

Sabatini, Fabio

Working Paper

The Behavioral, Economic, and Political Impact of the Internet and Social Media: Empirical Challenges and Approaches

IZA Discussion Papers, No. 16703

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Sabatini, Fabio (2023) : The Behavioral, Economic, and Political Impact of the Internet and Social Media: Empirical Challenges and Approaches, IZA Discussion Papers, No. 16703, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/282830>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION PAPER SERIES

IZA DP No. 16703

**The Behavioral, Economic, and Political
Impact of the Internet and Social Media:
Empirical Challenges and Approaches**

Fabio Sabatini

DECEMBER 2023

DISCUSSION PAPER SERIES

IZA DP No. 16703

The Behavioral, Economic, and Political Impact of the Internet and Social Media: Empirical Challenges and Approaches

Fabio Sabatini

Sapienza University of Rome and IZA

DECEMBER 2023

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

The Behavioral, Economic, and Political Impact of the Internet and Social Media: Empirical Challenges and Approaches

This paper presents a review of empirical methods used to assess the behavioral, economic, and political outcomes of Internet and social media usage. Instead of merely surveying the various impacts of the Internet, we examine the methods adopted to identify these impacts. We describe two main approaches for establishing causal effects, each with strengths and limitations. The first approach involves searching for exogenous sources of variation in the access to fast Internet or specific content. The second approach takes the form of field or laboratory experiments. In this paper, we focus on the first approach, delving into the methodological threats, empirical design, and main findings of the most prominent studies that exploit natural or quasi-experiments for identifying the causal impact of high-speed Internet or specific social media. This undertaking allows us to highlight the key empirical challenges in the field of Internet and social media economics while summarizing the causal relationships that the literature has uncovered so far.

JEL Classification: D71, D72, D74, D83, L82, L86, L88, L96, L98, Z13

Keywords: internet, social media, artificial intelligence, broadband infrastructure, politics, social capital

Corresponding author:

Fabio Sabatini
Sapienza Università di Roma
via del Castro Laurenziano 9
00161, Roma
Italy

E-mail: fabio.sabatini@uniroma1.it

1 Introduction

The advent and subsequent proliferation of the Internet and social media platforms have revolutionized the way we interact, communicate, and access information. While initial praises for these technologies highlighted their numerous advantages, we now see a growing focus on the potential negative impact they can have on human behavior, political dynamics, and societal structures. However, determining the causal effects of Internet usage poses significant methodological challenges.

Reverse causality is one of the major concerns, as it is hard to determine whether Internet use is the cause or effect of certain outcomes. For instance, individuals who are socially isolated may turn to the Internet as a way to connect with others, while others may become socially isolated as a result of excessive Internet use. Furthermore, unobserved personality traits can influence the predisposition towards Internet use and pre-existing behaviors and attitudes, raising the danger of incorrectly attributing these latter elements as consequences of online engagement. Finally, environmental factors further confound these relationships, creating an intricate web of cause and effect. The prevailing characteristics of the local population, the attributes of the residential area, and collective shocks could concurrently influence broadband deployment, online activities, and the presumptive outcomes of Internet use. This interplay of environmental nuances may lead to a common bias that further muddles the relationships among the treatment and the outcome variables. For instance, areas with higher levels of development typically have a greater demand for Internet services, thus rendering investments in broadband infrastructure more profitable. At the same time, development results in numerous behavioral and political consequences that might be inaccurately attributed to the diffusion of broadband services or social media.

The objective of this paper is to review the empirical literature on the behavioral and political outcomes of Internet and social media use, with a specific focus on the identification challenges. Instead of merely cataloging the manifold impacts of the Internet, our focus lies on surveying the methods used to identify these impacts. Our guiding questions include: how do economic studies establish the causal impact of Internet and social media use? How reliable is the existing evidence on the behavioral, economic, and political outcomes of the Internet? And what is the most appropriate approach for ascertaining causal effects in this field of research?

We identify two main approaches to assessing the causal impact of the Internet and social media. The first approach exploits exogenous circumstances that randomly assigned access to fast Internet, or specific content, to certain areas or users. Studies following this method often combine information about exogenous discontinuities in the provision or quality of Internet services with geolocated survey or administrative data measuring individual or aggregated behaviors. This approach is affected by the physiological lack of individual-level information on what people do

online, and even whether they actually use the Internet and social media. This limitation forces the analysis to adopt an intention-to-treat-like design framework where any user or area with potential Internet access is considered “treated.” The second approach involves field experiments or treatment manipulations in controlled environments that allow researchers to observe user responses to specific stimuli. Although these methods mitigate the lack of individual-level data on online access and behavior and control for confounding factors, they face potential external validity issues.

In this paper, we focus on the first approach, as we delve into studies that leverage natural and quasi-experiments to gauge the impact of the Internet and social media. The initial part of the review examines studies assessing the influence of high-speed Internet access, while the latter part analyzes those investigating the effects of social media. Both sets of research exploit a range of exogenous sources of treatment exposure, each dictating a specific empirical strategy. Studies on broadband access primarily exploit natural and quasi-experiments stemming from the supply side of Internet services. Conversely, literature on social media tends to leverage arguably exogenous variations in the demand for these platforms. The first section kicks off with an exploration of the technological landscape where Digital Subscriber Line (DSL) technology became a dominant method for offering quick Internet access across numerous nations. We will then delve into how the architecture of voice networks played a pivotal role in promoting the dissemination of DSL. During the early stages of broadband penetration, the proximity to a network node had a significant impact on high-speed Internet access. This geographical factor provided an exogenous variation in the quality of Internet connections across areas with very similar characteristics. Further into our review, we will address policy initiatives that resulted in quasi-experimental conditions for high-speed internet access. Finally, we will address the staggered roll-out of mobile networks, the deployment of undersea cables to Africa, and unexpected technical issues that impacted the provision of Internet services.

In the second section of our survey, we turn our focus to social media. Given the intangible nature of social networks’ architecture, identification issues become increasingly challenging. Traditional supply-side factors tend to diminish in importance, as social networks typically proliferated after broadband penetration was nearly complete, enabling anyone, anywhere, to easily sign up for these platforms. Nonetheless, several authors have exploited the intangible topology of online networks to develop new identification methods, under the assumption that the most central nodes could facilitate platform penetration and content accessibility. In this context, the dispersion of these central nodes could prove instrumental for identification, provided it’s feasible to identify the exogenous factors influencing the architecture of the network and to associate each node with a distinct geographical area. Several researchers maintain that network centrality exhibits a significant degree of temporal consistency. This enduring nature of network centrality potentially enables the initial geographic distribution of nodes across a certain region to act as an

exogenous source of variation, suitable for an instrumental variables approach.

Our review of the social media literature begins with a discussion of studies that used the initial geographical distribution of key network nodes as a source of variability in a platform’s penetration. Following this, we explore studies that harness online movements or trends that drove surges in the demand for social media. Then, we address studies exploiting the timing of publication of specific social media content to identify its impact on offline outcome variables. Lastly, we delve into research utilizing supply-side factors that generated geographical disparities in social media accessibility.

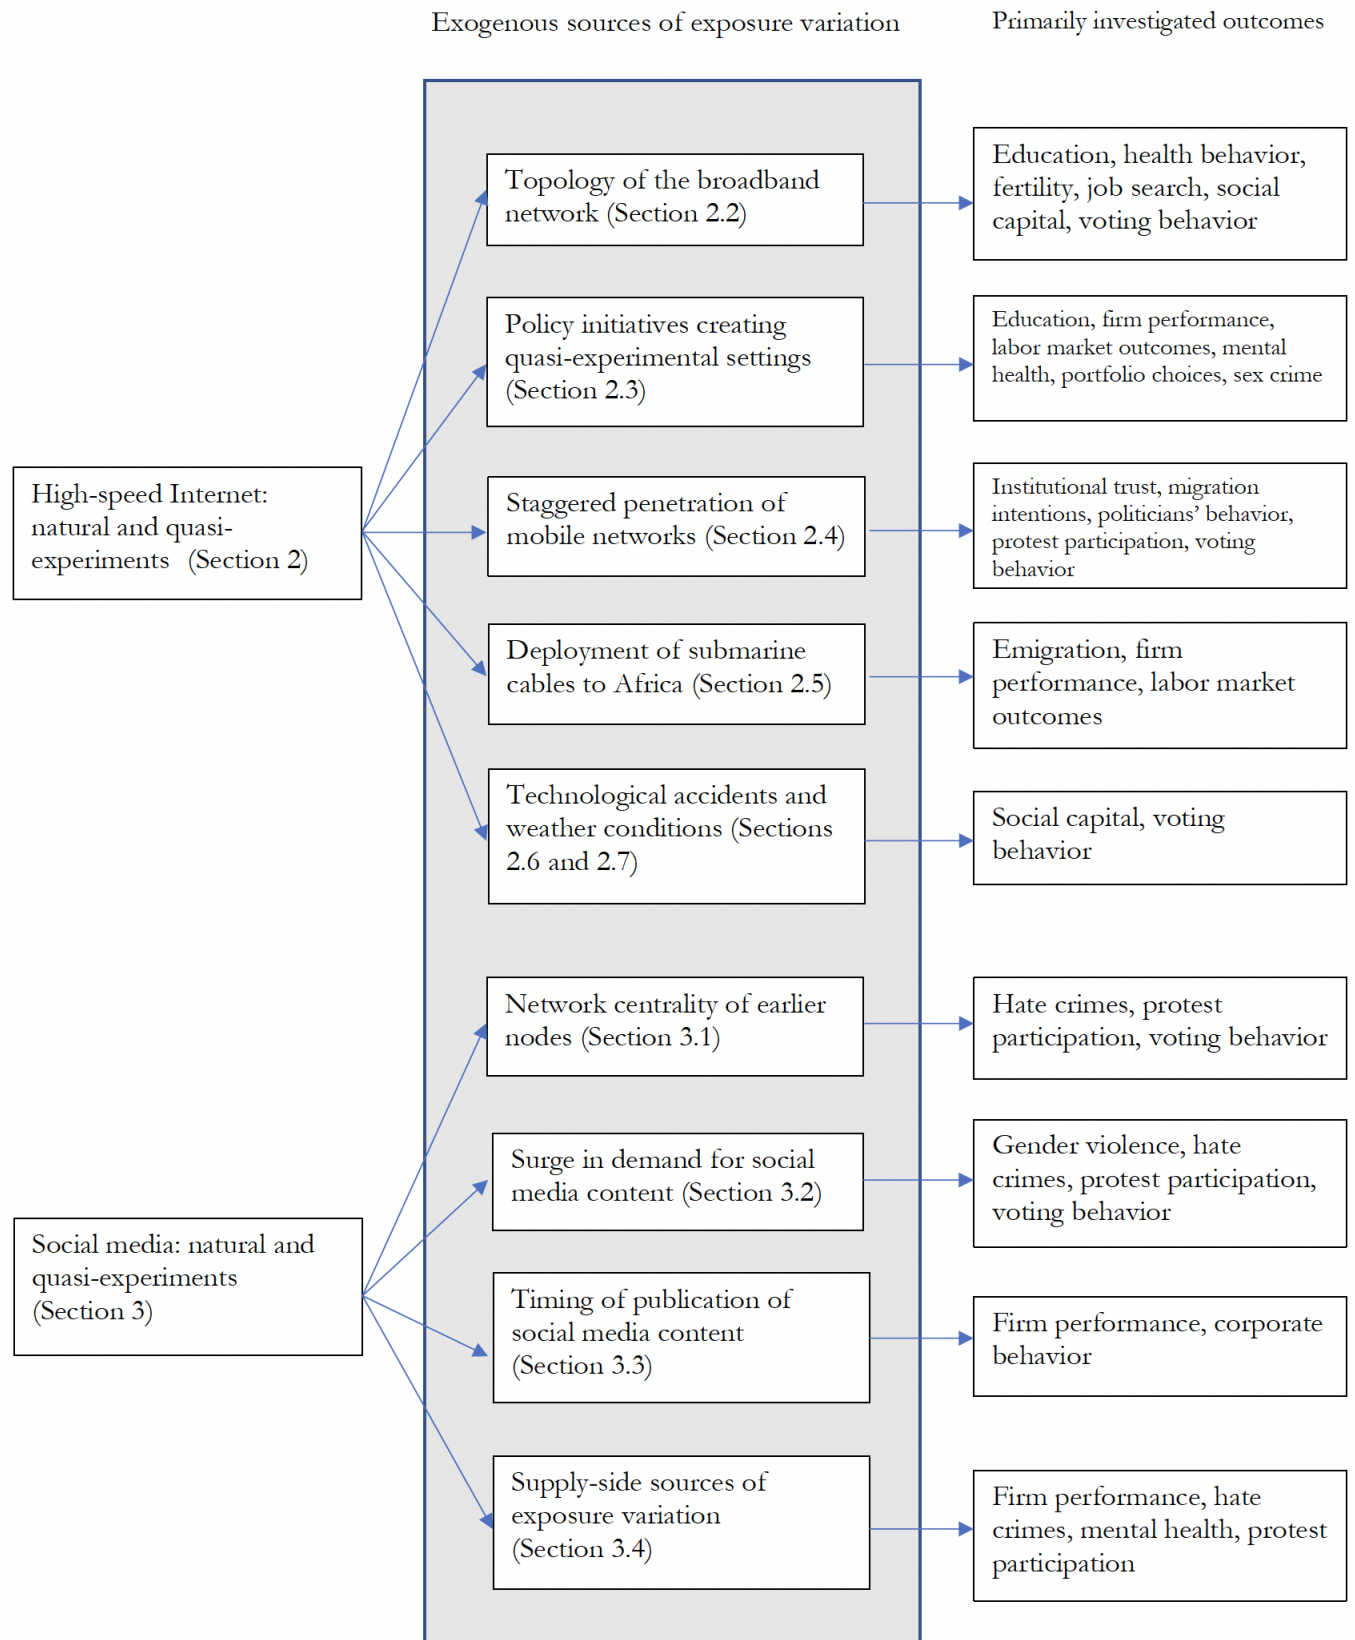
For the two strands of literature addressing the impact of high-speed Internet and social media, we allocate a subsection to each identification strategy, dissecting the methodological challenges, empirical design, and principal findings of the most prominent studies. Each subsection is complemented by a user-friendly table summarizing essential information about the studies discussed. Figure 1 outlines the structure of the paper as described above.

In most cases, identification relies on two-stage least squares estimates with the exogenous source of variation serving as an instrument for the explanatory variable. However, several studies, bolstered by highly granular data, adopt spatial regression discontinuity techniques. Other works adopt difference-in-differences frameworks, capitalizing on variations pre- and post-treatment, as well as spatial discontinuities in treatment assignment. While we detail the empirical outcomes of the considered studies, we forgo examining potential mechanisms of treatment effect transmission as such an inquiry extends beyond the scope of this paper.

Our literature review selects 60 studies, covering a vast geographical spectrum and time frame from 1995 to the present. The majority of the studies included are recent, dating from 2021 to 2023, which offers a comprehensive perspective on the cutting-edge research currently being conducted in this field. Overall, this undertaking allows us to highlight the key challenges and features of this literature while summarizing the causal relationships it has so far uncovered, providing an operative guide for researchers and non-academic specialists, such as journalists and policymakers, who want to delve into the behavioral, economic, and political impact of Internet penetration.

