiny by instituting related changes on their platforms. These changes can mitigate non-compliance issues by instituting platform rule changes that, to a greater extent, force firms on their platform to comply. For instance, Apple's App Tracking Transparency forced apps to request user opt-in consent for what Apple terms "tracking" as of April 2021. Some research examines the resulting consequences for apps and advertisers on Apple's platform (Kesler, 2022; Li & Tsai, 2022). As another example, YouTube eliminated all forms of personalization for child-directed content globally in January 2020 in reaction to its record US Children's Online Privacy Protection Act fine (Johnson et al., 2022a).

## 5.2 Privacy-enhancing technologies

Privacy-enhancing technologies (PETs) offer a potential solution for the tension between privacy and the data economy.. The United Kingdom defines PETs as "technologies that embody fundamental data protection principles by minimizing personal data use, maximizing data security, and empowering individuals" (ICO 2022). Examples of PETs include: differential privacy, federated learning, on-device computation, zero-knowledge proof, and secure multi-party computation.

PETs are now gaining practical use. For instance, the US Census will add noise to its data before computing its public statistics (i.e., differential privacy) in order to fulfill its legal obligation to not reveal information about individuals in the census. Some have argued

that PETs can aid in GDPR compliance efforts (e.g., Cummings & Desai, 2018). As well, Google's Privacy Sandbox proposes PETs as alternatives to browser cookies and mobile ad identifiers (Google, 2022). Still, privacy regulations and proposed regulations have seemingly ignored these developments to date. For instance, the FTC's request for public comment on "Commercial Surveillance and Data Security" only mentions PETs in passing.[15]

Economists can contribute to research on PETs. More fundamental research is required on how to design PETs. Economists in particular can help map out the privacy versus value-creation frontier of PETs (e.g., Hotz et al., 2022). PETs like differential privacy also create challenges for inference by transforming datasets (Komarova & Nekipelov, 2020). Economists can study the adoption and consequences of PETs just as they study other innovations like artificial intelligence and cloud computing (e.g., Zolas et al., 2021). PETs too can have competitive consequences, for instance, because smaller quantities of data are more likely to reveal an individual's data. In the case of online advertising, Johnson et al. (2022b) predict significant consequences of applying PETs for both practitioners and researchers.

# 6 Conclusion

The GDPR represents an opportunity for economists to understand the consequences of an economy-wide privacy regulation. However, the GDPR poses several challenges for economic research. First, the GDPR made a global impact as it covers both EU firms and non-EU firms that target EU residents. The GDPR also created substantial global spillovers, so researchers may struggle to find a suitable control group that both is excluded from the regulation and is comparable to the EU. Second, the variability of firm compliance and regulatory enforcement under the GDPR complicates the generalizations that we can draw from the literature. Third, the GDPR sought to limit personal data processing and to allow privacy-sensitive consumers to opt out of data processing. This, in turn, can limit empirical

[15]The request for comment contains 95 questions. The final question asks about the "potential obsolescence of any rulemaking" and references the privacy-related innovations in the online ad industry.

researchers' access to data and can introduce consent-based self-selection into the observed data.

The economic literature on the GDPR examines multiple facets of the regulation and its impact. The GDPR presented a novel opportunity for economists to empirically investigate long-held hypotheses like the consequences of privacy regulation for competition and innovation. Most GDPR research points to the GDPR hurting firm outcomes and disproportionately harming smaller and more data-dependent firms. For consumers, the literature illuminates objective improvements in privacy and surveys consumers for their views on the GDPR. The literature also explores the consequences of the GDPR's design elements including its international spillovers.

Looking back at the GDPR literature, one potential criticism is that the literature has documented the *unintended* consequences, but perhaps neglected the *intended* consequences of the GDPR. In particular, we want to better understand the privacy benefits to consumers and rigorously quantify these benefits. As well, we want to better understand and quantify the gains in data protection. To be fair, these are difficult subjects to evaluate convincingly and with the data at hand.

