Telecommunications Policy 45 (2021) 102083

Contents lists available at ScienceDirect

Telecommunications Policy

journal homepage: http://www.elsevier.com/locate/telpol

# The impact of the General Data Protection Regulation on internet interconnection

Ran Zhuo [a,*], Bradley Huffaker [b], kc claffy [b], Shane Greenstein [c]

[a] Department of Economics, Harvard University, United States
[b] Center for Applied Internet Data Analysis, University of California San Diego, United States
[c] Harvard Business School, United States

ARTICLE INFO

*JEL classification:*
L86
L00
L51

*Keywords:*
GDPR
Internet
Digital infrastructure
Networking
Interconnection

ABSTRACT

The Internet comprises thousands of independently operated networks, interconnected using bi-laterally negotiated data exchange agreements. The European Union (EU)'s General Data Protection Regulation (GDPR) imposes strict restrictions on handling of personal data of European Economic Area (EEA) residents. A close examination of the text of the law suggests signifi-cant cost to application firms. Available empirical evidence confirms reduction in data usage in the EEA relative to other markets. We investigate whether this decline in derived demand for data exchange impacts EEA networks' decisions to interconnect relative to those of non-EEA OECD networks. Our data consists of a large sample of interconnection agreements between networks globally in 2015–2019. All evidence estimates zero effects: the number of observed agreements, the inferred agreement types, and the number of observed IP-address-level inter-connection points per agreement. We also find economically small effects of the GDPR on the entry and the observed number of customers of networks. We conclude there is no visible short run effects of the GDPR on these measures at the internet layer.

## 1. Introduction

The Internet comprises thousands of independently owned, managed, and operated networks where network operators voluntarily exchange data via bilaterally-negotiated agreements (The Internet Society (2015)). The success of the Internet in creating economic surplus depends on these intercon-nection agreements. Hundreds of billions of dollars in transactions depend on the Internet's opera-tion in the US alone, and these revenues have been growing rapidly.[1] The European Union (EU)'s General Data Protection Regulation (GDPR) serves as a landmark privacy law, regulating the collec-tion, processing, and transfers of consumers' personal data that occur along with these transactions. Since the GDPR's approval in April 2016 and implementation in May 2018, it has inspired a wave of privacy regulation in countries such as Brazil, India, Japan, and the US (Goldberg et al. (2019)). The unprecedented scale and

* Corresponding author.
*E-mail addresses:* rzhuo@g.harvard.edu (R. Zhuo), bradley@caida.org (B. Huffaker), kc@caida.org (claffy), sgreenstein@hbs.edu (S. Greenstein).

[1] From 2012 to 2017, payments for access to wireline forms of Internet access reached $88.7 billion, growing more than 30% in those five years. Payments for access fees to wireless service reached over $90.0 billion, an increase of 57%. In 2017, online advertising contributed $105.9 billion in revenue to the GDP (Gross Domestic Product) among Internet Publishing and Broadcasting and Web Search Portals. That has grown 250% since 2012. The Census Bureau estimates electronic retailing at over $545 billion for just electronic shopping and mail order houses (NAICS 4541), a growth of 65% over the same period.

https://doi.org/10.1016/j.telpol.2020.102083
Received 13 August 2020; Received in revised form 3 December 2020; Accepted 4 December 2020
Available online 28 December 2020

scope of the GDPR makes it the most important privacy policy since the commercialization of the Internet in the 1990s and many hypothesize it would fundamentally change the Internet s operation and the digital economy.

This paper investigates whether the GDPR affected networks decisions to interconnect at the in-ternet layer. Consider this layer as analogous to the postal network. When consumers and providers of online content and services send each other letters containing digital data, networks ( post of-fices ) deliver the mail.[2] The GDPR restricts how and where content and service providers can collect, store, share, and monetize personal information contained in the mail, bringing increased cost and complexity for these application firms, which may in turn impact the demand for mail .

We investigate whether the networks ( post offices ) made fewer interconnections, or changed the types of interconnection agreements, post-GDPR in response to the decline in the demand for mail services. We will provide a precise definition of the internet layer in Section 2 and note features of the internet layer where the postal network analogy works less well.

Our study arrives against a backdrop of growing literature assessing the impacts of privacy regu-lation such as the GDPR on in-vestment in applications. To date there is little research on the impact of these policies at the transport or internet layer of the Internet due to a lack of high quality data at these layers. Comprehensive data on the volume and types of traffic across the Internet does not exist and our various measures capture only part of the networks interconnection activities. Still, our paper represents a rigorous first step towards an objective, data-driven analysis of the effects of the GDPR at the internet layer and adds to the body of research on the impact of online privacy regulation along this important margin.

We hypothesize that the GDPR shapes traffic and investment in connectivity at the internet layer through three different channels. First, we hypothesize that the GDPR operates as a tax on applica-tion firms and lowers the investment in applications. The GDPR raises the operational costs of online businesses that collect personal information as firms need to comply with a stringent set of obliga-tions. The prospective enforcement of rules and the uncertainty about how enforcement operates also raise the expectations of fines and ongoing negotiation. The GDPR s restrictions on collecting and processing personal data could also hamper application firms ability to monetize user partici-pation, which may in turn reduce these firms investment in content production or service provision supported by such monetization. This could lead to lower valuations for entrepreneurial startups in online commerce,[3] a lower supply of ap-plications, and lower traffic.[4]

Second, the GDPR influences user participation.[5] Participation reacts to rules limiting the col-lection, use, storage, and disposal of personal data, and limiting the resale and (re)disclosure of user data. If visitors value these privacy protections, then the GDPR may generate traffic from visits to online sites complying with the GDPR, and those visitors may engage more with the sites.[6] The GDPR simultaneously reduces the value of those visits because it lowers the effectiveness of tar-geted advertising and targeted sales.[7] The GDPR s restrictions on monetizing user participation and its impact on investment in applications supported by such monetization could also result in less user participation over time. These effects operate in opposite directions.

Moreover, we expect some of the GDPR s provisions to directly impact application firms ability to transmit data. The GDPR s requirements on encryption, pseudonymization, and data minimiza-tion could impact the size of data being collected and transmitted.[8] To prevent firms from simply moving personal data to a data haven with fewer restrictions, the GDPR restricts transfers of per-sonal data outside the EEA (European Economic Area). This could reduce data traffic directly.

We expect negative changes in traffic generated by the application layer, if there is any, to lead to a decline in connectivity at the internet layer. This effect operates through changing the bargaining incentives of the networks when the derived demand for data exchange between networks falls. A simple bilateral bargaining model between networks, such as one in Besen et al. (2001), formalizes this intuition.[9]

We then dive into the empirical analysis. Our data comes from various data sources collected by the Center of Applied Internet Data Analysis (CAIDA) at the University of California, San Diego, and represents the state-of-the-art in inferring the presence of inter-connection agreements and their types between networks on the world-wide scale, based on large collections of raw data on global network and IP address level topology of the Internet. Our data includes ownership information of all operating networks around the world, the number of observed agreements per network and the inferred type of each agreement. Using this network level data, we can estimate the number of networks that are customers to a given networks. By combing the topology, we can infer the number of interconnection points between pairs of networks with interconnection agreements on the level of IP addresses, the numerical labels assigned to unique devices connected to the Internet. We collect the datasets used in this paper quarterly, monthly or even daily. Most

[2] We thank Dennis Carlton for suggesting this simple and insightful analogy as an accessible introduction to the internet layer.

[3] Jia et al. (2019, 2020) found a reduction in entrepreneurial ventures and market share of smaller firms after the implementation of the GDPR.

[4] This is similar to Shiller et al. (2018) which found websites with larger proportions of visitors using ad blockers produced less content and had less traffic over time. Goldberg et al. (2019) found evidence of a decline in traffic at some existing firms, while Johnson and Shriver (2019) and Peukert et al. (2020) found evidence of a shift in traffic to the largest firms after the implementation of the GDPR.

[5] Miller and Tucker (2011, 2018) postulated a similar trade-off between privacy, participation and the costs of sup-plying services in the context of medical services.

[6] Empirical evidence on this hypothesis, however, is extremely lacking. Moreover, many provisions of the GDPR are motivated by views that these are intrinsic rights, and do not account for their consequences for online commerce. See Hoofnagle et al. (2019)).

[7] This is similar to Goldfarb and Tucker (2011). Goldberg et al. (2019) and Aridor et al. (2020) also hypothesized this effect.

[8] We thank an anonymous referee for suggesting mechanisms through which the GDPR impacts data flow and inter-connection more directly.

[9] We present the derivations and comparative statics in Appendix A for interested readers.

of the datasets go as far back as to early 2000s and are publicly accessible through CAIDA s website.[10]

Our data, however, is not without limitations, and we explain those limitations thoroughly in Sec-tion 4, where we outline the collection of the data and the construction of the variables.[11] Among the limitations, we highlight that our data only captures part of the networks activities the for-mation and termination of interconnection agreements, and the types of agreements and does not include important variables such as prices, capacity and actual data flow. Comprehensive data on those variables on the scale that we have for agreements has not been available for any academic research.

We begin by presenting descriptives which show persistent and similar growth in Internet inter-connection of EEA countries versus non-EEA OECD (Organization of Economic Cooperation and Development) countries,[12] though the levels of interconnectedness differ. We treat the GDPR s April 2016 approval and May 2018 enforcement as two cutoff dates for periods post policy treatment. We offer several reasons for this assumption and discuss them in detail in Section 5.

We then use a difference-in-differences approach, contrasting interconnection activities by net-works owned by organizations headquartered in the EEA (treatment group) and networks owned by organizations headquartered in other countries (control group) before and after the approval and implementation of the policy. Given the wide territorial scope of the GDPR, we find it important to discuss whether it is ever possible to have a reasonable control group. We motivate our choice of control group in several different ways and discuss them thoroughly in Section 5. We also acknowl-edge the limitations of our empirical approach in the same section.

Contrasting changes in EEA networks interconnection behavior before and after April 2016 and May 2018 relative to non-EEA OECD networks , we estimate zero effects across multiple measures. Networks in the EEA are similar to networks in non-EEA OECD countries in terms of the growth in the number of interconnecting parties and types of agreements reached. Networks affiliation with the EEA also does not affect the observed numbers of IP-address-level interconnection points between each pair of interconnecting networks. We also find economically small effects of the GDPR on the entry and the number of networks that are customers of networks in EEA countries relative to non-EEA OECD countries. Overall, we discover no discernible change in EEA networks interconnecting behavior across the measures we have. In Section 6, we present these results. We discuss several robustness checks to our main results in Section 6.7.

Our paper has an obvious policy implication: even stringent Internet privacy regulation that has evident negative impact at the application layer does not impact the incentive of network operators to interconnect and the short-run growth of interconnectivity. In the conclusion section of the paper, we discuss a number of possible reasons for this result.

Our paper also contributes by presenting data of unprecedented scale and scope.[13] While a theoretical literature tackles questions on network operators and interconnection agreements,[14] em-pirical research in Economics has been scant. Across the academic and policy arena, the lack of well-measured data describing the interconnectivity and traffic flow in the Internet has brought great attention, especially in issues such as net neutrality, international trade in digitally delivered goods, market power of big technology firms, and privacy regulations.[15] Our data may represent a small step towards filling the data gap. We think future works should keep tackling the issue of unmet data needs. Specifically, on the question of how privacy regulations impact the internet layer, if suitable data become available, future works may add additional results with respect to variables such as prices, capacity and the actual levels of data flow.[16]

We organize the rest of the paper as follows. Section 2 provides the background in network interconnection. Section 3 provides the background in the GDPR. Section 4 describes the data, vari-able construction and limitations. Section 5 presents the main regression specification and explains justifications for our empirical strategy. Section 6 presents results across a number of measures of the impact of the GDPR on interconnection. Section 7 concludes.

## 2. Internet interconnection

In Section 1, we use an analogy to the postal network to introduce the internet layer. We note this analogy, though useful, is not

[10] http://www.caida.org/data/overview/. For more information about the data sources used in this paper, please see the data appendix (Appendix B).

[11] We provide additional details about data in Appendix B.

[12] Please see Appendix Table B1 for lists of EEA countries and non-EEA OECD countries.

[13] To the best of our knowledge, the type of data used in our paper has only been used once in prior Economic literature, where D Ignazio and Giovannetti (2009) obtained data from the London Internet Exchange (LINX) of its member networks and one type of agreement (peer-to-peer) between the members. Our data represents a significant improvement from their data as it covers virtually all operating networks in the world, a large number of agreements of both peer-to-peer and provider-to-customer types, and is publicly accessible.

[14] See for examples, Binmore et al. (1986), Besen et al. (2001), Choi et al. (2015) and Laffont et al. (2001).

[15] See discussions in Weller and Woodcock (2013), US International Trade Commission (2014), Meltzer (2014), Nicholson and Giulia (2016) for a few examples.

[16] We note two data sources that offer partial pictures of network capacity and the actual levels of data flow. First, a number of networks self-report capacity associated with their peering agreements at public peering points (In- ternet exchange points) on PeeringDB (https://www.peeringdb.com/ ). Second, Packet Clearing House (PCH) col-lects traffic statistics (peak, average, trough) of a number of Internet exchange points and makes it available at https://www.pch.net/ixp/data. The two data sources do not cover capacity or data flows associated with peering agree-ments at private peering facilities, which handle large volumes of traffic, nor do they cover statistics associated with the other type of agreement, the provider-to-customer agreement. As the scale and scope of the two pieces of data are quite different from those of the measures currently in this paper, we do not include them in our present analysis. We thank an anonymous referee for pointing us to the PCH data for future research.

perfect. We offer a more precise definition of the internet layer in this section. The section explains the technicalities associated with the four layers of the Internet, and describes the demanders and suppliers associated with each layer and the flow of payment. We also discuss the contractual and institutional foundations behind interconnection, and argue networks can respond to policy changes by quickly changing the number or specifications of interconnection agreements.

The Internet was designed with four layers of data exchange in mind: application, transport, in-ternet, and link.[17] Each layer uses a specific set of protocols, shared state, and provides a connection for higher layers. Processes in each layer communicate both with the layer directly above and below, but also across the same layer through connections provided by lower layers. Fig. 1 provides a visual illustration of the four layers and how data exchange takes place between and across each layer.