Several previous literature reviews have summarized the effects of the Internet on specific domains of individual or public life. Castellacci and Tveito (2018) conducted a cross-disciplinary review of the literature on the relationships between Internet use and well-being, focusing on the potential channels of transmission of the impact of Internet exposure. Goldfarb and Tucker (2019) focused on the economic outcomes of Information and Communication Technologies (ICT) penetration. They underscored the role of ICTs in reducing several types of marginal costs.

Figure 1: Structure of the paper



In their review of the economics of media and social capital, Campante et al. (2022) explored the impact of Internet and social media usage on trust, civic engagement, and political participation. While stressing the importance of ICTs in addressing coordination problems and enhancing political protest participation, the authors also highlighted the risk of polarization linked with the use of online networking platforms. Zhuravskaya et al. (2020) provided a focused review of the political effects of the Internet and social media. The authors suggested that social media initially served as communication and coordination platforms fostering political participation and enabling citizens to hold their governments accountable. However, they also noted that autocratic governments and anti-democratic movements can harness these platforms to bolster regime support, disseminate propaganda, spread misinformation, and undermine democracy in rival nations. In a review exploring the drivers of populism, Guriev and Papaioannu (2022) posited that the inherent design of social media tends to foster polarization and the spread of fake news and stereotypes. This feature may allow autocratic regimes to exploit the platforms for raising the public’s concern towards polarizing issues, such as immigration, thereby providing a conduit for populist narratives. In their effort to organize the existing research on the economic consequences of the Internet in developing countries, Hjort and Tian (2022) developed a framework to categorize the possible mechanisms underpinning the influences of online connectivity. They partitioned these mechanisms into two primary groups: supply-side and demand-side influences. Supply-side effects focus on the Internet’s impact on business performance and labor productivity. Demand-side effects primarily pertain to enhanced market access for firms, workers and consumers, or to mitigating information asymmetries.

Our work differs from previous attempts to review Internet studies by adopting methodological approaches as the key criterion for critically assessing the literature. This effort contributes to the interdisciplinary debate on the impact of the Internet by providing a guide for evaluating the causal effects of Internet and social media use. As a result, our study’s coverage overlaps minimally or not at all with previously reviewed literature. Though the studies we account for in this review mainly belong to the fields of economics and political science, the resulting guidelines can prove useful from any disciplinary perspective, including those of sociology, psychology, and anthropology. For a broader audience, our work provides simple ground rules to gauge the reliability and external validity of empirical studies that frequently draw media and political attention. From a policy perspective, understanding the causal effect of new information and communication technologies (ICTs) is critical for developing strategies that maximize the benefits of the Internet and social media while mitigating their potential adverse effects. Policymakers must discern how ICTs can be harnessed for positive social change, while also considering their potential misuse to undermine democracy and human rights.

The remainder of this paper is organized as follows. Section 2 presents an overview of studies

utilizing natural and quasi-experiments to assess the economic, political, and societal implications of high-speed Internet. Section 3 shifts focus to social media. Lastly, Section 4 draws some concluding remarks.

2 Broadband Internet: Natural and quasi-experiments

Discerning the causal impact of Internet use is challenging for several reasons. The first hurdle stems from the possibility of reverse causality. Potential outcomes of Internet use, such as an individual’s propensity for social engagement, civic participation, and political involvement, may catalyze increased Internet use. This interplay may create ambiguity about what behavior comes first and what follows as a consequence. In addition, unobserved individual or local characteristics may simultaneously determine both the decision to use the Internet and social media and the hypothetical outcomes of such use, thereby creating a spurious correlation that may be mistaken for a causal relationship. Even the processes of broadband and fiber wires roll-out is inherently endogenous, as it depends on the profitability of the related investments, which are in turn connected to the time-variant and time-invariant characteristics of the area. Under these circumstances, access to high-speed Internet is contingent on factors that also influence various economic decisions that could be misconstrued as outcomes of Internet and social media use. As a result, it becomes increasingly complex to understand whether the observed outcomes are directly influenced by Internet usage or merely concurrent with it due to these underlying factors.

The primary approach to address these issues is to isolate some exogenous sources of variation in access to fast Internet (or specific content), which does not stem from individual decisions and is orthogonal to the potential outcomes of Internet use under investigation. In this setup, naturally occurring or external factors beyond the control of researchers randomly assign individuals to an experimental and a control condition, setting the stage for a natural experiment. This setup allows researchers to estimate the impact of the exogenous component of Internet access tackling many of the endogeneity issues at stake. The identifying assumption of this method is that the exogenous driver of exposure to the treatment influences the hypothetical outcome exclusively through its relationship with the treatment itself. Any other potential causal pathways from the instrument to the dependent variable would constitute violations of the exclusion-restriction principle, thereby invalidating the identification strategy. Alternatively, in certain circumstances, an experimental framework involving the random assignment of treatment to a group may unintentionally arise as a collateral effect of particular policy decisions. If the assignment mechanism is randomized and orthogonal to the outcome, the policy measure sets up the conditions for a “quasi-experiment” that

allows researchers to identify the treatment effect.

Assessing the credibility of Internet studies relying on natural or quasi-experiments requires a meticulous examination of the identifying assumptions. This section, therefore, scrutinizes the sources of variation employed in the existing literature, with a specific emphasis on the exclusion restrictions. The section starts with a brief overview of the technological background of the initial broadband penetration phase, which relied on the Digital Subscriber Line (DSL) in most countries. Subsequently, we delve into studies that leveraged the topology of the initial broadband network in Section 2.2. Following sections offer a concise review of studies utilizing policy initiatives, the staggered diffusion of mobile networks, the deployment of submarine cables, technological mishaps, and weather conditions as exogenous determinants of individuals' capacity to utilize high-speed Internet.

2.1 Technological background

The 1990s saw the copper wires of the voice telephone network enabling a low-speed Internet connection via dial-up—a modem that established connection to a service provider by dialing a telephone number. However, a significant evolution occurred around 1995 with the advent of Digital Subscriber Line technologies (DSL), and subsequently, the Asymmetric Digital Subscriber Line (ADSL). These technologies utilized a broader range of frequencies over the copper line, facilitating faster and cost-effective Internet access via the voice network. While the initial iterations of ADSL achieved only a modest speed and did not meet the broadband criteria, technological enhancements in the subsequent years greatly improved this situation (Nardotto et al., 2015). By the dawn of the 2000s, ADSL had achieved a broadband-qualifying speed of 2Mbit/s, marking a significant milestone in the process of broadband penetration (Ofcom, 2010; 2016).

Telephone companies upgraded all connections between the LEs, which account for the majority of the distance between the end-user and the Internet Service Provider (ISP), using fiber-optic wires. However, the connections between the LEs and the premises continued to rely on the existing copper infrastructure. This approach facilitated decent Internet speeds with relatively low investment costs, yet it inadvertently led to exogenous local differences in the quality of Internet access among households, dependent on the length of their “last mile”, i.e., the network section that connects individual premises with the local area’s serving node, also known as “local exchange” (LE), “Central Office” or “Main Distribution Frame” (MDF). The digital signal experiences substantial decay when transmitted over copper wire, with signal strength diminishing more than proportionally with the distance traveled. As a result, a longer distance from the LE led to slower connection speeds.

Installed post World War One in many countries, the LEs were originally designed to minimize average voice signal decay rather than for the transmission of digital signals. As a result, the LEs' catchment areas are irregularly shaped, often with the LEs not centrally located within them. Consequently, local access conditions can differ significantly even within relatively small areas. In an unforeseen turn of events, an unpredictable technological innovation unexpectedly transformed the topology of the post-WW1 telephone network into an instrumental, and exogenous, determinant of Internet connection quality. Primarily, the distance from the area's servicing LE became the key driver of connection speed. As we discuss in Section 2.2, several studies have leveraged the discontinuities in fast Internet access that the topology of the broadband network induced across households to examine the impact of the Internet on various individual and societal outcomes during the first phase of broadband penetration.

More recently, fiber optic cables enabled faster Internet access. However, the deployment of fiber lines necessitates substantial investment, thus presenting service providers with challenges in the swift and uniform dissemination of the technology. While new discontinuities in access to super-fast Internet have arisen, they can't be considered entirely exogenous since the deployment of the new technology primarily hinges on the profitability of related investments. Nevertheless, public policies supporting the rollout of new communication technologies have often inadvertently set up quasi-experimental situations. These circumstances randomly allocate specific areas or communities the opportunity to access the new technology. This has facilitated a recent surge in studies examining the impact of quicker Internet access on an array of individual and societal outcomes. We will succinctly review these studies in Section 2.5.

2.2 The topology of initial broadband networks

Several studies have used information on the topology of the voice network to ascertain the causal impact of access to fast Internet on individual and collective behaviors. As mentioned in Section 2.1, the efficacy of DSL connections diminishes with increasing distance from the network node serving the area, known as the local exchange (LE). Given that the telephone network was established prior to the advent of the Internet, when the voice signal quality bore no relation to the length of the copper wires, users were unable to strategically position themselves close to the network nodes. The irregular configuration of the catchment areas further enhanced the random allocation of the treatment, thereby establishing conditions conducive to a natural experiment in numerous countries.

Relying on distance from the local exchange as a source of exogenous variation in exposure to fast Internet presents several implications. Primarily, information regarding the architecture

of the network must be combined with high-resolution data measuring the hypothetical outcomes of Internet use. A precise geolocation of survey respondents is critical to estimate the distance between the end-users' telephone lines and the LE serving their area of residence. Municipality-level geolocation poses numerous challenges as larger cities host multiple local exchanges. This necessitates the construction of latent indicators to determine the distance between end users and local exchanges, which may give rise to measurement inaccuracies, reduced precision, and data loss. Information about the city block in which the survey respondent resides would provide an ideal setting to estimate the speed their Internet connections as a function of the distance from the LE. Nevertheless, the use of survey data rarely permits comparisons of individuals across the borders of LE catchment areas, mainly due to the difficulty in obtaining an adequate number of observations on both sides of the border.

An alternative approach involves measuring the behavioral outcomes of broadband access through the use of administrative data. Even in this instance, data granularity plays a pivotal role. While using administrative data might limit direct observation of individual behavior, it can facilitate a comparison of areas adjacent to the borders of catchment zones in a Regression Discontinuity Design (RDD) setting. This strategy is useful in leveraging geographical discontinuities for robust identification of broadband access effects (e.g., Ahlfeldt et al., 2017; Amaral-Garcia et al., 2022).

Secondly, due to the inherent shortage of data regarding online activities, studies that combine information on the broadband network's topology with survey or administrative data often rest on the assumption that all individuals are uniformly and effectively exposed to the treatment, as in an intention-to-treat (ITT) design. As we will see, virtually the entire literature employing naturally arising experimental settings to assess the impact of the Internet and social media needs to adopt this ITT approach.

Another complication stems from the waning significance of DSL technology with the rollout of alternative methods to access high-speed Internet. As a result, distance from the local exchange may no longer be an appropriate instrument for evaluating the impact of broadband access once these new technologies begin to proliferate. Additionally, evidence suggesting broadband availability may have influenced house prices in the latter stages of broadband penetration (Ahlfeldt et al., 2017) further amplifies self-selection concerns.

Despite these concerns, the technological shift does not undermine the validity of the exclusion restrictions in assessing the impact of high-speed Internet during its initial decade of penetration or beyond. Furthermore, findings suggest that behavioral changes triggered by the initial phase of high-speed Internet penetration may persist following the introduction of non-DSL Internet access methods (Geraci et al., 2022). However, the completion of the DSL penetration process calls for an examination of the effect of even faster Internet connections, such as those brought about by

the deployment of optical fiber cables or the proliferation of new mobile technologies. We will delve deeper into studies addressing the staggered penetration of ultra fast broadband and mobile technologies in Section 2.3 and 2.4, respectively.

Table 1 summarizes some key information about the studies that leveraged the topology of the voice network as an exogenous source of variation in Internet exposure. In the first, column, we report the authors and context of the study. In Column 2 (“Instrument”), we report the exogenous source of variation in the treatment that the authors exploit for identification and its level of aggregation in parentheses. Columns 3 and 4 contain information about the treatment and the outcome variable, and their level of aggregation in parentheses. The data sources for measuring these variables will systematically be mentioned in the body text.

The broad picture that emerges from this body of work is that the initial phase of broadband penetration triggered significant shifts in individual behavior across all aspects of life. These range from civic engagement and political participation to improvements in educational and labor market outcomes, and alterations in health behavior and fertility decisions.

In an early study on the impact of broadband penetration, Czernich (2012) examined the influence of broadband internet on political participation in the 2005 German Federal Election using municipal-level data. The author compared voting behavior during the DSL era with the pre-DSL era in a first-difference setting, focusing on municipalities without a local exchange (LE). The empirical analysis employs a two-stage least squares (TSLS) setting in which DSL diffusion, measured as the percentage of households in a municipality for which DSL is available, is instrumented with the distance between a municipality and its nearest LE. Although the second stage revealed a positive coefficient for the DSL rate on voter participation, it was not significant. Falck et al. (2014) expanded this framework by merging administrative data on election outcomes in Germany at various times with data on the share of households for which a DSL connection was technically available in the election year, across approximately 12,000 German municipalities. The authors used a dummy variable indicating municipalities with distances exceeding the threshold of 4,200 meters to instrument DSL availability in a TSLS setting that covered municipalities without an LE. The results showed a negative Internet effect on voter turnout in Western Germany, but not in Eastern Germany, with the impact primarily coming from non-local elections. Additionally, they noted a potential negative influence of DSL availability on extremist right-wing parties’ vote shares.

Similarly, Campante et al. (2018) investigated the impact of high-speed Internet diffusion on various forms of political participation, using municipal data from Italy from 1996 to 2013. The authors measured broadband access as the percentage of households with access to ADSL-based services. To tackle causality, they instrumented broadband access with the interaction between a municipality’s distance to the nearest underground station (UGS) and a dummy for the post-2001

period (i.e., after the introduction of broadband Internet) in a differences-in-differences (DID) approach. The findings suggested that broadband Internet significantly reduced turnout in Italian parliamentary elections until 2008. However, it was positively correlated with other forms of political participation, both online and offline, such as the rise of local online grassroots protest movements. The negative effect on parliamentary election turnout essentially reversed after 2008 when local grassroots movements evolved into the Five-Star Movement electoral list.

In a distinct political context, Miner (2015) examined the impact of the Internet on the semi-authoritarian incumbent government in Malaysia’s 2008 elections. The author measured Internet penetration as the number of IP addresses per voter between 2004 and 2008. To address the endogeneity of IP address density, the author employed the distances from a district’s centroid to Malaysia’s three largest ISP backbones as an instrumental variable in a TSLS setting.¹ The estimates suggest that that exposure to the Internet could account for a significant portion of the swing against the incumbent party in the 2008 election.

Several other studies have used distance from the network node as an exogenous source of variation of Internet exposure. Gurtzgen et al. (2021) and Gürtzgen et al. (2021) combined information on the geolocation of Germany’s broadband network nodes with the Integrated Employment Biographies of German workers provided by the Federal Employment Agency to assess the impact of the Internet on workers’ employment prospects between 1998 and 2008. The authors instrumented Internet availability, gauged as the share of households with the possibility to access a DSL connection in a municipality, with the distance of the nearest LE from each municipality’s center (population-weighted). This instrumental variable was set as a dummy assuming a value of one for municipalities with distances above the threshold of 4,200 meters, corresponding to slower connections. The authors’ results suggest that broadband Internet improves reemployment rates after the first months in unemployment for males but has a limited impact on the stability of new matches and their wages.

