

THE MISSING “CALIFORNIA EFFECT” IN DATA PRIVACY LAW

*Jens Frankenreiter**

ABSTRACT

The “California Effect” is a recurring trope in discussions about regulatory interdependence. This effect predicts that businesses active in multiple jurisdictions sometimes adopt the strictest standards that they face in any jurisdiction globally, even if the law does not require global compliance. There is a substantial literature that assumes the existence of California Effects both at the interstate level in the United States and the international level. However, empirical evidence documenting their existence and strength is scarce.

This paper investigates the existence of California Effects in data privacy law, a field in which these effects have been said to be particularly influential. Its main goal is to understand the extent to which EU law (which is usually described as comparably stringent) influences transactions between U.S. online services and consumers. Using a range of computational and traditional quantitative techniques, the paper tracks changes in almost 700 webpages’ privacy policies. The analysis covers two years starting in November 2017, a period that saw the enactment of a new, sweeping data privacy law in the EU. Contrary to what many assume, the analysis reveals that most U.S. online services treat U.S. consumers and EU consumers differently, with EU consumers enjoying higher levels of protection. This result indicates that the impact of EU law on the operations of U.S. online services is limited. Moreover, it suggests that California Effects driven by costs of differentiation might be less important than is commonly assumed, at least in data privacy law. The paper also discusses the implications of these findings for researchers and policymakers.

* Washington University in St. Louis School of Law. E-Mail: fjens@wustl.edu. I am grateful to participants in 25th Annual ISNIE / SIOE Conference, the Florida-Michigan-Virginia Virtual Law and Economics Workshop, the LEAP Text Analysis Conference at UC Berkeley, and audience members at the University of Chicago, Columbia, ETH Zurich, Oxford, USC, the University of Utah, and Yale for helpful comments. Particular thanks to Elliott Ash, Jack Balkin, Stefan Bechtold, Anu Bradford, Ryan Bubb, Dan Klerman, Meirav Furth-Matzkin, Yoan Hermstrüwer, Cathy Hwang, Dan Klerman, Mike Livermore, Florencia Marotta-Wurgler, Ruth Mason, Jonathan Masur, Justin McCrary, Tom Nachbar, Julian Nyarko, Sonja Starr, Paul Stephan, Lior Strahilevitz, Eric Talley, Pierre-Hugues Verdier, Mila Versteeg, and Tim Wu. All errors and omissions are solely mine.

TABLE OF CONTENTS

| | |
|---|----|
| INTRODUCTION..... | 4 |
| I. CALIFORNIA EFFECTS | 11 |
| A. Cost-Based California Effects..... | 13 |
| 1. Characteristics of Cost-Based California Effects | 13 |
| a. Trans-Jurisdictional Actors..... | 14 |
| b. Divergent Regulatory Standards..... | 16 |
| c. Costs of Differentiation and Global Compliance | 17 |
| 2. When do Cost-Based California Effects Occur? | 18 |
| 3. Distributional and Normative Implications | 19 |
| B. Other California Effects / Forms of Cross-Jurisdictional Influence | 20 |
| 1. Voluntary Compliance..... | 20 |
| 2. Diffusion of Laws..... | 21 |
| C. Cost-Based California Effects and the Internet..... | 21 |
| II. CONSUMER PRIVACY LAW IN THE UNITED STATES AND IN THE EU..... | 23 |
| A. The United States’ Market-Based Approach | 24 |
| B. Omnibus Regulation in the EU | 25 |
| C. The GDPR’s Legal Scope | 26 |
| III. COST-BASED CALIFORNIA EFFECTS IN DATA PRIVACY LAW?..... | 27 |
| A. General Considerations | 27 |
| B. Existing Empirical Evidence | 29 |
| IV. EMPIRICAL ANALYSIS | 31 |
| A. Research Design..... | 31 |
| 1. Empirical Approach..... | 31 |
| a. Focus on Privacy Policies..... | 31 |
| b. Changes Coinciding with the GDPR..... | 33 |
| c. Illustrations | 34 |
| d. Research Questions..... | 35 |
| 2. Data..... | 36 |
| B. Analysis and Results | 38 |
| 1. Computational Analysis | 38 |
| a. Outcome Measures | 38 |
| b. Results | 39 |
| i. The Development of U.S. Privacy Policies | 39 |
| ii. A Quantitative Test of Global Compliance | 44 |
| 2. Manual Coding | 46 |
| a. Sample Selection and Coding Scheme | 47 |

- b. Results 49
- 3. Determinants of Global Compliance 52
- 4. Other Potential Explanations 55
- C. Interpretation and Limitations 56
- V. IMPLICATIONS 57
 - A. Normative Implications..... 57
 - B. Implications for Regulatory Interdependence 61
 - C. Implications for Data Privacy Law 62
 - 1. Policy Implications 62
 - 2. The Role of the EU 63
- CONCLUSION..... 64

INTRODUCTION

In the spring of 2018, Google, Facebook, and several other leading tech companies announced major overhauls of their policies governing the handling of consumer data.¹ These changes were supposed to bring their data practices in line with the General Data Protection Regulation (GDPR),² a new data privacy law in the EU.³ Notably, Google’s and Facebook’s revised privacy policies applied to consumers everywhere, including in the United States.⁴

Even in our age of increasing regulatory interdependence,⁵ the global adoption of GDPR-compliant privacy policies by U.S. online services might seem startling. Compared to the GDPR, data privacy law in the United States imposes much less onerous obligations on online services,⁶ and it seems far from clear whether the United States will adopt comprehensive privacy legislation anytime soon. The GDPR itself does not legally apply to interactions between U.S. businesses and U.S. consumers.⁷ Besides, compliance with the GDPR’s various requirements is usually considered to be costly.⁸ Why then would Google and Facebook decide to extend these expensive protections to consumers in the United States?

And yet, the global nature of the privacy policy changes did not surprise most commentators. Some had for years predicted that stringent data privacy standards could spread between jurisdictions as a result of “California Effects”—a hypothesized process in which influential jurisdictions cause universal adoption of more stringent regulatory standards through unilateral

¹ Press Release, Facebook, *Complying With New Privacy Laws and Offering New Privacy Protections to Everyone, No Matter Where You Live* (Apr. 17, 2018), <https://about.fb.com/news/2018/04/new-privacy-protections/> [https://perma.cc/YP4A-54W3]; Press Release, Google, *Our Preparations for Europe’s New Data Protection Law*, (May 11, 2018), <https://blog.google/outreach-initiatives/public-policy/our-preparations-europes-new-data-protection-law/> [https://perma.cc/NP6H-9QZH].

² Parliament and Council Regulation 2016/679 (2016) [hereinafter: GDPR].

³ Katie Collins, *Google Makes Privacy Policy Clearer Than Ever to Comply with EU Law*, C|NET, MAY 11, 2018, <https://www.cnet.com/news/google-makes-privacy-policy-clearer-than-ever-to-comply-with-eu-gdpr-law/> [https://perma.cc/E9ZV-BSNS].

⁴ Press Release, Facebook, *supra* note 1; Press Release, Google, *supra* note 1. *See also* Michael L. Rustad & Thomas H. Koenig, *Towards a Global Data Privacy Standard*, 71 FLA. L. REV. 365 (2019), at 391-4.

⁵ *See* David Lazer, *Regulatory Interdependence and International Governance*, 8 JOURNAL OF EUROPEAN PUBLIC POLICY 474 (2001).

⁶ *See infra* notes 79–105 and accompanying text.

⁷ *See infra* notes 106–108 and accompanying text.

⁸ *E.g.*, Oliver Smith, *The GDPR Racket: Who’s Making Money From This \$9bn Business Shakedown*, FORBES, May 2, 2018, <https://www.forbes.com/sites/oliversmith/2018/05/02/the-gdpr-racket-whos-making-money-from-this-9bn-business-shakedown/?sh=75d181d134a2> [https://perma.cc/XBW7-GKD5].

policymaking.⁹ Proponents of this theory view the reactions by Google, Facebook, and their likes as evidence that it is often costly for online services to treat consumers in different jurisdictions differently. Hence, the global adoption of stringent standards is described as a cost-efficient—and rational—reaction for trans-jurisdictional businesses.¹⁰

The implications of this theory are momentous and reach far beyond data privacy law. In an increasingly interconnected world, we have become used to the fact that business regulation is a global enterprise: Many businesses’ activities span the globe, and national policymakers’ decisions affect outcomes far beyond the borders of their home jurisdiction. Yet the consequences of California Effects of the type described above differ from—and are arguably more drastic than—those of other forms of regulatory interdependence. The existence of California Effects in data privacy law would imply that high-standard jurisdictions like the EU can unilaterally force their regulatory standards on online transactions in other jurisdictions. Consequently, among others, the ability of U.S. policymakers to adopt alternative approaches in regulating transactions between U.S. businesses and their U.S. customers would be severely limited.

This paper challenges this narrative. It uses a novel dataset of privacy policies to investigate EU law’s influence on U.S. firms’ data practices on a larger scale than most existing studies. For this, it employs a range of empirical techniques, including state-of-the-art machine learning. This analysis shows that the GDPR prompted only few U.S. firms to adopt GDPR-compliant data practices globally. The analysis also suggests that California Effects in data privacy law are not, as many argue, driven by the costs of treating consumers in different jurisdictions differently. Instead, the evidence

⁹ See, e.g., JACK GOLDSMITH & TIM WU, WHO CONTROLS THE INTERNET? ILLUSIONS OF A BORDERLESS WORLD (2006), at 176. Note that the use of the term “California Effect” in the literature is somewhat inconsistent. Parts of the literature use it to describe the more general idea that trans-jurisdictional activity might help spread higher regulatory standards beyond the jurisdiction that initially enacted them. For example, consider David Vogel’s book *Trading Up*, which is commonly credited with coining the term “California Effect.” His account does not suggest that California’s rules prompted car manufacturers with sales in California to change the design of cars sold in other parts of the U.S. even in the absence of similar regulations there. Instead, Vogel describes how higher regulatory standards in some jurisdictions can incentivize trans-jurisdictional actors to lobby for the introduction of similar rules in other jurisdictions, tilting the political landscape in favor of more regulation. DAVID VOGEL, *TRADING UP: CONSUMER AND ENVIRONMENTAL REGULATION IN A GLOBAL ECONOMY* (1995), at 68-70.

¹⁰ E.g., ANU BRADFORD, *THE BRUSSELS EFFECT: HOW THE EUROPEAN UNION RULES THE WORLD* (2020), 142-3; Rustad & Koenig, *supra* note 4, at 391. See also Nitasha Tiku, *Europe’s New Privacy Law Will Change the Web, and More*, WIRED, Mar. 19, 2018, <https://www.wired.com/story/europes-new-privacy-law-will-change-the-web-and-more/> [<https://perma.cc/26G3-3D6P>].

points to other mechanisms—some of which have so far been largely ignored in the literature—as the main drivers of the decision by some to extend GDPR-style privacy protections to consumers in the United States. Consequently, this paper also casts doubt on claims about the power of the EU—or, for that matter, any other jurisdiction—to unilaterally impose rules on online transactions in the United States.

This result reminds us that, even in our age of incessant globalization, it is too early to declare national governance of business activities a relic of a bygone era. Nations remain the primary locus for politics and policymaking, and national borders have many significant consequences for the flow of labor, capital, and goods. A more accurate model of business regulation in the contemporary world recognizes that nations can be deeply embedded in a global context while retaining important areas of autonomy in which global influences are constrained.

Regulatory interdependence has long been recognized at both the domestic and international levels.¹¹ Domestically, these effects justify the federalization of certain areas of law.¹² Globally, regulatory interdependence is reflected in international trade law,¹³ networks of global banking

¹¹ See Lazer, *supra* note 5. In U.S. corporate law, scholars have for decades discussed whether permissive venue rules have led to an erosion of shareholder protections or instead, the emergence of more efficient corporate law. William L. Cary, *Federalism and Corporate Law: Reflections Upon Delaware*, 83 YALE L.J. 663 (1974); Roberta Romano, *The State Competition Debate in Corporate Law*, 8 CARDOZO L. REV. 709 (1987); Ralph K. Winter, *State Law, Shareholder Protection, and the Theory of the Corporation*, 6 JLEGAL STUD. 251 (1977). In environmental law, scholars have argued that emission standards set by Californian law have influenced not only the design of cars sold in all of the United States but in other nations as well. VOGEL, *supra* note 9. In all these situations, regulatory actions in one jurisdiction shape conduct in other jurisdictions. At the same time, such actions also impair the effectiveness of other jurisdictions' rules and their power to pursue their regulatory goals. Jack L. Goldsmith, *Against Cyberanarchy*, 65 U. CHI. L. REV. 1999 (1998), 1212.

¹² In environmental law, Congress cited the effects of regulatory interdependence (more precisely, the potential for detrimental competition between states) as a motivation to enact various statutes in this area. Richard L. Revesz, *Rehabilitating Interstate Competition: Rethinking the “Race-to-the-Bottom” Rationale for Federal Environmental Regulation*, N.Y.U. L. REV. 1210 (1992), 1226-1227. Considerations about harmful state competition also played a role in the enactment of New Deal legislation, and were cited by the Supreme Court in cases upholding such legislation. *United States v. Darby*, 312 U.S. 100 (1941), 115. See also Cass R. Sunstein, *Constitutionalism after the New Deal*, 101 HARV. L. REV. 421 (1987), 504-505. In corporate law, concerns about regulatory interdependence motivated similar calls for federalization. Cary, *supra* note 11. See also Lucien Arye Bebchuk, *Federalism and the Corporation: The Desirable Limits on State Competition in Corporate Law*, 105 HARV. L. REV. 1501 (1992).

¹³ See, e.g., Agreement on the Application of Sanitary and Phytosanitary Measures, 1867 U.N.T.S. 493.

regulators,¹⁴ and agreements concerning international tax reporting.¹⁵ In recent years, discussions about regulatory interdependence have taken on a new dimension in the United States as the emergence of other influential jurisdictions, most notably the European Union, has challenged its role as the primary exporter of legal rules.¹⁶

One recurring trope in discussions about regulatory interdependence is sometimes referred to as the California Effect. The hypothesis is that businesses active in multiple jurisdictions will sometimes adopt the strictest standards they face in any jurisdiction, even if the law does not mandate global compliance.¹⁷

One common explanation for this effect points to the costs of treating consumers in different jurisdictions differently. As the argument goes, California Effects often occur because firms find it less expensive to comply with the most stringent standard everywhere rather than provide different products to consumers in different jurisdictions.¹⁸ In her seminal work on the “Brussels Effect,” Anu Bradford identifies this mechanism as one of the main pathways through which the EU exerts influence globally.¹⁹ This version of California Effects, which I refer to as “Cost-Based” California Effects, is the main focus of this paper.

Costs of differentiation are, of course, not the only reason why firms might opt for global compliance with stringent regulatory standards. They might also do so because consumers in other jurisdictions are willing to pay higher prices for high-quality products or as a way to engage in virtue signalling.

Some in the literature treat different versions of the California Effect interchangeably.²⁰ However, the mechanisms giving rise to California Effects matter. Most importantly, Cost-Based California Effects have different normative implications than other versions of the California Effect and other forms of cross-jurisdictional influence. In the presence of Cost-Based

¹⁴ See Basel Comm. on Banking Supervision, *The Basel Framework*, BANK FOR INT’L SETTLEMENTS (2020), <https://www.bis.org/publ/bcbs189.pdf> [<https://perma.cc/9FGW-WYR2>].

¹⁵ *E.g.*, Agreement between the Government of the United States of America and the Government of Canada to Improve International Tax Compliance Through Enhanced Exchange of Information under the Convention Between the United States of America and Canada with Respect to Taxes on Income and on Capital, Feb. 5, 2014, T.I.A.S. 14-627.

¹⁶ See generally BRADFORD, *supra* note 10.

¹⁷ *Supra* note 9 and accompanying text.

¹⁸ *E.g.*, GOLDSMITH & WU, *supra* note 9, at 176.

¹⁹ BRADFORD, BRADFORD, *supra* note 10, at 142; Anu Bradford, *The Brussels Effect*, 107 Nw. U. L. Rev. 1 (2012).

²⁰ See, *e.g.*, BRADFORD, *supra* note 10, at 142-4 (describing how “de-facto Brussels Effects” can either be brought about by costs of differentiation or by consumer demand).

California Effects, trans-jurisdictional businesses comply with the most stringent standards globally even if—viewed in isolation—both businesses and consumers would profit from the application of local standards in low-protection jurisdictions.²¹ Similar concerns do not arise if the global compliance with the rules of one jurisdiction is motivated by businesses’ belief that they will profit from selling high-quality products in other jurisdictions.²² Consequently, a full assessment of the consequences and implications of California Effects requires differentiating between different versions of this phenomenon.

Although there is a substantial literature that assumes the existence of California Effects, relatively little work has been done to examine whether they are a widespread phenomenon. Most of the evidence that has been cited in support of their existence is anecdotal.²³ Systematic empirical studies are mostly absent from the literature. Also, little work has been done to distinguish the different mechanisms through which the laws of one jurisdiction affect outcomes elsewhere.

This paper contributes to an empirical understanding of California Effects by examining recent developments in data privacy law. With many online services catering to customers in a multitude of jurisdictions simultaneously, data privacy law has been hypothesized to be an area where Cost-Based California Effects are widespread.²⁴ The EU’s adoption of the GDPR raises the question of whether that legal change prompted widespread changes in U.S. online service providers’ data practices.

Because it is often impossible to observe the data practices of businesses directly,²⁵ my empirical strategy is to measure changes in publicly available websites’ privacy policies. The analysis relies on a longitudinal dataset consisting of the texts of the privacy policies of 695 websites, which I assembled with a co-author in another project.²⁶ The dataset contains one observation per week for the period between late November 2017 and October 2019. The analysis furthermore relies on a range of quantitative tools, including text analysis and machine learning.²⁷

The results of this analysis suggest, first, that the impact of EU data

²¹ See *infra* Sections I.A. 3 and V.A.

²² See *infra* Section I.B. 1.

²³ See, e.g., BRADFORD, *supra* note 10, at 143-6, 161-7; Rustad & Koenig, *supra* note 4, at 391-6.

²⁴ See *infra* Section III.A.

²⁵ But see Christian Peukert, et al., *European Privacy Law and Global Markets for Data* (ETH Zürich Center for Law & Econ. Working Paper 01/2020, 2020), <https://doi.org/10.3929/ethz-b-000406601>.

²⁶ See Jens Frankenreiter & Yoan Hermstrüwer, *Privacy’s Great Shock* (2021) (unpublished manuscript, on file with author).