As shown in Fig. 1, a consumer s personal computer (PC) or smart phone has applications like web browsers and gaming platforms working at the application layer, and an operating system handling the transport, internet, and link layers. Consumers use their applications to connect, us-ing lower layers, to other applications hosted on other devices remotely. The application layer is the layer where personal data is most relevant. A significant share of Internet traffic generated by applications may contain or depend on personal data due to, for examples, user authentication, third-party trackers,[18] product recommendations, bots, improved search results, and spam.[19] Application categories may also have significant heterogeneity regarding the degrees to which they monetize personal data based on their business models. Consumers may pay service and content providers directly, and/or provide their engagement and personal information to these application firms who resell it on to advertisers targeting content, services, and ads. Metrics of engagement include ad views, ad clicks, and purchases resulting from referrals.

When applications connect, data exchange happens between the consumer and content/service providers. Application layer communication relies on lower layers of Internet infrastructure and communication protocols. The transport layer makes sure data from applications arrives correctly and reliably between end point devices. Protocols at this layer break data into packets before handing them off to the internet layer. The internet layer maintains global routing state, routing data packets to their destination address by selecting the next closest router. At this layer, the Internet can be conceptualized as a collection of different networks, each with its own set of routers and routing policies. Routers connect multiple networks and forward data packets destined either for their own networks or other networks. In Fig. 1, the internet layer is visualized to facilitate mov-ing data from the consumer s network to intermediary ISPs (transport networks A and B) then to the service/content network, and the service/content network may send data back the same route. Below the internet layer, the link layer forwards data packets to immediately adjacent (the next hop ) routers.

Some descriptive statistics at the internet layer may help the readers to contextualize this layer in relation to the application layer.[20] By one estimate (Sandvine (2018)), the shares of Internet traffic of different application categories in 2018 were video streaming (58% downstream, 22% upstream), web (17%, 21%), gaming (8%, 3%), social (5%, 4%), marketplace (5%, 2%), file sharing (3%, 22%), messaging (2%, 8%), security (1%, 7%), storage (1%, 9%) and audio streaming (1%, 0%),[21] where video streaming has experienced particularly strong growth in recent years.

In order to reach other networks, individual networks make direct connections with each other, as well as indirect connections through other networks that transport data traffic on their behalf. Consumers and service/content firms pay ISPs to connect their networks to each other. ISPs in turn pay each other where necessary to complete or enhance reachability to the rest of the Internet.

We note an important difference between the internet layer and the postal network.[22] In the postal network, there is a complete separation between the postal service and its user base: the post offices deliver the mail but do not create them. At the internet layer, that is not the case. As some application firms have grown, they began to self-supply network services. Google, Apple and Netflix have followed this expansion path. Conversely, some network firms, such as Comcast and AT&T, have expanded into applications and content. Networks can therefore be operated by consumers and service/content firms Internet Service Providers (ISPs), or service/content firms themselves. Though outside of the scope of this paper, it is an interesting open question how this vertical integration affects the integrated firms responses to privacy regulations.

Network operators typically use a mix of agreements with different interconnection counterpar-ties. As described by the Internet Society (The Internet Society, 2015), we can broadly classify these agreements as one of two types:

Provider-to-customer (p2c) or customer-to-provider (c2p) is an agreement by which the provider network agrees to provide its customer network with connectivity to the rest of the Internet for a fee.

[17] In an official specification document for the Internet regarding requirements for Internet hosts, the Internet Engi-neering Task Force (RFC1122, 1989) describes the four layers and specifies protocols associated with each layer.

[18] A number of papers have found extensive third-party tracking activities associated with websites. Libert (2015) found nearly 90% of top one million websites by Alexa ranking leaked user data to third-parties of which the user was unlikely unaware; more than 60% websites spawned third-party cookies; and more than 80% websites loaded Javascript code from external parties onto users computers. Englehardt and Narayanan (2016) found over 80,000 third-party trackers on the top one million websites. Karaj et al. (2018) found 71% of traffic to 1330 highly visited websites in their data contained tracking.

[19] By a number of estimates ((Symantec (2010); MAAWG (2011); Cisco Talos Intelligence Group (2020); Rao & Reiley (2012)), more than 80% of worldwide email traffic is spam.

[20] We thank an anonymous referee for the suggestion of incorporating these high-level statistics.

[21] The numbers do not sum up to 100 due to rounding.

[22] We thank an anonymous referee for pointing out this important difference and for offering an insightful discussion from which this paragraph draws on.

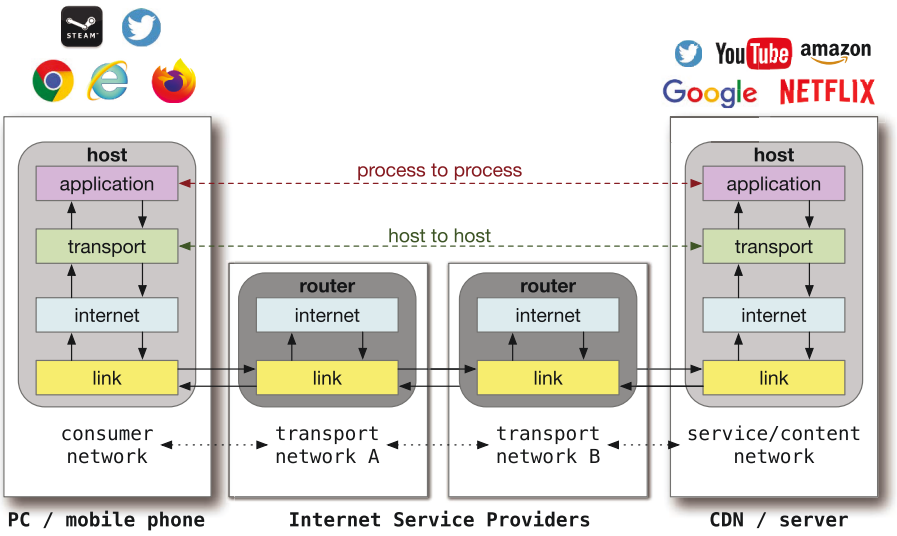


**Fig. 1.** Four layers of the Internet.

- Peer-to-peer (p2p) is an agreement by which two networks agree to a mutual exchange of traffic to and from their customer networks. Peering arrangements reduce the amount of traffic a network must send through its upstream transit provider network, lowering the average cost of traffic delivery. If the peers have similar negotiating power, they form a settlement-free agreement. Under an imbalance, the weaker network pays the other under a paid peering agreement.

Transition provider networks typically price p2c and c2p agreements as a metered service outside of the residential market on a per-megabit-per-second (Mbps) basis. Transit providers compete vig- orously, resulting in a strong declining trend for the prices for transition from 1998 until present.[23] The duration of p2c and c2p agreements can be as short as one month or as long as multiple years. Due to the strong declining trend in prices, networks usually renegotiate even multiple-year agree-ments yearly. Customers have strong incentives to renegotiate or switch to a different provider to bring down the unit transit prices to the current market price. Networks also commonly use ex-tremely short-term agreements, typically with no volume commits and with a duration of just one month, to fully capture the ever-decreasing market prices for transit (Norton, 2014).

p2p agreements may reduce the cost of traffic exchange even further when the volume of traffic is high. Potential peers typically negotiate p2p agreements on a case-by-case basis. Traffic volume of-ten represents "a key determinant" of whether a peering agreement is reached as "the decision hinges upon whether or not there is sufficient value from peering to justify spending time and money" (Norton, 2014). A portion of the cost of peering involves purchasing circuits of fixed capacities between the peers at the peering point and this cost scales with the capacities of the circuits. When a network does not have a Point of Presence (POP) at the agreed peering location, it incurs additional cost to bring its traffic to the peering point. Networks incur additional cost associated with colocation, equipment, and peering ports. The split of the cost is specific to the agreement and net payment between networks may occur, resulting in a paid peering agreement.[24]

In many cases, setting up an interconnection does not require the deployment of additional hard-ware (Norton, 2014), and can be done very quickly. The two parties may simply utilize existing assets, such as configuring an existing port or purchasing circuits between their existing POPs. The process to interconnect can take as little as minutes. When the physical assets for interconnection, such as optical fibers and undersea cables, are not present, it can take substantially longer to install the hardware to interconnect, often in years. As we will discuss in more detail in later sections, our empirical analysis contrasts interconnection activities of networks in the EEA versus networks in non-EEA OECD countries and the vast majority of interconnections that we study involve both par-ties in developed countries. We expect the availability of physical hardware to have little constraint on the incentives to interconnect, at least in the short run that we study, and therefore networks can establish or terminate interconnections reasonably quickly in response to policy changes.

## 3. The GDPR

In this section, we provide an overview of the GDPR. We closely examine the text of the law, discuss where networks fit into the

[23] Estimates based on a sample of US transit providers show that per Mbps transit prices averaged $12.00 in 2008 and averaged $0.63 in 2015 and yearly decreases between 2008 and 2015 ranged from 28% to 52% (Norton, 2014, Table 2–2).

[24] A very rough estimate of the total cost of a p2p interconnection with a 10Gbps capacity at a European peering point using cross-continent transport stands at $11,000 per month in 2014 (Norton, 2014, Table 5–1).

regulatory framework, and the mechanisms through which the law may impact them, presenting existing empirical evidence whenever appropriate. Combining the legal provisions and the empirical evidence with a discussion on the reaction of the popular media to the regulation, we find it easy to hypothesize significant negative impact of the regulation on networks investment, with EEA networks harder hit. We need sound empirical evidence to support or refute this hypothesis.

Approved on April 14, 2016 and effective on May 25, 2018, the GDPR applies to most ap-plication firms and networks with EEA end customers because it applies to any organization that processes personal data of EEA consumers.[25] The GDPR defines personal data broadly, as any information that might identify a consumer ( data subject ).[26] It also defines processing broadly, as any operation that is performed on personal data, whether or not by automated means.[27] The GDPR places the burden of responsibilities on organizations that determine the purposes and the means of processing of personal data (data controllers ), while organizations that process personal data on behalf of controllers (data processors ) also have to comply with a considerable portion of the GDPR.[28] Many application firms fall within the meaning of the GDPR as both data controllers and processors, while networks, routing data on behalf of application firms, fall within the meaning of the GDPR as data processors. The GDPR has a wide territorial scope. Even when a firm has no physical presence in the EEA, the GDPR applies if it is apparent that the firm envisages offering services to consumers located in the EEA.[29]

We hypothesize that the GDPR shapes traffic and investment in connectivity at the internet layer through three different channels. First, the GDPR may operate as a tax on application firms due to compliance costs, regulatory uncertainty and threat of a fine. This may have negative effects on firms ability to invest in applications that generate traffic. Under the GDPR, firms need to fulfill major obligations such as keeping detailed, account-like records of their processing activities,[30] in-corporating protection into the technical design for the services with data protection by design and by default, [31] developing Data Protection Impact Assessments (DPIA) for high-risk pro-cessing activities,[32] and so on (Hoofnagle et al. (2019)). If firms fail to comply, serious violations can trigger administrative fines of up to 20 million euros or up to 4% of the total worldwide an-nual turnover.[33] Regulatory uncertainties further add to the cost. The rules, especially those in the recitals, were written with open-ended features to provide regulators with the flexibility to respond to unexpected and unanticipated issues. Firms therefore need to operate under the presumption of ongoing communications with the Data Protection Authorities over the unresolved features of the rules (Hoofnagle et al. (2019)). Ernest & Young predicts the world s 500 largest companies would spend $ 7.8 billion to comply with the GDPR.[34] While a report by DataGrail, a privacy management platform, estimates 74% of small- and mid-sized organizations would spend more than $100,000 and 20% of them would spent more than $1 million.[35]

The GDPR may impose additional cost on application firms by hampering their ability to mon-etize user participation. This could further reduce application firms investment in applications supported by such monetization and reduce traffic. Under the GDPR, consent must be freely given, specific, informed, unambiguous, and revocable[36] (Hoofnagle et al. (2019)), preventing firms from using long and inaccessible consent processes to obtain personal data. The purpose limitation prin-ciple specifies that personal data shall be collected for specified, explicit and legitimate purposes and not further processed in a manner that is incompatible with those purposes. [37] This would limit application firms ability to repurpose data in unanticipated ways. The data minimization principle specifies that personal data should be limited to what is necessary in relation to the purposes for which they are processed and shall be kept in a form which permits identification of data subjects for no longer than is necessary for those purposes.[38] This accounting comes with increased cost and complexity, and by intent reduces the window for monetization.

[25] The GDPR was incorporated into the EEA Agreement on July 6, 2018, so its scope covers both EU member states and non-EU EEA member states (Iceland, Luxembourg and Norway). The GDPR was enforced in Iceland on July 15, 2018, in Norway on July 20, 2018, and in Luxembourg on August 20, 2018, all within three months of its enforcement in the EU. Iceland, Luxembourg and Norway collectively accounted for 1.4% of EEA population. The nationality of the consumer is not relevant, the relevant criterion is whether the person is located in the EEA (Hoofnagle et al. (2019)).

[26] GDPR Art. 4(1). Under this definition, not only a person s name and physical addresses, but also IP addresses, cookies, and similar data are personal data.

[27] GDPR Art. 4(2).

[28] GDPR Art. 4(7), (8). In principle, if data processors violate the GDPR, the data controller will be considered responsible and liable (Hoofnagle et al. (2019)).

[29] GDPR Recital 23.

[30] GDPR Art. 30.

[31] GDPR Art. 25(1), (2). Data protection by design refers to measures such as pseudonymization, which are designed to implement data-protection principles. Data protection by default means only personal data which are necessary for each specific purpose of the processing are processed and personal data are not made accessible without the individual s intervention to an indefinite number of natural persons.

[32] GDPR Art. 35(1). High-risk processing activities would include automated processing or profiling that leads to decisions that significantly affect people and sensitive data are processed on a large scale (Hoofnagle et al. (2019)).

[33] GDPR Art. 83(4), (5).

[34] Kahn, Jeremy, Stephanie Bodoni & Stefan Nicola. 2018. It ll Cost Billions for Companies to Comply With Europe s New Data Law. Bloomberg Businessweek.