Geraci et al. (2022) conducted an in-depth analysis of how broadband penetration influenced various forms of social capital in the UK between 1997 and 2017. They merged information on the topology of the British network with geocoded longitudinal data related to individual behaviors and opinions, obtained from the British Household Panel Survey (BHPS) from 1997–2008 and the UK Household Longitudinal Study (UKHLS) from 2009 to 2017.

Implementing a Difference-in-Differences (DID) approach, the authors assessed how the distance between each household’s Lower Layer Super Output Area (LSOA)² and the respective

¹The sample accounts for electoral districts in Peninsular Malaysia, exception made for the district of Kuala Lumpur, and excludes the insular regions of Sabah and Sarawak. In Peninsular Malaysia, there are 428 electoral districts.

²LSOAs are the second-narrowest geographical areas in the UK census, comprising on average 650 households and 1,500 inhabitants. In densely populated metropolitan areas, they correspond to portions of city blocks.

Local Exchange (LE) affected the social capital of survey respondents following the uptake of broadband penetration. The results revealed that the advent of high-speed Internet caused a significant decrease in civic and political engagement, suggesting that the use of fast Internet crowded out several aspects of social capital.

Billari et al. (2019) explored the impact of DSL access on individual fertility choices by combining information on the topology of the German broadband network with geolocated longitudinal survey data from the Socio-Economic Panel (SOEP). Utilizing a TSLS framework, they instrumented Internet access—defined as the availability of a DSL subscription at home—with a binary indicator capturing whether the respondent lived more than 4.2 km away from their local exchange. Their empirical analysis highlighted positive effects of broadband availability on the fertility of highly educated women. In a related study, Billari et al. (2018) demonstrated that broadband access could negatively impact sleep duration and satisfaction in Germany.

Adopting an alternative approach, some studies utilized administrative data to conduct Regression Discontinuity Design (RDD) analyses, assessing how Internet outcomes vary across postcodes situated near the invisible borders of catchment areas based on their distance from the serving local exchange. This method capitalizes on the irregular shape of catchment areas to compare highly similar regions on either side of the border, each served by a different local exchange.

A pioneering study in this regard is Faber et al. (2015), who coupled data on the topology of the UK broadband network with administrative test score records for English primary and secondary school students between 2000 and 2008. Their estimation sample comprised roughly 580,000 residential postcodes located within 1 km of an exchange station boundary, thereby yielding a sample of approximately 20,000 boundaries. After determining the boundary effect on jumps in distances from the LE and broadband access, the authors estimated the effect of distances between residential postcode centroids and their connected telephone exchange stations on student educational attainment. Their results suggest that faster Internet led to increased student consumption of online content but had no significant effects on the time students spent studying, either online or offline, or on their learning productivity. In a follow-up study, Sanchis-Guarner et al. (2021) narrowed the analysis to the period from 2005 to 2008, specifically focusing on 14-year-old students who sat for the Key Stage 3 (KS3) national examination. They employed a fuzzy spatial regression discontinuity (SRD) design to ascertain the correlation between connection speed and students' academic performance. This refined approach enabled the authors to revise the previously null result, leading to a more accurately estimated positive effect of broadband availability on test scores.

Table 1: Studies using information on the topology of the network

Authors	Instrument (level)	Treatment (level)	Outcome (level)
Ahlfeldt et al. (2017) England 1995-2010	Distance from the LE (postcode)	Connection speed (postcode)	Property prices (postcode)
Amaral-Garcia et al. (2022) UK 2000-2011		Distance from the LE (LSOA)	C-sections childbirths (individual)
Billari et al. (2018) Germany 2008-2012	Distance from the LE > 4.2Km (individual)	DSL subscription (individual)	Sleep duration (individual)
Billari et al. (2019) Germany 2008-2012	Distance from the LE > 4.2Km (individual)	DSL subscription (individual)	Fertility (individual)
Campante et al. (2018) Italy 1996-2013	Distance from the UGS (municipality)	DSL share (municipality)	Voter turnout, participation in 5SM meetings (municipality)
Czernich (2012) Germany 2002-2005	Distance from the LE (municipality)	DSL share (municipality)	Voter turnout (municipality)
De Stefano et al. (2023) UK 2000-2004	Distance from the LE < 5.5Km (firm)	DSL access (firm)	Firm performance (firm)
Donati et al. (2022) Italy 2001-2013	Distance from the UGS (municipality)	DSL share (municipality)	Mental health (municipality)
Faber et al. (2015) UK 2002-2008		Distance from the LE (individual)	Educational attainment (individual)
Falck et al. (2014) Germany 1995-2008	Distance from the LE (municipality)	DSL share (municipality)	Voting behavior (municipality)
Geraci et al. (2022) UK 1997-2017		Distance from the LE (LSOA)	Social capital (individual)
Gürtzgen et al. (2021) and Gurtzgen et al. (2021) Germany 1998-2008	Distance from the LE > 4.2Km (municipality)	DSL share (municipality)	Job search & matching (individual)
Miner (2015) Malaysia 2004-2008	Distance from the three largest ISP backbones (electoral district)	IP addresses per voter (electoral district)	Voting behavior (electoral district)
Sanchis-Guarner et al. (2021) England 2005-2008	Distance from the LE (postcode)	Connection speed (postcode)	Educational attainment (individual)

Ahlfeldt et al. (2017) merged geolocated information about English LEs with transaction data pertaining to mortgages granted by the Nationwide Building Society between 1995 and 2010 to assess the influence of broadband penetration on housing prices. In order to identify the effect of broadband availability, the authors exploited the abrupt variations in Internet speed over time and across LE boundaries, which were determined exogenously by the distance of a property from

its servicing LE. The empirical analysis demonstrated a significantly positive impact of high-speed Internet on property prices, albeit with diminishing returns to speed. In a similar vein, Amaral-Garcia et al. (2022) exploited cross-border spatial discontinuities in Internet quality to assess the impact of broadband penetration on childbirth procedure choices (Cesarean section or vaginal delivery) and subsequent health outcomes in England from 2000 to 2011. They combined data on the English broadband infrastructure’s layout with the UK Department of Health’s Hospital Episode Statistics (HES). Using a Regression Discontinuity Design, the authors compared childbirth choices of mothers in similar LSOAs separated by the invisible boundaries of LE catchment areas. This strategy enabled the authors to ascertain a causal impact of internet quality on the likelihood of opting for C-sections—a choice that could help mothers avoid numerous post-childbirth complications.

De Stefano et al. (2023) utilized the discontinuities in Internet access quality across catchment areas to evaluate the impact of DSL penetration on diverse facets of UK firms’ performance. The performance metrics were obtained from the Annual Business Inquiry (ABI), provided by the UK Census Bureau, for the period 2000 to 2004. The authors implemented a fuzzy discontinuity approach to ascertain whether a binary indicator—designated as one if the distance from the LE was less than 5.5 km—impacted firms’ performance on either side of the borders. The evidence gathered suggested that while broadband availability caused an increase in firm size, it did not have a significant impact on labor productivity.

Donati et al. (2022) investigated the impact of Internet penetration on adolescent mental health in Italy by merging municipal-level data on broadband access with hospital discharge records for the period from 2001 to 2013. They employed a TSLS framework, where municipal-level DSL coverage was instrumented with the municipality’s proximity to the nearest high-level telecommunication exchanges (Urban Group Stage, UGS), interacted with a dummy variable for the post-2001 period. Their empirical evidence suggests that access to high-speed Internet significantly influences the incidence of mental disorders among young cohorts, but not among older ones. Specifically, the availability of internet connectivity corresponds to an increase in diagnoses of depression, anxiety, drug abuse, and personality disorders for both genders, and to a rise in eating and sleep disorders exclusively for females.

With a few exceptions, the studies summarized above encompass a relatively limited timeframe, substantially corresponding to the phase of DSL broadband take up. As the process of DSL penetration is nearing completion in many countries, there arises the need to scrutinize the implications of even swifter Internet connections, including those resulting from optical fiber cables deployment or the spread of new mobile technologies for rapid Internet browsing.

2.3 Policy initiatives for broadband diffusion

A growing body of literature leverages policy initiatives as exogenous sources of variation in access to fixed-line Internet by citizens and firms. The underlying principle of these studies is that end users can only avail themselves of broadband connections if such services are accessible within their residence or operational area. Numerous governments have initiated plans to boost broadband penetration, particularly in rural regions, while encouraging more widespread adoption in urban areas. These efforts often encompass the establishment of necessary infrastructure, including the creation of new local exchanges for DSL connections and the deployment of optic fiber cables connecting local exchanges. In some cases, fiber is even brought directly to the users' premises. Other common public interventions focus on digitizing government services to incentivize citizens to enhance their digital skills and adopt home broadband connections.

This body of research often operates under the assumption that policy initiatives cannot guarantee a uniform and simultaneous nationwide deployment and adoption of new technologies. Instead, constraints such as limited funding, geographical characteristics, and technical challenges necessitate a phased approach, resulting in staggered technology penetration. Such constraints may set the stage for a quasi-experimental framework, enabling comparisons between areas that gained earlier access to broadband technologies and those still awaiting these services.

However, employing policy initiatives to evaluate the Internet's impact entails several implications. Notably, the phased roll-out of these publicly funded programs seldom follows a random trajectory. Areas with higher development and demand for Internet services often gain earlier access, introducing endogeneity concerns. Consequently, a crucial aspect of these studies lies in transparently verifying the randomness of the staggered funding and technology penetration. Additionally, akin to studies outlined in the preceding section, the dearth of granular information regarding individuals' online activities necessitates the adoption of intention-to-treat-like (ITT) designs. Furthermore, the use of high-resolution data is pivotal. Although cross-country and cross-region studies offer valuable insights, they may find it challenging to adequately address the common bias induced by numerous macro-level confounding factors.

In this section, we present a summary of works investigating the impact of fixed-line connections in relation to specific policy initiatives. The subsequent section will delve into the influence of new mobile technologies. Table 2 provides key details about studies that have utilized the policy-induced staggered penetration of new technologies for fixed-line connections. The first column reports the authors and context of the study. In columns 2, 3, and 4, we provide information about the exogenous source of variation in the treatment, the treatment itself, and the outcome variable, respectively. Their level of aggregation is in parentheses.

Bhuller et al. (2013) were among the first to exploit a policy measure to assess the impact of fast

Table 2: Studies using policy initiatives for fixed-line broadband

Authors	Instrument (level)	Treatment (level)	Outcome (level)
Akerman et al. (2015) / Akerman et al. (2022) Norway 2001-2008	Broadband availability rate (municipality)	Broadband adoption (firm)	Worker productivity / Exports (firm)
Arenas-Arroyo et al. (2022) Spain 2007-2019	Fiber lines at $t - 1$ (province)	Fiber lines at t (province)	Adolescent mental health (province)
Bhuller et al. (2013) Norway 2000-2008	Share of households covered by broadband at time $t - 1$ (municipality)	Share of households covered by broadband at time t (municipality)	Sex crime(municipality)
Bhuller et al. (2023a) Norway 2000-2014		Broadband availability rate (municipality)	Workers' labor market outcomes (firm)
Bhuller et al. (2023b) Norway 1997-2010	Share of households covered by broadband at time $t - 1$ (municipality)	Share of households covered by broadband at time t (municipality)	Newspapers firms performance (firm)
Boeri (2023) Italy 2012-2022		Share of households with access to speed $\geq 30Mbit/s$ (school catchment area)	Educational attainment (individual)
Cambini et al. (2022) Italy 2012-2017	Distance from the nearest $OLT \leq 10km$ (municipality)	Share of household with ultra fast broadband (municipality)	Educational attainment (individual)
Hvide et al. (2023) Norway 2000-2010	Share of households covered by broadband infrastructure (municipality)	Share with broadband subscription (municipality)	Portfolio choices (individual)
McDool et al. (2020) UK 2012-2017		Increase in average connection speed between waves (LSOA)	Youth mental health (individual)
Zuo (2021), U.S. 2012-2015		Discounted price broadband availability/eligibility (census block/individual)	Employment, income (individual)

Internet. Focusing on the case of Norway, the authors tested the relationship between broadband adoption rate and the incidence of sex crimes across Norwegian municipalities over 2000-2008. As an exogenous source of variation in the share of households possessing a broadband connection, the authors leveraged the National Broadband Policy, a Norwegian public program funding the expansion of the broadband infrastructure of Telenor, the state-owned telecom company. Local governments were eligible for funding from this program by submitting a project plan to be assessed by an independent committee of experts. Upon approval, the government financially supported the first phase of infrastructure deployment, thereby covering the initial fixed costs. Due to limited funding, the infrastructure was deployed gradually, meaning the necessary access points were established in different municipalities at different times. While the funds allocation mechanism indicates that the staggered penetration of this new technology did not follow a random trajectory, the authors took advantage of this spatial and temporal variation in the availability of broadband internet to identify the impact of fast Internet in an instrumental variables (IV) setting. Based on the assumption that timing of the roll-out is unlikely to co-vary with the key correlates of sex crime, they instrumented the fraction of households with broadband internet subscription with the fraction of households that were covered by broadband infrastructure in the previous year. Following this approach, the authors provided evidence of a significant and substantial association between Internet use and reported sex crimes.

The National Broadband Policy of Norway has served as a basis for numerous other studies. Akerman et al. (2015) delved into how the adoption of broadband Internet affected labor market outcomes and productivity among different types of workers between 2001 and 2007. The authors combined data on broadband subscriptions from a stratified random sample of firms, municipality-level information on household broadband availability, and administrative records concerning firms and employment relationships provided by Statistics Norway. Thus, they implemented an instrumental variable IV approach where the firm-level broadband adoption was instrumented with the municipality-level broadband availability. Their findings indicate that the integration of broadband resulted in an increase in the marginal productivity of skilled workers, while it decreased the marginal productivity of unskilled workers. Expanding on the same policy background and framework, Akerman et al. (2022) provided evidence that broadband adoption also made trade patterns counter-intuitively more sensitive to distance and economic size.

Leveraging the same policy context, Bhuller et al. (2023a) examined the influence of broadband penetration on the labor market outcomes of workers in Norway from 2000 to 2014. Their study yielded several significant findings. Firstly, the expansion of internet accessibility led to a surge in online job advertisements and an increase in the proportion of households scanning these ads online. This growing appeal to online connectivity notably bolstered the recruitment process, leading to a substantial decrease in the average vacancy duration. Unemployed job seekers saw

their prospects markedly improved as a result of enhanced internet accessibility. This improvement manifested as an increased likelihood of securing employment spells exceeding one year, along with a considerable reduction in the subsequent chances of experiencing repeated unemployment spells.

In the same spirit, Bhuller et al. (2023b) combined municipality-level data on the diffusion of broadband Internet plausibly prompted by the National Broadband Policy with data on newspaper firms and consumption of print media provided by Norwegian Media Businesses' Association (MBL) and the Norwegian Media Authority. The authors capitalized on the spatial and temporal variations in broadband availability across municipalities, instrumenting the percentage of households using broadband Internet at the start of a given year with the corresponding percentage of households having broadband access at the outset of the preceding year. Their empirical findings indicate that a significant portion of the decline in newspaper sales witnessed between 2000 and 2010 can be attributed to the rise in Internet adoption. Furthermore, the authors provided evidence that newspaper companies reacted to this shift by drastically slashing costs, either by diminishing labor inputs or reducing the physical size of their newspapers, thereby averting profit losses.