²⁷ *Infra* Section V.A. 2.a.

privacy law on the relationship between U.S. businesses and their U.S. customers might be more limited than is commonly assumed. While the analysis documents changes in a large share of U.S. online services’ privacy policies,²⁸ only a (small) minority of these services adopted GDPR-compliant data policies globally.²⁹ Instead, most companies that modified their policies to bring them in line with the requirements of the GDPR took active steps to limit the scope of at least some of the additional protections to consumers in the EU.³⁰ Second, the evidence presented in this paper raises serious doubts about the hypothesis that differentiation costs played a major role in some businesses’ decisions to roll out GDPR-style protections on a global basis.³¹

This paper’s findings speak to a range of different literatures, including the literatures on data privacy law, consumer contracts, and regulatory interdependence.

With regard to data privacy law, the paper challenges the common notion that Cost-Based California Effect are a widespread phenomenon. Over the past several years, California Effects have become an important topic in the global discourse on data privacy.³² One reason for this interest is that the United States and the EU pursue radically different regulatory approaches. In the United States, the scope of consumer privacy protections is largely a matter of contracting.³³ By contrast, particularly since the entry into force of the GDPR in 2018, the EU imposes strict limits on the gathering and processing of personal data.³⁴ Therefore, data privacy law is an area where California Effects would lead to quite different policy outcomes. Against this background, the empirical analysis contextualizes the true reach of EU data privacy law. Its results suggest that, if widescale changes in data privacy practices in the United States are warranted, they will likely only come about due to domestic economic and political forces, not actions in other jurisdictions.

A related discussion in data privacy law concerns the role that regulation at the state level can play in protecting consumers’ interests across the United States.³⁵ In this context, there is a widespread expectation that the California’s new data privacy law (the CCPA³⁶) will have nationwide effects.³⁷

²⁸ *Infra* Section IV.B. 1.b.i.

²⁹ *Infra* Section IV.B. 1.b.ii.

³⁰ *Infra* Sections IV.B. 2.b.

³¹ *Infra* Section IV.B. 3.

³² *See infra* Section III.A.

³³ *See infra* Section II.A.

³⁴ *See infra* Section II.B.

³⁵ *See also* BRADFORD, *supra* note 10, at 146.

³⁶ California Consumer Privacy Act of 2018 (CCPA), AB 375.

³⁷ Matt Chinworth, *Don’t sell my data! We Finally Have a Law for That*, WASH. POST,

Analogous to predictions about the extraterritorial effects of EU law, such expectations are primarily based on the assumption that it will be too costly for businesses to differentiate between consumers in different states. The results in this paper cast doubt on the validity of this assumption.³⁸

With regard to consumer contracts more generally, the analysis suggests that it is often technically feasible and economically viable for online services to tailor their products to regulatory standards of individual jurisdictions. At the same time, the paper also documents how, even in the absence of Cost-Based California Effects, some companies might still decide to extend stringent standards to consumers in different jurisdictions. Potential explanations for this effect include a desire on the part of businesses to create positive public relations effects or establish themselves as brands that offer high standards of privacy protection.

With regard to the literature on regulatory interdependence, the paper provides one of the first systematic quantitative investigations of California Effects in an area in which their existence is often treated as a given. Its findings imply that Cost-Based California Effects play a much smaller role than is often assumed in the literature. Of course, these findings leave open the possibility that such effects can be a widespread phenomenon in other legal areas. However, they suggest that anecdotal evidence might convey a misleading picture of the prevalence of this mechanism and point to the possibility that empirical studies might reveal a more nuanced picture in other areas as well.

The remainder of this paper is structured as follows. Section I describes different versions of California Effects in more detail. Section II provides an overview of the state of data privacy law in the EU and the United States, while Section III summarizes the state of the debate about Cost-Based California Effects in this area. Section IV contains the main contribution of this paper, namely an empirical analysis of the conditions under which U.S. online services adjust their privacy policies to the requirements of EU law. Section V discusses the implications of my findings, followed by a brief conclusion.

Feb. 12, 2020, <https://www.washingtonpost.com/technology/2020/02/06/ccpa-faq/?arc404=true> [<https://perma.cc/JVE3-UKTD>]; Aaron Holmes, *Here's Why Facebook, Google, and Every Other Major Tech Company Are Updating Their Privacy Policy in Time for 2020, and What It Means for You*, BUSINESS INSIDER, Jan. 10, 2020, <https://www.businessinsider.com/why-tech-companies-new-privacy-policy-2020-california-2019-12> [<https://perma.cc/W9KJ-PCK7>]; Kashmir Hill, *Want Your Personal Data? Hand Over More Please*, N.Y. TIMES, Jan. 15, 2020, <https://www.nytimes.com/2020/01/15/technology/data-privacy-law-access.html> [<https://perma.cc/P4JB-X9CN>].

³⁸ Note, however, that it could be easier for businesses to differentiate between consumers in different countries than it is for them to differentiate between consumers in different states.

I. CALIFORNIA EFFECTS

When California sets new emissions standards for cars, General Motors will build cars to the Californian standard for the entire United States. Its choice to do so depends, of course, on the fact that it is more expensive to create cars customized for California than just build one car for the entire country.

– Jack Goldsmith and Tim Wu³⁹

The exporter has an incentive to adopt a global standard whenever its production or conduct is nondivisible across different markets or when the benefits of a uniform standard due to scale economies exceed the costs of forgoing lower production costs in less regulated markets. Complying with just one regulatory standard allows a corporation to maintain a single production process, which is less costly than tailoring its production to meet divergent regulatory standards. A single standard also facilitates the preservation of a uniform global brand. Thus, unilateral regulatory globalization follows from the nondivisibility of a corporation’s production or conduct.

– Anu Bradford⁴⁰

California Effects are a recurring trope in discussions about regulatory interdependence and the regulation of trans-jurisdictional business activities. As the first cite demonstrates, California Effects are often associated with California’s role in promoting higher automobile emission standards across the United States.⁴¹ In recent decades, California’s laws have often required cars sold in this state to comply with higher emission standards than other U.S. states and federal rules. A common assumption in the literature is that, in response to the introduction of such standards, carmakers started selling low-emission cars in all of the U.S.

In recent years, similar effects have increasingly been described in connection with the EU’s regulatory activities. As the story goes, there are many regulatory areas in which EU rules promulgate more stringent standards than apply elsewhere, including in the U.S. Major global businesses operating in the EU have to apply these standards in their interactions with consumers there.⁴² With regard to consumers elsewhere, they face a choice.

³⁹ GOLDSMITH & WU, *supra* note 9, at 176.

⁴⁰ Bradford, *supra* note 19, at 17-18.

⁴¹ *E.g.*, GOLDSMITH & WU, *supra* note 9, at 176.

⁴² In some legal areas, EU law requires businesses active in the EU to structure their global operations in accordance with EU law. Maybe the most important example is antitrust law. For example, mergers and acquisitions involving major business organizations are often

They can either treat them differently from EU consumers or apply the EU’s standards across the globe. Observers assume that it is often beneficial for businesses to opt for global compliance. Areas in which California Effects are said to result in a global application of the EU’s standards include food safety,⁴³ chemical safety,⁴⁴ environmental law,⁴⁵ online hate speech,⁴⁶ and data privacy law.⁴⁷

Importantly, many observers assume that the most important driver of these effects is the costs of treating consumers in different jurisdictions differently. This is true in the context of car emission standards, where many ascribe the extra-jurisdictional reach of California’s laws to the costs of building two different versions of each car model (one compliant with California law, the other with the law applicable in other states) at the same time.⁴⁸ As the second cite above suggests, this is also true for the EU’s regulatory activities.⁴⁹

However, costs of differentiating between consumers in different jurisdictions are just one of a range of mechanisms by which stringent standards in one jurisdiction can affect outcomes elsewhere. While many in the literature treat California Effects caused by different mechanisms interchangeably, their normative implications and their implications for the reality of regulatory interdependence differ substantially. Therefore, this article distinguishes between Cost-Based California and other versions of the California Effect.

subject to antitrust approval in the EU (as well as in other jurisdictions in which at least two of the entities are active) irrespective of where the businesses are headquartered. Council Regulation 139/2004 (2004), art. 1. The reason is that the effects of a merger of two businesses based in one jurisdiction will often not be limited to this jurisdiction and affect operations elsewhere. At the same time, the scope of laws in many other areas is more limited. For example, in consumer law, EU law usually does not apply if neither the consumer nor the business is based in the EU.

⁴³ BRADFORD, *supra* note 10, at 179-187.

⁴⁴ *Id.*, at 196-199.

⁴⁵ *Id.*, at 213-221.

⁴⁶ *Id.*, at 160-167.

⁴⁷ *Id.*, at 142-147.

⁴⁸ *E.g.*, GOLDSMITH & WU, *supra* note 9, at 176. *See also* BRADFORD, *supra* note 10, at 65-66.

⁴⁹ BRADFORD, BRADFORD, *supra* note 10, at 179-187; Bradford, *supra* note 19, at 17-18. *See also* GOLDSMITH & WU, *supra* note 9, at 176; Gregory Shaffer, *Globalization and Social Protection: The Impact of EU and International Rules in the Ratcheting Up of U.S. Privacy Standards*, 25 YALE J. INT’L L. 78 (2000).

A. Cost-Based California Effects

1. Characteristics of Cost-Based California Effects

I define Cost-Based California Effects as situations in which differentiation costs compel trans-jurisdictional actors to comply with the most stringent standard they face in any jurisdiction globally.

To illustrate this concept, consider the following example: Widget Inc. (W) is the only manufacturer of widgets in its home jurisdiction Columbiana and neighboring East Atlantica. Widgets are traditionally made from a plastic compound that is considered by some as a health hazard for consumers. Alternatively, widgets can be made from steel, rendering them harmless to health. However, steel widgets are more expensive to manufacture, and they have no other advantages over plastic widgets. To protect its consumers, East Atlantica adopts a law that requires that all widgets sold in East Atlantica are made from steel. Similar legislative initiatives are unsuccessful in Columbiana, where consumers are also unwilling to pay more for steel widgets.

In situations like that, how will W respond? Probably the most straightforward response is to start manufacturing steel widgets for its customers in East Atlantica while continuing to market plastic widgets to customers in Columbiana. It is also possible that the increased production costs associated with manufacturing steel widgets make it unprofitable to continue serving consumers in East Atlantica. Then, W will cease its activities there.

However, there are situations in which W’s best response is to offer steel widgets to consumers in both East Atlantica and Columbiana, even though plastic widgets are still legal in Columbiana, and even though East Atlantica’s laws do not apply to the sale of widgets in Columbiana. Such situations can occur if technical or economic reasons make it costly for W to market different types of widgets simultaneously. For example, the production costs of all widgets could increase if A had to configure its factory to manufacture both plastic and steel widgets.

A decision by W to shift its global production to steel widgets in order to avoid the costs of differentiation is an example of what I refer to as Cost-Based California Effects.

The definition of Cost-Based California Effects can be broken down into three elements: First, a business or similar actor is involved in transactions subject to the laws of different jurisdictions. Second, some jurisdictions impose more stringent standards on transactions than others. Moreover, third, differentiation costs make it a rational choice for the trans-jurisdictional actor

to apply the same standard to every transaction.⁵⁰ In the following, I will describe each of these requirements in more detail.

a. Trans-Jurisdictional Actors

Cost-Based California Effects occur in situations in which more than one jurisdiction can set up and enforce binding rules for (parts of) the activities of a trans-jurisdictional actor.⁵¹ This requirement is ordinarily met whenever an actor is active in more than one jurisdiction, as jurisdictions are entitled to regulate conduct insofar as it takes place or has effects in their territory.⁵² For example, if a business sells goods or services in different jurisdictions, every one of these jurisdictions can usually determine the rules that apply to transactions in their territory. By contrast, Cost-Based California Effects ordinarily do not occur in situations where a business is active in only one jurisdiction.

However, there exist situations in which jurisdictions are not in a position to effectively regulate conduct taking place or affecting outcomes in their territory. Most importantly, some legal areas are characterized by rules that restrict the power of jurisdictions to regulate trans-jurisdictional actors.⁵³ For

⁵⁰ In her work on the Brussels effect, Anu Bradford identifies five conditions that have to be met for the EU to exert global power through unilateral regulation. These conditions are market size, regulatory capacity, stringent standards, inelastic targets, and non-divisibility of standards. BRADFORD, *supra* note 10. While there might be differences on the margin, this description and my definition of California Effects largely overlap.

⁵¹ Note that it is not required that all jurisdictions that have the power to regulate exercise this power. For example, in the example above, Columbiana does not impose any limitations on the sale of widgets. Still, W’s activities fall under the scope of both Columbiana and East Atlantica’s laws, as both jurisdictions could regulate (at least) transactions between W and consumers in the respective jurisdiction.

⁵² See Restatement (Fourth) of the Foreign Relations Law, §407 (2018).

⁵³ Rules restricting the regulatory reach of jurisdictions are often adopted to save businesses the costs of having to deal with multiple regulatory environments at the same time. In order to achieve this goal, the power to regulate transactions of a trans-jurisdictional actor is concentrated with one jurisdiction, usually the actor’s home jurisdiction.

In principle, such rules can either be rules of the jurisdiction itself, or rules adopted at a higher level. Examples of the first type of rules are rules on personal jurisdiction and conflict-of-law. At least in principle, however, these rules can be changed to extend the reach of a jurisdiction’s laws. *Cf.* Harold W. Horowitz, *The Commerce Clause as a Limitation on State Choice-of-Law Doctrine*, 84 HARV. L. REV. 806 (1971). There are numerous examples of the second type of rule at the interstate level in the U.S., where U.S. federal law imposes important limitations to the power of U.S. states to regulate trans-jurisdictional conduct. Besides federal legislation, such limits can flow from the Fourteenth Amendment’s due process clause and the (dormant) commerce clause. *Cf.* *International Shoe v. State of Washington*, 326 U.S. 310 (1945); *Pike v. Bruce Church, Inc.*, 397 U.S. 137 (1970). At the

example, consider the National Banking Act, which bars states from regulating certain aspects of credit agreements between their citizens and banks incorporated elsewhere.⁵⁴ Another example is the internal affairs doctrine in corporate law, which concentrates the power to set the rules governing interactions between shareholders and directors of a corporation with the jurisdiction in which the corporation is incorporated.⁵⁵ While these situations might result in other forms of jurisdictional interdependence,⁵⁶ they usually do not give rise to Cost-Based California Effects.⁵⁷

While Cost-Based California Effects require that a trans-jurisdictional actor is subject to the regulatory authority of more than one jurisdiction, this does not imply that individual transactions need to fall under the legal scope

international level, international law can impose (although usually comparably weak) limits on regulation. *Cf.* Goldsmith, *supra* note 11, at 1219. Limits applying to states worldwide can flow from customary international law and international treaties such as the General Agreement on Tariffs and Trade and other trade law instruments. *See* Restatement (Fourth) of the Foreign Relations Law, §407 (2018). *See also* Appellate Body Report, *United States — Measures Affecting the Cross-Border Supply of Gambling and Betting Services*, WTO Doc. WT/DS285/AB/R (adopted Apr. 20, 2005) (finding that the U.S. violated trade law in prohibiting providers of online gambling services based in Antigua to offer their services over the internet to customers based in the U.S.).

⁵⁴ *See* *Marquette Nat. Bank of Minneapolis v. First of Omaha Service Corp.*, 439 U.S. 299 (1978).

⁵⁵ The internal affairs doctrine is the traditional conflict-of-laws rule in Anglo-American countries. P. John Kozyris, *Corporate Wars and Choice of Law*, 1985 DUKE L.J. 3. Note that it is disputed whether the internal affairs doctrine is guaranteed by the U.S. constitution. *Compare* *VantagePoint v. Examen, Inc.*, 871 A. 2d 1108, 1113 (Del. 2005) with Richard M. Buxbaum, *The Threatened Constitutionalization of the Internal Affairs Doctrine in Corporation Law*, 75 CAL. L. REV. 35 and Fredeck Tung, *Before Competition: Origins of the Internal Affairs Doctrine*, 32 J. CORP. L. 33 (2006): at note 29. Be that as it may, the Supreme Court has, in the past, dismissed as unconstitutional certain state anti-takeover statutes that applied to corporations incorporated elsewhere. *Edgar v. MITE Corp.*, 457 U.S. 624 (1982). *Cf.* *CTS Corp. v. Dynamics Corp. of America*, 481 U.S. 69 (1987).

⁵⁶ As is widely discussed in the corporate law literature, these situations can result in a competition between jurisdictions that can have important ramifications for the standards of protection that apply. Bebachuk, *supra* note 12; Cary, *supra* note 11; Romano, *supra* note 11; Winter, *supra* note 11.

⁵⁷ Besides, there can also be factual barriers to regulation. Most importantly, jurisdictions might be unable to enforce their laws against a trans-jurisdictional actor. This situation can arise if a trans-jurisdictional actor does not have any physical presence or assets located in the respective jurisdiction. *See* Goldsmith, *supra* note 11, at 1217. Note that this obstacle's importance depends on whether a jurisdiction can rely on other jurisdictions to enforce its judgments. Within the U.S., the full faith and credit clause ensures that individual states' judgments that satisfy certain minimum requirements can be enforced nationwide. *See id.* In the international context, treaties on mutual judicial assistance allow jurisdictions to overcome some enforcement gaps. However, in most contexts, public policy exceptions allow countries to deny the enforcement of foreign judgments in conflict with their fundamental values. *Id.*, at 1219–1220.

of multiple laws at the same time. Instead, Cost-Based California Effects are characterized by “excessive” compliance with the laws that impose the strictest standards on certain types of transactions: Technical or economic factors rather than legal obligations compel a business or similar actor to apply stringent standards in its global operations, including in situations in which the law does not require compliance.

b. Divergent Regulatory Standards

Cost-Based California Effects furthermore require that some of the standards imposed by jurisdictions on a certain type of transaction are more stringent than others, or in other words, that the standards imposed by various jurisdictions diverge.

I use the term standard to refer to any requirement that the law imposes on transactions. These requirements can take on various forms. For example, they can relate to the substance of the transaction or its form. Examples of substantive requirements include the imposition of a price ceiling or the stipulation of mandatory product characteristics. Substantive requirements can also confer rights on one party that cannot be bartered away. Formal requirements include using a specific contractual form, disclosure requirements, and similar formalities that have to be fulfilled to make the transaction legal.⁵⁸

Standards of jurisdictions diverge whenever there are transactions that are legal under the laws of one jurisdiction, but illegal under the laws of other jurisdictions. More precisely, what is required is a hypothetical determination of whether the laws of different jurisdictions would treat the same transaction differently were it to fall under the scope of all laws simultaneously. In the example above, the standards that Columbiana and East Atlantica impose on the sale of widgets diverge. This is because the sale of plastic widgets is legal in Columbiana, but illegal in East Atlantica.

To determine which of these different standards constitutes the more stringent one, it makes sense to differentiate between two constellations. Consider first the case in which the standards of different jurisdictions have a “nested” relationship. Such a nested relationship exists between two jurisdictions insofar as a transaction in compliance with the first jurisdiction’s laws automatically complies with the laws of the second jurisdiction, while the opposite is not true. In this case, the standard imposed by the first jurisdiction is the more stringent standard.⁵⁹ The example above falls into this

⁵⁸ As described above, jurisdictions can also impose no specific requirements on a transaction type. *Supra* note 51.