[35] Lindsey, Nicole. 2019. Understanding the GDPR Cost of Continuous Compliance. CPO Magazine.

[36] GDPR Art 7(3).

[37] GDPR Art. 5(1)(b).

[38] GDPR Art. 5(1)(c), (e).

Available evidence from the empirical literature, though limited, largely supports the hypoth-esis that the GDPR is costly for application firms. Godinho de Matos and Adjerid (2019) studied the effectiveness of a campaign for obtaining GDPR-compliant consent for personal marketing and found such practices effective, though at additional cost to the firm to elicit such consent. Goldberg et al. (2019) found a 10% decrease in recorded e-commerce sales for a sample of EU firms after the GDPR s enforcement. Johnson and Shriver (2019) found that the week after the GDPR s enforce-ment, website use of web technology vendor fell by 15%. They also found websites were more likely to drop smaller vendors, which increased the relative concentration of the vendor market by 17%. Peukert et al. (2020) found similar effects and the magnitude of change was particularly large for websites with EU-specific top-level do-mains.[39] Aridor et al. (2020) found a 12.5% drop in observ-able consumers to a data analytics intermediary after the GDPR s enforcement and that resulted in declines in revenue from targeted ads for European travel platforms compared to their non-European counterparts. Lefrere et al. (2020) found the GDPR reduced the number of third-party cookies and tracking. However, they found no evidence that EU websites reduced content production relative to US website. This is in contrast to Shiller et al. (2018), which found websites with larger proportions of visitors using ad blockers produced less content and had less traffic over time.[40]

We also hypothesize that the GPDR influences traffic by changing user participation. Consumers receive many data-related rights under the GDPR, which may boost their participation and incen-tivize application firms to invest more. The GDPR specifies seven rights for consumers: the right 1) to access, 2) to data portability, 3) to rectify data, 4) to stop processing, 5) to object, 6) to erase data, and 7) to resist profiling and computerized decision-making processes[41] (Hoofnagle et al. (2019)). If consumers value these privacy protections, then the GDPR may generate more use of the content and services that comply with the law. Empirical evidence on this, however, is lacking. To the best of our knowledge, we are not aware of academic works that have found consumers actually increase their demand for online content or services in response to better privacy protection. In con-trast, Lefrere et al. (2020) found no effect of the GDPR on the amount of content that EU websites were able to publish, or the degree of average social media engagement and interaction with such content. Goldberg et al. (2019) found a large and significant 10% decline of recorded page views, visits and orders for a set of EU e-commerce firms after the GDPR became effective, suggesting the possibility of less user engagement with less personalized ad targeting and recommendations.[42] Given that the GDPR imposes various costs on application firms and the apparent lack of de-mand response from consumers to better privacy protection, one may expect the negative effects on investment to outweigh the positive, leading to lower overall investment in applications and lower traffic. The effect on investment may be especially pro-nounced for application firms located in the EEA.[43] Empirical evidence, though limited, supports this notion. Jia et al. (2019, 2020) show that the implementation of the GDPR strongly reduced venture capital investment in technol-ogy start-ups in Europe compared to their US counterparts and far away investors were more likely to respond negatively.

In addition to the GDPR s effect on investment in applications and user participation, we further expect some of the GDPR s provisions to directly impact the volume of data application firms are able to transmit. The GDPR requires that firms use encryption or pseudonymization to process personal data (data protection by design ), which may increase data size. In contrast, the GDPR s data minimization requirement could reduce the amount of data being collected and transmitted. Moreover, to prevent data controllers from simply moving personal data to a data haven with fewer or no restrictions, the GDPR only allows transfers of personal data outside the EEA when the destination country or organization upholds privacy protection to a comparable level of that specified in the GDPR.[44] Application firms either bear the significant cost of achieving GDPR-level personal data protection even outside the EEA or choose to reduce the amount of data they transfer outside the EEA.

[39] A top-level domain is the last segment of a domain name. Common top-level domains include.com,.org and.us.

[40] Additional empirical evidence suggests content and service providers alter their behavior significantly following the implementation of the GDPR. Libert et al. (2018) found the GDPR has led to a 22% decrease in third-party cookies on a set of EU news sites (third-party cookies are information stored in browsers used for tracking and advertising, sent from sites other than the one the user is currently visiting). Degeling et al. (2018) and Mohan et al. (2019) found extensive updates to websites and cloud services privacy policies. An exception is Iordanou et al. (2018), which found few changes in the amount of data flow associated with web tracking and in the percentage of this data flow attributed to tracking servers hosted in EU around the GDPR implementation window.

[41] GDPR Art. 4(3), 8, 16, 17, 20, 21(1), 22.

[42] Additional anecdotal evidence suggests consumers feel no better off under the GDPR. See for examples: Olenick, Doug. 2019. Consumers Feel Privacy is No Safer under GDPR. SC Media. Tesseras, Lucy. 2018. GDPR Three Months On: Most Consumers Feel no Better Off. MarketingWeek.

[43] The GDPR makes all EEA consumers costlier to serve, but a non-EEA application firm faces different costs for non-EEA consumers, and may not comply with the GDPR for their non-EEA consumers to lower their cost outside of the EEA. Non-EEA application firms may also choose not to comply, cutting out EEA consumers all together. We discuss in more detail the various evidence for noncompliance of non-EEA application firms and for differential impact of the GDPR at the application layer in Section 5, as support for our empirical strategy.

[44] GDPR Art. 45 47. Adequacy status are evaluated by the European Commission. The US in general does not achieve adequacy. US-based firms may choose to participate the EU-US Privacy Shield, which requires firms to commit to a GDPR-like level of protection.

We expect that a decline in traffic generated by the application layer, if there is any, would lead to a decline in connectivity at the internet layer. This effect operates through changing the networks bargaining incentives and gains from trade when the derived demand for data exchange between networks falls. A simple bilateral bargaining model between networks, such as one in Besen et al. (2001), formalizes this intuition.[45] The hypothesis that the GDPR s effects on the application layer may propagate to the internet layer coincides with the alarmist and negative discussion the popular media and opinion pieces have on the broad impact of the regulation.[46] As of this writing, these views continue to be the consensus. In extensive online search of news articles and editorials since the implementation of the GDPR, we have found no opinion or report to suggest any other impact on business than a costly impact, though views expressed in the news articles and editorials are often neither supported by systematic data collection, nor informed by a census of experience, and most of them stress the costs in unspecific terms. We need sound empirical works to support or refute the hypothesized impact and the uninformed discussions in the media.

## 4. Data

In this section, we describe our data, variable constructions and limitations. Our data comes from various data sources collected and compiled by the Center of Applied Internet Data Analysis (CAIDA) at the University of California, San Diego. Since 1998, CAIDA has been studying interconnectiv-ity of the Internet by actively probing the Internet using its many monitors placed at various van-tage points around the world. Its current flagship active measurement infrastructure, Archipelago, collects interconnectivity data on the IP-address-level from more than 200 monitors located on 6 continents in over 60 countries. CAIDA also collaborates with many organizations and compiles data collected from their monitors. Most notably, it collaborates with the Route Views Project at the University of Oregon and The Re seaux IP Europe ens Network Coordination Centre (RIPE NCC) in Europe to collect BGP routing tables that contain network-level interconnection paths announced across the Internet. Our main data on the network-level interconnection agreements comes from the routing tables, while our lower IP-address-level interconnection points for each agreement come from the active probes (Fig. 2 visualizes the different levels at which we collect data and their re-lationships). CAIDA also gathers records of network registration information from the world s five regional Internet registries (RIRs), allowing us to identify countries or territories of organizations owning individual networks (Center for Applied Internet Data Analysis, 2019b).[47]

Table 1 provides a summary of the variables used in this paper, describing their units of observa-tions, frequency, sources and definitions. Table 2 presents summary statistics of variables described in Table 1. In the remainder of this section we discuss the data collection process and the limitations of the data. For additional information, please refer to the data appendix (Appendix B).

As shown in Table 1, a number of our key variables come from a dataset called *AS Relationships* (Center for Applied Internet Data Analysis, 2019a). The dataset contains network-to-network level interconnection agreements extracted from routing tables contributed by Route Views and RIPE NCC. To correctly route data across the Inter-net, networks exchange routing and reachability information through a protocol called the Border Gateway Protocol (BGP). Each network router using the BGP protocol maintains a routing table. The table contains the connectivity information of the network and its immediate neighbors in the Internet and lists paths to particular network destinations. By placing monitors that peer directly with large networks, we can extract the full set of agreements used between the collecting networks and all visible destinations.

We then annotate the extracted agreements with algorithmically-inferred agreement types, as network operators consider the details of their business relationships as proprietary information and do not generally make them public. Our inference algorithm (Luckie et al., 2013) draws from a long literature of this type of inference including Gao (2001), Subramanian et al. (2002), Di Battista et al. (2003), Erlebach et al. (2002), Xia and Gao (2004), Dimitropoulos et al. (2005) and Dimitropoulos et al. (2007). It achieved over 98% accuracy of agreement type inference via direct validation with a set of network operators (Luckie et al., 2013). The algorithm succeeded in inferring 96% of the agreement types in our sample.

We compute the AS Relationships dataset monthly. We use data from January 2015 to June 2019 for our analysis. We first count the number of observed agreements each network $k$ has in this data. The variable $numAgNtwrk_{kt}$ represents this count. We then aggregate individual agreements to the number of agreements between networks owned by each pair of countries (or territories) $i$ and $j$. The variable $numAg_{i\,jt}$ represents this aggregate count of the number of agreements between country pairs $i\,j$. Breaking down the number of agreements between each country (or territory) pair by their agreement types, we make three variables $numProvAg_{i\,jt}$, $numPeerAg_{i\,jt}$, $numCustAg_{i\,jt}$ for when country (or territory) $i$ s networks are providers to, peers to, and customers of country (or territory) $j$ s networks respectively. We measure a network s centrality in the Internet by its customer cone, a commonly used measure of the number of networks that pay it directly or indirectly for transit. A network s customer cone is defined as itself and all the networks it was observed reaching following provider-to-customer agreements. Networks with larger customer cones have an especially important role in

[45] Appendix A shows the derivations and comparative statics for interested readers.

[46] A sampling from (the most credible) news sources gives a good sense of the range of concerns voiced as GDPR became binding. See for examples: Bershidsky, Leonid. 2018. Europe s Privacy Rules Are Having Unintended Consequences. Bloomberg. Cool, Alison. 2018. Europe s Data Protection Law Is a Big, Confusing Mess. New York Times. Downes, Larry. 2018. GDPR and the End of the Internet s Grand Bargain. Harvard Business Review. Eiss, Robert. 2020. Confusion over Data-Privacy Law Stalls Scientific Progress. Nature. Hern, Alex. 2018. Facebook Moves 1.5bn Users out of Reach of New European Privacy Law. Guardian. Kostov, Nick & Sam Schechner. 2019. GDPR Has Been a Boon for Google and Facebook. Wall Street Journal. Satariano, Adam. 2018. G.D.P.R., a New Privacy Law, Makes Europe World s Leading Tech Watchdog. New York Times. Trentmann, Nina. 2018. Companies Worry That Spending on GDPR May Not Be Over. Wall Street Journal.

[47] We present a complete list of countries and territories in our sample in Appendix Table B1.

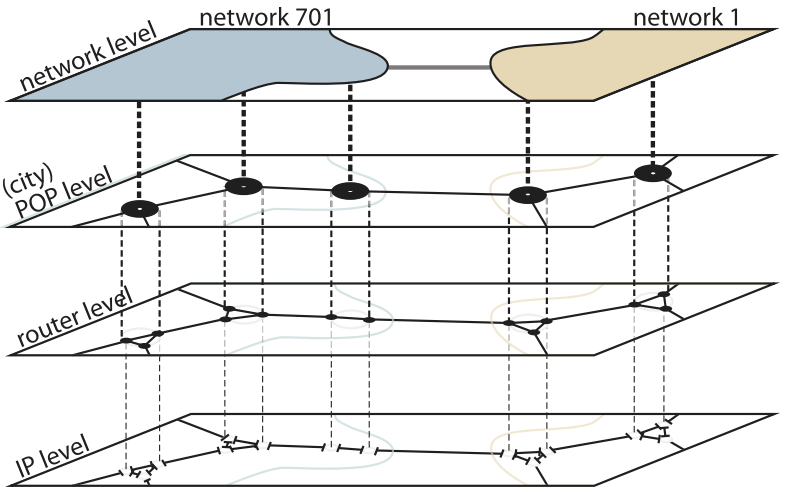


**Fig. 2.** Data collection at the internet layer. Notes: Our network-level interconnection agreements extracted from routing tables correspond to the topmost level in this figure. Our IP-address-level interconnection points for each agreement extracted from active probes correspond to the bottom level in this figure. Geolocating points of presence (PoP) and mapping routers to networks are challenging and open questions, therefore we do not use data on the middle levels.

**Table 1**
Description of variables.

| Variable | Unit of Observation | Frequency | Description | Source | Additional Notes |
|---|---|---|---|---|---|
| $numAg_{ijt}$ | $ctry_i$–$ctry_j$ | monthly | The number of interconnection agreements between pairs of net-works owned by the countries $i$ and $j$. | AS Relationships | Agreements are available on network–network level and are aggregated to country–country level. |
| $numProvAg_{ijt}$ | $ctry_i$–$ctry_j$ | monthly | The number of interconnection agreements where country $i$'s network is a provider to country $j$'s network. | AS Relationships | Same as above. |
| $numPeerAg_{ijt}$ | $ctry_i$–$ctry_j$ | monthly | The number of interconnection agreements where country $i$'s network and country $j$'s network are peers. | AS Relationships | Same as above. |
| $numCustAg_{ijt}$ | $ctry_i$–$ctry_j$ | monthly | The number of interconnection agreements where country $i$'s network is a customer of country $j$'s network. | AS Relationships | Same as above. Value is iden- tical to $numProvAg_{jit}$ for $ctry_j$– $ctry_i$. |
| $numAgIP_{klt}$ | $ntwrk_k$–$ntwrk_l$ | weekly | The number of IP-address-level interconnection points between network $k$ and network $l$, given $k$ and $l$ have an agreement. | IPv4 Prefix-Probing | Data is available daily and is ag-gregated to weekly. |
| $numAgNtwrk_{kt}$ | $ntwrk_k$ | monthly | The number of interconnection agreements network $k$ has. | AS Relationships | |
| $numNtwrk_{it}$ | $ctry_i$ | quarterly | The number of networks country $i$ owns. | AS Organizations | |
| $NtwrkCustCone_{kt}$ | $ntwrk_k$ | monthly | The number of networks net- work $k$ can reach through its cus-tomer connections alone. | AS Relationships | A measure of a network's impor- tance in Internet routing. |

Notes: In the computer science field, a network is referred to as an *Autonomous System* (AS). All data sources listed here are available through CAIDA's webpage http://www.caida.org/data/overview/.

interconnecting the global Internet. The variable $NtwrkCustCone_{kt}$ represents the customer cone size of network $k$.