Similarly, Hvide et al. (2023) matched data about municipal broadband availability with information from the Norwegian Central Securities Depository (NCSD) to ascertain the influence of Internet usage on individual portfolio choices between 2000 and 2010. The authors used the proportion of households covered by broadband infrastructure as an instrumental variable for the share of households with a broadband subscription in a specific year. Their findings indicate that Internet usage significantly bolstered participation in the stock market.

Arenas-Arroyo et al. (2022) applied an analogous analytical framework to determine the effect of ultra-fast Internet on adolescent mental health in Spain spanning the years 2007 to 2019. The identifying assumption in this study is that the staggered deployment of the fiber infrastructure across the country was due to a combination of regulatory reasons and providers' competitive strategies that, according to the authors, were largely unrelated to end users' propensity for online connectivity. The analysis employed the number of installed optic fiber lines (irrespective of usage) in each province at time $t - 1$ to instrument a series of dummy variables denoting five distinct levels of fiber penetration, within a two-stage least squares (TSLS) framework. This approach enabled the authors to discern a causal relationship between Internet usage and the incidence of behavioral and mental health (BMH) cases at the province level, as measured through hospital discharge diagnoses. Consistently with Donati et al. (2022), the outcomes of particular concern included anxiety, mood disorders, drug abuse, self-harm, and suicide attempts, with the majority of the effects attributed to girls.

Utilizing a similar methodology, Boeri (2023) exploited the phased roll-out of optic fiber cables under the National Ultra-Broadband Plan (NUPB) established by the Italian government

to investigate the influence of high-speed Internet on students' national test scores (INVALSI). Using census data, the author identified school catchment areas and determined the proportion of families within these zones with access to a minimum speed of 30 Mbit/s - a standard which constitutes fiber-to-the-node (FTNN) connectivity. This framework rests on the assumption that the FTNN broadband expansion was largely driven by the NUBP staggered implementation across contiguous territories, implying that the timing of deployment was independent of other potential factors impacting educational outcomes. This approach allowed the author to test for their impact on enrolled students' test scores. The study found no significant impact of 30 Mbit/s broadband on student performance in either numeracy or literacy. However, it uncovered a positive and significant effect among low-performers in the previous grade and students originating from socioeconomically advantaged backgrounds.

In the same spirit, Cambini et al. (2022) took advantage of the phased introduction of ultra-fast broadband in Italy to evaluate its impact on students' performance.³ The authors combined INVALSI test scores data from 2012 to 2017 with the proportion of households within a municipality having access to ultra-fast connections, ostensibly linked to the governmental program's rollout. Given the plausible endogeneity of ultra-fast broadband diffusion, the treatment variable was instrumented with a binary indicator, with a value of one denoting municipalities within a 10 km radius from the nearest Optical Line Terminal (OLT). This indicator served as an estimate of the cost of deploying optic fiber cables in the area. The empirical results imply that an increase in broadband availability negatively affects the academic performance of 8th-grade students, particularly boys, whereas it exerts no discernible effect on 2nd and 5th-grade students.⁴

McDool et al. (2020) adopted an alternative procedure to track the penetration of ultra fast broadband in the UK. Combining information on the average connection speed at the LSOA level with geolocated survey data from the UKHLS, the authors proxied the deployment of ultra fast Internet technologies by estimating the increase in connection speed observed between survey waves in each LSOA. The identifying assumption underlying the empirical analysis is that, during the period of analysis, ultra fast broadband penetration was driven by a public program that substantially distorted commercial demand driven broadband installation. The authors assume that, as a result broadband performance enhancement was not tied to area-level background variables

³Although Cambini et al. (2022) employ the municipality-level distance from the Optical Line Terminal (OLT) as an instrumental variable for the proportion of households with ultra-fast broadband access, their work is included in this section rather than in Section 2.2, which discusses the initial broadband network's topology. The rationale for this classification is that the distance from the OLT is used as a proxy for the cost of deploying optic fiber technology, not as a measure of signal strength, which is not affected by the distance traversed over the fiber cables.

⁴Another study that combined test scores and college application data with information on Internet availability is Dettling et al. (2018), who explored the correlation between the presence of at least one Internet provider with students' achievement across US postcodes. The identifying assumption of the study is that the binary measure of high-speed Internet availability is exogenous to a student's testing and application outcomes.

such as the unemployment rate, income per capita, and other socio-economic characteristics. After providing evidence supporting this assumption, the authors assessed the influence of emerging ultra-fast internet technologies on youth mental health between 2012 and 2017. In line with the findings of Arenas-Arroyo et al. (2022) and Donati et al. (2022) the empirical results indicate a negative correlation between Internet use and well-being across numerous domains. The most pronounced effect pertains to children’s self-perception of their appearance, with the impacts being more detrimental for girls than boys.

Zuo (2021) shifted the focus from policies that support broadband infrastructure to a U.S. initiative known as “Internet Essentials” encouraging broadband demand. Launched in 2012 by Comcast, the largest Internet service provider by subscriber count, this program aimed to offer broadband access at a reduced price for low-income families. The possibility to enroll in the program was determined by geographic availability, the timing of its introduction, and individual eligibility. Zuo, 2021 exploited these varying factors in a triple-difference framework, comparing the outcomes of eligible and ineligible individuals across areas with varying Comcast broadband coverage rates, before and after the program’s introduction. By combining data on local Comcast broadband coverage rates at the census-block level with individual-level labor market outcomes drawn from the American Community Survey (ACS), the author was able to estimate that the program significantly increased the likelihood of employment and household income.

2.4 From GSM to 5G: the staggered penetration of mobile Internet

The research we discuss in this section delves into the implications of second-generation (2G) and third-generation (3G) mobile technologies. The 3G protocol was designed to provide accelerated data access rates for mobile internet. There are several differences between 2G and 3G mobile network infrastructure. To begin with, 3G networks needs a greater number of antennas. Additionally, the equipment required in antennas to uphold 3G services (UMTS) is significantly different, implying the need to upgrade older nodes. Lastly, the 3G antennas necessitate a fiber-optic cable connection to the national backbone network infrastructure, typically installed underground. Considering these characteristics, the roll-out of 3G infrastructure necessitates substantial investments and could be curtailed by considerations of cost-effectiveness (Donati, 2023).

The uneven expansion of mobile Internet infrastructures, like that of fiber cables, has often resulted in a significant geographical variation in the ability to access fast mobile connections. This irregular distribution has allowed numerous researchers to estimate the causal effects of mobile Internet on a range of behavioral and political outcomes. Approaches that utilize the staggered timing and direction of mobile technology diffusion as a source of exogenous variation in its use

share similar assumptions with the studies we have summarized thus far. Due to the absence of microdata on individuals’ online activities, an ITT design is typically adopted. Most crucially, the central assumption in any attempt to discern the causal impact of mobile Internet based on the phased introduction of associated infrastructures rests on the randomness of end-user exposure to the new technology. However, given the cost of infrastructure deployment, private companies often introduce new technologies in areas with high Internet service demand, typically more developed and dynamic regions. This circumstance implies the convenience to retrieve exogenous interferences with deployment patterns. Finally, the high-definition of data plays a crucial role in giving the possibility of ruling out the potential effect of confounding factors.

In this section, we highlight a selection of studies that leveraged the phased deployment of mobile Internet technologies to evaluate their impacts. Essential characteristics of these studies are outlined in Table 3. As usual, the first column reports the authors and context of the study. Columns 2, 3, and 4 mention the exogenous source of variation in the treatment, the treatment itself, and the outcome variable, respectively. Their level of aggregation employed in the empirical analysis is in parentheses. These works primarily examined the political consequences of the expansion of third-generation (3G) mobile networks. 3G, being the inaugural generation of mobile broadband Internet that enabled users to access the web freely via their smartphones and stream or upload videos, acted as a pivotal catalyst for the swift proliferation of social media. The bulk of these analyses takes advantage of the phased deployment of mobile networks in the spirit of difference-in-differences analyses (DID), reinforcing the causal interpretation of their findings through two-stage least squares (TSLS) estimates that harness an exogenous source of variation in mobile coverage.

In their work on the political effects of mobile technology in Africa, Manacorda and Tesei (2020) focused on second-generation (2G) technology, which enables basic Internet connectivity. The authors combined information concerning local mobile phone coverage with three datasets on individual protest activities largely dependent on news sources⁵ to ascertain how mobile penetration influenced collective actions and political mobilization within geographic cells of $0.5^\circ \times 0.5^\circ \simeq 55km$ resolution. In order to mitigate causality concerns, the authors leveraged the observation that frequent electrostatic discharges during storms tend to damage mobile phone infrastructures, especially ground antennas responsible for local signal transmission, thereby adversely impacting connectivity. This interference affects both sides of the market, causing a reduction in supply due to the expensive nature of surge protection and diminished profitability from poor connectivity, alongside a decrease in demand as the prospect of inconsistent communications deters adoption

⁵Global Database on Events, Location and Tone (GDELT), the Armed Conflict Location and Event Data Project (ACLED), and the Climate Change and African Political Stability (CCPAS) Social Conflict Analysis Database (SCAD).

(Manacorda and Tesei, 2020).

Using National Aeronautics and Space Administration (NASA) satellite-generated data on the incidence of lightning, the authors showed that areas with higher than average incidences of lightning display slower adoption of mobile phone technology. The research design therefore involved TSLS estimates, instrumenting mobile phone coverage with the interaction between the average number of lightning strikes in a cell from 1995–2010 and a linear time trend representing the widespread increase in mobile phone uptake across the continent. The authors’ findings indicate that mobile phones play a crucial role in mass mobilization, which predominantly occurs during economic downturns when grievances become apparent and the opportunity cost of participation lessens.

Guriev et al. (2020) turned to the global expansion of third-generation (3G) networks to assess how mobile broadband Internet affected support for governments worldwide over 2007–2018. The authors assembled a dataset spanning 840,537 individual respondents in 13,004 region-year cells, across 2,232 subnational regions within 116 countries. In the dataset, information on 3G coverage in 1×1 kilometer binary grid cells (provided by Collins Bartholomew’s Mobile Coverage Explorer) are matched to individuals’ opinions drawn from Gallup World Poll surveys, and electoral data on 102 parliamentary elections in 33 European democracies. First, the authors performed a DID strategy exploiting the variation in the timing of 3G expansion across different subnational regions within countries. Building upon the IV framework introduced by Manacorda and Tesei (2020), the study addressed endogeneity concerns by using the regional variation in lightning strikes, with each strike weighted by the local population density, as an instrument for 3G coverage. The study revealed that an increase in 3G availability, on average, diminishes trust in government, heightens awareness of corruption, and weakens support for incumbent governments.

In the vein of Guriev et al. (2020), Adema et al. (2022) combined Collins Bartholomew’s maps of mobile coverage with data from the Gallup World Polls to explore the influence of mobile internet access on the aspiration and plans to emigrate across 2,120 subnational regions in 112 countries during the period 2008 to 2018. After exploiting variation in 3G penetration across regions over time in a DID framework, the authors adopted an IV framework wherein 3G penetration was instrumented with the variability in 2G coverage prior to 2006. The empirical results indicate a positive effect of the increases in mobile internet access on the desire and planning emigration.

Expanding upon the initial work on 2G by Manacorda and Tesei (2020), Manacorda et al. (2023) explored the political implications of the diffusion of 3G networks. The authors compiled data on the distribution of party votes from 2007 to 2017 across 82,094 municipalities and matched this information with proprietary data on the availability of 3G signal throughout twenty European countries.

Table 3: Studies assessing the impact of mobile Internet

Authors and context	Instrument (level)	Treatment (level) ⁶	Outcome (level)
Adema et al. (2022) 112 countries 2008-2018	2G network coverage \times linear time trend	3G coverage (region)	Desire and intention to migrate (individual)
Bessone et al. (2022), Brazil 2011-2014		3G coverage (municipality)	Municipality's mentions in politicians' Facebook posts and offline speeches Federal transfers (municipality)
Donati (2023) South Africa 2000-2016	Terrain ruggedness \times Post 2005 (electoral sub-district)	3G coverage (electoral sub-district)	Voter turnout, vote share of the incumbent, protest participation (electoral sub-district)
Guriev et al. (2020) 116 countries 2007-2018	Population-weighted frequency of lightning strikes (region)	3G coverage (region)	Trust in government (region) Vote share of the incumbent (region)
Manacorda and Tesei (2020) Africa 1995-2010	Average n. lightning strikes \times t trend mobile uptake ($0.5^\circ \times 0.5^\circ \simeq 55km$ cell)	2G coverage ($0.5^\circ \times 0.5^\circ \simeq 55km$ cell)	Protest participation ($0.5^\circ \times 0.5^\circ \simeq 55km$ cell)
Manacorda et al. (2023) 20 European countries 2007-2017	Log distance of municipality from the nearest birthplace of a TLC manager in office in preceding years \times country's growth rate in 3G coverage (municipality)	3G coverage (municipality)	Vote for populist parties (municipality)
Melnikov (2023) US 2007-2018	Population-weighted frequency of lightning strikes (postcode)	3G coverage (postcode)	Political orientation (individual) Voting behavior (county)

Initially, the analysis took advantage of the variable timing of significant local increases in mobile Internet coverage across municipalities to conduct a DID exercise. To strengthen the

⁶Measures of mobile technology coverage are population-weighted.

causal interpretation of results, the authors adopted a TSLS framework, where 3G coverage was instrumented using the interaction between the logarithmic distance of each municipality from the nearest birthplace of a telecommunication manager in office in preceding years, and the country’s growth rate in mobile phone coverage. The underlying hypothesis is based on evidence that telecommunication leaders tend to excessively invest in areas they are personally familiar with or derive personal benefits from, especially during times of high sectoral growth, hence greater shareholder inattention. The identifying assumption hinges on the absence of correlation between temporal changes in latent voting patterns and temporal changes in mobile Internet coverage across municipalities at various distances from managers’ birthplaces. The results suggests that 3G technology led to a surge in voters’ support for parties promoting extreme policy platforms with nationalistic and communitarian tendencies, such as opposition to minority rights, immigration, multiculturalism, and European integration.

Melnikov (2023) employed a similar analytical framework in the United States context by merging 3G maps from Collins Bartholomew’s Mobile Coverage Explorer, Gallup Poll data on individual opinions, and data pertaining to the 2008–2016 presidential elections, along with the 2008–2018 elections to the US House of Representatives. Initially, the study conducted a DID analysis by capitalizing on the temporal variation in the expansion of 3G networks across postal codes. Building upon Manacorda and Tesei (2020), the analysis incorporated an instrumental variable (IV) framework, using the frequency of lightning strikes as an exogenous determinant of mobile infrastructure roll-out. The findings indicate a significant increase in political polarization corresponding with access to 3G. Upon obtaining 3G Internet, Democratic voters demonstrated a shift towards more liberal political views and amplified their support for Democratic congressional candidates and policy priorities. Conversely, Republican voters exhibited a trend towards more conservative stances. At the aggregated level, after the arrival of 3G, voters in Republican leaning counties increased their support for Republican congressional candidates while proportionally decreasing the vote share of Democratic candidates; in Democratic-leaning counties, the relationship is reversed. On an aggregated level, following the introduction of 3G, electors in Republican-oriented counties bolstered their support for Republican congressional candidates while proportionately diminishing their votes for Democratic candidates; the pattern is inverted in Democratic-leaning counties.