Privacy regulation and the GDPR offer several more directions for research. First, Section 2 lists many elements of the GDPR that have received little attention so far. Second, more attention should be paid to understanding the strategic interactions between firms and regulators. We would like to better understand which enforcement strategies—e.g., fines, notices, choice of targets, establishing legal precedent—are effective in ensuring compliance. Third, the GDPR literature has so far neglected the GDPR's anticipated impact assessments like those of the European Commission (2012) as well as industry-funded studies like Christensen et al. (2013) and Deloitte (2013). These predictions identify lingering questions like the GDPR's impact on employment. Finally, we wish to better understand how to design effective privacy regulation and improve upon existing regulation like the GDPR. Existing work has highlighted many unintended consequences of the GDPR, and continued research

can explore how to limit these unintended harms.

Policy-makers and regulators around the globe continue to wrestle with how to regulate privacy effectively in the modern data economy. Continued research can illuminate their task. As the GDPR continues to evolve in practice, this will present more opportunities to study the law. New privacy laws worldwide also represent opportunities for research. Recent breakthroughs in commercializing privacy-enhancing technologies promise to limit certain tradeoffs between privacy and the data economy. More research is needed to understand the novel trade-offs that these technologies present as well as the economic consequences of adopting these technologies.

# References


Abis, S., Canayaz, M., Kantorovitch, I., Mihet, R., & Tang, H. (2022). *Privacy Laws and Value of Personal Data*. Technical report, EPFL.

Acquisti, A., Taylor, C., & Wagman, L. (2016). The economics of privacy. *Journal of Economic Literature*, 54(2), 442–92.

Aridor, G., Che, Y.-K., Nelson, W., & Salz, T. (2020). The economic consequences of data privacy regulation: Empirical evidence from GDPR. *Available at SSRN*.

Athey, S., Catalini, C., & Tucker, C. (2017). The digital privacy paradox: Small money, small costs, small talk. NBER working paper.

Autoriteit Persoonsgegevens (2019). AP: veel websites vragen op onjuiste wijze toestemming voor plaatsen tracking cookies.

Bessen, J. E., Impink, S. M., Reichensperger, L., & Seamans, R. (2020). GDPR and the importance of data to AI startups. *Available at SSRN 3576714*.

Blind, K., Niebel, C., & Rammer, C. (2022). The impact of the EU General Data Protection Regulation on innovation in firms. ZEW Discussion Paper.

Boroomand, F., Leiponen, A., & Vasudeva, G. (2022). Does the market value attention to data privacy? evidence from US-listed firms under the GDPR. Wharton Mack Institute working paper.

Bradford, A. (2020). *The Brussels effect: How the European Union rules the world*. Oxford University Press, USA.

Brill, J. (2011). The intersection of consumer protection and competition in the new world of privacy. *Competition Policy International*, 7, 7–313.

Cameron, A. C. & Trivedi, P. K. (2005). *Microeconometrics: methods and applications*. Cambridge university press.

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

Carrière-Swallow, Y. & Haksar, V. (2019). *The Economics and Implications of Data An Integrated Perspective*. Technical report, International Monetary Fund.

Chen, C., Frey, C. B., & Presidente, G. (2022). Privacy regulation and firm performance: Estimating the GDPR effect globally. The Oxford Martin Working Paper Series on Technological and Economic Change.

Christensen, L., Colciago, A., Etro, F., & Rafert, G. (2013). The impact of the data protection regulation in the EU. *Intertic Policy Paper, Intertic*.

Commission Nationale de l'Informatique et des Libertés (2019). Online targeted advertisement: what action plan for the CNIL? https://www.cnil.fr/en/online-targeted-advertisement-what-action-plan-cnil.

Congiu, R., Sabatino, L., & Sapi, G. (2022). The impact of privacy regulation on web traffic: Evidence from the GDPR. *Available at SSRN 4025033*.