⁵⁹ If there are more than two jurisdictions, a nested relationship need not exist between

first category. After the introduction of the law banning the sale of plastic widgets, East Atlantica’s are the more stringent standards because they ban a subset of transactions that are legal in Columbiana. At the same time, all legal sales of widgets in East Atlantica would also be legal if they took place in Columbiana.

Second, there are situations in which multiple jurisdictions’ standards are not nested, but in which a subset of transactions could pass under the laws of all jurisdictions.⁶⁰ To understand what that means, consider a modified version of the example above in which Columbiana introduces a ban on widgets that are circular in shape. In this situation, there are transactions that are legal under East Atlantica’s laws, but illegal under Columbiana’s laws (the sale of circular steel widgets). There are also constellations for which the opposite is true (the sale of quadratic plastic widgets). In situations like these, the most stringent standard is not the law of one jurisdiction, but a combination of all jurisdictions’ laws. In the example above, the emerging standard would be one that requires widgets to be both non-circular and made from steel.

c. Costs of Differentiation and Global Compliance

Cost-Based California Effects are situations in which a trans-jurisdictional actor complies with the most stringent standard globally to realize cost savings associated with treating customers in different jurisdictions alike. Put differently, global compliance has to be motivated by a desire to save costs of differentiation.

Costs of differentiation are any added production or transaction costs that businesses face if they treat consumers in different jurisdictions differently. Such costs can be related to different sources. First, they can concern the production of goods or services. For example, it can be costly to maintain different product lines for consumers in different jurisdictions. The example above falls into this first category. Second, differentiation costs can also emerge in the form of increased transaction costs if there are special requirements for contracts in some jurisdictions, but not in others. For example, jurisdictions can require different contractual formalities for certain

all of them. Instead, it is sufficient that there is one or more jurisdiction (in the latter case, with both jurisdictions imposing similar standards) that “dominate” all other jurisdictions.

⁶⁰ At least in theory, there can also be situations in which the standards imposed by different jurisdictions are mutually exclusive. Mutual exclusivity of standards implies that there cannot be any transactions of a particular type that would be considered legal in all jurisdictions. In the example above, consider that Columbiana passes a law that requires widgets to be made from wood. In these situations, trans-jurisdictional actors cannot offer the same product to customers in different jurisdictions. Instead, W’s only option is to offer different products to customers in Columbiana and East Atlantica.

types of transactions, or they can endow consumers with different contractual rights.

2. When do Cost-Based California Effects Occur?

Not every situation in which businesses face costs of differentiation will give rise to Cost-Based California Effects. This is because the decision to apply stringent standards globally will usually also result in added compliance costs. These added compliance costs have to be balanced against the benefits of treating all customers alike. Accordingly, Cost-Based California Effects only occur if the total cost savings from treating consumers across jurisdiction alike exceed the added costs of compliance.

To understand what this balancing of costs entails, consider first the costs of treating consumers everywhere in accordance with the most stringent standard. These costs will usually be a function of the regulatory requirement at hand and the amount of business a firm conducts in low-protection jurisdictions. All else equal, the total added compliance costs will be higher for firms with a higher share of consumers in these low-protection jurisdictions.⁶¹

Second, consider the costs of treating consumers in different jurisdictions differently. These costs will likely differ between different industries and depending on the legal requirement in question. For example, in case of a law imposing requirements on physical goods' product design, differentiation costs will often be comparably high. This is because a firm's decision to treat consumers in different jurisdictions differently would imply the simultaneous production of more than one product line. By contrast, firms should find it easier to restrict the application of a law requiring the granting of mandatory product warranties to just one jurisdiction.

Costs of differentiation can be either variable costs, fixed costs, or a combination of both. Importantly, different from the added costs of compliance, these costs need not be (positively) related to the share of a business's customers in low-standard jurisdictions. This is the case, first, insofar as the costs of differentiation are fixed costs. In the example above, imagine that the simultaneous manufacturing of steel and plastic widgets requires *W* to build a second production facility. Second, it seems possible that the differential treatment of consumers in different jurisdictions increases the costs of doing business in the high-standard jurisdiction as well. This can happen, for example, if product differentiation implies forgone economies of

⁶¹ This is true whenever some of the compliance costs are variable costs. Note that, insofar as compliance costs consist of fixed costs, these costs do not increase the costs of extending compliance with high standards to other jurisdictions.

scale that would have decreased per unit production costs everywhere.⁶²

These considerations suggest that, all else equal (and assuming that companies continue serving consumers in the high-standard jurisdiction), smaller firms are more likely than bigger ones to comply with more stringent standards across jurisdictions. Among businesses of similar size, Cost-Based California Effects will most likely influence the decision-making of those that derive more of their revenues from transactions in the high-standard jurisdiction.⁶³ The reason for both predictions is related to the amount of added costs of compliance that these firms face when deciding to extend the most stringent standards to consumers in other jurisdictions. These costs will usually be higher for firms with higher sales figures overall and for firms that conduct a larger share of their business in low-protection jurisdictions.

3. Distributional and Normative Implications

Cost-Based California Effects can have important distributional consequences. In the example above, W is not the only actor affected by the decision whether to sell steel widgets in Columbiana; instead, this decision also has implications for consumers in both Columbiana and East Atlantica. The reason for this is that an increase in production costs will often result in higher product prices, lower numbers of products sold, and a decrease in consumer welfare.⁶⁴ Under the assumption that the decision to offer different products will increase production costs everywhere, East Atlantica consumers will be better off if W sells steel widgets in Columbiana as well. Consumers in Columbiana will prefer the opposite decision, at least if plastic widgets can still be offered at a cheaper price compared to steel widgets.

More generally, if differentiation is costly, consumers in jurisdictions that impose the most stringent standards will usually benefit from this standard’s global application. At the same time, a decision in favor of global compliance can increase product prices in other jurisdictions. If the stringent standards benefit consumers, consumers in these jurisdictions might accept higher price tags.⁶⁵ However, this need not be the case. In particular, consumers in

⁶² See Bradford, *supra* note 19, at 17-18.

⁶³ This prediction also suggests that Cost-Based California Effects are most likely in the context of standards enacted by comparably large jurisdictions (e.g., California at the interstate level in the U.S., the U.S. and the EU at the international level). By contrast, if smaller jurisdictions enact similarly high standards, it is often rational for trans-jurisdictional actors to limit compliance to the extent required by the law. See also Anu Bradford, *Exporting Standards: The Externalization of the EU’s Regulatory Power via Markets*, 42 INT’L REV. L. ECON. 158 (2015), 161.

⁶⁴ Whether such consequences occur depends mainly on the competitive structure of a market and the number of companies that change their offerings due to California Effects.

⁶⁵ See also BRADFORD, *supra* note 10, at 239-240.

different jurisdictions might have different preferences regarding the appropriate level of regulation.⁶⁶ If this is the case, the consequences of Cost-Based California Effects can be normatively problematic.⁶⁷

B. Other California Effects / Forms of Cross-Jurisdictional Influence

Cost-Based California Effects are among several mechanisms by which stringent standards in one jurisdiction can affect outcomes in other jurisdictions. While these mechanisms seem to lead to similar outcomes on their face, their normative implications and their implications for the reality of regulatory interdependence differ substantially.

1. Voluntary Compliance

First, businesses might comply with stringent standards globally for reasons that are unrelated to costs of differentiation. Most importantly, businesses can offer high-standard products globally to increase their revenues, as consumers might be willing to pay more for such products.⁶⁸ In the example above, assume that there is a substantial percentage of Columbiana’s population that prefers steel widgets over plastic widgets, and that is willing to pay a higher price for the latter. In this case, even if there were no costs of differentiation, it would be rational for W to start selling steel widgets in Columbiana.⁶⁹

⁶⁶ See also Richard L. Revesz, *The Race to the Bottom and Federal Environmental Regulation: A Response to Critics*, 82 MINN. L. REV., 535, 536. Besides, it is also possible that the cost and benefits of certain types of regulation vary across jurisdictions. *Id.*, at 536-7.

⁶⁷ I discuss the normative consequences of California Effects at greater length below. *Infra* Section V.A.

⁶⁸ Bradford’s description of the Brussels effect includes instances in which businesses appear to have acted out of such motivation. *E.g.*, BRADFORD, *supra* note 10, at 144-145.

⁶⁹ At the same time, if consumer demand justifies global compliance with stringent standards, businesses ordinarily have incentives to offer consumers an option to purchase high-standard products even without mandatory laws in any jurisdiction. Then, why does a legal intervention in one jurisdiction lead to a change in transactions elsewhere? Aside from costs of differentiation, there are at least three mechanisms by which laws of one jurisdiction can bring about such a change. First, laws in one jurisdiction could help overcome market failures in other jurisdictions related to consumers’ inability to differentiate between high-standard and low-standard products. See George A. Akerlof, *The Market for “Lemons”: Quality Uncertainty and the Market Mechanism*, 84 Q. J. ECON. 488. In this situation, other jurisdictions’ laws and enforcement activities can play a role similar to that of private certification providers. Second, a legal change in one jurisdiction can lead to a shift of consumer preferences in another jurisdiction, for example, because it increases awareness about specific problems. Finally, if most of the costs required to comply with the new

While voluntary compliance and Cost-Based California Effects might appear to lead to similar outcomes, their consequences differ substantially. First, voluntary compliance usually implies that consumers in low-standard jurisdictions—at least in aggregate—benefit from the introduction of high-standard products. This implies that the distributional consequences described in Section I.A. 3 above are likely absent from instances of voluntary compliance.⁷⁰ Second, in the case of voluntary compliance, businesses can sell low-standard products alongside high-standard products if there is sufficient demand. Finally, in the absence of differentiation costs, businesses will typically be more easily able to revert their decision in favor of global compliance in the face of changing circumstances.

2. Diffusion of Laws

Second, a propagation of stringent standards across jurisdictions can also be the result of a diffusion of laws. In other words, other jurisdictions might decide to copy laws that implemented a particularly stringent standard elsewhere.⁷¹ These cases differ from Cost-Based California Effects on various dimensions. Most importantly, there is no direct relationship between the adoption of stringent laws in one jurisdiction and changes in transactions in other jurisdictions. As a consequence, any constraints that policy makers in other jurisdictions might experience in such a situation are different in nature than the constraints imposed by Cost-Based California Effects.⁷²

C. Cost-Based California Effects and the Internet

This article focuses on transactions between businesses and consumers

standard imposed by one jurisdiction are fixed costs, the expenditure of these costs can unlock more profitable business opportunities in other jurisdictions. Importantly, this is true even if the benefits of selling improved products in all jurisdictions are not high enough to justify the investment absent a legal obligation in at least one jurisdiction.

⁷⁰ This is because companies have incentives to offer whichever product maximizes the total surplus, which is divided between the company and its customers. Under normal circumstances, therefore, a voluntary decision on the part of a business to switch to high-standard products will maximize not only the business's profits, but also aggregate consumer welfare.

⁷¹ Parts of the literature on the regulatory interdependence describe this effect as an instance of the California effect. *Supra* note 41.

⁷² There are various ways in which the adoption of stringent standards in one jurisdiction can tilt the political landscape in other jurisdictions in favour of similar policy initiatives. For one, the former jurisdiction might attempt to exert pressure on other jurisdictions to adopt similar standards. For another, businesses active in multiple jurisdictions might lobby for the introduction of stringent standards everywhere, in particular because it might afford them advantages over local competitors. *See, e.g.*, VOGEL, *supra* note 9, at 68-70.

that take place on the internet. Different from traditional transactions, the actors' physical location often does not constrain interactions on the internet. At least in principle, content and services made available on websites and similar devices can be accessed everywhere. Also, the internet's architecture implies that it can be costly, and sometimes even impossible, for actors to ascertain the identity and physical location of a party with whom they interact.

Against this background, a naïve view might hold that differentiation costs are substantially higher for online service providers than for other businesses. If this were the case, Cost-Based California Effects would likely be more prevalent in the context of transactions on the internet than they are in traditional transactions. Taken to the extreme, if online services were generally unable to distinguish between customers in different jurisdictions,⁷³ such effects would be ubiquitous on the internet.⁷⁴

If this was ever an adequate description of online activities, it has been rendered obsolete by two parallel developments. The first development is the emergence of ever-better geo-identification technology.⁷⁵ This technology allows providers of online services to distinguish—with some degree of certainty—between customers located in different jurisdictions, and to offer different versions of their services to these customers.⁷⁶

The second development concerns the scope of laws regulating interactions between online service providers and their customers. Mostly, jurisdictions refrain from applying their laws to transactions between online service providers and consumers if the service provider has taken appropriate measures to prevent consumers in this jurisdiction from accessing a website

⁷³ This is equivalent to assuming infinite differentiation costs.

⁷⁴ There are several examples of cases before courts of various jurisdictions in which online service providers (unsuccessfully) argued for exemptions from regulation based on the argument that it would require them to change their operations in other jurisdictions as well. This argument also played a significant role in early academic debates about the regulation of online activities. Proponents of the cyberlibertarian movement in the 1990s in particular argued that the regulation of online activity would result in the simultaneous application of the laws of all jurisdictions simultaneously. *E.g.*, David R. Johnson and David Post, *Law and Borders—The Rise of Law in Cyberspace*, 48 STAN. L. REV. 1374. As Jack Goldsmith argued, advocates of this view overstated the extent of the ensuing problem because of limits in jurisdictions' power to enforce their laws against foreign actors. Goldsmith, *supra* note 11, at 1220.

⁷⁵ See GOLDSMITH & WU, *supra* note 9, at 60-62; DAN JERKER B. SVANTESSON, PRIVATE INTERNATIONAL LAW AND THE INTERNET (3rd ed. 2016), at 525-36.

⁷⁶ Differentiating between customers is even more straightforward if there are elements of a transaction that take place in the real world. For example, shopping websites, food delivery services and similar businesses can limit deliveries to specific jurisdictions or areas. Paid online content can be restricted to customers whose residence in a particular jurisdiction has been confirmed by a provider of payment services such as a bank or credit card company.

or service.⁷⁷

Together, these developments imply that it is generally feasible for online service providers to ascertain the laws that apply to a given transaction and modify their handling of the transaction according to these laws. As a result, there is little reason to assume that Cost-Based California Effects are inherently more common in online transactions than they are in traditional transactions.⁷⁸

II. CONSUMER PRIVACY LAW IN THE UNITED STATES AND IN THE EU

While data privacy laws in the United States and the EU have common intellectual roots⁷⁹ and developed in a similar direction during their early history,⁸⁰ their developments have taken strikingly different paths over the past decades.⁸¹ Today, these jurisdictions occupy what can be seen as the opposite poles of the spectrum of regulatory approaches to data privacy among liberal democracies.⁸² The EU has emerged as a forerunner in implementing so-called “omnibus” privacy laws which establish comprehensive, mandatory standards of protection that limit the collection and use of personal data by both public and private actors.⁸³ In the U.S., no such comprehensive set of rules exist. Federal (and until very recently, state) legislation targeting business is limited to narrow subfields such as education and credit reporting.⁸⁴ In most constellations, it is therefore left to the market to determine the scope of privacy protections for customers *vis-à-vis*

⁷⁷ See, e.g., *International League against Racism & Anti-Semitism (LICRA) and the Union of French Jewish Students (UEFJF) v. Yahoo! Inc.*, County Court of Paris, Nov. 20, 2005.

⁷⁸ See also Goldsmith, *supra* note 11, at 1200-1201.

⁷⁹ See Paul M. Schwartz, *The EU-U.S. Privacy Collision: A Turn to Institutions and Procedures*, 126 HARV. L. REV. 1966 (2013), at 1970–1971 (describing how the data privacy discourse in Germany had been influenced by early work on privacy law in the United States).

⁸⁰ See *id.*, at 1975 (describing how international harmonization even led some observers to hypothesize about a convergence of regulation).

⁸¹ See generally Paul M. Schwartz & Karl-Nikolaus Pfeifer, *Transatlantic Data Privacy Law*, 106 GEO L.J. 115 (2017).

⁸² See, e.g. Franz-Stefan Gady, *EU/U.S. Approaches to Data Privacy and the “Brussels Effect”: A Comparative Analysis*, 2014 GEO, J. INTL AFFAIRS 12 (2014), 15; Kimberly A. Houser & W. Gregory Voss, *GDPR: The End of Google and Facebook or a New Paradigm in Data Privacy?*, 25 RICH. J.L. & TECH. 1 (2018), 9.

⁸³ See Schwartz, *supra* note 79, at 1973–1974 (“the Directive has encouraged the rise of omnibus legislation throughout the EU and most of the world”); Paul M. Schwartz, *Global Data Privacy: The EU Way*, 94 N.Y.U. L. REV. 771 (2019), 777-778 (describing how EU data privacy law has inspired similar legislation elsewhere).

⁸⁴ E.g., Schwartz, *supra* note 79, at 1974-1975.

businesses.⁸⁵

A. *The United States’ Market-Based Approach*

One of the defining features of consumer privacy law in the U.S. is that businesses are by default free to gather, process, and share information that they obtain from their customers. Consumers enjoy legal protection only under a rather narrow set of circumstances. Limits on permissible data practices can follow from various legal sources. First, “sectoral” federal and state legislation restricts the gathering and use of information by specific businesses and concerning specific types of data.⁸⁶ Second, data practices can run afoul of rules that are applicable beyond data privacy law.⁸⁷ These rules include common law institutions such as contract law as well as statutory law. From a practical perspective, the most important example in this category is section 5 of the Federal Trade Commission Act, which has formed the basis for several enforcement actions brought by the FTC against data practices perceived as deceptive or unfair.⁸⁸ Finally, California recently adopted the CCPA, which imposes a range of obligations on most businesses that gather California consumers’ data.

In practice, whenever sectoral legislation is not applicable, data privacy is mostly a matter of contracting between customers and businesses.⁸⁹

⁸⁵ See Schwartz & Pfeifer, *supra* note 81, at 132 (“Unlike the EU’s data subject, U.S. law does not equip the privacy consumer with fundamental constitutional rights; rather, she participates in a series of free exchanges involving her personal information. In this legal universe, the rhetoric of bilateral self-interest holds sway.”); Shaffer, *supra* note 49, at 13 (“the United States . . . relies more on private ordering through market processes.”).

⁸⁶ E.g., Fair Credit Reporting Act (FCRA), 15 U.S.C. §§1681-1681x (2006 & Supp. V 2011); Family Educational Rights and Privacy Act (FERPA), 20 U.S.C. §1232g (2006 & Supp. V 2011).

⁸⁷ DANIEL J. SOLOVE & PAUL M. SCHWARTZ, *PRIVACY LAW FUNDAMENTALS* 2019 (2019), 134-144.