Donati (2023) examined the specific case of South Africa to ascertain the political consequences of 3G technology between 2000 and 2016. The research combined high-resolution data on 3G coverage with administrative records from sub-electoral districts on municipal elections. The empirical investigation adopted two complementary strategies. Initially, the author harnessed the spatial and temporal variation in 3G technology expansion within a DID framework. Subsequently, he utilized terrain ruggedness and its evolving impact over time on the rollout of the 3G infrastructure as an instrumental variable. The results indicate that the proliferation of mobile Internet stimulated

a significant increase in voter turnout and a decrease in the vote share of the incumbent party. Furthermore, the introduction of 3G seemed to intensify competition during elections and tripled the number of protests against political institutions by 2016. These findings suggest that mobile connectivity either intensified political discontent or facilitated the mobilization of the masses.

In a recent work addressing the case of Brazil, Bessone et al. (2022) leveraged the staggered penetration of the 3G network in the country to assess the impact of Facebook on politicians' behavior. In the absence of data on local-level Facebook adoption, the authors exploited within-municipality variation in 3G coverage to proxy for municipal-level exposure to the social medium. A unique aspect of Brazil's electoral system allowed the authors to further control for legislator-municipality fixed effects, since legislators run for elections across all municipalities of a given state. To examine the online and offline behaviors of politicians, the authors scraped the entirety of posts from all Facebook-active members of the Brazilian lower house during the 2011-2014 legislature from their public Facebook pages. They also collected data on their congressional speeches and earmarked transfers. This dataset was supplemented with information on the allocation of federal resources across municipalities and data regarding municipality-level 3G penetration provided by a telecommunications consultancy firm. The authors found that politicians' engagement on Facebook increased with 3G penetration. The online activity of the elected representatives was mainly directed towards strategically relevant municipalities. However, this attention appeared to displace legislators' focus on the same municipalities in their offline behavior, in terms of offline mentions and federal transfers. This result suggests that a substitution process occurred between online and offline activities. That is, more connected municipalities saw an increase in their presence in Facebook posts but experienced a decrease in legislative support for more tangible concerns.

Taken together, the body of literature examined in this section offers compelling evidence of the substantial political influence exerted by mobile Internet, with the specific effects being contingent upon the prevailing context. As a general observation, mobile connectivity tends to stimulate dissatisfaction towards incumbent political parties and enhance participation in protests. In fledgling or fragile democracies, as well as autocratic nations, this development is potentially positive in that it could empower citizens to demand greater accountability from their governments. However, in more established democracies, the proliferation of 3G connectivity appears to stimulate the rise of populism and political polarization.

2.5 Deployment of submarine cables to Africa

At the turn of the millennium, joint ventures involving private investors, African governments, and multilateral organizations were undertaken to construct submarine Internet cables linking Africa

with Europe. These cables reached coastal landing stations, often located on the outskirts of major cities, where a national backbone was typically meant to serve a single country.

While the distance travelled along submarine cables does not affect the quality of the digital signal, technological constraints typically emerge at the backbone level, where disparate networks intersect. To transmit content from a network to another, Internet Service Providers must sign collaborative agreements. In the absence of such agreements, African networks charge each other transit fees to exchange traffic. As a result, content hosted in a network must be routed through other continents to reach users (or countries) connected to a different network. (Hjort and Poulsen, 2019). This technological feature laid the foundation for a quasi-experimental setup based on the phased introduction of high-speed Internet across African nations. Since each country is typically serviced by a single backbone network and the advantages of a coastal nation’s submarine connection do not overflow to its neighboring countries, each country encounters a distinctive “treatment date”, generating potentially exogenous spatial and temporal variations in high-speed Internet access across countries (Hjort and Poulsen, 2019). Furthermore, once a land-based station in a country connects to the submarine cable, only locations linked to the national terrestrial cable network can access high-speed Internet. Unplugged locations primarily depend on the telephone network made of copper wires to connect to the cable network, making distance a crucial determinant of connection speed, as explained in Section 2.1.

As with the cases discussed in the preceding sections, the granularity of data significantly influences the empirical strategy. One of the primary challenges is the nationwide level of the treatment, primarily involving the connection of the submarine cable to a national backbone. In this scenario, obtaining detailed maps of the local cable network becomes indispensable to ascertain which locations are connected via copper wires and whether distance can influence broadband access. Nevertheless, all analyses leveraging the arrival of submarine cables must employ an intention-to-treat (ITT) design.

In this section, we highlight a selection of studies that leveraged the phased deployment of submarine cables to African coasts. We outline the essential characteristics of these studies in Table 4. As usual, the first column reports the authors and context of the study. Columns 2, 3, and 4 mention the exogenous source of variation in the treatment, the treatment, and the outcome variable, respectively. Their level of aggregation employed in the empirical analysis is in parentheses.

Table 4: Studies assessing the impact of submarine cables

Authors and context	Instrument (level)	Treatment (level)	Outcome (level)
D’Andrea and Limodio (2023) 37 coastal countries 2002-2018		Connection to a submarine cable (country) Availability of high speed Internet in the bank’s headquarters (country)	RTGS adoption, interbank transactions, lending (bank) Firms’ performance (firm)
Hjort and Poulsen (2019) 12 African countries 2006-2014	Connection to a submarine cable (country)	Distance from the country’s backbone network between 0 and 3,000m (individual)	Employment status, sector of employment, income (individual)
Houngbonobon et al. (2022)	Connection to a submarine cable (country) \times Broadband availability (municipality)	Internet access (individual and firm)	Entrepreneurship (individual), innovation (firm)
Strazzeri (2022)	Distance from the country’s backbone network \times dummy post-2010 (individual)	Internet access (individual)	Emigration (individual)

Hjort and Poulsen (2019) pioneered the evaluation of the economic outcomes of high-speed Internet access in Africa, with a particular focus on labor market outcomes. They merged data detailing the date and location of submarine cable landings, the geographical layout of backbone networks, and multiple survey sources encompassing information on business performance and employment status for twelve African nations. Their empirical framework compared individuals and firms close to the country’s backbone network—where high-speed Internet was accessible—to neighboring locations at a greater distance from the backbone network and likely excluded from broadband access. More specifically, the authors defined four treatment groups based on their proximity to the backbone, namely 0-500, 500–1,500, 1,500–2,500, and 2,500–3,500 meters. A control group comprised individuals residing between 3,500 and 10,000 meters from the backbone. These groups were then compared to gauge the differential effects of broadband access on those nearest to the backbone and those furthest away, both before and after the landing of submarine cables on the coast. The study’s underlying assumption was that these locations had parallel trends in labor market outcomes before the arrival of submarine Internet cables in Africa and were not subject to systematically divergent idiosyncratic shocks post-arrival. The findings from their empirical analysis indicates an increase in the likelihood of both employment in skilled positions and overall employment once high-speed Internet was made available. However, there was no

substantial impact on unskilled labor. Generally, the arrival of high-speed Internet seemed to have caused a shift in employment towards occupations with higher productivity. These results align with earlier studies conducted in Norway (Akerman et al., 2015; Bhuller et al., 2023a).

Strazzeri (2022) employed the same methodology to examine how the introduction of submarine cables influenced migration decisions from 2010 to 2016, focusing specifically on Nigeria, where the first cable was connected to a land-based station in Lagos in July 2010. The author linked migration and Internet access data from the geocoded Nigerian General Household Survey (GHS) with detailed maps of Nigeria’s terrestrial cable network prior to the submarine cables’ arrival. The configuration of the network before 2010 influenced connection speed based on the distance of users to the local exchange, as discussed in Section 2.1. The empirical analysis first leveraged the sequential penetration of high-speed Internet in a DID design. The second part of the study employed a TSLS model in which a survey-based measure of Internet access was instrumented using the interaction of users’ distance from the local exchange and a binary variable capturing the arrival of the cable. The findings suggest a significant positive influence of Internet usage on migration, with a more substantial effect on migration out of Africa and among individuals in the lower wealth distribution.

D’Andrea and Limodio (2023) focus on coastal countries to explore the influence of high-speed Internet’s arrival to Africa on banking and the adoption of financial technology between 2002 and 2018. The authors combined geographical and temporal data relating to the landing of submarine cables in Africa with balance sheet information from 629 African banks in 37 coastal countries, along with data concerning the adoption of the Real-Time Gross Settlement system (RTGS). The dataset was further enriched with firm performance measures derived from the World Bank Enterprise Surveys. The study used a DID approach, leveraging the phased arrival of submarine cables. In this setup, a bank is considered ‘treated’ when the country it operates in gains a submarine cable connection. Moreover, the authors leveraged an aspect of high-speed Internet connectivity for banks that was detached from local demand factors. A bank operating in a certain country may attain a high-speed Internet connection due to its headquarters, located in a different country, being connected. This measure of connectivity has the benefit of independence from the cable’s arrival in the country, which could influence the outcome variables by sparking changes in demand. The findings indicate that, after their countries connect to high-speed internet, banks adopt the RTGS more extensively, decrease inside liquidity, and amplify interbank transactions and lending. Interestingly, it appears that high-speed internet particularly bolsters firms in countries with weaker preexisting interbank markets. In a similar vein, Hounghonobon et al. (2022) studied the impact of Internet penetration on entrepreneurship and innovation among firms in ten African countries. The authors matched information about the deployment of fiber-optic Internet cables obtained from Africa Bandwidth Maps with survey data taken from the World Bank Enterprise

Surveys and the Living Standards Measurement Study, which contain information about firms' and households' Internet access. The empirical analyses instrumented Internet access with the interaction between two dummy variables capturing the arrival of a submarine cable to the country and broadband availability at the municipality level. The study found evidence that the arrival of high-speed Internet in a city increased the probability of becoming an entrepreneur and spurred innovation among firms.

2.6 Technological accidents: the case of East Germany

Numerous studies exploring the case of Germany (summarized in Section 2.2) employed a threefold identification strategy. Besides the DID and TSLS approaches covered in the preceding section, this body of work also capitalized on a “technological mistake” (Bauernschuster et al., 2014) that took place in East Germany following reunification. In the early 1990s, the German government, along with Deutsche Bundespost (the former state-owned telecommunication monopolist, later privatized into Deutsche Telekom), introduced a new telephone infrastructure - the so-called OPAL technology - utilizing fiber wires instead of traditional copper wires in East Germany. This technology represented the cutting-edge in voice communications at that time. The OPAL system was eventually deployed in 213 East German catchment areas, reaching about 11 percent of the East German population.

However, a few years later, with the advent of DSL technology which enabled the transmission of digital signals over copper wires, it became clear that DSL and OPAL were not compatible. In order to provide broadband internet access in OPAL areas, two significantly costly alternatives were available - replacing the OPAL wires in the access network with copper wires, or installing new hardware and software at the network nodes (the latter option being even more expensive). Consequently, areas covered by OPAL were effectively prevented from connecting to the broadband network. This particular situation enabled researchers to assess the impact of high-speed internet through a TSLS framework, wherein measures of internet access are instrumented via a binary indicator denoting whether the area was affected by the OPAL incident.

Early studies, such as those by Czernich (2012) and Falck et al. (2014), employed this technological hiccup as a source of exogenous variation in high-speed internet access. This approach was later adopted by others, including Bauernschuster et al. (2014), and Billari et al. (2018; 2019). An exemplary application of this method can be found in Bauernschuster et al. (2014), who combined geolocalized survey data drawn from the German Socio-Economic Panel (SOEP) with information on the historical roll-out of the OPAL technology. The deployment of the new cables led to the random assignment of a broadband-unfriendly technology to end users across over 1000 small-

scale catchment areas. This circumstance allowed the authors to use information on whether a household was located in an OPAL area before the advent of DSL technology as an instrument for determining the impact of broadband access on social capital. The results indicate significant positive effects of broadband Internet access on a latent indicator that combines measures of different social capital dimensions. Moreover, using East German municipality-level data, the authors' found no significant impact of broadband availability on voter turnout.

2.7 Weather conditions

A subset of research addresses the issue of endogeneity in broadband supply and demand by capitalizing on the well-documented impact of inclement weather on the reliability of Internet connections. This weather interference often leads to Internet outages, increases the costs of providing stable broadband, and can discourage end-users from purchasing and using these services. The identifying assumption underpinning this body of research is that adverse weather conditions can influence the outcomes under scrutiny solely through their impact on broadband's supply and demand. This presupposes that localities with varying weather conditions would have analogous political outcomes if not for the effect of broadband. However, a significant challenge exists with this identification strategy due to the wide range of outcomes weather conditions have been shown to impact in previous studies. In particular, numerous investigations have explored the effects of conditions like rain, temperature, wind speed, and sunlight on various outcomes across specific territories. These studies have revealed direct influences on phenomena such as political behavior, labor market outcomes, income, and mental health (Mellon, 2022).

Gavazza et al. (2019) examined the impact of internet penetration on electoral outcomes and local governments policy decisions in England between 2006 and 2010. The authors combined data from all local council elections held in England during this period with tax and expenditure data from 113 out of 125 English unitary Local Authorities (LAs). Additionally, they incorporated information about the quantity of broadband internet service providers, the number of subscribers each provider had, the availability of cable technology, the number of cable subscribers, and the catchment areas of each node, provided by the UK Office of Communications (Ofcom). To tackle endogeneity in broadband penetration, the authors utilized the influence of inclement weather. The average daily rainfall across electoral wards was thus employed as an instrument for broadband penetration. Utilizing this instrumental variables approach, the study revealed that increased internet penetration led to a decrease in voter turnout, and that areas with greater broadband diffusion had lower local government expenditures and taxes.

3 Social media: natural and quasi-experiments

Research seeking to determine the causal effects of social media usage encounters similar identification obstacles as those found in the literature on the influence of high-speed Internet penetration. From a demand perspective, social media use could correlate with a wide range of individual traits and local factors potentially impacting online interaction outcomes. On the supply side, access to social media might depend on elements such as mobile network penetration, which could correlate with local features affecting the profitability of infrastructure investments. Much like the studies examined in Section 2, the primary methodology employed in the economics literature involves isolating an exogenous source of variation in the demand for social media or in their availability.

The research in this area operates under the identifying assumption that specific drivers may make the demand or supply of a specific medium vary randomly, while influencing the potential outcome exclusively through their relationship with the treatment. Any additional causal pathways from the instrument to the dependent variable would constitute a violation of the exclusion-restriction principle. The majority of these studies primarily concentrate on the demand side, leveraging exogenous phenomena that might have randomly influenced public interest in specific social media platforms. For instance, certain studies utilize spatial discontinuities in the usage of particular media, such as Facebook, Twitter, or VKontakte, prompted by geographically precise events that might have stimulated usage. As with all the studies examined so far, the primary challenge lies in the treatment assignment mechanism, which needs to approximate randomness and be orthogonal to the outcome of interest. When these conditions are met, a natural experiment may arise, providing researchers an opportunity to identify the causal impact of social media usage.

Hence, the most methodologically significant distinction from the studies discussed in Section 2 pertains to the utilization of exogenous sources of variation in the demand for specific media as instruments to ascertain the causal effect of online interaction. Conversely, the studies on high-speed Internet access are exclusively reliant on supply-side instruments capturing geographical variations in the availability of such technologies. This difference makes the identifying assumptions more challenging, as the demand for a specific medium could be correlated with a multitude of factors influencing individual choices, from the personal attributes of users to the predominant characteristics of the area where they live.

We will commence by reviewing studies that capitalized on an area's association with the central nodes of a network. The underlying premise of these studies is that the propensity to use a specific medium is influenced by the spatial closeness to regions where the network initially proliferated more rapidly. For instance, during a medium's early diffusion stages, its utilization may reflect the geographic distribution of the birthplaces of the founder's classmates (Enikolopov et al., 2020a). Next, our focus shifts to studies utilizing online movements and events as exogenous sources of

surges in the demand for a medium. Some research, for example, leveraged the emergence of online movements such as #MeToo to discern the impacts of Twitter usage. Then, we will summarize studies exploiting the precise timing of the publication of specific content on a medium to identify the offline outcomes of the medium. Lastly, we will explore studies exploiting supply-side sources of exogenous variation in exposure to specific media. For example, the availability of a specific social medium in the local language may act as a source of variation in its usage at the local level. For each approach, we will discuss the identification challenges, empirical strategy, and the key findings of the most prominent studies.