Cummings, R. & Desai, D. (2018). The role of differential privacy in GDPR compliance. In *FAT'18: Proceedings of the Conference on Fairness, Accountability, and Transparency*.

D'Assergio, C., Manchanda, P., Montaguti, E., & Valentini, S. (2022). The race for data: Gaming or being gamed by the system? *Available at SSRN 4250389*.

Data Protection Commission (2020). *Report by the Data Protection Commission on the use of cookies and other tracking technologies*. Technical report, Data Protection Commission.

DataGrail (2020). *The Age of Privacy: The Cost of Continuous Compliance*. Technical report.

Deloitte (2013). *Economic impact assessment of the proposed General Data Protection Regulation*. Technical report, December.

Demirer, M., Hernández, D. J., Li, D., & Peng, S. (2022). Data, privacy laws, and firm production: Evidence from GDPR. Work in progress.

European Commission (2008). Flash eurobarometer 226: Data protection in the european union : Data controllers' perceptions. https://data.europa.eu.

European Commission (2012). *COMMISSION STAFF WORKING PAPER Impact Assessment Accompanying the document Regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation) and Directive of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data by competent authorities for the purposes of prevention, investigation, detection or prosecution of criminal offences or the execution of criminal penalties, and the free movement of such data*. Technical report, European Commission.

European Commission (2019). *Data protection rules as a trust-enabler in the EU and beyond – taking stock*. Communication from the Commission to the European Parliament and the Council, European Commission.

European Data Protection Board (2020). *Contribution of the EDPB to the evaluation of the GDPR under Article 97*. Technical report, European Data Protection Board.

Gal, M. S. & Aviv, O. (2020). The Competitive Effects of the GDPR. *Journal of Competition Law & Economics*. nhaa012.

Geradin, D., Karanikioti, T., & Katsifis, D. (2020). GDPR myopia: how a well-intended regulation ended up favouring large online platforms - the case of ad tech. *European Competition Journal*, (pp. 1–46).

Godinho de Matos, M. & Adjerid, I. (2022). Consumer consent and firm targeting after GDPR: The case of a large telecom provider. *Management Science*, 68(5), 3330–3378.

Goldberg, S., Johnson, G., & Shriver, S. (2022). Regulating privacy online: An economic evaluation of the GDPR. Available at SSRN 3421731.

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

Goldfarb, A. & Tucker, C. (2012). Privacy and innovation. *Innovation policy and the economy*, 12(1), 65–90.

Google (2022). The Privacy Sandbox: technology for a more private web. https://privacysandbox.com.

Greene, T., Shmueli, G., Ray, S., & Fell, J. (2019). Adjusting to the GDPR: The impact on data scientists and behavioral researchers. *Big Data*, 7(3), 140–162. PMID: 31033336.

Hils, M., Woods, D. W., & Böhme, R. (2020). Measuring the emergence of consent management on the web. In *Proceedings of the ACM Internet Measurement Conference*, IMC '20 (pp. 317–332). New York, NY, USA: Association for Computing Machinery.

Hotz, V. J., Bollinger, C. R., Komarova, T., Manski, C. F., Moffitt, R. A., Nekipelov, D., Sojourner, A., & Spencer, B. D. (2022). Balancing data privacy and usability in the federal statistical system. *Proceedings of the National Academy of Sciences*, 119(31), e2104906119.

Information Commissioner's Office (2019). *Update report into adtech and real time bidding*. Technical report.

Information Commissioner's Office (2021). *Guide to the General Data Protection Regulation (GDPR)*. Technical report, Information Commissioner's Office.

Information Commissioner's Office (2022). *Draft anonymisation, pseudonymisation and privacy enhancing technologies guidance*, chapter Privacy-enhancing technologies (PETs). Information Commissioner's Office.

International Association of Privacy Professionals (2017). Global 500 companies to spend $7.8b on GDPR compliance. https://iapp.org.