Our IP-address-level interconnection points within each agreement come from a dataset called *IPv4 Prefix-Probing* (Center for Applied Internet Data Analysis, 2019c). The dataset consists of daily traceroutes from a subset of our Archipelago monitors to every announced BGP routing prefix (a prefix is a block of IP addresses) in the Internet.

Each traceroute tries to reach each destination prefix and records the entire IP address-by-IP address path it takes. We then map each IP address to its network with the help of Route Views Prefix-to-AS mappings dataset (CAIDA, 2013, (Center for Applied Internet Data Analysis, 2019d)) and bdrmapIT tool (Marder et al., 2018), identify IP pairs that form inter-network links and label the observed interconnection links by their IP addresses and network identifiers. The IPv4 Prefix-Probing dataset is available since December 2015 on a daily basis from multiple monitors, so we use data from December 2015 to June 2019. We do two aggregations. First we aggregate

**Table 2**
Summary statistics.

| Variable | Observations | Mean | SD | Min | Max |
|---|---|---|---|---|---|
| *Panel A: unrectangularized variables* | | | | | |
| $numAg_{ijt}$ | 119,071 | 64.4 | 759.0 | 1 | 33,497 |
| $numProvAg_{ijt}$ | 121,369 | 44.8 | 681.8 | 1 | 31,485 |
| $numPeerAg_{ijt}$ | 62,241 | 30.8 | 140.1 | 1 | 4155 |
| $numCustAg_{ijt}$ | 121,369 | 44.8 | 681.8 | 1 | 31,485 |
| $numAgIP_{klt}$ | 19,413,597 | 9.8 | 144.8 | 1 | 172,481 |
| $numAgNtwrk_{kt}$ | 2,909,695 | 5.4 | 55.9 | 1 | 8391 |
| $numNtwrk_{it}$ | 3597 | 357.3 | 1754.3 | 1 | 24,887 |
| $NtwrkCustCone_{kt}$ | 2,909,695 | 7.8 | 263.3 | 1 | 37,061 |
| Panel B: *rectangularized variables* | | | | | |
| $numAg_{ijt}$ | 1,085,400 | 7.1 | 252.2 | 0 | 33,497 |
| $numProvAg_{ijt}$ | 2,160,000 | 2.5 | 161.9 | 0 | 31,485 |
| $numPeerAg_{ijt}$ | 1,085,400 | 1.8 | 34.3 | 0 | 4155 |
| $numCustAg_{ijt}$ | 2,160,000 | 2.5 | 161.9 | 0 | 31,485 |

Notes: Panel A presents the variables with the appropriate levels of aggregation from the raw data. For $numAg_{i\ jt}$, $numProvAg_{i\ jt}$, $numPeerAg_{i\ jt}$, $numCustAg_{i\ jt}$, we also rectangularize the variables by filling in zero values for country pairs and dates with no observed agreements from our raw data and present the rectangularized variables in Panel B.

daily captures from multiple monitors to weekly captures of unique IP-address-to-IP-address connections. Then we aggregate individual connections to the number of connections between each pair of networks *k* and *l*. The variable $numAgIP_{klt}$ represents the number of IP-address-level connections between networks *k* and *l*.

Although we know of no more rigorous data collection efforts of interconnection on the internet layer, we recognize that our data has limitations. First, we note that we are able to capture only part of networks activities the formation and termination of interconnection agreements, and the types of agreements. It is important to note that connectivity is not traffic, though there is evidence that IP address space advertised by BGP tables are strongly positively correlated with networks self-reported traffic volume for a large set of peer-to-peer interconnections (Lodhi et al., 2014). We do not know how much traffic exchange happens across an interconnection or how that traffic has changed over time. If major changes in traffic occurred purely through existing interconnections, causing increased or decreased investment in Internet infrastructure, it would be invisible in our data.[48]

Second, networks owned by organizations headquartered in a particular country or territory can have multiple points of presence (PoP) in many countries and locations within a country and a single Internet interconnection can represent multiple geographically distinct physical connections. Geolocating points of presence is a hard and an open question, so it is important to note the coun-try subscripts of our variables indicate network ownership by organizations headquartered in those countries or territories instead of the exact physical locations of the networks. This measure is es-pecially problematic for large global transit providers and content providers which have PoPs both within and outside the EEA. However, we note that though the relatively few large networks account for a substantial portion of global Internet traffic, the typical network is small and has limited geo-graphic reach beyond its country of origin.[49] Throughout this paper, we use unweighted measures of the number of networks and the number of interconnections. This to some extent alleviates the concern that the imperfect measurement of locations of a few large networks drives the results.

Moreover, the number of agreements we capture, though extremely large, is a subset of all agreements. Individual routers do not maintain a full set of Internet paths, but rather a set of best paths for each destination based on local preferences. Networks also do not announce their peer-to- peer paths to their providers so many peer-to-peer agreements are not observable in the data we use. A truly complete set of agreements would require collecting BGP tables and traceroute data from vantage points in the majority of Internet networks, while our data collection is limited to vantage points where we have our own or partner monitors. Over time, monitors were added at new vantage points, resulting in more visibility in parts of the Internet and hence a greater number of discoverable agreements. To keep visibility consistent throughout our sample periods, we extracted agreements only from a set of monitors that operated throughout our sample periods, January 2015 June 2019 for AS Relationships and December 2015 June 2019 for IPv4 Prefix-Probing.

Similarly, the number of interconnection points we capture is a subset of all interconnection points associated with each agreement. We do not have the ability to target an interconnection directly, but must instead target destinations and infer interconnections from paths that our monitors cross to reach those destinations. We thus miss interconnections not observed by our monitors.

We also note interconnection agreements are more complex than allowed for in our approach. The types of agreements between the same two networks can differ by peering location or even by prefix. Our inference algorithm oversimplifies these cases by assigning a

[48] This can happen when networks add or limit new capacity at existing interconnection points, utilizing the framework of existing agreements and the agreed terms of those agreements, or replacing old agreements by new ones with different terms.

[49] For reference, if we measure the combined value of an organization s users and content purely in terms of the number of IP addresses in its customer cone, an organization at the 95% percentile only accounts for 0.01% of the full routed IP address space, an organization at the 99% percentile accounts for 0.2%, while Amazon.com, Inc. accounts for 1.21%. The distribution of actual traffic across agreements may be even more skewed than the size distribution of the networks measured by customer cones.

single agreement type to each pair of networks (CAIDA, 2015-2019a).

Finally, sometimes technical problems occur with monitors, resulting in changes in visibility of some paths. In October 2018, configuration changes in three RIPENCC partner monitors placed in Amsterdam, Barcelona and Zurich caused permanent disappearance of around 2450 network-to- network interconnections from our sample. We dropped all of the affected interconnections through-out our sample.

## 5. Empirical strategy

In this section, we present our main regression specification and provide justifications for key as-sumptions in our empirical strategy.

Our empirical strategy, in short, constitutes using a difference-in-differences approach to com-pare interconnection activities of networks owned by organizations headquartered in the EEA (treat-ment group) versus networks owned by organizations headquartered in non-EEA OECD countries (control group) before and after the GDPR approval date (April 2016) and implementation date (May 2018). Our main regression specification is as follows:

(1)

where $m$ is the unit of observation of the outcome variable of interest. We take the log of the outcome variable to reflect estimated effects in percentage changes. $m$ can take country pair subscript $ij$, network pair subscript $kl$, country subscript $i$ or network subscript $k$. $POST_{e,mt}$ is an indicator variable equal to 1 if time $t$ is after the GDPR effective date. $POST_{a,mt}$ is an indicator variable equal to 1 if time $t$ is after the GDPR approval date. $EEA_{mt}$ is an indicator variable equal to 1 if the observation $m$ is in the treatment group, and equal to 0 if the observation $m$ is in the control group. A dummy $D_m$ for each unit of observation $m$ and a dummy $D_t$ for each time period $t$ are included. The difference-in-differences effect is identified by the coefficients on the interaction terms $POST_{e,mt}$ $EEA_{mt}$ and $POST_{a,mt}$ $EEA_{mt}$.

The validity of our approach hinges on the validity of our assumptions that any potential policy effect did not set in before the approval of the GDPR and that our control group is reasonable. We therefore focus our discussion on the justifications for these assumptions. We also acknowledge the limitations of our empirical approach.

We offer several justifications for our assumption about the timing of the effect: 1) a robustness check using an alternative cutoff date, 2) conversations with network operators on their decision horizon, and 3) empirical evidence at the application layer supporting stark cutoff dates.

We first discuss a robustness check that uses December 2015 as an alternative cutoff date to study whether networks responded to the law prior to its approval. Examining the timeline of the creation of the law, we think December 2015 is the earliest possible date for firms to respond to the future law. Consultation for the law began as early as 2009 and the European Commission published a proposal text in 2012. In 2013, the European Parliament adopted a compromised text, based on almost 4000 proposed amendments. In 2015, the Council of the European Union published its proposal for the GDPR and started negotiations with the European Parliament. The Parliament and Council reached agreement on the text of the GDPR in December 2015 (Hoofnagle et al. (2019)). Given the intensity of negotiation and the amount of changes the proposal went through, we think it was unlikely for firms to respond before the text of the law was fixed. Most of our variables are available well before December 2015, allowing us to use December 2015 as an alternative cutoff date and test whether networks responded in anticipation of the law. We discuss results from this robustness check in our results section.

Moreover, through conversations with network operators, we learn network operators respond to real-time changes in actual data flows at the application layer, rather than respond to potential changes on the longer time horizon, due to the fact that networks can establish and terminate inter-connection agreements relatively quickly. As such, we think they were unlikely to respond before changes happened at the application layer.

Empirical evidence at the application layer supports the notion that changes at the application layer happened after the GDPR effective date of May 2018. Jia et al. (2019) discuss that, within the two years between the GDPR s approval and effective dates, many organizations chose to roll out their compliance strategy only days and weeks before the effective date. Goldberg et al. (2019) show large declines in page views, visits, orders and revenue from EU consumers at a set of e-commerce sites relative to a control group within four weeks after the policy implementation date.[50] Johnson and Shriver (2019) show a cliff-like decline in websites use of web technology within thirty days after the policy implementation date.[51] Aridor et al. (2020) similarly show an immediate effect on consumer opt-out behavior and firm revenue.[52] We think such evidence of large and quick responses at the application layer helps to justify our choice of stark cutoff dates as well.

Now we move to discuss our choice of control group. Our treatment group consists of net-works owned by organizations headquartered in EEA countries. We choose networks owned by organizations headquartered in non-EEA OECD countries as the control group. We think this is a relevant comparison because networks in developed countries have similar growth rates of inter-connection prior to the GDPR. As we will show in a series of graphs later in the results section, the parallel pre-trends needed for the difference-in-

[50] See their Fig. 2 for details.
[51] See their Fig. 1 for details.
[52] See their Figs. 3 6 for details.

differences approach are visually apparent for the treatment and control groups across all outcome variables of interest. We exclude networks in non-EEA non-OECD countries and territories from the control group for worries that networks in developing countries might behave differently from networks in more developed countries prior to the GDPR. Appendix Table B1 presents complete lists of countries and territories in treatment group, in the control group), and are excluded.

We are well aware of the concern that, given the GDPR s global ambition and wide territorial scope, application firms in non-EEA OECD countries also need to incur substantial cost to comply with the law if they want to serve EEA consumers. This would in turn change the derived demand for data exchange associated with non-EEA OECD networks and bias our results towards zero given our choice of control group. We discuss thoroughly how we address this concern in five ways: 1) a robustness check using a first differences approach, 2) a discussion of the extent of compliance of non-EEA application firms, 3) an organization of our results based on our confidence of the validity of the control group across our eight outcome measures and various subsample breakdowns, 4) interpretation of our results as differential impact and empirical evidence of differential impact at the application layer, and 5) a straightforward acknowledgement of problems with our control group for some of our outcome measures and subsamples.

First, we note that we can perform a robustness check to our main difference-in-differences approach by simply first differencing our outcome variables within EEA subsamples and study-ing whether the approval or the implementation of the GDPR impacted the rate of interconnection growth within EEA countries or networks. This test helps to rule out the scenario under which the GDPR had significant and identical effect on EEA and non-EEA OECD networks. We discuss further the rationale and results of this robustness check in our results section.

We then discuss the extent of compliance to the GDPR among non-EEA application firms. We hypothesize that non-EEA consumers are unlikely to enjoy similar protection as EEA consumers fol-lowing the GDPR, even if these firms choose to comply. These firms may also choose not to comply, cutting out EEA consumers all together. As compliance to the GDPR can be extremely costly, there is incentive for non-EEA application firms to limit compliance to EEA consumers. While some non-EEA application firms allegedly improved privacy protection for non-EEA consumers follow-ing the GDPR, the degree of protection non-EEA consumers enjoyed fell far short from their EEA counterparts.[53] In fact, the California Consumer Privacy Act (CCPA) was strongly motivated by the goal to bring privacy protection of Californian residents on par with that of EEA residents under the GDPR as firms did not voluntarily do so. Peukert et al. (2020) suggest websites catering to non-EU audiences decreased their use of third-party web technology vendors following the GDPR. However, they found the magnitude of change for those websites at 2.2 3.6 percentage points in their most reliable specifications to be far smaller than the magnitude of change for websites catering to EU audience at 7.1 9.3 percentage points.[54] A lot of anecdotal evidence also suggests many content and service providers located outside the EEA simply stopped serving EEA consumers.[55] When EEA consumers were not blocked, they could be offered a very stripped-down version of the content.[56]

Our above discussion helps us to identify the outcomes and subsamples that are less likely to be plagued by the bias towards zero across our eight different outcome measures and various subsample breakdowns. We present our results in Section 6 with this consideration, showing first the measures we are the most confident about. We discuss the rationale for our confidence in those cases in Section 6.