3.1 Social network centrality

A recent set of studies has tackled the impact of social media by harnessing the topology of a social network as an exogenous driver of its geographic penetration. While there are numerous definitions of network centrality and the network literature is notably complex, studies on social media appear to implicitly utilize the concepts of degree and closeness centrality. The underlying assumption is that the nodes established first are destined to become central in one or both of these ways. Therefore, a social medium is more likely to gain traction earlier - and exert a greater impact - near the network nodes established first. Table 5 summarizes the studies that leveraged aspects of network centrality to identify the impact of social media. As usual, the first column reports the authors and context of the study. The following columns inform about the instrumental, treatment, and outcome variables, respectively.

This approach was employed to examine the political impact of social media in a significant case study - Russia, where the ongoing autocratic shift poses a significant threat to international peace and political stability. Leveraging the geographical distribution of early nodes of VKontakte (VK), the foremost Russian social network, Enikolopov et al. (2020a) determined its influence on protest participation during the 2011 Russian demonstrations, incited by supposed electoral frauds in parliamentary elections. VK, introduced by Pavel Durov in October 2006, began with Durov personally inviting Saint Petersburg State University (SPbSU) students via an online forum to join the platform. Consequently, VK's early adopters were predominantly his university peers. According to Enikolopov et al. (2020a), this mode of dissemination resulted in friends and family of these initial users, residing in their hometowns, being more likely to create a VK account, thus accelerating VK's expansion in these municipalities.

Table 5: Studies leveraging network centrality

Authors, social medium, and context	Instrument (level)	Treatment (level)	Outcome (level)
Bursztyn et al. (2020), VKontakte / Russia 2007-2015	Number of students in the same cohort with VK’s founder (municipality)	VK’s penetration (municipality)	Ethnic hate crimes (municipality)
Casanueva-Artis et al. (2022) Twitter / U.S. 2020 ⁷	Number of SXSU festival followers in 2007 (county)	Twitter new subscriptions in 2020 (county)	BLM protests (county)
Cagé et al. (2022) Twitter / France 2018-19	Tweet’s author network centrality*Twitter news pressure (event)	Twitter coverage of news events (event)	Newspapers coverage of news events (event)
Enikolopov et al. (2020a), VKontakte / Russia 2011	Number of students in the same cohort with VK’s founder (municipality)	VK’s penetration (municipality)	Incidence of protests, protest participation (municipality)
Enikolopov et al. (2020b), VKontakte / Russia 2011	Number of students in the same cohort with VK’s founder (municipality)	Size of VK’s municipality-based online protest groups	Protest participation (municipality)
Fujiwara et al. (2022), Twitter / U.S. 2016, 2020	Number of SXSU festival followers in 2007 (county)	Twitter usage (county)	Trump’s vote share in 2016 and 2020 presidential elections (county)
Müller and Schwartz (2023) Twitter / U.S. 2016	Number of SXSU festival followers in 2007 (county)	Twitter usage (county)	Hate crimes (county)

The authors initially demonstrated that the distribution of the home cities of students who studied simultaneously with Durov at SPbSU, predicts VK’s penetration in 2011⁸. Based on this evidence, the authors could employ the number of students who attended SPbSU within the same five-year cohort as VK’s founder and were born in a specific city, as an instrumental variable for

⁷The authors also used the SXSU instrument in a robustness check. Since their main specification employed a different source of variation in Twitter usage, we report about this study in more detail in Section 3.2.

⁸In a subsequent online corrigendum to the paper, the authors reported a computational error that affected the measurement of two control variables, namely the spherical distance from Moscow and Saint Petersburg. When these rectified control variables were included in the analysis, the instrument’s coefficient in the first stage regression reduced significantly in magnitude and lost statistical significance, in turn negatively influencing all ensuing instrumental variables estimates (Enikolopov et al., 2023). However, the authors maintained that the issue was not due to the instrument’s validity being compromised and showed that replacing the rectified control variables with measures of the driving distance from Moscow and Saint Petersburg, the first stage specification becomes stronger and the final results are very similar to those previously estimated.

VK penetration in that city in summer 2011, with penetration measured as the logarithm of the number of VK users in a city in the summer of 2011. Therefore, the study’s identifying assumption was that the variations in the number of students originating from diverse Russian cities to attend SPbSU bore no relation to unobserved city traits correlating with political outcomes. Leveraging this peculiarity of VK, the authors showed that a higher VK penetration significantly heightened the likelihood of a protest and the amount of protesters. Further findings imply that social media encouraged protest activity by diminishing coordination costs instead of disseminating government-critical information. On the other hand, an increase in VK penetration was associated with growing governmental support, hinting that while social media can potentially ease coordination efforts, its influence on information can swing in either direction. This outcome primarily hinges on whether the average social media content reflects a favorable view of the government, highlighting the potential for autocratic regimes to utilize social media platforms as conduits for propaganda, thereby moulding public sentiment. King et al. (2017), for instance, estimated that the Chinese government fabricates and circulates approximately half a billion social media comments annually. In contradiction to previous assumptions, the authors demonstrated that the Chinese regime’s tactic is to sidestep confrontations with party and government skeptics and to steer clear of contentious issues entirely.

In a related study, Enikolopov et al. (2020b) explored whether social image concerns can bolster protest participation levels during the same wave of demonstrations in Russia. This study leveraged another unique aspect of VK. During the protest period, VK facilitated the creation of online protest groups with a geographical focus aimed at planning and organizing offline demonstrations. The authors suggested that larger online protest groups provided participants an opportunity to signal their involvement to a wider audience. Hence, the size of these groups could serve as a proxy measure for social image motivations. Using the approach developed in Enikolopov et al. (2020a), the study employed the number of students from each city who were enrolled at SPbSU in the same five years as VK’s founder as an instrumental variable for the size of online groups. The empirical analysis showed that larger online protest groups positively influenced protest participation within a city, implying that social image considerations also contributed to fueling protest turnout.

Expanding on Enikolopov et al. (2020a), Bursztyn et al. (2020) harnessed the variations in the geographic distribution of SPbSU students to instrument the city-level penetration of VK, assessing its impact on hate crimes. The authors combined data on SPbSU students with hate crime statistics, which were independently compiled by the Russian NGO, SOVA, from 2007 to 2015. The empirical analysis established that higher social media penetration prompted an increase in ethnic hate crimes, but notably only in cities with a pre-existing high level of nationalist sentiment before the arrival of social media. This finding suggests that social media primarily serves as a coordination platform for individuals already harboring xenophobic tendencies. To further inves-

tigate whether social media merely coordinates or can also transform opinions, Bursztyn et al. (2020) executed an online survey experiment, which revealed a positive relationship between VK penetration and elicited ethnic hostility.

Parallel to studies addressing the role of VK in Russia, several authors have explored the influence of Twitter in the US by capitalizing on the geographical dispersion of early adopters as a source of exogenous variation in the platform’s penetration. Mirroring the studies on Russia, the underlying assumption is that early adopters may facilitate the spreading of the platform’s adoption within their their places of origin. This circumstance subsequently awards these places a position of centrality within the social network’s structure, thereby favoring the platform penetration and the unfolding of its effects. An event that has been recognized as crucial in the early penetration of Twitter is the Southwest (SXSW) Festival in March 2007. According to the company’s folklore, the platform received its inaugural boost of “social amplification” during this festival when a spectrum of bloggers, coders, and tech enthusiasts embraced the service, lured by tweets showcased on flat-panel screens in corridors—a novel feature at the time. Müller and Schwartz (2023) documented a rapid growth in Twitter activity in the weeks following SXSW. The surge was particularly noticeable in the home counties of SXSW attendees who had registered for Twitter in March 2007. The authors provided empirical evidence that this initial expansion made a lasting mark on the geographical spread of social media usage across the United States. The locations of Twitter’s early adopters at SXSW are a strong predictor of county-level Twitter usage today, even after controlling for the locations of SXSW followers that had already signed up prior to the festival.

This methodology hinges on the identifying assumption that areas evolving into key nodes within the network don’t possess specific characteristics that could simultaneously foster social media adoption and the outcome behaviors under scrutiny. Müller and Schwartz (2023) contended that, given Twitter was relatively unknown before SXSW and these counties’ lack of systematic differences in many observable characteristics, this assumption holds. Conversely, the relevance assumption necessitates the persistence of this social amplification effect, which would entail the enduring centrality of the early network nodes even a decade later.

To determine the causal impact of social media on hate crimes, Müller and Schwartz (2023) employed the number of SXSW followers per U.S. county in March 2007 as an instrument for Twitter usage. Much like the approach taken by Enikolopov et al. (2020a), the identifying assumption posits that the choice to follow the SXSW festival Twitter page in March 2007, instead of in the preceding months caused an increase in anti-Muslim sentiments exclusively by encouraging the use of the social medium. By combining Twitter usage data with FBI hate crime statistics, the authors discovered a disproportionate surge in anti-Muslim hate crime indicators in counties with more intense Twitter usage during the 2016 presidential campaign period. Their empirical analysis documents that a larger portion of this effect originated from areas with pre-existing higher

anti-minority bias, which aligns with the findings of Bursztyn et al. (2020). In order to isolate the direct impact of Trump’s tweets against minorities, the authors exploited an alternate instrument hinged on the observation that the former U.S. president predominantly tweeted about minorities during his golf outings. By using his golf sessions to instrument Trump’s tweets about Muslims within a time-series regression framework, the authors presented evidence supporting the notion that Trump’s tweets concerning Muslims triggered ripples of anti-Muslim sentiment.

Building on the previous study, Fujiwara et al. (2022) leveraged the number of SXSU followers per U.S. county in March 2007 as an instrument to explore the influence of Twitter usage in 2016 on the outcomes of the presidential elections in 2016 and 2020. Their estimates suggest that a 10% rise in the number of Twitter users within a county corresponded to a reduction of 0.2 percentage points in the vote share for Republican presidential candidate Donald Trump in both elections. In their attempt to delve into the potential transmission mechanism of the treatment effects, the authors documented that political content on Twitter has a pro-Democratic slant.

In instances where the spatial distribution of network nodes does not provide sufficient information to identify potential treatment effects, the timing of a specific content’s spread through central nodes may become crucial. Cagé et al. (2022) employed the concept of network centrality to discern the influence of social media on traditional media outlets’ news production decisions. Accurately establishing the causal impact of social media on the selection of news for newspaper coverage is challenging, as coverage could primarily arise from the inherent newsworthiness of the content, rather than reflecting the reciprocal influence between traditional and social media. Hence, the primary hurdle involves retrieving an exogenous source of variation in a story’s popularity on social media. In their preliminary work⁹, the authors gathered 31 million tweets and more than one million news articles. From the subset of events that received coverage on both Twitter and traditional media, they singled out 3,904 events that first appeared on the social platform to examine how their popularity on Twitter influenced their coverage in traditional media. To identify the impact of the social medium, the authors put forth a novel instrument that hinges on the interaction between a user’s centrality in the Twitter network and the “news pressure” on Twitter at the time of the event. Each term of the interaction, if viewed independently, could plausibly correlate with traditional media’s news production decisions. However, the authors argued that their interaction should only affect traditional news production through its influence on a tweet’s visibility on Twitter, measured through its impression count. In the study, network centrality is measured using the PageRank algorithm (Page et al., 1999), underpinning the idea that a node’s

⁹At the time of writing this review, the authors had conducted empirical analysis on a subset of their originally collected data. This approach was adopted to refine the set of hypotheses that could be empirically tested and to establish a research plan suitable for pre-registration. The final rendition of the paper will adhere to the pre-registered plan and carry out the empirical analysis on the remaining data.

significance is associated with the importance of the nodes linked to it. Twitter’s news pressure is quantified as the volume of interactions generated by all tweets posted in the hour preceding the first tweet relating to the event. Leveraging this methodology, Cagé et al. (2022) illustrated that a fifty-percent surge in the number of tweets posted before the first media article appears prompts a significant 17 percent uptick in the number of media articles covering the event.

3.2 Surges in the demand for social media or specific content

A potential source of variation in the demand for social media could be tied to trends or events, both offline and online, that captivate the public interest. These phenomena might drive potential users to sign up or invest more time in the platform, or, at the very least, pay closer attention to a specific ongoing trend within the platform. However, since these shifts in platform usage are completely demand-driven, it becomes challenging to isolate the causal influence of the platform from the impact of users’ personal traits and the myriad of confounding factors that might simultaneously affect the unfolding of pivotal events and the demand for social media.

This line of research generally strives to ascertain the impact of social media on a localized level, ranging from countries to counties. This task involves combining either observed or latent data on social media usage with administrative data measuring the hypothetical outcomes of social media. This approach inevitably grapples with the challenge of gauging the spatial distribution of social media usage. Various strategies have been deployed, such as analyzing the geolocation of posts or measuring online search trends related to specific phenomena prevalent on social media. Some studies attempted to circumvent endogeneity issues by exploiting certain potential drivers of the local demand for social media, like the subscription habits of local sports stars.

However, achieving accurate identification remains a formidable task. For studies that utilize the geolocation of posts, they must contend with the fact that only a small, unrepresentative subset of posts are geolocated. Studies using online search trends face the issue of low-definition data. Moreover, uncovering exogenous sources of variation in a medium’s local incidence becomes a challenging task more than ten years after their initial penetration, particularly when their dispersion has reached saturation in most regions. In every instance, when variations in the use of a medium are driven by demand, the study must depend on the identifying assumption that areas experiencing an upswing in social media usage do not have distinct characteristics that could simultaneously influence the outcome behaviors under investigation and encourage social media adoption independently on the enticing online event that the authors exploit for identification.

Table 6 reports the studies that have used surges in the demand for specific social media to determine their impact. As in previous tables, the first column lists the authors and the context

of each study. The subsequent columns provide details regarding the instrumental, treatment, and outcome variables, in that order.

Table 6: Studies leveraging demand surges

Authors, social medium, and context	Instrument (level)	Treatment (level)	Outcome (level)
Battisti et al. (2022), Twitter / U.S. 2014-16		GBV tweets per 100 cellphone Internet subscriptions (county)	GBV reported crime rates (county)
Cao et al. (2022), Twitter / U.S. 2020		Anti-Asian COVID-related tweets*political leaning (county)	Anti-Asian reported hate crimes (county)
Casanueva-Artis et al. (2022), Twitter / U.S. 2020	COVID exposure within a 50 km radius from county borders 6 weeks before George Floyd’s murder (county)	Twitter new subscriptions (county)	BLM protests (county)
Levy and Mattsson (2023), social media / 30 OECD countries 2017-18		Google searches for MeToo related keywords (country)	GBV reported crime rates (country)
Rotesi (2019), Twitter / U.S. 2008-2016	Degree of exposure to the NBA Draft (DMA region)	Number of accounts every 1000 inhabitants (DMA region)	Voter turnout, electoral performance, party donations (DMA region)

Levy and Mattsson (2023) conducted the first cross-country study to analyze the effects of social movements, focusing on how the MeToo movement impacted the reporting of sexual crimes. The authors combined country-level monthly Google search data for keywords signifying interest in the MeToo movement with statistics on crime reports received by authorities across 30 OECD countries. They classified countries into two categories: those with a strong MeToo movement, defined by interest levels above the OECD median in October 2017 (the inception of the movement), and those with a weak MeToo movement, characterized by below-median interest.