⁸⁸ Daniel J. Solove and Woodrow Hartzog, *The FTC and the New Common Law of Privacy*, 114 COLUMBIA LAW REVIEW 583 (2014). These activities focus on broken promises in privacy policies and other deceptive and unfair practices. *ibid.*, 627-643; SOLOVE & SCHWARTZ, *supra* note 87, at 143. The most important substantive constraints following from the FTC’s interpretation of section 5 FTC Act concern the implementation of adequate security practices to guard against data security breaches. Solove & Hartzog, *supra*, at 636-638, 643. Note that some in the literature have raised doubts about the effectiveness of this regime. Florencia Marotta-Wurgler & Dan Svirsky, *Do FTC Privacy Enforcement Actions Matter? Compliance Before and After US-EU Safe Harbor Agreement Actions* (NYU Law and Economics Working Paper, on file with the author).

⁸⁹ See Kevin E. Davis & Florencia Marotta-Wurgler, *Contracting for Personal Data*, 94 N.Y.U. L. Rev. 662 (2019), 663 (“To a large extent, the relationship between the business and user with regards to information privacy is contractual”). Note that it is

Businesses face no substantial constraints on their data practices as long as they provide consumers with an accurate and transparent description of these practices.⁹⁰ This is different only for California residents who, since the entry into force of CCPA, enjoy certain rights *vis-a-vis* businesses that collect information on them.

B. Omnibus Regulation in the EU

Since the late 1990s, EU law has offered consumers a uniform set of comprehensive protections against the collection and use of personal data by both public and private actors.⁹¹ Since then, businesses, by default, require a legal justification to gather, process, and share personal information about consumers in the EU. One important avenue for businesses to obtain authorization is to demonstrate that the “processing [of data] is necessary for the purposes of the[ir] legitimate interests,”⁹² which essentially delegates the decision about the scope of permissible data practices to the agencies and courts tasked with enforcing data privacy law. Authorization can also be obtained by securing the consumer’s consent.⁹³ Notably, however, EU law establishes both formal requirements for obtaining consent⁹⁴ and mandatory rights for consumers that cannot be contracted away.⁹⁵ Besides, EU law provides for the establishment of specialized enforcement agencies tasked with prosecuting data privacy violations.

While EU law has followed this general approach since the entry into

subject to dispute whether privacy notices outlining a business’s data practice should be treated as contracts in the legal sense. Compare SOLOVE & SCHWARTZ, *supra* note 87, at 136 with Omri Ben-Shahar & Lior Jacob Strahilevitz, *Contracting over Privacy: Introduction*, 45 J. LEGAL STUD. S7 (2016).

⁹⁰ Davis & Marotta-Wurgler, *supra* note 89, at 663 (“The protection of consumer information in the United States has followed a ‘Notice and Choice’ approach, where businesses outline their information privacy policies, which are typically incorporated by references in general Terms of Service contracts, to which users must agree”).

⁹¹ Before the EU started regulating privacy law, privacy law had been a domain of the EU member states, some of which had enacted comparably strong privacy protections even in the absence of EU law. See Schwartz, *supra* note 79, at 1969-1971.

⁹² GDPR art. 6(1)f. *Cf.* Council Directive 95/46 (1995), art. 7f [hereinafter: Data Protection Directive].

⁹³ GDPR art. 6(1)a. *Cf.* Data Protection Directive, art. 7a.

⁹⁴ Such formal requirements include restrictions on blanket provisions and certain forms of click-wrap contracts. See C-673/17, *Bundesverband der Verbraucherzentralen und Verbraucherverbände v. Planet49 GmbH*, C:2019:801 (2019), at 44-65 (failing to deselect pre-checked checkboxes does not imply consent).

⁹⁵ See also Schwartz & Pfeifer, *supra* note 81, at 139 (“EU data protection law establishes important areas of inalienable privacy, setting out bedrock data protection principles that are not subject to individual waiver and cannot be traded away in bargained-for exchanges.”).

force of the Data Protection Directive in 1995,⁹⁶ the GDPR substantially tightened the restrictions for businesses handling consumer data along multiple dimensions.⁹⁷ First, it scaled up the requirements that need to be met to legally handle consumer data in the first place. In particular, to obtain a consumer’s consent, businesses now need to provide them with a clear description of every intended use of their data.⁹⁸ Insofar as information is not needed to provide a good or service, it is generally impermissible to make interactions with consumers conditional on their consent with a business’s data practices.⁹⁹ Also, consumers can withdraw their consent at any time, rendering future processing of the data illegal.¹⁰⁰ Second, the GDPR extended the number and scope of rights that consumers enjoy *vis-a-vis* businesses that obtained information on them in the past. They can, *inter alia*, request information about the usage of their data,¹⁰¹ demand correction of any false information¹⁰² as well as deletion of information that is no longer needed,¹⁰³ and ask for a copy of the information obtained by the business in order to supply this information to another actor.¹⁰⁴ Finally, the GDPR also introduced the possibility for agencies to impose substantial monetary fines for violations.¹⁰⁵

In sum, EU law imposes substantial restrictions on businesses’ handling of consumer information that cannot be overridden by contractual agreement. Rather than trusting the market mechanism to determine the ideal scope of permissible data practices, the EU approach relies heavily on public actors such as enforcement agencies and courts.

C. The GDPR’s Legal Scope

While the territorial scope of the GDPR is comparably broad, it is not unlimited. Most importantly in the context of this article, EU privacy law is generally not applicable to transactions in which neither the business nor the

⁹⁶ See Data Protection Directive. Because of the choice of an EU Directive as the form of rulemaking, the Data Protection Directive required the adoption of national laws in the EU Member States to become fully effective. This is different for the GDPR, which (because it takes the form of an EU Regulation) is directly applicable.

⁹⁷ *E.g.*, Rustad & Koenig, *supra* note 4, at 376-379.

⁹⁸ GDPR art. 7(2), 13(1).

⁹⁹ GDPR art. 7(4).

¹⁰⁰ GDPR art. 7(3).

¹⁰¹ GDPR art. 15.

¹⁰² GDPR art. 16.

¹⁰³ GDPR art. 17.

¹⁰⁴ GDPR art. 20.

¹⁰⁵ Fines for violations of the GDPR can amount to up to 4% of an undertaking’s annual worldwide turnover.

consumers are physically present in the EU.

The application of the GDPR is triggered whenever one of two conditions is met. First, the GDPR covers all handling of personal data that is done by businesses or business units operating out of the EU.¹⁰⁶ Second, it also covers other businesses or business units’ data practices as their activities target consumers in the EU.¹⁰⁷ As a consequence, EU privacy law usually does not apply to interactions between businesses and consumers if none of them is located in the EU.¹⁰⁸

III. COST-BASED CALIFORNIA EFFECTS IN DATA PRIVACY LAW?

A. General Considerations

The de facto Brussels Effect is particularly strong in the domain of data privacy . . . Various examples suggest that, for today’s global digital companies, maintaining different data practices across global markets is often both difficult (due to technical non-divisibility) and costly (due to economic non-divisibility).

– Anu Bradford¹⁰⁹

As this cite demonstrates, Cost-Based California Effects are assumed to play an important role in data privacy law.¹¹⁰ Commentators have, for

¹⁰⁶ GDPR, art. 3(1). This norm establishes that all handling of consumer data that takes place “in the context of the activities of an establishment” in the EU is subject to the GDPR, even though the data processing itself might take place elsewhere. According to the case-law of the Court of Justice, this requirement is met whenever a business has a “branch or subsidiary” in one of the member states, and the use of consumer data is connected to the activities of this business unit. *Google LLC v. CNIL*, C:2019:772 (2019), at 48-52 (deciding that Google’s use of information about consumers to build its products together with the general “commercial and advertising activities” of Google’s French subsidiary was sufficient to fulfill this condition).

¹⁰⁷ GDPR, art. 3(2); EUROPEAN DATA PROTECTION BOARD, GUIDELINES 3/2018 ON THE TERRITORIAL SCOPE OF THE GDPR (ARTICLE 3) (2019).

¹⁰⁸ Given the scarcity of case law on Article 3, the precise territorial scope of the GDPR is still unclear. In particular, it is unclear whether the Court of Justice’s broad interpretation of Article 3(1) can result in a situation in which businesses that operate mostly outside the EU have to extend GDPR-style protections to consumers in non-EU countries. However, it seems unlikely that the Court of Justice will interpret the GDPR to cover situations in which the consumer, the business’s headquarters, and the business units involved in the transaction are located outside the EU. For example, in its decision in *Google LLC v. CNIL*, the Court of Justice displayed reluctance to extend the scope of rights established in the GDPR to situations that mostly involved actors in other jurisdictions. C-507/17, *Google LLC v. CNIL*, C:2019:772 (2019), at 53-72.

¹⁰⁹ BRADFORD, *supra* note 10, at 142-3.

¹¹⁰ Cost-Based California Effects are not the only channel through which the EU is said

decades, speculated about the existence of these effects. An early proponent of this idea was Gregory Shaffer, who in a 2000 article predicted that it would “be pragmatically difficult for businesses to employ two sets of data privacy practices, one for EU residents (providing for greater privacy protection) and one for U.S. residents (providing for less).”¹¹¹ In a 2006 book, Jack Goldsmith and Tim Wu described a similar concept as an example of “global laws.”¹¹² Finally, in her work on the “Brussels Effect,” Anu Bradford regularly describes data privacy law as one of the fields in which the EU extends its regulatory reach through Cost-Based California Effects.¹¹³

As these examples show, claims about Cost-Based California Effects in data privacy law precede the entry into force of the GDPR. However, when several prominent online services in 2018 announced the adoption of what they described as GDPR-compliant privacy policies on a global level, proponents of this theory viewed it as additional evidence in favor of Cost-Based California Effects.¹¹⁴ Proponents of this view conjecture that these effects caused many firms to globally adopt GDPR-compliant privacy standards,¹¹⁵ which also implies that EU data privacy law governs the relationship between many U.S. businesses and their customers in the United States.

Today, even commentators who are otherwise skeptical about the EU’s power to impose its data privacy laws on other jurisdictions sometimes

to have changed data practices beyond the territorial scope of the GDPR. According to some, the EU also exerts pressure on other jurisdictions to adopt data privacy laws similar to its own. In particular, the EU reportedly uses the “adequacy procedure” required for non-EU countries to receive general clearance that allows companies to transfer data gathered in the EU into these countries. *See* BRADFORD, *supra* note 10, at 149-150; Christina Lam, *Unsafe Harbor: The European Union’s Demand for Heightened Data Privacy Standards in Schrems v. Irish Data Protection Commissioner*, 40 B.C. INT’L & COMP. L. REV. 10 (2017). But see Schwartz, *supra* note 83. Besides, the obligation to implement strict data privacy standards in the EU could have provided trans-jurisdictional businesses with an incentive to lobby for the introduction of similar standards elsewhere. BRADFORD, *supra* note 10, at 148.

¹¹¹ Shaffer, *supra* note 49, at 78.

¹¹² GOLDSMITH & WU, *supra* note 9, at 173-177.

¹¹³ Bradford, *supra* note 19, at 25 (“Internet companies find it difficult to create different programs for different markets and therefore tend to apply the strictest international standards across the board. At times, it is technologically difficult or impossible to separate data involving European and non-European citizens. Other times it may be feasible but too costly to create special websites or data-processing practices just for the EU.”); Bradford, *supra* note 63, at 164 (“Technical non-divisibility often applies for the regulation of privacy. For example, the EU forces companies like Google to amend their data storage and other business practices to conform to European privacy standards. Facing a technical difficulty to isolate its data collection for the EU, Google is forced to adjust its global operations to the most demanding EU standard.”).

¹¹⁴ BRADFORD, *supra* note 10, at 142-145; Rustad & Koenig, *supra* note 4, at 389-396.

¹¹⁵ BRADFORD, *supra* note 10, at 142-143.

concede the possibility of Cost-Based California Effects.¹¹⁶

At the same time, there are reasons to doubt the widespread existence of Cost-Based California Effects in data privacy law. As I argue above, not every situation in which it is costly to treat consumers in different jurisdictions differently gives rise to such effects.¹¹⁷ Instead, these effects require that costs of differentiation outweigh the added costs of compliance related to treating consumers in low-standard jurisdictions in accordance with the more stringent standards established in a high-standard jurisdiction.¹¹⁸ Given the relative ease with which websites and other online services can be customized for users in different geographies,¹¹⁹ it seems at least questionable whether the costs of differentiation in this area outweigh the added costs of complying with stringent standards like the ones established in the GDPR worldwide.

B. Existing Empirical Evidence

While many in the literature seem to treat the existence of Cost-Based California Effects in data privacy law as a given, there is only limited empirical evidence to prove their existence and scope. Advocates of this theory mostly present anecdotal evidence of major online services that professed to align their global operations with EU privacy law either at the entry into force of the GDPR, or when confronted with EU regulators in other instances.¹²⁰ At the same time, authors like Anu Bradford acknowledge that (Cost-Based) California Effects are not ubiquitous and describe multiple examples of services that decided against extending GDPR-style protections to customers outside the EU. The picture is further complicated by the fact that even some companies whose reactions to the GDPR are often cited as examples of Cost-Based California Effects do not treat customers alike on every dimension.¹²¹

¹¹⁶ Anupam Chander, et al., *Catalyzing Privacy Law*, 42, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3433922&download=yes (“For the most part the GDPR has not had a (de jure) ‘California Effect’ on the U.S. federal government or U.S. states, but it has had a (de facto) “Brussels Effect” on companies operating in U.S. jurisdictions.”); Schwartz, *supra* note 83, at 780 (“Under Bradford’s factors, there is indeed much evidence that suggests a de facto unilateral Brussels Effect for privacy.”).

¹¹⁷ *Supra* Section I.A. 2.

¹¹⁸ *Id.*

¹¹⁹ *See supra* Section I.C.

¹²⁰ BRADFORD, *supra* note 10, at 142-145; Rustad & Koenig, *supra* note 4, at 389-396. *See also* GOLDSMITH & WU, *supra* note 9, at 175-176.

¹²¹ One example of disparate treatment of consumers despite a pledge to harmonize data

Several quantitative studies in different fields have investigated privacy-related changes to websites after the entry into force of the GDPR. While some studies document changes to websites that are likely not subject to the GDPR, none of these studies provide conclusive evidence of widespread Cost-Based California Effects in data privacy law.

A first group of studies analyzes changes in the text of privacy policies following the entry into force of the GDPR. One example is a recent paper by Thomas Linden and co-authors, who document substantial changes to the text of privacy policies obtained from websites they describe as “Global” websites (as opposed to “EU” websites).¹²² However, on its own, this finding is insufficient to show that websites extended GDPR-style privacy protections to consumers outside the EU, and it also does not suggest that online services located outside the EU treated consumers everywhere the same. In another recent paper, Kevin Davis and Florencia Marotta-Wurgler investigate changes in privacy policies in various industries between 2014 and 2018.¹²³ While they document that terms covered by the GDPR became more protective during that period, they also report relatively low absolute rates of compliance and variation across industries that suggests that market forces drive the adoption of more protective standards.¹²⁴

A second group of studies analyzes changes in privacy-relevant website

practices across jurisdictions concerns Facebook’s handling of consent for facial recognition technology. When Facebook introduced its new global data privacy in April 2018, it included a statement saying that it would use facial recognition technology only if users “turned on” this feature. Compl., 153, *United States v. Facebook* (D.D.C. 2019). However, while users in the EU had to opt in to activate this feature, it was automatically turned on for many users in the United States. Compl., 144-56, *United States v. Facebook* (D.D.C. 2019); Leo Kelion, *Facebook seeks facial recognition consent in EU and Canada*, BBC NEWS, Apr. 18, 2018, <https://www.bbc.com/news/technology> [<https://perma.cc/H542-CP35>].

In a similar vein, even if companies in principle treat customers in different jurisdictions alike, customers outside the EU will regularly not be able to rely on the GDPR’s enforcement mechanisms to protect their interests. In this context, it is interesting to note that Facebook restructured its legal relationship with customers in Africa, Asia, Australia, and the Middle East before the entry into force of the GDPR, replacing its European subsidiary with a U.S. entity as the provider of services for customers in these jurisdictions. Some commentators describe the elimination of potential enforcement actions related to the treatment of these customers outside the EU as the main reason for this move. See BRADFORD, *supra* note 10, at 145-6.

¹²² Thomas Linden et al., *The Privacy Policy Landscape After the GDPR* (2019), <https://arxiv.org/abs/1809.08396>.

¹²³ Davis & Marotta-Wurgler, *supra* note 89, at 695-700.

¹²⁴ *Id.*, at 702-3. Other studies focus on changes to privacy policies in the EU. *E.g.*, Martin Degeling et al., *We Value Your Privacy... Now Take Some Cookies: Measuring the GDPR’s Impact on Web Privacy* (2019), <https://arxiv.org/abs/1808.05096>. Because these studies provide little insight into the existence of Cost-Based California Effects, I do not describe them further.

features more broadly. While some document instances in which the GDPR affected transactions between online service providers and consumers outside the EU, others document a differential treatment of consumers in different jurisdictions. One example of the former category is a paper by Christian Peukert and co-authors.¹²⁵ Analyzing changes in the use of third-party services following the entry-into-force of the GDPR, they document changes for websites that target a non-EU audience and are therefore likely not subject to the GDPR.¹²⁶ A study by Adrian Dabrowski and co-authors belongs to the second category.¹²⁷ This study measures websites' use of cookies before and after the entry into force of the GDPR. Crucially, the authors access websites from locations inside and outside the EU,¹²⁸ which allows them to show that a substantial share of websites differentiates between consumers in the EU and elsewhere.¹²⁹

IV. EMPIRICAL ANALYSIS

A. *Research Design*

This Section investigates the existence of Cost-Based California Effects in data privacy law in the U.S. empirically. It uses a novel longitudinal dataset of privacy policies obtained from websites in the U.S. and the EU between 2017 and 2019 to investigate changes to privacy policies that occurred around the entry into force of the GDPR in May 2018.

1. Empirical Approach

In principle, I am interested in learning the extent to which EU law influences the handling of information on U.S. consumers by businesses in the U.S. For various reasons, it is challenging to measure this issue directly. Consequently, my empirical strategy builds on measuring changes in the privacy policies of websites owned by providers of online services in the U.S. around the time of the entry into force of the GDPR.

a. Focus on Privacy Policies

¹²⁵ Peukert et al., *supra* note 25.

¹²⁶ *Id.*, at 11-3.

¹²⁷ Adrian Dabrowski et al., *Measuring Cookies and Web Privacy in a Post-GDPR World*, in *Passive and Active Measurement 258* (David Choffnes & Marinho Barcellos eds., 2019).

¹²⁸ *Id.*, at 262.

¹²⁹ *Id.*, at 264-6.