Moreover, we are confident that our estimates, interpreted as the differential effects between EEA networks and non-EEA OECD networks, will be nonzero if the policy has any effect on decisions to interconnect. A number of empirical works have found significant and often large differential effects of privacy regulation at the application layer using a difference-in-differences approach and control groups similar to ours. For example, Goldfarb and Tucker (2011) found the EU s 2002 e- Privacy Directive reduced online display ad effectiveness in the EU relative to other countries. Jia et al. (2019) used EU-based technology ventures as their treatment group and US-based ventures as their primary control group and found EU venture deals declined by as much as 26.1% after the implementation of the GDPR. Jia et al. (2020) found the impact of the GDPR on EU venture investment relative to their US counterparts was larger when ventures and lead investors were not in the same state or union. Aridor et al. (2020) found the GDPR resulted in a 12.5% drop in trackable consumers on European travel platforms as compared to their non-European counterparts and resulted in declines in revenue for the European firms. These differential changes in derived demand for data exchange motivate our expectation of differential changes at the internet layer, if the policy does affect the internet layer.

In addition to all of above, we acknowledge the limitation of our difference-in-differences re-search design that it is not able to

[53] The Associated Press investigated Facebook s claim on global GDPR compliance and found its implementation of many GDPR provisions were vague for non-EEA consumers. While Facebook did not publicize the fact, the Associated Press also found users in six Asian countries did not get the protection through manual checks. See Jesdanun, Anick. 2018. How Google, Facebook will adapt to Europe s New Privacy Law. The Associated Press.

[54] We believe their estimates for websites catering to non-EU audience are overestimates. They explored four different definitions for websites that cater to EU audience : 1) the website has a top-level domain that is specific to a country in the EU (for examples,.de or.fr); 2) the website appears on Alexa s rank for any country in the EU; 3) the website returns content in any of the official languages of member countries of the EU, except English; 4) the website is visited by users in Germany but not users in the US in Nielsen clickstream data. In each of the four cases, websites that cater to non-EU audience are defined as the websites that did not meet the criterion. We note that all four definitions would misclassify a large number of European-based English language sites with common top-level domains such as.com,.org and.net as catering to non-EU audience.

[55] For example, Joseph O Connor, a web developer, compiled an list of 1361 websites (mostly US-based news sites) that blocked visitors from the EU after the GDPR effective date. See O Connor, Joseph. 2018 2019. Websites Not Available in the European Union after GDPR. https://data.verifiedjoseph.com/dataset/websites-not-available-eu-gdpr. As of March 2019, the last time the list was updated, 1129 websites on the list remained blocked.

[56] Sentance, Rebecca. 2018. GDPR: Which Websites are Blocking Visitors from the EU? Econsultancy.

estimate the absolute effect of the GDPR on interconnection decisions at the internet layer. We also acknowledge that the differential impact for some of our subsamples and outcome measures are more likely to tend towards zero while others less so. We discuss this in more detail when we present results by outcome measure and subsample in our next section.

## 6. Results

In this section, we present the regression results specific to each outcome variable. We also discuss results from various robustness checks at the end of this section.

### 6.1. The number of agreements between countries

In this subsection, we study the outcome variable $numAg_{ijt}$, the number of interconnection agree-ments between pairs of networks owned by the countries *i* and *j*. As the unit of observations is a country pair, we need to hold fixed the EEA membership status (or OECD status) of the counter-party of interconnection while we compare the outcomes for EEA countries (treatment group) and for countries in the OECD but not in the EEA (control group).

We therefore construct three subsamples based on counterparties: (a) the counterparties are non- EEA OECD countries, (b) the counterparties are non-EEA non-OECD countries, and (c) the counter-parties are EEA countries. Within each subsample, we then keep only observations where networks or countries are in the EEA (treatment group) or are in the OECD but not in the EEA (control group) and compare their outcomes. Fig. 3 illustrates visually the construction of the three subsamples.

We note that a bias towards zero is less likely to impact regression results for subsamples (a) and (b) than results for subsample (c). The control group of either subsample (a) or subsample (b) does not involve EEA countries. As we discussed in the previous section, we believe non-EEA application layer firms are far less likely to comply with the GDPR in markets outside the EEA or change their behavior in those markets due to the regulation. Their derived demand for data exchange at the internet layer from and to those markets therefore should change little. We are concerned that results for subsample (c) may be biased towards zero as non-EEA application firms need to comply with the GDPR in EEA markets or they may exit those markets, whichever would reduce derived demand for data exchange at the internet layer.

Fig. 4 shows a comparison of the total number of agreements in the EEA countries and in the non-EEA OECD countries, holding fixed the counterparties. We make a few observations. First, despite the differences in levels, EEA countries and non-EEA OECD countries exhibit remarkable parallel trends in setting up agreements with counterparties that are non-EEA OECD countries, non- EEA

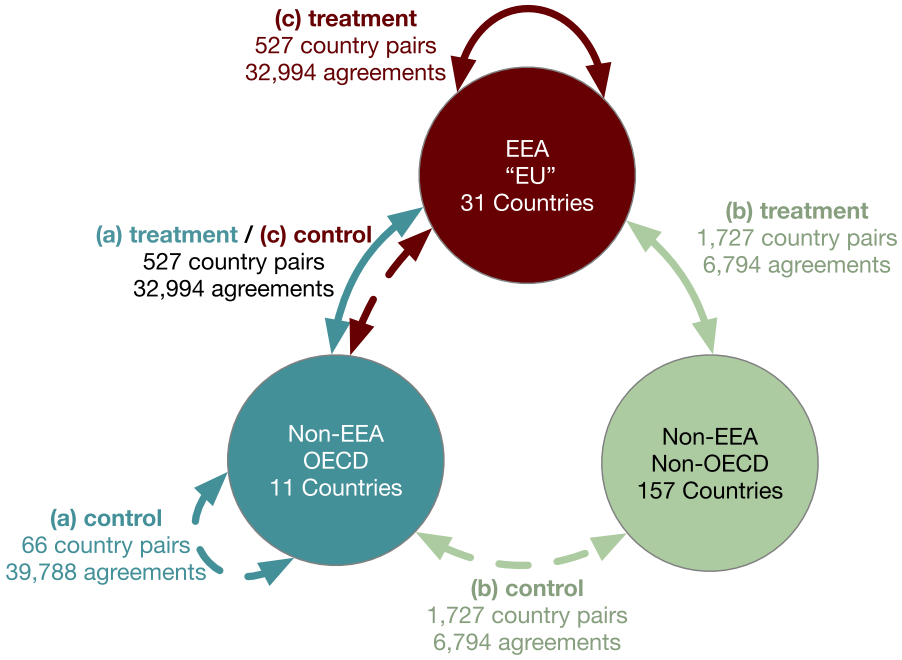


**Fig. 3.** Three subsamples for the analysis on the country pair level. Notes: Interconnections are bidirectional, as represented by the double-headed arrows. We present the total number of agreements in each treatment/control group in March 2016, the month before the GDPR s approval. Subsample (a) fixes non-EEA OECD countries (or territories) as interconnection counterparties. Subsample (b) fixes non-EEA non-OECD countries (or territories) as interconnec- tion counterparties. Subsample (c) fixes EEA countries (or territories) as interconnection counter- parties. Interconnections between EEA countries and non-EEA OECD countries contribute to both subsample (a) and subsample (c). For example, a country pair France  US contributes to both the treatment group in (a) and the control group in (c).

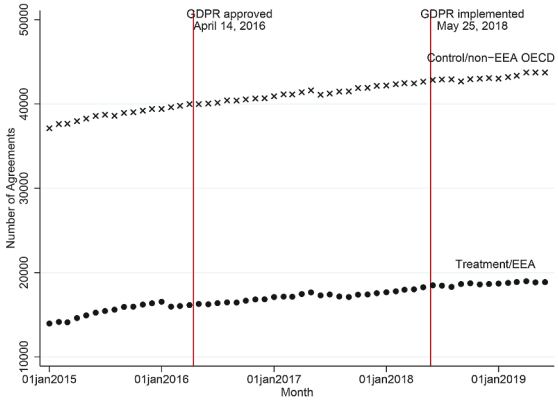


(a) Counterparties are non-EEA OECD countries

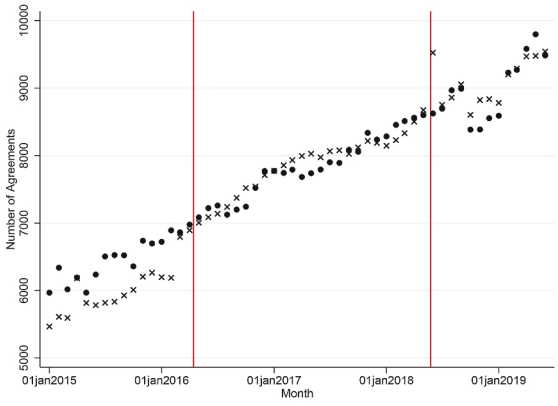


(b) Counterparties are non-EEA non-OECD countries

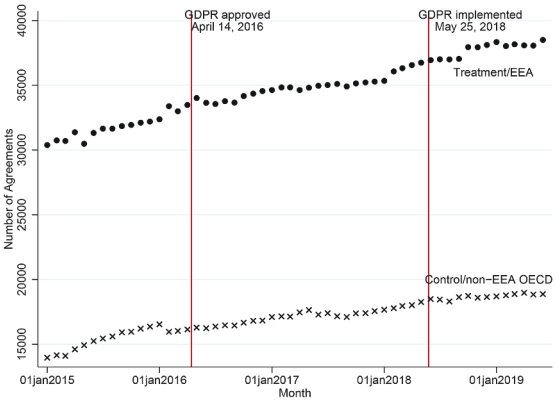


(c) Counterparties are EEA countries

**Fig. 4.** Number of interconnection agreements by EEA and non-EEA OECD countries by counterparty. Notes: The dots represent $\sum_{i\in EEA,t} numAg_{ijt}$, the total number of agreements by networks owned by EEA coun- tries when the counterparties are networks owned by (a) non-EEA OECD countries, (b) non-EEA non-OECD countries, (c) EEA countries. The crosses represent $\sum_{i\in OECD\wedge i\notin EEA,t} numAg_{ijt}$, the total number of agreements by networks owned by non-EEA OECD countries when the counterparties are networks owned by (a) non-EEA OECD countries, (b) non-EEA non-OECD countries, (c) EEA countries. The agreements of networks owned by non- EEA non-OECD countries when the counterparties are also networks owned by non-EEA non-OECD countries are not included in calculating these sums.

**Table 3**
The GDPR s impact on the number of agreements by EEA and non-EEA OECD countries, by counterparty.

| | Non-EEA OECD | Non-EEA Non-OECD | EEA |
|---|---|---|---|
| | (1) | (2) | (3) |
| $POST_e$ $EEA$ | 0.009 | 0.003 | 0.003 |
| | (0.029) | (0.005) | (0.016) |
| $POST_a$ $EEA$ | 0.007 | 0.017*** | 0.011 |
| | (0.024) | (0.005) | (0.017) |
| Group dummies | country pairs | country pairs | country pairs |
| Time dummies | months | months | months |
| Clusters | 418 | 6751 | 880 |
| $_R2$ | 0.991 | 0.948 | 0.987 |
| Observations | 22,572 | 364,554 | 47,520 |

Notes: The dependent variable is $log(numAg_{i\ jt}$ 1). The variable $numAg_{i\ jt}$ is rectangularized as described in Table 2 and we add one when we take the log to account for zero values. $POST_e$ is an indicator variable equal to 1 if the observation is made after the GDPR became effective. $POST_a$ is an indicator variable equal to 1 if the observation is made after the GDPR was approved. Column (1) includes observations when one party is a network owned by an EEA or non-EEA OECD country and the counterparty is a network owned by a non-EEA OECD country. Column (2) includes observations when one party is a network owned by an EEA or non-EEA OECD country and the counterparty is a network owned by a non-EEA non-OECD country. Column (3) includes observations when one party is a network owned by an EEA or non-EEA OECD country and the counterparty is a network owned by an EEA country. All regressions include month dummies and country pair dummies. All regressions cluster standard error by country pair. Standard errors are in parentheses. Significantly different from 0 in a two-tailed test at the *10% level, **5% level, ***1% level.

non-OECD countries, and EEA countries throughout the sample period. Second, agreements with developing countries or territories have a lot more noise in measurement compared to agree-ments within OECD countries or EEA countries.

We then run the regression specification in Equation (1) on each of the three subsamples. The outcome variable is ($numAg_{ijt}$ 1), where we add one to $numAg_{ijt}$, the number of agreements between countries *i* and *j* in month *t*, to account for zero values. The unit of observation *m* is country pair *i j*.

The results are shown in Table 3. The main effect, based on the coefficient on $POST_{e,i\ jt}$ $EEA_{i\ jt}$, is not significantly different from zero across the three subsamples. The only significant result in this table comes from the coefficient on $POST_{a,i\ jt}$ $EEA_{i\ jt}$ for the non-EEA non-OECD counterparty subsample and we test the robustness of this result. Table 3 clusters standard error by country pair. Alternatively, one might expect the interconnection decisions of one particular country to other countries to have correlated errors. This may be especially true for interconnection decisions from an EEA or OECD country to developing countries based on the EEA/OECD networks global interconnection strategy to remote and low demand areas. Therefore, we cluster standard error by EEA and OECD countries in the country pairs for the non-EEA non-OECD counterparty subsample as a robustness test, resulting in 43 clusters as compared to 6751 clusters in Column 2 of Table 3. The coefficient on $POST_{a,i\ jt}$ $EEA_{i\ jt}$ is no longer significant and is therefore likely a spurious result.