To identify the impact of MeToo, the authors applied a difference-in-differences (DID) strategy comparing sexual crimes reporting in countries with both high and low interest in the MeToo movement, before and after its emergence. The findings revealed that the movement sparked a 10% increase in reported sexual crimes during its first six months. For the assessment of long-term effects, the authors narrowed their focus to countries that initially demonstrated a strong interest in the movement. By employing a DID approach comparing sexual crimes to all other crimes over

time, their analysis found that the MeToo movement exerted a lasting influence, with a significant effect on reporting that extended to at least 15 months after its initiation.

Battisti et al. (2022) focused the influence of the MeToo movement within the United States. Instead of gauging online interest in the social movement as in Levy and Mattsson (2023), the authors quantified actual social media engagement in discussions surrounding gender-based violence (GBV) by developing a weekly measure of the prevalence of tweets on GBV per 100 cellphone internet subscriptions across federal states from 2014 to 2016. In order to assess the effect of the MeToo movement, they collated their Twitter data with data on crime incidences collected by the FBI. Their empirical analysis suggests that an increase in the prevalence of GBV-related tweets led to a decrease in GBV-related reported crime rates per 100,000 inhabitants. The authors provided evidence indicating that their results were most likely driven by behavioral changes among those who perpetrate GBV.

Cao et al. (2022) focused on the amplification effects of social media, specifically Twitter, in propagating the narrative by former U.S. President Donald Trump that COVID was a virus engineered by the Chinese government. On March 16, Trump publicly used the term “Chinese Virus” for the first time in a tweet about his intention to support industries impacted by the pandemic. This narrative shaped much of the President’s subsequent Twitter activity. To assess the influence of this political message propagated through social media, the authors examined the rates of hate crimes committed against Asian people in counties won by Donald Trump in the 2016 presidential election, compared to those won by Hilary Clinton utilizing a DID approach. The authors found that following Trump’s initial tweet, anti-Asian COVID-related tweets increased 235 percent more in Trump-favored counties than in those favoring Clinton. Concurrently, anti-Asian hate crimes escalated by over 4000 percent more in Trump-supporting counties than in counties that backed Clinton.

Casanueva-Artis et al. (2022) analyzed the role of Twitter in promoting participation in the Black Lives Matter (BLM) protests in May 2020. The BLM movement originated on Twitter in 2013, and in May 2020, #BlackLivesMatter had become the most used hashtag on Twitter, reaching a peak of 8.8 million daily mentions. Following evidence provided by social media platforms, the authors identified exposure to the novel coronavirus pandemic as a factor influencing social media usage. Notably, in the weeks preceding the protests, Twitter and Meta social networking sites reported a significant surge in new users and online activity, attributing this to the COVID outbreak. The authors then used this relationship to identify the impact of Twitter on BLM protests in a two-stage least squares (TSLS) framework. To measure exposure to COVID-19 at the county level, they harnessed cross-sectional variability in the number of Super Spreader Events (SSEs) within a 50 kilometer radius from county borders (but not within the counties) six weeks before George Floyd’s murder. By combining information on SSEs in the initial stages of the

outbreak with county-level data on new Twitter subscriptions and BLM protest participation, Casanueva-Artis et al. (2022) documented that rises in pandemic exposure led to notable increases in Twitter subscriptions, which in turn appeared to foster the emergence of BLM protests. In a robustness check, the authors also utilized the instrument introduced by Muller and Schwartz to demonstrate that BLM protests were more likely to take place in counties that had gained prominence in the social network earlier (for this reason, we include this study in both Table 5 and Table 6).

In an earlier study, Rotesi (2019) introduced an innovative method to identify the county-level influence of Twitter usage. To explore the platform’s effect on political participation, the author constructed an instrument that capitalized on fluctuations in the popularity of sports teams that hired new players with active Twitter accounts, thus enhancing the appeal of the social network to their fanbase. This approach took advantage of certain specificities of hiring regulations that render the destinations of these new signings quasi-random, once team and player-related factors are controlled for. Using this approach, Rotesi (2019) discovered that the influence of Twitter tends to be negative for the Democratic Party and positive for the Republican Party. Specifically, an observable decrease in the number of votes for the Democratic Party accompanied a marked increase in the volume of donations received by the Republican Party.

3.3 Timing of events

In cases where the structure of the social network fails to provide adequate information to discern potential treatment effects, and when it’s not feasible to geographically allocate a medium’s reach, alternative approaches may be required. One such method involves capitalizing on the exact moment of publication for specific content, and information regarding whether this content was previously or simultaneously disseminated by other media. This approach can facilitate a clearer understanding of a particular medium’s impact.

Enikolopov et al. (2018) examined the effects of blog posts authored by Alexey Navalny, a Russian opposition leader and anti-corruption activist, focusing on corruption within state-controlled corporations. Their research sought to determine the subsequent impact on stock market performance and corporate practices. To control for endogeneity, the authors capitalized on the specific timing of blog posts during a trading day, mapping their effects on the performance of companies with established political ties to the Russian government. They used data collected at five-minute intervals during trading hours, incorporating controls for company-specific trading days and different hours of the day.

To offset potential endogeneity arising from prior media coverage of these companies, the au-

thors narrowed their focus to blog posts that were not preceded by any public mentions of the companies in question from other sources. The empirical results showed a decline in stock returns of the mentioned companies within three hours of the blog posts being published. Additionally, the findings suggest an increased likelihood of management turnover and a reduced propensity for shareholder conflicts within the companies following the publication of Navalny’s blog posts.

3.4 Supply-side sources of exposure variation

This body of research employed varied supply-side sources of exogenous variation in the accessibility of social media to determine their influence on a wide array of outcomes, including mental health, firm performance, protest movements, and racial hate crimes. Given the inherent difficulty in discerning the impact of expected investment returns on the supply of an Internet product, these studies leveraged plausibly exogenous events, from Internet or social media outages, to policy reforms such as the abolition of roaming charges within the EU. In many of these cases, the studies applied a difference-in-differences methodology to examine the impact of a treatment activated by a plausibly exogenous event.

Table 7 provides an overview of this subset of literature. The first column lists the authors and the context in which each study was conducted. The second column details the exogenous source that influences the treatment exposure, while the third column describes the treatment itself. Finally, the fourth column highlights the outcomes under investigation. As in previous tables, the level of aggregation for the corresponding variables is denoted in parentheses.

Fergusson and Molina (2021) examined the influence of Facebook on the frequency of protest events on a global scale, incorporating a sample of 240 countries from 2000 to 2015. Their identification strategy hinged on the introduction of Facebook in specific languages as an exogenous variation in social media accessibility across countries, regions, and populations speaking those languages. The authors combined data from the Global Database of Events, Language, and Tone (GDELT) - a source which catalogues various types of collective actions daily - with information regarding the dates when Facebook became available in all 81 unique languages up until March 2016, including beta versions. The main explanatory variable was a “Facebook speakers” measure at the country level, representing the proportion of the country’s population that could access the platform in their primary language. The variable was constructed by interacting a dummy indicating whether a Facebook version in a certain language was available, with the percentage of the country’s population that speaks that language. Employing this treatment, the authors performed two-way fixed-effects regressions based on the identifying assumption that, in the absence of Facebook’s introduction in those languages, countries with different proportions of speakers of

Table 7: Studies leveraging supply-side sources of exposure variation

Authors, social medium, and context	Exogenous source of variation (level)	Treatment (level)	Outcome (level)
Braghieri et al. (2022), Facebook / U.S. 2004-2006	Staggered penetration of Facebook in American colleges	Facebook availability (college)	Mental health, depression (individual)
Donati (2022), Tripadvisor / Rome 2015-2019	Policy reform making mobile Internet abroad cheaper (EU)	Post policy reform*tourist exposure (restaurant / ZIP code)	Revenues, size, closure, quality upgrading (restaurant / ZIP code)
Fergusson and Molina (2021), Facebook / 240 countries 2000-2015	Facebook availability in the local language (country / region)	Share of country's (or region's) population that can access Facebook in their 1st language.	Protest events (country / region)
Müller and Schwartz (2021), Facebook / Germany 2015-2017	Facebook outages (country), Internet outages (municipality)	AFD users*saliency of anti-refugee hate speech (municipality*AFD Fb page)	Anti-refugee hate crimes (municipality)

the respective languages would have displayed similar trends in collective action. In an additional specification, the authors exploited within-country variation to control for a comprehensive set of country-time fixed effects. The empirical analysis indicated that the introduction of Facebook significantly amplified the number of protests, both across countries and regions. This finding further supports the external validity of prior research on the subject.

To assess the causal impact of Facebook, Braghieri et al. (2022) took advantage of its gradual rollout across U.S. colleges during the mid-2000s as a source of quasi-experimental variation in exposure to the medium. Facebook was founded at Harvard in February 2004 and only opened to the general public in September 2006. Between these dates, the platform expanded incrementally across U.S. colleges, resulting in rapid and substantial penetration among students once each college was granted access to the Facebook network. This staggered and swift introduction of the platform provides the opportunity for causal identification, especially when combined with student or college-level data from the same period. The authors assembled a dataset by combining the dates of Facebook's introduction at 775 US colleges with responses from seventeen consecutive waves of the National College Health Assessment (NCHA), a survey administered semi-annually to college students by the American College Health Association (ACHA). Given that the NCHA survey does not contain any questions related to social media use, the authors implemented an intention-to-treat design and used a generalized DID strategy. This strategy allowed them to examine variations based on the colleges attended by the students, as well as whether the survey

was completed before or after the introduction of Facebook at those colleges. Their empirical analysis revealed that the introduction of Facebook at a college led to a significant decline in student mental health. Additionally, the authors found that, subsequent to the introduction of Facebook, students were more likely to report that poor mental health had adversely affected their academic performance. Lastly, further evidence suggested that these results are consistent with the notion that Facebook usage exacerbates students' tendencies to engage in unfavorable social comparisons.

In their study of the impact of Facebook on hate crimes against refugees in Germany, where the platform was introduced nationwide in 2006, Müller and Schwartz (2021) adopted a different approach, leveraging the exact timing of nationwide Facebook outages and local Internet disruptions, which in turn reduced exposure to social media posts. The empirical analysis involved a DID strategy, with the key treatment variable being a measure of the prominence of anti-refugee hate speech on social media. This was quantified based on the activity on the Facebook page of "Alternative für Deutschland" (Alternative for Germany, or AfD), and then interacted with the number of the page's followers across different municipalities. The authors discovered that nationwide disruptions of Facebook service led to a reduction in local hate crimes, particularly in areas with a large number of AfD followers. Conversely, while hate crimes tended to increase during periods of heightened visibility for refugee issues, this correlation evaporated for municipalities experiencing an Internet outage. The authors also documented that other salient news events, such as the UEFA European Football Championship, could crowd out the salience of anti-refugees hate speech, thereby muting the relationship between social media exposure and hate crimes, though temporarily.

In a focused and potentially widely applicable case study, Donati (2022) investigated the impact of Tripadvisor on the performance of restaurants situated in a heavily touristic area, namely the province of Rome. The author leveraged a natural experiment prompted by a policy change which suddenly eliminated mobile roaming charges in the EU. This policy change caused an arguably exogenous variation in the costs for travelers to access online reviews. Donati (2022) combined data from Tripadvisor with administrative employer-employee records sourced from the Italian National Social Security Institute (INPS) and the Italian Business Registry. For approximately 5,500 matched restaurants in the province of Rome, the dataset includes time-invariant information such as name, address, price category, type of cuisine and other covariates. It also encompassed dynamic information like the number of Tripadvisor reviews (sorted by origin of reviewer and device), rating, opening and closure dates, employee count, contract types, wages, and comprehensive employment histories of workers at monthly intervals between 2007 and 2019. To identify the impact of Tripadvisor the study exploited the temporal variation introduced by the policy change and the spatial variation in tourist demand at the restaurant and ZIP code levels. The empirical

analysis revealed that high-rated restaurants experienced a significant annual revenue increase after the policy change, while revenues for low-rated restaurants remained the same. The study also highlighted a spatial reallocation of demand, with restaurants situated in “hidden alleys” growing more than those positioned directly in front of tourist attractions. For low-rated establishments, the likelihood of exiting the market doubled post-policy compared to the baseline period. On the whole, the probability of hiring a worker with prior experience in the restaurant industry increased by 10% relative to the pre-policy mean.

4 Discussion and conclusions

The Internet and social media platforms have revolutionized our interactions, communications, and access to information. As these transformations continue to evolve, understanding their impacts remains imperative. In this light, our paper has reviewed the empirical literature on the behavioral, economic, and political effects of online connectivity.

The analysis of existing research reveals a web of effects that operate at individual and societal levels. At the individual level, high-speed Internet access has shown to bring substantial benefits, such as better access to health information, enhanced job search, and improved educational attainment and productivity for workers and businesses. Conversely, increased Internet access has superseded certain offline activities, specifically those relating to civic engagement and political participation. Online connectivity also poses potential risks to aspects of well-being, including mental health and sleep quality. At an aggregate level, faster Internet and social media use have been found to increase protest occurrence and participation, support for populist leaders, and hate crimes, while worsening voter turnout and mental health, among other outcomes. This array of outcomes underscores the multi-dimensional and context-dependent impacts of the Internet and social media platform, highlighting the need for a comprehensive understanding, taking into account the diverse ways these platforms can be exploited by different actors within varying political contexts.

However, this paper goes further than merely cataloging the outcomes associated with Internet and social media usage. It delves into the empirical challenges researchers face in determining the impacts of Internet and social media use. We have critically evaluated various identification strategies, highlighting their respective weaknesses and strengths.

Our review emphasizes the importance of identifying an exogenous source of variation for the treatment, be it exposure to fast Internet or specific content on social media platforms. Such an instrument must be orthogonal to the dependent variable under study, ensuring a quasi-random distribution of the treatment and the presence of a comparable control group, encompassing in-

dividuals, groups, or territorial clusters. Retrieving such experimental conditions in nature is challenging. The rise of DSL technology in the mid-2000s allowed scholars to leverage the topology of the early copper wire infrastructure to assess the impact of initial broadband penetration. This method exploits the distribution of network nodes that was established before or shortly after the Second World War, making it exogenous to modern-day supply decisions and consumer preferences.

Several studies have adopted different instruments to identify the impact of Internet access. For instance, the incidence of lightning strikes and the prevalence of certain weather patterns have been leveraged as exogenous variables affecting mobile Internet coverage and broadband access. However, these identification strategies face significant challenges, especially since weather conditions can affect a wide range of outcomes, potentially undermining the orthogonality condition crucial for experimental validity. When employing this strategy, it is essential to provide evidence that the particular weather phenomenon under consideration does not directly impact the outcome variable. Another approach involves capitalizing on the quasi-experimental setups created by policy interventions that facilitate fiber cable deployment or promote the purchasing of faster connections. While these strategies are promising, a careful evaluation is needed to verify the quasi randomization of cable deployment trajectories, as infrastructure tends to be prioritized in areas of high demand, and local administrative planning often plays a significant role.