The main reason to focus on privacy policies is that there are usually only limited opportunities to obtain direct information about the handling of consumer data by businesses. Most of the storing and processing of customer data are hidden from public view. By contrast, privacy policies are available for everyone to inspect on almost all major websites on the internet.¹³⁰ They describe—in varying degrees of detail—what information is stored, when and how it is processed, and when and how it is transferred to servers in other jurisdictions and/or third parties.

There are two main reasons why privacy policies can provide insights into the existence of Cost-Based California Effects in data protection law. First, it seems reasonable to assume that privacy policies are a relatively sound proxy for businesses’ actual handling of consumer data. Notably, a failure to disclose data practices adequately can result in legal consequences. This is true not only in the EU, where such a failure would render the data processing illegal. Although privacy policies are not generally mandated by federal law in the U.S., a failure to comply with a privacy policy can result in enforcement actions by the FTC based on section 5 of the FTC Act.¹³¹

Second, the structure and content of a privacy policy can provide insights into whether businesses attempt to be GDPR-compliant. For online services that fall under the scope of EU privacy law, the GDPR imposes an extensive set of requirements regarding the contents of privacy policies.¹³² These requirements differ markedly from the requirements set up in the GDPR’s predecessor, the Data Protection Directive.¹³³ It seems reasonable to assume that almost all businesses that wanted to comply with the GDPR’s

¹³⁰ In the EU, privacy policies have long been mandatory for all websites that collect their visitors’ information. GDPR, art. 13; Data Protection Directive, art. 10. In the U.S., while federal law does not mandate the universal use of privacy policies, most websites feature a privacy notice either to comply with state law (for example, the California Online Privacy Protection Act of 2003 requires providers of websites that collect personally identifiable information from California residents to include a privacy policy on their website), or because it is required by third-party services whose tools are implemented on a website.

¹³¹ *Supra* Section IV.A. 1. One example of an instance in which a failure to comply with a GDPR-inspired privacy policy in the U.S. led to FTC enforcement actions is Facebook’s handling of consent for its face recognition feature. *Supra* text accompanying notes 180-182. These actions were part of the alleged misconduct that resulted in a USD 5B settlement between the FTC and Facebook in 2019. *See FTC Imposes \$5 Billion Penalty and Sweeping New Privacy Restrictions on Facebook*, FTC, Jul. 24, 2019, <https://www.ftc.gov/news-events/press-releases/2019/07/ftc-imposes-5-billion-penalty-sweeping-new-privacy-restrictions> [<https://perma.cc/T436-G6JC>]. Note that some in the literature have raised doubts about the effectiveness of the FTC’s enforcement actions in the field of data privacy. Marotta-Wurgler & Svirsky, *supra* note 88.

¹³² GDPR, art. 13.

¹³³ Data Protection Directive, art. 10.

requirements had to change their privacy policy before the entry into force of this regulation.

b. Changes Coinciding with the GDPR

As a matter of principle, it is challenging to measure the effects of laws without observing variation over time. This is because, without such variation, it is usually impossible to obtain an estimate for how the same actors observed in the study would behave in the absence of the law. For this reason, I focus on changes to privacy policies around the entry into force of the GDPR. This allows me to compare the state of the world before and after this point in time. Changes in privacy policies that can be attributed to the entry into force of the GDPR suggest that these businesses are—at least de facto—under the influence of EU law.

While this study can, therefore, exploit changes over time, it lacks a second feature that is usually considered an essential prerequisite for measuring causal effects: Because the potential reach of EU privacy law is not confined to a specific set of websites, there exists no untreated control group of privacy policies, i.e., a group of policies that could not have been affected by the GDPR. This implies that it is challenging to attribute any observed changes in privacy policies to the entry into force of the GDPR. After all, other factors might have effected similar changes even in the absence of the GDPR.

My response to this challenge is twofold. First, the dataset’s structure allows me to compare the changes observed during the entry into force of the GDPR to changes during other periods. Therefore, I can determine whether changes similar to the ones observed around the entry into force of the GDPR also occurred in other periods. While such tests cannot completely rule out the possibility that some other factor caused changes observed during the entry into force of the GDPR, they have the potential to render such an alternative explanation unlikely. Second, I investigate the quality of the changes in detail, measuring whether they implement the specific requirements introduced by the GDPR.¹³⁴ While this is no perfect remedy for the lack of a control group, either, a finding that privacy policies conform closely with Article 13 of the GDPR renders alternative explanations rather unlikely.

¹³⁴ See generally Marion Dumas & Jens Frankenreiter, *Text as Observational Data*, in *Law as Data: Computation, Text, and the Future of Legal Analysis* 59 (Daniel Rockmore & Michael A. Livermore eds., 2019) (discussing new opportunities offered by textual data in the exploration of causal processes).

c. Illustrations

To illustrate this approach, consider the following examples. Both Google and Facebook are often cited as examples of companies that offer GDPR-style protections to consumers worldwide.¹³⁵ Already before the entry into force, Google’s and Facebook’s websites displayed (at least essentially) the same privacy policy to users accessing their websites from the EU and the U.S. (including to those customers who accessed country-specific versions of Google). On or shortly before the GDPR’s enactment, both Google and Facebook changed the content of their privacy policies for users everywhere, again offering (at least essentially) the same privacy protections to all users. This observation, while not offering definitive proof that EU law de facto governs the handling of all personal data by Facebook and Google, seems to support the claim that European data privacy law affects the relationship between these two companies and their U.S. customers.¹³⁶

However, not all services are like Google and Facebook. One important counterexample is Amazon. Until May 2018, customers accessing amazon.com from the U.S., amazon.co.uk from the UK, and amazon.de from Germany were shown privacy policies that contained essentially the same information. On May 22, 2018, Amazon changed the privacy policies on its EU websites, but not the privacy policy on its U.S. website. Subsequently, the EU website’s privacy policy differed markedly from the one Amazon used in the U.S. Among other things, the revised privacy policy in the EU suggests that Amazon stopped using e-mail tracking in the EU, and also stopped using location-based services in case a consumer accessed the website using a mobile device. If Amazon adopted these changes to conform with what it perceived as requirements imposed by the GDPR, then the fact that it did not change its U.S. privacy policy suggests that EU law did not influence its relations with consumers based in the United States.¹³⁷

Amazon is not the only service that “forked” its privacy policy.

¹³⁵ E.g., BRADFORD, *supra* note 10, at 143.

¹³⁶ Of course, along the lines of the challenges described in Sections V.A. 2.a and V.A. 2.b above, there are several reasons why this observation does not offer full proof of the proposition that EU law governs the relationship between Google/Facebook and their U.S. customers. First, as we do not know whether Google and Facebook would have adopted similar changes in the absence of the GDPR, we cannot exclude the possibility that the entry into force of the GDPR did not cause the observed change in privacy policies. Second, it seems at least possible that Google and Facebook changed only their privacy policies, but not their handling of personal data.

¹³⁷ Of course, the fact that Amazon did not change the privacy policy on its U.S. website offers no proof that it did not start handling data concerning U.S. customers differently in reaction to the GDPR’s enactment, either. However, there is no apparent reason why Amazon would not change its U.S. privacy policy as well.

Whatsapp, a Facebook subsidiary since 2014, acted similarly. In late April of 2018, it posted a new privacy policy on the German version of its website. In this privacy policy, it addressed at length the rights users enjoyed under the GDPR. By contrast, the U.S. version of Whatsapp’s website did not change its privacy policy between August 2016 and January 2020. As a result, Whatsapp’s U.S. privacy policy did not contain comparably protections.

Other services changed the privacy policy's text for all users, but explicitly limited the rights following from the GDPR to EU citizens. One example of this approach is Envato’s U.S. privacy policy adopted around the entry into force of the GDPR. This policy states that certain rights mandated by the GDPR would only be available to European consumers. The text of the provision is as follows:

If you’re a user or visitor in the European Economic Area these rights also apply to you: . . . [y]ou are also entitled to ask us to port your personal information (i.e. to transfer in a structured, commonly used and machine-readable format, to you), to erase it, or restrict its processing.

Importantly, when websites adopt this type of provision, consumers in the EU and consumers in the U.S. enjoy different protection levels despite being shown identical privacy policies. Furthermore, the existence of this type of provision points to a fundamental limitation of studies that seek to measure the effect of the GDPR on privacy protections in the U.S. solely by documenting changes in the text of privacy policies of U.S. websites. Even if such changes implement requirements of EU law, they might leave the level of protection for U.S. consumers untouched.

d. Research Questions

The examples in the preceding Section illustrate that there is no simple answer to the question of whether the owners of U.S. websites extend EU-style privacy protections to U.S. customers. Some do, others don’t. Against this background, this empirical investigation has three main goals: It seeks to determine, first, how widespread the adoption of GDPR-style protections is. Put very simply, do most U.S. websites resemble Google, Facebook, and other GDPR-compliant websites in their reactions to the entry into force of the GDPR, or do they look more like the U.S. version of Amazon? Second, for those websites that adopt GDPR-style privacy protections, it examines whether these rights are limited to EU consumers. Finally, it explores whether the observed patterns of responses allow for insights into businesses’ motivations to extend GDPR-style privacy protections to consumers in other jurisdictions. I am particularly interested in obtaining evidence about the existence of Cost-Based California effects.

2. Data

In the analysis, I use a longitudinal dataset consisting of the texts of the privacy policies of 695 websites, with one observation per week between late November 2017 and October 2019.¹³⁸ The dataset was assembled in two steps. The core of the dataset (covering 271 websites, with a majority of websites in the EU) consists of privacy policies that were downloaded weekly during that period.¹³⁹ The dataset was amended in January 2020 using snapshots¹⁴⁰ of other websites’ privacy policies obtained from archive.org.

The dataset consists of two parts. The first part contains most of the most frequented websites in the United States (here referred to as *U.S. websites* and *U.S. privacy policies*). This dataset includes most websites that appear in the Top 500 ranking in Alexa’s Top Sites service.¹⁴¹ For various reasons, I exclude some of the websites that appear in this ranking, including all websites operated by online services located in the EU.¹⁴² As a result, this dataset consists of privacy policies for 357 websites. In assembling this part of the dataset, I used additional measures to ensure that the dataset does not mistakenly contain a privacy policy exclusively shown to EU consumers visiting the website.¹⁴³ For this, these websites’ privacy policies were either downloaded from locations within the U.S. or using a VPN client.¹⁴⁴

The second part of the dataset, which serves mostly as a control group, consists of some of the most important websites in the EU (*EU websites* and *EU privacy policies*). The dataset contains all websites among the Alexa Top 500 for the U.K. and Germany that meet one of three conditions: (1) They are operated by services located in the EU, (2) they use a European top-level domain (.de, .uk), or (3) they feature a separate version of the website that is

¹³⁸ More details on the dataset can be found in the Online Appendix (available here: <https://bit.ly/3hUyodM>). See also Frankenreiter & Hermstrüwer, *supra* note 26.

¹³⁹ Websites were downloaded using a python script.

¹⁴⁰ Where available, I obtained weekly snapshots. For some websites, the intervals at which privacy policies are available are considerably longer than that.

¹⁴¹ <https://www.alexa.com/topsites>.

¹⁴² The reason for this last decision is that services located in the EU are under a legal obligation to treat all consumers in line with the provisions in the GDPR. *Supra* Section IV.B. 1. Therefore, the incentives that these sites face in their treatment of U.S. customers are different from those of service providers based in the U.S. Besides, I limited the dataset to websites with privacy policies that are available in English, and excluded websites that use the same privacy policy as another website in the sample.

¹⁴³ In principle, it is possible that consumers in different jurisdictions are being shown different versions of a website. *Supra* Section III. It seems unclear whether online service providers have used this opportunity to display country-specific privacy policies to different customers. See also Peukert et al., *supra* note 25.

¹⁴⁴ NordVPN.

explicitly directed at consumers in the EU, for example, a German version (in the last case, I use the version directed at EU consumers).¹⁴⁵ For EU privacy policies, I took similar steps like the ones described above to ensure that the dataset contains the version of the privacy policy displayed to EU consumers. Overall, the second part of the dataset consists of 278 websites from Germany and 60 websites from the U.K.

The final dataset contains over 60,000 privacy policies, with a structure similar to that of a (balanced) panel data set with $N = 695$ and $T \sim 100$. To make these privacy policies amenable to further analysis, I removed those that were duplicates of the same website’s privacy policy at $T-I$,¹⁴⁶ and used an array of tools to extract the text of the actual privacy.¹⁴⁷ I then inspected all non-duplicate texts manually inspected to ensure that they contained the privacy policy’s actual text.¹⁴⁸

In addition to the texts of the privacy policies, I obtained a range of background variables for all U.S. websites. On the basis of information collected from alexa.com, I obtained a measure for the relative share of users visiting a website from an EU member country (*Pct_EU_Users*). From similarweb.com, I obtained the average number of users per month (in the analysis, I use the logarithmic version of this measure, *Log_Total_Users*) as well as the type of service provided by the website.

Besides, I determine whether a website explicitly targets EU consumers alongside consumers in the U.S. (*EU_target*), or whether there is a separate version of the website available that is directed at EU consumers (*EU_twin*).¹⁴⁹ In the latter case, I also check whether the privacy policy of the EU version of the website is featured among the EU privacy policies in my sample.

¹⁴⁵ Besides, I excluded websites that made no privacy policy available in either English or German.

¹⁴⁶ For this, I used a python script that compared the occurrence of the most frequent words with more than three letters in the text of different privacy policies.

¹⁴⁷ Because custom methods for boilerplate removal such as boiler pipe produced unsatisfying results, I used a custom-made algorithm trained to “predict” the beginning and end of the text of a privacy policy.

¹⁴⁸ The resulting corpus consists of 3,904 texts.

¹⁴⁹ One example of a website targeting EU consumers alongside U.S. consumers is Facebook.com, which is available in German. Amazon is an example of a service offering different versions of its website to consumers in the U.S. and in the EU. For most websites, I obtained information on the service provider from Wikipedia.

B. Analysis and Results

1. Computational Analysis

In this Section, I analyze the development of privacy policies around the time of the entry into force of the GDPR using measures obtained by way of automated text analysis.

a. Outcome Measures

Automated text analysis comprises a range of techniques that make text amenable to quantitative research.¹⁵⁰ Put very simply, these tools convert text into numerical representations without the need for human coders. Because these measures are calculated automatically, I can obtain these measures for every privacy policy in the sample.

In the analysis, I use three different measures.¹⁵¹ First, I calculate the length (measured by the number of words) of each policy in the sample (*num_words*). The second measure captures the amount of text added between two versions of the same privacy policy (*compare_docs*). The measure is based on the distribution of tri-grams in both documents. This measure resembles a simple plagiarism detector, with the difference that I am mostly interested in the parts of the text that were not copied from another source.¹⁵²

Finally, I also include a measure of the use of GDPR-specific vocabulary (*GDPR_vocab*) obtained through topic modeling. Topic modeling is a machine learning technique that can be used to measure the semantic content of documents. In order to do so, topic modeling identifies groups of co-occurring words and groups them into topics. Topic modeling is an

¹⁵⁰ See generally Jens Frankenreiter & Michael A. Livermore, *Computational Methods in Legal Analysis*, 16 ANN. REV. L. SOC. SCI. 39 (2020).

¹⁵¹ Note that there is a plethora of candidate measures available. The choice of these three measures does not necessarily imply that they are better suited than others to illustrate the effects in question. Rather, I decided to use these three measures because they represent three fundamentally different approaches to track changes in privacy policies. In a series of robustness checks, I replicated the analyses in this subsection using a large range of alternative measures. The results of these analyses are not substantially different from the ones presented in this paper. The results of several robustness checks are reported in the Online Appendix.

¹⁵² The measure ranges from 0 to 1, with 1 indicating a privacy policy that was completely rewritten. I obtained the measure using the following steps: (1) Calculate the tri-gram frequency vectors for the earlier and the later document; (2) subtract the vector representing the earlier document from the vector representing the later document; (3) set all negative values to 0; (4) divide the sum of the resulting vector by the sum of the tri-gram vector representing the later document.

unsupervised technique: Contrary to other text analysis tools, it does not require training data. In other words, topic modeling can discern the structure of a corpus of documents without any input that guides its decisions.¹⁵³ To obtain *GDPR_vocab*, I estimate a structural topic model with $K = 39$ topics and calculate the sum of all topics whose average prevalence increased by at least 100% during the entry into force of the GDPR.¹⁵⁴ The measure ranges from 0 to 1, with 0 indicating no use of GDPR-specific language.

b. Results

i. The Development of U.S. Privacy Policies

In a first step, I use the measures described above to obtain a bird’s-eye view of the development of U.S. privacy policies around the entry into force of the GDPR. Remember that the U.S. privacy policies were obtained from websites operated by U.S. and other non-EU service providers.¹⁵⁵ These websites were under no *legal* obligation to apply the GDPR in interactions with U.S. customers,¹⁵⁶ and the policies were obtained in ways that make sure that the dataset only contains policies that were used for these customers. Against this background, if the GDPR only affected interactions that fall under its legal scope, one might expect to see no or only a few unusual changes in the texts of these policies around the time of the entry into force of the GDPR.

I begin by analyzing the timing of changes in the texts of privacy policies. For this, I use the *compare_docs* measure described above.¹⁵⁷ Figure 1 depicts the development of this measure graphically. The figure features two panels, with the upper panel representing U.S. websites and the lower panel EU websites. A dark blue line represents each website in the sample. A measure close to 0 indicates no or only minimal changes between a privacy policy at the dates shown on the x-axis and the same website’s privacy policy seven days before. A measure close to 1 indicates that almost the entire text of the privacy policy was revised. The orange line displays the average

¹⁵³ The output of a topic model consists of two main components. The first component is a set of distributions of topics over documents. Simply said, each document is assigned a numerical vector (whose components add up to 1), indicating the influence of each of the topics on the document. The second part is the topics themselves. Topics are also represented by numerical vectors adding up to 1. In the case of topics, these numerical vectors represent probability distributions over words.

¹⁵⁴ The Online Appendix contains detailed information on the procedure used to determine K and on descriptions of the topics that were included in *GDPR_vocab*.

¹⁵⁵ *Supra* Section IV.A. 2.