### 6.2. The number of agreements between countries by agreement type

In this subsection, we further break down the number of agreements between country pairs to provider-to-customer, peer-to-peer, and customer-to-provider types. As before, we prioritize results for subsamples where interconnection counterparties are non-EEA OECD countries or non-EEA non-OECD countries.

Fig. 5 shows a comparison of the total number of agreements in the EEA countries and in the non-EEA OECD countries, by agreement type. We still observe EEA countries and non-EEA OECD countries have remarkable parallel trends by agreement type throughout the sample period. Based on visual evidence, the GDPR does not have heterogeneous effects on different types of agreements. We then run the regression specification in Equation (1) on each agreement type for each of the three counterparty subsamples. The outcome variables are ($numProvAg_{i\ jt}$ 1), ($numPeerAg_{i\ jt}$ 1) and ($numCustAg_{i\ jt}$ 1). We add one to $numProvAg_{i\ jt}$, $numPeerAg_{i\ jt}$ and $numCustAg_{i\ jt}$, the number of provider-to-customer, peer-to-peer and customer-to-provider agreements between countries *i* and *j* in month *t*, to account for zero values. The unit of observation *m* is country pair *ij*.

We show the results in Table 4. We see a few significant results in the non-EEA non-OECD counterparty subsample. As previously, once we cluster standard error by EEA and OECD countries in the country pairs for the non-EEA non-OECD counterparty subsample as a robustness test, the significance of these results disappears. We also note these results, though sometimes significant, lack systematic patterns and are economically small in magnitude.[57]

[57] To illustrate how economically small the implied effect based on the coefficients is, we take for example the coef-ficient 0.038 on $POST_e$ *EEA* from column (8) of Table 4, the largest significant result in the table. The dependent variable for the regression in column (8) is $log(numPeerAg_{i\ jt}$ 1). It has a mean of 0.109 and an SD of 0.544. Therefore, being in the treatment group post GDPR effective date has an effect which is a tiny fraction of one standard deviation of the outcome.

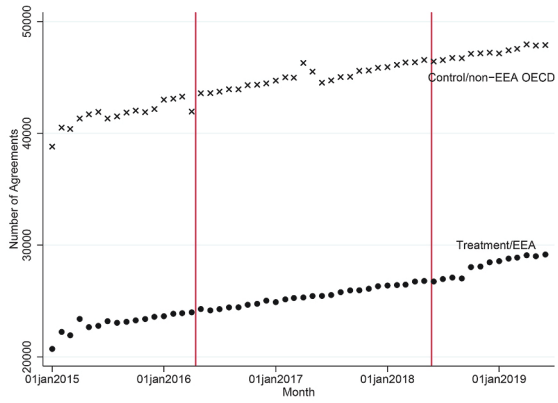


(a) Provider-to-customer (p2c) agreements

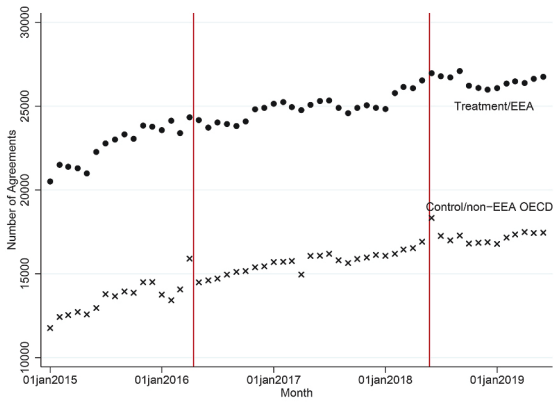


(b) Peering (p2p) agreements

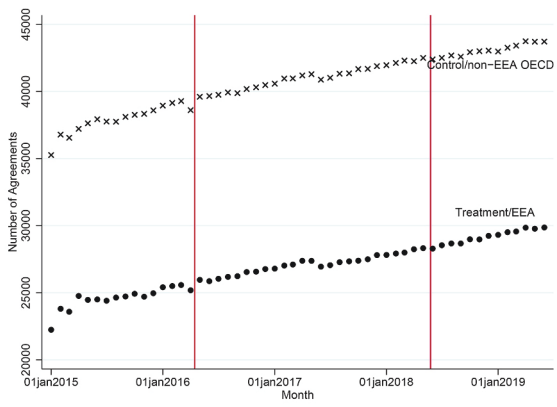


(c) Customer-to-provider (c2p) agreements

**Fig. 5.** Number of agreements by EEA and non-EEA OECD countries by agreement type. Notes: (a) The dots represent $\sum_{i\in EEA,t} numProvAg_{ijt}$, the total number of agreements by networks owned by EEA countries where these networks are providers. The crosses represent $\sum_{i\in OECD\wedge i\notin EEA,t} numProvAg_{ijt}$, the total number of agreements by networks owned by non-EEA OECD countries where these networks are providers. (b) The dots represent $\sum_{i\in EEA,t} numPeerAg_{ijt}$. The crosses represent$\sum_{i\in OECD\wedge i\notin EEA,t} numPeerAg_{ijt}$. Peering agree-ments between EEA countries and non-EEA OECD coun- tries contribute to counts in both series. (c) The dots represent $\sum_{i\in EEA,t} numCustAg_{ijt}$. The crosses represent $\sum_{i\in OECD\wedge i\notin EEA,t} numCustAg_{ijt}$.Agreements be-tween net-works owned by non-EEA non-OECD countries are not in-cluded in calculating these sums.

### *6.3. The number of IP-address-level interconnection points per agreement*

Our earlier results suggest the GDPR did not change whether agreements were made and what types of agreements were made between networks. One hypothesis for the absence of behavior change is that setting up an agreement is such a substantial decision that

**Table 4**
The GDPR s impact on the number of agreements by EEA and non-EEA OECD countries, by counterparty and agreement type.

| | Counterparty is non-EEA OECD | | | Counterparty is non-EEA non-OECD | | | Counterparty is EEA | | |
|---|---|---|---|---|---|---|---|---|---|
| | Provider | Peer | Customer | Provider | Peer | Customer | Provider | Peer | Customer |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
| $POST_e$ $EEA$ | 0.032 | 0.040 | 0.020 | 0.002 | 0.007 | 0.013*** | 0.007 | 0.038** | 0.005 |
| | (0.023) | (0.042) | (0.022) | (0.003) | (0.004) | (0.004) | (0.012) | (0.016) | (0.013) |
| $POST_a$ $EEA$ | 0.021 | 0.031 | 0.025 | 0.011*** | 0.007* | 0.006** | 0.027* | 0.002 | 0.021 |
| | (0.026) | (0.035) | (0.027) | (0.004) | (0.004) | (0.003) | (0.016) | (0.017) | (0.014) |
| Group dummies | ctry pairs | ctry pairs | ctry pairs | ctry pairs | ctry pairs | ctry pairs | ctry pairs | ctry pairs | ctry pairs |
| Time dummies | months | months | months | months | months | months | months | months | months |
| Clusters | 473 | 418 | 473 | 6751 | 6751 | 6751 | 1376 | 880 | 1376 |
| $_R2$ | 0.984 | 0.984 | 0.985 | 0.941 | 0.930 | 0.925 | 0.978 | 0.980 | 0.977 |
| Observations | 25,542 | 22,572 | 25,542 | 364,554 | 364,554 | 364,554 | 74,304 | 47,520 | 74,304 |

Notes: The dependent variable is $log(numProvAg_{i\,jt}$ 1) for columns (1), (4), (7), $log(numPeerAg_{i\,jt}$ 1) for columns (2), (5), (8), and $log(numCustAg_{i\,jt}$ 1) for columns (3), (6), (9). The dependent variables are rectangularized as described in Table 2 and we add one when we take the log to account for zero values. $POST_e$ is an indicator variable equal to 1 if the observation is made after the GDPR became effective. $POST_a$ is an indicator variable equal to 1 if the observation is made after the GDPR was approved. Columns (1), (2), (3) include observations when the treat-ment/control party is a network owned by an EEA/non-EEA OECD country and is the provider, peer, or customer to the counterparty network owned by a non-EEA OECD country. Columns (4), (5), (6) include observations when the treatment/control party is a network owned by an EEA/non-EEA OECD country and is the provider, peer, or customer to the counterparty network owned by a non-EEA non-OECD country. Columns (7), (8), (9) include observations when the treatment/control party is a network owned by an EEA/non-EEA OECD country and is the provider, peer, or customer to the counterparty network owned by an EEA country. All regressions include month dummies and country pair dummies. All regressions cluster standard error by country pair. Standard errors are in parentheses. Significantly different from 0 in a two-tailed test at the *10% level, **5% level, ***1% level.

changes in usage and bargaining friction due to the GDPR are small in comparison. Networks may only change the capacity associated with each interconnection in response to lower usage instead of cancelling an agreement altogether. If that is the case, we are unlikely to observe effects of the GDPR on the extensive margin. The GDPR s impact may be on how dense the two networks interconnection is. Motivated by this consideration, we examine how the GDPR affected the number of IP-address-level interconnection points two networks had, conditional on them having an agreement.

The outcome variable we study in this section is $numAgIP_{klt}$, the number of IP-address-level interconnection points between network $k$ and network $l$, given $k$ and $l$ have an agreement. As the unit of observations is a network pair, we hold fixed the EEA membership status (or OECD status) of the counterparty of interconnection while we compare the outcomes for EEA countries (treatment group) and for countries in the OECD but not in the EEA (control group).

As previously, we construct three subsamples based on counterparties: (a) the counterparties are in non-EEA OECD countries, (b) the counterparties are in non-EEA non-OECD countries, and (c) the counterparties are in EEA countries. Within each subsample, we then keep only observations where networks or countries are in the EEA (treatment group) or are in the OECD but not in the EEA (control group) and compare their outcomes.

The control group of either subsample (a) or subsample (b) does not involve networks in EEA countries. As before, we are less concerned about the results for these subsamples than those for subsample (c). Results from subsample (c) may be biased towards zero as non-EEA application firms need to comply with the GDPR in EEA markets or they may exit those markets, whichever would reduce derived demand for data exchange at the internet layer.

Fig. 6 shows a comparison of the average log number of interconnection points per agreement in the EEA countries and in the non-EEA OECD countries, holding fixed the interconnection coun-terparties. We first note that observed interconnection points with developing countries have a lot of noise in our measurement while observed interconnection points among EEA and OECD coun-tries are quite precisely measured, reflecting the large number of vantage points inside developed countries. When interconnection points are well-measured, we observe that, despite the differences in levels, EEA countries and non-EEA OECD countries still exhibit remarkable parallel trends in terms of the number of interconnection points per agreement throughout the sample period.

Given this particular data source only started in December 2015, close to the GDPR approval date, we do not include the interaction term $POST_{a,klt}$ $EEA_{kl\,jt}$ in our regression. Therefore, instead of Equation (1), we run the following regression on each of the three subsamples,

(2)

We take the log of the outcome variable $numAgIP_{klt}$ to reflect estimated effects in percentage changes. $POST_{e,klt}$ is an indicator variable equal to 1 if time $t$ is after the GDPR effective date. $POST_{a,klt}$ is an indicator variable equal to 1 if time $t$ is after the GDPR approval date. $EEA_{klt}$ is an indicator variable equal to 1 if the network pair $kl$ is in the treatment group for the subsample, and equal to 0 if the network pair $kl$ is in the control group for the subsample. A dummy $D_{kl}$ for each network pair $kl$ and a dummy $D_t$ for each week $t$ are included. The difference-in-differences effect is identified by the coefficient on the interaction term $POST_{e,klt}$ $EEA_{klt}$.

The results are shown in Table 5 and are in no case significantly different from zero. We include agreements present for at least 150 weeks for our regressions in Table 5. Given we study the intensive margin, alternatively we keep only agreements present for all of 169 weeks between December 2015 and June 2019. Doing so substantially reduces the sample size and the results are similar to those in

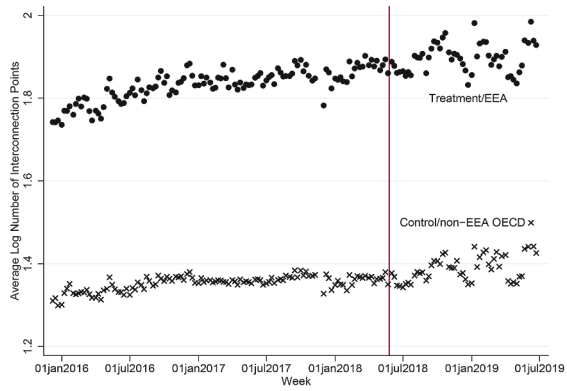


(a) Counterparties are non-EEA OECD countries

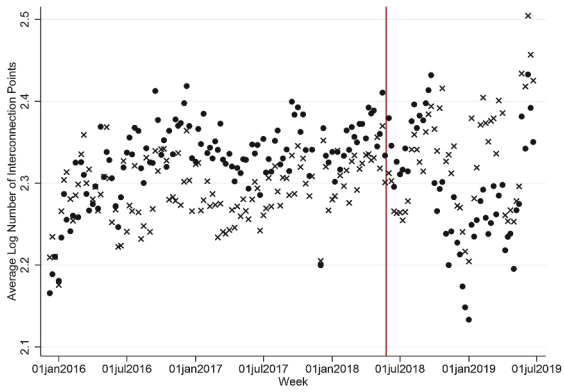


(b) Counterparties are non-EEA non-OECD countries

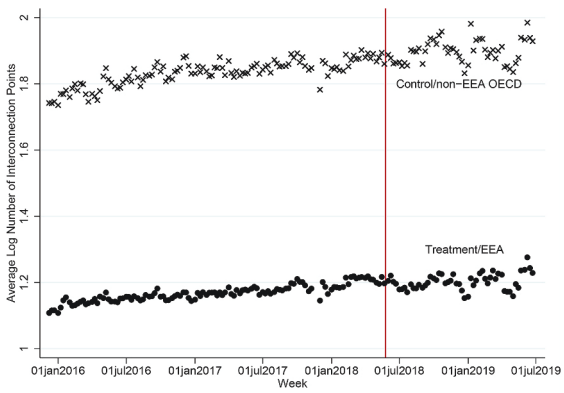


(c) Counterparties are EEA countries

**Fig. 6.** Average log number of IP-address-level interconnection points per agreement by EEA and non-EEA OECD countries by counterparty. Notes: The dots represent $\overline{log(numAgIP_{k\in EEA,lt})}$, the log number of IP-address-level interconnection points aver- aged among agreements by networks owned by EEA countries when the counterparties are networks owned by (a) non-EEA OECD countries, (b) non-EEA non-OECD countries, (c) EEA countries. The crosses represent $\overline{log(numAgIP_{k\in OECD\wedge k\notin EEA,lt})}$, the log number of IP-address-level interconnection points averaged among agreement by networks owned by non-EEA OECD countries when the counterparties are networks owned by (a) non-EEA OECD countries, (b) non-EEA non-OECD countries, (c) EEA countries. Only agreements present throughout Dec 2015 – June 2019 are used to take the averages.