International Association of Privacy Professionals & Ernst & Young (2017). *IAPP-EY Annual Privacy Governance Report 2017*. Technical report.

Janssen, R., Kesler, R., Kummer, M. E., & Waldfogel, J. (2022). *GDPR and the lost generation of innovative apps*. Technical report, National Bureau of Economic Research.

Jia, J., Jin, G. Z., & Wagman, L. (2020). GDPR and the localness of venture investment. *Available at SSRN 3436535*.

Jia, J., Jin, G. Z., & Wagman, L. (2021). The short-run effects of the General Data Protection Regulation on technology venture investment. *Marketing Science*, 40(4), 661–684.

Johnson, G., Lin, T., & Cooper, J. (2022a). COPPAcalypse? The YouTube settlement's impact on kids content creation. University of Pennsylvania Economics of Digital Service Blog.

Johnson, G., Runge, J., & Seufert, E. (2022b). Privacy-centric digital advertising: Implications for research. *Customer Needs and Solutions*, 9(1), 49–54.

Johnson, G., Shriver, S., & Goldberg, S. (2022c). Privacy & market concentration: Intended & unintended consequences of the GDPR. *Available at SSRN 3477686*.

Johnson, G. A., Shriver, S. K., & Du, S. (2020). Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Science*, 39(1), 33–51.

Jones, M. L. & Kaminski, M. E. (2020). An American's guide to the GDPR. *Denver Law Review*, 98(1).

Ke, T. T. & Sudhir, K. (2022). Privacy rights and data security: GDPR and personal data driven markets. *Available at SSRN 3643979*.

Kesler, R. (2022). The impact of apple's app tracking transparency on app monetization. *Available at SSRN 4090786*.

Kircher, T. & Foerderer, J. (2021). Does EU-consumer privacy harm financing of US-app-startups? Within-US evidence of cross-EU-effects. In *Proceedings of the 42nd International Conference on Information Systems (ICIS)* (pp. 12–15).

Komarova, T. & Nekipelov, D. (2020). Identification and formal privacy guarantees. *arXiv preprint arXiv:2006.14732*.

Koski, H. & Valmari, N. (2020). Short-term impacts of the GDPR on firm performance. ETLA Working Papers.

Lambrecht, A. (2017). *E-privacy provisions and venture capital investments in the EU*. Technical report, working paper.

Lefrere, V., Warberg, L., Cheyre, C., Marotta, V., & Acquisti, A. (2022). The impact of the GDPR on content providers: A longitudinal analysis.

Li, D. & Tsai, H.-T. (2022). Mobile apps and targeted advertising: Competitive effects of data exchange. *Available at SSRN 4088166*.

Li, H., Yu, L., & He, W. (2019). The impact of GDPR on global technology development. *Journal of Global Information Technology Management*, 22(1), 1–6.

Libert, T., Graves, L., & Nielsen, R. K. (2018). Changes in third-party content on European news websites after GDPR.

Lukic, K., Miller, K. M., & Skiera, B. (2021). The impact of the General Data Protection Regulation (GDPR) on the amount of online tracking. Working paper.

Martin, N., Matt, C., Niebel, C., & Blind, K. (2019). How data protection regulation affects startup innovation. *Information Systems Frontiers*, 21(6), 1307–1324.

McDowall, D., McCleary, R., & Bartos, B. (2019). *Interrupted Time Series Analysis*. Oxford University Press.

Miller, A. R. (2022). Privacy of digital health information. In A. Goldfarb & C. Tucker (Eds.), *Economics of Privacy* chapter 3. University of Chicago Press.

Miller, A. R. & Tucker, C. (2009). Privacy protection and technology diffusion: The case of electronic medical records. *Management Science*, 55(7), 1077–1093.

Miller, A. R. & Tucker, C. E. (2011). Encryption and the loss of patient data. *Journal of Policy Analysis and Management*, 30(3), 534–556.

noyb (2022). noyb aims to end "cookie banner terror" and issues more than 500 GDPR complaints. https://noyb.eu/en/noyb-aims-end-cookie-banner-terror-and-issues-more-500-gdpr-complaints.