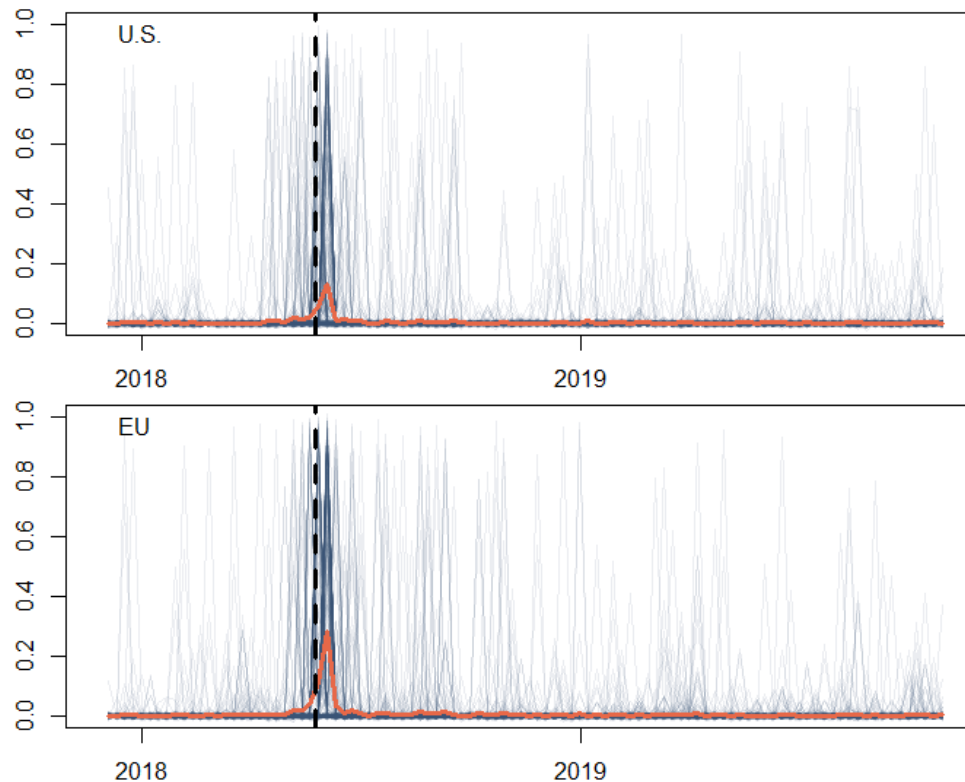
¹⁵⁶ *See supra* Section II.C.

¹⁵⁷ *Supra* Section IV.B. 1.a.

amount of text added across the websites in a jurisdiction. The day of the GDPR’s enactment, May 25, 2018, is marked by a black, dashed vertical line.

The focus here is on the upper panel, which depicts changes for U.S. websites. As this graph shows, privacy policies change at various times in the period under observation. However, there is a flurry of activity around the entry into force of the GDPR. In the two weeks surrounding this event, U.S. websites added an average of almost 20% of new text to their privacy policies. This change is far bigger than any other change observed during the time under observation.

Figure 1: Newly added text per week




Notes. Amount of newly added text to privacy policies by week. Blue lines represent individual websites in the sample. x-axis: date. y-axis: amount of newly added text (variable *compare_docs*). Measures close to 0 indicate limited or no changes. Measures close to 1 indicate a full revision of the privacy policy. Orange line depicts the sample mean. Black dashed line: Date of the entry into force of the GDPR.

However, it merits mention that the reactions observed for different websites in the sample differ considerably. 132 out of 357 websites in the sample (37%) added no new text to their privacy policies between April 2018 and July 2018 (Amazon’s U.S. website is among this group). 177 out of 357

U.S. websites in the sample (49.6%) added 10% or more new text to their privacy policies between April 2018 and July 2018. Changes of a similar magnitude are unusual under normal circumstances; for example, between November 2017 and February 2018, such changes could only be observed for 20 websites (5.6% of the sample). Only 66 websites (18.4% of the sample) changed their privacy policies to the same extent as Google and Facebook, whose privacy policies featured more than 75% newly added text in July 2018.

Figure 2. Length of privacy policies



Notes. Distribution of the length of privacy policies (measured in number of words) at different points in time. x-axis: date. y-axis: number of words in privacy policy. y-axis uses a logarithmic scale. Light grey areas represent the area between the 2.5th and 97.5th percentile. Dark grey areas represent the area between the 25th and 75th percentile. Black line represents the sample median. Black dashed line: Date of the entry into force of the GDPR.

An analysis of the length of privacy policies yields similar results. Figure 2 depicts the distribution of the lengths of all policies in the sample at any given point in time (using a logarithmic scale on the y-axis). The light grey area indicates the range that includes 95% of all websites (excluding the

upper and the lower 2.5% of the distribution). The dark grey area indicates the area between the 25th and the 75th percentile of the distribution (distribution (the range that comprises 50% of websites centered around the median). The black line shows the median length at a given point in time.

As can be seen, the length of privacy policies increased substantially in the weeks around the entry into force of the GDPR. On April 2, 2018, they averaged 3,405 words. By July 2, 2018, they had grown to an average of 3,966 words, an increase of around 16.5% compared to April 2, 2018. The rate of growth spiked around the entry into force of the GDPR. In the two-week period starting on May 21, 2018 (the week of the entry into force of the GDPR), the average length of privacy policies increased by 410 words. This increase is more than 1,000% bigger than any increase observed for any two weeks outside May and June 2018.

The distribution of changes mirrors the one for *compare_docs*. The privacy policies of 135 websites (37.8%) did not increase in length between April and July 2018. 139 websites (38.9%) showed changes of a magnitude that occurred only rarely in other periods. And only 67 websites (18.8%) showed changes in the order of magnitude of Google and Facebook, which increased their privacy policies by more than 1,500 words.

A similar picture also emerges when focusing on the use of GDPR-specific language, *GDPR_vocab*.¹⁵⁸ As shown in Figure 3, such language was almost absent from U.S. privacy policies before the entry into force of the GDPR. For most sites (263, or 73.7% of the sample), such language represented below .5% of the total vocabulary used in privacy policies. This changes after the entry into force of the GDPR. By July 2, 2018, the median privacy policy in the U.S. used around 1% of GDPR specific language. At the same time, 152 privacy policies (among them that of Amazon’s U.S. site) still featured less than .5% of GDPR-specific language.

Overall, these changes seem to suggest that a considerable number of U.S. privacy policies changed in reaction to the enactment of the GDPR. This evidence is a powerful refutation of the hypothesis that the GDPR did not affect online transactions outside its legal scope.


At the same time, these results on their own do not allow for the conclusion that many online services introduced GDPR-compliant privacy standards globally. Rather, it is interesting to observe that the changes observed for most privacy policies (~80%) were a lot less pronounced than those of Google’s and Facebook’s.

The limited nature of the changes observed for U.S. privacy policies is also evident from comparing U.S. privacy policies and EU privacy policies. As Figure 1–3 illustrate, the changes for most U.S. websites were

¹⁵⁸ See *supra* Section IV.B. 1.a.

considerably smaller than those of average EU websites.

Figure 3: Use of GDPR-specific language over time



Notes. Distribution of the use of GDPR-specific language (*GDPR_vocab*) at different points in time. x-axis: date. y-axis: use of GDPR-specific language. y-axis uses a logarithmic scale. Light grey areas represent the area between the 2.5th and 97.5th percentile. Dark grey areas represent the area between the 25th and 75th percentile. Black line represents the sample median. Black dashed line: Date of the entry into force of the GDPR.

Consider first the *compare_docs* measure. The lower panel of Figure 1 represents changes observed for EU websites. In principle, the pattern looks similar to that observed for U.S. websites. However, the changes appear to be of a bigger magnitude than the changes observed for U.S. websites. In fact, the responses observed for Facebook and Google (which were in the top quintile of U.S. websites) seem fairly typical for EU websites. 147 EU websites (43.5%) showed changes of the same magnitude or bigger.

In a similar vein, the changes observed for the numbers of words used in U.S. privacy policies and the use of GDPR-specific language appear modest compared to the changes observed in the EU. As described above, the average length of U.S. privacy policies increased by 16.5% between April 2, 2018, and July 2, 2018. As can be seen from the lower panel of Figure 2, these

changes pale in comparison with the changes observed for EU privacy policies: These privacy policies increased by an average of 1,741 words (65%), with a mean increase of 1,396 words in the two weeks after May 21, 2018, alone. As Figure 3 illustrates, the same picture emerges for changes in the amount of GDPR-specific language.

The differences between U.S. and EU websites persist when one limits the analysis to U.S. websites operated by service providers known to interact with consumers in the EU ($EU_target = 1$ and/or $EU_twin = 1$). For example, while the length of these websites’ privacy policies grew somewhat stronger than that of other U.S. privacy policies (the average growth in length for this subsample was 716 words or 18.2%), this increase is still more than 1,000 words less than that observed for EU websites.

ii. A Quantitative Test of Global Compliance

In a second step, I test more formally whether U.S. privacy policies changed in a way that suggests a GDPR-compliant treatment of U.S. consumers by U.S. online services.

Similar to the informal comparisons between U.S. privacy policies and EU privacy policies in Section IV.B. 1.b.i above, this test uses the changes observed for EU privacy policies as a baseline against which it compares the changes observed for U.S. privacy policies. Remember that the EU websites were under a *legal* obligation to treat consumers in line with the GDPR. Against this background, if Cost-Based California Effects forced U.S. online service providers to treat their U.S. customers in line with the GDPR, one would expect U.S. privacy policies (at least those used by services that also interacted with consumers in the EU) to show patterns of change similar to those of EU privacy policies.

However, unlike the informal comparisons presented above, this test uses only the subsample of U.S. websites that feature a separate website version directed at EU consumers whose privacy policy is included in the EU part of my sample.¹⁵⁹

The reason to restrict the sample in this way is that, on their own, differences between the U.S. and EU samples are not sufficient to conclude that the GDPR affected U.S. and EU online services differently. This is because of the potential role of differences in the two samples. As described above, EU policies were sampled with an eye on ensuring that all websites in this sample were under a legal obligation to treat EU consumers according to EU data privacy law.¹⁶⁰ No emphasis was placed on ensuring that the EU

¹⁵⁹ See *supra* Section IV.A. 2.

¹⁶⁰ *Id.*

websites in the dataset were comparable to the U.S. websites. Against this background, it seems possible that some or even all of the observed differences between U.S. policies and EU policies are not due to general differences in the way U.S. websites and EU websites react to the entry into force of a new law like the GDPR. Instead, these differences could be explained by differences in the characteristics of the websites in both samples. Figuratively speaking, the analysis above might not be comparing apples to apples.

Restricting the sample to services whose EU policies are included in the dataset provides an effective way to tackle these (and related) concerns. This is because the same service providers operate both the U.S. privacy policies in this sample and their EU counterparts. Therefore, any differences in their observed reactions cannot be attributed to unobserved website characteristics. Even more, one can directly observe how many individual service providers implement similar changes to their U.S. and EU privacy policies. If Cost-Based California Effects forced online service providers to treat consumers in all jurisdictions alike, one would expect these changes to be identical.

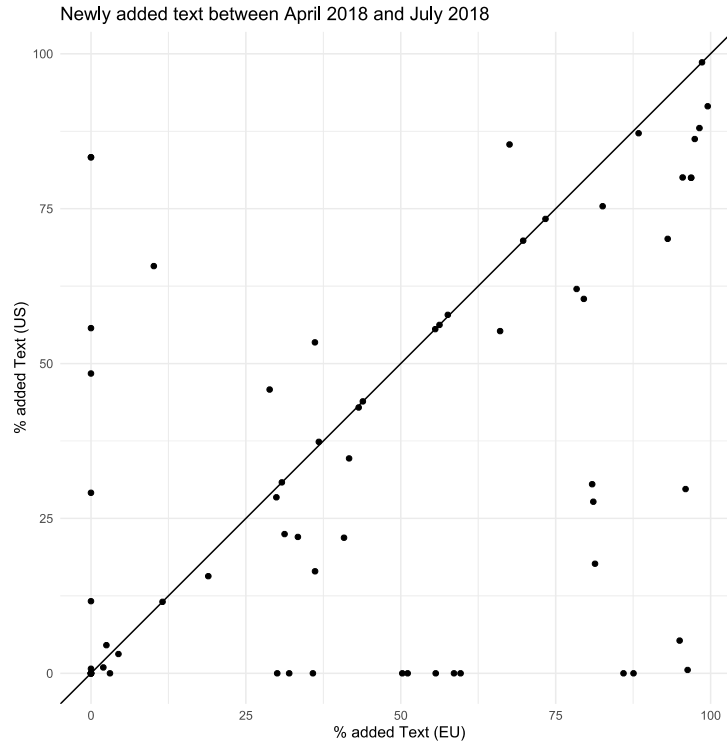
The resulting dataset consists of 68 U.S. privacy policies and their EU counterparts. Examples of online service providers in this sample include Facebook and Amazon, whose reactions to the entry into force of the GDPR I described above.¹⁶¹

Figure 4 reports how the service provider’s policies in the U.S. and the EU changed around the entry into force of the GDPR. More specifically, it reports the amount of newly added text between April 2 and July 2, 2018 (a version of the *compare_docs* measure described above). Each point represents one service provider. The amount of added text to a provider’s U.S. privacy policy is shown on the y-axis. The x-axis depicts how much text the same service provider added to its EU privacy policy. If service providers changed the texts of their U.S. privacy policies in roughly similar ways to the texts of their EU privacy policies, the points would be clustered around the diagonal line running from the bottom left to the upper right corner of the graphic. Instead, the graph shows that a substantial number of providers added considerably more text to their EU privacy policies than to their U.S. privacy policies. Overall, 32 out of 68 websites in the matched sample added at least five percentage points more new text to their EU privacy policy than they did to their U.S. privacy policy. On average, EU privacy policies grew by 11.4 percentage points more than the respective U.S. privacy policies. A paired samples Wilcoxon test indicates that these differences are statistically

¹⁶¹ *Supra* Section IV.A. 1.c.

significant, with a p-value of .0017.¹⁶²

Figure 4: Changes for services with U.S. and EU privacy policies



Notes. Scatterplot showing the amount of newly added text between April 2, 2018, and July 2, 2018 for a service’s EU (x-axis) and U.S. privacy policy (y-axis). Each point represents one service provider.

This result suggests the existence of systematic differences in how U.S. businesses with operations in Europe adjusted the privacy policies of U.S. websites and EU websites in reaction to the entry into force of the GDPR. In other words, it points to the possibility that a sizeable share of U.S. online services with operations in the EU did not follow the example of Google and Facebook in adopting a global privacy policy that extended the rights established in the GDPR to consumers in the U.S. Instead, this effect suggest

¹⁶² Because the distribution of the difference between the added text is not normal, a Wilcoxon test is more appropriate than the standard t-test. A paired t-test yields a p-value of .0097. In the Online Appendix, I replicate this analysis using *num_words* and *GDPR_vocab* as outcome variables. These tests confirm that the existence of systematic differences in how online service providers changed U.S. privacy policies and EU privacy policies in reaction to the entry into force of the GDPR.

that numerous websites might not have granted U.S. consumers the same privacy protections they offered to EU consumers post-GDPR.

2. Manual Coding

While the computational analysis suggests that U.S. websites reacted differently to the entry into force of the GDPR than EU websites, this method is ultimately unable to determine the degree to which U.S. consumers profited from the rights established in the GDPR. One important reason for this is the existence of privacy protections the scope of which is limited to consumers in the EU, a fact that would arguably be missed by most available automated text analysis tools.¹⁶³ Therefore, in this part of the analysis, I analyze the contents of privacy policies using a manually coded subsample.¹⁶⁴

a. Sample Selection and Coding Scheme

The hand-coded sample comprises two privacy policies for each of 246 randomly selected websites in the dataset. The first privacy policy for each website is the one that was in place on April 2, 2018. The second is the one from October 1, 2018. Given the focus of this project, the sampling scheme prioritized U.S. over EU privacy policies. 150 privacy policies are from U.S. websites, 96 from EU websites (82 from Germany, and 14 from U.K.).

Websites were coded according to a coding scheme that attempts to capture whether privacy policies satisfy a range of requirements of the GDPR. As described above, the GDPR contains a set of rather specific requirements that have to be met before businesses can legally obtain consumer data. Among others, businesses have to have a privacy policy that contains a description of the legal bases for gathering data under EU law, and that communicates to the consumer the various rights she has against the business.¹⁶⁵ The coding scheme distills these requirements into nine items that—at least in principle—have to be present to achieve compliance with the GDPR. Seven of the items in coding scheme represent rights that the consumers has against the business; two concern the legal basis for gathering data.

For each of the nine items, three different responses were allowed under the coding instructions: (1) compliance (indicating that the requirement established by the GDPR was met); (2) no compliance (the privacy policy failed to implement the requirement); and (3) compliance limited to EU citizens (the policy contained the provision required by the GDPR, but

¹⁶³ See also *supra* Section IV.A. 1.c.

¹⁶⁴ The hand-coding was done in the context of a parallel project with a co-author at Max Planck Bonn. See Frankenreiter & Hermstrüwer, *supra* note 26.

¹⁶⁵ GDPR, art. 13. See also *supra* Section II.B. .

stipulated that the provision would not apply to U.S. citizens).

The following examples illustrate the use of the coding scheme. The GDPR requires businesses to provide consumers with information about “the existence of the right to request from the controller . . . erasure of personal data.”¹⁶⁶ The coding scheme asks whether websites conform with this requirement. One example of a compliant privacy policy (coded as a “1”) is Airbnb’s U.S. privacy policy adopted in April 2018. The privacy policy contains the following provision:

We generally retain your personal information for as long as is necessary for the performance of the contract between you and us and to comply with our legal obligations. If you no longer want us to use your information to provide the Airbnb Platform to you, you can request that we erase your personal information and close your Airbnb Account.

In January 2019, Airbnb updated its privacy policy. From that point on, the respective paragraph in the privacy policy read as follows:

We generally retain your personal information for as long as is necessary for the performance of the contract between you and us and to comply with our legal obligations. In certain jurisdictions, you can request to have all your personal information deleted entirely.

Compared with the previous provision, this change suggests that AirBnb would reserve the right to reject requests for data deletion if such request were not made by consumers protected by the GDPR. Therefore, this provision would have been coded as a “3” under the coding scheme.

From the coding, I construct a compliance score reporting the number of items for which the privacy policy corresponds with the requirements of the GDPR. The score ranges from 0 to 9, with 9 indicating compliance with all items. For websites in the U.S., I calculate two versions of this score. The first version tracks compliance with the GDPR from the perspective of U.S. customers (*compl_UScust*). In other words, this measure captures whether the policy promises a treatment of U.S. customers that is in line with the requirements of the GDPR. Second, I measure compliance from the perspective of EU customers, who, as described above, might profit from additional rights granted exclusively to them (*compl_EUcust*). For EU websites, which generally do not differentiate between customers from different jurisdictions, I only calculate *compl_EUcust*.

¹⁶⁶ GDPR, art. 13(2)(b).

b. Results

Like in the preceding Section, it can be helpful to contrast the changes of U.S. privacy policies to those of EU privacy policies, which were under a legal obligation to comply with the GDPR. Therefore, I first report results for the latter sample. Figure 5 reports compliance scores (variable *compl_EUcust*) for EU websites before and after the entry into force of the GDPR. The figure indicates that the measure for GDPR-compliance increased dramatically over the six months between April 2018 and October 2018. Only four websites in the sample featured a privacy policy that fulfilled most of the GDPR’s requirements in early April already. In October 2018, by contrast, most EU privacy policies seemed to, by and large, comply with the GDPR: The number of websites that complied with the majority of the requirements of the GDPR had gone up to 73 (76% of the sample).¹⁶⁷

The substantial shift in GDPR compliance can also be illustrated by comparing mean *compl_EUcust* scores before and after the entry into force of the GDPR. In April 2018, EU websites had an average *compl_EUcust* score of 1.57. In October 2018, this score had increased to 6.13.¹⁶⁸ These results suggest that the coding scheme captures changes induced by the GDPR in a meaningful way.