**Table 5**
The GDPR s impact on the number of IP-address-level interconnection points per agree-ment by EEA and non-EEA OECD countries, by counterparty.

| | Non-EEA OECD | Non-EEA Non-OECD | EEA |
|---|---|---|---|
| | (1) | (2) | (3) |
| $POST_e$ *EEA* | 0.039 | 0.003 | 0.032 |
| | (0.023) | (0.049) | (0.024) |
| Group dummies | network pairs | network pairs | network pairs |
| Time dummies | weeks | weeks | weeks |
| Clusters | 128 | 522 | 307 |
| $_R2$ | 0.871 | 0.827 | 0.867 |
| Observations | 2,593,805 | 494,374 | 1,886,031 |

Notes: The dependent variable is $log(numAgIP_{klt})$. $POST_e$ is an indicator variable equal to 1 if the observation is made after the GDPR became effective. $POST_a$ is an indicator variable equal to 1 if the observation is made after the GDPR was approved. Column (1) includes observations when one party of the agreement is a network owned by an EEA or non-EEA OECD country and the counterparty is a network owned by a non- EEA OECD country. Column (2) includes observations when one party of the agreement is a network owned by an EEA or non-EEA OECD country and the counterparty is a network owned by a non-EEA non-OECD country. Column (3) includes observations when one party of the agreement is a network owned by an EEA or non-EEA OECD country and the counterparty is a network owned by an EEA country. Only agreements present for at least 150 weeks are used. The GDPR approval date Apr 2016 is close to the sample starting date Dec 2015, so $POST_a$ *EEA* is not included in the regressions. All regressions include week dummies and network pair dummies. All regressions cluster standard error by country pair. Standard errors are in parentheses. Significantly different from 0 in a two-tailed test at the *10% level, **5% level, ***1% level.

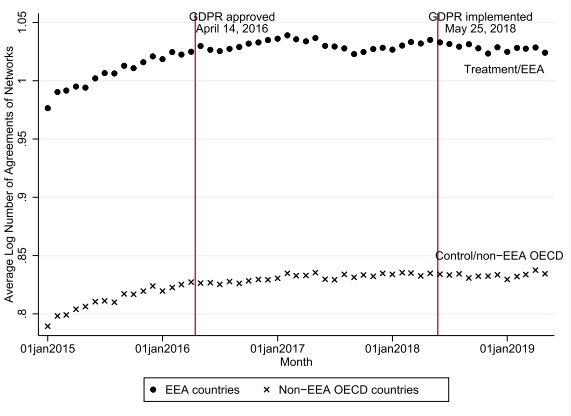


**Fig. 7.** Average log number of interconnection agreements by networks in EEA and non-EEA OECD countries. Notes: The dots represent $\overline{log\ numAgNtwrk_{k\ EEA\ t}}$, the log number of agreements averaged among networks owned by EEA countries. The crosses represent $\overline{log\ numAgNtwrk_{k\ OECD\ k\ EEA\ t}}$, the log number of agreements averaged among networks owned by non-EEA OECD countries. Non- EEA and non-OECD countries networks are not included in taking the averages. Only networks present throughout Jan 2015 June 2019 are used to take the averages. The first red vertical line represents 14 April 2016, the approval date of the GDPR. The second red vertical line represents 25 May 2018, the implementation date of the GDPR. Regression including month and network fixed effects has the coefficient on $POST_e$ *EEA* .004 (se 0.007, clustered by country) and the coefficient on $POST_a$ *EEA* 0.006 (se 0.007, clustered by country). Both are insignificant at conventional levels of significance.

Table 5.

### 6.4. The number of agreements by networks

In addition to interconnection behavior between pairs of countries or networks, we study how the GDPR might have impacted the number of agreements per network, the number of networks per country and the sizes of the customer cones of each network. We stress that these estimates are differential impact between EEA networks and non-EEA OECD networks.

Fig. 7 shows a comparison of the average log number of agreements by networks in the EEA countries and in the non-EEA OECD countries. We observe visually apparent parallel trends be-tween the two groups prior to the approval of the GDPR, between the approval and implementation of the GDPR, as well as after the implementation of the GDPR.

We then run the regression specification in Equation (1). The outcome variable is $numAgNtwrk_{kt}$, the number of agreements network $k$ has in month $t$. The unit of observation $m$ is network $k$.

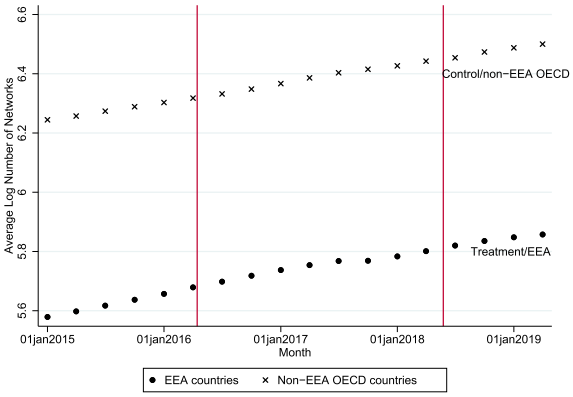


**Fig. 8.** Average log number of networks in EEA and non-EEA OECD countries. Notes: The dots represent $\overline{log\ numNtwrk_{i\ EEA\ t}}$, the log number of networks averaged among EEA countries. The crosses represent $\overline{log\ numNtwrk_{i\ OECD\ i\ EEA\ t}}$, the log number of networks averaged among non-EEA OECD countries. Non-EEA and non-OECD countries networks are not included in taking the averages. Regression including quarter and country fixed effects has the coefficient on $POST_e$ *EEA* .002 (*se* 0.017, clustered by country) and the coefficient on $POST_a$ *EEA* 0.016 (*se* 0.024, clustered by country). Both are insignificant at conventional levels of significance.

We show the results below Fig. 7. The coefficient on $POST_e$ *EEA* and the coefficient on $POST_a$ *EEA* are both insignificant at conventional levels of significance. This result suggests the GDPR does not have differential impact on the number of interconnection agreements on EEA networks relative to their non-EEA OECD counterparts.

### 6.5. The number of networks

Fig. 8 shows a comparison of the average log number of networks per country in the EEA coun-tries and in the non-EEA OECD countries. Again, we observe visually apparent parallel trends between the two groups prior to the approval of the GDPR, between the approval and implementa-tion of the GDPR, as well as after the implementation of the GDPR.

We then run the regression specification in Equation (1). The outcome variable is $numNtwrk_{it}$, the number of networks country *i* has in quarter *t*. The unit of observation *m* is country *i*.

We show the results below Fig. 8. The coefficient on $POST_e$ *EEA* and the coefficient on $POST_a$ *EEA* are both insignificant at conventional levels of significance. This result suggests the GDPR does not differentially impact the number of networks in EEA countries compared to non- EEA OECD countries.

### 6.6. Customer cone of networks

Fig. 9 shows a comparison of the average log customer cone of networks in the EEA countries and in the non-EEA OECD countries. We observe visually apparent parallel trends between the two groups prior to the approval of the GDPR, between the approval and implementation of the GDPR, as well as after the implementation of the GDPR.

We then run the regression specification in Equation (1). The outcome variable is $NtwrkCustCone_{kt}$, the size of network *k* s customer cone in month *t*. The unit of observation *m* is network *k*.

We show the results below Fig. 9. Though both the coefficient on $POST_e$ *EEA* and the coefficient on $POST_a$ *EEA* are significantly different from zero, their magnitudes are economically very small, suggesting the GDPR has little impact on the centrality of networks in EEA countries compared to non-EEA OECD countries.

### 6.7. Robustness checks

In this subsection, we discuss our robustness checks. We first replace the logged outcome variables in all of our specifications with their original unlogged values. We present the results in Appendix C.1. As shown, all of the results are qualitatively similar to results with logged outcome variables.

This alleviates the concern that our zero estimates are driven by taking the log of the outcome variables.

We then perform a robustness check by redefining $POST_a$ to equal to 1 if the observation is made after December 2015. As discussed in our empirical strategy section, it is possible that networks invested in interconnection decisions in anticipation of the enforcement of the GDPR even before the law was approved. We think the earliest possible date for the firms to respond in anticipation is December 2015, the time when the text of the law was fixed. The results from this robustness check for our various outcomes are almost identical to our main results presented in Tables 3 5 and Figs. 7 9. This alleviates the concern that our main specifications did not capture possible effects due to anticipation.

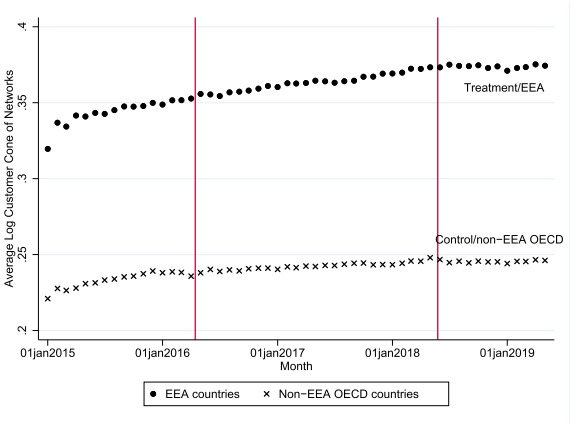


**Fig. 9.** Average log customer cone of networks in EEA and non-EEA OECD countries. Notes: The dots represent $\overline{log\ NtwrkCustCone_{k\ EEA\ t}}$, the log customer cone averaged among networks owned by EEA countries. The crosses represent $\overline{log\ NtwrkCustCone_{k\ OECD\ k\ EEA\ t}}$, the log customer cone averaged among networks owned by non-EEA OECD countries. Non-EEA and non-OECD countries networks are not included in taking the averages. Only networks present throughout Jan 2015 June 2019 are used to take the averages. Regression including month and network fixed effects has the coefficient on $POST_e$ *EEA* .007* (*se* 0.004, clustered by country) and the coefficient on $POST_a$ *EEA* 0.011*** (*se* 0.004, clustered by country). Though both are significantly different from zero, their magnitudes are economically small.

Lastly, we perform a first differences regression for each outcome using only the subsample consisting of the treatment group in our main regression. The rationale for this robustness check is as follows. Consider a scenario in which the GDPR did have a nonzero effect on networks. This would mean the trends before and after the policy cutoff dates were different for the treatment group (and possibly also for the control group). Also suppose that the effects of the policy were identical for the treatment and control groups, or in other words, the trends for the treatment group and control group changed for the exact same amount post policy cutoffs. Then our difference-in-differences approach would not be able to detect the effect of the policy. We therefore use a first differences approach to test whether there were differential trends before and after the GDPR for the treatment group alone. We will be able to rule out the above scenario if the first differences estimates are zero.

Specifically, we run the following regression:

( ) ( ) (3)

where $m$ is the unit of observation of the outcome variable of interest. $m$ can take country pair subscript $ij$, network pair subscript $kl$, network subscript $k$, or country subscript $i$. $POST_{e,mt}$ is an indicator variable equal to 1 if time $t$ is after the GDPR effective date. $POST_{a,mt}$ is an indicator variable equal to 1 if time $t$ is after the GDPR approval date. A dummy $D_m$ for each unit of obser-vation $m$ is included. We do not include a time dummy as such a dummy would absorb any effect of the GDPR. The effect on first differences is identified by the coefficients on the terms $POST_{e,mt}$ and $POST_{a,mt}$. We show these results in Appendix C.2. As shown, the results are very precisely estimated and, in all cases, do not exceed 0.6 percentage points. These results suggest that the rate of growth in interconnection in the EEA after the GDPR did not change when compared to that before the GDPR. We can therefore rule out the scenario under which the GDPR had significant and identical impact on EEA and non-EEA OECD networks.

## 7. Conclusion

The effectiveness of the Internet in creating economic surplus depends on efficient interconnections bilaterally negotiated by independently operated networks. In this paper, we investigate whether the approval and implementation of the GDPR affects the growth in interconnection of the Internet in Europe. Despite evidence that the GDPR so far had significant effects at the application layer on European firms, we find no visible consequences at the infrastructure layer across the multiple measures we have. Occasionally we estimate statistically significant effects, which prove to be not robust. Our robustness checks suggest that our results are not driven by taking the log of outcome variables or our choice of April 2016 as the first policy cutoff date. Using the first differences approach as an additional robustness check, we show that EEA networks had similar growth rates in interconnection agreements before and after the GDPR approval or implementation, therefore our main results are not driven by our choice of control group.

A number of possible reasons could have contributed to this finding. First, the lack of discernible short-run effect on interconnection could have arisen from slow investment and behavioral changes at the internet layer. This seems unlikely because renegotiations of interconnection agreements happen frequently and we observe continued growth across all network connections.

It is also possible that despite the evident behavioral changes at the application layer due to the GDPR, the effect is small compared to other considerations in negotiating interconnection agree-ments. That could happen if, for example, the regular growth in data due to growth in many applica-tions overwhelms any short-run impact of the GDPR. In that case, network operators may rationally expect the long run effect of the GDPR to be small even at the application layer.

Another possibility is that the GDPR may change behaviors at the internet layer only along the margins that our measures do not capture. For example, if changes in derived demand for data flow only prompt network operators to change capacity at existing interconnection points, then these changes are not visible to us.