Peukert, C., Bechtold, S., Batikas, M., & Kretschmer, T. (2022). Regulatory spillovers and data governance: Evidence from the GDPR. *Marketing Science*, 41(4), 318–340.

Phillips, N. (2019). Keep it: Maintaining competition in the privacy debate. Remarks for Internet Governance Forum.

Polinsky, A. M. & Shavell, S. (2000). The economic theory of public enforcement of law. *Journal of Economic Literature*, 38(1), 45–76.

Prasad, A. & Perez, D. R. (2020). The effects of GDPR on the digital economy: Evidence from the literature. *Informatization Policy*, 27(3), 3–18.

Presthus, W. & Sønslien, K. F. (2021). An analysis of violations and sanctions following the GDPR. *International Journal of Information Systems and Project Management*, 9(1), 38–53.

Presthus, W. & Sørum, H. (2019). Consumer perspectives on information privacy following the implementation of the GDPR. *International Journal of Information Systems and Project Management*, 7(3), 19–34.

Presthus, W. & Sørum, H. (2021). A three-year study of the GDPR and the consumer. In *14th IADIS International Conference Information Systems 2021*.

Ravichandran, D. & Korula, N. (2019). *Effect of disabling third-party cookies on publisher revenue*. Technical report, Google Inc.

Schmitt, J., Miller, K. M., & Skiera, B. (2021). The impact of privacy laws on online user behavior. *arXiv preprint arXiv:2101.11366*.

Skiera, B., Miller, K., Jin, Y., Kraft, L., Laub, R., & Schmitt, J. (2022). *The impact of the General Data Protection Regulation (GDPR) on the online advertising market*.

Sobolewski, M. & Paliński, M. (2017). How much consumers value on-line privacy? welfare assessment of new data protection regulation (GDPR). University of Warsaw Faculty of Economics Sciences Working Paper.

Sørum, H. & Presthus, W. (2020). Dude, where's my data? the GDPR in practice, from a consumer's point of view. *Information Technology & People*, 34(3), 912–929.

TrustArc (2018). *GDPR Compliance Status: A Comparison of US, UK and EU Companies*. Technical report, TrustArc.

Urban, T., Tatang, D., Degeling, M., Holz, T., & Pohlmann, N. (2020). Measuring the impact of the GDPR on data sharing in ad networks. In *Proceedings of the 15th ACM Asia Conference on Computer and Communications Security (ASIA CCS '20)*: ACM.

Venkatesan, R., Arunachalam, S., & Pedada, K. (2022). Short run effects of generalized data protection act on returns from AI acquisitions.

Wang, X., Xu, F., & Zhang, F. (2022). Consumer privacy in online retail supply chains.

World Bank (2021). *World Development Report 2021: Data for Better Lives*. The World Bank.

Yom-Tov, E. & Ofran, Y. (2022). Implementation of data protection laws in the European Union and in California is associated with a move of clinical trials to countries with fewer data protections. *Frontiers in Medicine*, 9.

Yuan, B. & Li, J. (2019). The policy effect of the General Data Protection Regulation (GDPR) on the digital public health sector in the European Union: An empirical investigation. *International Journal of Environmental Research and Public Health*, 16(6).

Zhao, Y., Yildirim, P., & Chintagunta, P. K. (2021). Privacy regulations and online search friction: Evidence from GDPR. *Available at SSRN 3903599*.

Zhuo, R., Huffaker, B., kc claffy, & Greenstein, S. (2021). The impact of the General Data Protection Regulation on internet interconnection. *Telecommunications Policy*, 45(2), 102083.

Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., Goldschlag, N., Foster, L., & Dinlersoz, E. (2021). *Advanced technologies adoption and use by US firms: Evidence from the annual business survey*. Technical report, National Bureau of Economic Research.