Figure 6 reports GDPR compliance for U.S. websites. Consider first the dark grey bars. These plots report the level of protection that U.S. consumers enjoyed under the respective privacy policy. It can be seen that the level of protection enjoyed by U.S. consumers increased somewhat between April 2018 and October 2018. 44 out of 150 websites increased the level of protection offered to U.S. consumers, while 10 websites reduced the level of protection. These changes are substantial enough that they cannot be explained by chance.¹⁶⁹ Yet only a small minority of websites (10, or 6.7%) complied with more than half of the requirements of the GDPR captured by the coding scheme.

At the same time, it can be seen that the level of protection afforded to EU consumers visiting the same websites (depicted as light grey bars) changed to a much larger degree. In October 2018, a substantial share of privacy policies (52 websites or 34.7% of the sample) had a policy in place that complied with the majority of requirements captured by the coding scheme. This result also suggests that numerous U.S. websites assumed that

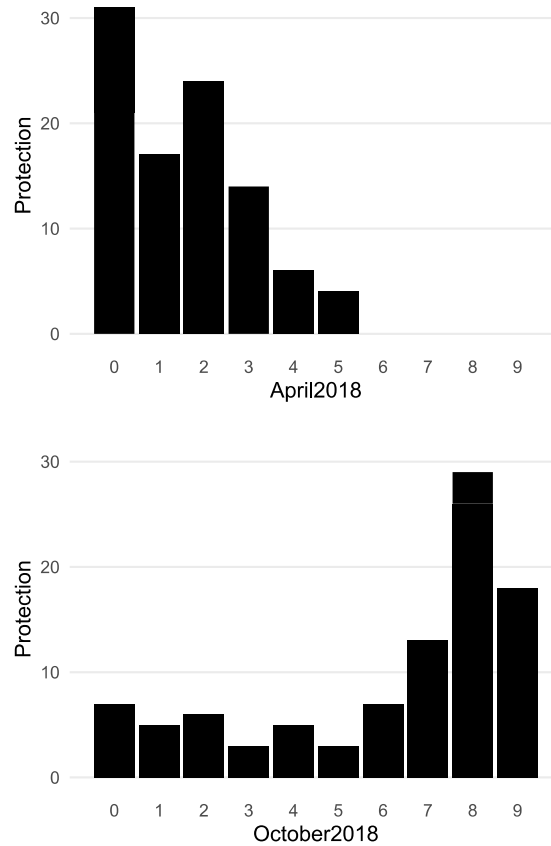
¹⁶⁷ This finding also suggests that the changes to the texts of privacy policies documented above were undertaken to achieve GDPR compliance. Besides, this result can be seen as a confirmation that the coding scheme captures GDPR-compliance in a meaningful way.

¹⁶⁸ These changes are highly significant (a one-sample t-test yields a p-value of $< .0001$).

¹⁶⁹ A t-test yields a p-value of $< .0001$.

they would fall under the scope of the GDPR (at least insofar as they dealt with EU consumers).

Figure 5: Compliance scores for EU websites



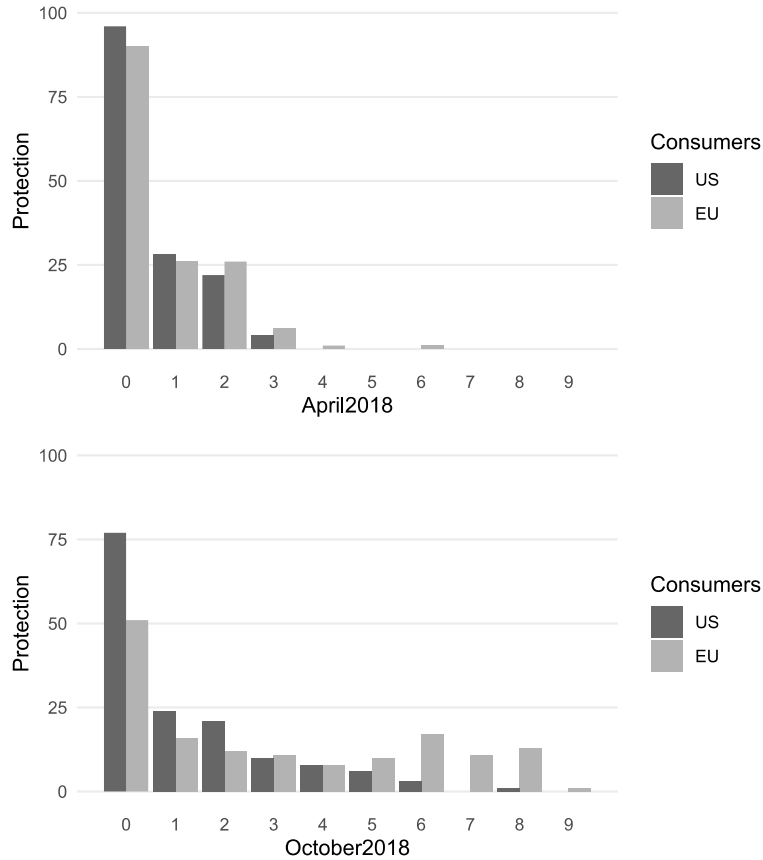
Notes: Histograms depicting the distribution of compliance scores (*compl_EUcust*) for EU privacy policies. The upper panel reports scores for privacy policies in use on April 2, 2018. The lower panel reports scores for privacy policies from October 1, 2018. Values further to the right indicate a higher degree of GDPR-compliance.

Overall, in October 2018, 58 out of 150 U.S. privacy policies established a preferential treatment of EU consumers. The differences of protections granted to EU consumers and U.S. consumers were substantial. The mean *compl_EUcust* score for U.S. websites was 2.98, the mean *compl_UScust* score 1.23.

These differences appear even more pronounced when one focuses on U.S. websites that introduced stronger privacy protections to their privacy policies between April 2018 and October 2018. Of the 70 websites that meet

this criterion, 52 (74.3%) treated EU consumers favorably. The *compl_EUcust* score for these websites averaged 5.66. The mean *compl_UScust* score for the same sample, by comparison, was 2.11.

Figure 6: Compliance scores for U.S. websites



Notes. Histograms depicting the distribution of compliance scores for U.S. privacy policies. Dark grey bars report compliance vis-à-vis U.S. consumers (*compl_UScust*), light grey bars compliance vis-à-vis EU consumers (*compl_EUcust*). The upper panel reports scores for privacy policies in use on April 2, 2018. The lower panel reports scores for privacy policies from October 1, 2018. Values further to the right indicate a higher degree of GDPR-compliance.

These results confirm that most U.S. websites that changed their treatment of EU consumers in the wake of the entry into force of the GDPR did not treat consumers in all jurisdictions identically. By contrast, the majority of websites in this group took active measures that allowed them to differentiate between consumers in different jurisdictions and, in particular, ensure that consumers outside the EU would not profit from the stringent privacy protections introduced in the GDPR.

Furthermore, these results suggest that analyses that rely solely on measuring changes to the texts of U.S. privacy likely overstate the GDPR’s impact on the relationship between U.S. websites and their U.S. customers. Many U.S. websites modified the text of the privacy policies they used in their relationship with U.S. customers. However, many of the changes did not substantially alter the legal status of consumers based in the U.S., but profited exclusively EU consumers.

3. Determinants of Global Compliance

While the results above suggest that only a minority of U.S. websites started offering GDPR-style privacy protections to U.S. consumers after the entry into force of the GDPR, it offers only limited insights into the mechanisms at work. In particular, the analysis does, on its own, not allow for the conclusion that Cost-Based California Effects are absent from data privacy law. Some online services in the sample did extend the protections introduced in the GDPR to consumers in the U.S. Is it possible that these online services faced differentiation costs that were higher than those of other websites?

In this part of the analysis, I shed some light on this question. For this, I use regression analysis to analyze which website characteristics predict the adoption of a (more) GDPR-compliant privacy policy that applies equally to consumers in the U.S. and the EU. My dependent variable is a dummy variable capturing whether a website offered the same privacy protections to consumers in the U.S. and the EU in October 2018. I restrict the sample to all U.S. websites that introduced stronger privacy protections (for any type of consumer) to their privacy policies between April 2018 and October 2018 (N = 70).

In the analysis, I focus on two variables in particular. The first variable is *Pct_EU_Users*, a measure of the share of visitors accessing a website from the EU. As discussed above, if Cost-Based California Effects are at play, organizations that do a lot of business in a high-standard jurisdiction are more likely to apply the rules of this jurisdiction globally than others.¹⁷⁰ Therefore, if costs of differentiation are responsible for the global adoption of GDPR-compliant privacy policies, we would expect the probability of the adoption of such a policy to increase with the share of consumers accessing the website from the EU.

The second variable is *Log_Total_Users*, the average number of monthly visits to the website. Suppose global compliance is mainly due to Cost-Based California Effect, and some of the costs of applying different standards of

¹⁷⁰ See *supra* Section I.A. 2.

protection across jurisdictions are fixed costs. In that case, larger websites should more easily be able to treat consumers in different jurisdictions differently. For example, consider that holding consumer data apart might require the development of systems that document where the data was obtained. For small companies, these investments might not be worth the costs, because the potential benefits from processing the data of consumers from low-standard jurisdictions without constraints are comparably small. By contrast, bigger companies might more easily be able to make this investment. Accordingly, the probability of adopting a uniform GDPR-compliant privacy policy should decrease with the number of visitors to a website.

Figure 7 shows the relationship between these two variables and my dependent variable, the adoption of a uniform privacy policy with a higher level of protection than the one in place before the entry into force of the GDPR. In both panels, each website is represented by a black dot. The y-axis represents the (jittered) dependent variable. The x-axis of the left panel shows the percentage of website visitors from the EU. The right panel displays the total number of website visitors per month.

Several results are immediately apparent. First, none of the variables predicts the global application of EU privacy rights perfectly. For example, one can find both websites with very few and substantial numbers of visitors from the EU among the websites that extend EU-style privacy rights to U.S. customers. Second, both measures are positively correlated with the dependent variable. This result is particularly surprising for the number of website visitors. As described above, if differentiation costs were a major factor in the adoption of globally compliant privacy standards, one would expect to see a higher share of adopters among the smaller websites in the sample. Figure 7 indicates that the opposite is the case. The more visitors a website has, the more likely it is to extend EU-style privacy rights to consumers from other jurisdictions.

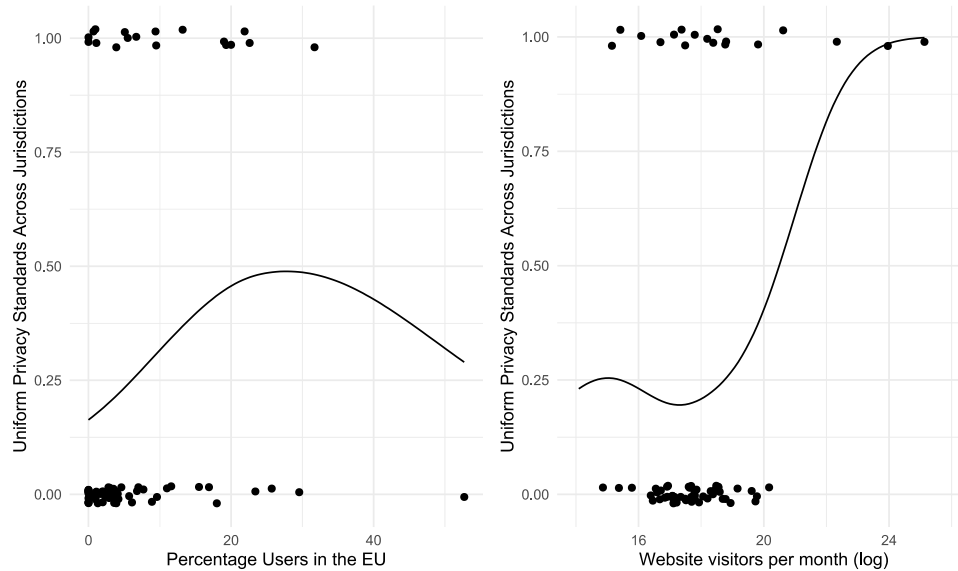
In the following, I use regression analysis to investigate the relationship between these variables more closely. In addition to my variables of interest, I include a categorical variable that captures the industry in which the website is active. I estimate all regressions using both linear probability (OLS) and probit models. I also estimate a probit model that uses a “Heckman correction” to deal with potential concerns about selection effects.¹⁷¹ Table 1 reports results.

It can be seen that the relationship between the share of users in the EU

¹⁷¹ To understand these concerns, recall that I use only those websites that introduced additional privacy protections between April 2018 and October 2018 in the analysis. As a result, the sample used in the analysis does not constitute a random subsample of all policies, giving rise to potentially biased results.

and the dependent variable is not statistically significant. Moreover, its size decreases substantially when additional variables are included in the analysis.¹⁷² By contrast, the relationship between the number of users and the adoption of global privacy standards is significant across specifications and changes comparably little with the inclusion of additional variables.¹⁷³

Figure 7: Website characteristics and uniform privacy standards



Notes. Predictors of the adoption of a uniform privacy policy in October 2018 for all websites that adopted more protective privacy policies between April 2018 and October 2018 ($N = 70$). Websites are represented by black dots. y-axis: dummy variable for whether the privacy policy granted the same rights to U.S. consumers and EU consumers. x-axis: percentage of users in the EU (left panel); logarithmic version of the number of website visitors per month (right panel). Black lines: predicted probabilities obtained from smoothing splines.

These results raise doubts about the importance of differentiation costs in bringing about the global application of the GDPR. There is no evidence of a systematic relationship between the share of users in the EU and the global adoption of a (more) GDPR-compliant privacy policy. Furthermore, contrary to what one would expect if Cost-Based California Effects were at play,

¹⁷² This result persists when various transformations of the variable are used.

¹⁷³ The size of this effect is also substantial. According to the model estimates in Column (9), an average e-commerce website in the baseline category with 10% users in the EU and a number of monthly visitors at the upper end of the first quartile (~22.1M visitors) is predicted to adopt a global, GDPR-compliant privacy policy with a probability of ~23.5%. A similar website with a number of visitors at the upper end of the third quartile (~113M visitors) does so with a probability of ~40.2%.

bigger websites are considerably more likely to treat consumers in different jurisdictions alike than smaller websites.

Table 1: Regression Analysis

| | Dependent variable: Binary Variable indicating global adoption of GDPR-compliant privacy policy | | | | | | | | |
|-----------------------------------|---|------------------|------------------|-------------------|--------------------|-------------------|------------------|-------------------|---------------------------|
| | (1) OLS | (2) OLS | (3) OLS | (4) OLS | (5) Probit | (6) Probit | (7) Probit | (8) Probit | (9) Probit+ Heckman |
| <i>Pct_EU_Users</i> | .009 (.194) | - | .007 (.316) | .003 (.556) | .026 (.154) | - | .021 (.264) | .009 (.601) | .002 (.906) |
| <i>Log_Total_Users</i> | - | .057** (.003) | .051* (.012) | .071** (.001) | - | .209* (.016) | .195* (.030) | .314** (.007) | .290* (.022) |
| <i>Category:</i> | | | | | | | | | |
| <i>Computers & Technology</i> | - | - | - | -.110 (.364) | - | - | - | -.524 (.225) | -.508 (.229) |
| <i>Dating & Adult</i> | - | - | - | .450* (.030) | - | - | - | 1.52* (.013) | 1.36* (.040) |
| <i>E-Commerce</i> | - | - | - | -.036 (.894) | - | - | - | -.225 (.769) | -.153 (.833) |
| <i>Education</i> | - | - | - | -.124 (.548) | - | - | - | -.318 (.680) | -.238 (.741) |
| <i>Entertainment</i> | - | - | - | -.118 (.442) | - | - | - | -.261 (.622) | -.290 (.563) |
| <i>_Intercept</i> | .190** (.004) | -.746* (.027) | -.702* (.046) | -.989** (.006) | -.863*** (.000) | -4.39** (.006) | -4.32* (.010) | -6.31** (.004) | -5.42* (.045) |
| <i>N</i> | 70 | 70 | 70 | 70 | 70 | 70 | 70 | 70 | 70 |

Notes. p-values based on robust standard errors included in parentheses. $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

4. Other Potential Explanations

If the global adoption of GDPR-compliant privacy policies is not primarily driven by differentiation costs, what else explains this phenomenon? The analysis above allows for some preliminary insights into potential alternative explanations.

First and foremost, the results seem to suggest that consumer demand plays a major role in the decision by some services to extend GDPR-style privacy rights to consumers in other jurisdictions. Most importantly, the analysis reveals that websites in the Adult & Dating category are substantially more likely to adopt GDPR-compliant privacy policies on a global level than

other websites.¹⁷⁴ There is little reason to believe that websites in this category face higher differentiation costs than other websites. Instead, as others have argued, it seems reasonable to assume that consumers are more likely to use these services if they trust their privacy protections.¹⁷⁵ Therefore, it seems plausible to assume that these services adopted GDPR-compliant privacy policies to signal high standards of privacy protections to their customers.

The (positive) relationship between the number of visitors to a website and the probability of the global adoption of a GDPR-compliant privacy policy presents a bigger puzzle. One potential explanation also points to consumer demand: Maybe consumers worry more about the treatment of their personal data by organizations they perceive as powerful. If this conjecture is right, the voluntary adoption of more stringent privacy protections might allow these organizations to increase demand for their products, while similar decisions would not entail any increased demand for the products offered by smaller online services.

Alternatively, it also seems possible that major online services adopted GDPR-compliant privacy policies everywhere to deflect regulatory scrutiny by government agencies in other jurisdictions, particularly in the United States. The business practices of companies like Google and Facebook have come under increased public scrutiny in recent years. One hotly debated topic is whether additional privacy protections are needed in the United States to protect consumers in their interactions with these services. Against this background, the decisions by these services and some of their prominent peers to extend GDPR-style privacy protections to consumers in the EU could have been an attempt to convince regulators and the public that such additional regulation is unnecessary.

C. Interpretation and Limitations

Overall, the analysis suggests that the GDPR’s influence on U.S. businesses’ operations outside the EU is limited at best. The privacy policies of a sizeable share of U.S. websites show no attempt to become GDPR-compliant at all. Even among U.S. websites that change their data practices

¹⁷⁴ The analysis suggests that the differences between websites in this group and other websites are substantial. To understand the magnitude of the predicted effect, consider again a website with a 10% share of EU users and an average number of visitors per month of 22.1M. As described above (*supra* note 173), an e-commerce website with these characteristics would be predicted to adopt a GDPR-compliant privacy policy with a probability of ~23.5%. By contrast, a dating website with similar characteristics would adopt such a privacy policy with a probability of 80.6%.