In instances where identifying an exogenous source of variation for treatment assignment is non viable, researchers might opt for designing field or laboratory experiments. Conducting experiments within controlled settings offers a feasible means to evaluate the effects of specific online phenomena, such as the tone of online discourse (e.g., Antoci et al., 2019). While developing field experiments to measure the effects of high-speed Internet access can be complex and, at times, ethically questionable, such methods are especially suited for assessing the impacts of distinct social media content. Although outside the scope of this review, the literature does include a number of randomized trials that either expose participants to social media or restrict access to it (e.g., Allcott et al., 2020). The part of our paper dedicated to social media, however, concentrates on studies employing natural and quasi-experimental designs.

These investigations typically leverage the unique diffusion characteristics of particular social media platforms. The core challenge is to establish a treatment assignment mechanism that approximates randomness and is orthogonal to the outcome being studied. When these conditions are met, a natural experiment may arise, giving researchers the opportunity to discern the causal effects of social media usage.

A remarkable approach has been employed to examine one of the most notable case studies with implications for international peace and political stability: the influence of social media in authoritarian regimes, specifically Russia. Enikolopov et al. (2020a) leveraged the number of

students from the same university as VKontakte’s founder who originated from a given city serves as an instrumental variable to identify the political outcomes of the platform’s penetration. This method brilliantly assesses the impact of social media in a critical context, yet its generalizability is limited.

A broader methodology hinges on exploiting diverse supply-side factors that induce exogenous variation in social media access. Some authors have capitalized on specific settings provided by Facebook accessibility, ranging from its phased introduction across American colleges to language-specific availability and instances of platform outages. Although these strategies have primarily focused on Facebook, the underlying principle could be extended to analyze other platforms with phased or interrupted service, like TikTok or Meta’s Threads.

To summarize, whether investigating fast Internet or social media usage, our review suggests that future research should prioritize identifying exogenous factors that inadvertently randomize exposure to the treatment. This approach enables the comparison between outcomes of treated and untreated individuals or locations, ensuring a balanced control group. When such natural randomization is not available, field or laboratory experiments emerge as the most viable alternatives.

Similar observations can be made regarding the emerging field of research on the behavioral and economic impacts of artificial intelligence (AI). While a comprehensive review of these studies falls beyond the scope of this paper, noteworthy similarities are evident. AI, predominantly utilized online, can be empirically treated as a form of specific content accessible through high-speed Internet. The phased introduction or occasional interruptions of AI-based services present the opportunity to assess the economic impact of AI. For instance, Bertomeu et al. (2023) and Kreitmer and Raschky (2023) leveraged the temporary ban of Chat GPT in Italy, imposed by the Italian Data Protection Authority, to assess its impact on firm-level and individual productivity.

Reflecting the methodologies of many studies covered in this paper, research that capitalizes on variations in AI access must adopt designs akin to intention-to-treat analysis. An alternative, yet potentially effective, approach might involve field experiments in which subjects, whether they are firms or individuals, are randomly assigned the chance to utilize specific forms of AI.

In conclusion, while this review is purposefully focused in scope, we hope it serves as a practical guide for researchers embarking on empirical studies of new communication technologies, indicating the challenges that must be overcome and the methods that have been employed in the literature to date. However, we envision this work also as a helpful guide for a broader audience, including journalists and policymakers. In highlighting the importance of validating specific assumptions to ascertain the impact of Internet or social media use on particular outcomes, we encourage non-specialists to acknowledge the distinctions among various studies. Recognizing that it may be difficult for those outside the field to discern which studies have effectively tackled endogeneity issues, our paper offers some assistance. We advocate for a cautious and thorough approach when

interpreting the findings in this intricate area of study.

References

- Adema, J. A. H., Aksoy, C. G., and Poutvaara, P. (2022). Mobile Internet access and the desire to emigrate. CESifo Working Paper No. 9758.
- Ahlfeldt, G., Koutroumpis, P., and Valletti, T. (2017). Speed 2.0: Evaluating access to universal digital highways. *Journal of the European Economic Association*, 15(3):586–625.
- Akerman, A., Gaarder, I., and Mogstad, M. (2015). The skill complementarity of broadband Internet. *Quarterly Journal of Economics*, 130(4):1781–1824.
- Akerman, A., Leuven, E., and Mogstad, M. (2022). Information frictions, broadband Internet and the relationship between distance and trade. *American Economic Journal: Applied*, 14(1):133–163.
- Allcott, H., Braghieri, L., Eichmeyer, S., and Gentzkow, M. (2020). The welfare effects of social media. *American Economic Review*, 110(3):629–676.
- Amaral-Garcia, S., Nardotto, M., Propper, C., and Valletti, T. (2022). Mums go online: Is the Internet changing the demand for healthcare? *The Review of Economics and Statistics*, 104(6):1157–1173.
- Antoci, A., Bonelli, L., Paglieri, F., Reggiani, T., and Sabatini, F. (2019). Civility and trust in social media. *Journal of Economic Behavior and Organization*, 160:83–99.
- Arenas-Arroyo, E., Fernández-Kranz, D., and Nollenberger, N. (2022). High speed Internet and the widening gender gap in adolescent mental health: Evidence from hospital records. *IZA DP 15728*.
- Battisti, M., Kauppinen, I., and Rude, B. (2022). Twitter and crime: The effect of social movements on gender-based violence. CESifo Working Paper No. 381/2022.
- Bauernschuster, S., Falck, O., and Woessmann, L. (2014). Surfing alone? The internet and social capital: Evidence from an unforeseeable technological mistake. *Journal of Public Economics*, 117:73–89.
- Bertomeu, J., Lin, Y., Liu, Y., and Zhenghui, N. (2023). Capital market consequences of generative AI: Early evidence from the ban of ChatGPT in Italy.
- Bessone, P., Campante, F., Ferraz, C., and Souza, P. (2022). Social media and the behavior of politicians: Evidence from Facebook in Brazil. NBER Working Paper N. 30306.

- Bhuller, M., Ferraro, D., Kostøl, A. R., and Vigtel, T. C. (2023a). The Internet, search frictions and aggregate unemployment. *NBER WP 30911*.
- Bhuller, M., Havnes, T., Leuven, E., and Mogstad, M. (2013). Broadband Internet: An information superhighway to sex crime? *Review of Economic Studies*, 4(1):1237–1266.
- Bhuller, M., Havnes, T., McCauley, J., and Mogstad, M. (2023b). How the Internet changed the market for print media. *NBER WP 30939, forthcoming in American Economic Journal: Applied*.
- Billari, F. C., Giuntella, O., and Stella, L. (2018). Broadband Internet, digital temptations, and sleep. *Journal of Economic Behavior and Organization*, 153:58–76.
- Billari, F. C., Giuntella, O., and Stella, L. (2019). Does broadband Internet affect fertility? *Population Studies*, 73(3).
- Boeri, F. (2023). High-speed broadband, school closures and educational achievements. LSE Geography and Environment Discussion Paper Series.
- Braghieri, L., Levy, R., and Makarin, A. (2022). Social media and mental health. *American Economic Review*, 112(11):3660–3693. mimeo.
- Bursztyn, L., Egorov, G., Enikolopov, R., and Petrova, M. (2020). Social media and xenophobia: evidence from Russia. CEPR Discussion Paper DP14877.
- Cagé, J., Herve, N., and Mazoyer, B. (2022). Social media influence mainstream media: Evidence from two billion tweets. CEPR Discussion Paper DP17358.
- Cambini, C., Sabatino, L., and Zaccagni, S. (2022). The faster, the better? The effect of ultra-fast broadband on students' performance. CEBI Working Paper 14/22.
- Campante, F., Durante, R., and Sobbrío, F. (2018). Politics 2.0.: The multifaceted effect of broadband Internet on political participation. *Journal of the European Economic Association*, 16(4):1094–1136.
- Campante, F., Durante, R., and Tesei, A. (2022). Media and social capital. *Annual Review of Economics*, 14:69–91.
- Cao, A., Lindo, J. M., and Zhong, J. (2022). Can social media rhetoric incite hate incidents? Evidence from Trump's "Chinese virus". NBER Working Paper N. 30588.
- Casanueva-Artis, A., Avetian, V., Sardoschau, S., and Saxena, K. (2022). Social media and the broadening of social movements: Evidence from Black Lives Matter. IZA Discussion Paper 15812.
- Castellacci, F. and Tveito, V. (2018). Internet use and well-being: A survey and a theoretical framework. *Research Policy*, 47(1):308–325.

- Czernich, N. (2012). Broadband Internet and political participation: Evidence for Germany. *Kyklos*, 65(1):31–52.
- D’Andrea, A. and Limodio, N. (2023). High-speed Internet, financial technology, and banking. *Management Science*, <https://doi.org/10.1287/mnsc.2023.4703>.
- De Stefano, T., Kneller, R., Timmis, J., et al. (2023). The (fuzzy) digital divide: the effect of universal broadband on firm performance. *Journal of Economic Geography*, 23(1):139–177.
- Dettling, L. J., Goodman, S., and Smith, J. (2018). Every little bit counts: The impact of high-speed Internet on the transition to college. *Review of Economics and Statistics*, 100(2):260–273.
- Donati, D. (2022). The end of tourist traps: A natural experiment on the impact of Tripadvisor on quality upgrading. CESifo Working Paper No. 9834.
- Donati, D. (2023). Mobile internet access and political outcomes: Evidence from South Africa. *Journal of Development Economics*, 162(103073).
- Donati, D., Durante, R., Sobbrío, F., and Zejcirović, D. (2022). Lost in the net? Broadband Internet and youth mental health. IZA Discussion Paper 15202.
- Enikolopov, R., Makarin, A., and Petrova, M. (2020a). Social media and protest participation: Evidence from Russia. *Econometrica*, 88(4):1479–1514.
- Enikolopov, R., Makarin, A., and Petrova, M. (2023). Online corrigendum to "social media and protest participation: Evidence from Russia". *Econometrica*, (<https://doi.org/10.3982/ECTA20824>).
- Enikolopov, R., Makarin, A., Petrova, M., and Polishchuk, L. (2020b). Social image, networks, and protest participation. Available at SSRN: <http://dx.doi.org/10.2139/ssrn.2940171>.
- Enikolopov, R., Petrova, M., and Sonin, K. (2018). Social media and corruption. *American Economic Journal: Applied*, 10(1):150–174.
- Faber, B., Sanchis-Guarner, R., and Weinhardt, F. (2015). Ict and education: evidence from student home addresses. *NBER WP 21306*.
- Falck, O., Gold, R., and Heblich, S. (2014). E-lections: Voting behavior and the Internet. *American Economic Review*, 104(7):2238–65.
- Fergusson, L. and Molina, C. (2021). Facebook causes protests. Documento CEDE No. 41.
- Fujiwara, T., Muller, K., and Schwartz, C. (2022). The effect of social media on elections: Evidence from the United States. NBER Working Paper N. 28849. Conditionally accepted, Journal of the European Economic Association.
- Gavazza, A., Nardotto, M., and Valletti, T. (2019). Internet and politics: Evidence from UK local elections and local government policies. *Review of Economic Studies*, 86(5):2092–2135.

- Geraci, A., Nardotto, M., Reggiani, T., and Sabatini, F. (2022). Broadband Internet and social capital. *Journal of Public Economics*, 206(104578).
- Goldfarb, A. and Tucker, C. (2019). Digital economics. *Journal of Economic Literature*, 57(1):3–43.
- Guriev, S., Melnikov, N., and Zhuravskaya, E. (2020). 3G Internet and confidence in government. *Quarterly Journal of Economics*, 136(4):2533–2613.
- Guriev, S. and Papaioannu, E. (2022). The political economy of populism. *Journal of Economic Literature*, 60(3):753–832.
- Gurtzgen, N., Diegmann, A., Pohlen, L., and van den Berg, G. J. (2021). Do digital information technologies help unemployed job seekers find a job? Evidence from the broadband Internet expansion in Germany. *European Economic Review*, 132(103657).
- Gürtzgen, N., Lochner, B., Pohlen, L., and van den Berg, G. J. (2021). Does online search improve the match quality of new hires? *Labour Economics*, 101981.
- Hjort, J. and Poulsen, J. (2019). The arrival of fast Internet and employment in Africa. *American Economic Review*, 109(3):1032–1079.
- Hjort, J. and Tian, L. (2022). The economic impact of Internet connectivity in developing countries. Columbia University, mimeo.
- Houngbonobon, G. V., Mensah, J., and Traore, N. (2022). The impact of internet access on innovation and entrepreneurship in Africa. World Bank Policy Research Working Paper n. 9945.
- Hvide, H. K., Meling, T. G., Mogstad, M., and Vestad, O. L. (2023). Broadband Internet and the stock market investments of individual investors. *NBER WP 30383*, forthcoming in the *Journal of Finance*.
- King, G., Pan, J., and Roberts, M. E. (2017). How the chinese government fabricates social media posts for strategic distraction, not engaged argument. *American Political Science Review*, 111(3):484–501.
- Kreitmer, D. H. and Raschky, P. A. (2023). The unintended consequences of censoring digital technology - Evidence from Italy’s ChatGPT ban.
- Levy, R. and Mattsson, M. (2023). The effects of social movements: Evidence from metoo. Available at SSRN: <http://dx.doi.org/10.2139/ssrn.3496903>.
- Manacorda, M., Tabellini, G., and Tesei, A. (2023). Mobile Internet and the rise of communitarian politics. CEPR Discussion Paper No. 18063.
- Manacorda, M. and Tesei, A. (2020). Liberation technology: Mobile phones and political mobilization in Africa. *Econometrica*, 88(2):533–567.
- McDool, E., Powell, P., Roberts, J., and Taylor, K. (2020). The Internet and children’s psychological wellbeing. *Journal of Health Economics*, 69.

- Mellon, J. (2022). Rain, rain, go away: 192 potential exclusion-restriction violations for studies using weather as an instrumental variable. *American Journal of Political Science*, *conditionally accepted*.
- Melnikov, N. (2023). Mobile Internet and political polarization. Nova School of Business and Economics, mimeo.
- Miner, L. (2015). The unintended consequences of internet diffusion: Evidence from Malaysia. *Journal of Public Economics*, 132:66–78.
- Müller, K. and Schwartz, C. (2021). Fanning the flames of hate: Social media and hate crime. *Journal of the European Economic Association*, 19(4):2131–2167.
- Müller, K. and Schwartz, C. (2023). From hashtag to hate crime: Twitter and anti-minority sentiment. *American Economic Journal: Applied Economics*, forthcoming.
- Nardotto, M., Valletti, T., and Verboven, F. (2015). Unbundling the incumbent: evidence from UK broadband. *Journal of the European Economic Association*, 13(2):330–362.
- Ofcom (2010). UK broadband speeds 2010: The performance of fixed-line broadband delivered to UK residential consumers. Technical report, Ofcom, London, UK.
- Ofcom (2016). Clarity for businesses baffled by broadband speeds. Technical report, Ofcom.
- Page, L., Brin, S., Motwani, R., and Winograd, T. (1999). The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab.
- Rotesi, T. (2019). The impact of Twitter on political participation. University of Lausanne, mimeo.
- Sanchis-Guarner, R., Montalban, J., and Weinhardt, F. (2021). Home broadband and human capital formation. CESifo Working Paper No. 8846.
- Strazzeri, M. (2022). Internet usage and migration decisions: Evidence from Nigerian micro data. University of Bern, mimeo.
- Zhuravskaya, E., Petrova, M., and Enikolopov, R. (2020). Political effects of the Internet and social media. *Annual Review of Economics*, 12:415–438.
- Zuo, G. W. (2021). Wired and hired: Employment effects of subsidized broadband Internet for low-income Americans. *American Economic Journal: Economic Policy*, 13(3):447–482.