In addition, the GDPR has seen limited and heterogeneous enforcement until the end of our sample period and it may not materially reduce Internet interconnections. While regulators have issued large fines during the period,[58] they have issued relatively few of them[59] and in many cases have stopped short of enforcement.[60] Enforcement actions may directly and materially affect data flow and interconnections.[61]

Finally we only observe the short run, so we cannot rule out that more gradual changes due to the GDPR may surface in the longer run, which is an open question. If we are able to observe a much longer period of time, we will be able to use the data from additional periods and the same methodology to study the effect of the GDPR in the longer run.

Our results have immediate policy implications. As many countries are contemplating imple-menting their own versions of privacy and data protection regulations, there are concerns about whether such regulations may negatively impact the growth of the Internet, reduce technology firms incentives in operating and innovating, reduce the use of the Internet in productivity enhancing ac-tivities, and reduce the economic surplus generated through the use of the Internet in the country and beyond. Our results suggest limited effects of such regulations at the internet layer for the measures we are able to capture.

Our results also speak to the debate on the allocation of rents generated through the successful commercialization of the Internet. The enormous rents associated with the exploitation of Web 2.0 and mobile web represent a large portion of the private returns to innovation in the 21st century. These rents have been overwhelmingly captured by players at the application layer, notably the big tech companies, while firms at the internet layer captures little of the rents. Our study is consistent with the view that the cost of the GDPR has been a shock to rents, and the costs have been borne by the application layer, paid out of the rents from innovation.

Our results also mask the potential heterogeneity in the GDPR s impact on Internet data flow across application categories. The GDPR affects application firms to different degrees, depending on their business models and how much those models depend on monetization of personal data. Similarly, applications differ in the amount of data traffic generated. The interaction of the two affects the extent to which the GDPR affects the derived demand for data exchange at the inter-net layer. Video traffic includes Netflix and Youtube, and may be less affected by GDPR due to subscription-based business models, and could grow even as other traffic drops. This combination may explain the visible effects of the GDPR to certain application firms but little evidence of effects on interconnection decisions at the internet layer.

We also note that current empirical works, including this paper, study the impact of the GDPR on suppliers of Internet services and content at various layers of the Internet. Empirical evidence on consumers responses to privacy regulation is extremely lacking. As policy makers strive to enhance consumer welfare through better privacy protection while trying to minimize such laws impact of the digital economy, evaluating the laws effect on consumers is an important direction for future research to allow for the overall welfare analysis.

In addition to policy implications, our paper presents data consisting of virtually all operating networks in the world and a large number of interconnection agreements among them across many years, which opens the possibility of investigating a range of eco-nomic- and policy-relevant ques-tions. We acknowledge that important pieces of data are still missing, notably the actual data flows between networks on a similar scale. Future works should keep bridging the unmet data needs.

**Acknowledgements**

The authors thank Erik Bohlin, Tim Bresnahan, Dennis Carlton, Roderick Fanou, Samuel Goldberg, Avi Goldfarb, Ginger Zhe Jin, Garrett Johnson, Stephen Strowes, and three anonymous referees for helpful suggestions. We thank Dan Andersen for technical assis-tance. The authors are grateful to the Doctoral Office at Harvard Business School for financial support for field work.

**Appendix A. Supplementary data**

Supplementary data to this article can be found online at https://doi.org/10.1016/j.telpol.2020.102083.

[58] French Data Protection Authority (CNIL) issued a e50 million fine, the largest GDPR fine to date, to Google Inc. on January 21, 2019 due to issues with Google s consent practices when users configure their Android phones (CNIL, 2019).

[59] Regulators have issued a total of 57 GDPR fines up to June 2019.

[60] The Irish Data Protection Commission conducted a sweep on 38 data controllers between August 2019 and Decem-ber 2019. They found a range of violations but allowed for a six-month period for compliance before taking enforcement actions (Data Protection Commission (2020)).

[61] The UK Information Commissioner s Office (ICO) found that implementing its own best practices regarding opt-in consent on its own website reduced data flows to Google Analytics by 90%. (Cross, Tim. 2019. The ICO s Cookie Consent Rate Dropped 90 Percent After Implementing its Own Best Practices. Video Ad News.)

## References


Aridor, G., Che, Y.-K., & Tobias Salz. (2020). *The economic consequences of data privacy regulation: Empirical evidence from GDPR*. Available at: SSRN 3522845.

Besen, S., Paul, M., Mitchell, B., & Padmanabhan, S. (2001). Advances in routing technologies and internet peering agreements. *The American Economic Review, 91*(2), 292 296.

Binmore, K., Rubinstein, A., & Wolinsky, A. (1986). The Nash bargaining solution in economic modelling. *The RAND Journal of Economics*, 176 188.

Internet Engineering Task Force. (1989). *RFC 1122: Requirements for internet hosts communi- cation layers*. In R. Braden (Ed.). https://tools.ietf.org/html/rfc1122.

Center for Applied Internet Data Analysis. (2013). *Routeviews Prefix-to-AS mappings (pfx2as) for IPv4 and IPv6*. San Diego: University of California. http://data.caida.org/datasets/routing/routeviews-prefix2as/README.txt. (Accessed 21 September 2019).

Center for Applied Internet Data Analysis. (2015 2019). *AS relationships*. San Diego: University of California. http://www.caida.org/data/as-relationships/ accessed June-July 2019.

Center for Applied Internet Data Analysis. (2015 2019). *Inferred AS to organization map- ping dataset*. San Diego: University of California. http://www.caida.org/data/as-organizations/ accessed June-July 2019.

Center for Applied Internet Data Analysis. (2015 2019). *IPv4 prefix-probing traceroute dataset*. San Diego: University of California accessed June-July 2019 https://www.caida.org/data/active/ipv4 prefix probing dataset.xml.

Center for Applied Internet Data Analysis. (2015 2019). *Macroscopic internet topology data kit (ITDK)*. San Diego: University of California accessed June-July 2019 http://www.caida.org/data/Internet- topology-data-kit/.

Choi, J. P., Jeon, D.-S., & Kim, B.-C. (2015). Net neutrality, business models, and internet interconnection. *American Economic Journal: Microeconomics, 7*(3), 104 141.

Cisco Talos Intelligence Group. Email and spam data (september 2020). Retrieved https://talosintelligence.com/reputation center/email rep. (Accessed 8 October 2020).

Commission Nationale de l Informatique et des Libertés (CNIL.). (2019). *The CNIL s restricted committee imposes a financial penalty of 50 million*. Euros Against GOOGLE LLC. https://www.cnil.fr/en/cnils-restricted-committee-imposes-financial-penalty-50-million-euros-against-google-llc#::text On%2021%20January%202019%2C%20the,consent%20regarding%20the%20ads%20personalization.

Data Protection Commission, The. (2020). *Report by the data protection commission on the use of cookies and other tracking technologies*. https://dataprotection.ie/en/news-media/publications/report-dpc-use-cookies-and-other-tracking-technologies.

Degeling, M., Utz, C., Lentzsch, C., Henry, H., Schaub, F., & Holz, T. (2018). *We value your privacy... Now take some cookies: Measuring the GDPR s Im- pact on web privacy*. arXiv preprint arXiv:1808.05096.

Di Battista, G., Patrignani, M., & Pizzonia, M. (2003). Computing the types of the relationships between autonomous systems. In *Ieee INFOCOM 2003. Twenty-second annual joint conference of the IEEE computer and communications societies (IEEE cat. No. 03CH37428)* (Vol. 1, pp. 156 165). IEEE.

Dimitropoulos, X., Krioukov, D., Bradley, H., & Riley, G. (2005). Infer- ring AS relationships: Dead end or lively beginning?. In *International workshop on ex- perimental and efficient algorithms* (pp. 113 125). Berlin, Heidelberg: Springer.

Dimitropoulos, X., Krioukov, D., Fomenkov, M., Bradley, H., Young, H., & Riley, G. (2007). AS relationships: Inference and validation. *ACM SIGCOMM - Computer Communication Review, 37*(1), 29 40.

D Ignazio, A., & Giovannetti, E. (2009). Asymmetry and discrimination in internet peering: Evidence from the LINX. *International Journal of Industrial Organization, 27*(3), 441 448.

Englehardt, S., & Narayanan, A. (2016). Online tracking: A 1-million-site measure- ment and analysis. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security* (pp. 1388 1401).

Erlebach, T., Hall, A., & Schank, T. (2002). *Classifying customer-provider relationships in the internet*. TIK-Report 145.

Gao, L. (2001). On inferring autonomous system relationships in the internet. *IEEE/ACM Transactions on Networking, 9*(6), 733 745.

Godinho de Matos, M., & Adjerid, I. (2019). *Consumer consent and firm targeting after GDPR: The case of a large telecom provider*. Working paper).

Goldberg, S., Johnson, G., & Scott, S. (2019). *Regulating privacy online: The early impact of the GDPR on European web traffic & E-commerce outcomes*. Available at: SSRN 3421731.

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

Hoofnagle, C. J., Bart van der, S., & Frederik Zuiderveen, B. (2019). The European union general data protection regulation: What it is and what it means. *Information and Communications Technology Law, 28*(1), 65 98.

Iordanou, C., Smaragdakis, G., Poese, I., & Laoutaris, N. (2018). Tracing cross border web tracking. In *Proceedings of the internet measurement conference 2018* (pp. 329 342). ACM, 2018.

Jia, J., Jin, G. Z., & Wagman, L. (2019). *The short-run effects of GDPR on technol- ogy venture investment*. National Bureau of Economic Research Working. Paper No. w25248.

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

Johnson, G. A., & Shriver, S. K. (2019). *Privacy & market concentration: Intended & unintended consequences of the GDPR*. Available at: SSRN 3477686.

Karaj, A., Macbeth, S., Berson, Rémi, & Pujol, J. M. (2018). *WhoTracks. Me: Shedding light on the opaque world of online tracking*. arXiv preprint arXiv:1804.08959.

Laffont, J.-J., Scott, M., Patrick Rey, & Jean, T. (2001). Interconnection and access in telecom and the internet. *AEA Papers and Proceedings, 91*(2), 287 291.

Lefrere, V., Warberg, L., Cheyre, C., Marotta, V., & Acquisti, A. (2020). *The impact of the GDPR on content providers*.

Libert, T. (2015). *Exposing the hidden web: An analysis of third-party HTTP requests on 1 million websites*. arXiv preprint arXiv:1511.00619.

Libert, T., Graves, L., & Kleis Nielsen, R. (2018). *Changes in third-party con- tent on European news websites after GDPR*. Factsheet: Reuters Institute for the Study of Journalism Reports.

Lodhi, A., Larson, N., Dhamdhere, A., & Dovrolis, C. (2014). Using peer- ingDB to understand the peering ecosystem. *ACM SIGCOMM - Computer Communication Review, 44*(2), 20 27.

Luckie, M., Bradley, H., Dhamdhere, A., & Giotsas, V. (2013). AS rela- tionships, customer cones, and validation. In *Proceedings of the 2013 conference on Internet measurement conference* (pp. 243 256). ACM, 2013.

Marder, A., Luckie, M., Dhamdhere, A., Bradley, H., Jonathan, M., & Smith. (2018). Pushing the boundaries with bdrmapIT: Mapping router ownership at internet scale. In *Proceedings of the internet measurement conference 2018* (pp. 56 69). ACM, 2018.

Meltzer, J. (2014). *The importance of the internet and transatlantic data flows for US and EU trade and investment*. Global Economy and Development Working Paper 79.

Messaging Anti-Abuse Working Group (MAAWG). (2011). *Email metrics program: The network operator s perspective*. Report 14 http://www.maawg.org/sites/maawg/files/news/MAAWG 2010 Q3Q4 Metrics Report 14.pdf.

Miller, A. R., & Tucker, C. E. (2011). Can health care information technology save babies? *Journal of Political Economy, 119*(2), 289 324.

Miller, A. R., & Tucker, C. (2018). Privacy protection, personalized medicine, and genetic testing. *Management Science, 64*(10), 4648 4668.

Mohan, J., Wasserman, M., & Chidambaram, V. (2019). *Analyzing GDPR com- pliance through the lens of privacy policy*. arXiv preprint arXiv:1906.12038.

Nicholson, J. R., & McHenry, G. (2016). *Measuring cross-border data flows: Data, lit- erature, and considerations*. US Department of Commerce. https://www.ntia.doc.gov/other- publication/2016/measuring-cross-border-data-flows-data-literature-and-considerations.

Norton, W. B. (2014). *The 2014 internet peering playbook: Connecting to the core of the internet*. DrPeering Press.

Peukert, C., Bechtold, S., Batikas, M., & Kretschmer, T. (2020). *European privacy law and global markets for data*. CEPR Discussion Paper 14475.

Rao, J. M., & Reiley, D. H. (2012). The economics of spam. *The Journal of Economic PerSpectives, 26*(3), 87 110.

Sandvine. (2018). The global internet phenomena report. https://www.sandvine.com/hubfs/downloads/phenomena/2018-phenomena-report.pdf.

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

Subramanian, L., Agarwal, S., Rexford, J., & Katz, R. H. (2002). Characterizing the internet hierarchy from multiple vantage points. In *Proc. IEEE infocom 2002*.

Symantec. (2010). *Messagelabs intelligence: 2010 annual security report*. https://retelur.files.wordpress.com/2007/10/symantec-messagelabsintelligencefinal2010-110111043129-phpapp02.pdf.

The Internet Society. (2015). *Policy brief: Internet interconnection*. https://www.Internetsociety.org/policybriefs/Internetinterconnection/.
US International Trade Commission. (2014). Digital trade in the US and global economies, Part 2. Available at: https://www.strtrade.com/media/publication/7238pub4485.pdf.
Weller, D., & Woodcock, B. (2013). *Internet traffic exchange market development and policy challenges*. OECD Digital Economy Papers. https://doi.org/10.1787/5k918gpt130q- en
Xia, J., & Gao, L. (2004). On the Evaluation of AS Relationship Inferences [Inter- net reachability/traffic flow applications]. In *IEEE global telecommunications conference, 2004. GLOBECOM 04* (Vol. 3, pp. 1373 1377). IEEE.