¹⁷⁵ Davis & Marotta-Wurgler, *supra* note 89.

in response to the entry into force of the GDPR, most limit the bulk of privacy protections to customers located in the EU. Besides, the apparent ease with which many businesses differentiate between consumers also raises questions about the sustainability of the commitment of websites that implemented global privacy policies. For example, it seems possible that the introduction of additional privacy protections in the EU could further increase the added costs of global compliance, thereby tipping the balance in favor of differentiation for those websites.¹⁷⁶

At the same time, the results of this study are in important ways limited. The analysis focuses exclusively on protections reflected in the texts of privacy policies, with a particular focus on provisions that endow consumers with enforceable rights *vis-à-vis* the business (for example, the right to request deletion of one’s data).

As a result of this approach, the analysis might miss some ways in which the GDPR increased the privacy protection levels enjoyed by U.S. consumers. This is because it is arguably relatively easy for service providers to restrict rights, such as the right to request data deletion, to consumers in certain jurisdictions. Insofar as the GDPR required businesses to make other changes to their privacy practices, it might have been more costly for them to treat consumers in different jurisdictions differently. In particular, the benefits of changes that require modifications to a website’s structure or design are likely harder to restrict to a subgroup of consumers. One potential example of such a change to a website’s structure concerns the reliance on third-party providers, which reportedly decreased globally following the entry into force of the GDPR.¹⁷⁷ Arguably, it is impossible to measure the full extent of such effects by studying privacy policies.

V. IMPLICATIONS

A. *Normative Implications*

The evidence obtained above indicates that Cost-Based California Effects are less common in data privacy law than is often assumed. Here, I address the question of whether this result is good or bad. This question is directly related to the more general question about the normative desirability of Cost-Based California Effects.¹⁷⁸

¹⁷⁶ The example of AirBnB’s introduction of limits on the rights to request deletion of personal data suggests that this is more than a theoretical possibility. *Supra* Section IV.B. 2.a.

¹⁷⁷ Peukert et al., *supra* note 25.

¹⁷⁸ By contrast, the existence of California Effects does not have direct legal

There are striking differences in how different observers comment on the normative desirability of Cost-Based California Effects. Many view this phenomenon as inherently problematic because other jurisdictions’ rules de facto govern activities taking place in one jurisdiction.¹⁷⁹ By contrast, much of the literature on “California” and “Brussels Effects” paints this phenomenon in a more positive light. For example, Bradford acknowledges that Cost-Based California Effects might undermine “the ability of foreign governments to serve their citizens in accordance with their democratically established preferences.”¹⁸⁰ Nevertheless, she argues that Cost-Based California Effects do not necessarily thwart the democratic process elsewhere, because they might override rules that “are too permissive, too weakly enforced, or otherwise suboptimal.”¹⁸¹

As these different views suggest, there is no easy answer to the question of whether Cost-Based California Effects are normatively desirable or not. As I argue below, the answer to this question ultimately depends on

implications: There is no rule in international law barring jurisdictions from regulating transactions in situations in which their rules give rise to such effects.

While there is no (global) multilateral treaty governing questions of jurisdiction, it is commonly assumed that customary international law imposes some limits on jurisdictions’ powers to regulate activities taking place elsewhere. *See, e.g.*, Restatement (Fourth) of the Foreign Relations Law, §407 (2018). However, these limits are comparably lax. *See also* Goldsmith, *supra* note 11, at 1219 (“In contrast to the domestic interstate context, customary international law imposes few enforceable controls on a country’s assertion of personal jurisdiction, and there are few treaties on the subject.”). Jurisdictions can impose rules on activities as long as there is “a genuine connection between the subject of the regulation and the state seeking to regulate.” Restatement (Fourth) of the Foreign Relations Law, §407 (2018). *See also* Goldsmith, *Unilateral Regulation of the Internet: a Modest Defence*, 11 EUROPEAN J. INT’L. L. 135 (2000), 138 (“It is well accepted today that international law permits a nation to regulate the harmful effects of foreign conduct”). The fact that the same activity also falls under other jurisdictions’ laws does not render the exercise of jurisdiction by the first state illegal, even if laws impose contradictory requirements on actors. *See* Restatement (Fourth) of the Foreign Relations Law, §407 cmt. d (2018). *But see* Restatement (Third) of the Foreign Relations Law, §403 cmt. e (1987).

Further limits on the regulation of commercial activity can follow from areas such as trade law. While the details differ for different kinds of products and different kinds of regulations, trade law focuses on national measures that discriminate against foreign products or services. *See, e.g.*, General Agreement on Tariffs and Trade 1994 (GATT 1994), 1867 U.N.T.S. 187, arts. 1, 2; General Agreement on Trade in Services (GATS), 1869 U.N.T.S. 183, arts. 2, 17. The mere fact that one jurisdiction imposes stricter standards on products or services than other jurisdictions in which the same products or services are sold, by contrast, does usually not constitute a violation of trade law. For a detailed analysis of whether earlier versions of EU privacy law complied with trade law, *see* Shaffer, *supra* note 49, at 46-55.

¹⁷⁹ Goldsmith, *supra* note 11.

¹⁸⁰ BRADFORD, *supra* note 10, at 250.

¹⁸¹ *Id.*, at 251.

assumptions about the capacity of the political process in different jurisdictions to produce rules that conform with the preferences of their citizens (or meet certain objective standards such as efficiency). In short: Those who consider data privacy rules adopted in most jurisdictions as inefficiently lax might lament the absence of Cost-Based California Effects. All others should view this outcome more positively.

Maybe the most important reason to be skeptical about Cost-Based California Effects is their potential to work against some of the most important benefits of decentralized rulemaking. Mandatory laws, like many of the provisions of the GDPR, invariably impose costs on some actors. Members of a population will almost always disagree about whether the benefits of a mandatory rule outweigh its costs. Whenever the preferences of discernible subpopulations differ, it can make sense to implement different rules for these subpopulations that reflect their particular distribution of preferences.¹⁸² Besides, variation in rules can also be valuable because they provide an opportunity to learn about the effects of different types of rules.¹⁸³

In the presence of Cost-Based California Effects, many of these benefits are weakened or disappear altogether. The main reason is that the effects of rules are not limited to the jurisdiction that adopts them. Under these circumstances, decentralization cannot ensure that the rules that effectively apply in a jurisdiction correspond to the local population’s preferences.¹⁸⁴

¹⁸² Revesz, *supra* note 66, at 536. Besides, the costs and benefits of regulation can vary depending on the circumstances under which these laws apply. This can provide another justification for applying different rules to different subpopulations. *Id.*, at 536-537.

Note that, in principle, all these benefits do not depend on the decentralization of political authority. In practice, however, the ability of central authorities to supply different rules to different subpopulations is limited. Maybe the most important reason for this is that central authorities usually do not have the information needed to customize rules to local populations’ preferences. Therefore, it is reasonable to assume that the decentralization of political authority often constitutes a precondition for reaping many benefits of variation in rules.

¹⁸³ Decentralized decision-making might also offer additional benefits. In particular, citizens might have more opportunities to participate in debates about rules that affect them. See Robert O. Keohane, et al., *Democracy-Enhancing Multilateralism*, 63 INT’L ORG. 1 (2009), 8. See also Pascal Langenbach & Franziska Tausch, *Inherited Institutions: Cooperation in the Light of Democratic Legitimacy*, 35 J.L. ECON. & ORG. 364.

¹⁸⁴ See also BRADFORD, *supra* note 10, at 247 (“In particular, many consumers in developing country markets likely view the trade-off between product safety and cost differently than Europeans but are denied these preferences when the Brussels Effect steers companies toward more stringent regulation also in those markets.”). The weakening of the link between preferences about regulation and effective rules is not the only way California Effects counteract the benefits of decentralization. For example, to the extent that this effect results in the dominance of the rules of just one jurisdiction, there is also less opportunity to learn from the effects of different rules in different jurisdictions. Besides,

Potentially even more concerning is the tendency of Cost-Based California Effects to support rules that are systematically biased towards more stringent standards. From the viewpoint of public choice theory, regulatory standards should (and tend to) be chosen so that they lie near the middle of the distribution of preferences of a population.¹⁸⁵ As described above, Cost-Based California Effects compel actors to conform, at a minimum, with the rules which impose the most stringent requirements on the activity at hand.¹⁸⁶ This promotion of comparably extreme rules across jurisdictions appears problematic, because these rules will often be further away from the middle of the distribution than more moderate rules.¹⁸⁷

At the same time, there can also be scenarios in which Cost-Based California Effects appear beneficial.¹⁸⁸ The first scenario concerns fields in which the political process everywhere tends to produce rules that are not sufficiently protective of vulnerable actors. In this case, the inherent tendency of Cost-Based California Effects to promote rules that set high standards of protection can act as a healthy counterweight to biases in the political process. One potential reason for the tendency to undershoot the desirable standard of protection is that the interests of the vulnerable actors’ interests are less concentrated than their counterparties’ interests. In data privacy law, which is mostly concerned with conflicts of interests between consumers and businesses, this is not a far-fetched assumption.

A second, related scenario concerns situations in which the political process in some jurisdictions is biased in a way that leads to inefficiently low protection levels. In these situations, Cost-Based California Effects can promote the trans-jurisdictional application of rules originating in jurisdictions that do not suffer from similar problems.¹⁸⁹

regulatory externalities also strip the population of at least some jurisdictions of the ability to participate in the rule-making process, which might be a source of disutility. *See* Keohane et al., *supra* note 183, at 8. *See also* Langenbach & Tausch, *supra* note 183.

¹⁸⁵ More precisely, under stylized assumptions about the rulemaking process in democracies, regulatory standards will usually be set at the median voter’s preferences.

¹⁸⁶ *Supra* Section I.A. 1.b.

¹⁸⁷ For advocates of theories that view efficiency as the goal of rule-making, selecting outlier rules usually leads to bad outcomes whenever it can be assumed that rulemakers in different jurisdictions all strive to meet that standard, but fail because of uncertainty.

¹⁸⁸ All arguments discussed above implicitly rely on the assumption that individual jurisdictions’ political process is—at least in principle—unbiased. In other words, these arguments assume that, in the absence of California Effects, the rules of a jurisdiction in expectation meet certain standards, either corresponding to a measure of the distribution of preferences in the population (in the case of public choice theory), or converging towards an objective measure such as efficiency. In reality, this assumption is often unwarranted.

¹⁸⁹ In her defense of the “Brussels Effect,” Anu Bradford invokes this latter scenario to justify the EU’s extraterritorial exercise of power. BRADFORD, *supra* note 10, at 250-251

B. Implications for Regulatory Interdependence

Besides these normative implications, the results also have several different implications for the role of traditional national and sub-national governance in a globalizing world. The activities of businesses and similar organizations increasingly transcend jurisdictional boundaries, a reality that poses various challenges to the regulatory power of countries and sub-national jurisdictions. Cost-Based California Effects are one among a range of mechanisms that can contribute to this effect.

Against this background, the finding that Cost-Based California Effects are (at least mostly) absent from data privacy law suggests that nations—even in an age of incessant globalization—retain important areas of autonomy in which global influences are constrained. This finding is particularly noteworthy for at least two reasons. First, data privacy law is an area where the existence of widespread Cost-Based California Effects has often been treated as a given.¹⁹⁰ And second, online services are provided in a context in which traditional jurisdictional boundaries appear particularly porous.

At the same time, the implications of this case study are also limited in important ways. Most importantly, the finding that Cost-Based California Effects are less common in data privacy than expected appears to say little about the prevalence of these effects in other legal areas. The main reason for this is that the costs of differentiation in exchanges involving digital goods or services are likely different from the costs of differentiation in transactions involving physical goods.

Still, this paper’s results offer several lessons for our thinking about Cost-Based California Effects in general. Most importantly, the results provide a powerful confirmation that Cost-Based California Effects cannot be expected in every situation in which trans-jurisdictional actors interface with customers in different jurisdictions, and in which there appear to be potential cost savings from treating them uniformly. Besides, the case study also points to potential pitfalls of using anecdotal evidence to support claims about the existence of Cost-Based California Effects. Almost always, anecdotes will involve companies or other actors that are in some way unusual. Therefore, there is often limited reason to believe that behavior reported in such anecdotes is representative of the behavior of most other actors in the field.

In fact, even in the case of California’s role in promoting higher car emission standards across the U.S., there are indications that Strong

(“The Brussels Effect may . . . have the effect of balancing the alleged overrepresentation of business interests in American public life by empowering consumers”).

¹⁹⁰ BRADFORD, *supra* note 10, at 142-3.

California Effects might be more limited than some suggest.¹⁹¹ For example, according to guidance issued by the California Department of Motor Vehicles in January 2020, “many manufacturers make vehicles . . . with smog equipment that meets federal emission standards, but not California standards.”¹⁹² Heated political battles in other states about the adoption of California-style emission rules similarly suggest that California’s rules are not sufficient to induce all car manufacturers to change their production lines for all of the U.S.¹⁹³ Tellingly, advocates speculated in 2005 about whether the introduction of similar rules in several additional states might stop producers from producing different cars for high-standard and low-standard states.¹⁹⁴

C. Implications for Data Privacy Law

1. Policy Implications

The absence of widespread Cost-Based California Effects in data privacy law also has many important implications for policymakers and privacy advocates in the U.S.

First, the absence of Cost-Based California Effects has important implications for U.S. policymakers' ability to regulate data privacy law in accordance with local preferences. If Cost-Based California Effects compelled most major online service providers to comply with EU law globally, legislative initiatives in data privacy would face important constraints. Whenever a proposed law fell short of the protections of the GDPR, online service providers would have to comply with the latter. As a result, policymakers seeking to adopt a regulatory model different from that of the GDPR could be prevented from doing so, at least insofar as they rely exclusively on national regulatory instruments. In this situation, the most effective way for U.S. policymakers to change the effective standards of

¹⁹¹ As described above, David Vogel’s book *Trading Up*, which coined the term “California Effect,” does not include descriptions of California Effects in the sense as I use the term here. *Supra* note 9. Other descriptions of the mechanisms through which California’s laws impacted outcomes in other states are somewhat ambiguous. *See* Lazer, *supra* note 5, at 477.

¹⁹² *Fast Facts 29. Buying a Vehicle From Out of State — Can You Register It in California?* CAL. DEP’T MOTOR VEHICLES (Jan 2020), <https://www.dmv.ca.gov/portal/file/buying-a-vehicle-from-out-of-state-can-you-register-it-in-california-ffvr-29-pdf/> [https://perma.cc/5DN8-HZHM].

¹⁹³ Danny Hakim, *Battle Lines Set as New York Acts to Cut Emissions*, N.Y. TIMES, Nov. 26, 2005, <https://www.nytimes.com/2005/11/26/nyregion/battle-lines-setas-new-york-acts-to-cut-emissions.html> [https://perma.cc/H3SQ-BFPB].

¹⁹⁴ *Id.*

protection would often be through international negotiations. By contrast, this paper’s findings suggest that policymakers in the U.S. face comparably few external constraints in their pursuit of regulatory strategies in data privacy law.

Second, the findings imply that sustainable changes in data privacy practices in the U.S. will likely only come about due to domestic economic and political forces, not actions in other jurisdictions. It seems hard to imagine a setting in which the data privacy law of another jurisdiction would have had a better chance to influence U.S. businesses’ global data practices than did the GDPR. Apart from the U.S., the EU is commonly regarded as the most potent regulator capable of affecting major businesses’ global operations.¹⁹⁵ The GDPR also has a broad geographical scope, applying to all businesses that target consumers in the EU irrespective of where the former are based. Nevertheless, the analysis shows that the GDPR had only limited effects on the relationship between U.S. online service providers and their customers in the U.S.

Finally, the findings also have potential implications for the impact that legislative and regulative initiatives at the U.S. state level will have on the privacy protections enjoyed by consumers across the U.S. In principle, Cost-Based California Effects could occur at the intrastate level in the U.S., making the most stringent data privacy law in any state the *de facto* law of the land. In fact, when the CCPA entered into force in January 2020, observers predicted that companies would extend CCPA-style protections to all U.S. consumers.¹⁹⁶ However, the costs of differentiating between consumers in different U.S. states would need to be substantial for that to happen. Given the apparent ease with which many businesses differentiate between customers in different countries, it seems at least possible that businesses will also find it worthwhile to treat customers in different states differently.

2. The Role of the EU

Finally, the findings also have implications for our understanding of the EU’s role in data privacy law worldwide. As described above, while observers mostly agree that EU data privacy law has influenced data privacy law on a global level, there is less agreement about the mechanisms behind this effect. Some describe the global impact of EU law as a unilateral exercise of power by the EU.¹⁹⁷ According to these accounts, Cost-Based California

¹⁹⁵ See generally *id.*, at 31-7.

¹⁹⁶ Kashmir Hill, *supra* note 37.

¹⁹⁷ *E.g.*, BRADFORD, *supra* note 10, at 22-26, 132-155; Christopher Kuner, *The Internet*

Effects are one of the primary mechanisms by which the EU asserts its global influence in data privacy law.¹⁹⁸ Others paint a different picture, describing the spread of EU privacy law as a story of “success in the marketplace of regulatory ideas”¹⁹⁹ rather than the result of unilateral action.²⁰⁰

The findings in this paper seem to offer some support for the latter camp’s position. However, they do not offer any evidence about other channels through which the EU could have unilaterally imposed its regulatory model on other nations.

CONCLUSION

Data privacy law is often cited as a prime example of a legal area in which businesses that operate across jurisdictions have to comply with the strictest set of rules everywhere because of an inability to offer differentiated sets of protections to consumers in different jurisdictions. This is one reason why the EU is said to play an outsize role in regulating the data practices of online services worldwide, including in the United States. The results of the analysis in this paper, however, suggest that this form of cross-jurisdictional influence (often referred to as a “California” or “Brussels Effect”) is less widespread than is commonly assumed in the literature. Focusing on changes in the privacy policies of a sample of U.S. websites at the time of the entry into force of the GDPR, the paper documents that most websites do not adjust their policies in a way that would suggest a desire to achieve GDPR-compliance everywhere. This finding has important implications, among other things, for the available regulatory strategies for U.S. legislators and regulators in data privacy law.

and the Global Reach of EU Law (University of Cambridge Legal Studies Research Paper Series 24/2017, 2017), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2890930, at 15-18, 21.

¹⁹⁸ *E.g.*, BRADFORD, *supra* note 10, at 142-147.

¹⁹⁹ Schwartz, *supra* note 83, at 775.

²⁰⁰ Schwartz, *supra* note 79, at; Schwartz, *supra* note 83; Schwartz & Pfeifer, *supra* note 81.